Expectations and Experience Effects: New Facts From an Empirical Model

Heiner Beckmeyer* and Ahmed Guecioueur†

August 28, 2025

(First version: February 2023)

Abstract

We develop an empirical model of how households form beliefs about stock market returns, utilizing advances in machine learning tools. Our approach estimates interpretable experience weights that vary across households while capturing a shared expectation formation process. These estimated experience effects explain close to one-fifth of the variation in subjective return expectations extracted from survey microdata. Interpreting the estimated weights reveals four stylized facts that offer new empirical foundations for theories of expectation formation. Most notably, we uncover a role for expected *future* states to influence the retrieval of past experiences, and illuminate households' perceived links between asset prices and the business cycle.

Keywords: Expected Returns, Household Surveys, Disagreement, Machine Learning
JEL Codes: D83, G11, G12, G41, G51

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

We would like to thank Philip Bond, Stefano Cassella, Josefina Cenzón, Tim de Silva, Theodoros Evgeniou, Roberto Gómez-Cram, Naveen Gondhi, Maurizio Montone, Rodolfo Prieto, Joshua D. Rauh, Stephan Siegel, Mark M. Westerfield, Trevor Young 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. Gueciouer gratefully acknowledges the support of the Henry Grunfeld Foundation. This paper was previously circulated under the title "Expectations and Attention to Experience."

1 Introduction

Beliefs serve as crucial inputs to financial and economic decisions. Their importance has driven rapid growth in a literature seeking to understand expectation formation, with a particular emphasis on models in which beliefs depend on past experiences and memories.[1] While the theoretical diversity of this literature reflects a richness of potential mechanisms, a key challenge is identifying which systematic patterns a general expectation formation model should incorporate in order to accurately capture investor behavior.

We address this challenge by developing and estimating a flexible *empirical* model that uncovers systematic patterns in how households rely on past experiences as they form beliefs about the stock market's return. We establish a benchmark with strong empirical performance, explaining nearly one-fifth of variation in surveyed beliefs, including out-of-sample. We then interpret the parameters of our benchmark in order to highlight the mechanisms that theoretical models should jointly capture to match our explanatory power in the data.

Our methodology builds on the broad class of experience-based expectation formation models. Rather than imposing restrictive assumptions about which experiences are more influential, we adapt machine learning techniques to our setting in order to reveal these patterns from household survey data. Our data-driven approach overcomes a common challenge faced by existing models: these typically require researchers to make functional form assumptions (e.g., exponential decay) or to specify how households assess similarity between experiences, which aspects of past events are most salient, or what contextual factors trigger memory retrieval, to give just a few model-dependent examples. By flexibly estimating household-specific experience weights, we can identify patterns to inform future theoretical work. Our approach also incorporates observed household characteristics (which do not necessarily enter into existing theoretical models) to understand their potential role.

Four Facts Linking Experiences and Return Expectations We establish four key empirical facts that any comprehensive theory of expectation formation should accommodate in order to match the explanatory power of our empirical model.

First, we replicate the well-documented recency bias within our framework, finding that

[1] Theoretical models of expectation formation include extrapolative expectations (Barberis, Greenwood, Jin, and Shleifer, 2015), lifetime experience effects (Malmendier and Nagel, 2011), diagnostic expectations (Bordalo, Gennaioli, Porta, and Shleifer, 2019), salience & attention (see the review by Bordalo, Gennaioli, and Shleifer (2022)), and memory & association (e.g., Mullainathan, 2002; Gennaioli and Shleifer, 2010; Bordalo, Gennaioli, and Shleifer, 2020; Wachter and Kahana, 2024; Enke, Schwerter, and Zimmermann, 2024). These varied models share a common foundation: that certain experiences should carry more weight than others in shaping beliefs. Our paper empirically documents how such differential weighting manifests when households form expectations about the stock market return.

experience weights follow an approximate exponential decay pattern, with a high degree of extrapolation from recent experiences. However, this pattern of mechanical decay explains only about one-third of the variation in our estimated experience weights, indicating that additional mechanisms are at work, and motivating our investigation of new patterns.

Second, we demonstrate that demographic characteristics and macroeconomic beliefs are important drivers of how households weight past stock market experiences. By allowing experience weights to vary based on these instruments, our framework captures time-varying experience effects that account for the majority of our model's explanatory power.[2] These instruments enable us to explain both temporal variation and cross-sectional differences in return expectations, via corresponding variation in experience weights across households and over time.

Our finding that experience weights vary systematically across households supports Meeuwis, Parker, Schoar, and Simester's (2022) argument that belief disagreement arises because households hold different models of the world at a given point in time. We exploit the occurrence of the 2016 Presidential election during our sample to show that households' differing mental models result at least in part from local, heuristic thinking: households shift their experience weights based on their own political affiliations, and this effect cannot be explained purely by changes to their beliefs about macroeconomic conditions as a result of the election outcome.

Third, we show that recession periods – including recessions in the distant past – are salient events that exert an outsize impact on the beliefs that we form today. Past recession episodes receive systematically higher weights in expectation formation than expansion periods, even after controlling for recency effects. Furthermore, this effect strengthens whenever expectations are elicited during a recession period, suggesting a role for similarity (or association) in macroeconomic conditions. This empirical fact is consistent with long-term belief scarring, and exerts a potentially opposing force to the recency effect; explaining both effects jointly is therefore a useful target for theories of expectation formation.

Fourth, we uncover a novel similarity-based weighting mechanism: we show that households weight past stock market experiences more heavily when historical macroeconomic conditions during those episodes resemble their present beliefs about *future* economic conditions. This indicates that expectation formation involves an active process of pattern matching between past and expected future states, rather than simply retrieving experiences based upon the state of the world today. This future state-dependent weighting mechanism has important

[2]Three-quarters of our richest specification's explanatory power is due to time-varying experience decays. Importantly, this time-variation is due purely to demographic composition changes and variation in households' common beliefs about the macroeconomy, rather than reduced-form time effects. The remaining one-quarter of our explanatory power is due to within-cohort heterogeneity, which includes the role played by disagreement about the state of the macroeconomy.

implications for theories of expectation formation, since the effect is not necessarily captured by models in which *current* cues or contexts drive the process by which an investor retrieves (or is otherwise influenced by) past experiences. Additionally, this stylized fact reinforces the importance of including households' beliefs about the macroeconomy as inputs into the process by which they form stock market beliefs.

Microdata on Subjective Return Expectations Our paper begins by introducing a rich source of microdata on subjective expectations about the return of the market portfolio. We show how to extract households' subjective return expectations 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 methodological innovation is to recognize that two decades' worth of monthly responses to a consistently-elicited survey question can be recast as return expectations. We show that a simple probit transformation of survey responses is consistent with two commonly-assumed perceived data generating processes for returns.[3]

We extract subjective return expectations that behave similarly to others that have been examined in the literature: these beliefs are extrapolative (Malmendier and Nagel, 2011), procyclical (Greenwood and Shleifer, 2014), positively associated with beliefs about the macroeconomy (Amromin and Sharpe, 2014; Giglio, Maggiori, Stroebel, and Utkus, 2021) and covary with demographic characteristics (Das, Kuhnen, and Nagel, 2020; Giglio, Maggiori, Stroebel, and Utkus, 2021). The subjective expectations that we extract are also consistent predictors of household demand, at both the individual household level and the level of the representative household, and thus drivers of stock market flows (Gabaix and Koijen, 2021).

This dataset has several unique advantages. 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 extract its subjective return expectations about the market portfolio, as *jointly* perceived by that individual household. We are thus able to use this dataset to estimate a general expectation formation process, that allows for subjective within-household links between perceptions of the stock market and the business cycle.

[3]A similar parametric assumption has been employed by the prior literature: see Dominitz and Manski, 2007; Ben-David, Graham, and Harvey, 2013; Lochstoer and Muir, 2022.

Machine Learning-Based Methodology This paper's main methodological contribution is to introduce a machine learning-based methodology for estimating a systematic expectation formation process used by households, while allowing for heterogeneous variation in outcomes across households. Our framework generalizes the experience-based learning framework of Malmendier and Nagel (2011), which captures the influence of historical stock market returns on the return expectations that households hold today, using 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. We introduce a customized Transformer neural network architecture (Vaswani et al., 2017) that flexibly estimates household-level experience effect weights to best explain individual households' beliefs about the stock market.[4] We employ the Michigan survey microdata described earlier in our estimation.

Our richest estimated specification is able to explain almost a fifth of both aggregate and cross-sectional variation in households' subjective return expectations over our two-decade sample period. This high degree of explanatory power confirms the usefulness of experience effects as an organizing empirical framework for understanding households' shared expectation formation process. Reassuringly, our model-implied beliefs are very similar to those produced by the state-of-the-art (yet reduced-form) Instrumented PCA technique (Kelly, Pruitt, and Su, 2019); we thus do not sacrifice explanatory power for the sake of economic interpretability. We also confirm our model's performance is robust during the unprecedented Covid market conditions that ensued in our out-of-sample holdout period.

We next document several facts about households' expectation formation process, taking advantage of the economic interpretability our model's parameters offer. We find that, remarkably, the estimated experience effects exhibit an exponential-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. However, only a third of variation in estimated experience weights can be attributed to such a mechanical decay. Indeed, our estimated experience weights exhibit non-monotonicities and substantial heterogeneity, which lead us to investigate the presence of additional mechanisms.

[4]The economic restriction that experience weights must be non-negative and sum to unity both allows us to interpret our model parameters as experience effects, and 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.

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 better understand this 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.

We conclude our analysis by shedding light on cross-household disagreement (see Meeuwis, Parker, Schoar, and Simester (2022)) and the delayed response of household expectations to the onset of recessions (see Gómez-Cram (2022)).

Related Literature Our findings contribute to a growing literature that emphasizes the importance of establishing stylized facts to guide theoretical development in subjective expectation formation. Similar to the approach taken by Giglio, Maggiori, Stroebel, and Utkus's (2021) analysis of return expectations, we provide a systematic characterization of empirical patterns that can inform the development and application of expectation formation theories.

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). We introduce a measure of households' subjective return expectations that predicts individual households' equity holdings, and produce an empirical model of households' expectation formation process that is based upon their experiences and characteristics.[5] Our paper also sheds 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 char-

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

acteristics. 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.[6]

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 explain and understand human behavior, rather than to produce objective predictions 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 are generated by a particular memory model. We focus on general experience effects that may nest multiple mechanisms, and households' jointly-held expectations of returns and economic conditions.

Finally, this paper provides empirical evidence on households' perceived links between asset prices and the state of the macroeconomy, adding to findings by Amromin and Sharpe (2014) and Giglio, Maggiori, Stroebel, and Utkus (2021). Linking stock market returns with business cycles has been a focal point for research since the seminal work of Grossman and Shiller (1981), and our results can inform the development of macro-finance models that couple expected returns to the state of the macroeconomy (Cochrane, 2017).

Outline The remainder of this paper proceeds as follows. Section 2 describes our survey microdata and some empirical patterns. 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.

[6]To the extent that Large Language Models (LLMs) such as ChatGPT 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)).

2 Subjective Expectations From Survey Microdata

2.1 Extracting Beliefs About the Stock Market Return

To study how investors form expectations about stock market returns, we exploit two decades (2002-2022) of detailed monthly microdata from the Michigan Surveys of Consumers, which are a nationally-representative and broad-based survey of households. Our sample period covers the full business cycle, including the global financial crisis, as well as the unprecedented Covid-19 shock.

Our main variable of interest stems from the following question posed to Michigan survey respondents throughout the two-decade sample period: *"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 stock market return:

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

Proposition 1 shows that we can extract an even more useful subjective return expectation from the investor's response by applying a simple probit transformation that converts each survey response into units of (risk-adjusted) return:

Proposition 1 (Subjective Return Expectation). *Assume that an individual i perceives the market's annual return to be conditionally Normally distributed, $r_{t+1} \sim \mathcal{N}(\mu_{it}, \sigma_{it})$. Then:*

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

Validating the Assumption Internet Appendix A discusses the assumption underlying Proposition 1, empirically verifies it is reasonable for our sample period and setting, and discusses an alternative assumption that nevertheless leads to a similar identity to the above.

Validating Extracted Subjective Return Expectations By placing households' survey responses into units of (risk-adjusted) return, the one-to-one mapping given by (2) allows us to conveniently interpret households' survey responses as subjective return expectations. As one would expect from a return expectation, we show in Internet Appendix B that the quantity given by Eqn. (2) is a significant predictor of investor demand, both at the level of an individual household i, as well as at the level of the aggregate household sector (measured both as the aggregate equity share and as flows into the stock market); we confirm that this predictive

power is not due to mechanical variation in the stock market's valuation. These results all align with basic portfolio choice intuition.[7]

The aggregate behavior of these subjective return expectations is also consistent with the prior literature: the timeseries shown in Figure 1 is the cross-sectional mean over time, and highlights procyclical variation over the business cycle, consistent with a key finding by Greenwood and Shleifer (2014). In addition, the cross-sectional dispersion series shown in Figure 2 indicate the presence of substantial heterogeneity in subjective return expectations among investors surveyed in the same year-month, consistent with a key finding by Giglio, Maggiori, Stroebel, and Utkus (2021).

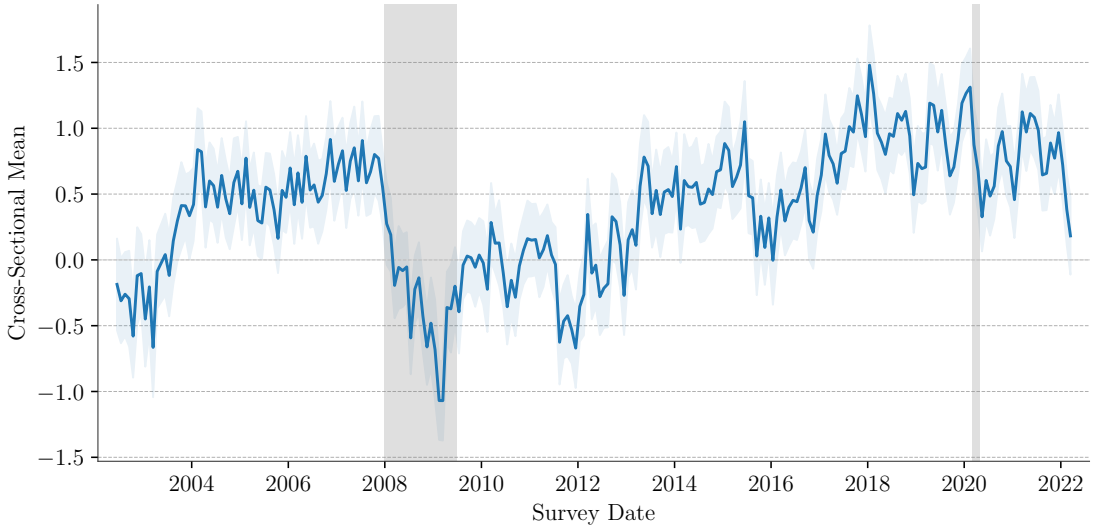


Figure 1: Mean of Survey Subjective Return Expectations

The time series shown in this chart 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 (see Proposition 1), at a monthly frequency. Blue-shaded areas are the ±2 standard error range. Gray-shaded areas are NBER recession periods.

Subjective Perceptions vs. Realizations Focussing for now on the time series aggregates, we present additional patterns that are consistent with the prior literature. Figure 3 shows that, on average, subjective return expectations diverge from the future realizations, over the same 1-year horizon. In this sense, subjective expectations are poor *ex ante* predictors of the future

[7]Survey waves that directly ask for expected returns and match these to comprehensive demographic data (e.g., Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014), while valuable, are also infrequent. To the best of our knowledge, other than our own dataset we are not aware of any other nationally-representative survey that consistently asks a question about the stock market's return during its monthly waves over a multi-decade-period. Another advantage of our data lies in its ability to predict investor demand: Internet Appendix C shows that our measure (2) is not subsumed by other subjective expected return measures when they are available – in fact, the quantity we extract is a comparatively stronger predictor (within a household) of the household's own actions.

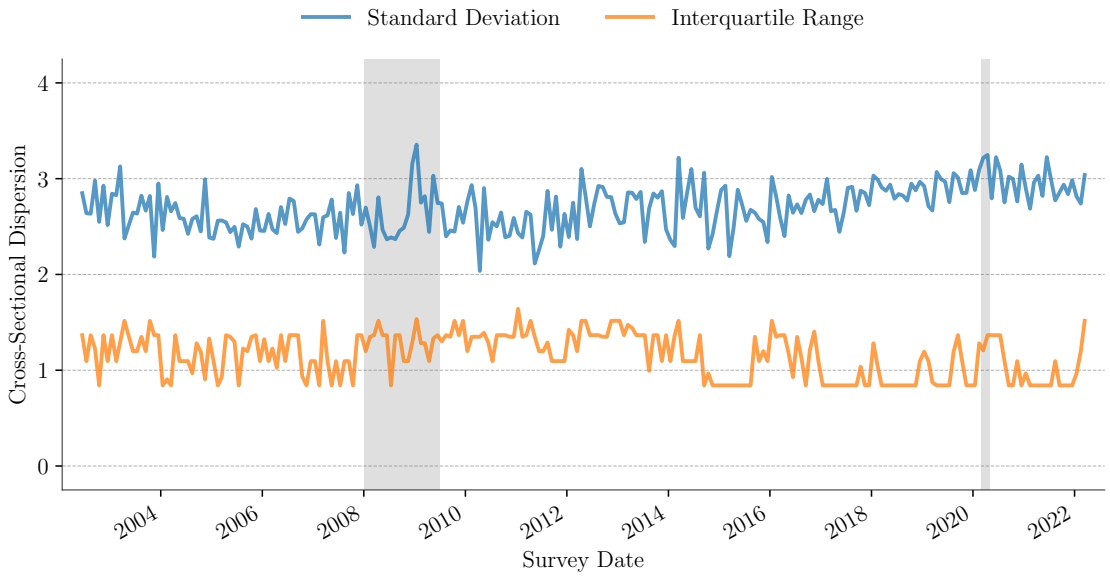


Figure 2: Dispersion of Survey Subjective Return Expectations

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.

ex post realizations. Figure 4 points to one reason why: the cyclicality of our subjective return expecations time series matches that of *past* 1-year realizations of the S&P 500 risk-adjusted return. The magnitude of this correlation is 0.11, and it is significant at the 1% level. On the one hand, this positive correlation is consistent with the large body of evidence that subjective return expectations 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.

2.2 Beliefs About the Business Cycle, and Other Variables

The Michigan surveys also record respondents' beliefs concerning general economic conditions, besides the stock market. Since these beliefs tend to be associated within a household (Amromin and Sharpe, 2014), such 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.

The Michigan surveys also record information about the demographic characteristics of

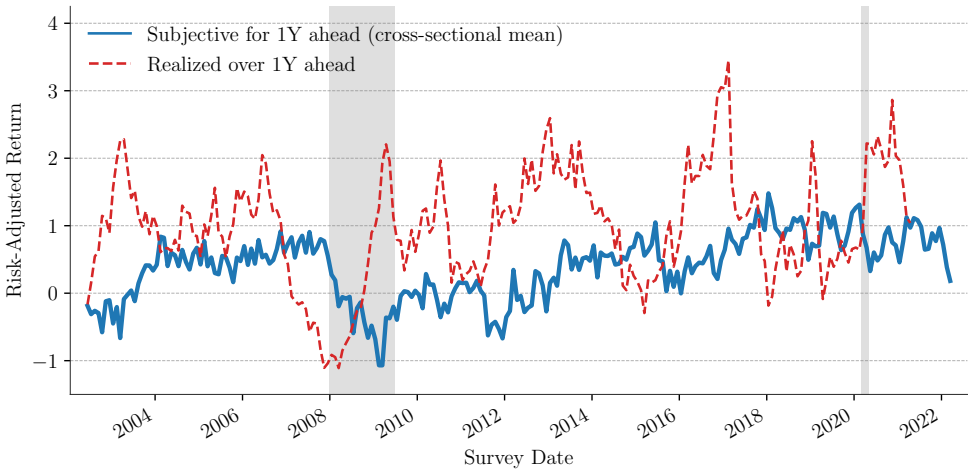


Figure 3: Mean Survey Subjective Return Expectations for Next Year vs. Actual Future Realization

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 (see Proposition 1). Secondary series (dashed red) is the actual, realized future twelve-month ahead risk-adjusted return of the S&P 500 at each survey date. The realized risk adjusted return is computed by the ratio of the mean daily returns within that year to the standard deviation of those daily returns, with an annualization factor. Shaded areas are NBER recession periods.

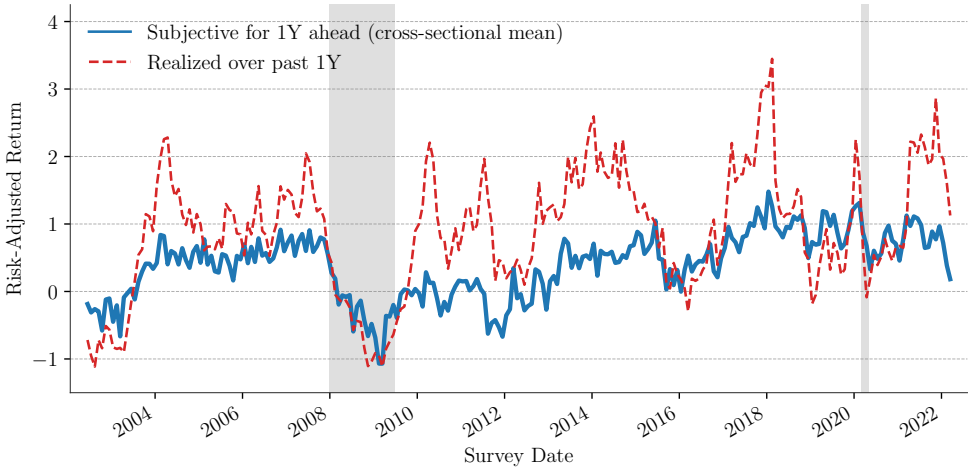


Figure 4: Mean Survey Subjective Return Expectations for Year Ahead vs. Previous Year Realization

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 (see Proposition 1). Secondary series (dashed red) is the realized rolling twelve-month risk-adjusted return of the S&P 500 looking back from each survey date. The realized risk adjusted return is computed by the ratio of the mean daily returns within that year to the standard deviation of those daily returns, with an annualization factor. 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

value.

Our primary 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.[8] After this filtering step, our survey microdata comprises 81,879 responses. Since our analyses condition on households who are also stock market investors, we will use the terms "household" and "investor" interchangeably. Table 1 presents summaries of the Michigan survey variables used in our study.

2.3 Relationship Between Macroeconomic and Stock Market 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 return when they are optimistic about macroeconomic conditions, over the same short-run horizon of the year ahead.

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

This table shows within household correlations between its subjective return expectation 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	$\rho(\cdot, SRE)$
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***

[8] 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.

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 effects, 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 return expectation perceived by an individual i surveyed at time t, $\text{SRE}_{i,t}$, is formed based on an exponentially-weighted average of all the lagged stock market realizations that she has experienced during her lifetime. We implement these experienced realizations using the lagged risk-adjusted returns RAR_{t-k} of the S&P 500 over her lifetime:

$$\text{SRE}_{i,t} = \sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \text{age}_{i,t}) \times \text{RAR}_{t-k} \tag{3}$$

$$\text{where} \quad w_{i,t}(k, \text{age}_{i,t}) = \frac{(\text{age}_{i,t} - k)^{\lambda}}{\sum_{k=1}^{\text{age}_{i,t}-1} (\text{age}_{i,t} - k)^{\lambda}} \tag{4}$$

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.[9]

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 this framework.

[9] 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 $\lambda > 1$ 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.

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 H. Much like the classic exponential weighting framework described in Eqns. (3)-(4), we model the subjective expectation held by each individual i surveyed at time t as a weighted sum of her past experienced return realizations RAR_{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:

$$\text{SRE}_{i,t} = \sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \times \text{RAR}_{t-k} \tag{5}$$

$$\text{where} \quad \mathbf{w}_{i,t} = g(\text{age}_{i,t}, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \tag{6}$$

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

$$\sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) = 1 \quad \text{and} \quad w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \geq 0 \ \ \forall \ \ k \in [1, \text{age}_{i,t} - 1]. \tag{7}$$

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. We will describe the estimation procedure after first making some brief remarks on the above setup.

Risk-Adjusted Returns Both the forward-looking target variable $\text{SRE}_{i,t}$ and the set of backward-looking stock market realizations RAR_{t-k} are in identical units of a risk-adjusted return: $\text{SRE}_{i,t}$ is the probit-transformed Michigan survey response described by Proposition 1, and RAR_{t-k} is the appropriately-annualized ratio of $\frac{\text{Mean of Daily Returns}}{\text{Std. Dev. of Daily Returns}}$ within the corresponding year $t-k$. There are a number of advantages to specifying the experience-based learning problem in this manner. Conceptually, both quantities $\text{SRE}_{i,t}$ and RAR_{t-k} are interpretable as return expecta-

tions and realizations rather than up-move-probabilities. Practically, historical realized returns RAR_{t-k} in high-volatility periods (such as recessions) do not unduly influence the estimation as they are dampened. Furthermore RAR_{t-k} observations exhibit considerable variation, which would not have been the case for binary indicators of whether the stock market had experienced an up-move (which would have led to substantial collinearity).

Macroeconomic Beliefs In addition to the household's demographic characteristics entering into Equation (5), its beliefs about macroeconomic conditions also play a key role. Namely, each household will attend to or retrieve 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 $\mathbf{b}_{i,t}$. This assumption is consistent with prior findings: Von Gaudecker and Wogrolly (2022, Figure 1) surveyed households on what they most thought about when forming subjective stock market return probabilities, and found the "state of the economy" was the dominant response.

3.3 Estimation Procedure

We implement our empirical framework using a customized Transformer neural network.

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. 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 returns that realized over the lifetime of respondent i, $\mathbf{RAR}_{(t-\text{age}_{i,t}-1):t}$. The database is characterized by keys $\mathbf{k}$ and values $\mathbf{v}$. $\mathbf{k}$ is a nonlinear transformation of $\mathbf{RAR}_{(t-\text{age}_{i,t}-1):t}$, which is directly learned from the data:[10]

$$\mathbf{k}_{i,t} = \mathbf{W}_2^{\mathbf{k}} \times \text{GELU}\left(\mathbf{W}_1^{\mathbf{k}} \times \mathbf{SR}_{(t-\text{age}_{i,t}-1):t} + \mathbf{a}_1^{\mathbf{k}}\right) + \mathbf{a}_2^{\mathbf{k}}. \tag{8}$$

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 of past return realizations, 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

[10] 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.

across and nonlinearities within different demographic characteristics and various macroeconomic beliefs. The details of this step are given in Appendix H. After this interaction step, we obtain queries $\mathbf{q}$ through a linear transformation of these interaction effects $\widehat{\mathbf{d}}_{i,t}, \widehat{\mathbf{b}}_{i,t}$:

$$\mathbf{q}_{i,t} = \mathbf{W}^{\mathbf{q}} \times [\widehat{\mathbf{d}}_{i,t}, \widehat{\mathbf{b}}_{i,t}], \tag{9}$$

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 appealing for the extraction of experience effects:

$$\mathbf{Attn}_{i,t} = \frac{\mathbf{q}_{i,t} \times \mathbf{k}'_{i,t}}{\sqrt{d^{\text{Attn}}}} \tag{10}$$

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 realizations of the stock market. 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 stock market experiences (indexed by k) are most informative for explaining the subjective return expectations of respondents who share a given combination of demographics and other beliefs:[11]

$$w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) = \frac{e^{\text{Attn}_{i,t}(k)}}{\sum_{k=1}^{\text{age}_{i,t}-1} e^{\text{Attn}_{i,t}(k)}}, \quad \forall\ k \in [1, \text{age}_{i,t} - 1] \tag{11}$$

In words, we assign a larger weight to past stock market returns 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 $\widehat{\text{SRE}}_{i,t}$ values best explain the variation in households' observed SREs. To do this, we minimize the mean squared error of model-predicted return expectations across subjects i and months t:

$$\text{MSE} = \sum_{i}^{I} \sum_{t}^{T} \left(\widehat{\text{SRE}}_{i,t} - \text{SRE}_{i,t} \right)^2 \tag{12}$$

[11] 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 (7).

The primary output of our Transformer-based estimation procedure is a panel of estimated experience effects $\widehat{w}_{i,t,k}$ for which the following conditions hold:

$$\widehat{\text{SRE}}_{i,t} = \sum_{k=1}^{\text{age}_{i,t}-1} \widehat{w}_{i,t,k} \times \text{RAR}_{t-k} \tag{13}$$

$$0 \leq \widehat{w}_{i,t,k} \quad \forall\, i,t,k \tag{14}$$

$$1 = \sum_{k=1}^{\text{age}_{i,t}-1} \widehat{w}_{i,t,k} \quad \forall\, i,t \tag{15}$$

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.

3.4 Explanatory Power of Our Empirical Model

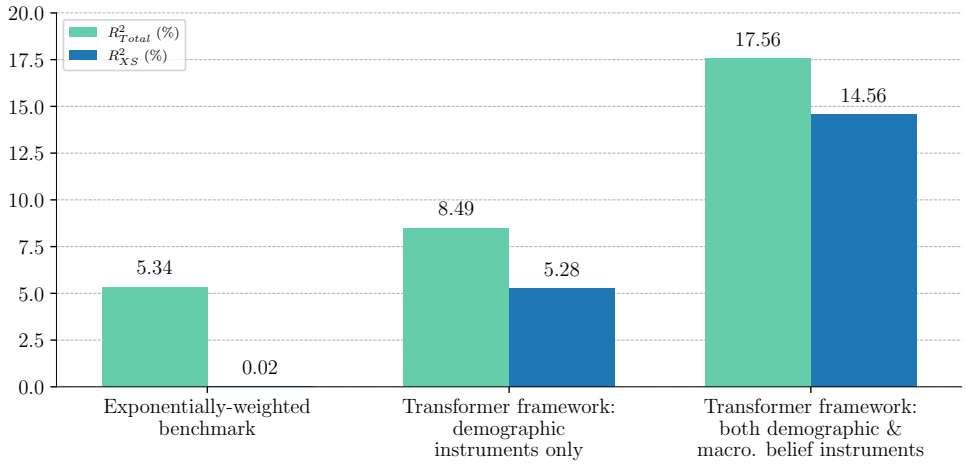


Figure 5: Explanatory Power Using Different Models Of Experience Effects

This figure compares the proportion of variance in respondents' subjective return expectations 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} (16), while green columns focus on the fraction (%) of cross-sectional variance explained R^2_{XS} (17).

As a first result, we measure how well our models of experience effects perform in explaining respondents' subjective return expectations, $\text{SRE}_{i,t}$. Figure 5 displays the performance of Transformer-based experience effects models (Equation (5)), alongside a classic exponentially-weighting baseline.[12] Our two Transformer-estimated models differ solely in which variables are used as instruments for experience weights in Equation (6): 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 $\text{SRE}_{i,t}$ can be explained in comparison to a benchmark of the cross-sectional mean in month t, using an R^2_{XS} metric that is equivalent to the "within-R^2" of a panel regression with time effects:[13]

$$R^2_{Total} = 1 - \frac{\sum_{i=1}^{N}(\text{SRE}_{i,t} - \widehat{\text{SRE}}_{i,t})^2}{\sum_{i=1}^{N}(\text{SRE}_{i,t} - \overline{\text{SRE}})^2} \tag{16}$$

$$R^2_{XS} = 1 - \frac{\sum_{i=1}^{N}\left((\text{SRE}_{i,t} - \overline{\text{SRE}}_t) - (\widehat{\text{SRE}}_{i,t} - \overline{\widehat{\text{SRE}}}_t)\right)^2}{\sum_{i=1}^{N}(\text{SRE}_{i,t} - \overline{\text{SRE}}_t)^2} \tag{17}$$

The R^2 results of Figure 5 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 5 – 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

[12] The baseline model is described in Equation (3). We set the optimal value of $\widehat{\lambda} = 10.39$ for this exponentially weighted baseline using a grid search, see Appendix F.

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

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 overall explanatory power, we also verify that the performance decay from our training period ("in-sample") to an out-of-sample hold-out period is reasonable. To do this, we assess explanatory power during the unprecedented Covid-related economic conditions that occurred from 2020 onwards; recall that our model was trained up to 2019. Figure 6 compares the total R^2 across both periods. This shows the explanatory power of our richest specification (which includes macroeconomic beliefs as instruments) decays by only one-fifth out-of-sample.[14]

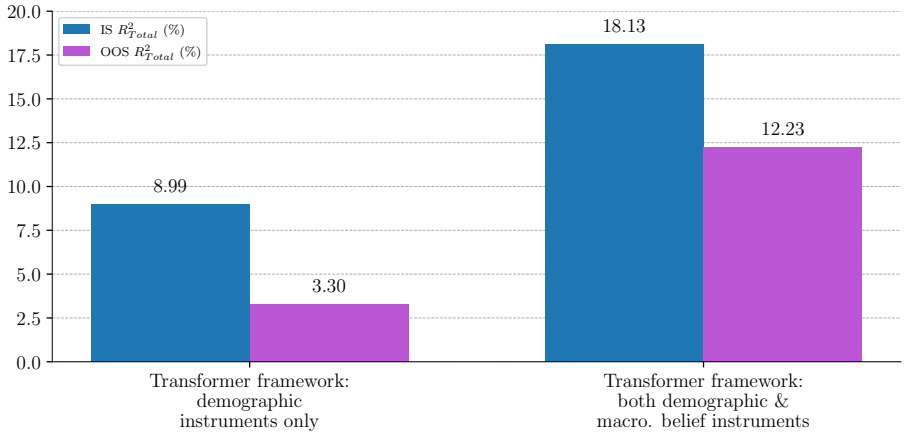


Figure 6: Out-of-Sample Explanatory Power

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

3.5 Comparison to an IPCA benchmark

As a final validation check, we compare our main model's explained subjective return expectations (SREs) with the portion of SREs that are explained by the state-of-the-art Instrumented

[14] As an example to put this in context, Campbell and Thompson (2008, Table 1B) show that one-year-ahead realized returns can be predicted by the D/P ratio with economic restrictions imposed, and report a steeper out-of-sample R^2 decay to about one-half of the in-sample magnitude However, we caution that the two samples differ, and that we do not focus on realized returns.

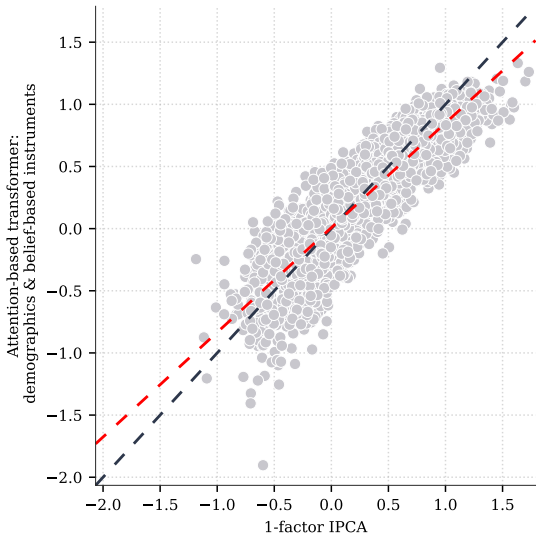


Figure 7: Predictions of the Latent 1-Factor Model vs. Transformer-based Experience Effects

The figure shows the predictions of the 1-factor IPCA, specified in Equation (14) 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.

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 G. 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 means it can serve as a useful benchmark for our own approach.

Figure 7 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 8 visualizes these experience weights, and highlights the substantial heterogeneity in households' models of the stock market.

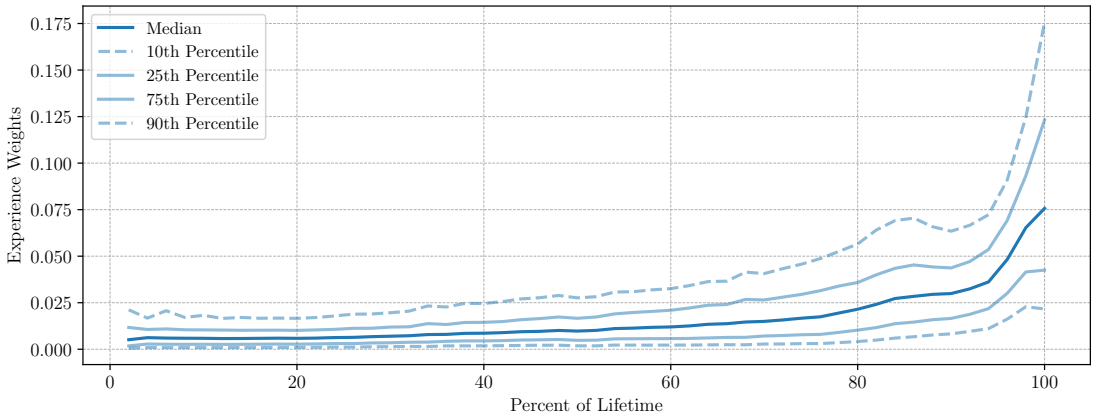


Figure 8: Heterogeneity in Estimated Experience Weights

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 The estimated experience weights in Figure 8 exhibit an exponential-like decay; remarkably, this pattern is produced even though we did not specify an explicit functional form in 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)).

Variance Decomposition of Experience Weights We now investigate the sources of such 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 8). 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. Thus, a mechanical decay effect (capturing a recency bias, but ignoring the role of the similarity of economic conditions) captures a minority of total variation. Figure 9 Panel A shows the results of this projection: the middle chart reveals that considerable timeseries variation remains.

To allow for the possibility that experience weights vary with time (as shown by Cassella and Gulen (2018)), we next project the estimated experience weights onto doubly-interacted Survey Date $t \times$ Lag $k \times$ 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 interpret this result as showing the importance of including common perceptions of economic conditions for capturing additional time-variation in households' stock return expectations. Figure 9 Panel B shows the results of this projection: the middle chart highlights the absence of any remaining timeseries variation.

Finally, we note that the remaining variation must be attributable to cross-sectional variation in experience effects. This variation includes both heterogeneity in sociodemographic characteristics within a given cohort, together with differing perceptions of the state of the business cycle (which may arise for idiosyncratic reasons). Since no timeseries variation remains at this stage, we furthermore conclude that any remaining disagreement in beliefs about the business cycle feeds through into disagreement about the stock market return (i.e., cross-sectional heterogeneity).

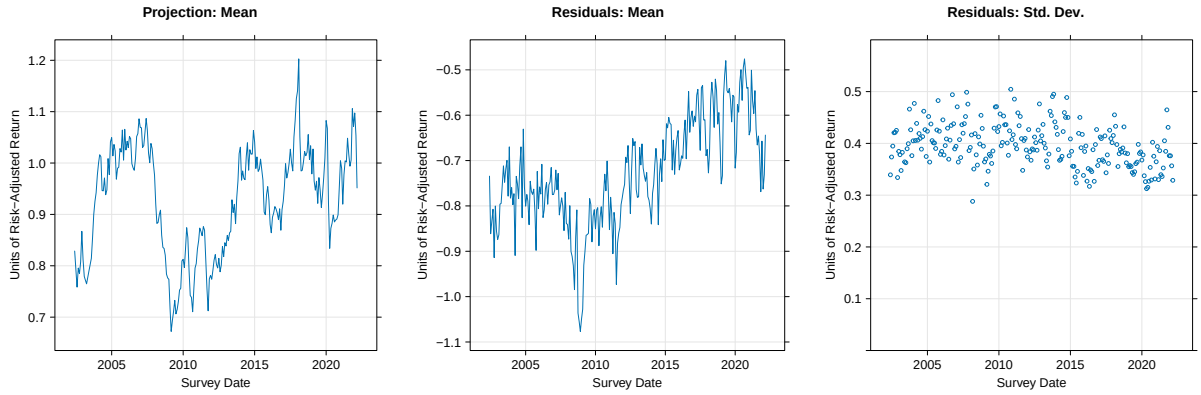
This variance decomposition analysis show our empirical model encompasses multiple dimensions of heterogeneity (both timeseries & cross-sectional).

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

Panel A: Projecting Modeled Beliefs onto Lag k × Age$_{i,t}$ Fixed Effects



Panel B: Projecting Modeled Beliefs onto Survey Date t × Lag k × Age$_{i,t}$ Fixed Effects

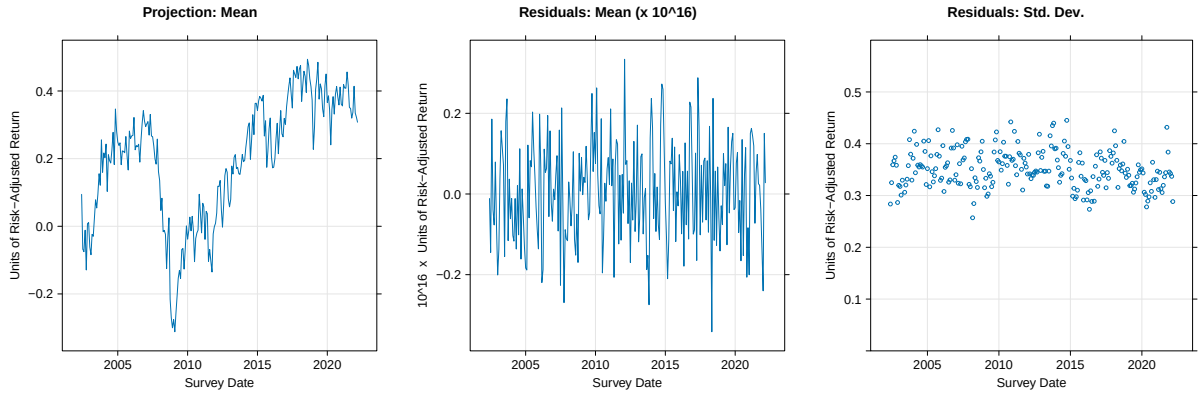


Figure 9: Interpreting Modeled Beliefs About the Stock Market Return Using Projections onto Fixed Effects

Note that "residuals" refer to residuals from the fixed effects projection, not residuals from our Transformer-based empirical model itself: all this "residual" variation is captured by our Transformer-based framework.

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 subjective return expectations, stock market return realizations experienced during the household's lifetime, and the household's 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,

$$\begin{aligned}\text{Macro Pessimism}_{i,t} \;=\; (\,&\text{Unemployment Direction}_{i,t} - \text{Business Conditions Direction}_{i,t} \\ &- \text{Interest Rates Direction}_{i,t}\,) \,/\, 3,\end{aligned} \tag{18}$$

and include it as an interaction term in the panel regression. 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[15] between a vector of each household's macroeconomic beliefs $\mathbf{b}_{i,t}$ and a vector of each past year's change in equivalent observed macroeconomic variables $\mathbf{m}_{t-k}$, using multiple measures. We then repeat the exercise for the cosine similarity between the future conditions that are yet to realize $\mathbf{m}_{t+1}$ and the same past realizations $\mathbf{m}_{t-k}$. Including both in a panel regression, we find a positive coefficient on $\cos(\mathbf{b}_{i,t}, \mathbf{m}_{t-k})$ in Table 5, meaning 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 going forward. This effect controls for the similarity $\cos(\mathbf{m}_{t+1}, \mathbf{m}_{t-k})$ between actual past realizations and actual future realizations. Furthermore, the included time effects in the panel regression reveal the effect is cross-sectional and independent of the objective state of the macroeconomy.

Local Thinking Since households appear to be trying to predict the future return by looking up past returns, we next investigate to what extent their 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 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).

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

[15]The cosine similarity between two vectors $\mathbf{u}$ and $\mathbf{v}$ is $\cos(\mathbf{u}, \mathbf{v}) = \frac{\mathbf{u}\cdot\mathbf{v}}{||\mathbf{u}||\,||\mathbf{v}||} \in [-1, +1]$.

Table 3: Experience Weights on Past Recession Periods

The dependent variable in each regression is $\widehat{w}_{i,t,k}$, the estimated experience weight on a lag k for a household i surveyed at time t in order to best explain that household's beliefs about the stock market return; see Eqn. (13). The experience effects estimates are produced using our richest Transformer-based specification. The independent variables are an indicator for whether that particular lifetime experience fell into a regression, as well as an indicator for whether the Michigan survey was run during a recession: note that neither are inputs into the Transformer estimation procedure.

	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

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 any mechanical weighting effects, 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 10. Interpreting the chart, we observe a lagged response of experience effects, which is consistent with Gómez-Cram (2022)'s argument.

Table 4: Experience Weights and Pessimism About the Macroeconomy

The dependent variable in each regression is $\widehat{w}_{i,t,k}$, the estimated experience weight on a lag k for a household i surveyed at time t in order to best explain that household's beliefs about the stock market return; see Eqn. (13). The experience effects estimates are produced using our richest Transformer-based specification. The macro pessimism independent variable is computed according to Eqn. (18).

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

*Standard errors are clustered at the level of the fixed-effects shown. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 5: Experience Weights and the Similarity of Macroeconomic Conditions

The dependent variable in each regression is $\widehat{w}_{i,t,k}$, the estimated experience weight on a lag k for a household i surveyed at time t in order to best explain that household's beliefs about the stock market return; see Eqn. (13). The experience effects estimates are produced using our richest Transformer-based specification. The explanatory variables are cosine similarity measures that proxy for the distance of a household i's macroeconomic beliefs expressed at time t to either the past experience that aligns with the estimated experience weight at time $t-k$, or to the future realization at time $t+1$. The macroeconomic beliefs used in the cosine similarity calculations are ternary variables for the direction of Unemployment, Business Conditions, and Interest Rates, provided by the respondent i at time t regarding a one-year ahead horizon; this is the same horizon over which stock market beliefs are elicited. The variables used to specify the objective macroeconomic conditions in each cosine similarity measure vary from columns (1) to (4): Diebold business conditions index (ADS) or Chicago Fed National Activity Index (CFNAI) for business conditions; and Fed Funds Rate (FFR) or 10-year Treasury yield minus FFR (T10) for interest rates.

	Experience Weight$_{i,t,k}$ (%)			
	(1)	(2)	(3)	(4)
Similarity of Beliefs to Past$_{i,t,k}$	0.2705***	0.2980***	0.3602***	0.3491***
	(0.0069)	(0.0075)	(0.0064)	(0.0068)
Similarity of Truth to Past$_{t,k}$	-0.3824***	-0.4553***	-0.1773***	-0.2129***
	(0.0127)	(0.0147)	(0.0109)	(0.0115)
Similarity calculations				
Business conditions measure	ADS	CFNAI	ADS	CFNAI
Interest rates measure	FFR	FFR	T10	T10
Fixed-effects				
Lag k × Age$_{i,t}$	✓	✓	✓	✓
Survey Date t	✓	✓	✓	✓
N	2,496,013	2,269,763	2,588,781	2,382,453
R^2	0.34782	0.34225	0.35343	0.34980

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

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 6: Experience Weights and Political Affiliation

The dependent variable in each regression is $\widehat{w}_{i,t,k}$, the estimated experience weight on a lag k for a household i surveyed at time t in order to best explain that household's beliefs about the stock market return; see Eqn. (13). The experience effects estimates are produced using our richest Transformer-based specification. The sample is split into respondents who identified as either Republicans or Democrats at t. The independent variable is an interaction term for whether the Michigan survey wave t was run immediately following the 2016 Presidential election results interacted with an indicator for whether the experience weight aligns with a lifetime experience that fell in a recession; baseline effects are subsumed by the fixed effects.

	Experience Weight$_{i,t,k}$ (%)			
Sub-sample:	Republicans		Democrats	
	(1)	(2)	(3)	(4)
$\mathbb{1}\{\text{Survey}_t$ POST Election$\}$	-1.035**	-0.7000**	0.3024*	0.2843*
$\times\ \mathbb{1}\{\text{Experience}_{t,k}$ Overlaps NBER Recession$\}$	(0.0236)	(0.0110)	(0.0370)	(0.0257)
Controls				
Baseline Effect	✓	✓	✓	✓
Demographic Characteristics	✓	✓	✓	✓
Macroeconomy Beliefs		✓		✓
Fixed-effects				
Lag $k \times$ Age$_{i,t}$	✓	✓	✓	✓
Survey Date t	✓	✓	✓	✓
N	8,080	3,878	7,588	3,881
R^2	0.76215	0.83095	0.86477	0.90109

*Standard errors are clustered at the level of the fixed-effects shown. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

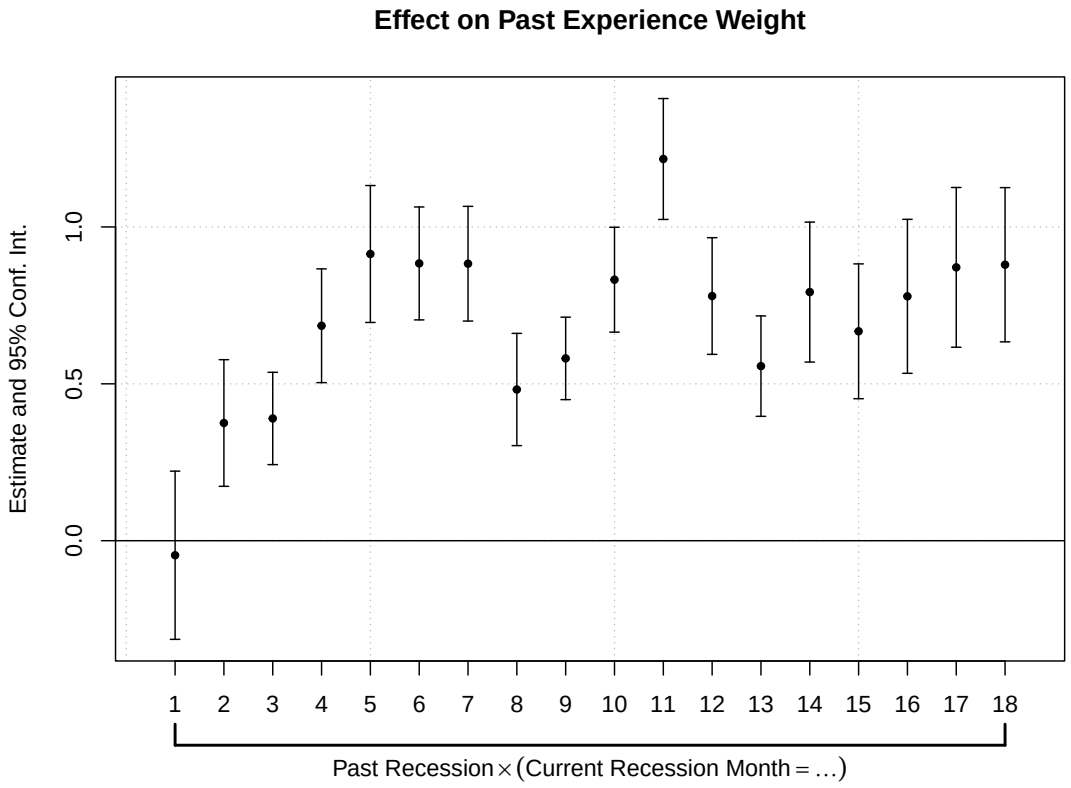


Figure 10: Weighting of Past Recessions During a Current Recession

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

5 Conclusion

This paper introduces a new source of subjective risk-adjusted expectations data about the market portfolio. We contribute a set of stylized facts about investors' subjective return expectations, 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) beliefs about the short-run evolution of the macroeconomy, but furthermore that other factors unrelated to the macroeconomy may play a role. We characterize systematic patterns that can inform the development of theoretical models of belief formation and 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.

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.

Bordalo, P., Gennaioli, N., Porta, R. L., and Shleifer, A. (2019). Diagnostic expectations and stock returns. *Journal of Finance*, *74*(6), 2839–2874.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2020). Memory, attention, and choice. *Quarterly Journal of Economics*, *135*(3), 1399–1442.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2022). Salience. *Annual Review of Economics*, *14*(1), 521–544.

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*.

Campbell, J. Y., and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, *21*(4), 1509–1531.

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.

Enke, B., Schwerter, F., and Zimmermann, F. (2024). Associative memory, beliefs and market interactions. *Journal of Financial Economics*, *157*, 103853.

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

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

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.

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.

Mullainathan, S. (2002). A memory-based model of bounded rationality. *Quarterly Journal of Economics*, *117*(3), 735–774.

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

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.

Wachter, J. A., and Kahana, M. J. (2024). A retrieved-context theory of financial decisions. *Quarterly Journal of Economics*, *139*(2), 1095–1147.

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 Return Expectation). *Assume that an individual i perceives the market's annual return to be conditionally Normally distributed, $r_{t+1} \sim \mathcal{N}(\mu_{it}, \sigma_{it})$. Then:*

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

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 $\operatorname{erf}(-z) = -\operatorname{erf}(z)$ to write

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

and then rearrange as

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

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

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

□

Internet Appendix

A Mapping Survey Responses to Subjective Return Expectations . . . 2
A.1 Probit Transformation . . . 2
A.2 Assumption of a Gaussian Return . . . 2
A.3 Alternative Proposition for Lognormality . . . 4
B Subjective Return Expectations Predict Household Demand . . . 5
B.1 Micro Demand: Equity Shares in the Cross-Section of Households . . . 5
B.2 Macro Demand: Direct and Indirect Equity Holdings by the Aggregate Household 5
B.3 Macro Demand: Fund Flows Into the Stock Market . . . 6
C A Comparison to Directly-Elicited Subjective Expected Returns . . . 8
D Cross-Sectional Variation in Subjective Return Expectations . . . 10
E Interpreting Subjective Return Expectations as Sharpe Ratios . . . 12
E.1 Assumption of a Zero Riskless Rate . . . 12
E.2 A Rational Benchmark for the Conditional Sharpe Ratio . . . 13
F Estimating the Malmendier and Nagel (2011) Benchmark . . . 17
G Estimating the IPCA Benchmark . . . 19
H Overview of the Transformer Framework . . . 22
References in Internet Appendix . . . 24

A Mapping Survey Responses to Subjective Return Expectations

A.1 Probit Transformation

Following directly from Proposition 1, one alternative interpretation of our subjective return expectation 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.

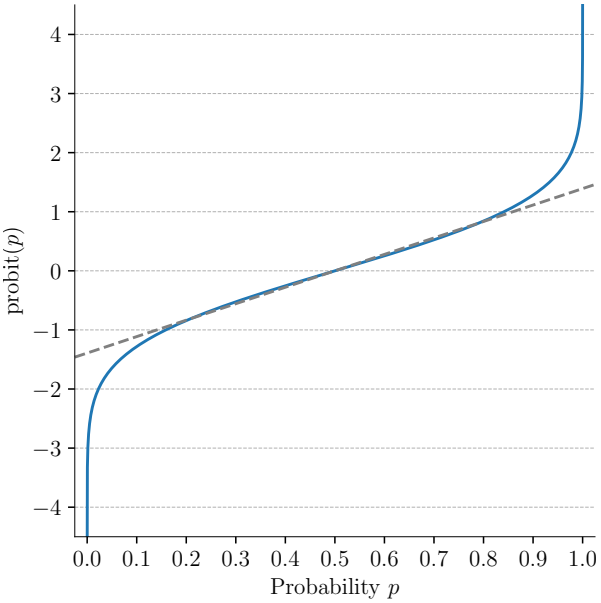


Figure IA.1: Probit transformation (y-axis) of example survey responses (x-axis).

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 a Gaussian distribution can well approximate investors' perceived return distribution. In making such an assumption, we build upon the existing subjective expectations literature. 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.

More importantly, statistical tests support our parametric assumption of normality: We directly test whether annual returns can be modeled as Gaussian by collecting annual returns for the S&P 500 index since 1925, and conducting three tests for Normality. Under each test, the Null hypothesis is that the returns are drawn from a Normal distribution. The results are presented in Table IA.1. For each test, p-values are high, and therefore the Gaussian Null cannot be rejected, supporting our assumption of Gaussian annual returns.

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 Alternative Proposition for Lognormality

We now show that an alternative lognormality assumption for returns also enables us to transform surveyed subjective return expectations to units of risk-adjusted return. The transformation is identical to that implied by Proposition 1, and the proof is also similar:

Proposition 2. *Assume that an individual i perceives the market's annual log-return to be conditionally Normally distributed,* $\log(1+r_{t+1}) \sim \mathcal{N}(\mu_{it}, \sigma_{it})$*. Then:*

$$\frac{\mu_{it}}{\sigma_{it}} \;=\; \text{probit}\big(\mathbb{P}_{it}(r_{t+1}>0)\big). \tag{9}$$

Proof. To simplify the notation, we omit all subscripts i. The proof proceeds identically to that of Proposition 1 (see Appendix), leading to the identity

$$\frac{\mu_t}{\sigma_t} = \text{probit}\big(\mathbb{P}_t\big(\log(1+r_{t+1})>0\big)\big). \tag{10}$$

As a final step, note that positive simple returns and positive log-gross returns coincide: $r_{t+1} > 0 \iff 1+r_{t+1} > 1 \iff \log(1+r_{t+1}) > 0$. Therefore, the elicited probabilities are also equal: $\mathbb{P}_t(r_{t+1}>0) = \mathbb{P}_t\big(\log(1+r_{t+1})>0\big)$, and we arrive at the same expression as in Proposition 1:

$$\frac{\mu_t}{\sigma_t} \;=\; \text{probit}\big(\mathbb{P}_t\big(\log(1+r_{t+1})>0\big)\big) \;=\; \text{probit}\big(\mathbb{P}_t(r_{t+1}>0)\big). \tag{11}$$

□

B Subjective Return Expectations Predict Household Demand

This section shows that households' subjective return expectations determine both individual and aggregate equity holdings by households.

B.1 Micro Demand: Equity Shares in the Cross-Section of Households

Table IA.2 examines the behavior of individual households i, and finds that each household's subjective return expectation predicts the share of wealth it holds in stocks. Each household's wealth ($) is proxied by the sum of its home value ($) and stock market investment ($). Time effects control for market conditions, including returns of the market portfolio.

Table IA.2: Equity Shares in the Cross-Section of Households Are Predicted by Subjective Expected Returns

Dependent Variable:	Wealth Fraction Invested in Stocks$_{i,t}$ (%)			
	(1)	(2)	(3)	(4)
Subjective Return Expectation$_{i,t}$	0.9147***	0.6522***	1.229***	0.7532***
	(0.1550)	(0.0752)	(0.1434)	(0.1188)
Intercept	56.70***		57.13***	
	(1.213)		(1.704)	
Controls				
Demographic Characteristics			✓	✓
Fixed-effects				
Survey Date t		✓		✓
N	62,614	62,614	27,271	27,271
R^2	0.00177	0.31996	0.02580	0.08617

Double-clustered std. errs. in parentheses

B.2 Macro Demand: Direct and Indirect Equity Holdings by the Aggregate Household

We now confirm that the previous micro-level predictive relationship also holds empirically at the aggregate level.

We are able to measure wealth shares held in equities – the fraction x in **??** – at an aggregate level, thanks to US national accounts data.[16] Figure IA.2 compares the timeseries variation in

[16]The results shown in this section make use of the FRED series "Households and Nonprofit Organizations;

the aggregate household's equity share with the timeseries variation in the subjective return expectation 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 over time: the regression results displayed in Table IA.3 show that the positive association survives when controlling for the level of the S&P 500, and is significant even when standard errors are clustered by year-month.

Table IA.3: Equity Share of the Aggregated Household Sector is Predicted by the Asset-Weighted Mean Subjective Return Expectation (SRE_t) of the Representative Household.

Dependent Variable:	Wealth Fraction Invested in $Stocks_t$ (%)			
	(1)	(2)	(3)	(4)
Constant	22.72^{***}	22.72^{***}	13.80^{***}	13.80^{***}
	(0.5488)	(0.7737)	(0.4186)	(0.5859)
SRE_t	3.567^{***}	3.567^{***}	2.453^{***}	2.453^{***}
	(1.113)	(1.252)	(0.6995)	(0.8382)
S&P 500 $Level_t$	0.0042^{***}	0.0042^{***}	0.0036^{***}	0.0036^{***}
	(0.0003)	(0.0003)	(0.0002)	(0.0003)
Wealth	Financial	Financial	Total	Total
Std. Errs.	Clustered by Time	Newey-West	Clustered by Time	Newey-West
N	79	79	79	79
R^2	0.83780	0.83780	0.86760	0.86760

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

B.3 Macro Demand: Fund Flows Into the Stock Market

As discussed by Greenwood and Shleifer (2014), *changes* to subjective return expectations 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.

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.

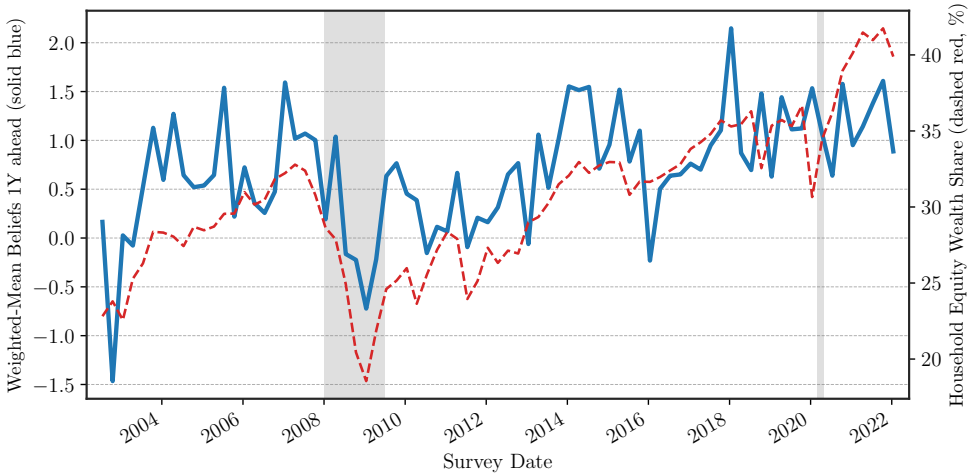


Figure IA.2: Aggregated Subjective Return Expectations for Year Ahead vs. Households' Share of Wealth Held in Equities

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.

We measure the representative investor's subjective return expectation simply by the mean of individual investors' values at each point in time. Aggregate flows into the stock market are calculated using Investment Company Institute data on fund flows:[17]

$$\text{Flows}_{t\to t+1} = \frac{\$\text{ Domestic Equity Fund Flows}_{t\to t+1}}{\$\text{ Domestic Equity Market Capitalization}_t} \times 100\%. \tag{12}$$

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.

[17]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.

C A Comparison to Directly-Elicited Subjective Expected Returns

In this appendix, we compare our subjective expectations measure to that of Amromin and Sharpe (2014) and Greenwood and Shleifer (2014), while retaining our focus on the Michigan survey dataset. Our own approach to measuring subjective expectations – described in Section 2.1 – has the advantage of relying on a survey question with a long timeseries and wide cross-section that has been elicited throughout our two-decade sample period. An empirical comparison shows that the subjective expectations measures are similar, yet ours is a more accurate predictor of investors' stock market risk-taking at the level of the individual household. Coupled with our large dataset, this result justifies the usefulness of our measure.

We begin by replicating the Michigan survey-based expected returns that were collected by Amromin and Sharpe (2014) and Greenwood and Shleifer (2014). These are based on questions that *directly* ask respondents to provide their expected returns for the stock market over either a 2-year or 3-year horizon (the exact wording of the question varies between different Michigan survey waves). This dataset begins in September 2000 and ends in October 2005. The frequency varies from semi-annual to monthly; we found 22 waves in total. These microdata are available from the ICPSR repository, in which the expected return variable is named V674, V542, or V562 (depending on the exact survey wave).

Table IA.4 uses two subjective expected return variables to explain stock market risk-taking at the household level. To make our comparison, in this table we focus only on survey waves where a directly-elicited expected return variable was elicited, and use this variable to predict stock market risk-taking together with the main measure we use in our own paper (which we label the "extracted" expected return). We find in column (1) that the sign of the directly-elicited expected return coefficient is positive (consistent with basic portfolio choice intuition), but that the magnitude is low and thus the result cannot significantly be distinguished from zero using this small sample. On the other hand, our extracted measure is a significant predictor in column (2), even with a smaller sample. Furthermore, this result survives when including both expected return measures in column (3). We conclude from this exercise that our extracted measure of subjective return expectations is a useful and strong predictor of household risk-taking, including in the sample focused on by Amromin and Sharpe (2014).

Table IA.4: Individual Household Stock Market Risk-Taking vs. Subjective Expected Returns

Dependent Variable:	Log Stock Holdings$_{i,t}$		
	(1)	(2)	(3)
Directly-elicited return expectation$_{i,t}$	0.0007		-0.0042
	(0.0052)		(0.0081)
Our extracted return expectation$_{i,t}$		0.2140***	0.1317**
		(0.0465)	(0.0559)
Year-month fixed-effects	✓	✓	✓
N	2,840	1,822	1,478
R^2	0.01457	0.01632	0.01697

Standard errors are clustered by year-month, and displayed in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

D Cross-Sectional Variation in Subjective Return Expectations

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.3 plots various timeseries of the cross-sectional means of respondents' subjective return expectations. Younger investors, male investors, and investors who support the current President's party are most optimistic when compared to older investors, female investors, and those who oppose the current President's party.[18] Moreover, their expectations align with their investments in the stock market: investors with above-median investments in stocks are more optimistic 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.

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

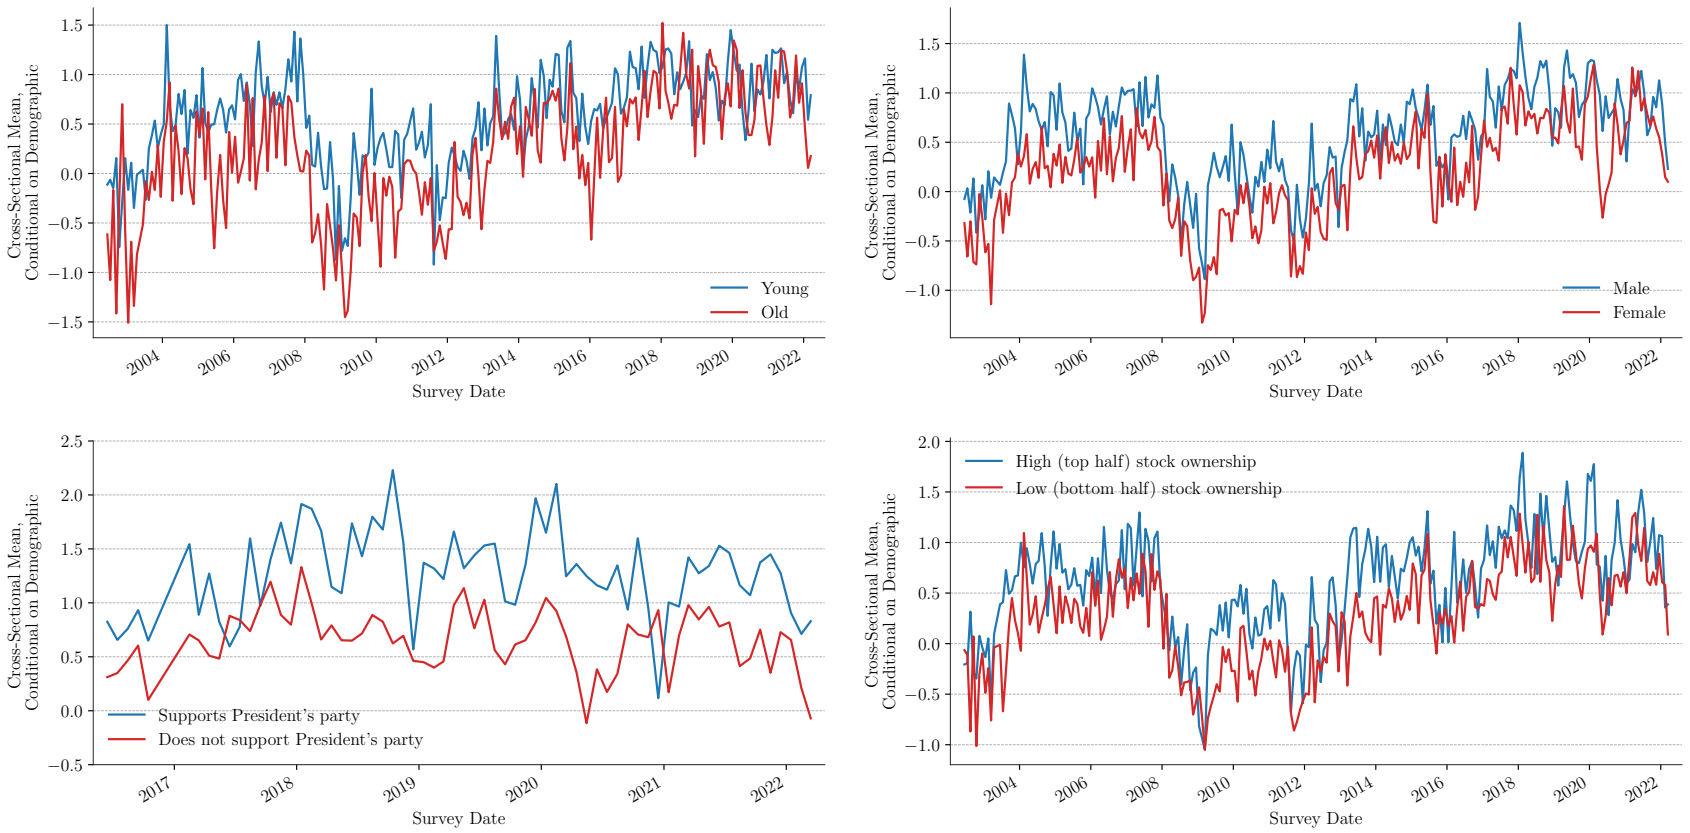


Figure IA.3: Mean Survey Subjective Sharpe Ratio Conditional on Various Attributes of Respondents

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.

E Interpreting Subjective Return Expectations as Sharpe Ratios

The subjective return expectations that Proposition 1 allows us to extract are in units of risk-adjusted return, much like a Sharpe ratio. The main difference relates to the role of the risk-free rate. Recall that the Sharpe ratio for a single asset is defined as

$$\frac{\mu_t}{\sigma_t} - \frac{r_f}{\sigma_t}, \tag{13}$$

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.

Our subjective return expectation measure is thus a valid proxy for Sharpe ratios whenever the household can only access (or perceive) a zero risk-free rate. This is not an unreasonable assumption when studying households: cash holdings in a retail current account facility insured by the FDIC are likely to provide a negligible return.

E.1 Assumption of a Zero Riskless Rate

We now quantify the potential measurement error of an incorrect assumption of a zero risk-free rate in the case of interpreting our subjective return expectations as subjective Sharpe ratios.

Table IA.5 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. To understand the consequences of an agent who perceives a non-zero riskless rate, Table IA.5 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.4 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.5: 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

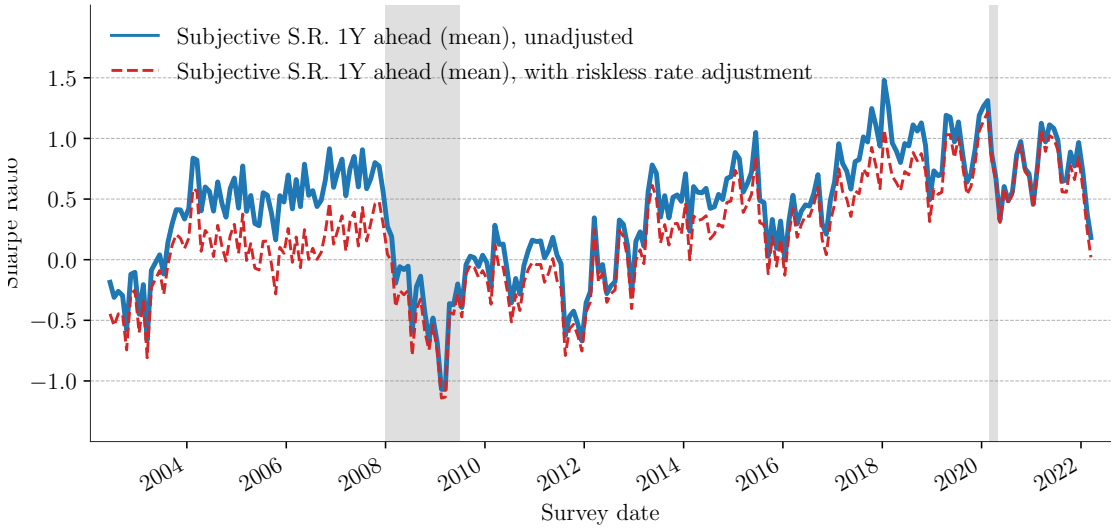


Figure IA.4: Mean Survey Subjective Sharpe Ratio With Riskless Rate Adjustment

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.

E.2 A Rational Benchmark for the Conditional Sharpe Ratio

Continuing with our interpretation of subjective return expectations as subjective Sharpe Ratios, we now highlight 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.

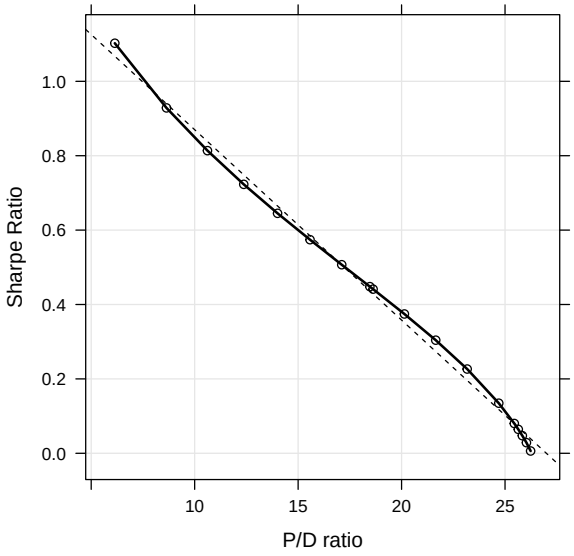


Figure IA.5: Rational model-implied relationship between the Sharpe Ratio and P/D ratio

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.

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:[19] 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.

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

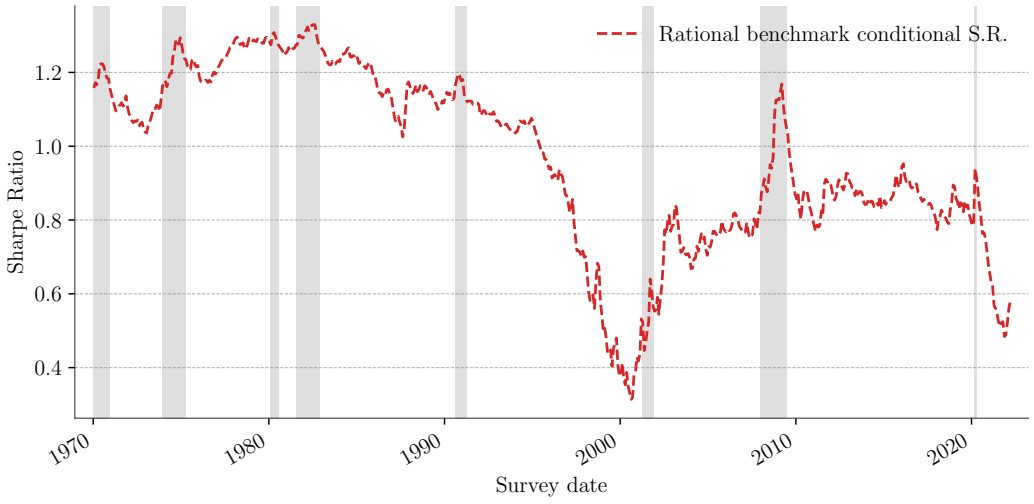


Figure IA.6: Empirical benchmark for the conditional Sharpe Ratio

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.

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 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 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

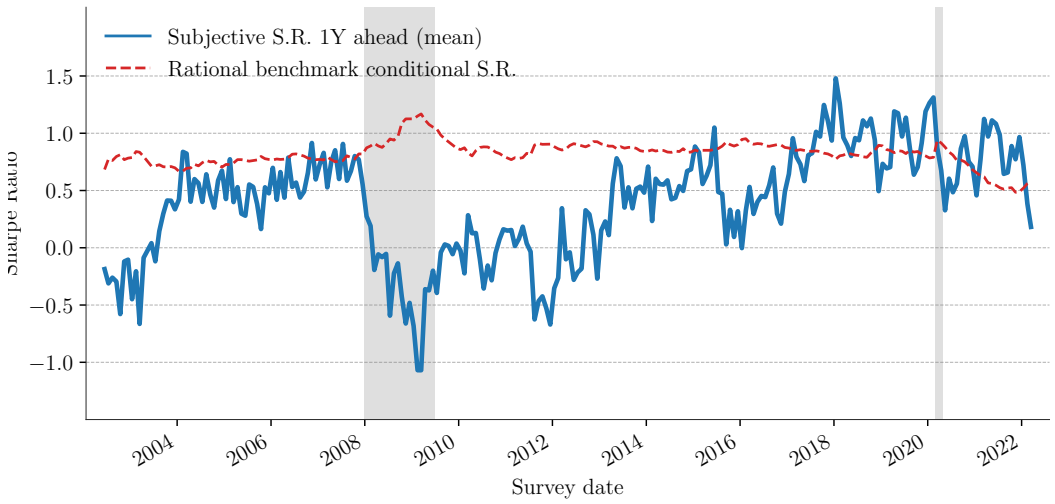


Figure IA.7: Mean Survey Subjective Sharpe Ratios for Year Ahead vs. a Rational Benchmark

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.

and fall in good times.

F 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 (3) and (4) on our dataset of risk-adjusted returns, choosing λ so that subjects' past experiences best explain their current expectations. Historical risk-adjusted returns RAR_{t-k} are based on the annualized ratio of S&P 500 realized returns divided by realized volatilities within each experienced year (similarly to our machine learning-based model).

The goal of the optimization is to pick the value of the decay parameter λ in Eqn. (4). We specify an objective of minimizing the Mean Squared Error (and thus maximizing the R^2) between the values predicted by the exponential weighting procedure, in Eqn. (3), as compared to the subjective return expectations 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.8 visualizes the objective function minimized over the grid search procedure.

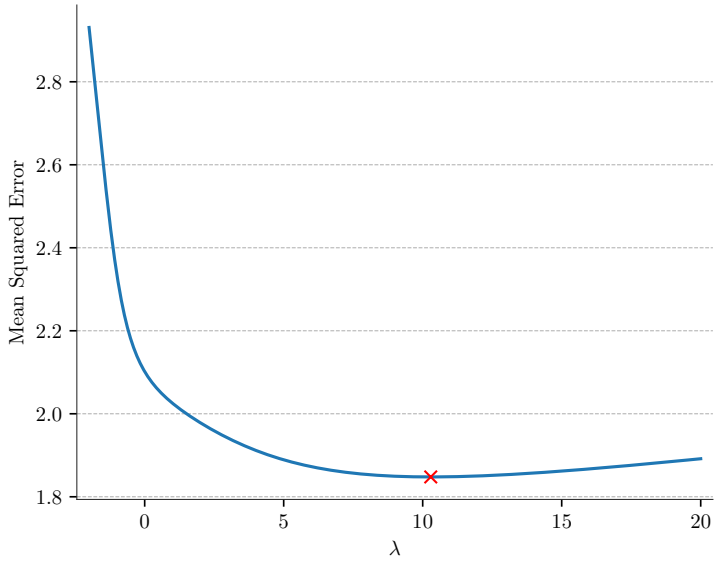


Figure IA.8: Objective function during the grid search procedure for estimating the λ parameter of the Malmendier and Nagel (2011) exponential weighting function.

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 expec-

tations about the future stock market return. The experience weights that correspond to this value of λ are illustrated in Figure IA.9.

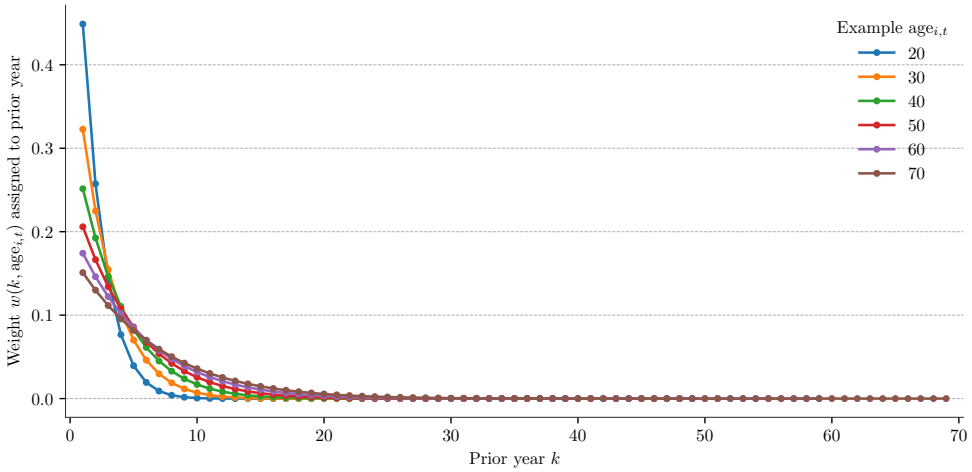


Figure IA.9: Weights $w(k, \text{age}_{i,t})$ assigned by the Malmendier and Nagel (2011) experience effects model to prior years k for various example ages of respondents. The optimal value $\widehat{\lambda} = 10.39$ is used.

G 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:

$$\text{SSR}_{i,t} = \alpha_{i,t} + \beta_{i,t} f_{t+1} + \varepsilon_{i,t+1}, \qquad \beta_{i,t} = [\mathbf{d}, \mathbf{b}]'_{i,t} \mathbf{\Gamma}^{AC}_{\beta}; \quad \alpha_{i,t} = [\mathbf{d}, \mathbf{b}]'_{i,t} \mathbf{\Gamma}^{AC}_{\alpha} \tag{14}$$

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 (14), 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).

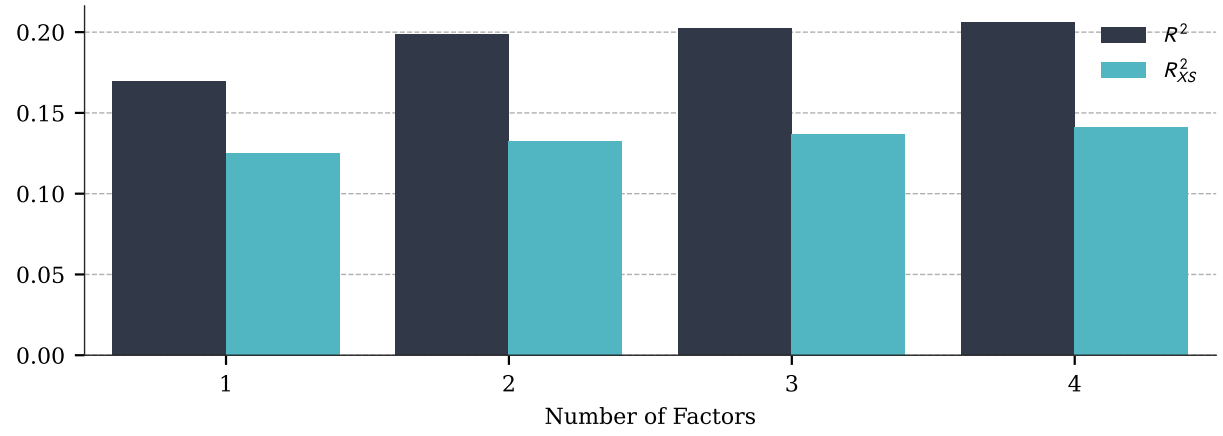


Figure IA.10: R^2 of IPCA with 1 to 4 Factors

The figure shows the R^2 and R^2_{XS} of IPCA models with 1 to 4 factors, following the specification in Equation (14) with $\alpha_{i,t} = 0$.

When estimating Equation (14), 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.6: $\mathbf{\Gamma}_\beta$ Estimates for the IPCA Factor Model

This table shows the $\mathbf{\Gamma}_\beta$ estimates for a one-factor IPCA model following Equation (14) 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. (17). Figure IA.10 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.6 presents estimates for the $\mathbf{\Gamma}_\beta$ vector of Equation (14) 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.3), 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.

H Overview of the Transformer Framework

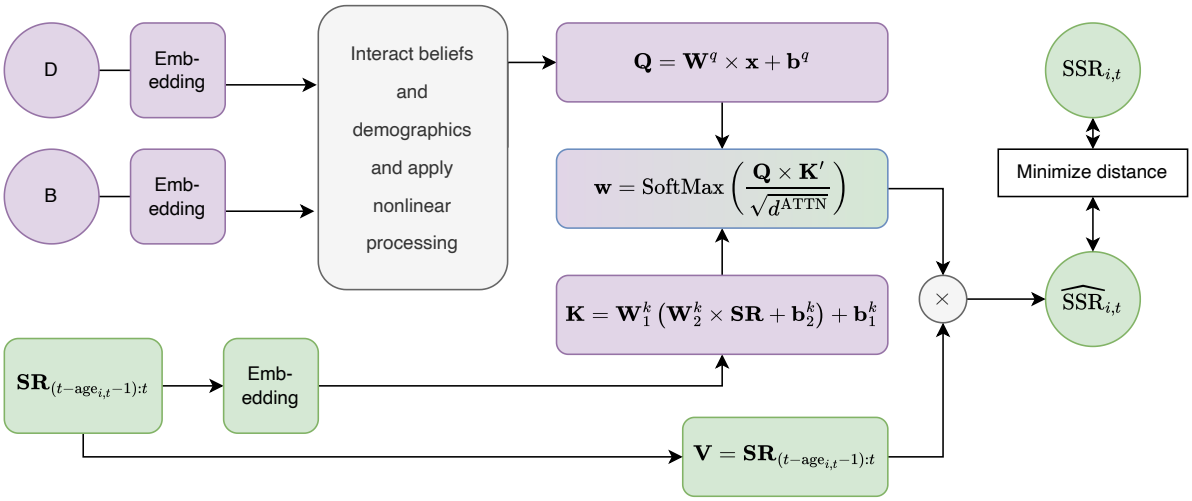


Figure IA.11: Model Setup

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:

$$\mathbf{x}^{\text{emb}} = \mathbf{x} \times \mathbf{W}, \tag{15}$$

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 (9), 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.