Expectations and Attention to Experience

Heiner Beckmeyer* and Ahmed Guecioueur†

August 30, 2023

Abstract

We measure which past experiences determine investors' expectations about the market's future Sharpe Ratio. We first introduce a simple method to recover individuals' subjective Sharpe Ratios from a rich source of survey microdata. These subjective expectations are procyclical, extrapolative, cross-sectionally correlated with individual demographic characteristics, and well explained by a low-dimensional latent factor structure. We then use a customized machine learning estimation technology to estimate an economic model of experience effects that generalizes the lifetime weighting scheme of Malmendier and Nagel (2011). The model includes the influence of demographic characteristics in how past experiences determine individuals' future expectations, and succeeds in explaining a larger fraction of survey belief heterogeneity. We find that households' aggregate wealth share held in equities is strongly correlated with the share of investors who have experienced positive Sharpe Ratios (as measured by our model), thus confirming that experience effects drive investor flows. We also contribute a new set of facts on the role of demographic characteristics and the outsized influence of past recessions on how individuals learn from their experiences.

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

*University of Münster. Email: heiner.beckmeyer@wiwi.uni-muenster.de
†INSEAD. Email: ahmed.guecioueur@insead.edu

First version: February 16, 2023. We would like to thank Stefano Cassella, Josefina Cenzón, Theodoros Evgeniou, Naveen Gondhi, Maurizio Montone (discussant), Rodolfo Prieto, Joshua D. Rauh, Trevor Young (discussant), as well as participants at the 2023 China International Conference in Finance, the 2023 Financial Management Association European Conference, the 2023 Financial Markets and Corporate Governance Conference and the University of Münster for their helpful comments. Guecioueur gratefully acknowledges the support of the Henry Grunfeld Foundation.

1 Introduction

Beliefs play a crucial role in economic models. Financial economists and macroeconomists alike have thus sought to empirically understand the mechanisms by which individuals form their heterogeneous expectations. The evidence suggests individuals' very recent experiences are important (Greenwood and Shleifer, 2014), as are their lifetime experiences (Malmendier and Nagel, 2011, 2016), their demographic attributes (Dominitz and Manski, 2007; Das, Kuhnen, and Nagel, 2020) and even unobservable individual characteristics (Giglio, Maggiori, Stroebel, and Utkus, 2021). At present, there is little understanding of how individuals' characteristics and experiences interact; in the words of Giglio, Maggiori, Stroebel, and Utkus (2021), "a likely explanation is that individual beliefs reflect a combination of many demographic characteristics and experiences, without a single dominant explanation."

This paper examines the interaction between these determinants of individuals' subjective expectations. We focus on the mechanism of experience effects and show that demographic characteristics play a key role in how past experiences shape individuals' expectations about the future. An influential study by Malmendier and Nagel (2011) shows that incorporating individuals' lifetime experiences into an adaptive expectation formation procedure (Evans and Honkapohja, 2001) generates differences in beliefs, but does not include a role for demographic characteristics besides age. To examine how demographics and experience effects interact in general, we develop an empirical approach that measures how individual investors with different demographic attributes form expectations differently based on past experiences; for example, we capture the potential for male vs. female investors to pay different levels of attention to different parts of their experiences, even if they share the exact same lifetime experience. Our method explains a large fraction of cross-sectional heterogeneity in subjective expectations, thereby supporting a unified mechanism for experience effects in which other variables such as demographic attributes play an important role, without being direct determinants of individuals' subjective 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 recover investors' subjective perceptions of the Sharpe Ratio of the market from a long-running and well-known source of survey microdata – the Michigan Surveys of Consumers. Classic portfolio choice frameworks suggest two reasons an investor's perceived Sharpe Ratio is an important quantity to measure: first, it proportionally determines the size of her investment; and second, it is a sufficient statistic for both the perceived first and second moments of the return distribution of the risky asset. Our approach therefore unlocks a better understanding of investors' subjective expectations over and above their perceived first moments, as Brunnermeier et al. (2021, pp. 2128-9) have called

for.

Using two decades of monthly survey samples, we document four empirical facts about investors' subjective Sharpe Ratios that complement the existing literature. First, subjective Sharpe Ratios are procyclical, in contrast to the conditional Sharpe Ratios implied by a rational model such as Campbell and Cochrane (1999)'s. Second, they are extrapolative (Greenwood and Shleifer, 2014), and depend on the full history of household experiences, with more recent experiences exerting a stronger influence (Malmendier and Nagel, 2011). Third, we provide evidence of cross-sectional variation of investors' subjective Sharpe Ratios with their individual demographic characteristics (Giglio, Maggiori, Stroebel, and Utkus, 2021). And fourth, we show that a factor model in which households have time-varying exposures to a single latent factor can explain a large fraction of across-household variation in subjective Sharpe Ratios. We estimate this factor model using Instrumented PCA (Kelly, Pruitt, and Su, 2020).

We then study how individual investors' past lifetime experiences interact with their demographic characteristics to determine their subjective expectations about the Sharpe Ratio of the market. We contribute a new machine learning-based methodology for estimating individual experience effects. Building upon the classic approach by Malmendier and Nagel (2011), we allow demographic attributes to flexibly instrument for the fraction of attention that individuals pay to their past experiences. At the same time, we loosen the assumption of exponentially-weighted past experiences that has so far been used in the literature. Our goals thus require us to estimate an empirical model that permits a flexible functional form – while taking sequential dependencies among time series of experienced realizations into account – and that explicitly captures restrictions on historical data based on the modeled individuals' lifetime experiences, allows individuals' characteristics to enter into their experience weights, and that admits a clear interpretation of estimated parameters to the econometrician. Our estimation technology is a customized attention-based transformer neural network (Vaswani et al., 2017) that meets this challenging set of criteria.

By design, our attention-based transformer framework estimates individuals' expectation formation processes as a linear combination of attention weights and lifetime experiences, with attention weights (nonlinearly) instrumented by demographic characteristics. Since the estimated attention weights are non-negative and sum to unity, they are directly interpretable as lifetime experience weights, akin to the model of Malmendier and Nagel (2011); the structure of our neural network architecture thus embeds an economic mechanism. We interrogate the weights estimated by the model and show empirically that individuals pay greater attention to recent experiences, consistent with the prior findings of Malmendier and Nagel (2011, 2016) and Greenwood and Shleifer (2014). Remarkably, the decay in individuals' experience weights is quite smooth and somewhat similar in form to an exponential for the most recent lags, although

no functional form was specified at the outset. Importantly, our estimated experience weights exhibit non-monotonicities: we find that investors pay extra attention to market conditions during periods of recession that they have personally experienced, consistent with the notion of belief scarring. Our results also suggest that investors are attentive to multiple moments of the return distribution. Finally, we characterize heterogeneity in the relative extent to which individuals are influenced by their lifetime experiences based on their differing demographics, exploiting the fact that estimated attention weights vary with demographic characteristics other than age.

Our model of experience effects is more effective in explaining individuals' subjective expectations than traditional models of exponential weighting. Our richest specification explains almost a fifth of both aggregate and cross-sectional variation in investor expectations. These results strengthen the case for experience effects as an explanation for expectation formation, as our findings indicate that experience effects explain a much larger share of the belief heterogeneity seen in survey data when additional demographic characteristics are used to specify heterogeneity in experience weights.

Furthermore, our most comprehensive model of experience effects explains a similar fraction of variation in subjective expectations as a statistical latent factor model. We compare the two models' predictions and find them to be highly correlated. Our experience effects models pick up similar information as reduced-form statistical methods, while simultaneously capturing an important expectation formation mechanism. The benefit of our approach is that it enables us to measure the relative importance of prior experiences and their interaction with demographic attributes by analyzing estimated experience weights; these weights have an economic interpretation.

We next turn to the question of whether experience effects can explain the household wealth shares held in equities, as implied by a simple portfolio choice framework. We first find that an aggregated subjective Sharpe Ratio correlates strongly with the share of wealth held in equities (both directly & indirectly) by the representative household. As an empirical corollary, changes to expectations correlate with aggregate flows to the market. We then measure the aggregate effect of individual experience effects using our estimated model: by predicting each individual's subjective Sharpe Ratio, we extract a measure of the aggregate share of investor wealth held by optimistic investors who have a positive lifetime-experienced Sharpe Ratio. This measure is strongly correlated with the household aggregate wealth share held in equities; in fact, it is even more strongly correlated than aggregated subjective Sharpe Ratios as a whole. While individual heterogeneity makes explaining expectations challenging (Giglio, Maggiori, Stroebel, and Utkus, 2021), the portion of subjective expectations that we *can* explain by experience effects is an important driver of aggregate investor holdings in the

stock market. Given that flows into the market have a strong effect on the market valuation (Gabaix and Koijen, 2021), these results suggest that experience effects may be an important driver of asset prices.

Interpreting investors' estimated experience weights provides some additional insights on the effect of recessions on investor behavior. While individuals are influenced both by recent experiences and the more distant past, we also find that experience weights are not simply exponential, but that past recessions exert an outsized influence on how expectations are formed today. Furthermore, individuals' tendencies to extrapolate the recent past is exacerbated during current recession periods. This result suggests that extrapolative expectations are asymmetric according to current economic conditions, rendering subjective expectations even more procyclical in bad than in good times. Our findings thus pose a challenge for constant gain learning representative agent models (Nagel and Xu, 2022) that may not capture these effects; further work is therefore necessary to understand the aggregate asset pricing implications of the micro-level phenomena that we uncover.

In our final analysis, we explore whether stock market experiences can also distinguish between households' beliefs about other macroeconomic state variables. We find this is indeed the case for a number of belief variables, suggesting our measure of stock market experience effects is informative for understanding expectation formation more broadly.

The remainder of this paper proceeds as follows. We first review the relevant literature. Section 2 describes our survey microdata and method for recovering subjective Sharpe Ratios, and presents a number of key stylized facts. Section 3 explains our machine learning-based procedure for estimating experience effects. Section 4 presents estimation results and characterizes cross-sectional heterogeneity in experience effects. Section 5 shows that aggregate subjective Sharpe Ratios and experience effects drive the household equity share. Section 6 investigates to what extent beliefs are integrated. Section 7 concludes. The Appendix contains all proofs, additional results, and further methodological details.

Related Literature

We contribute to a recent literature in financial economics that studies the determinants of individuals' return expectations using survey data (Malmendier and Nagel, 2011; Dominitz and Manski, 2011; Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014; Giglio, Maggiori, Stroebel, and Utkus, 2021).[1] Like Malmendier and Nagel (2011, 2016), we focus on experience

[1]Additionally, work by Da, Huang, and Jin (2021) and Andonov and Rauh (2022) documents that even professional investors' beliefs about returns are extrapolative. Other related studies analyze the potential roles played by memory (Mullainathan, 2002; Wachter and Kahana, 2019; Bordalo, Gennaioli, and Shleifer, 2020) and affect (Kuhnen and Knutson, 2011; Bodoh-Creed, 2020) in belief formation.

effects, and extend this mechanism by allowing for experiences to influence the expectation formation process differently, based on individuals' demographic attributes.

Our standpoint is that investors' demographic characteristics matter for asset pricing. An extensive literature has examined the implications of age, especially, for asset prices (e.g. Abel (2001) and Poterba (2001)) and return predictability (Cassella and Gulen, 2018; Geanakoplos, Magill, and Quinzii, 2004; Goyal, 2004; Favero, Gozluklu, and Tamoni, 2011). We contribute by showing that other demographic characteristics contribute to the mechanism of experience effects and thereby impact how investors incorporate past realizations into their expectations about the future.

We build on a small literature studying investors' beliefs about volatility (Lochstoer and Muir, 2022; Sheikh, 1993). Our measure of investor expectations captures both the first and second moments of the return distribution and can be mapped to a subjective Sharpe Ratio, allowing us to study a key quantity for investors' portfolio choices.

Our paper joins an extensive literature on household finance (Gomes, Haliassos, and Ramadorai, 2021). We use survey data collected from households and extract insights into individual respondents' expectation formation processes using machine learning techniques.

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), firm earnings (van Binsbergen, Han, and Lopez-Lira, 2022; de Silva and Thesmar, 2021) or stock prices (Cao, Jiang, Wang, and Yang, 2021) by comparing them to machine learning forecasts — however, we stand apart in two key ways. First, we estimate a machine learning model upon the forecasts produced by individuals, rather than on the actual realizations: our use of machine learning techniques is to directly predict and understand human behavior, rather than set up a benchmark for comparison. And second, we investigate determinants of the expectation formation of ordinary households, rather than of professional forecasters.

The estimation technology that we employ is a form of neural network that is capable of approximating any functional form (Cybenko, 1989) that might underlie the formation of experience weights, thus freeing us from parametric and functional form assumptions. Moreover, the attention-based transformer architecture that we employ has proven to be adept at capturing sequential relationships (for example, in the machine translation domain, Vaswani et al., 2017), and we exploit this architecture to capture the sequential nature of historic time-series of experienced realizations.

We add to a small but burgeoning literature in financial economics that applies attention-based neural networks for empirical studies: Barbon (2019) studies the determinants of market microstructure effects, and Beckmeyer and Wiedemann (2022) analyze the effects of missing firm-level characteristics. To our knowledge, we are the first to exploit attention-based neural

network architectures to study the behavior of economic agents in any field of economics. Non-attention-based neural networks have been used to forecast the stock market return (Kelly, Malamud, and Zhou, 2022) and realized inflation (Nakamura, 2005) based on macroeconomic state variables.

2 Expectations About the Market Sharpe Ratio

2.1 Recovering Subjective Sharpe Ratios from Survey Microdata

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

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

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

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

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

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

Under the above set of assumptions, which we discuss in more detail shortly, there is thus a simple one-to-one mapping between the elicited quantity $\mathbb{P}_{it}(r_{t+1} > 0)$ and the investor's perceived Sharpe Ratio of the market portfolio. An investor's beliefs about the Sharpe Ratio are a key determinant of her portfolio choice decision, under a range of horizons and preference structures: in a choice between a risky and a riskless asset, the investor's optimal portfolio share

of the risky asset is typically directly proportional to her perceived Sharpe Ratio of the risky asset. For example, an investor with mean-variance preferences (with risk aversion parameter γ_i) who chooses which fraction x of her wealth to allocate to the risky market portfolio returning r_{t+1} seeks to maximize

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

and therefore chooses to allocate a fraction of her wealth

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

to the risky market portfolio, in direct proportion to her perception of the Sharpe Ratio $\frac{\mu_{it}}{\sigma_{it}}$.

Proposition 1 allows us to recover a crucial economic quantity. A few remarks on the underlying assumptions are now in order: first, the assumption of a zero riskless rate being available to (or perceived by) the investor is realistic in our setting of individual investors; for example, cash holdings in a retail current account facility insured by the FDIC are likely to provide a negligible return. Table 1 in Appendix B 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. Second, we build upon the existing subjective expectations literature by using the parametric assumption of normality to link survey responses to our economic quantity of interest. Notably, Dominitz and Manski (2007, pp. 376) and Lochstoer and Muir (2022, pp. 1060) make similar assumptions to measure the first and second moments (respectively) of investors' perceived distributions of the market return. Ben-David, Graham, and Harvey (2013, pp. 1553) also rely on this assumption.

Furthermore, while the probit function that follows from the conditional normality assumption is nonlinear, it is well-approximated by a linear mapping for the majority of survey responses (illustrated by Figure 3 in Appendix B). This suggests our paper's main results can alternatively be interpreted as linking together experience effects and investors' 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.

2.2 Subjective Sharpe Ratios are Procyclical

Figure 1 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

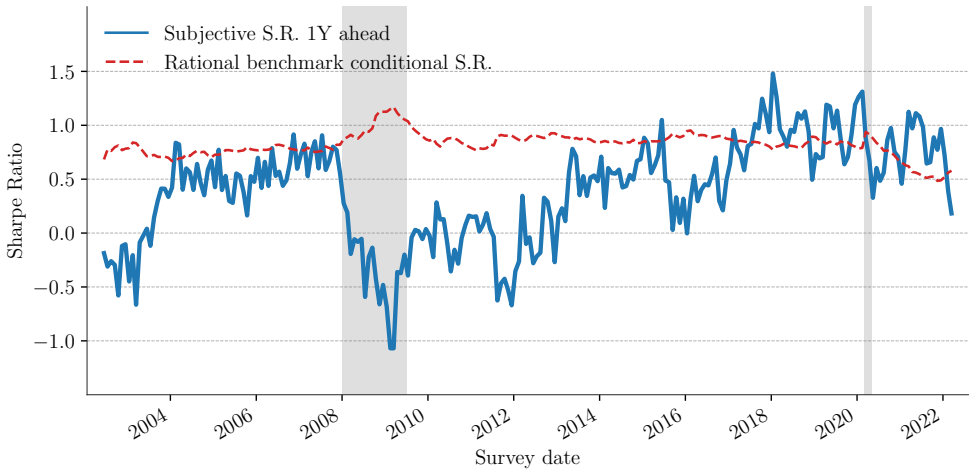


Figure 1: Mean Survey Subjective Sharpe Ratios for Year Ahead vs. a Rational Benchmark

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.

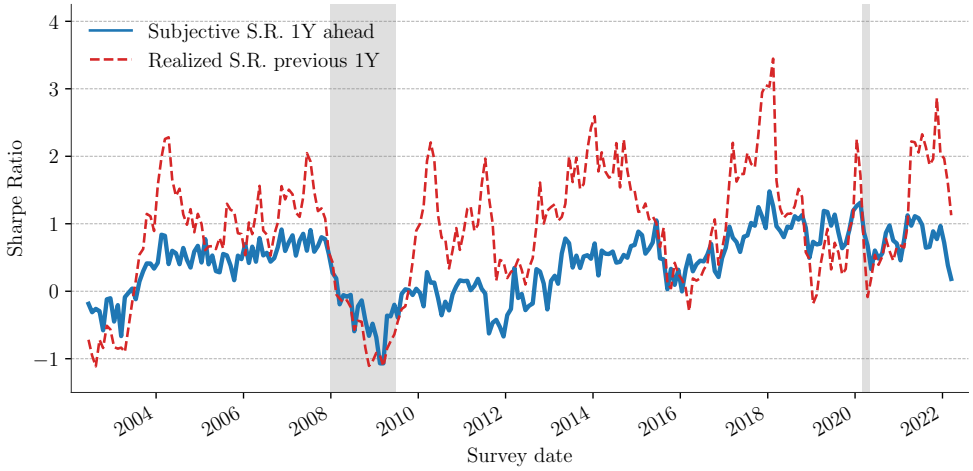


Figure 2: Mean Survey Subjective Sharpe Ratios for Year Ahead vs. Previous Year Realizations

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

similar procyclical patterns in subjective expectations collected from other surveys.

It is worth highlighting that time-varying Sharpe Ratios can be produced by rational asset pricing models, but these tend to be countercyclical in order to match the equity premium. Figure 1 also overlays (in red) the conditional Sharpe Ratios implied by Campbell and Cochrane

(1999)'s benchmark asset pricing model, which by contrast increase in bad times and fall in good times. Appendix C provides further details of the benchmark calculation.

2.3 Subjective Sharpe Ratios are Extrapolative

We now provide evidence that survey respondents' subjective expectations are driven by experience effects. As a first step, Figure 2 plots the aggregate subjective Sharpe Ratio (in blue) and this time overlays the prior year's realized Sharpe Ratio (in red). The figure visually illustrates that the two series are positively correlated: expectations about the future depend (positively) on recent past realizations; the magnitude of the correlation is 0.11, which is significant at the 1% level.

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

To therefore understand the relative contributions of both recent and distant past realizations, we follow the approach of Malmendier and Nagel (2011) and model agents' experiences as an exponentially weighted average of all their past lifetime experiences:

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

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

Each individual i surveyed at time t is assumed to hold a subjective expectation that is a weighted sum of her past experienced Sharpe Ratios SR_{t-k} using positive scalar weights $w_{i,t}(k, \lambda)$ that 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.[2]

We estimate this adaptive expectation formation model on our dataset of Subjective Sharpe

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

Ratios, choosing λ so that subjects' past experiences best explain their current expectations. Subjects' experienced Sharpe Ratios SR_{t-k} are specified as the lagged annual S&P 500 Sharpe Ratios that were realized during their lifetimes; Appendix D provides further details. We estimate $\widehat{\lambda} = 10.39$, suggesting that subjects place a relatively high weight on their recent past experiences (compared to their early lifetime experiences) when forming expectations about the future Sharpe Ratio of the stock market. The experience weights that correspond to this value of λ are illustrated in Figure 3.

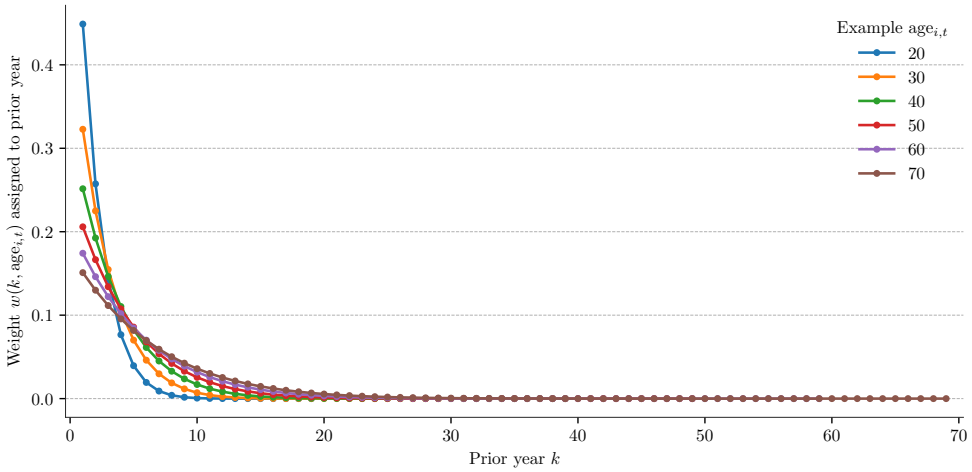


Figure 3: Weights $w(k, \text{age}_{i,t})$ assigned by the Malmendier and Nagel (2011) experience effects model to prior years k for various example ages of respondents. The optimal value $\widehat{\lambda} = 10.39$ is used.

2.4 Subjective Sharpe Ratios Vary Cross-sectionally with Demographics

2.4.1 Additional Microdata Variables

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

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

value.

In addition, the Michigan surveys record respondents' beliefs concerning general economic conditions, and other economic variables besides the stock market. These can be used to infer respondents' general levels of optimism or pessimism. We therefore make use of the following belief-related variables: probability of an increase in personal income, direction of business conditions, direction of interest rates, direction of unemployment, and changes in local house prices, all elicited for a horizon of one year.

Our main data processing step is to include only survey responses in which the respondent provides her subjective probability of a positive market return over the next year – our main variable of interest – and in which the respondent reported that she or a member of her household own any stocks.[3] After this filtering step, our survey microdata comprises 81,879 responses. Table 1 presents summaries of the Michigan survey variables used in our study.

2.4.2 Expectations and Demographics

The prior literature finds that demographic attributes have explanatory power for the cross-section of subjective expectations. The same is true in our setting. Figure 4 plots various timeseries of the cross-sectional means of respondents' subjective Sharpe Ratios. Younger investors, male investors, and investors who support the current President's party all expect higher subjective Sharpe Ratios for the market in the year ahead, when compared to older investors, female investors, and those who oppose the current President's party.[4] Moreover, their expectations align with their investments in the stock market: investors with above-median investments in stocks expect a higher Sharpe Ratio than those with below-median investments in stocks. Our goal in this paper is to understand how demographic variables interact with experiences, for which we provide a new methodology in Section 3.

2.5 Factor Structure of Subjective Sharpe Ratios

To better understand the cross-section of subjective Sharpe Ratios, we next investigate whether they share a common factor structure across households.[5] Since the Michigan Surveys of Consumers are repeated cross-sectional surveys, which do not allow us to track the same

[3]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.

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

[5]Kamdar and Ray (2022) show that multiple expectations *of one household* have a strong factor structure. In contrast, we wish to understand the factor structure of expectations *across households*.

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

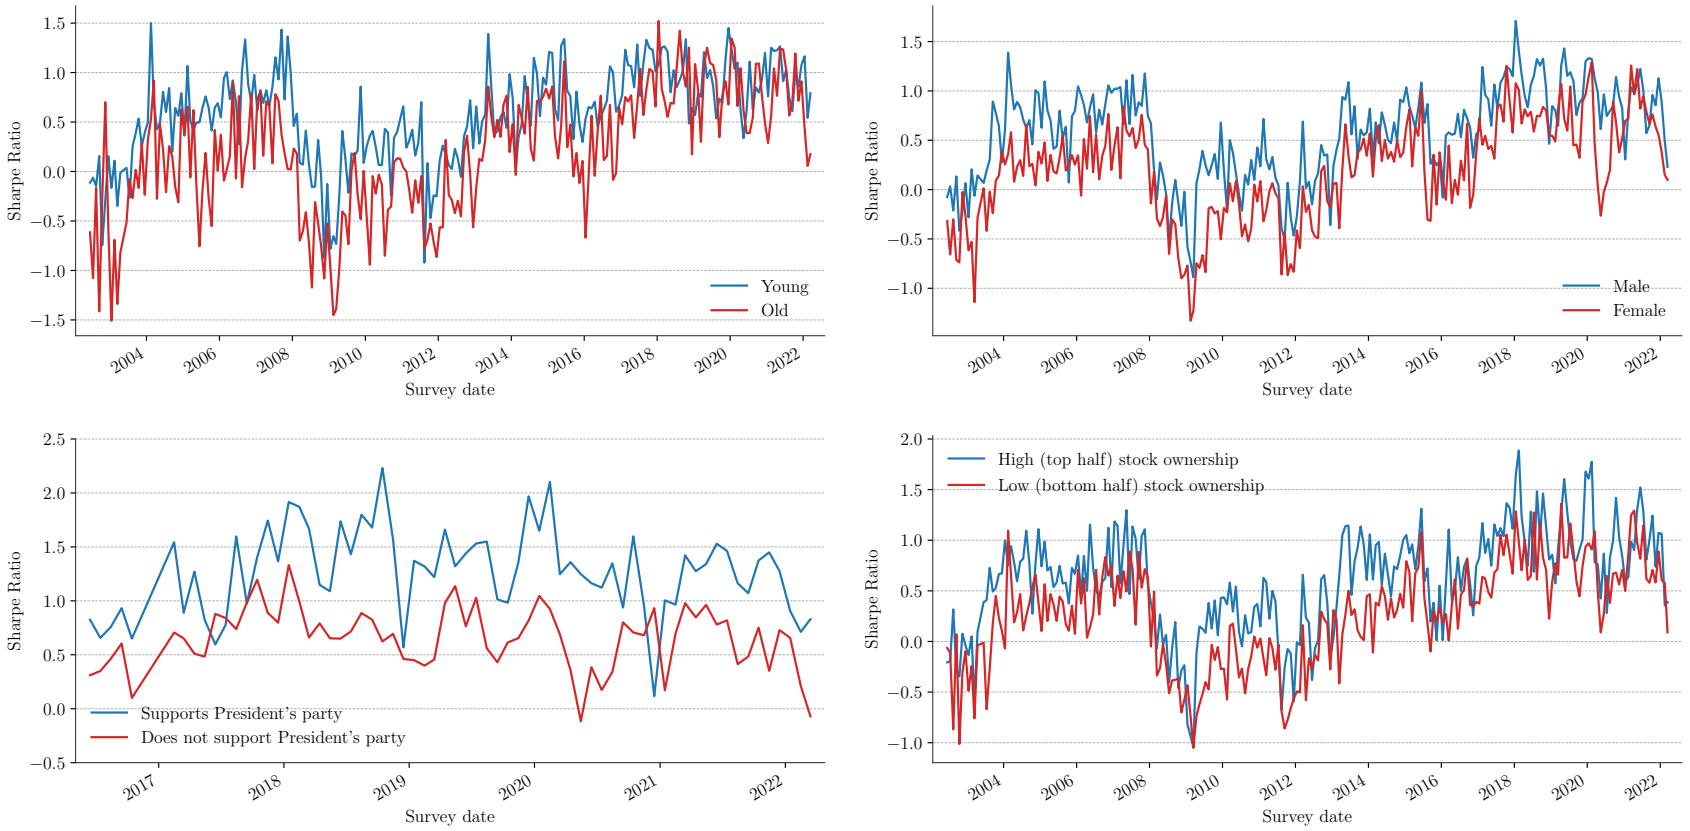


Figure 4: Mean Survey Subjective Sharpe Ratio Conditional on Various Attributes of Respondents

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.

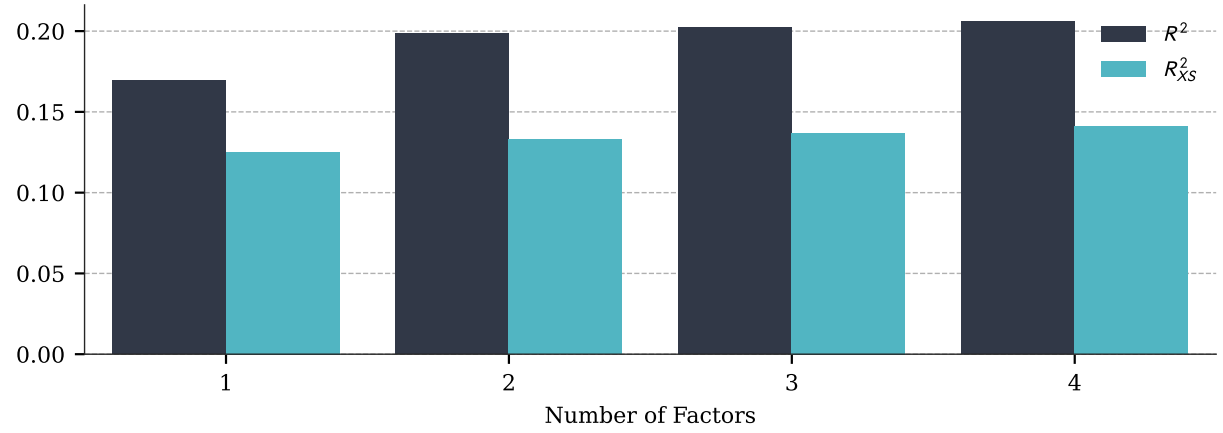


Figure 5: R^2 of IPCA with 1 to 4 Factors

The figure shows the R^2 and R^2_{XS} of IPCA models with 1 to 4 factors, following the specification in Equation (7) with $\alpha_{i,t} = 0$.

household over time, we cannot apply standard principal component analysis. We therefore 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:

$$\text{SSR}_{i,t} = \alpha_{i,t} + \beta_{i,t} f_{t+1} + \varepsilon_{i,t+1}, \qquad \beta_{i,t} = [\mathbf{d}, \mathbf{b}]'_{i,t} \mathbf{\Gamma}^{AC}_{\beta}; \quad \alpha_{i,t} = [\mathbf{d}, \mathbf{b}]'_{i,t} \mathbf{\Gamma}^{AC}_{\alpha} \tag{7}$$

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 (7), 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).

When estimating Equation (7), 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 begin by reporting the classic fraction of total variance R^2_{Total} explained by a $K \in [1, 4]$-factor IPCA model. We then also assess how much of the

cross-sectional variation of $\text{SSR}_{i,t}$ can be explained in comparison to a benchmark of the cross-sectional mean in month t, using an R^2_{XS} metric that is equivalent to the "within-R^2" of a panel regression with time effects:[6]

$$R^2_{Total} = 1 - \frac{\sum_{i=1}^{N}(\text{SSR}_{i,t} - \widehat{\text{SSR}}_{i,t})^2}{\sum_{i=1}^{N}(\text{SSR}_{i,t} - \overline{\text{SSR}})^2} \tag{8}$$

$$R^2_{XS} = 1 - \frac{\sum_{i=1}^{N}\left((\text{SSR}_{i,t} - \overline{\text{SSR}_t}) - (\widehat{\text{SSR}}_{i,t} - \overline{\widehat{\text{SSR}}_t})\right)^2}{\sum_{i=1}^{N}(\text{SSR}_{i,t} - \overline{\text{SSR}_t})^2}, \tag{9}$$

where $\widehat{\text{SSR}}_{i,t} = \beta_{i,t} f_{t+1}$, when $\alpha_{i,t}$ is restricted to zero.

Figure 5 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 beliefs that is not already captured by a household's exposure to the latent systematic factor.

Following this insight, Table 2 presents estimates for the $\mathbf{\Gamma}_\beta$ vector of Equation (7) 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 4), 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

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

Table 2: Γ_β Estimates for the IPCA Factor Model

This table shows the Γ_β estimates for a one-factor IPCA model following Equation (7) 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: Other Beliefs	
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

important in instrumenting its exposure to this factor structure. As with any latent specification, we lack a clear understanding of what these factors represent. Therefore, in Section 4.5 we investigate the commonality between ICPA-predicted subjective Sharpe Ratios and those obtained from flexibly incorporating experience effects using our machine learning method.

3 Modeling Experience Effects

In order to understand how experience effects are influenced by demographics other than age, this paper introduces a flexible experience weighting function that builds on prior work by Malmendier and Nagel (2011).

To estimate how a household's past experience in the stock market relates to expectations of the market's future performance, we consider three types of information. From the Michigan Surveys of Consumers, we include information about a respondent's demographics, such as her age or income decile, and her beliefs about other macroeconomic state variables. A breakdown of the included variables is provided in Table 1. The third input consists of past realized Sharpe Ratios of the S&P 500 (SR_t) over the investor's lifetime, for which we extract appropriate weights given the investor's demographics and beliefs about macroeconomic state variables.

There are a number of reasons why differences in demographics and various other beliefs are informative about differences in stock market expectations. For one, macroeconomic expectations may distinguish pessimists from optimists if expectations are tightly integrated – something we will investigate in more depth in Section 6. Indeed, Kamdar and Ray (2022) find that individuals' subjective expectations of different quantities tend to be linked through a factor structure.

At the same time, life-cycle considerations may be an important driver of cross-sectional differences. Anecdotal investment advice suggests that young people should invest more heavily in stocks, while older people should shift more towards bonds (Bali, Demirtas, Levy, and Wolf, 2009). These life-cycle considerations will also depend on the current financial well-being of the respondent, highlighting the need for modeling interaction effects between different demographic indicators. Financial literacy is another important factor in understanding the expectation formation process. While hard to observe empirically, we try to approximate it through a household's stock holdings and obtained degree, as well as their interaction with other demographics.

We now describe our 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 E. Much like the classic exponential weighting framework described in Eqns. (5)-(6), we model the subjective expectation held by each individual i surveyed at time t as a weighted sum of her past experienced Sharpe Ratios SR_{t-k}, with $k \in [1, \text{age}_{i,t} - 1]$. We depart in using the respondent's demographics $\mathbf{d}$ and beliefs about macroeconomic state variables $\mathbf{b}$ to instrument for the across-household variation in experience effects.

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

$$\textsc{Experience}(\mathbf{d}_{i,t}, \mathbf{b}_{i,t}) = \sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \times \text{SR}_{t-k} \tag{10}$$

$$\text{where} \quad \mathbf{w}_{i,t} = g(\text{age}_{i,t}, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \tag{11}$$

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

$$\sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) = 1 \quad \text{and} \quad w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \geq 0 \ \ \forall \ \ k \in [1, \text{age}_{i,t} - 1]. \tag{12}$$

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.

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

$$\mathbf{k}_{i,t} = \mathbf{W}_2^{\mathbf{k}} \times \text{GELU}\left(\mathbf{W}_1^{\mathbf{k}} \times \mathbf{SR}_{(t-\text{age}_{i,t}-1):t} + \mathbf{a}_1^{\mathbf{k}}\right) + \mathbf{a}_2^{\mathbf{k}}. \tag{13}$$

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

$$\mathbf{q}_{i,t} = \mathbf{W}^{\mathbf{q}} \times [\widehat{\mathbf{d}}_{i,t}, \widehat{\mathbf{b}}_{i,t}], \tag{14}$$

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

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

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

appealing for the extraction of experience effects:

$$\textsc{Attn}_{i,t} = \frac{\mathbf{q}_{i,t} \times \mathbf{k}'_{i,t}}{\sqrt{d^{\textsc{Attn}}}} \tag{15}$$

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

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

$$w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) = \frac{e^{\textsc{Attn}_{i,t}(k)}}{\sum_{k=1}^{\text{age}_{i,t}-1} e^{\textsc{Attn}_{i,t}(k)}}, \quad \forall\ k \in [1, \text{age}_{i,t} - 1] \tag{16}$$

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

We estimate the Sharpe Ratio weights such that the extracted total stock market experience best explains the variation in households' Subjective Sharpe Ratios,

$$\widehat{\text{SSR}}_{i,t} = \textsc{Experience}(\mathbf{d}_{i,t}, \mathbf{b}_{i,t}), \tag{17}$$

and minimize the mean squared error of model-predicted Subjective Sharpe Ratios across subjects i and months t:

$$\text{MSE} = \sum_{i}^{I} \sum_{t}^{T} \left(\widehat{\text{SSR}}_{i,t} - \text{SSR}_{i,t} \right)^2 \tag{18}$$

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.

[9]Note that the SoftMax normalization does not impose an exponential decay of experience weights, but is merely a convenient way to normalize weights and abide by the restrictions in Equation (12).

4 Estimated Experience Effects

4.1 Explaining Subjective Expectations with Lifetime Experiences

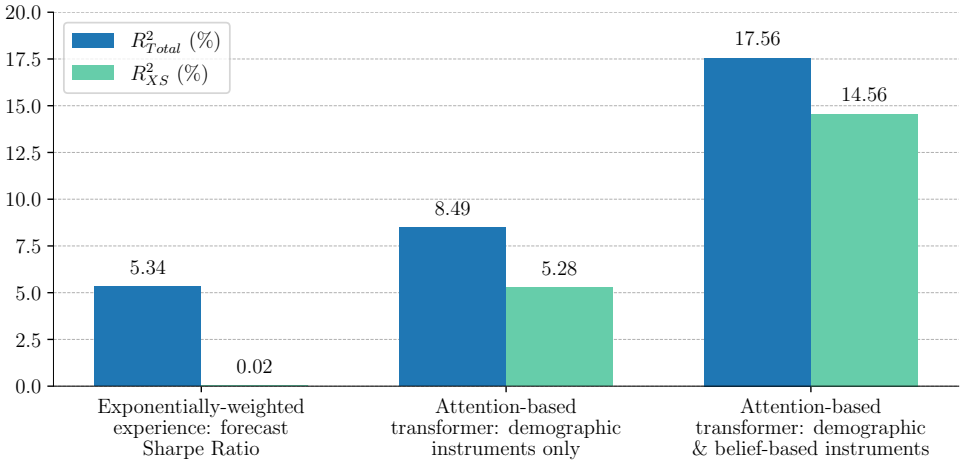


Figure 6: Explanatory Power Using Different Models Of Experience Effects

This figure compares the proportion of variance in respondents' subjective Sharpe Ratios that is explained by three empirical models of experience effects. These models incorporate past experiences in different ways. In the first model, individuals adaptively weight their past experiences according to Malmendier and Nagel (2011). The second two experience effects models are estimated by attention-based Transformer frameworks, with different sets of instruments used for computing attention weights to past experience. Blue columns report the overall fraction (%) of variance explained R^2_{Total} (8), while green columns focus on the fraction (%) of cross-sectional variance explained R^2_{XS} (9).

As a first result, we investigate how well our models of experience effects perform in explaining respondents' Subjective Sharpe Ratios, $\text{SSR}_{i,t}$. Figure 6 displays the performance of neural network-based models of experience effects (Equation (10)), alongside a classic exponentially-weighting baseline.[10] Our two neural network-based models differ in which variables are used as instruments for attention weights in Equation (11): the first uses only demographic variables $\mathbf{d}_{i,t}$ (including age), and the second augments these with beliefs about macroeconomic state variables $\mathbf{b}_{i,t}$.

The R^2 results of Figure 6 show that our neural network-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 demographic and macroeconomic belief-related variables are used as instruments for the attention weights. Although not intended as an exhaustive comparison between approaches, we note that a baseline that exponentially weights past experiences

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

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 6 – they merely instrument the weights nonlinearly in the Transformer estimation framework.

When excluding other beliefs from the 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%. The dropoff in explanatory power suggests that the excluded variables are useful in capturing cross-sectional heterogeneity in individual optimism or pessimism.

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 vary by respondents' demographic characteristics and variables that capture their optimism or pessimism in general.

4.2 Demographic-driven Differences in Experience Weights

The differences in explanatory power in Figure 6 stem from the ability of individual demographic characteristics beyond age to capture heterogeneity in experience effects. By contrast, a baseline approach of exponentially-weighted experience effects does not capture cross-sectional variation in lifetime experience effects. We therefore illustrate how various demographic characteristics affect how investors pay attention to different parts of their experiences. To do so, we condition on investor age and then interrogate the experience weights estimated by our richest model specification.

Figure 7 visually illustrates variation in experience weights across investors who differ in selected individual characteristics other than age (which we fix at the median of 53 years). No matter which variable we condition on, weights on past experiences are broadly decaying, albeit with non-monotonicities: experiences in the recent past matter far more than earlier, more distant, experiences in determining investors' expectations about the future. Although no functional form was explicitly specified at the outset (besides the constraint that weights should be non-negative and sum to unity), this broad decay suggests that Malmendier and Nagel (2011)'s core framework of decaying past weights is justified in principle.

It is apparent that investors' education levels play a role in experience effects: investors without any post-high school education pay less attention to recent experiences and more attention to their formative (teenage) experiences. Gender also plays a role in experience effects: female investors pay more attention to the most recent (5 years) part of their lifetime experiences, while male investors pay more attention to the two decades before that; they do not vary in early-lifetime experience effects. Stock ownership plays a role too: investors with larger stock investments (in dollar terms) pay more attention to recent experiences and

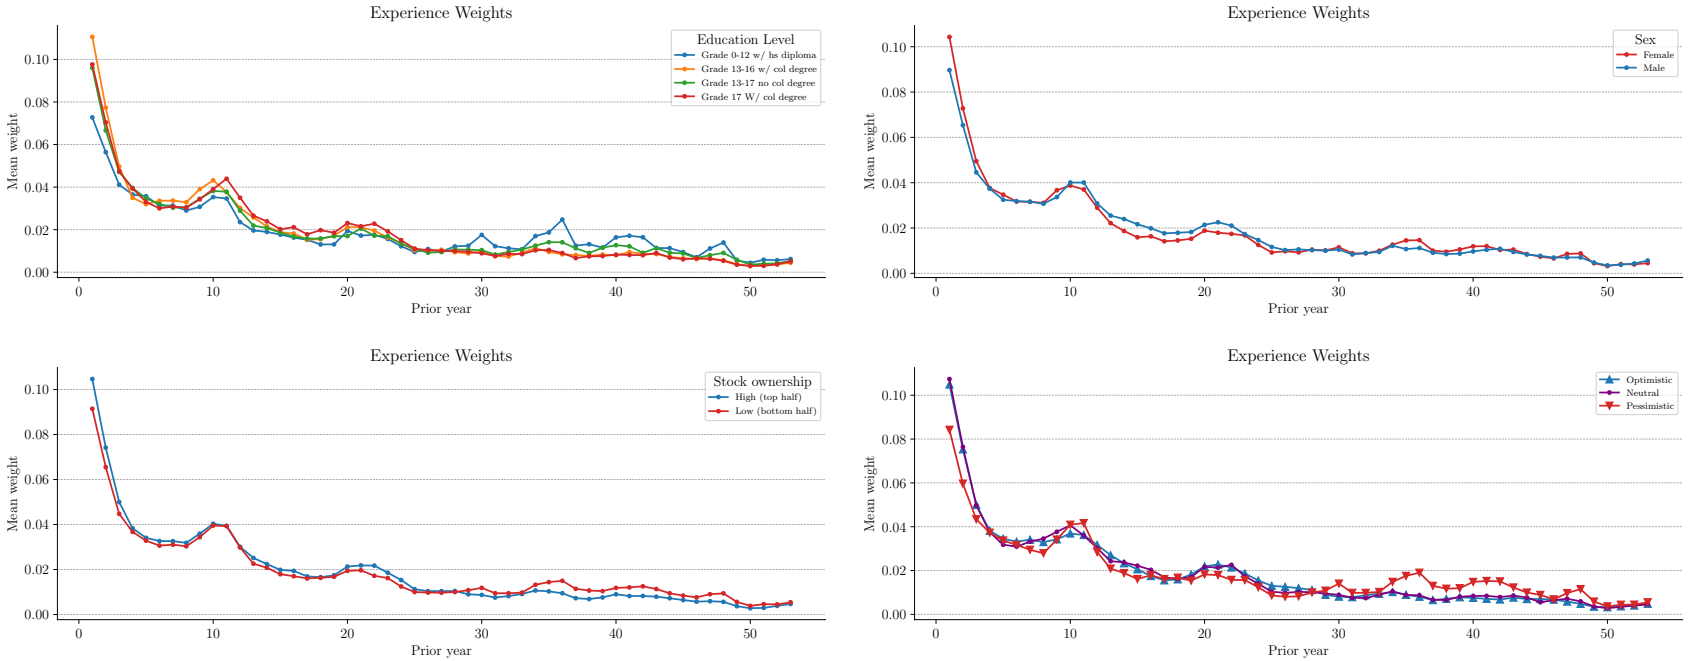


Figure 7: Estimated Experience Weights Conditional on Various Attributes of Respondents (for Respondents of Median Age)

Each chart shows mean individual-level experience weights estimated by our Transformer-based estimation technology (in its richest specification), for the subset of respondents aged 53 years old (the median age in the overall Michigan sample), conditioned by a single demographic attribute besides age. The upper-left chart conditions on education level. The upper-right chart conditions on gender. The bottom-left chart conditions on the amount invested in the stock market; more specifically, on whether an investor is in the top half or bottom half of investors by size (calculated cross-sectionally). The bottom-right chart conditions on individuals' general levels of optimism or pessimism, as proxied by the individual level optimism index of Equation (19).

less to distant experiences earlier in their lifetimes. Finally, beliefs about macroeconomic state variables appear to successfully capture variation in individual-level optimism or pessimism: an individual-level index of beliefs about macroeconomic variables,

$$\begin{aligned}\text{Optimism Index}_{i,t} \;&= \text{sign}\big(\text{Forecast Probability of Increase in Personal Income}_{i,t} - 0.5\big)\\ &\quad + \text{Forecast Direction of Business Conditions}_{i,t}\\ &\quad - \text{Forecast Direction of Interest Rates}_{i,t}\\ &\quad - \text{Forecast Direction of Unemployment}_{i,t},\end{aligned} \tag{19}$$

reveals that individuals who are generally pessimistic (i.e. with a negative index value) pay less attention to recent experiences and rely instead on their earlier lifetime experiences when forming expectations about the future.

Table 3 shows the results of baseline panel regressions conducted on the Transformer-estimated experience weights of individuals i surveyed at time t on a prior year k. These annual experience weights are regressed against a dummy that takes the value 1 for prior years $1 \leq k \leq 10$, i.e., for all experience weights on the decade prior to the survey date. Columns (2) & (3) include the respondents' ages as a control variable, since equal weights (as a hypothetical scenario) would increase linearly with the length of respondents' lifetimes. Finally, columns (3) & (4) include time effects. These unconditional results are consistent across all four specifications: respondents place a significantly larger weight on experiences within the most recent decade. Taking the estimates for the richest specification in column (4) as an example, respondents place an additional $10 \times 3.241 = 32.41$ percentage points of weight on their past decades' worth of experience when forming subjective expectations about the market's Sharpe Ratio.

Table 4 repeats the richest regression specification of Table 3's column (4), and now includes interactions with the four demographic attributes used as conditioning variables in Figure 7. There are a few advantages of conducting regressions over producing plots: first, one can include age as a control rather than conditioning on a subset of respondents with a certain age. Second, coefficient estimates are interpretable in the units of experience weights. And finally, the significance of demographic interactions can be statistically ascertained (including by taking cluster heterogeneity into account). The coefficient estimates in the first row of Table 4 indicate that all respondents continue to (unconditionally) place more weight on years in the most recent decade of their experiences. The last eight rows of the table capture the incremental effects (via interacted covariates) of various demographic attributes on the tendency of investors to place more weight on recent years. By interpreting the coefficient estimates, we can see that women place an additional $10 \times 0.290 = 2.90$ percentage points of weight than men do on the most recent decade of their experiences. Similarly, a twenty-

Table 3: Experience Weights on Recent Years

This table presents the results of OLS regressions of annual experience weights (response variable) versus an indicator for the prior decade of lifetime experience. Variants of this baseline specification replace the intercept with survey wave fixed effects, or include a control for the lengths of investors' lifetimes.

Dependent Variable:	Experience Weight$_{i,t,k}$ (%)			
Model:	(1)	(2)	(3)	(4)
Variables				
Prior Decade$_k$	3.399***	3.241***	3.396***	3.241***
	(0.0471)	(0.0490)	(0.0471)	(0.0490)
Age$_{i,t}$		-0.0317***		-0.0318***
		(0.0003)		(0.0003)
(Intercept)	1.317***	3.062***		
	(0.0098)	(0.0246)		
Fixed-effects				
Survey year-month			✓	✓
Fit statistics				
Observations	3,804,951	3,804,951	3,804,951	3,804,951
R^2	0.19941	0.21735	0.19974	0.21738

Clustered (Survey year-month) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

percentile-rank increase in the stock holdings of an investor is associated with the investor placing an additional $20 \times 10 \times 0.0058 = 1.16$ percentage points of weight on the most recent decade of his/her personal experiences. An investor who supports the current President's party likewise places an additional $10 \times 0.1562 = 1.56$ percentage points of weight on his/her most recent decade of lifetime experience (compared to an investor who does not support the current President's party). Finally, the results in the last five rows of column (4) compare investors' incremental experience weights by education level against a baseline of an investor who has a high school diploma and is educated to the level of grades 0-12: investors with lower levels of education than this baseline place a lower weight on their most recent decade's worth of experiences, and investors with higher levels of education than the same baseline place a higher weight on their most recent decade's worth of experiences.

4.3 Lifetime Experience Effects and Age

In a similar exercise to the previous section, we now interrogate the model to reveal the estimated functional form of experience weights by the age of the individual. Figure 8 plots the mean experience weight per lag over all individual-level weights, this time conditioning on

Table 4: Experience Weights on Recent Years Conditioning on Demographics

This table presents the results of OLS regressions of annual experience weights (response variable) versus an indicator for the prior decade of lifetime experience, interacted with demographic attributes, and with a control for lifetime length. The baseline level of the categorical education code variable used in column (4) is grades 0-12 with a high school diploma; i.e. a high school graduate with no further education.

Dependent Variable:	Experience Weight$_{i,t,k}$ (%)			
Model:	(1)	(2)	(3)	(4)
Variables				
Prior Decade$_k$	3.110***	2.967***	3.281***	2.588***
	(0.0451)	(0.0581)	(0.0927)	(0.0518)
Age$_{i,t}$	-0.0318***	-0.0316***	-0.0316***	-0.0318***
	(0.0003)	(0.0003)	(0.0005)	(0.0003)
Female$_i$	-0.0655***			
	(0.0046)			
Holdings Percentile$_{i,t}$		-0.0016***		
		(6.48×10^{-5})		
Supports President$_{i,t}$			-0.0337***	
			(0.0074)	
Educ. Grades 0-8, No HS Diploma$_{i,t}$				0.0671***
				(0.0231)
Educ. Grades 9-12, No HS Diploma$_{i,t}$				0.1126***
				(0.0138)
Educ. Grades 13-16, Has College Degree$_{i,t}$				-0.2170***
				(0.0050)
Educ. Grades 13-17, No College Degree$_{i,t}$				-0.1290***
				(0.0043)
Educ. Grade 17, Has College Degree$_{i,t}$				-0.1624***
				(0.0052)
Prior Decade$_k$ × Female$_i$	0.2908***			
	(0.0235)			
Prior Decade$_k$ × Holdings Percentile$_{i,t}$		0.0058***		
		(0.0003)		
Prior Decade$_k$ × Supports President$_{i,t}$			0.1562***	
			(0.0372)	
Prior Decade$_k$ × Educ. Grades 0-8, No HS Diploma$_{i,t}$				-0.2710**
				(0.1088)
Prior Decade$_k$ × Educ. Grades 9-12, No HS Diploma$_{i,t}$				-0.0234
				(0.0726)
Prior Decade$_k$ × Educ. Grades 13-16, Has College Degree$_{i,t}$				0.9768***
				(0.0236)
Prior Decade$_k$ × Educ. Grades 13-17, No College Degree$_{i,t}$				0.6342***
				(0.0208)
Prior Decade$_k$ × Educ. Grade 17, Has College Degree$_{i,t}$				0.6966***
				(0.0256)
Fixed-effects				
Survey year-month	✓	✓	✓	✓
Fit statistics				
Observations	3,804,951	3,155,122	1,687,823	3,797,179
R^2	0.21774	0.22414	0.23714	0.21939

Clustered (Survey year-month) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

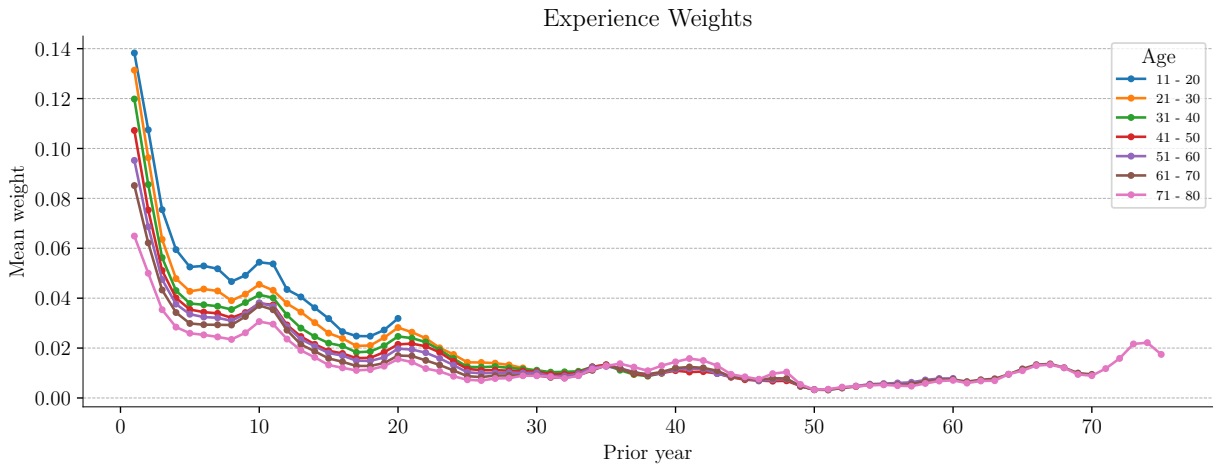


Figure 8: Estimated Experience Weights On Prior Years by Age

These charts show the means of estimated attention weights paid to experienced Sharpe Ratios, taken over groups of respondents with similar ages, with experienced years presented as an offset from the survey date. The weights are estimated by our customized Transformer framework, using both demographic attributes and other beliefs as instruments.

cohort (age range) instead of other demographic attributes.

There are a number of similarities and differences to the weights implied by the Malmendier and Nagel (2011) model, as seen earlier in Figure 3. First, both experience weighting schemes imply that younger investors place a proportionately higher weight on recent experiences as compared to earlier experiences. Second, the Transformer-estimated weights exhibit non-monotonicities that may capture important common experiences across individuals. And third, the decay from recent to distant past experiences is much less pronounced in the Transformer-estimated specification.

4.4 Influence of Past Recession Periods

We now consider the persistence with which past recessions that a respondent has lived through influence her expectations about the future. To do this, we regress respondents' individual-level experience weights against a dummy for whether the experience weights intersect with a past recession period, and an additional dummy for whether investors are experiencing a recession today. Table 5 shows the results.

The positive and significant coefficients in the first row of Table 5 indicate that investors are always more attentive to years in past recession periods than to non-recession periods. While recessions may make up only a fraction of investors' overall lifetimes, experiences of recessions exert an outsized influence today. The flexible form of how past experiences shape today's expectations in Equation (10) allows the model to identify important clusters of experiences

Table 5: Experience Weights on Past Recession Periods

This table presents the results of OLS regressions of annual experience weights (response variable) versus an indicator for whether the prior year fell into a recession period, and an interaction with another indicator for whether investors' subjective expectations are elicited during a recession period too. Columns (3) & (4) are for reduced samples of experience weights. All specifications include a control for the lengths of investors' lifetimes.

Dependent Variable:	Experience Weight$_{i,t,k}$ (%)			
Model:	(1)	(2)	(3)	(4)
Variables				
Experience Weight During Recession$_{t,k}$	0.8273***	0.7496***	0.7695***	0.6470***
	(0.0519)	(0.0510)	(0.0375)	(0.0355)
Age$_{i,t}$	-0.0486***	-0.0485***	-0.0353***	-0.0279***
	(0.0002)	(0.0002)	(0.0003)	(0.0005)
Experience Weight During Recession$_{t,k}$		0.9115***	0.0785	-0.1532**
× Surveyed During Recession$_{i,t}$		(0.1929)	(0.1035)	(0.0617)
Fixed-effects				
Survey year-month	✓	✓	✓	✓
Sample				
Included experience weights k	All	All	$k > 5$	$k > 10$
Fit statistics				
Observations	3,804,951	3,804,951	3,425,406	3,045,861
R^2	0.05430	0.05564	0.05818	0.04834

Clustered (Survey year-month) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

that would remain invisible under the assumption of a more restrictive parametric decay.

The coefficients estimated in the third row of Table 5 test whether a recession happening at the time investors stated their subjective expectations has an additional effect on how much attention they pay to past periods of recession. The positive and significant coefficient in column (2) indicates that this indeed appears to be the case: investors are more attentive to past periods of recession when they are currently living through a recession. The coefficients in the last row of columns (3) & (4) tell a subtler story: when the past 5 years (column 3) and then 10 years (column 4) are excluded from the sample, the effect slowly reverses. That is, while investors pay extra attention to past recession periods when they are currently living through a recession, they concentrate this in the most recent few years; i.e. they extrapolate from the current ongoing recession. Still, the positive coefficients throughout the first row of Table 5 confirm that investors unconditionally always pay more attention to past years when those past years are recession years.

There are two key implications of these results. The first is that past recessions exert an

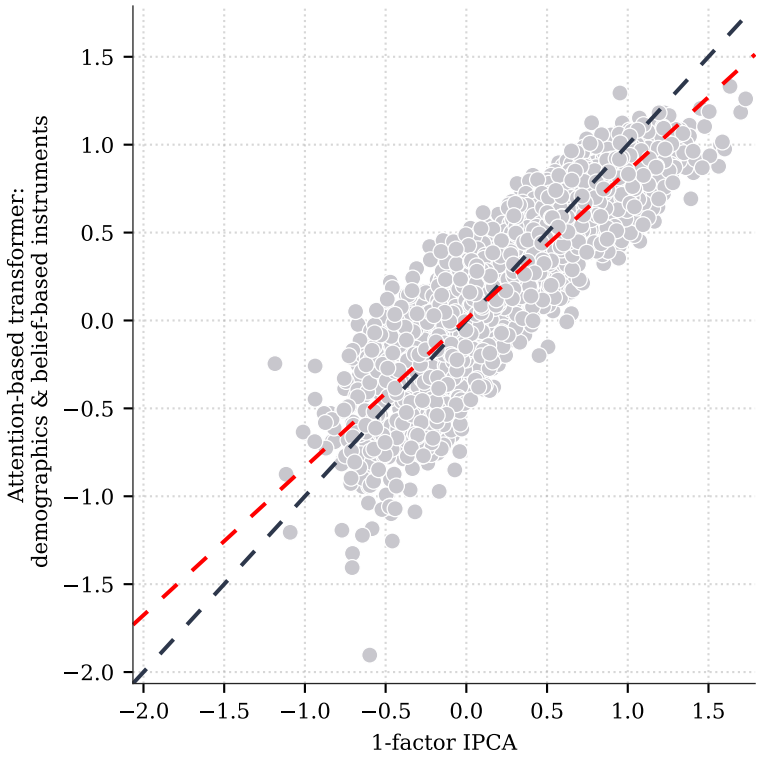


Figure 9: Predictions of the Latent 1-Factor Model vs. Transformer-based Experience Effects

The figure shows the predictions of the 1-factor IPCA, specified in Equation (7) 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.

outsized influence on investors' experiences, irrespective of the current state of the economy. This suggests that constant gain learning models along the lines of Nagel and Xu (2022) may not be sufficient to capture this empirical phenomenon. The second implication is that investors become even more extrapolative during recession periods. This sheds light on the mechanism underlying the observed procyclicality in investors' subjective expectations.

4.5 Experience Effects and Latent Factors

In Section 2.5 we find that a single latent factor is sufficient to drive out the direct explanatory power of household demographics and macroeconomic beliefs, and explains roughly 17% of the variation in subjective Sharpe Ratios. This latent factor has no direct interpretation, but is instead the result of a reduced-form statistical exercise. Our Transformer-estimated experience effects model, which captures an economic mechanism, also succeeds in explaining 17.56% of the variation in households' subjective expectations. But how well do these experience effects pick up the reduced-form explanatory power of the latent IPCA model?

Figure 9 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.

5 Consequences for the Household Equity Share

This section shows that investors' subjective Sharpe Ratios determine aggregate holdings in equities. Furthermore, the mechanism of experience effects appears to drive this relationship.

5.1 Expectations and Equity Shares

5.1.1 Direct and Indirect Equity Holdings by Households

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

We are able to measure wealth shares held in equities – the fraction x in Equation (4) – at an aggregate level, thanks to US national accounts data.[11] Figure 10 compares the timeseries variation in the aggregate household's equity share with the timeseries variation in the subjective Sharpe Ratio perceived by this representative investor. The two series covary strongly: the correlation coefficient of the two series is 0.66 (and is significant at the 1%-level). Furthermore, the effect is not driven by a mechanical increase in the market's value value over time: in a regression framework (unreported), the positive association survives controlling for the level of the S&P 500, and is significant even when standard errors are clustered by year-month.

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

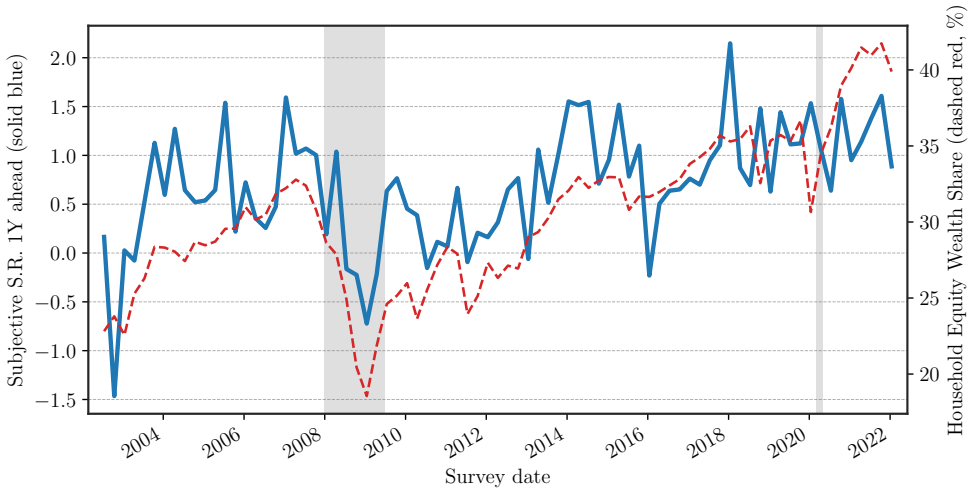


Figure 10: Mean Survey Subjective Sharpe Ratios for Year Ahead vs. Households' Share of Wealth in Equities

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

5.1.2 Fund Flows

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

We measure the representative investor's Subjective Sharpe Ratio by the mean of individual investors' values at each point in time. Aggregate flows into the stock market are calculated using Investment Company Institute data on fund flows:[12]

$$\text{Flows}_{t\to t+1} = \frac{\$\text{ Domestic Equity Fund Flows}_{t\to t+1}}{\$\text{ Domestic Equity Market Capitalization}_t} \times 100\%. \tag{20}$$

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

[12]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.

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.

5.2 Experience Effects and Equity Shares

As a final step, we analyze to what extent experience effects drive the relationship between investors' subjective expectations and their holdings. To answer this question, we use our richest Transformer-estimated model of experience effects to predict each household's subjective expectation based on her past experiences, $\widehat{\text{SSR}}_{i,t}$. We then multiply these by the dollar amount invested in the stock market and compute the share of all investment wealth in our sample for which the experienced subjective Sharpe Ratio $\widehat{\text{SSR}}_{i,t}$ is positive.[13] This quantity measures the aggregate share of wealth in the household sector that has experienced a positive Sharpe Ratio according to a weighted sum of lifetime realizations. We emphasize that this measure results solely from aggregated individual investor-level predictions by a model that itself made use only of microdata from the Michigan surveys (and no aggregate quantities).

The results are shown in Figure 11. The share of optimistic wealth is positively correlated with the aggregate equity share of the household sector, as reported in the US national accounts. The magnitude of the correlation is 0.77 (and is highly significant), surpassing the observed correlation between the aggregate equity share and the subjective Sharpe Ratios recovered from the Michigan surveys. This strengthening of the association suggests that experience effects are an important driver of overall holdings in equity asset markets.

We consider this result to be a striking endorsement of the importance of experience effects for explaining household portfolio behavior, and therefore aggregate asset prices.

6 How Integrated are Household Expectations?

Malmendier and Nagel (2016) show that the decay parameter in Equation (5), which governs how much weight respondents place on more recent vs. more distant experienced returns or inflation, is similar for inflation and stock market expectations. Following this evidence, we investigate to what extent the extracted stock market experience of single households, estimated as in Equation (10), can differentiate more optimistic from more pessimistic households along other macroeconomic dimensions. For this, we group respondents by their total stock

[13] Simply computing the share of investors with an experienced positive Sharpe Ratio (without weighting by invested amounts) yields similar results.

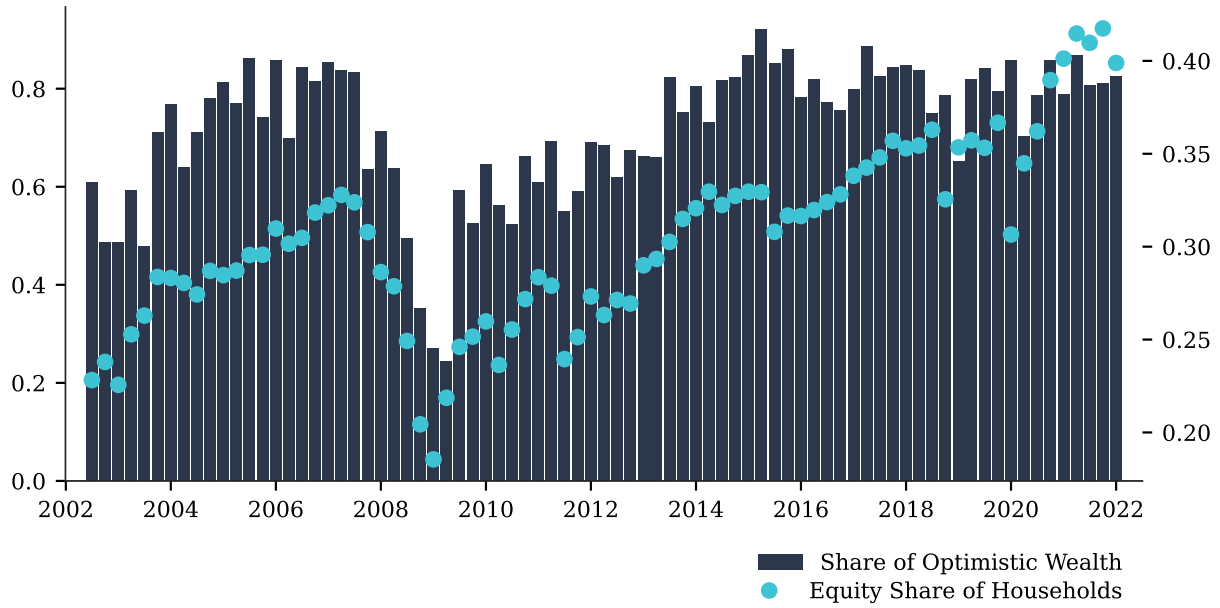


Figure 11: Equity Share of Households vs. Stock Wealth Held by Optimists

The figure shows the equity share of households in blue dots and the share of stock wealth held by optimists as gray bars. Optimists are defined as respondents who have experienced a positive Sharpe Ratio. The left y-axis is the share of optimistic wealth, and the right y-axis is the equity share of the household sector.

market experience, defined as their experience-weighted lifetime Sharpe Ratio. We dub the group of investors who experienced a positive Sharpe Ratio "optimists" and the group which experienced a negative Sharpe Ratio "pessimists." We assess whether systematic differences in demographics exist between these two groups. The average age amounts to roughly 50 years for both groups, optimists are slightly less invested in the stock market, have lower household incomes, and slightly lower education. Women are slightly less optimistic than men, which squares well with abundant research on overconfidence and relative optimism between genders (for example, Barber and Odean, 2001, and Jacobsen, Lee, Marquering, and Zhang, 2014). Furthermore, we find no systematic differences between Republican and Democratic party supporters.

Next, we investigate if differences in experienced stock market Sharpe Ratios translate into differences in expectations of other macroeconomic variables. We include expectations about future business conditions, future unemployment rates, personal income growth, changes in house prices and expectations of the future direction of interest rates. In Figure 12 we show that stock market experiences can distinguish more optimistic from more pessimistic expectations among households for other macroeconomic dimensions: we find a pronounced cross-sectional wedge for subjective Sharpe Ratios, future business conditions, a household's personal income and expectations about future unemployment rates. These results show that our approach is

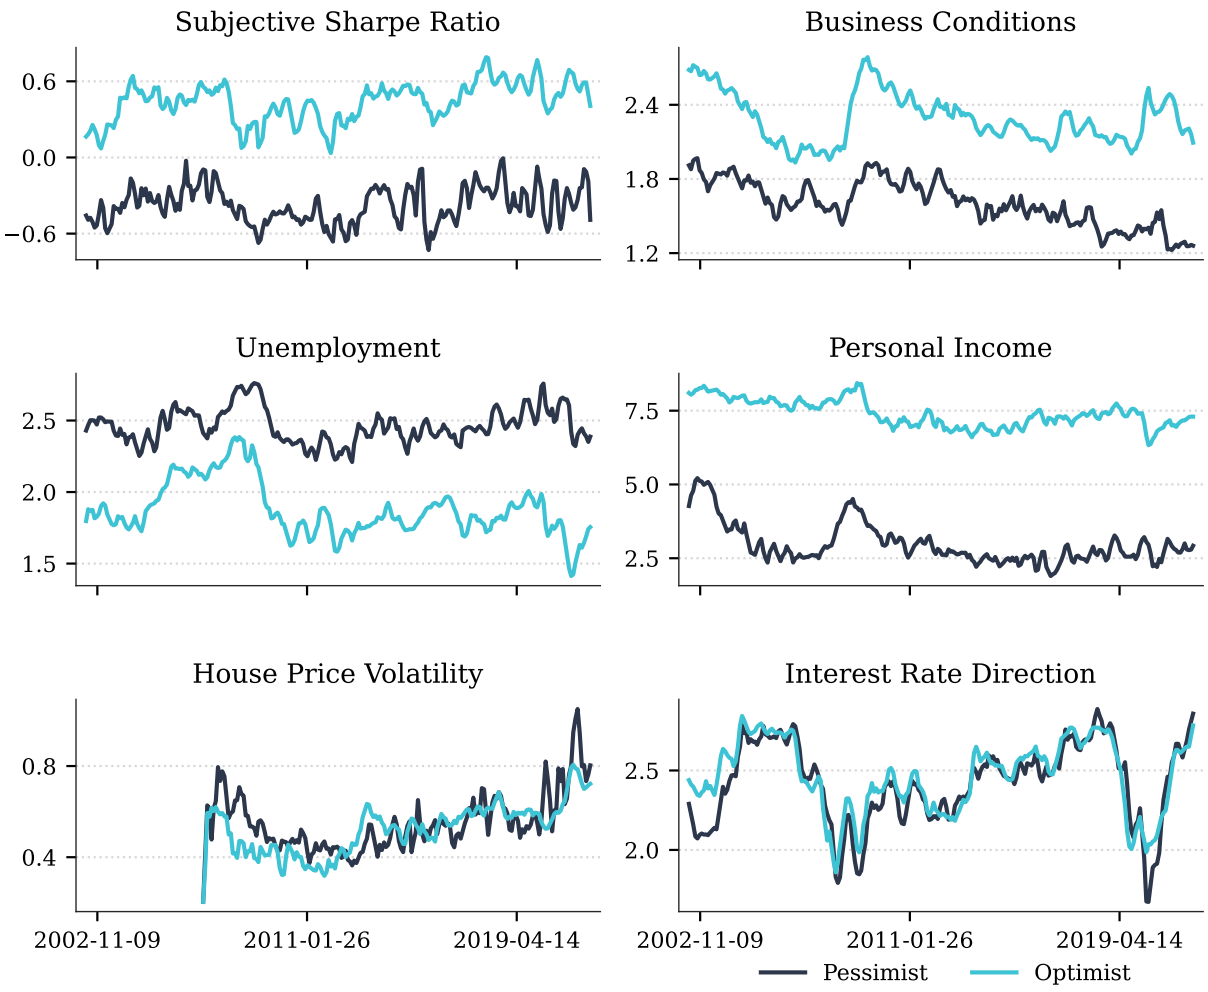


Figure 12: Cross-Experience Effects using Instrumented Experience Weights

The figure shows average expectations of more optimistic and more pessimistic households, where optimists are identified as those for which the experienced Sharpe Ratio exceeds the cross-sectional median. We separately show the spread in a number of macroeconomic expectations, including the probability of an up-move in the stock market, expectations about future business conditions, unemployment rates, household income, house prices, and the direction of interest rates.

able to extract meaningful experience patterns from demographic information and beliefs. As we only use stock market experience for the identification of high vs. low expectations about other macroeconomic indicators, this suggests that stock market experiences are important for explaining the overall expectation formation process of households. On the flip-side, we cannot distinguish between high and low expectations regarding future house price volatility, nor the direction of future interest rates. These expectations seem to be uncoupled from experiences made in the stock market.

Our model is able to extract stock market experience effects, which allows us to dissect the

cross-section of households into those with more optimistic and those with more pessimistic future outlooks concerning a number of macroeconomic indicators. We now compare the results using our flexible experience weighting function with the baseline approach of experience weights that decay at a fixed rate. Figure 2 in the Appendix has the results: consistent with the evidence in Figure 6 that age alone is an insufficient conditioning variable to explain variation *across* households, not instrumenting cross-sectional and temporal variation in experience effects with demographics and beliefs fails to produce the carry-over effect we document for our machine learning approach. We find little to no differences of other expectations for households with high or low experienced stock market returns. For personal income, we even find the opposite effect for the first part of the sample: households who experienced lower life-time market returns expect more favorable growth of total household income.

7 Conclusion

This paper introduces a new source of subjective expectations data about the conditional Sharpe Ratio of the market portfolio. We contribute a set of stylized facts about investors' Subjective Sharpe Ratios, building upon recent work documenting the properties of other subjective expectations. To understand how these subjective expectations are driven by experience effects, we extend a classic adaptive model of experience effects along two dimensions. First, we allow for complex functional forms. Second, we allow for attributes of individuals other than age (such as demographics) to instrument for the weights they place on their past experiences. Our empirical model performs well in explaining the cross-section of subjective expectations about the stock market's conditional Sharpe Ratio. Furthermore, the success of our approach suggests that models of experience effects can explain cross-sectional variation in beliefs that would otherwise be attributed simply to immutable demographic characteristics. We link estimated experience effects to the share of household wealth held in equities, suggesting that the mechanism of learning from experience is an important driver of this procyclical aggregate quantity.

References

Abel, A. B. (2001). Will bequests attenuate the predicted meltdown in stock prices when baby boomers retire? *Review of Economics and Statistics*, *83*(4), 589–595.

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.

Andonov, A., and Rauh, J. D. (2022). The return expectations of public pension funds. *Review of Financial Studies*, *35*(8), 3777–3822.

Bali, T. G., Demirtas, K. O., Levy, H., and Wolf, A. (2009). Bonds versus stocks: Investors' age and risk taking. *Journal of Monetary Economics*, *56*(6), 817–830.

Barber, B. M., and Odean, T. (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics*, *116*(1), 261–292.

Barbon, A. (2019). Focusing at high frequency: An attention-based neural network for limit order books. http://www.abarbon.com/papers/focusing-at-high-frequency

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.

Bodoh-Creed, A. L. (2020). Mood, memory, and the evaluation of asset prices. *Review of Finance*, *24*(1), 227–262.

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

Brunnermeier, M., Farhi, E., Koijen, R. S., Krishnamurthy, A., Ludvigson, S. C., Lustig, H., Nagel, S., and Piazzesi, M. (2021). Perspectives on the future of asset pricing. *Review of Financial Studies*, *34*(4), 2126–2160.

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.

Cao, S., Jiang, W., Wang, J. L., and Yang, B. (2021). From man vs. machine to man + machine: The art and AI of stock analyses. *NBER Working Paper*, *28800*.

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.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, *2*(4), 303–314.

Da, Z., Huang, X., and Jin, L. J. (2021). Extrapolative beliefs in the cross-section: What can we learn from the crowds? *Journal of Financial Economics*, *140*(1), 175–196.

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

de Silva, T., and Thesmar, D. (2021). Noise in expectations: Evidence from analyst forecasts. *NBER Working Paper*, *28963*.

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.

Evans, G. W., and Honkapohja, S. (2001). *Learning and expectations in macroeconomics*. Princeton University Press.

Favero, C. A., Gozluklu, A. E., and Tamoni, A. (2011). Demographic trends, the dividend-price ratio, and the predictability of long-run stock market returns. *Journal of Financial and Quantitative Analysis*, *46*(5), 1493–1520.

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

Geanakoplos, J., Magill, M., and Quinzii, M. (2004). Demography and the long-run predictability of the stock market. *Brookings Papers on Economic Activity*, *2004*(1), 241–325.

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.

Goyal, A. (2004). Demographics, stock market flows, and stock returns. *Journal of Financial and Quantitative Analysis*, *39*(1), 115–142.

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*.

Jacobsen, B., Lee, J. B., Marquering, W., and Zhang, C. Y. (2014). Gender differences in optimism and asset allocation. *Journal of Economic Behavior & Organization*, *107*, 630–651.

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

Kelly, B. T., Malamud, S., and Zhou, K. (2022). The virtue of complexity in machine learning portfolios. *Available at SSRN*, *3984925*.

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*.

Kuhnen, C. M., and Knutson, B. (2011). The influence of affect on beliefs, preferences, and financial decisions. *Journal of Financial and Quantitative Analysis*, *46*(3), 605–626.

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.

Malmendier, U., and Nagel, S. (2016). Learning from inflation experiences. *Quarterly Journal of Economics*, *131*(1), 53–87.

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

Nagel, S., and Xu, Z. (2022). Asset pricing with fading memory. *Review of Financial Studies*, *35*(5), 2190–2245.

Nakamura, E. (2005). Inflation forecasting using a neural network. *Economics Letters*, *86*(3), 373–378.

Poterba, J. M. (2001). Demographic structure and asset returns. *Review of Economics and Statistics*, *83*(4), 565–584.

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

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

van Binsbergen, J. H., Han, X., and Lopez-Lira, A. (2022). Man vs. machine learning: The term structure of earnings expectations and conditional biases. *Review of Financial Studies*.

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*.

Wachter, J. A., and Kahana, M. J. (2019). A retrieved-context theory of financial decisions. *NBER Working Paper*, *26200*.

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

Appendix

A Additional Figures and Tables

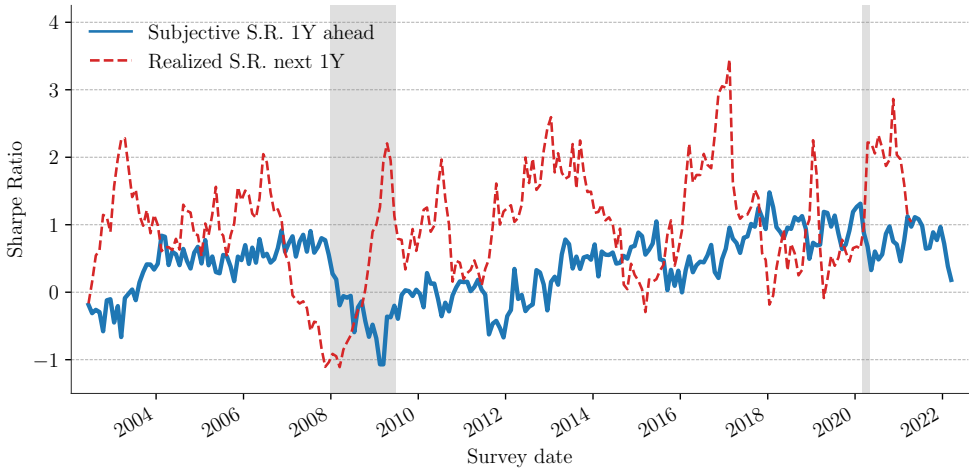


Figure 1: Mean Survey Subjective Sharpe Ratio for Next Year vs. Actual Future Realization

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

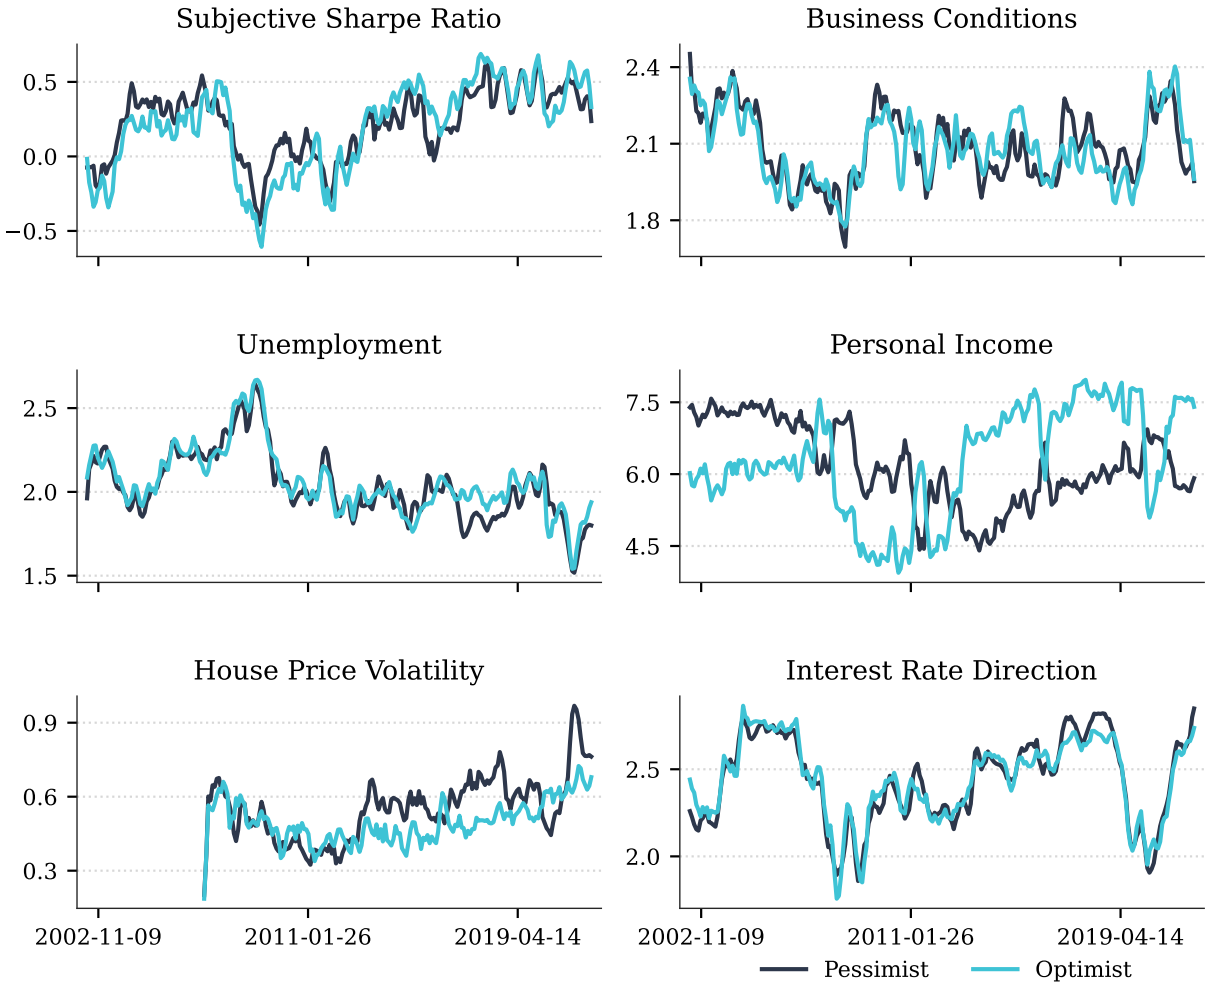


Figure 2: Cross-Experience Effects using the Malmendier and Nagel, 2011 Approach

The figure shows average expectations of more optimistic and more pessimistic households, where optimists are identified as those for which the experienced market return using the fixed-decay approach of Malmendier and Nagel, 2011 exceeds the cross-sectional median. We separately show the spread in a number of macroeconomic expectations, including the probability of an up-move in the stock market, expectations about future business conditions, unemployment rates, household income, house prices, and the direction of interest rates.

B Mapping Survey Responses to Subjective Sharpe Ratios

In this appendix we derive Proposition 1 and then examine the nature of the mapping.

B.1 Derivation

Recall that the error function is defined as

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

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:

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

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

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

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

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

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

in terms of the error function erf as

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

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

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

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

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

and then rearrange as

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

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

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

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

B.2 Empirical Behavior

Figure 3 illustrates the probit transformation. While it is nonlinear, it is well approximated by a linear relationship over most of the input domain.

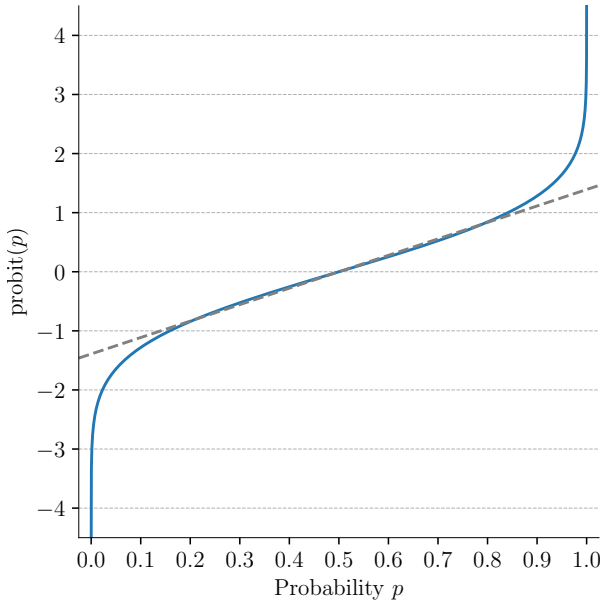


Figure 3: Probit transformation of example survey responses (x-axis) to subjective Sharpe Ratios (y-axis), following Proposition 1.

We now examine the extent to which the assumption of a zero riskless rate perceived by

the agent adds an approximation error to our measure of the recovered Sharpe Ratio. Recall that the Sharpe Ratio for a single asset is defined as

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

where μ_t and σ_t are the perceived first and second moments of the asset's return distribution, and r_f is the perceived riskless rate. Our assumption in deriving Proposition 1 implies that the second term r_f/σ_t is zero. To understand the consequences of an agent who perceives a non-zero riskless rate, Table 1 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.

Table 1: 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

C A Rational Benchmark for the Conditional Sharpe Ratio

In our paper we draw a contrast between the procyclical nature of elicited subjective Sharpe Ratios and the countercyclical nature of a rational benchmark. This appendix explains the construction of the rational benchmark, which is based on a classic consumption-based asset pricing model.

C.1 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 4 plots the relationship between the P/D ratio and the Sharpe Ratio:[14] 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.

C.2 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 5 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.

[14] These results are presented separately in Figures 3 and 6 of the Campbell and Cochrane (1999) paper. Our Figure 4 combines the two.

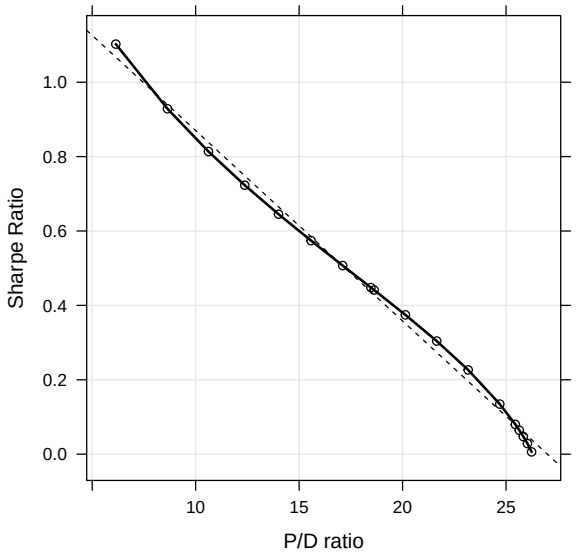


Figure 4: Relationship between the Sharpe Ratio and P/D ratio of the dividend claim in Campbell and Cochrane (1999)'s numerically solved asset pricing model. The dashed gray line is a linear approximation.

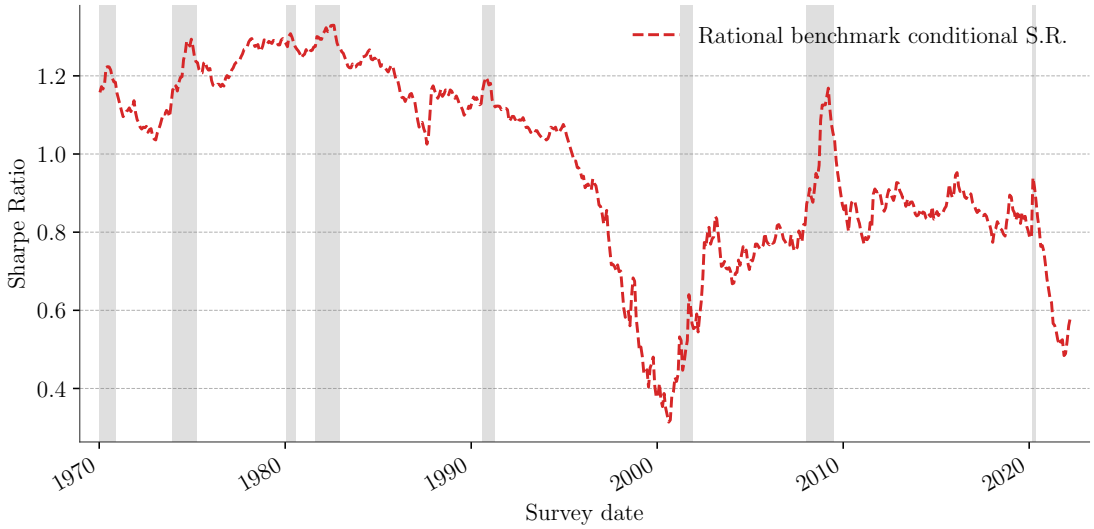


Figure 5: Empirical benchmark for the conditional Sharpe Ratio, as implied by a rational model in which the P/D ratio negatively predicts the Sharpe Ratio, such as the habit formation model of Campbell and Cochrane (1999). Shaded areas are NBER recession periods.

D Estimating the Malmendier and Nagel (2011) Model

This appendix provides further details on how the adaptive weighting model of Section 2.3 was estimated.

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

We run a grid search, similar to Malmendier and Nagel (2011). Over the range $\lambda \in [-2, +20]$, we find the optimal value to be $\widehat{\lambda} = 10.39$. Figure 6 visualizes the objective function minimized over the grid search procedure.

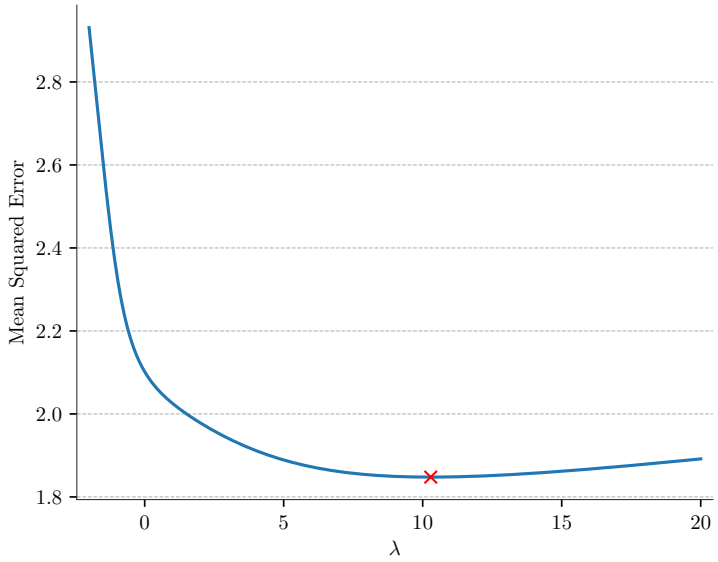


Figure 6: Objective function during the grid search procedure for estimating the λ parameter of the Malmendier and Nagel (2011) exponential weighting function.

E An Overview of the Model Setup

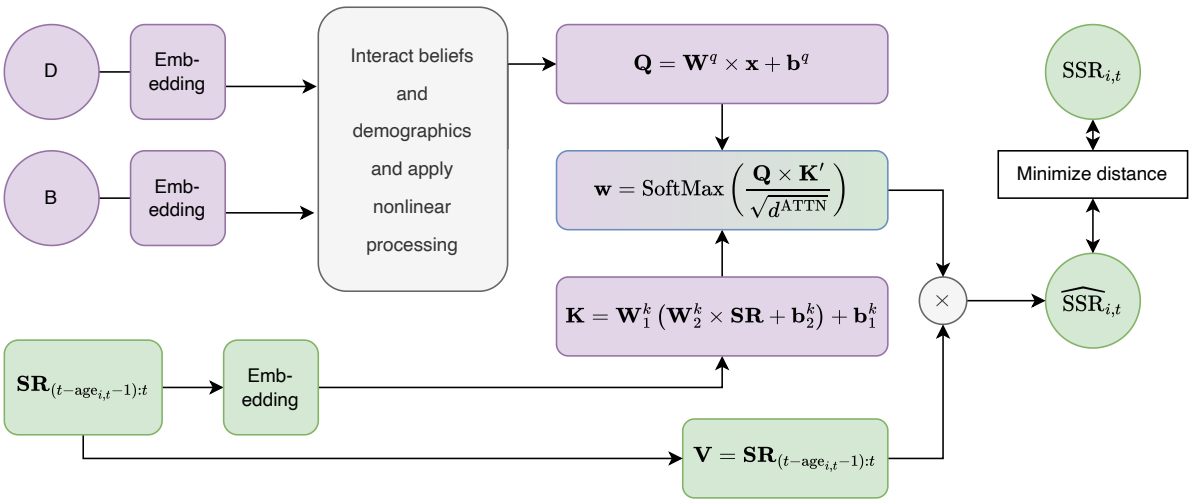


Figure 7: Model Setup

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:

$$\mathbf{x}^{\text{emb}} = \mathbf{x} \times \mathbf{W}, \tag{10}$$

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.

Interact 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 (14), 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 $\widehat{\mathbf{b}}_{i,t}$.

Extract 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.