Seven Facts About Beliefs and Portfolios

Stefano Giglio* Matteo Maggiori† Johannes Stroebel‡ Steve Utkus§

December 2018

Abstract

We administer a newly-designed belief survey to a large panel of retail investors who have substantial wealth invested in financial markets. The survey elicits beliefs that are crucial for macroeconomics and finance, and matches each respondent with administrative data on her portfolio composition and trading activity. We establish seven facts in these data: (i) beliefs are reflected into portfolio allocations, but the effect is small; (ii) this sensitivity of portfolios to beliefs is heterogeneous across agents and depends on investor wealth, attention, trading frequency, and confidence; (iii) it is hard to predict when investors trade, but conditional on trading, beliefs changes affect both the direction and the magnitude of trades; (iv) beliefs are mostly characterized by large and persistent individual heterogeneity; (v) demographic characteristics do not explain why some individuals are optimists and some are pessimists; (vi) variation in both (non-rational) expected cash flows and discount rates contribute to movements in asset prices; and (vii) expected returns and the probability of rare disasters are negatively related. Our results challenge the rational expectation framework for macro-finance, and provide important guidance for the design of alternative behavioral models.

JEL Codes: G11, G12, R30.
Keywords: Surveys, Expectations, Sentiment, Behavioral Finance, Discount Rates, Rare Disasters.

*Yale School of Management, NBER, CEPR: stefano.giglio@yale.edu.
†Department of Economics, Harvard University, NBER, CEPR: maggiori@fas.harvard.edu.
‡Stern School of Business, New York University, NBER, CEPR: johannes.stroebel@stern.nyu.edu.
§Vanguard: steve_utkus@vanguard.com.

We thank David Laibson for advice and for connecting us with Vanguard. We thank Nick Barberis, John Campbell, Ben Enke, Emmanuel Farhi, Xavier Gabaix, John Geanakoplos, Nicola Gennaioli, Andrei Shleifer, and Jeremy Stein for useful comments.

I Introduction

Researchers are increasingly turning to survey data to test macro-finance models. The unique benefit of survey data is that it can provide direct evidence on the beliefs of different agents about future returns and cash flows. These beliefs play a fundamental role in both rational expectation models and behavioral models of macro-finance, yet they are not observed in standard macro-finance data. Survey evidence fills this central gap. Despite the potential for survey data to improve our understanding of macro-finance theory, its use has become the subject of criticism on many fronts. Critics have argued that survey data is often based on small and unrepresentative samples, that it asks qualitative questions that are not informative for models, and that it has no guarantee of revealing those beliefs on which agents truly base their actions.

In this paper, we provide new evidence on the link between beliefs elicited through surveys and real actions taken by survey respondents. To do this, we administered a newly-designed belief survey to a large panel of retail investors who have substantial wealth invested in financial markets. The survey was explicitly designed to address the primary prevailing criticisms of existing survey data. The survey elicits investors' beliefs about future stock returns, GDP growth, and bonds; the survey design trades off asking quantitative questions about moments that are crucial for economic theory with keeping the questions sufficiently simple that they could be answered by non-specialist investors. The survey was administered to U.S. based retail clients of Vanguard, one of the world's largest investment management firms with about $5.1 trillion in assets under management. The respondents are a sample of individuals relevant for macro-finance models: they participate in the stock and bond markets and have substantial assets. The survey has been administered every two month since February 2017. The first ten waves have generated more than 2,000 responses per wave. Many individuals respond to multiple survey waves, generating a substantial panel component to our data. We are able to link each respondent's survey answers to her investment portfolio, for which we observe all holdings and transactions in high quality administrative data. This allows us to explore the relationship between the elicited beliefs and real-world high-stakes investment behavior. While survey data is by its very nature imperfect, we believe that the design and sample of our survey allows us to overcome many of the common criticisms of existing survey data.

Our most general finding is that survey data, despite all its shortcomings, is very informative to understand how individuals form their portfolios. We find strong and statistically significant relationships between expectations and portfolios both across individuals and within individuals over time. In this sense, we conclude that survey-based evidence is "here to stay" and that it has to be confronted and explained by theory going forward. While this view has been increasingly accepted by the behavioral finance literature, a large part of the profession has remained on the fence about the eventual usefulness of survey data. By building a large-scale quantitative survey and directly linking it to portfolios, we establish a rich set of facts that we hope will push the conversation forward.

We organize our empirical results around seven facts. These facts highlight empirical patterns

both within our survey (e.g., in the relation between beliefs about different topics) and between beliefs and portfolios that are particularly informative about macro-finance models. For each fact, we discuss the implications for macro-finance theory, pointing out how it challenges the rational-expectation paradigm and provides further guidance for the existing behavioral models.

Before describing our facts, let us provide more details on the survey. Every two months, a random subset of Vanguard clients receive an email from Vanguard asking them to answer a brief survey. Previous survey respondents are usually recontacted in addition to a refreshed sample of randomly chosen clients. If potential respondents click on the survey link, they are redirected to a Vanguard-branded online survey portal. They are then asked three blocks of questions: one block is about the stock market, one block is about GDP growth, and a final block is about interest rates. The questions are short and concrete, and do not focus on hypotheticals or lotteries; the average time to complete the survey is about ten minutes. The questions on stock returns and GDP growth are very similar: we ask about subjective expected stock returns over the next year (over the next three years for GDP growth), we elicit the distribution of these outcomes by asking respondents to put subjective probabilities into various buckets of possible outcomes, and we ask about expected stock returns and GDP growth over the next ten years. We randomize the order in which respondents see these two blocks of questions. The questions about interest rates always come last and ask respondents about the expected return of a ten-year bond over the next year and about the level of interest rates they expect to prevail one year in the future. After each block of questions, we ask the respondents whether they feel that they understood the questions in the previous block and how confident they are about their answers.

The average response rate for newly contacted individuals is about 4%. For previous respondents, the re-response rate is about 20% in each subsequent survey wave. As a result, we obtain an unbalanced panel of responses. We observe many instances of multiple responses from the same individuals over time. On average, we obtain 2,000 responses per wave. For each respondent, we observe administrative data on their entire portfolio history at Vanguard. This includes all positions across all accounts and all trading activity down to the security and transaction level. We have no information on the investments that respondents hold outside of Vanguard. For each person we observe demographic information as well as details about the number of years for which she has been a client, and how often she logs into the Vanguard website to monitor her investments or trade. Importantly, our survey respondents are relatively wealthy, with mean assets of \$467k and median assets of \$201k.

Armed with this joint panel of survey measures of subjective beliefs and administrative portfolio data, we proceed to establish seven facts that we hope will help shape new advances in understanding and modeling portfolio formation and macro-finance.

Fact 1 establishes that there is a strong relationship between beliefs and portfolios, but that the average sensitivity of an investor's portfolio share in equity to that investor's expected stock market returns and perception of stock market risk is an order of magnitude lower than predicted by benchmark frictionless macro-finance models. We establish this fact by regressing the contem-

poraneous share of each individual's portfolio that is invested in equities on her subjective expectations for stock market returns, controlling for the perceived variance of returns (constructed using the survey responses about the returns distribution) and probability of a large drop of the market. We find a statistically strong, but economically small, relationship. Controlling for the two measures of risk, a 1% increase in expected returns over the next year increases the equity share by 0.7%; this relationship is not driven by measurement errors in beliefs and the associated attenuation bias. We find that the perceived variance of stock returns has both an economically and statistically weak relationship with portfolios, and that a better measure of risk is the probability of a large stock market drop (a rare disaster). However, even this tail probability has a muted effect on portfolios with an increase of 1% in the probability of a stock market drop of more than 30% only decreasing the portfolio equity share by 0.1%. Both these sensitivities are an order of magnitude too small compared to benchmark frictionless portfolio models.

This relatively small response of equity shares to beliefs about stock returns is qualitatively and quantitatively consistent with evidence documented across a number of other studies that link equity market participation and equity shares to expected stock market returns (e.g., Vissing-Jorgensen, 2003; Dominitz and Manski, 2007; Kézdi and Willis, 2011; Amromin and Sharpe, 2013; Ameriks et al., 2016; Drerup, Enke and Von Gaudecker, 2017), though many of these studies had to rely on self-reported equity shares or a binary indicator for holding stocks rather than administrative data. Our contribution in this paper is to not only establish this fact on a large sample of high quality data, but also to document the empirical importance of a number of possible explanations for the relatively low observed sensitivity. Specifically, we quantify (i) the role of attenuation bias due to measurement error in beliefs, (ii) the role of infrequent trading, and (iii) the role of respondents' confidence in their own beliefs. We also explore the margins of substitution that drive the baseline result, analyzing whether more optimistic individuals substitute towards equity from cash or from fixed income investments. Finally, we analyze the role of other moments of the belief distribution (e.g., tail events), the role of stock market expectations over different horizons, and the role of beliefs about bond returns and GDP growth.

We document that wealth, trading frequency, and investor confidence are the most important sources of heterogeneity in the sensitivity of portfolios to beliefs. We summarize our findings in Fact 2: the sensitivity of portfolios to beliefs is increasing in wealth; it is also increasing in investors' trading frequency, attention, and confidence in their own beliefs. An idealized investor who is wealthy, pays attention to the portfolio, trades often, and is confident in her beliefs exhibits a sensitivity that is about four times larger than the average sensitivity, and, as a result, much closer to the range of sensitivities generated by frictionless benchmark models.

We next investigate the role that belief changes play in explaining trading activity. Fact 3 establishes that belief changes have minimal explanatory power for predicting when trading occurs (the extensive margin), but explain both the direction and magnitude of the trading conditional on a trade occurring (the intensive margin). At a broad level, these results confirm that beliefs as elicited by surveys play an important role in explaining actual investment decisions such as trad-

ing. This is complementary to the analysis of portfolio shares across individuals, since trading focuses on the changes rather than the levels of the portfolios. We find it natural that predicting whether trading occurs over a short horizon is difficult, since trading is often caused by factors other than beliefs about returns and risk such as liquidity needs induced by life events. This is consistent with models of infrequent trading that generate either endogenously or exogenously a flat hazard function of trading based on belief changes. We find it interesting and encouraging, however, that conditional on a trade occurring, beliefs play an important role in explaining both the direction, a buy or a sale of equity, and the magnitude.

At a fundamental level, this first set of results challenges the notion, most forcefully articulated by Cochrane (2011, 2017), that survey data and the associated critiques of rational expectations model of macro-finance can simply be discarded because respondents do not understand the questions or because their answers reveal beliefs that do not actually correspond to their investments. While all belief surveys are imperfect and subject to valid concerns about multiple dimensions of measurement error, we clearly document that our survey answers are relevant to understanding individual behavior. As a result, our findings support the view, expressed most prominently by Greenwood and Shleifer (2014) and Barberis et al. (2015), that the challenges that survey data exposes for macro-finance models are to be taken seriously.

Our results also provide guidance for refining existing behavioral models of macro-finance that have being built explicitly to match newly uncovered survey evidence (see, for example, the recent work in this direction by Barberis et al., 2015). These models often embed a portfolio demand with a similar sensitivity to beliefs as rational-expectations models, but rather than imposing rational expectations, they impose expectations that are more consistent with survey evidence. These models' predictions for asset prices rest on two building blocks: (i) that beliefs are volatile and change over time, and (ii) that individual portfolios react strongly to changes in these beliefs. Our evidence shows that, for the average investor, the second channel is much weaker than implied by common calibrations of these models. Indeed, for the majority of investors, infrequent trading and lack of confidence in beliefs substantially reduces the pass-through from beliefs to portfolios. As a result, the frictionless behavioral models, in their current form, are likely to overstate the power of expectation changes in explaining asset price movements. Our research opens a new challenge to improve these models: a successful model has to match both expectations and portfolio dynamics jointly with asset prices, while the current literature, which is already a major advancement, focuses primarily on matching expectations and asset prices.

The richness of our data helps us identify possible ways to improve the fit of models to the data. As highlighted by Barberis (2018), behavioral macro-finance models often feature two types of investors (e.g., rational arbitrageurs and behavioral investors), each modeled with its own demand for stocks. Asset prices are then determined by the dynamics of expectations of the behavioral agents, modulated by the relative wealth shares of the two agents and their relative demand sensitivities. One possible avenue for the behavioral agents' expectations to have a quantitative impact on the equilibrium price is to have the behavioral group be substantially wealthier (in

the aggregate) than the rational arbitrageurs. This would induce the changes in demand by the behavioral agents, even with a relatively insensitive demand, to be larger for the arbitrageurs to absorb and therefore generate quantitatively realistic asset price movements. Alternatively, one could make the demand of the behavioral agents more sensitive than that of the arbitrageurs. Belief surveys like ours could be used in the future to explore the correlation between deviations from rationality and the sensitivity of beliefs and portfolio allocations.

After analyzing the relationship between investors' beliefs and their portfolios and trading behavior, we further explore the richness of our beliefs data. We start by decomposing the variation in beliefs across individuals and over time, since this plays a crucial role in macro-finance models. Fact 4 establishes that individual beliefs are mostly characterized by heterogenous individual fixed effects. In our panel, between 50% and 60% of all beliefs' variation is due to individual fixed effects, 1% to common time series variation, and the rest is accounted for by idiosyncratic individual variation or measurement error. Some individuals are optimistic and some are pessimistic; their beliefs are persistent and very far apart. We find that optimism (and pessimism) are correlated across questions: if a respondent is optimistic about the stock market, she is also optimistic about economic growth and vice versa. While there is some co-movement across all individuals over time, the associated shift in the average belief only accounts for a small part of the variation in beliefs across survey answers.

This provides a powerful reduced-form characterization of beliefs across agents which can be compared against the assumptions of macro-finance models. Consider two (stylized) characterization of existing theoretical models. At one extreme, consider a model in which all individuals hold the same beliefs and the only variation is in the time series. At the other extreme, consider instead a model in which agents have permanent and constant differences in beliefs. The first class of models generates variation in asset prices directly with swings in individual beliefs over time, which then translate immediately in changes in the representative agent's beliefs. The second class of models generates variation in asset prices by movements in the wealth-weighted aggregate belief. Intuitively, this requires several steps: individual heterogeneity in beliefs induces heterogeneity in portfolios; shocks over time redistribute wealth across individuals given the differences in portfolios, ultimately changing the composition of aggregate asset demand and affecting asset prices.[1]

Our research highlights that both classes of models miss core elements of the data. The first class of models is silent on a major feature of the data: the individual heterogeneity; it obtains all its action from what appears to be a small component (1%) of the total variation in beliefs, the time-series. In light of our results, we believe that the literature would benefit from explicitly modeling the aggregation of the heterogeneity in individual beliefs. The second class of models captures the individual static heterogeneity, but gets all its action by having this heterogeneity strongly reflected in portfolios. Our Fact 1 highlights that this second step is challenging: belief

[1]For example, optimists have portfolios that load more than pessimists' portfolios on good economic outcomes; if these outcomes are realized, then optimists relative wealth increases compared to pessimists.

differences are reflected in portfolios in a statistically reliable manner, but their impact on portfolios is quantitatively limited. If individuals with vastly different beliefs do not on average have vastly different portfolios, then shocks do not substantially redistribute wealth between optimists and pessimists, thus vastly reducing the variation in aggregate wealth weighted belief.

Having established the importance of individual fixed effects in explaining the panel-variation in beliefs, it is a natural question to ask which observables explain why some individuals are optimists and some are pessimist. Fact 5 establishes that the heterogeneity in beliefs is not explained by observable demographic characteristics such as gender, age, wealth, attention (days with visit to the Vanguard website), confidence, and geographic location. These characteristics have a strong statistical relationship with beliefs but their explanatory power is extremely limited, a fact that we dub the *belief disconnect*.[2] We provide evidence that this is not the result of measurement error in eliciting beliefs with the survey. A likely explanation is that these individual fixed effects reflect combinations of individual characteristics and experiences that economic research is yet to discover. Indeed, it is an exciting challenge for future research to establish the determinants of the large differences in individual beliefs.

Another important theoretical question our survey speaks to is what are the most important variables investors disagree about. Existing theoretical models vary according to whether investor disagree about cash flows or expected returns (see, for example, Cutler, Poterba and Summers, 1990; De Long et al., 1990*a*,*b*; Barberis et al., 2015). Using the results of our survey (which elicits expectations about both objects), we document the importance of taking into account the implicit links between these variables. In particular, the Campbell and Shiller (1988) identity links expected returns, cash flows, and expected future price-dividend ratios. If investors disagree about one element of the decomposition, they must disagree conversely on the other elements as well. We find in Fact 6 that investors disagree about all three of these elements. Across individuals, higher expected cash flows are partially offset by higher expected returns (i.e., higher discount rates) in both the short and long run, but not enough to be consistent with the same current prices (which is observed, so there can be no disagreement about it). Implicitly, therefore, investors must also disagree about the long-term expected price-dividend ratios. In addition, while our sample offers extremely limited aggregate time series variation, we find evidence consistent with expected return extrapolation as in Greenwood and Shleifer (2014). Overall, our results show that beliefs about both cash-flow and expected returns are important to understand the data and that, at least for quantitative evaluations, their dynamics have to be modeled jointly.

Our last fact, Fact 7, is that when individuals expect large stock market declines ("rare disasters") to occur with higher probability, they also expect stock market returns to be lower. This holds both across individuals and within individuals over time. An important strand of the macro-

[2]Our belief disconnect results are closely related to the literature linking expectations to socioeconomic status and personal experience. It is not uncommon in this literature to find strong statistical relationships but low explanatory power between expectations and variables such as wealth, gender, and place of birth, location, own past experience, or friends' past experiences (e.g., Vissing-Jorgensen, 2003; Choi et al., 2009; Malmendier and Nagel, 2011; Greenwood and Shleifer, 2014; Kuchler and Zafar, 2015; Das, Kuhnen and Nagel, 2017; Bailey et al., 2017, 2018).

finance literature has emphasized that expectations of rare but potentially catastrophic negative events, sometimes referred to as rare disasters, can help explain portfolio holdings and asset prices (Rietz, 1988; Barro, 2006; Gabaix, 2012). This literature stresses that the expected probability and size of rare disasters is one of the most important moments to understand risks and returns.

In standard equilibrium models with rare disasters, expected returns and the probability of disaster are positively related. The intuition is that a higher probability of disaster induces individuals to demand for a higher compensation for holding the stock market and this makes equilibrium expected returns higher. The relationship in the data appears with an opposite sign relative to the theory, thus posing a challenge to rational-expectations models of rare disasters.

However, our results are not inconsistent with the importance of rare disasters for macro-finance more generally. In fact, it is likely that modifications to existing models that incorporate at least some behavioral agents can more closely align model predictions and data. One possible approach would be to model some agents as over-confident or as thinking in partial equilibrium. When these agents believe that rare disasters have become more likely, they do not take into account that other agents might think the same, and therefore do not incorporate the fact that stock prices might already be lower to account for the increased risk. Rather, they think that prices are still high, and expect returns to be lower going forward. They therefore sell stocks to reduce their exposure, by more than they would if they anticipated that the compensation for holding on to the risk is increasing in equilibrium. In turn, this force may induce more selling pressure (and therefore a bigger price effect) in response to an increase in the disaster probability. Of course, there might also be rational agents who take the opposite side of these trades; these rational agents require higher returns to compensate for the risk and to absorb the stocks that the behavioral agents are selling. It is beyond the scope of this paper to provide a full model of the behavioral approach to rare disasters, but we expect that it will prove an interesting area of future research.

II Survey Description

To explore the structure of investors' beliefs and the relationship between those beliefs and investors' portfolio allocations, we designed a new online survey of U.S.-based individual investors who hold accounts at Vanguard, one of the world's largest investment management firms with about $5.1 trillion in assets under management. We first provide a high-level overview of the survey questions. The exact phrasing and the survey interface are presented in the Appendix. We then explore the survey sample, the response rates, and the demographics of respondents and non-respondents, which allow us to analyze the dimensions of selection into responding.

II.A Survey Design

The survey includes questions on three broad topics: expected stock market returns, expected GDP growth rates, and expected bond returns. In the implementation of the survey, we randomized whether individuals were first asked about their expectations on stock returns or GDP growth rates. The questions on bond returns were always asked last.

Expected Stock Market Returns. In the first set of questions, we asked respondents about their expectations for the return of the U.S. stock market. We elicited point estimates for the expected annualized returns over the coming year and the coming ten years, as well as subjective probabilities that the return over the next year would fall into one of five buckets: less than -30%, between -30% and -10%, between -10% and 30%, between 30% and 40%, and more than 40%. These buckets were chosen such that most of the historical mass would fall into the middle bucket, allowing individuals to express their expected beliefs about the probabilities of relatively extreme events.[3] The ordering of the buckets (i.e., lowest to highest vs. highest to lowest) was randomized across survey respondents, and the survey enforced the assigned probabilities to add up to 100%. As shown in the Appendix, to help individuals answer the questions, the survey interface also showed histograms to represent the survey responses as they were entered.

Expected Real GDP Growth Rates. In the second set of questions, we asked respondents about their expectations for the annualized growth rate of real GDP. We elicited point estimates for the expected growth rates over the coming three years and the coming ten years. In addition, we asked about the probabilities that individuals assigned to the annualized GDP growth rate over the coming three years falling into one of five buckets: less than -3%, between -3% and 0%, between 0% and 3%, between 3% and 9% and more than 9%; as before, these were chosen to allow individuals to report their expectations over relatively extreme positive and negative events. As with the buckets for expected stock returns, we randomized the ordering of buckets across respondents in the same way (high to low or low to high).

Expected Bond Returns. The final set of questions elicited respondents' expectations about U.S. government bond returns and interest rate changes. We first asked respondents for their expectations of the 1-year return of a 10-year zero coupon U.S. government bond. In addition, we asked individuals about their expectations for the yield curve in one year: in particular, we asked what respondents expected the yields-to-maturity of 1-year, 5-year and 10-year zero coupon bonds to be one year in the future. Prior to answering the question, individuals were shown the contemporaneous yield curve to familiarize them with the concept.

Difficulty and Confidence. At the end of every block of questions (stock, GDP, and bonds), the survey asked individuals how confident they were about their answers (a five-point scale from "not at all confident" to "extremely confident"), and how difficult they found the questions (a five-point scale from "not at all difficult" to "extremely difficult").

[3] Between 1927 and 2014, annualized stock market returns fell into the lowest bucket 3.7% of the time, the second bucket 11.6% of the time, the middle bucket 65.9% of the time, the fourth bucket 13.8% of the time, and the highest bucket 5% of the time.

II.B Survey Sample and Response Rate

The online survey is conducted every two months among U.S.-based Vanguard investors. In this paper, we explore the first ten waves of the survey, starting in February 2017 and going through August 2018. In the first wave, 40,000 Vanguard clients were invited by email to participate in the survey. These clients were randomly selected such that 80% of the sample was drawn from clients in Vanguard's RIG division (which includes retail portfolios and individually-managed tax-advantaged retirement accounts such as IRAs), and 20% from clients of Vanguard's IIG division (which includes institutionally-managed retirement accounts).[4] Additional requirements to be potentially invited to participate in the survey are that users: (i) had to have opted into receiving Vanguard statements via email, (ii) had to be between the ages of 21 and 74, and (iii) had to have total Vanguard assets of at least $10,000.[5] If individuals respond to the survey in any wave, they are recontacted in every subsequent wave (unless they opt out). Those individuals who do not respond to the first wave in which they are contacted are then recontacted in two subsequent waves (unless they opt out). If they respond in neither of these waves, they are dropped from the sample. In the second wave, an additional 25,000 clients were invited to participate (in addition to those carried over from wave 1). Waves 3 to 5 invited 13,000 new clients each; from wave 6 onwards, the steady-state number of new clients contacted in each wave was increased to 14,500. The Vanguard clients that are newly contacted each wave are also selected using the criteria described above.

Response Rates. Figure I shows the response rates for each of the ten waves, where we count only fully completed surveys as "responses."[6] The orange circle-line shows that the response rate among individuals contacted for the first time was relatively stable at 3%-4% across waves. The response rates among individuals who were previously contacted but had not yet responded, given by the blue diamond-line, were between 1.5% and 2% across waves. The green square-line shows the response rate among individuals who had responded to at least one previous survey (recall that we re-contact these individuals in all subsequent waves, unless they opt out of receiving further emails). The steady-state re-response rate of these individuals is between 15% and 20%. It declines somewhat over time, though much of this is driven by compositional effects: in later survey waves, the average time since the last response of individuals who have previously responded is increasing.

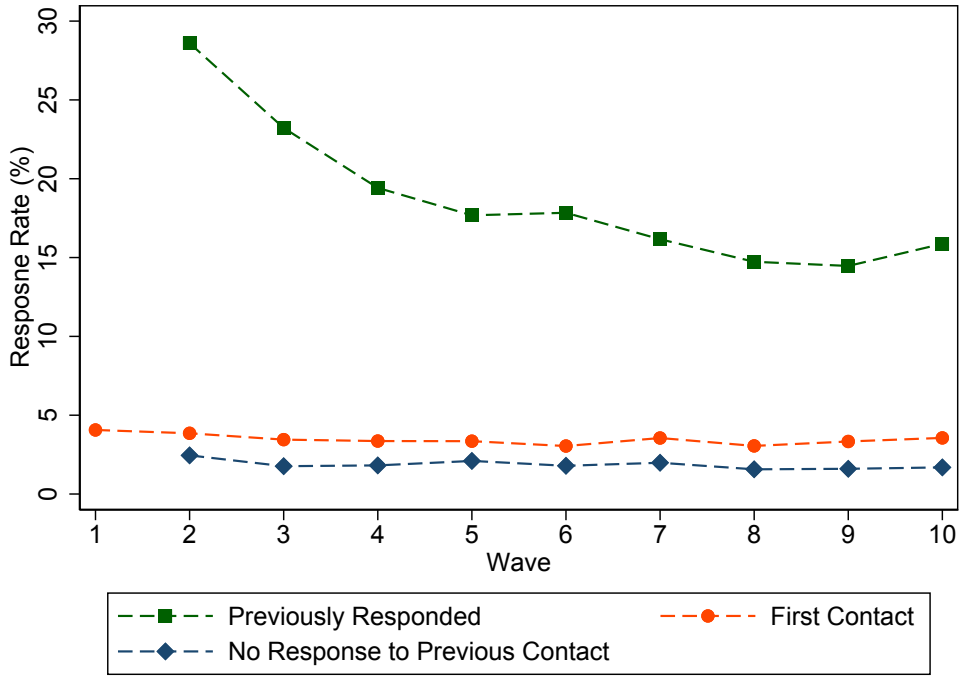
These response rates translate into about 2,000 survey responses per wave; across the ten waves, we received 20,207 complete responses. The top panel of Figure II reports the number of responses we received in each wave, with different shades of gray keeping track of the first wave

[4] In the selection of this sample, individuals who hold both types of accounts at Vanguard were counted as being part of the group in which they hold the most assets.

[5] In addition, some Vanguard clients that had been concurrently contacted by studies other than ours were excluded from the pool of possible clients for our survey.

[6] In the analysis below, we drop a number of partially completed surveys. Most of these answer at most one question before exiting the survey. These are treated as non-responses in the statistics in this section.

Figure I: Survey Response Rates

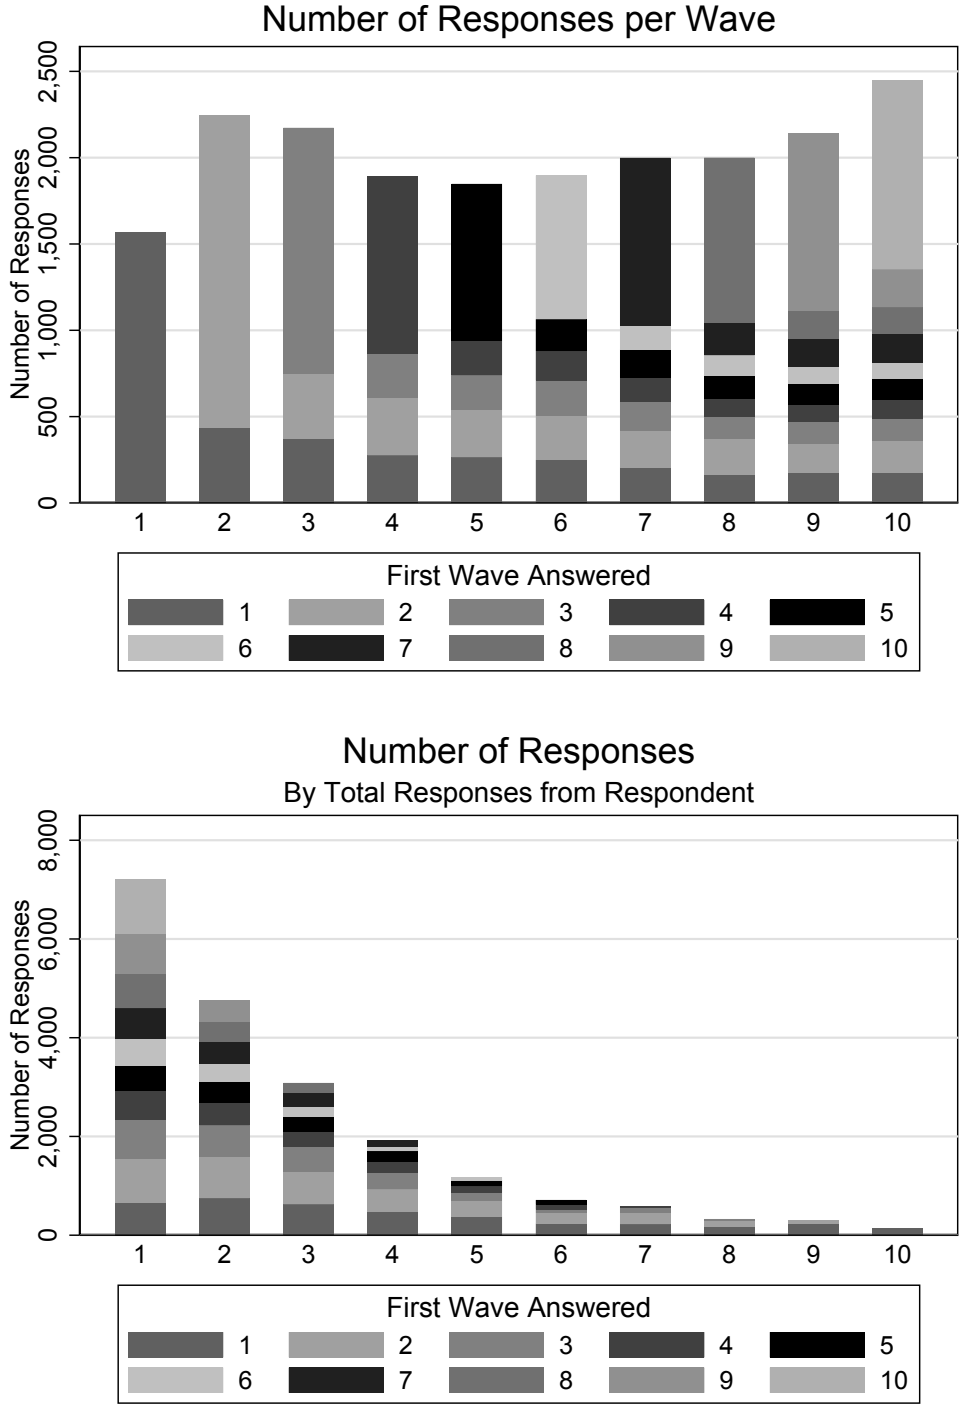


Note: Figure shows the response rates to the Vanguard Investor Expectations Survey in each of the ten waves between February 2017 and August 2018. The orange line (circle) shows the response rates for individuals contacted for the first time. The blue line (diamonds) shows the response rates for individuals that were contacted in previous waves, but who had not yet responded. The green line (squares) shows the response rate for individuals who had previously responded.

in which an individual responded. The figure shows that, starting in wave 5, we receive more responses from individuals who are re-respondents than we receive from individuals who are responding for the first time. The bottom panel reports a histogram of the number of responses we obtained from each respondent. The bottom panel shows that about 35% of responses come from individuals who have responded to one survey only (though some of these may still end up responding to future surveys). Over 25% of responses come from individuals who have responded to at least four survey waves, and 10% of responses come from individuals who have responded to at least six survey waves. As a result of the relatively high re-response rate, our survey data contain a large and growing panel dimension.

Demographics of Respondents. In addition to the information provided by survey respondents directly, we observe data on demographic characteristics such as age, gender, and location for all individuals that were invited to take the survey. We also observe administrative data on their portfolios, including details on total portfolio value, portfolio allocation, portfolio returns, and trading behavior. We combine the CUSIP-level information on holdings with information from Morningstar to calculate the portfolio share held in equities, fixed income instruments, cash, and other investments. Table I shows summary statistics on these demographics and portfolio charac-

Figure II: Number of Responses by Wave



Note: Figure shows number of responses to the Vanguard Investor Expectation Survey. The top panel shows the number of responses per wave. The bottom panel shows the total number of responses separately by how many different survey a person has responded to. In both panels, differently colored bars correspond to different waves in which these individuals answered.

Table I: Demographics - Survey Respondents and Non-Respondents

	Survey Respondents							Non-respondents	Difference
	Mean	SD	P10	P25	P50	P75	P90	Mean	
Age (years)	58.7	13.4	37	50	61	69	74	51.4	7.39***
Male	0.68	0.5	0	0	1	1	1	0.54	0.13***
Region									
Northeast	0.23	0.42	0	0	0	0	1	0.24	-0.01***
Midwest	0.21	0.41	0	0	0	0	1	0.19	0.02***
South	0.31	0.46	0	0	0	1	1	0.31	0.00
West	0.25	0.44	0	0	0	1	1	0.25	0.00
Total Vanguard Wealth ($)	467,139	778,740	26,294	67,618	201,577	545,689	1,144,906	249,010	218,130***
Length of Vanguard Relationship (Years)	15.07	8.82	3	7	15	21	27	12.83	2.24***
Active Trades / Month	1.44	3.86	0	0.10	0.44	1.51	3.46	0.88	0.57***
Monthly Portfolio Turnover (%)	2.25	6.47	0.00	0.10	0.62	2.15	5.00	1.95	0.30***
Days with Log-Ins / Month	4.40	6.08	0.05	0.38	1.43	5.86	14.76	0.88	3.52***
Total Time Spent / Month (Minutes)	32.3	77.3	0.50	3.46	11.17	31.36	76.66	4.9	27.34***
Portfolio Shares (%)									
Equity	68.5	26.4	31.3	53.6	72.3	90.0	100.0	71.3	-2.85***
Fixed Income	21.3	21.9	0.0	1.5	15.7	33.8	51.0	17.8	3.54***
Cash	8.7	18.9	0.0	0.0	1.1	6.7	26.2	9.0	-0.29
Other/Unknown	1.5	7.3	0.0	0.0	0.0	0.0	2.1	1.9	-0.40***

Note: Table shows summary statistics on both the respondents and non-respondents to the survey. We present respondent age, gender, location, total wealth at Vanguard and Length of Vanguard relationship as of August 2018. The other variables are presented as monthly averages between January 2017 and August 2018, covering the period of the ten survey waves.

teristics for both the sample of survey respondents and the sample of non-respondents.

Like in all surveys, our survey respondents are a selected sample from the pool of individuals invited to take the survey. One strength of our setting is that we observe administrative data on both respondents and non-respondents, allowing us to quantify the selection. Our average respondent is about 59 years old, which is about 7.4 years older than the average non-respondent. Respondents are substantially wealthier than non-respondents, with an average wealth held at Vanguard of $467k for respondents relative to $249k for non-respondents. 68% of respondents are male, while this number is 54% among non-respondents. The average respondent has been a Vanguard client for over 15 years, relative to about 13 years for non-respondents. Respondents are also likely to transact more frequently (an average of 1.4 trades vs. 0.9 trades a month during our sample), and their monthly portfolio turnover is larger. In fact, respondents turn over about 2.25% of their portfolio every month, while non-respondents have a portfolio turnover of 1.95% per month. Respondents are also more likely to log into their Vanguard accounts (an average of 4.4 respondents vs. 0.9 days with log-ins per month for non-respondents).[7] Portfolio allocations of respondents and non-respondents are relatively similar, though respondents are somewhat more likely to hold fixed income assets in their portfolios (21.3% vs. 17.8% fixed income share) and somewhat less likely to hold equity (equity share of 68.5% for respondents vs. 71.3% for non-respondents). The Appendix shows that all of the univariate differences between respondents and non-respondents documented in Table I also persist qualitatively in a multivariate analysis.

[7]Responding to the survey does not involve respondents logging into their Vanguard accounts; the process of answering the survey does therefore not lead to a mechanical increase in log-ins.

Overall, while the sample of respondents is not entirely representative of the overall Vanguard population, the direction of the selection is potentially helpful: in particular, since asset prices are likely to be disproportionately affected by the beliefs of wealthier people, the sample-skew towards wealthy Vanguard clients over-represents this particularly interesting group of individuals.

While there are differences between non-respondents and respondents, there is also large heterogeneity across the sample of respondents. While the average age is 59 years, this ranges from 37 years at the 10th percentile of the distribution to 74 years at the 90th percentile. The 10-90 percentile range of assets held at Vanguard is $26k to $1.14m. The respondents come from all over the United States, with 31% coming from the Southern Census Region, 23% coming from the North Eastern Census Region, and 21% coming from the Midwestern Census Region. Activity on the Vanguard site, both in terms of log-ins and in terms of trading activity, also differs substantially across our survey respondents. At the 10th percentile of the distribution, respondents spend an average of 0.5 minutes per month on the Vanguard side, while this number is 76.7 minutes at the 90th percentile of the distribution. There are also substantial cross-sectional differences in portfolio allocations among our respondents. At the 10th percentile of the distribution, the equity share is 31.3%, while at the 90th percentile it is 100%. Similarly, at the 10th percentile of the distribution, respondents have none of their portfolios invested in fixed income securities, while at the 90th percentile of the distribution, respondents have more than 50% invested in fixed income assets. While the average cash allocation in respondents' portfolios is 8.7%, this ranges from 0% at the 10th percentile of the distribution to 26.2% at the 90th percentile of the distribution.

III Survey Responses: Summary Statistics

Table II shows summary statistics across the 20,208 survey responses that completed the survey flow entirely.[8] The average respondent takes about 517 seconds (8.6 minutes) to answer all survey questions. The 10-90 percentile range for the total time to respond is 433 seconds (7.2 minutes) to 834 seconds (13.9 minutes). Therefore, all respondents spend a significant amount of time answering the question, rather than carelessly clicking through the survey; this is consistent with the

[8] Even among the responses that completed the full survey flow, some respondents skipped a few questions. We verified that redoing our analysis to the sample of respondents to those who provided answers to every question does not affect our conclusions. We also explore the prevalence of extreme outlier responses, such as individuals reporting that the expected return on the U.S. stock market over the coming year was 400% or -100%. Since such outlier responses have an extreme effect on our analysis, in our baseline analysis, we set extreme outlier answers for each unbounded expectation question equal to missing. In particular, we set answers to missing if individuals report expected 1-year stock returns outside of the range -20% to +20%. The accepted ranges for the annualized expected return on the stock market over the next ten years was -10% to +20%. For the annualized GDP growth rate over the next three and ten years, we allow answers in the range -15% to +15%. For expected bond returns over the coming year we allow a range from -15% to +15%. For each question, this approach to dealing with extreme responses sets about 2% of answers to missing; it is often the case that the same individuals report multiple answers outside the accepted ranges. Naturally, there are some critical judgment calls involved in selecting these cutoffs, which involves trading off wanting to capture true extreme beliefs from not wanting to include answers from individuals who misunderstood the question or the units. We have done extensive robustness checks to confirm that our results are robust to a wide range of choices for the cutoff values. We also confirmed that the results are robust to winsorizing extreme answers rather than setting them equal to zero and to to dropping all answers of individuals who report extreme answers to at least one questions.

non-compensated nature of the survey requiring a certain intrinsic interest from participants.

The average expected stock market return over the coming year was given as 5.23%, while the average annualized expected return over the next ten years was given as 6.32%. These numbers compare to an average annual stock market return of 11.9% between 1927 and 2014. There is substantial heterogeneity in these numbers across responses. At the 10th percentile of the distribution, individuals reported a 1-year expected stock return of 0.5%, while at the 90th percentile they expected a return of 10%. The across-responses standard deviation of expectations of 1-year returns is nearly twice as large, at 5.32%, as the standard deviation of expectations of annualized 10-year returns, which is 2.96%. This suggests that individuals expect some medium-run mean reversion in stock returns. Panels A and B of Figure III show the distributions of responses to these two questions. When we ask individuals about their expectations of annualized GDP growth, the mean across the next three years and the next ten years is quite similar, at 2.83% and 2.95% respectively. The distribution of answers over the two horizons also have a similarly sizable standard deviation at about 1.9%. Panels C and D of Figure III show the full distribution of answers for the annualized GDP growth over the next three and ten years, respectively. Overall, this substantial heterogeneity in beliefs about stock market returns and GDP growth across survey respondents is consistent with the patterns in many of the surveys that we reviewed in the introduction; in Section VI we exploit the panel dimension of our data to explore the reasons for this heterogeneity.

Table II also shows that respondents put substantial probabilities on relatively large short-run stock market declines and GDP declines. The average individual assigns a 5.1% chance to the 1-year return of the stock market being less than -30%; the median respondent assigns a 3% chance of such an event (these numbers are quite close to the fraction of years between 1927-2014 that had such a drop, which was 3.7%).[9] As with the the other answers, there is substantial across-answer heterogeneity. Individuals at the 25th percentile of the distribution assign a zero percent chance to returns lower than -30%, while at the 90th percentile of the distribution, individuals assign a 10% probability to such events. Panel E of Figure III shows the full distribution of answers; the bunching of reported probabilities at multiples of five is consistent with reporting behavior across many surveys. Similarly, individuals assign an average (median) probability of an annualized decline in GDP of more than 3% over the coming three years of 4.5% (2%), though at the 25th percentile of the cross-sectional distribution individuals assign zero probability to such an event.

Individuals were also asked about the expected one-year return of 10-year zero coupon bonds. The average response was 1.38%, with a 10-90 percentile range of -2% to 4% (see Panel F of Figure III for the full distribution of answers). Lastly, individuals were asked about their expectations for the yield curve in one year. Table II provides summary statistics on these answers for completeness, though they are not the subject of this paper. On average, individuals expected the yield

[9] Recall that we randomize the ordering whether the buckets are ordered smallest-to-largest or largest-to-smallest. We find that this ordering does have a small effect on the probabilities assigned to the smallest bucket. Among those individuals where the smallest bucket is presented first, they assign 5.8% probability to a return lower than -30%, while this number is 4.4% for individuals who are shown the largest bucket first. We verified that the conclusions in this paper are independent of the order in which the buckets are presented.

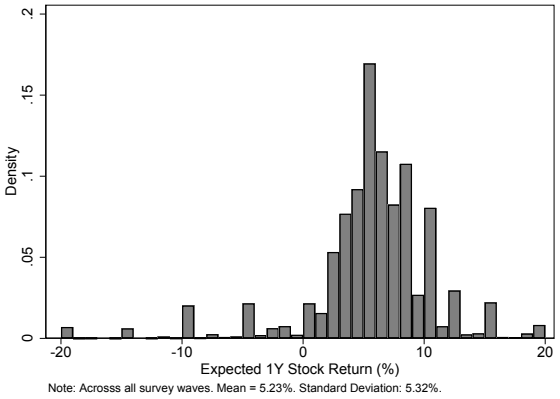
Table II: Summary Statistics - Survey Responses

	Mean	SD	P10	P25	P50	P75	P90
Expected Stock Returns							
Expected 1Y Stock Return (%)	5.23	5.32	0.5	4	5	8	10
Expected 10Y Stock Return (% p.a.)	6.32	2.96	3	5	6	8	10
Probability 1Y Stock Return in Bucket (%)							
Less than -30%	5.1	8.2	0	0	3	5	10
-30% to -10%	13.6	13.1	0	5	10	20	30
-10% to 30%	70.8	22.5	40	60	75	90	100
30% to 40%	7.6	10.8	0	0	5	10	20
More than 40%	2.9	6.5	0	0	0	5	10
Expected GDP Growth							
Expected 3Y GDP Growth (% p.a.)	2.83	1.92	1.5	2	2.5	3	4
Expected 10Y GDP Growth (% p.a.)	2.95	1.94	1.5	2	2.5	3	5
Probability p.a. 3Y GDP Growth in Bucket (%)							
Less than -3%	4.5	7.7	0	0	2	5	10
-3% to 0%	12.0	12.0	0	3	10	20	25
0% to 3%	57.8	26.1	20	40	60	80	90
3% to 9%	22.4	23.3	0	5	15	30	60
More than 9%	3.3	8.7	0	0	0	5	10
Expected Bond Returns							
Expected 1Y Return of 10Y zero coupon bond (%)	1.38	3.13	-2	1	2	3	4
Expected yield-to-maturity of zero coupon in 1Y (%)							
1Y zero coupon bond	1.83	0.96	1	1	2	2.2	3
5Y zero coupon bond	2.59	1.08	1.5	2	2.5	3	4
10Y zero coupon bond	3.30	1.42	2	2.5	3	4	5
Difficulty ("Not at all difficult", ..., "Very difficult")							
Expected Stock Returns	2.31	0.97	1	2	2	3	3
Expected GDP Growth	2.42	0.97	1	2	2	3	4
Expected Bond Returns	2.85	0.99	2	2	3	3	4
Confidence ("Not at all confident", ..., "Very confident")							
Expected Stock Returns	3.06	0.85	2	3	3	4	4
Expected GDP Growth	3.00	0.84	2	3	3	3	4
Expected Bond Returns	2.61	0.84	2	2	3	3	4
Time of responding to survey (seconds)	517	433	238	307	414	579	834

Note: Table shows summary statistics across all 20,208 complete survey responses across the ten waves between February 2017 and August 2018. The possible answers for "difficulty" are 1 = "Not at all difficult", 2 = Not very difficult", 3 - "Somewhat difficult", 4 - "Very Difficult, 5 - "Extremely Difficult." The possible answers for confidence are 1 = "Not at all confident", 2 = "Not very confident", 3 = "Somewhat confident", 4 = "Very confident", and 5 = "Extremely Confident."

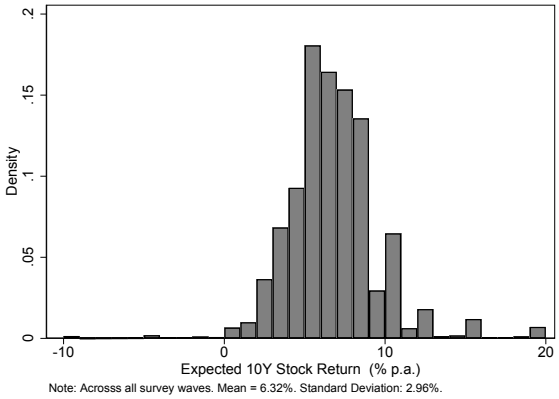
Figure III: Distribution of Responses to Expectation Survey

(A) 1Y Stock Returns



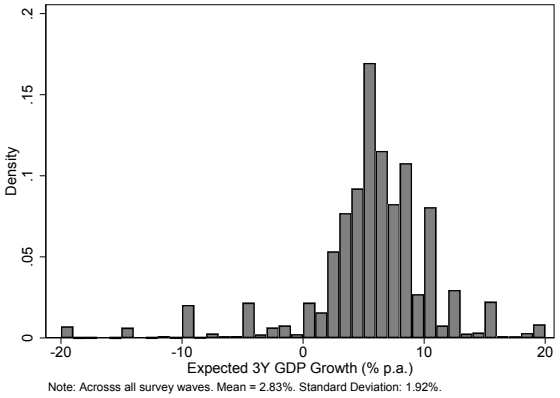
Note: Acrosss all survey waves. Mean = 5.23%. Standard Deviation: 5.32%.

(B) 10Y Stock Returns



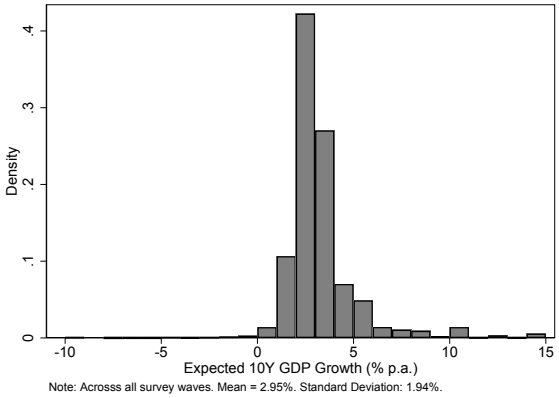
Note: Acrosss all survey waves. Mean = 6.32%. Standard Deviation: 2.96%.

(C) 3Y GDP Growth



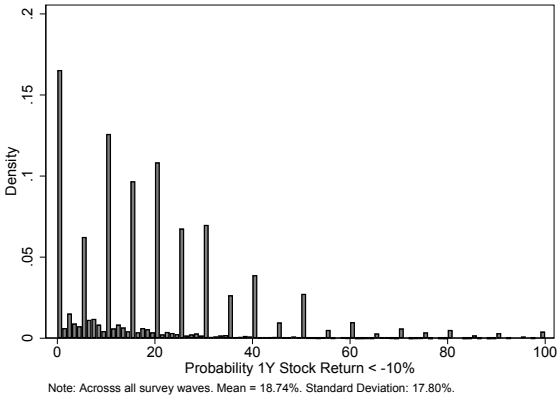
Note: Acrosss all survey waves. Mean = 2.83%. Standard Deviation: 1.92%.

(D) 10Y GDP Growth



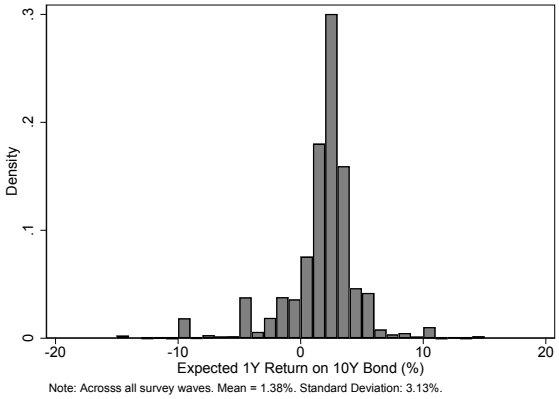
Note: Acrosss all survey waves. Mean = 2.95%. Standard Deviation: 1.94%.

(E) Probability Stock Disaster



Note: Acrosss all survey waves. Mean = 18.74%. Standard Deviation: 17.80%.

(F) 1Y Bond Returns



Note: Acrosss all survey waves. Mean = 1.38%. Standard Deviation: 3.13%.

Note: Figure shows histograms of the answers across complete survey responses across the ten waves between February 2017 and August 2018.

curve to remain upward-sloping in the future.

Most individuals report finding the survey questions relatively easy to understand, though the questions on bond returns were perceived to be more difficult than the questions on expected stock market returns or expected GDP growth. There also appears to be a relatively wide range of confidence that individuals have in their answers. For each of the three survey blocks, individuals at the 10th percentile of the distribution report being "not very confident" in their answers, while individuals at the 90th percentile reported being "very confident." In the following sections, we will explore the importance of individuals' confidence in their expectations on the relationship between beliefs and portfolio allocations.

IV Beliefs and Portfolios

In this section, we explore the relationship between respondents' beliefs and their portfolio allocations. Our primary focus is on understanding the asset allocation, that is, the share of each respondents' portfolio that is invested in stocks, fixed income instruments, and cash.[10] We first explore the role of expectations about the 1-year stock return before analyzing the relationship between other beliefs and portfolio allocations.

IV.A Expected 1-Year Stock Returns and Portfolios

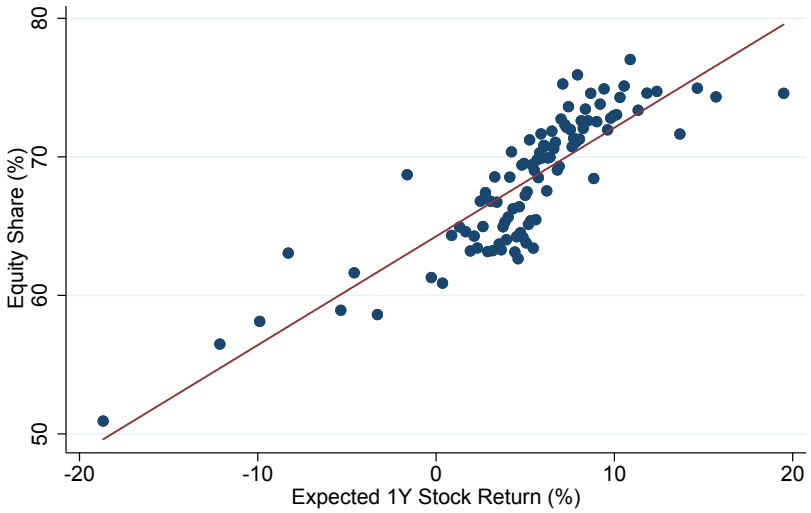
Figure IV shows a binned scatter plot of the relationship between an individual's beliefs about expected 1-year stock returns and the share of her portfolio held in equities. The underlying data is all responses to the survey, and the figure conditions on respondents' demographic characteristics such as age, gender, region, and wealth; it also includes survey wave fixed effects. There is a clear positive relationship: respondents who report higher expected stock returns also hold more equity in their portfolio. Quantitatively, moving the 1-year expected stock return from 0% to 10% increases the conditional expectation of the equity share from about 64% to about 72%.
To obtain our estimates of the associated sensitivity of portfolio shares to beliefs, we run the following regression:

$$EquityShare_{i,t} = \alpha + \beta ERetur\mathrm{n}1Y_{i,t} + \gamma X_{i,t} + \psi_t + \epsilon_{i,t} \tag{1}$$

The unit of observation is a survey response by individual i in wave t. The dependent variable is the equity share in the individual's Vanguard portfolio at time t. The variable of interest is β, which gives the increase in an individuals' equity share for each percentage point increase in the expected 1-year stock return. We continue to control for demographic characteristics such as age,

[10] The classification uses Vanguard's data on individual clients' portfolio compositions. Investments in mutual funds are apportioned depending on the portfolio composition of each fund (e.g., 60% equity and 40% fixed income). Cash includes cash-equivalent investments such as short-term paper and money-market funds. A category "other investments" groups alternative investments such as commodities, real estate, and some derivatives; the vast majority of investors hold no such investments.

Figure IV: Expected 1-Year Stock Returns and Equity Share



Note: Figure shows a conditional binned scatter plot of survey respondents' expected 1-year stock returns and the equity share in their portfolios, conditional on the respondents' age, gender, region, wealth, and the survey wave.

gender, wealth, and region, as well as survey-wave fixed effects. Column 1 of Table III shows that, consistent with Figure IV, the effect of an extra percentage point of expected return is associated with a 0.79 percentage point increase in respondents' equity share.[11] Combined with the wide heterogeneity in beliefs across individuals, variation in beliefs is thus associated with a significant variation in equity shares. A one-standard-deviation increase in expected one-year equity returns is associated with a $\frac{0.785 \times 5.32}{26.2} = 0.16$ standard deviation increase in equity shares.

While survey answers have a clear and statistically strong relationship with individuals' portfolios, the estimated sensitivity of equity shares to beliefs is low compared to predictions from benchmark theoretical models. For example, Merton (1969) provides a benchmark relation between expected returns and the share of the portfolio invested in equity, for a power-utility investor:

$$\omega_{E,i} = \frac{1}{\gamma} \frac{E_i[R] - R_f}{Var_i[R]}, \tag{2}$$

where $\omega_{E,i}$ is the fraction of an individual's portfolio invested in equity, γ is the individual's coefficient of relative risk aversion, $E_i[R]$ is the individual's expected stock return, R_f is the risk-free rate, and $Var_i[R]$ is the individual's subjective variance of equity returns. In the classic work of Merton (1969), individuals have rational expectations, while here we allow individuals to have arbitrarily formed expectations as long as the variance is non-zero and finite. We measure $\omega_{E,i}$ in the Vanguard data, and $E_i[R]$ with the survey answer to 1-year expected stock returns. For a back

[11] While the estimates of β are the primary object of interest, the coefficients on some of the control variables are also interesting. Males and females do not have significantly different equity shares. Equity shares are strongly declining in age, with individuals above 70 years of age having about 20 percentage points lower equity share than individuals below the age of 40. Equity shares also do not differ significantly across regions or across wealth quintiles.

Table III: Expected Returns and Portfolios

	Equity Share (%)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Expected 1Y Stock Return (%)	0.784***	0.952***	0.769***					
	(0.035)	(0.080)	(0.074)					
Expected 1Y Stock Return (%) x (Assets < $100k)				0.432***				
				(0.144)				
Expected 1Y Stock Return (%) x (Assets >= $100k)				1.191***				
				(0.096)				
Expected 1Y Stock Return (%) x Monthly Turnover < 0.5%					0.518***			
					(0.131)			
Expected 1Y Stock Return (%) x Monthly Turnover ∈ [0.5%,2%]					1.355***			
					(0.153)			
Expected 1Y Stock Return (%) x Monthly Turnover > 2%					1.642***			
					(0.151)			
Expected 1Y Stock Return (%) x Monthly Vanguard Visits ∈ (0,1)						0.605***		
						(0.130)		
Expected 1Y Stock Return (%) x Monthly Vanguard Visits ∈ (1,4)						0.861***		
						(0.144)		
Expected 1Y Stock Return (%) x Monthly Vanguard Visits ∈ (4,31)						1.477***		
						(0.142)		
Expected 1Y Stock Return (%) x Low Confidence							0.364	
							(0.208)	
Expected 1Y Stock Return (%) x Medium Confidence							1.033***	
							(0.091)	
Expected 1Y Stock Return (%) x High Confidence							1.565***	
							(0.337)	
Expected 1Y Stock Return (%) x Not Idealized								0.942***
								(0.081)
Expected 1Y Stock Return (%) x Idealized								4.302***
								(1.160)
Demographic Controls	Y	Y	Y	Y	Y	Y	Y	Y
Wave Fixed Effects	Y	Y	Y	Y	Y	Y	Y	Y
Specification	OLS	IV	IV, Interior	IV	IV	IV	IV	IV
R-Squared	0.107	0.105	0.138	0.104	0.114	0.104	0.106	0.106
N	19,478	19,301	16,343	19,301	17,904	19,301	19,301	19,301

Note: Table shows summary results from regression 1. The dependent variable is the equity share. "Low confidence" corresponds to individuals who report ed being "not at all confident" or "not very confident" in their answers about expected stock returns; "medium confidence" corresponds to individuals who report being "somewhat confident" or "very confident" about their answers; and "high confidence" corresponds to individuals who report being "extremely confident." "Idealized" respondents in column 8 are those whose behavior most closely corresponds to that of the assumptions in the frictionless model: they have average monthly portfolio turnover of at least 4%, they have at least four log-ins a month, and they report to be extremely confident in their beliefs.

of the envelope calculation, we assume that individuals have a common measure of the variance $Var_i[R] = Var[R]$.[12] Similarly, we assume a common and constant coefficient of relative risk aversion. In this setting, the β estimated in Table III corresponds to $\beta = \frac{1}{\gamma Var[R]}$. The historical standard deviation of stock market returns is around 16% a year. The simple model thus implies that a β of 0.785 requires a coefficient of relative risk aversion of $\gamma = 50$. This is considerably higher than most estimates in the empirical literature, which usually find values of γ between 3 and 10. To obtain a realistic coefficient of risk aversion, let us say around 4, we would need a β estimate in Table III of around 10. In other words, the estimated sensitivity is an order of magnitude too small to align with the simplest frictionless model.[13]

This relatively small response of equity shares to beliefs about stock returns is qualitatively and quantitatively consistent with evidence documented across a number of other studies that link equity market participation and equity shares to expected stock market returns (e.g., Vissing-Jorgensen, 2003; Dominitz and Manski, 2007; Kézdi and Willis, 2011; Amromin and Sharpe, 2013; Ameriks et al., 2016; Drerup, Enke and Von Gaudecker, 2017), though many of these studies had to rely on self-reported equity shares or a binary indicator for holding stocks rather than administrative data. In the following, we contribute to the literature in a number of dimensions. First, we directly explore the empirical importance of a number of possible explanations for the relatively low observed sensitivity. Specifically, we quantify (i) the role of attenuation bias due to measurement error in beliefs, (ii) the role of infrequent trading, and (iii) the role of respondents' confidence in their own beliefs. We also explore the margins of substitution that drive the baseline result, analyzing whether more optimistic individuals substitute towards equity from cash or from fixed income investments. Finally, we analyze the role of other moments (like tail events) of the belief distribution, of stock market expectations over different horizons, and of beliefs about bond returns and GDP growth.

Dealing with Measurement Error. A first possible explanation for the relatively low estimated sensitivity of equity shares to beliefs is that large measurement error in beliefs may induce attenuation bias in our estimates of β. To explore whether this can explain the finding, we exploit the fact that we have two "independent" estimates of the same explanatory variable. The first is the expected return as reported directly by respondents in the survey. The second is the implied mean of the distribution over possible 1-year returns reported by individuals.[14] Both of these elicit the

[12] We later relax this assumption and measure the variance of 1-year expected stock returns that is implied by the distribution question. We find the simple calculation with a common variance across individuals parametrized to the historical variance to be appealing for several reasons: (i) in many models, it is easy to learn the variance of returns but hard to learn the mean; (ii) equation 2 is particularly sensitive to measurement errors in the denominator; and (iii) model misspecification is likely and other moments (like tail event probabilities) are more important than the variance.

[13] Note that in this model, γ drives both the sensitivity of the portfolio share to changes in expected returns, and the unconditional level of the equity share. In particular, for $\gamma = 4$, average risk premia of 6% and standard deviation of 16%, we obtain an average share of 58%, close to the average share we see in the data (for example, 68% for respondents as shown in table I). When instead $\gamma = 50$, the equity share drops to 4.7%. So high risk aversion can explain the low sensitivity of portfolios to beliefs, but at the cost of missing the average level of the portfolio share.

[14] To construct the implied mean from the distribution, we take the mid-points of the closed buckets. For the open-ended buckets, we take -40% for the lowest bucket ($\leq -30\%$) and 50% for the highest bucket ($\geq 40\%$). Our findings

same object using very different styles of question, and are therefore likely to have independent measurement errors. This setting therefore allows us to minimize attenuation bias by instrumenting for the reported expected return with the implied expected return from the distribution (the correlation across the two measures is 0.43). Consistent with this procedure minimizing attenuation bias, the estimates in column 2 show that the sensitivity increases by about 20%. This suggests that classical measurement error, and the associated attenuation bias, while present, do not contribute substantially to the relatively low estimated sensitivity. One caveat here is that the instrumental variables procedure does not address any potential non-classical measurement error that is correlated across the different elicitation methods; indeed, any such measurement error is observationally equivalent to true differences in beliefs.

Corner Solutions. The preceding analysis uses all portfolio data, including those individual portfolios that own either 0% or 100% in equity. Since Vanguard customers generally cannot short stocks or take substantial leverage within the account, the boundaries of 0% and 100% represent natural constraint on the investment set. One concern, therefore, could be that we measure a low sensitivity because individuals who are extremely optimistic (pessimistic) cannot express their view by taking positions in excess of 100% equity (taking negative equity positions). Column 3 of Table III shows that this is not the case; to the contrary, excluding the extreme portfolios lowers the measured elasticity which goes from 0.952 in the unrestricted Column 2 to 0.769 in Column 3. This occurs because, on average, very optimistic and very pessimistic agents do in fact hold 100% and 0% of their portfolios in equity, respectively. Contrary to the conjecture, and despite the constraints, these individuals have portfolios that are more in line with theirs beliefs (i.e., more sensitive to the beliefs) than those investors with interior portfolio allocations. For the remainder of the paper, we include the extreme portfolios from the analysis, but treat them separately whenever it is informative to do so as (e.g., in the analysis of trading).

Heterogeneity by Investor Wealth. As we discuss below, the relatively low sensitivity of portfolios to beliefs provides a potential challenge for a number of standard asset pricing models. However, in many of these models it is the wealth-weighted beliefs that drive asset prices, rather than beliefs that are equally weighted across all investors. We therefore next explore whether the sensitivity of portfolios to beliefs is different for wealthier individuals. Column 4 of Table III shows that respondents with more than $100k in assets have a sensitivity that is about 2.5 times as large as that of individuals with a lower wealth; in unreported results, we show that the sensitivity does not increase substantially beyond $\beta = 1.2$ as Vanguard assets grow further.[15] This means

are robust to picking other values for the open-ended buckets.

[15]We also explored heterogeneity of the sensitivity of portfolio shares with respect to stock market beliefs along a number of other dimensions. We find that the equity shares of the portfolios of older individuals are particularly sensitive to their 1-year expected returns. Similarly, the portfolios of male respondents are substantially more sensitive to beliefs than the portfolios of female respondents. We did not find significant heterogeneity in the responsiveness of portfolios to beliefs based on the region where individuals are located, or based on how long individuals take to answer the surveys.

that wealth-weighted sensitivity may be higher than the equally-weighted sensitivity, but not by nearly enough to generate the quantity movements required by frictionless asset pricing models. We thus turn to exploring two possible frictions that may help explain the muted sensitivity: infrequent trading, and lack of confidence in own beliefs.

Infrequent Trading. A first plausible explanation for the low estimated sensitivity of equity shares to beliefs is that investors only adjust their portfolios infrequently. To the extent that investors change their beliefs over time and report their current belief in the survey, the contemporaneous portfolios may thus not be fully reflective of current beliefs. Through this channel, infrequent trading could thus help explain the attenuated relationship between beliefs and portfolios relative to the predictions from frictionless models. Researchers have focused on several complementary explanations for infrequent portfolio adjustments. The first is the cost of monitoring portfolio allocations, which can cause investors to only infrequently pay attention to their portfolios. The second explanation is that even after paying attention to their portfolios, a number of additional costs may prevent investors from trading; these can include physical transaction costs coming from brokerage commissions and capital gains taxes, and other mental transaction costs coming from the need to determine optimal behavior based on beliefs and current portfolios.[16]

We next explore whether infrequent trading can help us understand the low sensitivity of portfolio allocations to beliefs. We first split individual respondents into three groups depending on their trading behavior during the sample period. Specifically, we classify individuals by the average monthly turnover in their portfolios, but our results are robust to other definitions of "infrequent trading", such as the average monthly number of unique trades. Column 5 of Table III shows that individuals with a monthly portfolio turnover of at least 2% have a sensitivity of equity shares to beliefs that is about three times as large as that of individuals with a monthly portfolio turnover of less than 0.5%.[17] However, even for those investors who have a sizable portfolio turnover, the sensitivity remains at 1.64, well below the benchmark value of 10 from the frictionless model discussed above. Nevertheless, these findings suggest that trading frequency and intensity are important determinants of how strongly (and quickly) beliefs are reflected in portfolio holdings, and that the reasons that can explain the infrequent trading activity need to be understood.

We next explore the role that a specific reason for infrequent trading, namely investor attention, plays in explaining this relationship. We measure investor attention by the frequency with which investors log into their Vanguard accounts.[18] Column 6 of Table III shows that individu-

[16] A number of researchers have studied environments in which investors only infrequently pay attention to their portfolios and face costs in changing their portfolios (e.g., Duffie, Sun et al., 1990; Gabaix and Laibson, 2001; Madrian and Shea, 2001; Agnew, Balduzzi and Sunden, 2003; Hirshleifer and Teoh, 2003; Peng and Xiong, 2006; Abel, Eberly and Panageas, 2007; Alvarez, Guiso and Lippi, 2012).

[17] In other words, we explore the role of trading as a static individual characteristic by measuring how often and how much each respondent trades on average. Section V further explores the relationship between trading and beliefs by exploiting the time series for each individual to measure both the extensive margin (whether beliefs predict trading) and the intensive margin (pass-through of beliefs to portfolios conditional on trading).

[18] Logging into one's Vanguard account is correlated with, but different from actual trading activity. Obviously one

als who log into their Vanguard accounts more frequently also have a higher sensitivity of equity share to beliefs.[19] Indeed, individuals who log in more than four times per month have an sensitivity that is about 2.5 times as large as that of individuals who log in less than once per month. These results suggest that investor attention to their portfolios is an important driver of infrequent trading, and therefore of the attenuated relationship between beliefs and portfolio allocations.[20]

Confidence. A further mechanism that is potentially important in understanding how differences in beliefs translate into portfolio holdings is the confidence that individuals have in their own beliefs. Indeed, a large literature suggests that individuals that are more confident in their own beliefs are more likely to trade on them (e.g., De Long et al., 1990*a*; Kyle and Wang, 1997; Daniel, Hirshleifer and Subrahmanyam, 1998; Gervais and Odean, 2001). To explore the role that investor confidence has on the extent to which their beliefs are reflected in their portfolios, we exploit the fact that the survey directly elicits how confident individuals are about their answers. The confidence of individuals in their beliefs is essentially uncorrelated with trading volume or how often individuals log into their portfolios (the correlation with average monthly portfolio turnover is 0.04, and the correlation with the number of visits per month is 0.06). Indeed, Figure V confirms that while individuals that are more confident log in or trade slightly more often, most of the variation in trading and attention is within individuals with the same reported confidence. This means that any variation in sensitivity by confidence is picking up a conceptually different object than variation in sensitivity by trading frequency or attention.

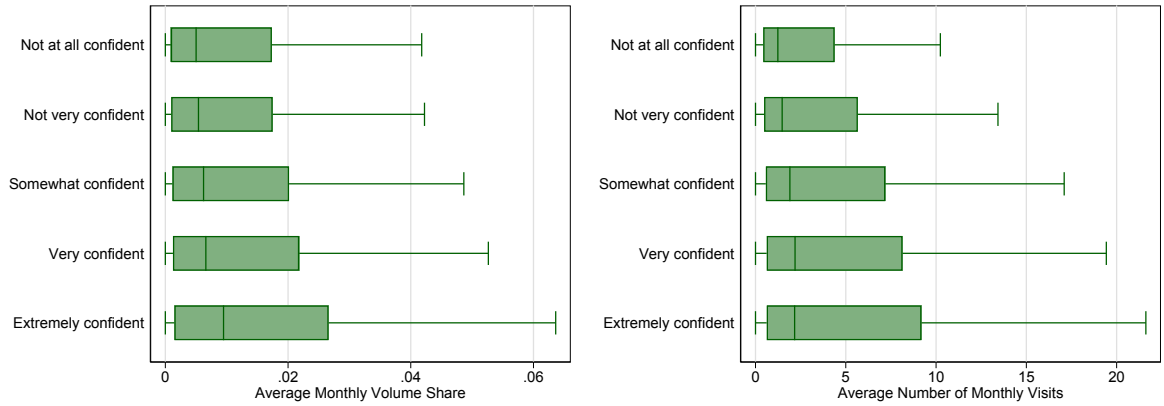
In Column 7 of Table III, we explore how individuals' sensitivity of equity shares to their beliefs varies with the confidence in their beliefs about the stock market. We find that individuals who are more confident in their beliefs have a much higher sensitivity of equity share to beliefs. In particular, individuals who report being "extremely confident" in their stock market beliefs have a 5 times higher sensitivity of portfolio shares to beliefs than individuals who report being "not at all confident" or "not very confident." However, we find that even the highest measured sensitivity for the most confident group is 1.56, substantially below the benchmark of 10 described above. Figure VI shows how equity shares vary with beliefs for individuals with different confidence levels using (unconditional) local polynomial smooths. Interestingly, the largest differences are among relatively pessimistic individuals where large reductions in equity share are primarily observed among the most confident individuals. Our results are qualitatively consistent with Drerup, Enke and Von Gaudecker (2017) who find that individuals who are more confident in their beliefs about the stock market are more like to participate (i.e. report that they hold at least some

can log in and not trade, but also one can trade on the phone without logging into the online account.

[19] In this specification, we treat "paying attention" as a time-invariant individual characteristic, and split individuals by the average attention they paid during our sample. We find similar heterogeneity when we divide individuals by the number of minutes that they spend on the Vanguard platform in a representative month.

[20] Note that while the frequency of trading and log-ins strongly affect the relationship between beliefs and portfolios, they are not correlated with the level of beliefs themselves. Indeed, the correlation between the average monthly volume share and the expected 1-year stock market return is 0.02, and the correlation between the average monthly number of Vanguard log-ins and the expected 1-year stock market return is 0.05.

Figure V: Trading Frequency and Log-Ins, By Confidence



Note: Figure shows the distribution of average monthly volume share (left panel) and average monthly Vanguard visits (right panel) separately by individual respondents' answer to how confident they were in their answers about expected stock returns. The box plots show the 5th, 25th, 50th, 75th and 95th percentiles.

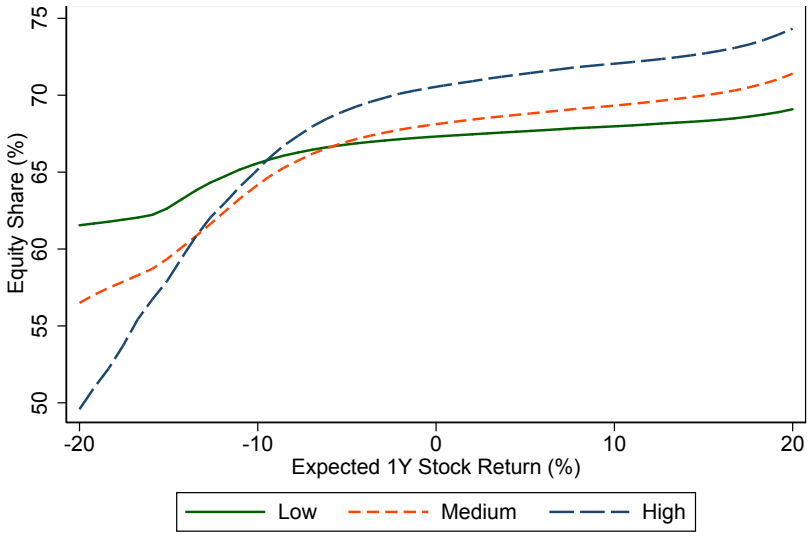
stocks) in that market.

Mapping confidence as reported in the survey to economic theory models is not immediate. One possibility is that confidence is about the degree of certainty that individuals have about the entire distribution of outcomes (GDP growth and stock returns). Individuals who are less confident think that there is a higher probability that the outcomes will be drawn from a distribution different from the one that they reported in the survey.[21]

Unobservable Heterogeneity: Labor Income, Other Investments, Risk Aversion. There are a number of other possible explanations for the relatively low sensitivity of equity shares to stock market beliefs. One possible set of explanations falls under the category of "optimists take risks outside of Vanguard portfolios". For example, more optimistic respondents might have other accounts, at a different firm, and predominantly take risks in those accounts. Note that to explain our results it is not enough that all respondents choose a safer portfolios at Vanguard than in their other accounts; it would have to be that pessimists are disproportionally likely to do so. Another example would be that more optimistic respondents choose job occupations that are riskier; since their labor income is risky, then they tilt their portfolio investment toward safer assets. Another explanation is that more optimistic respondents are substantially more risk averse. This seems unlikely since the implied variation in risk aversion across individuals would have to be extraordinarily large. In addition, Ameriks et al. (2016) observe one cross-section of both expected returns and risk-aversion (elicited via lottery type-questions) and conclude that risk-aversion heterogeneity is not sufficiently strong to explain the low estimated sensitivity of portfolios to beliefs.

[21] One possibility that we ruled out is that confidence simply reflects individuals' uncertainty about the expected outcome. While confidence is inversely related to the standard deviation of outcomes implied by the distribution questions, the relation is far from perfect, and the effect of confidence on actions such as portfolio risk-taking goes well beyond the effect induced by the standard deviation.

Figure VI: Expected Returns and Equity Share, By Confidence



Note: Figure shows local polynomial smooths of equity shares and expected 1-year stock returns and the equity share separately for individuals with different levels of confidence in their own beliefs. "Low confidence" corresponds to individuals who report ed being "not at all confident" or "not very confident" in their answers about expected stock returns; "medium confidence" corresponds to individuals who report being "somewhat confident" or "very confident" about their answers; and "high confidence" corresponds to individuals who report being "extremely confident."

Summary. The evidence above has documented that investors' portfolios systematically vary with their beliefs. However, quantitatively, the average relationship is smaller than what would be predicted by simple frictionless models. In the above analysis, we have identified a number of deviations from this frictionless model that can help explain the overall low sensitivity. Individuals who trade more extensively, who pay more attention to their portfolios, and who are more confident in their beliefs all have a higher sensitivity of equity shares to stock market expectations. However, the individuals who are most similar to the frictionless benchmark on each of these three dimensions still do not have a sensitivity that approximates that of the frictionless model. In column 8 of Table III, we explore the behavior of those respondents whose behavior comes closest to that described by the frictionless model on all three dimensions: individuals who are extremely confident in their beliefs, who pay substantial attention, and whose trading volume is significant. For respondents in that group, we estimate a β of 4.3, though the standard error around this estimate is quite large.[22] This estimate is much closer to the benchmark number of 10 from the simple frictionless model described above; the point estimate itself would imply a γ of 9.1, which is at the upper end of values considered in the literature. The remaining gap to the more commonly used values of 3 to 5 can easily be explained by a combination of residual measurement error in beliefs and model misspecification; indeed, in Section IV.B we show that other beliefs, such as the 10-year stock market beliefs and beliefs about disaster probabilities, also affect the equity share.

[22]Since confidence, attention, and trading are only weakly correlated, there are only 71 individuals who have a monthly portfolio turnover larger than 4%, who log in an average of four times per month or more, and who are "extremely confident" in their answer.

Overall, our findings suggest that investor attention, adjustment costs, and confidence are important features that mediate the transmission from beliefs into quantities, and should therefore play a more prominent role in the description of asset pricing models going forward. We sum up these findings in the following facts on the relationship between investor beliefs and portfolios.

Fact 1. *The average sensitivity of an investor's portfolio share in equity to that investor's expected stock market returns is lower than predicted by frictionless asset pricing models.*

Fact 2. *The sensitivity of portfolios to beliefs is increasing in wealth. It is also increasing in investor trading frequency, investor attention, and investor confidence in their own beliefs.*

IV.B Other Beliefs and Portfolios

In the previous analyses, we explored the role that an individual's expectations of one-year stock returns have on her portfolio allocation. We next explore how beliefs about other moments of stock returns as well as beliefs about bond returns and GDP growth affect these portfolio allocations.[23] In column 1 of Table IV, we additionally control for the standard deviation of expectations of 1-year stock returns. A higher standard deviation of expected returns has a statistically significant negative effect on the equity share. However, the estimated sensitivity to expected returns decreases when controlling for the variance. Indeed, the sensitivity decreases from 0.952 in Column 2 of Table III to 0.889 in Column 1 of Table IV. This occurs because expected returns and variance are negatively related at the individual respondents level.

In column 2, we also control for the probability that individuals assign to a stock market decline of more than 30%. Asset pricing theory assigns a prominent role to the probability of these rare disasters in shaping both portfolio formation and, in general equilibrium, asset returns (Rietz, 1988; Barro, 2006; Gabaix, 2012). It is therefore ex-ante likely to be a major source of model misspecification for equation 2 and the associated predictions for the sensitivity of portfolios to stock market beliefs. Indeed, we find that a higher probability of a rare disaster is associated with declines in the equity portfolio share even when controlling for the variance of returns. In fact, the effects of variance on portfolio formation is cut in half and becomes statistically insignificant once we control for the probability of disaster. A one-standard deviation increase in the expected stock market disaster probability is associated with about a one percentage point lower equity share. Interestingly, controlling for the perceived probability of rare disasters lowers the sensitivity of portfolios to 1-year expected returns. This occurs because at the individual level expected returns and the probability of a rare disaster are negatively related, a point that Section VIII returns to in more detail.

The sensitivity of the portfolio choice to the probability of disaster is in itself an interesting parameter. To provide a model-based benchmark to interpret this coefficient, we solve for the

[23] We are not able to instrument for these other beliefs; this means that our estimates on the effects of these other beliefs may be affected by attenuation bias, though our analysis above suggests that such a bias may be relatively small.

static portfolio choice under power utility with a benchmark risk aversion of 4, assuming each investor faces the 5-point distribution they reveal in the survey. Using this model, we obtain a counterfactual, model-predicted equity share, and then run a regression of the portfolio share on expected returns, the implied standard deviation, and the probability of the lowest bin. Not surprisingly, we find that model-based portfolio shares are about 10 times more sensitive to changes in the probability of disaster than what we find in the data (with a coefficient of -2.3).

In column 3, we include an individual's beliefs about the average annualized stock return over the coming ten years in addition to the beliefs about the expected stock returns over the coming year. Short horizon and long horizon stock-market return expectations are positively correlated (see Section VII), but long-run expectations matter for portfolio formation even after controlling for short-run expectations. Interestingly, the magnitude of the effects are similar for long and short horizon expectations. These results suggest that individuals choose their portfolio for the long-run, particularly since they do not adjust it frequently, and do not behave myopically by only focusing on their short-run expectations.[24]

In column 4, we also include controls for a respondent's beliefs about the 1-year return of a ten-year risk-free bond. Holding fixed beliefs about equity returns, increased optimism on bond returns is associated with a lower equity share. Finally, column 5 also includes controls for GDP growth expectations, but these do not have an effect on portfolio shares over and above the stock market and bond market expectations.

IV.C Substitution Patterns: Equity, Fixed Income, and Cash

The previous discussion explored the relationship between various beliefs (about stock returns, bond returns, and GDP growth) on the equity share in investors' portfolios. In columns 6 and 7 of Table III, we instead use the fixed income share and the cash share as the dependent variables, allowing us to explore the substitution between stocks, bonds, and cash. The vast majority of increase in equity shares of individuals who expect higher stock market returns comes from individuals substituting away from cash rather than individuals substituting away from fixed income securities. This is despite the fact that the average fixed income share is substantially larger than the average cash share. Similarly, we find that increases in the expected bond returns are associated with increases in the fixed income share. Quantitatively, however, the sensitivity of the bond portfolio shares to bond expected returns in even lower than the corresponding case for equities.[25] We conclude that fixed income offers a similar, even more extreme, picture to equities: portfolios co-move with beliefs but to a much lesser extent than benchmark portfolios theories imply. About

[24] A "myopic" behavior does not necessarily have to be behavioral; for example, an agent with logarithmic utility only cares about short-horizon returns in forming her portfolio.

[25] A back of the envelope calculation is again very illustrative. We can apply the Merton formula in equation 2 by replacing equities with bonds and using the historical standard deviation of long-term Treasury-bond returns around 5%. The estimate of $\beta = 0.465$ from column 6 of Table III then implies a coefficient of relative risk aversion of 860. Recall that, in this setting, $\beta = \frac{1}{\gamma Var[R]}$; hence, an asset with low volatility should have a high sensitivity of expected returns into portfolios since a given move in the expected return corresponds to a larger movement in the asset's Sharpe ratio.

Table IV: Beliefs and Portfolios: Long-Run Returns, Variance, Tail Risk, and Bond Returns

	Equity Share (%)					Fixed Income Share (%)	Cash Share (%)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Expected 1Y Stock Return (%)	0.889***	0.719***	0.554***	0.640***	0.671***	-0.106	-0.731***
	(0.082)	(0.099)	(0.123)	(0.130)	(0.142)	(0.120)	(0.111)
Standard Deviation 1Y Stock Return (%)	-0.124***	-0.056	-0.081	-0.068	-0.052	0.025	0.003
	(0.036)	(0.045)	(0.045)	(0.045)	(0.046)	(0.039)	(0.036)
Probability 1Y Stock Return < 30% (%)		-0.106**	-0.080*	-0.062	-0.063	0.047	-0.003
		(0.041)	(0.040)	(0.041)	(0.042)	(0.035)	(0.033)
Expected 10Y Stock Return (%)			0.735***	0.710***	0.758***	-0.666***	-0.007
			(0.106)	(0.108)	(0.115)	(0.097)	(0.090)
Expected 1Y Return of 10Y bond (%)				-0.190**	-0.173*	0.468***	-0.279***
				(0.069)	(0.071)	(0.060)	(0.055)
Expected 3Y GDP Growth (% p.a.)					-0.09	-0.14	0.220
					(0.179)	(0.151)	(0.140)
Expected 10Y GDP Growth (% p.a.)					-0.007	0.053	-0.125
					(0.153)	(0.129)	(0.120)
Demographic Controls	Y	Y	Y		Y	Y	
Wave Fixed Effects	Y	Y	Y		Y	Y	
Specification	IV	IV	IV		IV	IV	
R-Squared	0.108	0.109	0.114	0.113	0.115	0.084	0.028
N	16,918	16,918	16,552	16,241	15,816	15,816	15,816

Note: Table shows summary results from regression 1. The dependent variable in columns 1 to 5 is the equity share; in column 6 it is the fixed income share, and in column 7 it is the cash share.

two thirds of this increase comes in form of reductions in the cash share, while about one third comes from reductions in the equity share.

IV.D Beliefs and Portfolios: Implication for Current Behavioral Models

The above section highlights that surveys of expected returns provide a fruitful testing ground for models of asset pricing. In particular, our survey is large scale, designed to elicit beliefs about moments that are particularly informative for theory, matched to high quality administrative data on portfolios and trading, and focused on a panel of individuals with substantial wealth invested in the markets. At a fundamental level, our results challenge the notion, most notably articulated by Cochrane (2011, 2017), that survey data and the associated critiques of rational expectations model of asset pricing can simply be discarded because respondents do not understand the questions or because their answers reveal beliefs that do not actually correspond to their investments. While all belief surveys are imperfect and subject to valid concerns about multiple dimensions of measurement error, we clearly document that our survey answers are relevant to understanding individuals' investment behavior.

These results can also guide the construction and calibration of behavioral asset pricing mod-

els. While some of the recent models in this literature are explicitly built to match some moments obtained from surveys (see, for example, the recent work by Barberis et al., 2015), our results provide new and direct empirical evidence on the link between beliefs and portfolios that plays an important role in these models. More specifically, the existing models often build on a modification of Mertonian portfolio demand in CARA-normal set ups, but rather than imposing rational expectations, they impose expectations that are more consistent with survey evidence. These models' predictions for asset prices rest on two modeling blocks: (i) that beliefs are volatile and change over time, and (ii) that individual portfolios react strongly to changes in these beliefs.[26] Our evidence shows that, for the average investor, the second channel is much weaker than in common calibrations of these models. Indeed, most behavioral models imply the same sensitivity of portfolios to beliefs as the benchmark rational-expectations models. Our results show that the sensitivity assumed in these models may only describe the behavior of a select group of investors who pay attention to their portfolios, who trade regularly, and who are confident in their beliefs. For the majority of investors, infrequent trading and lack of confidence in beliefs reduces the passthrough from beliefs to portfolios. As a result, the frictionless behavioral models, in their current form, are likely to overstate the power of expectation changes in explaining asset price movements. Our research opens a new challenge to improve the models: a successful model has to match both expectations and portfolio dynamics jointly with asset prices, while the current literature, which is already a major advancement, focuses primarily on matching expectations and asset prices.

It is beyond the scope of this paper to fully specify such a model, but we find it useful to describe some possible avenues arising from our work. First, it appears important to better understand and model investor attention (and other determinants of trading frequency) as well as investors' confidence in their own beliefs. Since these variables are central to mediating the relationship between beliefs and portfolios, we expect that future belief surveys would benefit from jointly eliciting beliefs, confidence, attention, and portfolios. One further possibility for adjustments to the basic models is to explore the relative wealth and demand sensitivity of different classes of agents. For example, behavioral models such as those reviewed by Barberis (2018), often feature two types of investors (e.g., rational arbitrageurs and behavioral investors), each modeled with its own demand for stocks similar to equation 2. Asset prices are then determined by the dynamics of expectations of the behavioral agents, modulated by the relative wealth shares of the two agents and their relative demand sensitivities. One possible avenue therefore is to have the behavioral group be substantially wealthier (in the aggregate) than the arbitrageurs. This would induce the changes in demand by the behavioral agents, even with a relatively insensitive demand, to be larger for the arbitrageurs to absorb and therefore generate bigger asset price movements. Alternatively, one could make the demand of the behavioral agents more sensitive

[26] The first block is familiar from a long literature in finance going back to Shiller (1981) that has attempted to explain the excess volatility of asset prices compared to fundamentals via the excess volatility of agents' expectations (of either fundamentals or returns) compared to rational expectations. The second block is more nuanced: for movements in expectations to move asset prices in general equilibrium, these same movements of expectations have to move asset demand in partial equilibrium. The intuition is that if expectations move asset demand, fixed asset supply implies that prices have to adjust for markets to clear.

than that of the arbitrageurs. Belief surveys like ours could then be used to explore the correlation between deviations from rationality and the sensitivity of beliefs and portfolio allocations.

V Trading and the Pass-Through of Beliefs To Portfolios

In the previous section, we provided evidence that infrequent trading is at least partially responsible for the low measured sensitivity of portfolios to beliefs in a cross-section of survey respondents. Indeed, Section IV showed that respondents who trade more frequently tend to have portfolios that are more strongly reflective of their beliefs. In this section we inspect this mechanism further, by looking at investors' trading behavior over time. First, we explore the *extensive* margin of trading: can we predict the incidence of trading with changes in beliefs or other observables? Second, we inspect the *intensive* margin of trading: do changes in beliefs affect the direction and magnitude of trading, conditional on a trade taking place? Third, we inspect the extent to which the allocation of new funds that come into an account in the form of cash reflects investors' beliefs.

Before proceeding with the analysis, we provide a brief overview of how we measure trading; we refer the reader to the Online Appendix for full details. We observe administrative data on all transactions for each client with a retail account. These include, among other things, money moving in and out of the Vanguard accounts, purchases or sales of securities (e.g., bonds, equities, ETFs), and purchases, sales or exchanges of shares in mutual funds. The data include information on the date of the transaction, the type of transaction, and the quantity and price involved. Since we observe beliefs only when an investor answers the survey, we aggregate all trades that occur between two survey responses. We can then relate trading behavior to the beliefs at the time of each survey as well as the change in beliefs between surveys. We also aggregate trades within each asset class: equity, fixed income, cash and cash-equivalents, other investments, and unknown. For each window between two surveys we compute the changes in portfolio shares across these asset classes that are induced by active trading that occurs during the window.[27] For example, if a respondent with 100 dollars in her account holds 60% of her portfolio in equity at the beginning of the window, say January 15th 2018, and on January 20th she buys 10 more dollars of equity using cash already available in the account, then we predict *based on trading* the end of window equity share to be 70%. This is different from the actual change in portfolio shares during each window, because changes in market values also affect the final portfolio shares. The advantage of using trading data is precisely that we can separately identify the component of the change in portfolio shares that can be attributed to active trading decisions, as opposed as the component due to fluctuations in market values.

We begin our analysis by exploring how trade depends on the investors' beliefs. For every investor, we identify the windows of time between each pair of consecutive survey answers. These windows differ across investors and within the same investor over time since each investor might

[27] Trades that occur within an asset class, for example selling one stock to buy another, do not contribute to changes in asset class shares, but we still count them as trades in measures that are purely based on number or volume of trades.

Table V: Trading Analysis

	% Equity Trade		Trading Probability		Buy Probability		% Equity Trade, Conditional	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Change in Expected 1Y Stock Return (%)	0.101***	0.180***			1.110***	1.007*	0.367***	0.484***
	0.021	0.050			0.313	0.397	0.087	0.127
Expected 1Y Stock Return (%)	0.092***	0.131*	0.283*	0.032	1.290***	1.301**	0.411***	0.415**
	0.020	0.056	0.127	0.260	0.316	0.433	0.087	0.139
Equity Share (%)	-0.029***	-0.065***	-0.055*	-0.002	-0.299***	-0.415***	-0.184***	-0.239***
	0.003	0.009	0.028	0.058	0.064	0.089	0.017	0.027
Absolute Change in Expected 1Y Stock Re			0.256	0.187				
			0.172	0.311				
Constant	1.341***	3.134***	26.856***	44.156***	60.095***	66.756***	8.869***	12.118***
	0.245	0.699	2.051	4.426	4.446	6.331	1.192	1.935
# Waves between responses	2 or less	>2	2 or less	>2	2 or less	>2	2 or less	>2
R-Squared	0.016	0.032	0.002	0.000	0.029	0.041	0.095	0.115
N	5,929	1,809	4,997	1,510	1,273	677	1,298	700

Note: Table shows the results from regression 3. The unit of observation is a window between two consecutively answered surveys. The table has 5 panels, each of two columns. In each panel, the first column uses only trading between surveys that are at most 2 waves apart, the second column uses only trading between surveys that are more than 2 waves apart. Columns 1 and 2 regress the change in the equity share due to trading between waves on the expected 1-year stock return as of the beginning of the window, the change in this expectation during the window, and the equity share at the beginning of the window. Columns 3 and 4 regress an indicator of trading on the expected 1-month return as of the beginning of the window, the absolute change in this expectation during the window, and the equity share at the beginning of the window. Columns 5 and 6 regress an indicator of whether the individual increased her portfolio share in equity during the window on the expected 1-month return as of the beginning of the window, the change in this expectation during the window, and the equity share at the beginning of the window. Columns 7 and 8 adjust the analysis of 1 and 2, but condition on the individual trading during the window.

answer one survey, then do not answer a number of the following waves of the survey, and then answer again, thus creating irregular time intervals between any two completed surveys. For each window, we compute the total trading activity during that window, the beliefs at the beginning of the window and their change during the window. We also track the investors' beginning-of-window portfolio allocation. For example, if an investor has answered waves 1, 2 and 5 of the survey, we would identify two windows: the 2-months period between wave 1 and wave 2, and the 6-months period between waves 2 and 5. Each window would appear as a separate observation in the analysis below. Finally, note that the length of the window potentially matters for the results, since individuals are mechanically more likely to trade over longer windows. Therefore, for each of the results in this analysis, we report separately short windows (2 waves or less) and long windows (3 or more waves long).

Columns 1 and 2 of Table V report the results of a regression of the change in equity portfolio share due to trading for individual i in window w, given as $\Delta share_{i,w}^{trading}$, on the expected 1-year return as of the beginning of the window, the change in this expectation during the window, and the equity share at the beginning of the window:

$$\Delta share_{i,w}^{trading} = \alpha + \beta_{beg} ER_{i,w}^{beg} + \beta_{change} \Delta ER_{i,w} + \beta_{lev} share_{i,w} + e_{i,w}. \tag{3}$$

Column 1 of Table V reports the results for short windows, and column 2 for long windows. The table shows overall low sensitivity of trading to beliefs and their changes. For example, for short windows, an increase of 1% in expected returns at the beginning of the window predicts a 0.09% increase in the equity share due to trading over the following period; a 1% change in belief over the window predicts a 0.1% change in the equity share. While these sensitivities are statistically significant, they are quantitatively small, and significantly smaller than what we obtained in Section IV from the cross-sectional analysis. Since these regressions are looking at the changes over time in portfolio choice, rather than levels as in Section IV, they make use of a different source of variation; while one would expect these two approaches to produce similar results in a frictionless world, this is not necessarily the case if trading frictions are present. Finally, the table also shows that investors with high equity shares at the beginning of the window tend to subsequently trade to reduce their equity exposures, potentially a sign of rebalancing of their positions.

The low sensitivities in this table could reflect two different mechanisms. First, they could simply reflect the fact that individuals trade infrequently, so that the average sensitivity to beliefs appears low (extensive margin); they could also reflect a low sensitivity of trading to beliefs even conditional on investors trading (intensive margin). Our trading data allows us to disentangle the two explanations. We analyze them below.

The Extensive Margin of Trading. A voluminous academic literature has aimed to explain volume in financial markets via a mix of changes in beliefs and overconfidence (e.g., Harrison and Kreps, 1978; Hong and Stein, 1999; Scheinkman and Xiong, 2003; Hong and Stein, 2007). Similarly, the household finance literature has studied models of the optimal frequency and size of trading in the presence of limited information and transaction costs (e.g., Duffie, Sun et al., 1990; Gabaix and Laibson, 2001; Abel, Eberly and Panageas, 2007; Alvarez, Guiso and Lippi, 2012). A natural question in our setup, therefore, is whether changes in beliefs are associated with trading. Columns 3 and 4 of Table V report estimates of a regression of an indicator of whether the investor has traded at all during the window on the beginning-of-window beliefs, the beginning-of- window equity share, and the absolute change in beliefs over the window. As before, the two columns correspond to separate regressions split by the length of the window. These regressions condition on the initial portfolio share being strictly inside the interval (0,100), as we would naturally expect to see little predictable trading as a function of beliefs for investors that are at these boundaries.

The table reports that beliefs, their absolute changes, and the beginning-of-window portfolio shares do seem to marginally help predict whether an investor trades or not, at least over short windows. However, the predictive ability is statistically small: the R^2 of this regression is 0.002 for short windows and 0.0002 for longer windows. Effectively, these variables help only minimally in predicting whether trading occurs in our sample. We find similar results when we also control for gender, region, age, and wealth.

On the one hand, this difficulty with predicting who trades over a given time horizon might not be surprising, since trading happens infrequently and is often driven by liquidity needs or

life events (like job status) that we do not observe. On the other hand, these results support models with a probability of trading that does not vary much with beliefs. Such models range from imposing a flat hazard function exogeneously (i.e., agents are selected at random to trade with a constant probability per period), to deriving it as the optimal equilibrium choice.

The Intensive Margin of Trading. We next condition on trading actually occurring and ask whether the direction and the magnitude of the trade can be explained by belief changes.[28] We start by predicting the direction of trade for those windows in which investors trade. Columns 5 and 6 of Table V regress an indicator of whether an individual has increased her equity share during a window on the beliefs at the beginning of the window, the belief change during the window, and the equity share at the beginning of the window. This regression conditions only on having adjusted the equity share, in any direction, during the window. We classify a window as having had trades if the equity share changed by at least 1% (we therefore ignore very small changes in the portfolios), and we similarly define whether an investor has bought equities if the equity share increased by at least 1%.

The results show that beliefs predict the direction of trade conditional on a trade actually occurring: an investor who is more optimistic about future returns by 1% is around 1% more likely to buy equities in a given window. This coefficient is very similar for the beginning-of-period beliefs and for the change in beliefs over the window. Since we do not know exactly at which point in the window the change in beliefs occurred, we do not interpret these results strongly in light of a particular model of trading behavior. Rather, it is likely that both the beginning-of-window beliefs and the change during the window help capture the individuals' true beliefs at the time of trade.

Finally, columns 7 and 8 explore the magnitude of trading conditional on a trade occurring. The dependent variable of the regression is the change in equity share due to trading: that is, the regression is the same as for columns 1 and 2, but the results are conditional on trading taking place (again, measured by the equity share changing by at least 1% in any direction due to trading during a window). The table reports that conditional on trading, the sensitivity of trading to beliefs increases substantially, by about 4 times, compared to the unconditional results of columns 1 and 2: a 1% increase in investors' beliefs about future returns corresponds to a 0.4-0.5% change in the equity share due to trading. When we condition on larger trades (at least a 5% change in equity share, not reported in the table) we find even bigger magnitudes, of around 0.7%-1%.

These results show that, consistent with the analysis in previous sections, trading activity occurs infrequently, but when it does occur, it lines up with beliefs with economically and statistically significant magnitudes.

The allocation of new funds. A final question we investigate is how respondents allocate "new funds" across the different asset classes depending on their beliefs. We measure these new funds

[28] We condition on trades that change the shares allocated to different asset classes. We are therefore excluding trades that involve buying and selling securities within the same asset class.

as cash entering a Vanguard account from outside sources; more specifically, we look at all cases in which during a window we see a net inflow of outside money into the account that is at least as large as 20% of the existing total Vanguard assets. Of course, we do not observe whether these are new funds, such as labor income, or proceeds from other asset sales outside of Vanguard.[29] Since investors devote some time to deciding an allocation for funds when they first transfer them to their trading account, this represents a particularly informative window to observe how beliefs affect portfolio formation. We repeat the regression of columns 7 and 10 of Table V, but conditioning on a large inflow occurring during the window. We find that when investors actually trade during that window (that is, they actively allocate the new money), the sensitivity of equity shares to beliefs increases significantly, to 0.5-0.7%, for both levels and changes. On the contrary, when little or no trading occurs during the window, we find that the sensitivity drops to around 0.01%, consistent with the simple fact that many investors who bring cash into the account but do not immediately invest it.

We summarize the results of this section in the following fact:

Fact 3. (Intensive and Extensive Margin of Trading). *While belief changes have little to no explanatory power for predicting when trading occurs (extensive margin of trading), they explain both the direction and magnitude of trading conditional on a trade occurring (intensive margin of trading).*

Other trading results. We have explored a variety of additional trading results, that we briefly summarize here. First, the results in Table V are robust to including measures of risk, like the probability of extremely negative returns, though this probability itself does not appear significantly. Second, we have repeated the analysis by looking at the 10-year expected returns instead of (or in addition to) the 1-year expected return. When both the short-term and long-term expectations are included in the regression, they both appear significant determinants of trading, with similar magnitudes. Third, we have explored the allocation to other asset classes, where the results are statistically weaker. For example, trading in bonds seems to have a statistically weak relation with beliefs about bond returns, even conditioning on trading.

VI Variance Decomposition of Beliefs

Section II showed substantial heterogeneity in investors' beliefs. In this section, we further explore this heterogeneity by decomposing the variance of beliefs into three components: fixed individual characteristics, common variation in individual beliefs over time, and a residual component that combines idiosyncratic individual time variation and measurement error. This decomposition allows us to establish two additional fact about beliefs. First, beliefs variation is mostly summarized by individual fixed effects, that is, individuals tend to have very large and persistent differences in their views, that strongly dominate the time variation in beliefs (Fact 4), at least during our sample

[29] We exclude direct transfers of securities, such as equities, bonds, or shares in mutual funds into Vanguard since in those cases we know that the investment already existed in the agent's portfolio.

period. Second, this difference in beliefs is difficult to explain with observable characteristics of the individuals (Fact 5).

VI.A The Dominance of Individual Fixed Effects

We let $B_{i,t}$ be the belief, i.e., the survey answer, expressed by individual i at time t. For example, for the question about 1-year expected stock market return we have $B_{i,t} = E_{i,t}[R_{t+1}]$. For the (unbalanced) panel of these beliefs we begin by running the following regressions:

$$B_{i,t} = \chi_t + \epsilon_{1,i,t}, \tag{4}$$

$$B_{i,t} = \phi_i + \epsilon_{2,i,t}, \tag{5}$$

$$B_{i,t} = \phi_{3,i} + \chi_{3,t} + \epsilon_{3,i,t}. \tag{6}$$

The regression in equation (4) estimates a set of time (i.e., survey wave) fixed effects, χ_t, that absorb the common variation over time of respondents' beliefs. The regression in equation (5) estimates a set of individual fixed effects, ϕ_i, that absorb the average belief over time of each respondent. The regression in equation (6) estimates both individual and time fixed effects together. We estimate these regressions both unweighted and weighted by respondent wealth, and only include individuals who have answered at least 3 waves.

Table VI: Decomposing the Variation in Beliefs: Individual and Time Fixed Effects

	R^2(%) – Unweighted				R^2(%) – Wealth-weighted			
	Reg (4)	Reg (5)	Reg (6)	N	Reg (4)	Reg (5)	Reg (6)	N
Expected Stock Ret (1yr)	1.0	62.3	63.0	8,015	1.4	62.0	63.0	7,481
Expected Stock Ret (10yr)	0.3	58.7	58.9	7,920	0.6	59.9	60.4	7,394
Prob. Ret<-10%	0.2	56.7	56.8	8,194	0.4	56.9	57.1	7,657
Std. Expected Ret. (1yr)	0.3	62.3	62.6	8,194	0.8	64.6	65.0	7,657
Confidence (Stock Qs)	0.2	66.1	66.2	8,100	0.4	68.5	68.6	7,559
Expected ΔGDP (3yr)	0.6	53.2	53.7	8,079	1.9	55.0	56.2	7,554
Expected ΔGDP (10yr)	0.5	50.9	51.2	7,949	1.3	49.6	49.9	7,430
Prob. ΔGDP<0%	0.3	54.2	54.4	8,173	0.7	57.3	57.7	7,639
Std. Expected ΔGDP (3yr)	0.2	43.7	43.9	8,156	1.7	43.5	43.8	7,622
Confidence (GDP Qs)	0.1	67.8	67.8	8,083	0.3	68.2	68.5	7,544
Expected Bond Ret (1yr)	0.4	56.9	57.3	7,981	2.5	58.4	59.1	7,464
Confidence (Bond Qs)	0.1	66.2	66.3	7,970	0.2	67.6	67.7	7,439

Note: Table shows R^2s corresponding to the three regressions 4, 5, and 6. Each row corresponds to a different question in the survey. The left panel uses unweighted regressions, while the right panel uses wealth-weighted regressions.

Table VI reports the R^2 of each of the three regressions (unweighted on the left side and wealth-weighted on the right side), for a subset of questions from our survey. The table shows that most of the variation in belief is absorbed by individual fixed effects. Consider for example the first row, corresponding to the question on 1-year expected stock market return. Time fixed effects capture about 1% of the total variation in the panel, whereas individual fixed effects capture more than 62% of the total variation. This large difference in explanatory power is common across all beliefs of interest: expected long-run stock market returns, GDP growth in the short and

long run, stock market or GDP disasters, bond returns, and confidence. The same is true when we decompose the heterogeneity in individuals' confidence in their beliefs: most of the variation is across individuals rather than over time. These conclusions change little when beliefs are weighted by the wealth of the respondent.

Fact 4. (Large and persistent heterogeneity). *Individual beliefs as well as their confidence in their own beliefs are mostly characterized by heterogeneous individual fixed effects. In our panel, between 50% and 60% of all variation is due to individual fixed effects and 1% to common time series variation. The remaining variation is accounted for by idiosyncratic individual variation over time or measurement error.*

One worry with this analysis is that individual fixed effects might appear more important than they truly are because of our short time span (10 waves) and the relatively small number of individuals that answer multiple waves. This might occur, for example, if individuals' beliefs followed a persistent autoregressive process. We investigated this possibility in two different ways. First, we added to the panel regression an autoregressive term. This allows us to potentially distinguish whether the observed persistence of beliefs is due to a true "fixed effect" of each individual (i.e., a long-run mean belief), as opposed to a temporary but persistent fluctuation of an individual's beliefs around his or her long-run mean. For all the specifications reported above, this autoregressive term was small, statistically insignificant, and did not account for much of the variation, thus confirming the dominance of fixed effects in explaining the panel variation in our sample. As a second robustness test, we vary from 3 to 6 the minimum number of waves that an individual has had to have answered to be included in the sample; this allows us to address concerns that the high explanatory power of individual fixed effects arises mainly from in-sample over-fitting of noise. Panel A of Table VII shows how the R^2 changes for each question as we increase the requirement on the number of answers per individual. Panel B reports the number of observations in each case. The R^2 are remarkably stable across columns: our results show at most a modest deterioration in the importance of individual fixed effects (as measured by the R^2) as we increase the minimum number of answers. We conclude that in-sample over-fitting of noise does not drive the explanatory power of the individual fixed effects.

While we cannot ultimately rule out that, as the survey continues in different economic circumstances (recession, stock market crash), the importance of individual fixed effects will change, their overwhelming importance compared to time fixed effects in our sample is striking.

VI.B Beliefs and Demographics: The Beliefs Disconnect

Having established the overwhelming importance of individual fixed effects in explaining the panel-variation in beliefs, it is a natural question to ask which observables explain why some individuals are optimistic and others are pessimistic. In our data, observable individual characteristics have little or no explanatory power for beliefs, a fact we dub the *Beliefs Disconnect*. This is true despite including measures such as: gender, age, wealth, attention (days with visit), con-

Table VII: Decomposing the Variation in Beliefs: Robustness

Panel A: R^2 **(total,%)**	#Resp≥3	#Resp≥4	#Resp≥5	#Resp≥6
Expected Stock Ret (1yr)	62.3	59.6	58.2	54.9
Expected Stock Ret (10yr)	58.7	56.9	57.5	57.4
Prob. Ret<-10%	56.7	55.9	56.1	55.5
Std. Expected Ret. (1yr)	62.3	61.4	61.8	61.3
Confidence (Stock Qs)	66.1	64.4	63.2	64.3
Expected ΔGDP (3yr)	53.2	51.5	49.9	48.8
Expected ΔGDP (10yr)	50.9	50.1	50.7	51.6
Prob. ΔGDP<0%	54.2	52.4	50.9	50.8
Std. Expected ΔGDP (3yr)	43.7	39.7	37.5	37.1
Confidence (GDP Qs)	67.8	65.4	63.4	63.4
Expected Bond Ret (1yr)	56.9	54.2	53.0	52.8
Confidence (Bond Qs)	66.2	64.7	64.0	63.1

Panel B: N. of obs.	#Resp≥3	#Resp≥4	#Resp≥5	#Resp≥6
Expected Stock Ret (1yr)	8,015	4,970	3,138	1,998
Expected Stock Ret (10yr)	7,920	4,965	3,113	2,003
Prob. Ret<-10%	8,194	5,128	3,208	2,038
Std. Expected Ret. (1yr)	8,194	5,128	3,208	2,038
Confidence (Stock Qs)	8,100	5,040	3,144	2,014
Expected ΔGDP (3yr)	8,079	5,061	3,173	2,018
Expected ΔGDP (10yr)	7,949	4,982	3,118	2,003
Prob. ΔGDP<0%	8,173	5,131	3,211	2,051
Std. Expected ΔGDP (3yr)	8,156	5,120	3,192	2,037
Confidence (GDP Qs)	8,083	5,047	3,139	2,029
Expected Bond Ret (1yr)	7,981	5,023	3,135	2,015
Confidence (Bond Qs)	7,970	4,991	3,119	1,999

Note: Panel A of the table reports the R^2 statistics corresponding to regression 5. In each column, going from left to right, we increase from 3 to 6 the minimum number of responses for an individual to be included in the sample. Panel B reports the number of observations. Each row corresponds to a different question in the survey.

fidence, and geographic location.[30] Similarly, we know from the analysis of Section II that while these individual fixed effects are reflected in portfolio choices, the pass-trough of beliefs to portfolios is much smaller than canonical theories would imply. To establish the belief disconnect, we run the following econometric analysis:

$$\phi_{3,i} = \alpha + \Gamma X_i + \epsilon_i, \tag{7}$$

where $\phi_{3,i}$ are the individual fixed effects estimated in regression 6, and X_i are individual static characteristics.[31] Table VIII focuses on the explanatory power of these regressions – the R^2 – while the Appendix provides the details of the estimated coefficients. The observed characteristics have only minimal explanatory power, with values for the R^2 in the order of 2%-5%. Importantly, when we restrict the analysis to explaining fixed effects that are estimated on more observations,

[30] We are in the process of constructing past individual returns on the Vanguard portfolio to verify whether those respondents the portfolios of whom have outperformed in the past tend to be more optimistic.

[31] For dynamic variables such as age and wealth we take the average over the sample. For geographic location and gender we take the value of the most recent observation.

and which should thus be more precise, there is only a modest increase in R^2.[32] We conclude that classic measurement error in beliefs cannot be responsible for the belief disconnect. Despite the low overall explanatory power of demographics for beliefs, the Appendix shows that some characteristics have a statistically significant relationship with the beliefs. For example, focusing on explaining 1-year expected stock returns in a multivariate regression, we find that: (i) wealthier respondents are less optimistic, (ii) older respondents are more optimistic, (iii) males are more pessimistic, (iv) respondents who pay more attention (log-in more often) are more optimistic, and (v) respondents who are more confident are more optimistic. Despite these effects, and their statistical significance, the overall explanatory power of these demographics for the variation in beliefs is negligible.

Table VIII: Beliefs Heterogeneity and Demographics

R^2	#Resp≥1	#Resp≥2	#Resp≥3	#Resp≥4	#Resp≥5
Expected Stock Ret (1yr)	2.7	4.0	5.0	4.9	5.6
Expected Stock Ret (10yr)	2.3	3.3	4.2	5.4	8.2
Prob. Ret<-10%	3.1	5.0	7.1	9.2	11.9
Std. Expected Ret. (1yr)	8.2	10.7	13.3	16.3	17.4
Expected ΔGDP (3yr)	2.0	3.6	3.8	4.1	9.3
Expected ΔGDP (10yr)	2.7	4.6	4.0	5.5	11.1
Prob. ΔGDP<0%	6.1	8.7	9.4	12.3	10.8
Std. Expected ΔGDP (3yr)	3.3	5.0	5.4	7.5	9.6
Expected Bond Ret (1yr)	3.0	4.1	3.5	4.2	7.6

Note: The table reports the R^2 statistics corresponding to regression 7. In each column, going from left to right, we increase from 1 to 5 the minimum number of responses for an individual to be included in the sample. Each row corresponds to a different question in the survey.

Fact 5. (Belief Disconnect). *The heterogeneity and persistence in individual beliefs documented in Fact 4 is not explained by observable characteristics such as gender, age, wealth, attention, confidence, and geographic location. These characteristics have a strong statistical relationship with beliefs but their explanatory power is extremely limited.*

These results relate to the literature linking expectations to socioeconomic status and personal experience. It is not uncommon in this literature to find strong statistical relationships but low explanatory power between expectations and variables such as wealth, gender, and place of birth, location, own past experience, or friends' past experiences (Malmendier and Nagel, 2011; Kuchler and Zafar, 2015; Das, Kuhnen and Nagel, 2017; Bailey et al., 2017, 2018).[33] Our results above highlight that the low R^2 in this type of regressions is unlikely to be primarily due to measurement errors in beliefs. An exciting possibility is that these individual fixed effects therefore reflect combinations of individual characteristics and experiences that economic research is yet to discover.

[32] Some increase in the R^2 is to be expected since the number of observations decreases quickly and we include a large number of regressors.

[33] We have thus far not included direct measures of personal experience, like own stock market performance, in our regressions, but are in the process of doing so. Some elements of personal or cohort experience are potentially captured by age and geographic location.

If one takes the heterogeneity in individual fixed effects to be correctly measured, and considers the poor explanatory power of the demographic characteristics we observe, then an important challenge for future research is to find what determines such large individual differences.

VI.C Implications of Beliefs Heterogeneity and Disconnect For Asset Pricing Theory

The preceding section has analyzed the variation in beliefs across individuals and over time. We next discuss the lessons for asset pricing theory that arise from this analysis. We find it useful to consider two stylized classes of models used in the theoretical literature. At one extreme, consider a model in which all individuals hold the same beliefs and the only variation is in the time series. At the other extreme, consider a model in which agents have permanent and constant differences in beliefs. The first class of models aims to generate variation in asset prices by swings in individual beliefs over time, which then translate immediately in swings in the representative agent's beliefs. The second class of models aims to generate variation in asset prices by movements in the wealth-weighted aggregate belief. Intuitively, this requires several steps: first, individual heterogeneity in beliefs induces heterogeneity in portfolios; second, shocks over time redistribute wealth across individuals given the differences in portfolios;[34] in turn, this affects equilibrium prices.

Our research highlights that these two stylized models miss core elements of the data. The first class of models is silent on a major feature of the data, the individual heterogeneity, and gets all its action from what appears to be a small component (1%) of the total variation in beliefs. In light of our results, we believe that the literature would benefit from explicitly modeling the aggregation of the heterogeneity in individual beliefs.

The second class of models captures the individual static heterogeneity, but gets all its action by having this heterogeneity strongly reflected in portfolios. The findings discussed in Section IV highlight that this second step is challenging. Beliefs differences are reflected in portfolios in a statistically reliable manner, but their impact on portfolios is quantitatively limited. When we regress, as we do in Column 1 in Table III, portfolios on beliefs and individual characteristics we obtain an R^2 around 11%, but controls alone have an R^2 of 8%, thus highlighting the limited sense in which optimists have riskier portfolios.[35] If individuals with vastly different beliefs do not on average have vastly different portfolios, then shocks do not substantially redistribute wealth between optimists and pessimists, thus vastly reducing the variation in aggregate wealth weighted belief. In Section IV we discussed a number of possible adjustments to the second class of models that might help better fit the data, including taking explicit account of infrequent trading and confidence in beliefs. Whether models with these adjustments can generate large swings in asset pricing is an interesting area of future research.

[34] For example, optimists have portfolios that load more than pessimists' portfolios on good economic outcomes; if these outcomes are realized, then optimists relative wealth increases compared to pessimists.

[35] Among the controls the key contributor to the explanatory power is age. We find that older people have significantly less equity in their portfolio.

VII Is it Variation in Discount Rates or Variation in Expected Cash Flows?

In this section, we explore how beliefs about different quantities relate to each other. Whereas our survey asks separately about stock market returns and economic growth, these quantities are naturally linked in equilibrium. To relate these variables, we make the imperfect but useful abstraction that considers the economy's GDP to be the "dividend" paid to the holders of the stock market, so that we can proxy for beliefs about dividend growth using beliefs about GDP growth. This allows us to apply the Campbell and Shiller (1988) decomposition in our context. The decomposition states that an identity links prices, expected stock market returns, and expected cash-flow growth according to:

$$pd_t^i = \frac{1-\rho^n}{1-\rho}k + E_t^i \sum_{j=0}^{n-1} \rho^j (\Delta d_{t+1+j} - r_{t+1+j}) + \rho^n E_t^i pd_{t+n}, \tag{8}$$

where pd_t is the logarithm of the price dividend ratio, Δd_{t+1} is the growth in dividends between time t and $t+1$, and r_{t+1} is the return of the stock market between between time t and $t+1$.[36] This decomposition has to hold individual by individual, so that we use the individual expectation operator $E_t^i[.]$ to stress that this reflects individual beliefs about future growth and returns. In addition, as noted by Campbell (2017), the decomposition does not require that individuals have rational expectations, but only that the expectations do not violate mathematical identities. The current price-dividend ratio is in theory observed by all investors. Therefore, in what follows we assume that all investors agree on the observed value, and we impose: $pd_t^i = pd_t$.
We map the survey answers to the terms in equation 8 by imposing that:

$$Q_{3d,t}^i = E_t^i \Delta d_{t+1} = E_t^i \Delta d_{t+2} = E_t^i \Delta d_{t+3}, = E_t^i \Delta d_S \tag{9}$$

$$Q_{r1,t}^i = E_t r_{t+1} = E_t r_S, \tag{10}$$

$$Q_{d10,t}^i = \frac{3}{10} E_t \Delta d_S + \frac{7}{10} E_t \Delta d_L, \tag{11}$$

$$Q_{r10,t}^i = \frac{1}{10} E_t \Delta r_S + \frac{9}{10} E_t r_L. \tag{12}$$

Equation 9 imposes that the average annualized expected growth rate of GDP over the next three years, i.e., the answer to the question about short-term growth, maps into the expected growth rate of dividends in each of the next three years ($E_t^i \Delta d_S$). Equation 10 simply sets up the notation for short-term (i.e., one year) expected stock market returns ($E_t r_S$). Equation 11 maps the survey question about long-run economic growth ($Q_{d10,t}^i$) into long-run dividend growth ($E_t \Delta d_L$), defined as dividend growth between years 3 and 10. Equation 12 similarly maps the survey question about

[36] ρ and κ are log-linearization constants; the derivation is in Campbell and Shiller (1988). Here it suffices to think of ρ as a number close to 1, and κ as a positive constant.

ten-year stock market returns ($Q^i_{r10,t}$) into long-run (between year 1 and 10) stock market returns ($E_t \Delta r_L$). The Campbell-Shiller decomposition requires that dividend growth and returns are either both nominal or both real. In our survey, however, the stock returns questions are nominal and the GDP growth questions are real.[37] In mapping the answers to this decomposition, we make the assumption that expected inflation is constant, in which case equation 8 still applies with an adjustment to the constant.

From the answers to the four survey questions about long- and short-run economic growth and stock market returns, we therefore obtain estimates of the terms: $E^i_t \Delta d_S$, $E^i_t \Delta d_L$, $E^i_t r_S$, $E^i_t \Delta r_L$. We rewrite equation 8 by grouping the terms of interest:

$$pd_t = \frac{1-\rho^{10}}{1-\rho}k + \frac{1-\rho^3}{1-\rho}E^i_t \Delta d_S + \frac{\rho^3-\rho^{10}}{1-\rho}E^i_t \Delta d_L - E^i_t r_S - \frac{\rho-\rho^{10}}{1-\rho}E^i_t r_L + \rho^{10}E^i_t pd_{t+10}. \tag{13}$$

Since the term pd_t is observable at each point in time, and having calibrated $\rho = 0.96$, we use this identity to back out the expectation of each individual about the price-dividend ratio ten years into the future ($E^i_t pd_{t+10}$).[38] Once we have obtained this "residual" term, we can then interpret this equation as follows. All investors agree on the (observable) price-dividend ratio pd_t. However, they disagree about its determinants. Some investors, for example, might expect higher dividend growth, but to square it with the observed price-dividend ratio they need to also expect either higher returns or a lower price-dividend ratio ten years from now. This identity, therefore, disciplines the relationship between the survey answers given by each individual.

Table IX: Within-Response Correlation Across Answers

Corr/std	$E^i_t r_S$	$E^i_t r_L$	$E^i_t \Delta d_S$	$E^i_t \Delta d_L$	$E^i_t pd_{t+10}$
$E^i_t r_S$	0.05				
$E^i_t r_L$	0.21	0.03			
$E^i_t \Delta d_S$	0.25	0.17	0.02		
$E^i_t \Delta d_L$	0.07	0.23	0.41	0.02	
$E^i_t pd_{t+10}$	0.34	0.76	-0.20	-0.36	0.40

Note: Table shows the correlation across answers given by the same individual in the same survey wave. The diagonal includes standard deviations measured as the dispersion in beliefs across respondents.

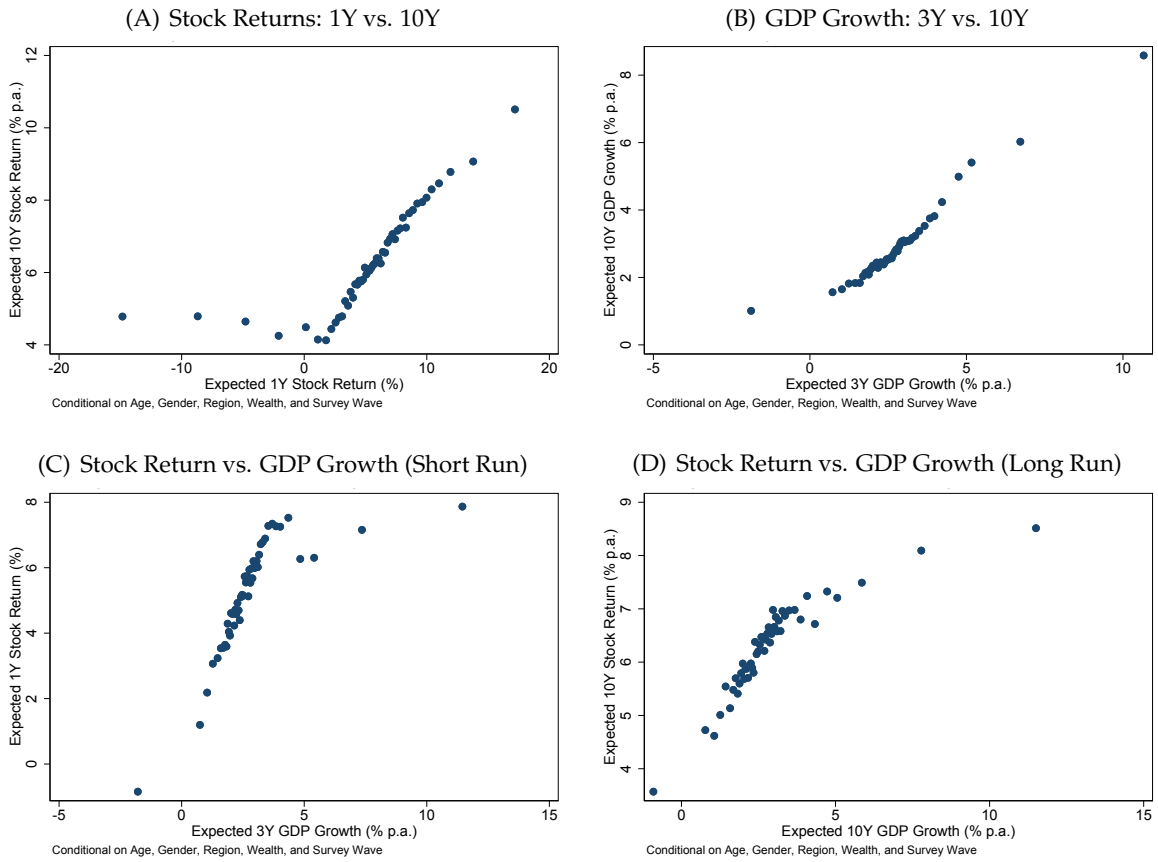
Table IX reports the correlation of the variables $E^i_t \Delta d_S, E^i_t \Delta d_L, E^i_t r_S, E^i_t \Delta r_L$ and $E^i_t pd_{t+10}$, with standard deviations on the diagonal. Short-run and long-run expected returns are positively correlated (correlation of 0.21). Panel A of Figure VII further illustrates this point by presenting conditional binscatter plots showing the relationship between short-run and long-run expectations: interestingly, even those respondents who expect negative returns over the next year, expect long-term returns to be positive. Similarly, short-term and long-term dividend growth are positively

[37] Our decision was driven by a desire to ask about objects that are more salient to the respondents. In the general press, stock market returns are almost always reported in nominal terms, while economic growth is reported in real terms.

[38] We calibrate ρ using the average log price-earning ratio over the period 1881-2018, using data provided by Professor Shiller on his website.

correlated (correlation of 0.41), a pattern that is also evidenced by Panel B of Figure VII. Table IX and Panels C and D of Figure VII also show that expectations of stock market returns and economic growth are positively related at all horizons. In unreported results, we confirm that the same correlations hold when we compare beliefs within individuals over time: when individuals become more optimistic about short-run returns (cash flows), they also become more optimistic about long-run returns (cash-flows). And when they become more optimistic about returns, they also become more optimistic about cash-flows.

Figure VII: Relationships Among Different Beliefs Within the Same Survey



Note: Figure shows conditional binned scatterplots across different answers from within the same survey response.

Table IX also highlights that the standard deviation (across respondents) of short-term expected returns is twice as large as the standard deviation of (annualized) long-term expected returns and short-term and long-term dividend growth rates (see also Table II). All standard deviations, however, are dominated by that of the residual term $E_t^i pd_{t+10}$. This indicates that the expected returns and dividends obtained from the survey – taken alone – are not consistent with a zero or constant residual as implied by the Campbell-Shiller decomposition. The term $E_t^i pd_{t+10}$ simply picks up the unexplained difference between the actual price-dividend ratio and that implied by the quantities we measure in the survey.

We next use the Campbell-Shiller decomposition to perform a variance decomposition. Since equation 13 is an identity, the variance of any term can be decomposed into the fractions explained by the covariation with the remaining terms. The contribution of each term to total variation is not required to be positive, but the sum of all individual terms contributions is always 100%.[39]

Table X: A Behavioral Campbell-Shiller Variance Decomposition

	$-E_t^i r_S$	$-\frac{\rho-\rho^{10}}{1-\rho}E_t^i r_L$	$\frac{1-\rho^3}{1-\rho}E_t^i \Delta d_S$	$\frac{\rho^3-\rho^{10}}{1-\rho}E_t^i \Delta d_L$	$\rho^{10}E_t^i pd_{t+10}$	$-pd_t$
$-E_t^i r_S$	1	-0.06	0.27	0.04	0.08	-0.06
$-\frac{\rho-\rho^{10}}{1-\rho}E_t^i r_L$	-0.73	1	0.68	0.43	0.66	-0.03
$\frac{1-\rho^3}{1-\rho}E_t^i \Delta d_S$	0.22	0.04	1	-0.19	0.04	0.02
$\frac{\rho^3-\rho^{10}}{1-\rho}E_t^i \Delta d_L$	0.14	0.13	-0.90	1	0.17	-0.16
$\rho^{10}E_t^i pd_{t+10}$	1.40	0.89	0.93	0.75	1	1.23
$-pd_t$	-0.04	0.00	0.02	-0.02	0.04	1

Note: Variance decomposition of equation 13. In columns we report the term the variance of which is being decomposed using the variation in each of the terms in the rows.

Table X reports the variance decomposition of each of the terms, one per column, in terms of all the remaining terms reported in the rows. The interpretation of each column is that positive numbers absorb variation while negative numbers add further variation for the remaining terms to account for. Consider, for example, the first column which reports the decomposition of the variance of short-term returns $E_t^i r_S$. The covariance between short-term expected returns and short-term dividend growth, for example, helps explains 22% of the variance of short-term expected returns. Similarly, the covariance between short-term expected and long-term dividend growth helps explain another 14% of the total variance of short-term expected returns. Going back to equation 13, we can interpret this as follows. Investors that have higher expectations of future returns also tend to have higher expectations of dividend growth. This helps match the price-dividend ratio, since part of the positive effect (higher current price-dividend ratio) of higher expected future cashflows is offset by discounting these cashflows at higher rates. Instead, expectations about long-term expected returns explain a *negative* portion of the total variance of short-term expected returns (-73%). Short and long-run expected returns are positively correlated; in the variance decomposition, this means that long-run expected returns add variance to short-run expected returns – holding the current price-dividend ratio fixed. In equation 13, when short-run returns are high, long-run returns are also high, thus creating a larger deviation from the current price dividend ratio that needs to be accounted for by the remaining terms in the equation. Expected dividend growth in both the short and the long run offsets some of this variation, but this is not enough to completely balance the equation (i.e. the sum of the variance contribution

[39] Generally, if $X + Y + Z = 0$, then one can write $-X = Y + Z$, and $Var(X) = Cov(-X, -X) = Cov(-X, Y) + Cov(-X, Z)$. Scaling by the variance of X, we have: $\frac{Var(X)}{Var(X)} = \frac{Cov(-X,Y)}{Var(X)} + \frac{Cov(-X,Z)}{Var(X)} = 100\%$. The total (100%) variance of X can be decomposed into the sum of two components, a fraction due to the covariation of X with Y, the term $\left(\frac{Cov(-X,Y)}{Var(X)}\right)$, and a fraction due to the covariation of X with Z, the term $\left(\frac{Cov(-X,Z)}{Var(X)}\right)$. The fractions sum up to 100% by definition.

of long-run returns and short- and long-run dividend growth does not sum to 1). The slack is absorbed by the residual term: the expectation of the price dividend ratio in ten years.

A similar pattern occurs in Column 3 of Table X which decomposes the variance of short-term growth expectations. Long-term growth expectations are positively correlated with short-term ones, thus adding to the total variance (-90%). Expected returns at all horizons absorb some of the variation, but not enough to completely offset it, so that the slack is again absorbed by the residual term: the expectation of the price dividend ratio in ten years. We collect the results in this subsection thus far in Fact 6 below.

Fact 6. *Both expected returns and expected cashflows contribute to variation in asset prices. Across individuals, higher expected cash flows are partially offset by higher expected returns (i.e., by higher discount rates) in both the short and long run, and vice-versa. However, variation in expected returns and expected cashflows does not entirely offset, given the current price dividend ratio, and the slack is absorbed by a residual term: the expectation of the price dividend ratio in ten years.*

Finally, Column 6 of Table X decomposes the variance of the price dividend ratio, which reflects only time-series variation. With only 10 survey waves (spaced two months apart each), our time series variation is extremely limited. With that caveat in mind, the results suggest that higher price-dividend ratios are associated with higher expected returns at all horizons. This effect is consistent with the return extrapolation patterns documented by Greenwood and Shleifer (2014), and runs contrary to the implications of rational asset pricing models, which usually explain higher price-dividend ratios with lower expected returns. Similarly, higher price-dividend ratios are associated with lower expected growth (in the long run). This runs contrary to the behavioral finance literature that explains higher price-dividend ratios with higher expected cash flows.

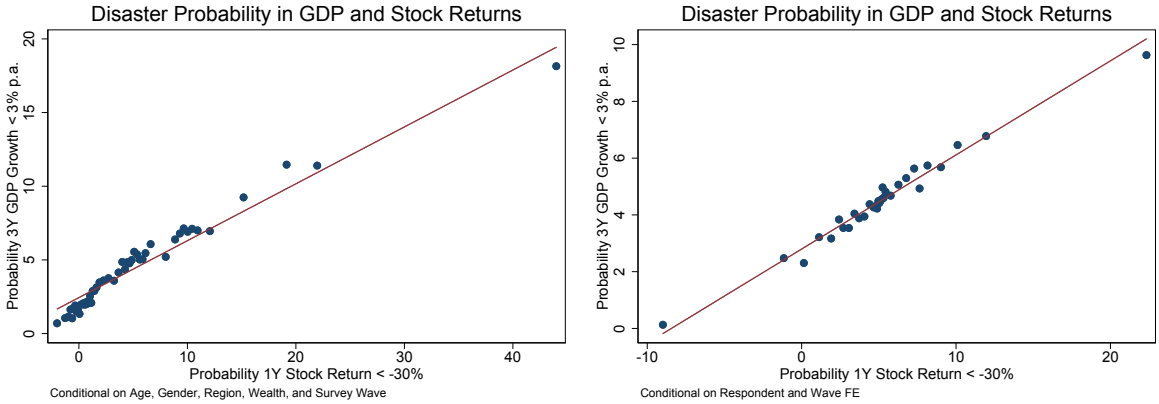
Implications for Asset Pricing Models. The behavioral finance literature has alternated between modeling cash-flow expectations and return expectations (Cutler, Poterba and Summers, 1990; De Long et al., 1990*a*,*b*; Barberis et al., 2015). Our research shows that both dimensions are important to understand the data and that, at least for quantitative evaluations, their dynamics have to be modeled jointly. Across individuals, we find that higher expected cash flows come together with higher expected returns, thus in part mitigating the effect on prices. However, the effect is not strong enough, and as a result much of the variation remains unexplained and is absorbed in our setting by the expectation of the price dividend ratio ten years into the future. As our survey keeps running and more time-series evidence accumulates, we aim to further characterize how changes in expected returns and in expected cashflows contribute to changes in the price-dividend ratio.

VIII Going Beyond the Mean and the Variance: A Behavioral Approach to Rare Disasters

In the previous section, we focused on those moments of the belief distribution that have been of most central interest to the asset pricing literature, such as the mean and variance of expected returns and growth rates. However, an important strand of this literature has emphasized that expectations of rare but potentially catastrophic negative events, sometimes called rare disasters, can help explain portfolio holdings and asset prices (Rietz, 1988; Barro, 2006; Gabaix, 2012). This literature stresses that the expected probability and size of rare disasters is one of the most important moments to understand risks and returns. In this Section, we explore a number of features of these disaster probabilities, and discuss their implications for the theoretical literature.

Our survey was specifically designed to elicit the expectations of rare disasters both in stock market returns and GDP growth. The distribution questions for both 1-year stock market returns and 3-year GDP growth rates included extremely negative buckets (returns less than -30% for 1-year stock returns and less than -3% for annualized 3-year GDP growth) that are routinely identified as rare disaster thresholds. Section IV.B confirmed the importance of rare-disaster information in understanding portfolio formation. Indeed, holding fixed the mean and standard deviation of beliefs, respondents who reported a higher perceived probability of stock market disasters also held a lower equity share (see column 3 of Table IV).

Figure VIII: Disaster Probability in GDP and Stock Returns

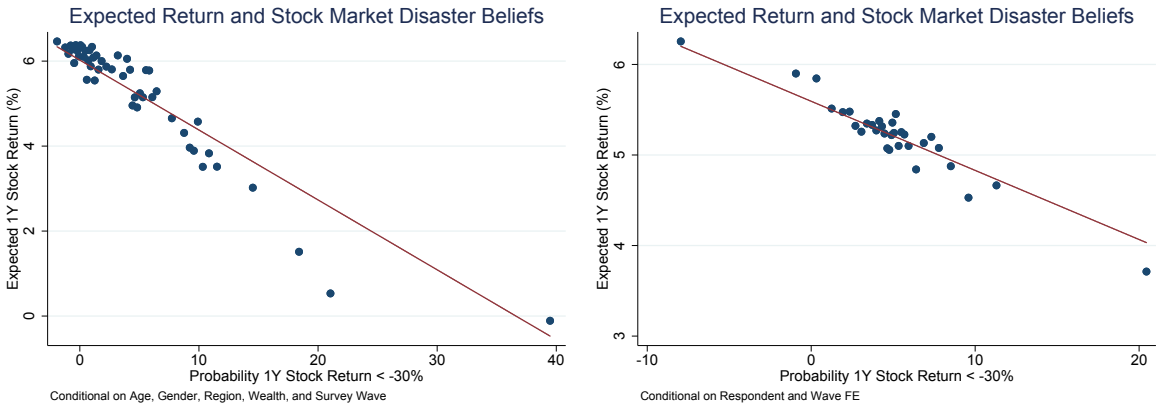


Note: The left panel shows a conditional binned scatter plot of survey respondents' subjective probability that the 1-year stock returns are below 30% and their expectation that annualized average GDP growth over the next three years is below -3%. The left panel conditions on the respondents' age, gender, region, wealth, and the survey wave; the right panel conditions on respondent and wave fixed effects.

We first explore the relationship between individuals' expectations of the probabilities of stock market and GDP disasters. Figure VIII shows that expectations of the two types of disasters are positively related at the individual level. Indeed, the left panel shows that those individuals who expect the economy to shrink more than 3% a year for the next three years are also more likely to report a high probability that the stock market will crash by more then 30% over the next year

(the slope of the regression line is 0.38).[40] The right panel conditions on individual fixed effects; it shows that when the same individual perceives an increased stock market disaster probability, this is associated with an increased perceived GDP disaster probability (the slope of the regression line is 0.33). These findings suggest that expectations of rare stock market disasters come with expectations of lower cash flows and are not just purely the result of expecting higher future returns (i.e., stock market disasters are not purely due to discount rate variation).

Figure IX: Disaster Probability in GDP and Expected Stock Returns



Note: The left panel shows a conditional binned scatter plot of survey respondents' subjective probability that the 1-year stock returns are below 30% and their expectation that annualized GDP growth over the next three years is below -3%. The left panel scondition on the respondents' age, gender, region, wealth, and the survey wave; the right panel conditions on respondent and wave fixed effects.

We next explore the relationship between expected returns and expected disaster probabilities. We establish the following fact:

Fact 7. *When individuals expect large stock market declines ("rare disasters") to occur with higher probability, they also expect stock market returns to be lower.*

Figure IX shows that individuals with a higher subjective probability of stock-market rare disasters have lower expected stock-market returns. The associated regression coefficient in the left panel, which uses the entire panel variation, is a precisely estimated -0.16; this estimate implies that a 5 percentage point increase in individuals' subjective probability of a stock market disaster reduces their expected return by 0.8 percentage points. The same pattern holds true when we compare across answers within individuals (right panel, that adds individual fixed effects), though with a slightly lower regression coefficient of -0.07. The negative relationship between expected return and disaster probability is extremely robust. It is independent of (i) the ordering of buckets for possible stock market outcomes; (ii) the wealth of respondents; (iii) the time it takes respondents to answer the questions; (iv) how easy individuals perceive the stock questions to be; and (v) the confidence of investors in their stock market beliefs. It is also robust to exclud-

[40] A similar picture emerges when considering less extreme outcomes, for example a less than 0% growth rate and a 10% stock market crash.

ing individuals with a zero probability assigned to stock market disaster, as well as to excluding individuals who report very large probabilities of stock market disasters.

Connection to Rare Disaster Models. In equilibrium models with rare disasters, expected returns and the probability of disaster are positively related. The intuition is that a higher probability of disaster induces individuals to demand a higher compensation for holding the stock market and this makes equilibrium expected returns higher. The closest mapping between our empirical results and theoretical models is between the right panel of Figure IX and the time-varying rare-disaster framework of Gabaix (2012). In that framework, a representative agent with rational expectations prices assets in an economy affected by time-varying rare disasters.[41] When the representative agent expects a higher probability of a disaster she also expects returns to be higher. The relationship in the data appears with an opposite sign relative to the theory.[42] This poses a challenge to rational expectations models of rare disasters.

However, we do not think of our results as inconsistent with the importance of rare disasters for asset pricing more generally. In fact, it is likely that modifications to existing models that incorporate at least some behavioral agents can more closely align model predictions and data. One possible approach would be to model some agents as over-confident or as thinking in partial equilibrium. When these agents believe that rare disasters have become more likely, they do not take into account that other agents might think the same, and therefore do not incorporate the fact that stock prices might already be lower to account for the increased risk. Rather, they think that prices are still high, and expect returns to be lower going forward. They therefore sell stocks to reduce their exposure, by more than they would if they anticipated that the compensation for holding on to the risk has increased in equilibrium. In turn, this force may induce more selling pressure (and therefore a bigger price effect) in response to an increase in the disaster probability. Of course, there might also be rational agents who take the opposite side of these trades; these rational agents require higher returns to compensate for the risk and to absorb the stocks that the behavioral agents are selling. To summarize, in this simple sketch of a model: (i) rational expected returns are higher when the probability of disaster is higher, but (ii) most people are behavioral and answer the exact opposite in surveys. Part of the reason for the effect in point (i) is that behavioral selling induced by point (ii) is a priced risk. It is beyond the scope of this paper to provide a full model of the behavioral approach to rare disasters, but we expect that it will prove an interesting area of future research.

[41] The model specifies jointly the dynamics of the probability and size of disaster, which we map for simplicity into pure variation in probability.

[42] Comparing our estimates with the model of Barro (2006) is more indirect since, in that model, the disaster probability and expected (rational) returns are both constant over time. A rough comparison can be obtained by studying comparative statics of the model: again, a higher probability of disaster corresponds to a higher expected return.

REFERENCES

Abel, Andrew B, Janice C Eberly, and Stavros Panageas. 2007. "Optimal inattention to the stock market." *American Economic Review*, 97(2): 244–249.

Agnew, Julie, Pierluigi Balduzzi, and Annika Sunden. 2003. "Portfolio choice and trading in a large 401 (k) plan." *American Economic Review*, 93(1): 193–215.

Alvarez, Fernando, Luigi Guiso, and Francesco Lippi. 2012. "Durable consumption and asset management with transaction and observation costs." *American Economic Review*, 102(5): 2272–2300.

Ameriks, John, Gábor Kézdi, Minjoon Lee, and Matthew D Shapiro. 2016. "Heterogeneity in Expectations, Risk Tolerance, and Household Stock Shares." *Ann Arbor*, 1001: 48109–1248.

Amromin, Gene, and Steven A Sharpe. 2013. "From the horse's mouth: Economic conditions and investor expectations of risk and return." *Management Science*, 60(4): 845–866.

Bailey, Michael, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel. 2017. "House price beliefs and mortgage leverage choice." National Bureau of Economic Research.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018. "The economic effects of social networks: Evidence from the housing market." *Journal of Political Economy*, 126(6): 000–000.

Barberis, Nicholas C. 2018. "Psychology-based Models of Asset Prices and Trading Volume." National Bureau of Economic Research.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2015. "X-CAPM: An extrapolative capital asset pricing model." *Journal of financial economics*, 115(1): 1–24.

Barro, Robert J. 2006. "Rare disasters and asset markets in the twentieth century." *Quarterly Journal of Economics*, 121(3): 823–866.

Campbell, John Y. 2017. *Financial decisions and markets: a course in asset pricing*. Princeton University Press.

Campbell, John Y, and Robert J Shiller. 1988. "The dividend-price ratio and expectations of future dividends and discount factors." *The Review of Financial Studies*, 1(3): 195–228.

Choi, James J, David Laibson, Brigitte C Madrian, and Andrew Metrick. 2009. "Reinforcement learning and savings behavior." *The Journal of Finance*, 64(6): 2515–2534.

Cochrane, John H. 2011. "Presidential address: Discount rates." *The Journal of finance*, 66(4): 1047–1108.

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

Cutler, David M, James M Poterba, and Lawrence H Summers. 1990. "Speculative Dynamics and the Role of Feedback Traders." *The American Economic Review*, 80(2): 63–68.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam. 1998. "Investor psychology and security market under-and overreactions." *the Journal of Finance*, 53(6): 1839–1885.

Das, Sreyoshi, Camelia M Kuhnen, and Stefan Nagel. 2017. "Socioeconomic status and macroeconomic expectations." National Bureau of Economic Research.

De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann. 1990*a*. "Noise trader risk in financial markets." *Journal of political Economy*, 98(4): 703–738.

De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann. 1990*b*. "Positive feedback investment strategies and destabilizing rational speculation." *the Journal of Finance*, 45(2): 379–395.

Dominitz, Jeff, and Charles F Manski. 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.

Drerup, Tilman, Benjamin Enke, and Hans-Martin Von Gaudecker. 2017. "The precision of subjective data and the explanatory power of economic models." *Journal of Econometrics*, 200(2): 378–389.

Duffie, Darrell, Tong-sheng Sun, et al. 1990. "Transactions costs and portfolio choice in a discrete-continuous-time setting." *Journal of Economic dynamics and control*, 14(1): 35–51.

Gabaix, Xavier. 2012. "Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance." *Quarterly Journal of Economics*, 127(2): 645–700.

Gabaix, Xavier, and David Laibson. 2001. "The 6D bias and the equity-premium puzzle." *NBER macroeconomics annual*, 16: 257–312.

Gervais, Simon, and Terrance Odean. 2001. "Learning to Be Overconfident." *The Review of Financial Studies*, 14(1): 1–27.
Greenwood, Robin, and Andrei Shleifer. 2014. "Expectations of returns and expected returns." *The Review of Financial Studies*, 27(3): 714–746.
Harrison, J Michael, and David M Kreps. 1978. "Speculative investor behavior in a stock market with heterogeneous expectations." *The Quarterly Journal of Economics*, 92(2): 323–336.
Hirshleifer, David, and Siew Hong Teoh. 2003. "Limited attention, information disclosure, and financial reporting." *Journal of accounting and economics*, 36(1-3): 337–386.
Hong, Harrison, and Jeremy C Stein. 1999. "A unified theory of underreaction, momentum trading, and overreaction in asset markets." *The Journal of finance*, 54(6): 2143–2184.
Hong, Harrison, and Jeremy C Stein. 2007. "Disagreement and the stock market." *Journal of Economic perspectives*, 21(2): 109–128.
Kézdi, Gábor, and Robert J Willis. 2011. "Household stock market beliefs and learning." National Bureau of Economic Research.
Kuchler, Theresa, and Basit Zafar. 2015. "Personal experiences and expectations about aggregate outcomes."
Kyle, Albert S, and F Albert Wang. 1997. "Speculation duopoly with agreement to disagree: Can overconfidence survive the market test?" *The Journal of Finance*, 52(5): 2073–2090.
Madrian, Brigitte C, and Dennis F Shea. 2001. "The power of suggestion: Inertia in 401 (k) participation and savings behavior." *The Quarterly journal of economics*, 116(4): 1149–1187.
Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression babies: do macroeconomic experiences affect risk taking?" *The Quarterly Journal of Economics*, 126(1): 373–416.
Merton, Robert C. 1969. "Lifetime portfolio selection under uncertainty: The continuous-time case." *The review of Economics and Statistics*, 247–257.
Peng, Lin, and Wei Xiong. 2006. "Investor attention, overconfidence and category learning." *Journal of Financial Economics*, 80(3): 563–602.
Rietz, Thomas A. 1988. "The equity risk premium a solution." *Journal of monetary Economics*, 22(1): 117–131.
Scheinkman, Jose A, and Wei Xiong. 2003. "Overconfidence and speculative bubbles." *Journal of political Economy*, 111(6): 1183–1220.
Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *The American Economic Review*, 71(3): 421–436.
Vissing-Jorgensen, Annette. 2003. "Perspectives on behavioral finance: Does" irrationality" disappear with wealth? Evidence from expectations and actions." *NBER macroeconomics annual*, 18: 139–194.