American Economic Review 2021, 111(5): 1481–1522
https://doi.org/10.1257/aer.20200243

Five Facts about Beliefs and Portfolios[†]

By Stefano Giglio, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus*

We study a newly designed survey administered to a large panel of wealthy retail investors. The survey elicits beliefs that are important for macroeconomics and finance, and matches respondents with administrative data on their portfolio composition, their trading activity, and their login behavior. We establish five facts in these data. (*i*) *Beliefs are reflected in portfolio allocations. The sensitivity of portfolios to beliefs is small on average, but varies significantly with investor wealth, attention, trading frequency, and confidence.* (*ii*) *Belief changes do not predict when investors trade, but conditional on trading, they affect both the direction and the magnitude of trades.* (*iii*) *Beliefs are mostly characterized by large and persistent individual heterogeneity. Demographic characteristics explain only a small part of why some individuals are optimistic and some are pessimistic.* (*iv*) *Expected cash flow growth and expected returns are positively related, both within and across investors.* (*v*) *Expected returns and the subjective probability of rare disasters are negatively related, both within and across investors. These five facts provide useful guidance for the design of macro-finance models.* (*JEL* D83, E23, G11, G12, G41, G51)

Researchers are increasingly turning to survey data to calibrate and test macro-finance models. The unique benefit of survey data is that they can provide direct evidence on the beliefs of different agents about future economic outcomes such as stock returns and economic growth. These beliefs play a central role in both rational expectation models and behavioral models of macroeconomics and finance. Despite the potential for survey data to shed light on previously unobservable elements of macro-finance theories, their use has been criticized on many fronts. Critics

*Giglio: Yale School of Management, NBER, CEPR (email: stefano.giglio@yale.edu); Maggiori: Graduate School of Business, Stanford University, NBER, CEPR (email: maggiori@stanford.edu); Stroebel: Stern School of Business, New York University, NBER, CEPR (email: johannes.stroebel@stern.nyu.edu); Utkus: Vanguard (email: steve_utkus@vanguard.com). Emi Nakamura was the coeditor for this article. Stephen Utkus is employed at Vanguard in a research capacity. Giglio, Maggiori, and Stroebel are unpaid consultants at Vanguard in order to access the anonymized data. Vanguard provided anonymized portfolio and survey data as well as survey research services for this project. We thank Klaus Adam, Isaiah Andrews, Nick Barberis, John Campbell, Darrell Duffie, Ben Enke, John Geanakoplos, Nicola Gennaioli, Gary Gorton, Camelia Kuhnen, Arvind Krishnamurthy, Hanno Lustig, Yueran Ma, Emi Nakamura, Stefan Nagel, Frank Schilbach, Alexi Savov, Amit Seru, and Andrei Shleifer, as well as numerous seminar and conference participants. The authors would also like to thank Catherine Clinton, Sophia Bunyaraksh, and Jean Young at Vanguard for their help with the project.
[†]Go to https://doi.org/10.1257/aer.20200243 to visit the article page for additional materials and author disclosure statements.

have argued that survey data are often based on small and unrepresentative samples, that they are ridden with measurement error, that they ask qualitative questions that are not informative for models, and that they may not reveal those beliefs on which agents actually base their actions.

In this paper, we provide new evidence on the link between beliefs elicited through surveys and actions taken by survey respondents. To do this, we study a newly designed online expectations survey of a large panel of individual retail investors with substantial wealth invested in financial markets. The survey elicits the investors' beliefs about future stock returns, GDP growth, and bond returns, and was designed to address prevailing criticisms of existing survey data. The survey design trades off asking quantitative questions about moments that are crucial for macro-finance theory with keeping the questions sufficiently simple that they can be answered by nonspecialists. The survey is also short, in order to not discourage respondents from participating repeatedly over time.

The survey was administered by Vanguard, one of the world's largest asset management firms, to an almost random sample of its US-based clients. About 80 percent of the investors in the sample have retail trading accounts at Vanguard, while the remaining 20 percent have employer-sponsored retirement accounts. The respondents are individuals relevant for macro-finance models: they participate in financial markets and they have substantial wealth, with the average respondent holding more than half a million dollars of assets at Vanguard. The survey has been conducted every two months since February 2017. In this paper, we study the first 21 survey waves, covering the period February 2017 to June 2020, which generated a total of 46,419 responses. Many individuals responded to multiple survey waves, providing a substantial panel component to the data. We link survey responses to anonymized administrative data on the respondents' investment holdings and transactions at Vanguard. This allows us to explore the relationship between the elicited beliefs and real-world high-stakes investment behavior. Since Vanguard clients are potentially more likely to be buy-and-hold investors, whenever possible we confirm that the patterns in our survey data are consistent with the corresponding patterns in other surveys covering different investor populations.

Our most general finding is that survey data are highly informative about individuals' portfolio decisions. Specifically, we find a robust relationship between beliefs and portfolio allocations, both across individuals and within individuals over time. In this sense, we conclude that survey-based evidence is "here to stay," and that theoretical work has to continue to confront such evidence. We organize our findings around five facts that highlight empirical patterns about beliefs as well as their relationships with portfolios. We believe that these facts, which are robust to including or excluding the period around the March 2020 stock market decline induced by the COVID-19 crisis, can guide future empirical and theoretical work in macro-finance.

Fact 1 summarizes our main findings on the relationship between beliefs and portfolios. We first document a statistically strong relationship between beliefs and portfolio composition. However, the average sensitivity of an investor's equity share to that investor's subjective expected stock return is substantially lower than predicted by benchmark frictionless macro-finance models. More specifically, we find that a 1 percentage point increase in expected stock returns over the next year is associated with a 0.7 percentage point higher equity share, which is one order of magnitude smaller than implied by standard calibrations of the frictionless Merton (1969) model. We rule out that this relatively low magnitude is primarily the result of attenuation bias from classical measurement error in beliefs. We also 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 subjective probability of a large stock market drop (a rare disaster).

This relatively small response of equity shares to beliefs about stock returns is consistent with evidence documented across several other studies that link retail investors' equity market participation and equity shares to expected stock returns (e.g., Vissing-Jorgensen 2003, Dominitz and Manski 2007, Kézdi and Willis 2011, Amromin and Sharpe 2014, Ameriks et al. 2016). Our contribution to this literature is twofold. First, we use administrative data to confirm this fact for a large sample of wealthy investors, while accounting for key dimensions of measurement error. Second, we show that investors are heterogeneous in their sensitivity along several economically interesting dimensions. The sensitivity of portfolios to beliefs is increasing in wealth; it is also higher in tax-advantaged retail accounts, and increasing in investors' trading frequency, investors' attention to their portfolios, and investors' confidence in their own beliefs. We find that an idealized investor who holds a tax-advantaged retail account, pays attention to her portfolio, trades often, and is confident in her beliefs has a sensitivity that is about three times larger than the average sensitivity, and gets close to the sensitivities generated by frictionless benchmark models.

We next investigate the role of belief changes in explaining trading activity. Fact 2 establishes that an individual's belief changes have little or no explanatory power for predicting when trading occurs (the extensive margin), but help explain both the direction and magnitude of trading conditional on a trade occurring (the intensive margin). Our findings are thus consistent with models of infrequent trading that generate a flat hazard function of trading based on belief changes, with heterogeneous trading probabilities across people.

These first two facts are informative about a central element of both rational and behavioral macro-finance models: the transmission channel from beliefs to portfolio choices. Our results show that for the average investor, this pass-through is positive but weak, which might dampen the effects of belief changes on equilibrium prices and quantities in theoretical models. This cautions against commonly used calibrations of many representative agent behavioral models that use variation in survey-based beliefs to explain asset price movements. Specifically, the variation in asset demand implied by these models, many of which are based on Mertonian portfolio demand, is far too high relative to the pass-through from beliefs to portfolios observed in the data. At the same time, we show that this pass-through is heterogeneous across investors along economically interesting dimensions. Incorporating this heterogeneity into macro-finance models should help these models to jointly match data on beliefs, quantities, and asset prices.

We next decompose the variation in beliefs across individuals and over time. Fact 3 establishes that individual beliefs are mostly characterized by heterogeneous and persistent individual fixed effects: some individuals are optimistic and some are pessimistic, and their beliefs are persistent and far apart. While there is some comovement in beliefs across individuals over time, the time variation in average beliefs only accounts for about 5 percent of the total variation in beliefs in the panel. Instead, between 40 percent and 60 percent of all panel variation in beliefs is captured by individual fixed effects, while the rest is due to idiosyncratic individual variation and measurement error. We also find that the heterogeneity in beliefs is not well explained by observable respondent characteristics such as gender, age, wealth, attention, confidence, past returns, and geographic location. These characteristics sometimes have strong statistical relationships with beliefs, but their joint explanatory power is limited. We provide evidence that this is not the result of measurement error in eliciting beliefs. Instead, a likely explanation is that individual beliefs reflect a combination of many demographic characteristics and experiences, without a single dominant explanation.

Fact 3 provides a simple but powerful description of the panel variation of investor beliefs: investors disagree strongly and persistently about expected cash flows and returns. This contrasts with much of the existing literature that builds on survey evidence on beliefs, which has focused on representative agent models disciplined by matching the time-series behavior of average beliefs. This approach misses a more prominent feature of the data, the persistent individual heterogeneity. As a result, models that explicitly feature heterogeneous agents with different beliefs are likely to offer a fruitful starting point for future work. Indeed, incorporating persistent belief heterogeneity into macro-finance models is not only a way to better match the survey evidence, but might also allow for interesting model dynamics, for example coming from the redistribution of wealth between optimists and pessimists as shocks are realized.

We next explore how beliefs about different objects correlate. Fact 4 establishes that higher expectations of GDP growth are associated with higher expectations of future stock returns, at both short and long horizons. In the cross-section, investors who expect higher cash flow growth also tend to expect higher returns. In the time series, when an investor becomes more optimistic about cash flow growth, she also becomes more optimistic about expected returns. The correlation between expected returns and cash flow growth is an informative moment for macro-finance analysis. Indeed, the Campbell and Shiller (1988) decomposition shows that expected cash-flow growth and expected returns have opposite effects on current valuations. Models that specify belief dynamics for one process (either cash flows or returns) imply equilibrium beliefs about the other process, and our work can be used to verify whether the resulting joint distribution of beliefs is qualitatively and quantitatively consistent with the data.

Fact 5 establishes that when individuals perceive large stock market declines to be more likely, they also expect stock returns to be lower. This relationship holds both across individuals and within individuals over time. This finding relates to an important strand of the macro-finance literature, which has emphasized that expectations of rare but potentially catastrophic events, so-called rare disasters, can help explain portfolio holdings and asset prices (see Rietz 1988, Barro 2006, Gabaix 2012). Our results are consistent with versions of the rare disaster framework, like the model by Chen, Joslin, and Tran (2012), that allow for heterogeneous beliefs about disaster probabilities and a willingness of investors to "agree to disagree."

We conclude this introduction by summarizing the desired characteristics of a model that would be consistent with our five facts. This represents just one way one could write such a model, based on our evidence, rather than the only possible model. The proposed model would have three key ingredients: (i) large and highly persistent heterogeneity in beliefs about both expected returns and cash flows, with the two beliefs positively related; (ii) a willingness to "agree to disagree" that allows for trading based on disagreement; and (iii) infrequent trading with an exogenous probability of trading that differs across agents. It is an interesting open question how well such a model would perform in quantitatively matching aggregate asset prices in addition to the main features of beliefs and portfolios documented in this paper.

Related Literature.—Our paper contributes to a growing literature that focuses on exploring the role of beliefs in explaining a large number of economic outcomes (see DellaVigna 2009, Benjamin 2019, for a review). In this literature, Manski (2004) was among the first and most prominent to argue for using survey data about expected equity returns and risks to better understand individuals' investment behaviors. Over time, a series of papers has connected survey expectations to the behavior of respondents. For example, Ameriks et al. (2015a, b, 2016, 2017, 2018) have provided recent advances by linking survey evidence to retirement choices. As part of this agenda, Ameriks et al. (2016) find a low sensitivity of retail investors' equity investment to stock market expectations, a fact also documented by Vissing-Jorgensen (2003); Dominitz and Manski (2007); Kézdi and Willis (2009, 2011); Amromin and Sharpe (2014); Arrondel, Calvo Pardo, and Tas (2014); Merkle and Weber (2014); Choi and Robertson (2020); and Drerup, Enke, and von Gaudecker (2017), with related work by Hurd, Van Rooij, and Winter (2011); Hudomiet, Kézdi, and Willis (2011); and De Marco, Macchiavelli, and Valchev (2018).[1] Our work builds on this literature by exploring a quantitative survey of a large panel of wealthy investors, which is matched to administrative data on these investors' portfolios and trading behaviors. Our survey, which was designed to inform theoretical models, allows us to discover new facts and deepen our understanding of existing patterns, both quantitatively and in terms of their variation across individuals and over time.

We also contribute to a macro-finance literature that has debated the trade-offs between behavioral and rational modeling approaches, with survey evidence providing an important input (see Cochrane 2011, 2017; Greenwood and Shleifer 2014; Adam, Matveev, and Nagel 2018). Among the many proposed equilibrium models, the most relevant for our work are those that directly incorporate survey evidence (e.g., Barberis et al. 2015; Adam, Marcet, and Beutel 2017; Bhandari, Borovička, and Ho 2016) and those that feature heterogeneous belief (e.g., Scheinkman and Xiong 2003, Geanakoplos 2010, Caballero and Simsek 2017, Martin and Papadimitriou 2019).

Our work also relates to a literature that has explored the role of beliefs in other settings. For example, in the housing market (e.g., Piazzesi and Schneider 2009; Case, Shiller, and Thompson 2012; Cheng, Raina, and Xiong 2014; Kuchler and Zafar 2019; Burnside, Eichenbaum, and Rebelo 2016; Glaeser and Nathanson 2017; Adelino, Schoar, and Severino 2018; Bailey et al. 2018; and Bailey et al. 2019) as well as the role of firm expectations (e.g., Cummins, Hassett, and Oliner 2006; Bacchetta, Mertens, and van Wincoop 2009; Coibion and Gorodnichenko 2012;

[1] A related literature has used surveys to explore the relationship between risk aversion and portfolio holdings (e.g., Guiso, Sapienza, and Zingales 2018).

Gennaioli, Ma, and Shleifer 2016; Landier, Ma, and Thesmar 2017; Bachmann et al. 2018; Bordalo et al. 2018; García-Schmidt and Woodford 2019; Fuhrer 2018; Bailey et al. 2018; and Bailey et al. 2019). A further related literature has explored how individuals with different political convictions respond to political events (see Mian, Sufi, and Khoshkhou 2015; Kempf and Tsoutsoura 2018; Meeuwis et al. 2018).

I. Survey Description

To explore the structure of investors' beliefs and their relationship with investors' portfolio allocations, we study a new online survey of US-based individual investors who hold accounts at Vanguard, one of the world's largest asset management firms with more than $6 trillion in assets under management. We refer to this survey as the GMSU-Vanguard survey in the rest of this paper. We first provide a high-level overview of the survey questions (we report the exact phrasing and the survey interface in the online Appendix). We then explore the survey sample, the response rates, and the demographics of respondents and nonrespondents, which allow us to analyze the dimensions of selection into responding.

A. *Survey Design*

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

Expected Stock Market Returns.—The survey asks respondents about their expectations for the return of the US stock market. It elicits point estimates for the expected annualized returns over the coming year and the coming ten years. It also elicits subjective probabilities that the return over the next year would fall into one of five buckets: less than -30 percent, between -30 percent and -10 percent, between -10 percent and 30 percent, between 30 percent and 40 percent, and more than 40 percent.[2] The ordering of the buckets (i.e., lowest to highest or highest to lowest) is randomized across survey respondents, and the survey enforces that the assigned probabilities add up to 100 percent. As shown in the online Appendix, the survey interface also presents real-time histograms of the survey responses as they are entered, which helps individuals visualize the probability distributions implied by their numerical answers.

Expected Real GDP Growth Rates.—The survey also asks respondents about their expectations for the annualized growth rate of real GDP. It elicits point estimates for the expected growth rates over the coming three years and the coming ten

[2] These buckets were chosen such that the tails correspond to extreme events that still have substantial probability mass based on historical frequency. Between 1927 and 2014, the share of 1-year stock returns in each bucket was: 3.7 percent (lowest bucket), 11.6 percent (second bucket), 65.9 percent (third bucket), 13.8 percent (fourth bucket), and 5 percent (highest bucket).

years. It also elicits subjective probabilities that the annualized GDP growth rate over the coming three years would fall into one of five buckets: less than −3 percent, between −3 percent and 0 percent, between 0 percent and 3 percent, between 3 percent and 9 percent, and more than 9 percent.[3] The ordering of buckets (i.e., lowest to highest or highest to lowest) is randomized across respondents.

Expected Bond Returns.—The final block of questions elicits respondents' expectations about the 1-year return of a 10-year US government zero coupon bond.[4]

Difficulty and Confidence.—At the end of every block of questions (i.e., about expected stock returns, expected GDP growth, and expected bond returns), the survey asks individuals how confident they are about their answers (on a five-point scale from "not at all confident" to "extremely confident"), and how difficult they found the questions (on a five-point scale from "not at all difficult" to "extremely difficult").

B. *Survey Sample and Response Rate*

The online survey is conducted every two months by Vanguard among its US-based individual customers. In this paper, we explore the first 21 waves of the survey, covering February 2017 to June 2020. In the first wave, 40,000 clients were invited by email to participate in the survey. These clients were randomly selected such that 80 percent of the sample were retail investors and 20 percent were investors in defined contribution plans.[5] Overall, the sample of individuals who are potentially contacted represents about $2 trillion in assets. If individuals respond to the survey in any wave, they are recontacted in each subsequent wave. Individuals who do not respond to the first wave in which they are contacted are recontacted for two subsequent waves. If they respond in neither of these waves, they are dropped from the sample. Individuals can opt out of the study at any point and are, in this case, never contacted again. 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 onward, the number of new clients contacted in each wave increased to 14,500.

Response Rates.—The left panel of Figure 1 shows the response rates for the first 21 waves, where we count only fully completed surveys as "responses." The orange-circle line shows that the response rates among individuals contacted for the first time were relatively stable at 2.5–4 percent 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 percent and 2 percent across waves. The

[3] Between 1929 and 2015, the share of annualized US real three-year GDP growth in each bucket was: 4.2 percent (lowest bucket), 5.4 percent (second bucket), 33.3 percent (third bucket), 50.8 percent (fourth bucket), and 6.3 percent (highest bucket).

[4] A last question also elicits expectations about future yield curves; we do not use this question in the current paper.

[5] Additional requirements to be potentially invited to take the survey are that clients (i) had opted into receiving account statements via email, (ii) were older than 21 years, and (iii) had total Vanguard assets of at least $10,000.

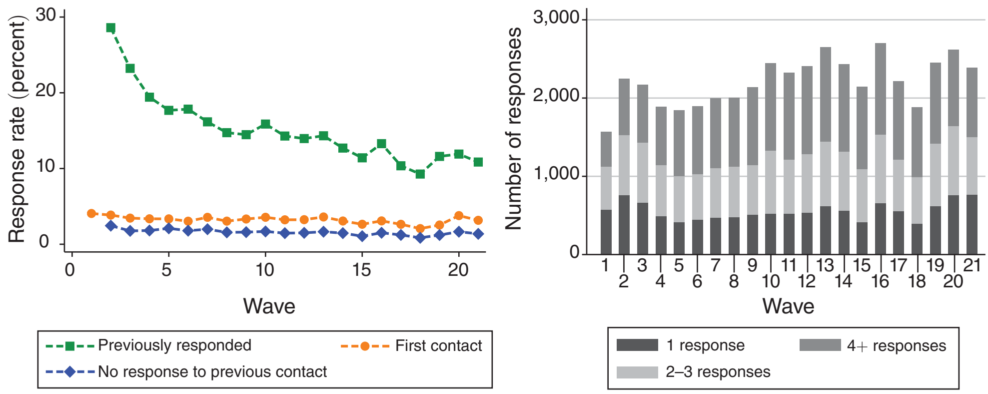

Figure 1. Survey Responses

Notes: Figure shows the responses to the GMSU-Vanguard Survey in each of the 21 waves between February 2017 and June 2020. The left panel shows response rates. The orange line (circles) shows the response rates for individuals contacted for the first time. The blue line (diamonds) shows the response rates for individuals who were contacted in previous waves, but who had not yet responded. The green line (squares) shows the response rates for individuals who had previously responded. The right panel shows the number of responses per wave. It splits out responses that come from individuals who only respond to one of the waves, from individuals who respond to two or three waves, and from individuals who respond to at least four waves.

green-square line shows the response rates among individuals who had responded to at least one previous survey. The steady-state re-response rate of these individuals is between 10 percent and 15 percent. It declines somewhat over time, though much of this decline is driven by compositional effects: in later survey waves, the average time since the last response of individuals who have previously responded is higher. These response rates translate into more than 2,000 survey responses on average per wave. Across the 21 waves, we received 46,419 total responses.

The right panel of Figure 1 shows the number of responses in each wave, split out by how many overall survey waves the respondents participated in. Overall, about 25 percent of responses come from individuals who have responded to one survey only (though some of these may end up responding to future survey waves). Over 40 percent of responses come from individuals who have responded to at least four survey waves, and more than 25 percent come from individuals who have responded to at least six waves. Online Appendix Section A.1 provides additional details on our response rates.

Demographics of Respondents.—Table 1 presents summary statistics on the demographics, portfolio composition, and trading behavior of the survey respondents as well as the nonrespondents. The average survey respondent is about 60 years old, while 69 percent of respondents are male. Respondents hold assets with an average value of \$520,000 at Vanguard, while the 10–90 percentile range of assets held at Vanguard is \$23,000 to \$1.3 million. The average respondent logs into her Vanguard account about 3.7 times a month, has 1.5 active trades per month, and turns over 2.3 percent of her portfolio every month. Activity on the Vanguard site, both in terms of logins and in terms of trading activity, varies across survey respondents. At the

Table 1—Demographics: Survey Respondents and Nonrespondents

	Survey respondents				Non-respondents	
	Mean	P10	P50	P90	Mean	Difference
Age (years)	60.1	39.0	63.0	76.0	52.3	7.86
Male	0.69	0	1	1	0.54	0.16
Region						
Northeast	0.23	0	0	1	0.24	−0.01
Midwest	0.21	0	0	1	0.20	0.01
South	0.31	0	0	1	0.31	0.00
West	0.25	0	0	1	0.25	0.00
Total Vanguard wealth (k$)	520.0	23.2	227.8	1301.1	254.5	265.4
Length of Vanguard relationship (years)	17.11	5	17	29	14.32	2.78
Active trades/month	1.54	0.02	0.55	3.71	0.91	0.63
Monthly portfolio turnover (percent)	2.30	0.00	0.91	4.85	1.96	0.34
Days with log-ins/month	3.71	0.09	1.53	10.96	1.81	1.90
Total time spent/month (minutes)	31.0	1.1	11.9	72.5	12.3	18.72
Portfolio shares (percent)						
Equity	67.5	30.8	70.7	99.9	71.5	−4.02
Fixed income	20.9	0.0	15.8	48.4	17.3	3.54
Cash	10.1	0.0	1.7	31.4	10.1	−0.01
Other/unknown	1.6	0.0	0.0	4.2	1.1	0.48
Number of unique assets	7.55	1	4	17	4.65	2.90
Number of mutual funds	4.82	1	3	11	3.31	1.51
Number of ETFs	0.81	0	0	2	0.43	0.38
Number of stocks	1.80	0	0	4	0.87	0.92
Number of bonds	0.13	0	0	0	0.04	0.09

Notes: Table shows summary statistics on both the survey respondents and nonrespondents. Age, gender, location, total wealth at Vanguard, length of Vanguard relationship, and number of assets are measured as of June 2019. Other variables are presented as monthly averages between January 2017 and June 2019.

tenth percentile of the distribution, respondents spend essentially no time on the Vanguard site, while at the ninetieth percentile of the distribution, respondents log in every third day. The average respondent has 7.6 unique assets in her portfolio, and holds 67.5 percent of her portfolio in equity, 20.9 percent in fixed income assets, and 10.1 percent in cash.[6] There are substantial differences in portfolio allocations across our survey respondents. At the tenth percentile of the distribution, the equity share is 30.8 percent, while at the ninetieth percentile, it is 99.9 percent.

Sample Selection.—Like in all surveys, the sample of respondents is likely selected on a number of dimensions. We explore two types of selection: (i) selection into who is a Vanguard client, and (ii) selection into which Vanguard clients answer the survey.

We first compare the characteristics of respondents and nonrespondents. Table 1 shows that survey respondents are older and more likely to be male than nonrespondents. Respondents are also substantially wealthier, with average wealth held at Vanguard of $520,000 for respondents relative to $254,000 for

[6] The CUSIP-level information on individual security holdings is matched with information from Morningstar for mutual funds to calculate the portfolio share held in equities, fixed income instruments, cash, and other investments. Investments in mutual funds are apportioned depending on the portfolio composition of each fund (e.g., 60 percent equity and 40 percent fixed income). Cash includes cash-equivalent investments such as money-market funds. The category "other investments" includes alternative investments such as commodities, real estate, and derivatives.

nonrespondents. Respondents trade more frequently and their monthly portfolio turnover is larger; they also log into their Vanguard accounts more frequently than nonrespondents.[7] Portfolio allocations of respondents and nonrespondents are relatively similar, though the average respondent holds more unique assets. Overall, our sample over-represents individuals who are wealthier and who trade more often, and whose beliefs are thus more likely to affect asset prices.

We also analyze whether Vanguard clients are similar to the overall pool of retail investors. One concern is that Vanguard's investment philosophy of focusing on passive and low-fee investments attracts a selected sample of investors. Indeed, there are some differences between Vanguard clients and other retail investors. For example, Cogent Wealth Reports (2018) compared Vanguard retail clients to a representative sample of investors with at least \$100,000 in investable assets. Vanguard clients were more likely to be older and richer than the comparison sample; they also held a larger portfolio share in risky assets and a larger share in passive-like instruments. These differences, however, do not mean that our sample is uninteresting or not quantitatively relevant. Vanguard manages more than \$6 trillion in investments (with the potential survey respondents holding around \$2 trillion), so the investors targeted in our study own a nontrivial fraction of global investable wealth. With the rising popularity of low-fee investment strategies, our sample is also likely to become even more relevant to understanding investments and asset prices. In addition, while there are some notable differences in characteristics between Vanguard investors and retail investors more generally, the two investor groups are similar on other important dimensions.

First, we compare the trading intensity of the survey respondents to the trading intensity observed in other investor samples studied in the literature. A recent paper by Meeuwis et al. (2018) analyzes typical retirement investors at a "large US financial institution" that is not Vanguard; they report that 29.5 percent of those investors made an active trade in the past year.[8] In our sample, which comprises both retail and retirement accounts, 67 percent of non-respondents and 78 percent of respondents made at least one active trade in the year 2017. We conclude that despite Vanguard's focus on passive buy-and-hold strategies, individuals in our sample do not appear to trade less frequently than representative retail investors at other firms.[9]

Second, we explore whether Vanguard investors are less likely than other investors to follow a "flow-performance" pattern, whereby they increase their positions in mutual funds that just experienced high returns (see Coval and Stafford 2007). Online Appendix Section A.2 shows that the flow-performance sensitivity in the population of all investors in Vanguard funds as well as in the population of

[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 logins.

[8] Meeuwis et al. (2018) report that their data contains the characteristics and individual portfolio holdings of millions of anonymized households covering trillions of dollars in investable wealth. Their research focuses on a subsample that is representative of "typical" American investors with retirement saving, a group that holds 41 percent of household investable wealth in the United States.

[9] On the other hand, we would expect that some of our results might not generalize to a broader set of investors. For example, clients at brokerage firms that target day-traders or professional hedge funds are likely to have higher portfolio turnover. Indeed, Barber and Odean (2000) document the trading behavior of about 80,000 investors at a large discount brokerage firm in the United States, and find that the median person in their sample has a portfolio turnover of about 2.7 percent per month, relative to 0.91 percent in our sample.

investors who respond to our survey is very similar to the flow-performance sensitivity documented in the literature.

Third, online Appendix Section A.3 shows that the level and time-series variation in our survey responses are similar to those from other surveys that target different investor populations during the same period. Ultimately, it is a worthy pursuit for the future literature to extend our results to other contexts and investor populations to explore possible sources of interesting heterogeneity.

C. *Survey Responses: Summary Statistics*

Table 2 shows summary statistics across the 46,419 survey responses; online Appendix Section A.4 presents the full distribution of the responses to the key questions.[10] The average respondent takes about 8.5 minutes to answer all survey questions. The 10–90 percentile range for the total time to respond is 3.9 minutes to 13.5 minutes. Therefore, all respondents spend a sizable amount of time answering the questions, rather than carelessly clicking through the survey; this is consistent with the noncompensated nature of the survey requiring a certain intrinsic interest from participants.

The average expected 1-year stock market return is 4.64 percent, while the average annualized expected 10-year stock market return is 6.64 percent. There is substantial heterogeneity in the expected 1-year stock market return across responses. At the tenth percentile of the distribution, individuals reported a 1-year expected stock return of -1 percent, while at the ninetieth percentile, they expected a return of 10 percent. The across-responses standard deviation of expected 1-year returns is 6.08 percent, much larger than the standard deviation of expectations of annualized 10-year returns, which is 3.85 percent. This suggests that individuals anticipate some medium-run mean reversion in stock returns. When we ask individuals about their expectations of annualized GDP growth, the means for the next three years and the next ten years are quite similar at 2.77 percent and 3.15 percent, respectively. The average respondent expected 1-year returns of 10-year US Treasury zero coupon bonds to be 1.74 percent, with a 10–90 percentile range across respondents of -1 percent to 4 percent.

Table 2 also shows that respondents put substantial probabilities on relatively large short-run stock market declines and GDP declines. The average individual assigns a 5.3 percent chance to the 1-year return of the stock market being less than -30 percent, while the median respondent assigns a 3 percent chance to such an event. As with the other answers, there is substantial across-answer heterogeneity.

[10] Even among the respondents who completed the entire survey flow, some respondents skipped a few questions. We verified that restricting our analysis to the sample of respondents who provided answers to every question does not affect our conclusions. We also explored the presence of extreme outlier responses, such as individuals reporting that the expected return on the US stock market over the coming year was 400 percent or -100 percent. Since such outliers have extreme effects on the analysis, in our baseline analysis we set extreme outlier answers (below the bottom percentile, and above the top percentile) for each unbounded expectation question equal 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 with selecting these cutoffs, which involve trading off retaining true extreme beliefs with excluding answers from individuals who probably misunderstood the question or the units. We have done extensive sensitivity analysis 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 missing, and to dropping all answers of individuals who report extreme answers to at least one question.

Table 2—Summary Statistics: Survey Responses

	Mean	SD	P10	P25	P50	P75	P90
Expected stock returns							
Expected 1Y stock return (percent)	4.64	6.08	−1	3	5	8	10
Expected 10Y stock return (percent p.a.)	6.64	3.85	3	5	6	8	10
Probability 1Y stock return in bucket (percent)							
Less than −30 percent	5.3	8.5	0	0	3	8	10
−30 percent to −10 percent	14.5	14.1	0	5	10	20	30
−10 percent to 30 percent	70.1	23.0	40	60	75	90	100
30 percent to 40 percent	7.3	10.6	0	0	5	10	20
More than 40 percent	2.7	6.3	0	0	0	5	10
Expected GDP growth							
Expected 3Y GDP growth (percent p.a.)	2.77	2.16	1	2	3	3	4
Expected 10Y GDP growth (percent p.a.)	3.15	2.74	2	2	3	3	5
Probability p.a. 3Y GDP growth in bucket (percent)							
Less than −3 percent	4.9	8.5	0	0	2	5	10
−3 percent to 0 percent	13.3	13.0	0	5	10	20	30
0 percent to 3 percent	58.7	25.8	20	40	60	80	90
3 percent to 9 percent	20.0	21.9	0	5	10	25	50
More than 9 percent	3.1	8.5	0	0	0	5	10
Expected bond returns							
Expected 1Y return of 10Y zero coupon bond (percent)	1.74	2.84	−1	1	2	3	4
Difficulty (“*not at all difficult*,” …, “*very difficult*”)							
Expected stock returns	2.35	0.98	1	2	2	3	4
Expected GDP growth	2.46	0.98	1	2	2	3	4
Expected bond returns	2.84	1.00	1	2	3	3	4
Confidence (“*not at all confident*,” …, “*very confident*”)							
Expected stock returns	3.04	0.85	2	3	3	4	4
Expected GDP growth	2.98	0.85	2	3	3	3	4
Expected bond returns	2.62	0.86	2	2	3	3	4
Time of responding to survey (seconds)	507	429	235	302	407	567	810

Notes: Table shows summary statistics of the answers across responses to the first 21 waves of the GMSU-Vanguard survey. 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.”

Answers at the twenty-fifth percentile of the distribution correspond to a 0 percent chance of returns lower than −30 percent, while those at the ninetieth percentile of the distribution correspond to a 10 percent probability of such events. Similarly, in the case of GDP growth, individuals assign an average probability of an annualized decline in GDP of more than 3 percent over the coming three years of 4.9 percent.

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 and 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 tenth percentile of the distribution report being “not very confident” in their answers, while individuals at the ninetieth percentile reported being “very confident.”

II. Beliefs and Portfolios

In this section, we analyze the relationship between respondents' beliefs and their portfolio allocations. In the main body of the paper, we explore the role of expectations about 1-year stock returns in determining equity shares. In online Appendix Section A.6, we analyze the role of other moments of the belief distribution (e.g., beliefs about the probability of large stock market declines), the role of stock market expectations over longer horizons, and the role of beliefs about bond returns and GDP growth. To estimate the sensitivity of equity shares to beliefs, we run the following regression:

$$EquityShare_{i,t} \;=\; \alpha + \beta E_{i,t}\big[R_{1y}\big] \;+\; \gamma X_{i,t} + \psi_t + \epsilon_{i,t}. \tag{1}$$

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. Since most Vanguard investors find it hard to short-sell or obtain leverage, the equity share is essentially censored at both 0 percent and 100 percent. We thus estimate regression (1) using Tobit models. The coefficient β captures the increase in an individual's equity share for each percentage point increase in the expected 1-year stock return.

Column 1 of Table 3 shows estimates from this regression without controlling for any additional covariates. An extra percentage point of expected 1-year stock returns is associated with a 0.67 percentage point increase in respondents' equity shares.[11] In column 2 of Table 3, we control for demographic characteristics such as age, gender, wealth, and region of residence, as well as survey-wave fixed effects. Figure 2 shows a conditional binscatter plot of the resulting relationship between expected returns and equity shares. The estimated sensitivity of portfolio shares to beliefs is similar to that in column 1.[12] Given the wide heterogeneity in beliefs across individuals, our estimates imply substantial belief-driven variation in equity shares: quantitatively, a 1 standard deviation increase in expected 1-year stock returns is associated with a 0.16 standard deviation increase in equity share.[13]

Figure 2 suggests that the estimated relationship between beliefs and portfolios might be sensitive to beliefs at the two extremes. Therefore, we next run regression (1) on a sample of respondents that report expected returns between 0 percent and 15 percent. This drops about 10 percent of our responses. Column 3 of Table 3

[11] Online Appendix Section A.6 shows that about one-half of the increase in equity shares of individuals who expect higher stock market returns comes from substituting away from cash rather than substituting away from fixed income securities.

[12] While the estimates of β are the primary object of interest, the coefficients on 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, and are only weakly declining with wealth.

[13] In online Appendix Section A.6 we explore the effect of other moments of the belief distribution on equity shares. We find that the subjective risk of large stock market declines has a more significant effect on portfolio allocations than the subjective variance. We also show that long-run stock market beliefs matter in addition to short-run beliefs, and highlight that beliefs about other investments, including fixed income investments, also influence the optimal equity share. Expected GDP growth does not appear to significantly affect the equity share once we control for expected stock returns.

Table 3—Expected Returns and Portfolio Equity Shares

	Equity share (percent)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Expected 1Y stock return (percent)	0.672 (0.034)	0.690 (0.034)	1.164 (0.061)	0.634 (0.053)	0.818 (0.042)	1.177 (0.062)	1.188 (0.069)	1.146 (0.071)
Expected 1Y stock return (percent) × assets > \$225k				0.114 (0.067)				
Expected 1Y stock return (percent) × above median time							−0.021 (0.056)	
Expected 1Y stock return (percent) × closest prior trade 2 weeks before								0.479 (0.356)
Expected 1Y stock return (percent) × closest prior trade 1 week before								0.460 (0.220)
Expected 1Y stock return (percent) × closest next trade 1 week after								0.384 (0.245)
Expected 1Y stock return (percent) × closest next trade 2 weeks after								0.188 (0.266)
Controls	N	Y	Y	Y	Y	Y	Y	Y
ORIV	N	N	N	N	N	Y	Y	Y
Sample			*E*(return) 0–15 percent		Feb 2017–Feb 2020			Retail accounts
Observations	44,595	44,565	39,296	44,565	39,859	44,235	44,235	41,142

Notes: Table shows results from regression (1). The unit of observation is a survey response. The dependent variable is the equity share. Columns 2–8 control for the respondents' age, gender, region of residence, wealth, and the survey wave. For the interaction specifications in columns 4, 7, and 8, we also include dummy variables for the respondent characteristics that we estimate the sensitivity for. Standard errors are clustered at the respondent level.

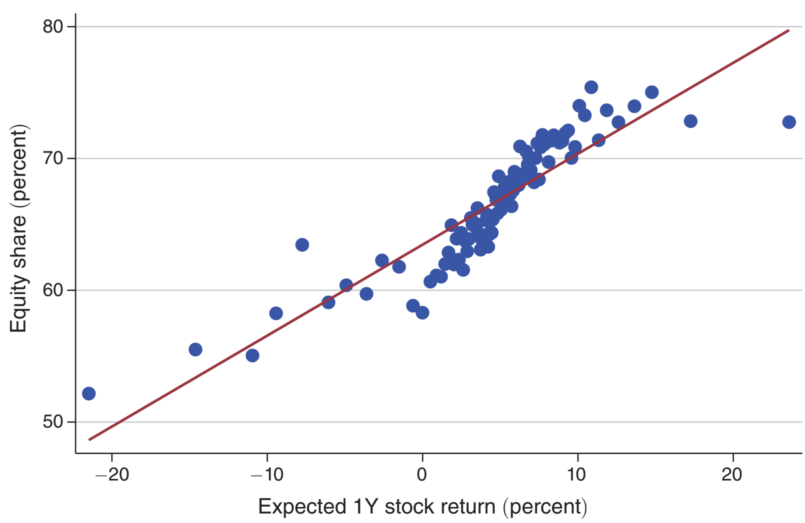

Figure 2. Expected 1-Year Stock Returns and Equity Share

Note: Figure shows a conditional binscatter 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.

shows that the sensitivity estimate in this restricted specification is almost 70 percent higher, suggesting that outliers have some effect on the relationship.

To highlight the economic magnitude of the estimated β-coefficient, we perform a back-of-the-envelope calculation using the Merton (1969) model, which shows that for power-utility investors:

$$EquityShare_{i,t} = \frac{1}{\gamma}\frac{E_i[R] - R_f}{\mathrm{var}_i[R]}. \tag{2}$$

Here, 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 $\mathrm{var}_i[R]$ is the individual's subjective variance of equity returns. We measure $EquityShare_{i,t}$ in the Vanguard data, and $E_i[R]$ with the survey answer to 1-year expected stock returns. For the back-of-the-envelope calculation, we assume that individuals have a common measure of the variance $\mathrm{var}_i[R] = \mathrm{var}[R]$.[14] Similarly, we assume a common and constant coefficient of relative risk aversion. In this simplified setting, the β estimated in Table 3 corresponds to $\beta = 1/(\gamma \mathrm{var}[R])$. The historical standard deviation of stock market returns is around 20 percent a year. The simple model thus implies that a β of 0.69 requires a coefficient of relative risk aversion of $\gamma = 36$. This is considerably higher than most estimates in the experimental literature, which usually finds values of γ between 3 and 10. To obtain a realistic coefficient of relative risk aversion, let us say around 4, we would need an estimate of β of around 6.25; for $\gamma = 6$, we would require $\beta = 4.2$. In other words, the sensitivity estimated in column 2 is an order of magnitude too small to align with the simplest frictionless model.[15] This relatively small response of equity shares to beliefs about stock returns is consistent with evidence documented across a number of other studies that link measures of 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 2014; Ameriks et al. 2016; Drerup, Enke, and Von Gaudecker 2017).

In many models, asset prices are driven by wealth-weighted beliefs, rather than beliefs that are equally weighted across all investors. We thus explore whether the sensitivity of portfolios to beliefs is different for wealthier individuals. Column 4 of Table 3 shows that respondents with more than \$225,000 in assets, which corresponds to our sample median, have a sensitivity that is marginally larger than that of individuals with lower wealth; in unreported results, we find that the sensitivity does not increase substantially for even higher levels of wealth. These results show

[14] In online Appendix Section A.6, we relax this assumption and measure the variance of 1-year expected stock returns that is implied by the distribution question; this does not affect the estimated effect of changes in expected equity returns on equity shares. We find the simple calculation with a common variance across individuals parameterized 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 (e.g., tail event probabilities) may also be important.

[15] Note that in this model, γ drives both the sensitivity of the equity portfolio share to changes in expected returns as well as the unconditional level of the equity share. In particular, for $\gamma = 4$, an average risk premium of 6 percent, and a standard deviation of 20 percent, we obtain an average equity share of 38 percent. When $\gamma = 30$, the equity share drops to 5 percent. This means that high risk aversion can explain the low sensitivity of portfolios to beliefs, but at the cost of grossly missing the average level of the equity portfolio share.

that even the wealth-weighted sensitivity would not be large enough to generate the quantity movements implied by Mertonian portfolio demand.

One natural question concerns the extent to which our findings are driven by the large drop and subsequent recovery of the stock market resulting from the COVID-19 pandemic in the first half of 2020.[16] As we show below, following the stock market crash of March 2020, average beliefs in the April 2020 wave fell dramatically, and had only partially recovered by the June 2020 wave. In column 5 of Table 3, we restrict our analysis to the first 19 waves of the GMSU-Vanguard survey, ending near the all-time high of the S&P 500 in February 2020. The estimated sensitivity of portfolios to beliefs over this period, which generally saw rising stock prices, is somewhat higher, with a $\beta \approx 0.82$. This finding suggests that the substantial decline in expected returns following the crash was associated with only a modest active reduction in equity shares (see also Giglio et al. 2020). However, even the somewhat higher elasticity during the relatively calm period in the stock market between February 2017 and February 2020 remains substantially below that implied by the frictionless model. This results highlights that the weak estimated relationship between beliefs and portfolios is not primarily the result of the large stock market crash during our sample. Indeed, online Appendix Section A.12 highlights that all results in this paper are robust to both including and excluding the COVID-19 stock market crash in the sample.

There are a number of possible explanations for this relatively low estimated average sensitivity of equity shares to expected stock returns. The first set of explanations involves measurement error in the key measure of beliefs, $E_{i,t}[R_{1y}]$, and the associated attenuation bias that such measurement error would entail. The second set of explanations centers around possible frictions in the transmission of beliefs to portfolios. Indeed, the Merton (1969) model is based on a number of strong assumptions, including that investors continuously pay attention to their portfolios, that they continuously rebalance them, that they are confident in their beliefs, and that there are no other frictions to trading, such as the tax implications from realizing capital gains.[17] Any deviation of investor behavior from these assumptions suggests that the high sensitivity in the Merton (1969) model is likely to be an upper bound for real world applications.

[16] After one of the longest and most pronounced stock market booms on record during 2009–2019, the US stock market experienced a sudden crash starting on Monday, February 24, following increasingly negative news about the COVID-19 pandemic. On March 23, the S&P 500 reached its lowest point at 34 percent below the February peak, before recovering to about 8 percent below the peak by the time of the June 2020 wave of the GMSU-Vanguard survey.

[17] Another possible set of alternative 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. Such an explanation is inconsistent with the evidence in Ameriks et al. (2016), who field a survey to show that equity shares observed in the Vanguard data are quite representative of equity shares across Vanguard users' wider range of accounts. The same study also helps mitigate the concern that more optimistic respondents might be more risk averse: that survey observes one cross-section of both expected returns and risk-aversion (elicited via lottery-type questions) and concludes that risk-aversion heterogeneity is not sufficient to explain the low estimated average sensitivity of portfolios to beliefs.

A. *Possible Explanation I: Measurement Error*

We start by exploring whether measurement error can account for the relatively low estimated average sensitivity of portfolio allocations to beliefs. A first possible explanation is that classical measurement error in beliefs may induce attenuation bias in our estimates of β.[18] To deal with such measurement error, we exploit the fact that, from each survey response, we obtain two separate estimates of the same explanatory variable, $E_{i,t}\left[R_{1y}\right]$. The first measure is the expected return as reported directly by the survey respondents. The second measure is the implied mean of the distribution over possible returns reported by each respondent.[19] The correlation across the two measures is 0.49, and the different elicitation methods likely have measurement errors that are not perfectly correlated. This setting thus allows us to exploit recent advances from the econometrics literature on instrumental variables (IV) approaches to reduce the bias from measurement error.

In principle, the attenuation bias from classical measurement error could be addressed by instrumenting for one of the estimates of expected return with the other, though there is no theoretical guidance as to which estimate should be the instrumented variable and which should be the instrument. Our approach follows the Obviously Related Instrumental Variables (ORIV) strategy proposed by Gillen, Snowberg, and Yariv (2019), which consolidates the information from these different formulations to provide an estimator that is more efficient than either of the two IV strategies alone. Column 6 of Table 3 shows that this ORIV approach increases the estimated sensitivity by more than 70 percent relative to column 2, to $\beta = 1.18$. Classical measurement error, therefore, accounts for a nontrivial component of the low sensitivity. This finding highlights the value for future surveys to include various ways of eliciting the same beliefs, thereby allowing researchers to use ORIV techniques to reduce the attenuation bias associated with measurement error that is imperfectly correlated across elicitation methods.

Nevertheless, even the sensitivity obtained by using ORIV techniques remains far below that implied by the frictionless Merton (1969) model. In this light, it is important to emphasize that if measurement error is positively correlated across the two elicitations, something that is not unlikely in our setting, then instrumented coefficients will still be biased downward, although less so than without instrumenting. We next take a number of steps to determine whether such correlated measurement error explains a substantial part of the remaining gap between our estimates and the predictions from frictionless models. In the end, we find it unlikely that correlated measurement error is an important contributor to the low estimated sensitivity. Instead, we find that frictions in the transmission of beliefs to portfolios explain much of the remaining gap between our estimates and the quantitative predictions of the frictionless benchmark model.

[18] We refer to classical measurement error as the concern that the reported belief is a noisy measure of individuals' true beliefs, $E_{i,t}\left[R_{1y}\right] = E_{i,t}\left[R_{1y}\right]^{True} + \epsilon_{i,t}$, where $\epsilon_{i,t}$ represents i.i.d. and mean zero measurement error.

[19] To construct the implied mean from the distribution, we first compute, for each bucket, the average historical return conditional on the return being in that bucket, and then we weight these estimates by the subjective probabilities of each bucket reported by the respondent. Our results are unchanged if we take the mid-points of the buckets, and assign a value of -40 percent for the lowest open-ended bucket (expected 1-year stock returns $\leq -30\%$) and a value of 50 percent for the highest open-ended bucket (expected 1-year stock returns $\geq 40\%$).

To further explore the role of measurement error, we next analyze the hypothesis that the time spent by individuals to answer the questions may allow us to identify individuals who are more or less subject to various types of measurement error. Column 7 of Table 3 shows that we obtain similar estimates of the sensitivity among people in the top and bottom halves of the time-spent distribution. Similarly, in unreported results we find that cutting out the top and bottom 10 percent of the time-spent distribution has little effect on the estimated sensitivity.

We also investigate whether beliefs that are elicited close to when people trade are less noisy and, therefore, more closely related to respondents' portfolios. This test is motivated by the model of Azeredo da Silveira and Woodford (2019), which predicts that beliefs should be most closely aligned with portfolios just before an agent trades. Column 8 of Table 3 shows that respondents who traded in the four weeks around the survey date have a higher sensitivity of beliefs to portfolios, a feature we will explore more extensively in the following section. However, the coefficients are estimated with substantial measurement error and do not allow us to determine with any precision whether individuals who traded the week before versus the week after the survey have a differential sensitivity of portfolios to beliefs.[20]

B. *Possible Explanation II: Heterogeneous Frictions*

We next show that deviations from the frictionless benchmark model of Merton (1969) can help us account for much of the remaining difference between our estimates and the predictions from that model.

Capital Gains Taxes.—A first friction that can reduce the pass-through from changes in beliefs to portfolios is the presence of capital gains taxes that can arise in the rebalancing process. To test for the importance of this friction, we exploit that some survey respondents have both standard and tax-advantaged individually managed accounts. Columns 1 and 2 of Table 4 focus on these individuals, thus controlling for potential differences in individuals' preferences such as their aversion to realize gains and losses. In column 1, we study the equity share in their standard retail accounts, while in column 2, we focus on the equity share in their tax-advantaged retail accounts (usually IRAs). We find that, for the same individuals, the equity share in the tax-advantaged accounts is more aligned with the individuals' beliefs than the equity share in the standard brokerage accounts. This evidence suggests that capital gains taxes can be an important friction that inhibits the transmission of beliefs to portfolios relative to the predictions from frictionless models.

Default Options in Defined Contribution Plans.—As described in Section I, our survey sample includes investors holding two types of tax-advantaged accounts: individually managed tax-advantaged retirement accounts and employer-sponsored retirement accounts such as defined contribution plans (though very few investors hold both types of accounts). Investments in the first account usually represent an active decision of the investor. Within defined contribution plans, it is increasingly

[20] We also do not find evidence of an increase in the probability of trading following the survey.

Table 4—Expected Returns and Portfolios: Heterogeneity

	Equity share (percent)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Expected 1Y stock return (percent)	1.198 (0.140)	1.419 (0.100)	1.270 (0.178)				
Expected 1Y stock return (percent) × monthly turnover $< 0.5\%$				0.737 (0.133)			
Expected 1Y stock return (percent) × monthly turnover $\in [0.5\%, 4\%]$				1.330 (0.084)			
Expected 1Y stock return (percent) × monthly turnover $> 4\%$				1.758 (0.231)			
Expected 1Y stock return (percent) × monthly Vanguard visits $\in (0,1)$					1.048 (0.116)		
Expected 1Y stock return (percent) × monthly Vanguard visits $\in [1,7)$					1.106 (0.072)		
Expected 1Y stock return (percent) × monthly Vanguard visits $\in [7,31)$					1.640 (0.135)		
Expected 1Y stock return (percent) × low confidence						0.810 (0.127)	
Expected 1Y stock return (percent) × medium confidence						1.229 (0.072)	
Expected 1Y stock return (percent) × high confidence						1.448 (0.268)	
Expected 1Y stock return (percent) × not idealized							1.338 (0.075)
Expected 1Y stock return (percent) × idealized							3.553 (1.329)
Controls + fixed effects	Y	Y	Y	Y	Y	Y	Y
Sample	Retail account	Retail account tax adv.	Defined contribution plans				Retail account tax adv.
Observations	19,921	19,921	4,112	41,252	44,235	44,235	34,486

Notes: Table shows results from regression (1), estimated using ORIV. The dependent variable is the equity share. In column 1 it is the equity share in individually managed retail accounts, in column 2 it is the equity share in individually managed tax-advantaged retail accounts, and in column 3 it is the equity share in institutionally managed retirement plans (defined contribution plans). In columns 5 and 6, it is pooled across the three types of accounts, while column 4 uses data from both types of retail accounts only. The sample in columns 1 and 2 is restricted to respondents holding both types of retail accounts. In column 6, *low confidence* corresponds to individuals who reported 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 7 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 percent, they have at least seven logins a month, and they report to be extremely confident in their beliefs. For the interaction specifications in columns 4, 5, and 6, we also include dummy variables for the respondent characteristics that we estimate the sensitivity for. Standard errors are clustered at the respondent level.

common to automatically enroll employees at prespecified contribution rates and into prespecified assets.[21] A robust empirical finding is that these default investments are very sticky (e.g., Madrian and Shea 2001, Beshears et al. 2009). Among Vanguard investors, Clark, Utkus, and Young (2015) found that 89 percent of participants

[21] Indeed, by the end of 2017, 46 percent of Vanguard plans had adopted automatic enrollment (about one-half of those enrolled all eligible employees, and the other half enrolled newly eligible employees only).

under automatic enrollment remained 100 percent invested in the default option after 12 months. Many investors in defined contribution plans thus appear to make relatively few active portfolio allocation decisions that could reflect their beliefs. Consistent with this, comparing columns 2 and 3 of Table 4 shows that the average sensitivity of portfolios to beliefs in institutionally managed defined contribution plans is somewhat smaller than it is in individually managed tax-advantaged accounts, though the differences are not statistically significant. This suggests a role for another deviation from the assumption of the frictionless benchmark model, namely that a substantial amount of wealth is invested through sticky default options rather than through active allocations.

Infrequent Trading.—Another plausible contributor to the low estimated sensitivity of portfolios to beliefs is that even those investors who actively choose their portfolios only adjust them infrequently (e.g., Duffie and Sun 1990; Gabaix and Laibson 2001; Agnew, Balduzzi, and Sunden 2003; Peng and Xiong 2006; Abel, Eberly, and Panageas 2007; Alvarez, Guiso, and Lippi 2012; Adam et al. 2015). To the extent that investors change their beliefs over time and report their current beliefs in the survey, the contemporaneous portfolios may thus not be fully reflective of current beliefs. Prior research has focused on several complementary explanations for infrequent portfolio adjustments. The first explanation 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 from brokerage commissions and capital gains taxes (see above), and mental costs from the need to determine optimal behavior based on beliefs and current portfolios.

We next explore whether infrequent trading can help explain the low sensitivity of portfolio allocations to beliefs. We first split respondents with retail accounts 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 trades. Column 4 of Table 4 shows that individuals with a monthly portfolio turnover of at least 4 percent have a sensitivity of equity shares to beliefs that is more than twice as large as that of individuals with a monthly portfolio turnover of less than 0.5 percent. These findings suggest that trading intensity is an important determinant of how strongly (and quickly) beliefs are reflected in portfolio holdings.

We also analyze the role of a specific motivation behind infrequent trading, namely investor attention, in explaining this relationship (see DellaVigna and Pollet 2009; Barber and Odean 2013; Ouimet and Tate 2017; and Arnold, Pelster, and Subrahmanyam 2018 for discussions of investor attention). We measure investor attention as the average frequency with which investors log into their Vanguard accounts during our sample period. Online Appendix Figure A.3 shows that this measure is correlated with, but different from actual trading activity. Column 5 of Table 4 shows that individuals who log into their Vanguard accounts more frequently have a higher sensitivity of equity shares to beliefs: individuals who log in more than seven times per month have a 56 percent higher sensitivity than individuals who log in less than once per month.[22] This result suggests that investor attention is an important driver of the attenuated relationship between beliefs and portfolios.

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 who are more confident in their own beliefs are more likely to trade on them (e.g., De Long et al. 1990; Kyle and Wang 1997; Daniel, Hirshleifer, and Subrahmanyam 1998; Odean 1999; Gervais and Odean 2001; Barber and Odean 2001; Grinblatt and Keloharju 2009; Hoffmann and Post 2016; Drerup, Enke, and von Gaudecker 2017). To explore the effect of investor confidence on the extent to which beliefs are reflected in portfolios, we exploit the fact that the survey directly elicits how confident individuals are about their answers.[23] While individuals who 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 (see online Appendix Figure A.3). 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. Column 6 of Table 4 shows that individuals who report being "extremely confident" in their stock market beliefs have an almost two times higher sensitivity of portfolio shares to beliefs than individuals who report being "not at all confident" or "not very confident."

The Idealized Frictionless Investor.—There is substantial heterogeneity across individuals in the sensitivity of portfolios to beliefs, and those individuals who are most similar to the frictionless benchmark on a number of dimensions have the highest sensitivities. In column 7 of Table 4, we explore the sensitivity of those respondents whose behavior comes closest to the frictionless model on all four dimensions jointly: individuals who are actively investing in tax-advantaged retail accounts, who are very confident in their beliefs, who pay substantial attention, and whose trading volume is significant. For respondents in that group, we estimate a β of 3.6, though the standard error around this estimate is quite large. This estimate gets quite close to the benchmark of $\beta = 6.25$ implied by the frictionless Merton (1969) model with $\gamma = 4$, and matches the predictions from that model with a $\gamma \approx 7$, a value squarely within the range considered in the asset pricing literature. This powerful result shows the importance of frictions for quantitatively explaining the deviations of observed average sensitivities from the benchmark Merton (1969) model.

[22] Note that while the frequency of trading and logins strongly affects the relationship between beliefs and portfolios, they are not strongly correlated with either the level of beliefs or the equity share themselves.

[23] Mapping confidence as reported in the survey to economic theory is not straightforward. One possibility is that confidence captures the degree of certainty that individuals have about the entire distribution of outcomes. 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 report in the survey. 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.

C. *Summary of Explanations for Low Sensitivity*, *and Their Implications*

This section documents that investors' portfolios systematically vary with their beliefs. However, the average sensitivity is smaller than that predicted by frictionless models. Some of the low estimated sensitivity can be attributed to various forms of measurement error, but this accounts for only a small part of the observed gap. Instead, we identified a number of frictions that can help explain the low average sensitivity. Overall, our findings suggest that heterogeneous investor attention, adjustment costs, capital gains taxes, and confidence are important mediators of the transmission from beliefs to portfolio allocations, and should therefore play a more prominent role in the design of macro-finance models going forward. We sum up these findings in the following fact.

Fact 1: Portfolio shares vary systematically with individuals' beliefs. However, the average sensitivity of an investor's portfolio share in equity to that investor's expected stock market returns is lower than predicted by simple and frictionless asset pricing models. This sensitivity is higher in tax-advantaged retail accounts and is increasing in wealth, investor trading frequency, investor attention, and investor confidence.

Our results on the relatively low average sensitivity of portfolios to beliefs speak to a large class of both rational and behavioral macro-finance models that explicitly account for survey evidence on beliefs. These models' predictions for asset prices usually rest on two modeling blocks: (i) beliefs that change over time in a way that is consistent with survey data, and (ii) individual portfolios that react strongly to changes in these beliefs, often by building on modifications of Mertonian portfolio demand in CARA-normal setups. In contrast to this assumption, we find that for the majority of investors in our sample, infrequent trading, inattention, and lack of confidence in beliefs reduce the pass-through from beliefs to portfolios relative to the frictionless benchmark.

It is an open question whether a model in which agents have a lower sensitivity of portfolio demand to beliefs can match asset prices without further adjustments; our research suggests that successful models should match expectations and portfolio dynamics together with asset prices. Motivated by our finding of strong heterogeneity in investor behavior, one possibility for adjustments to current models is to explicitly account for the heterogeneity in terms of wealth and sensitivity of portfolios to beliefs. For example, behavioral models such as those reviewed by Barberis (2018) often feature two types of investors, for example, behavioral investors and rational arbitrageurs, 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. In the context of these models, equilibrium prices might still be significantly driven by the behavioral investors' beliefs if these investors tend to have higher sensitivity of their portfolios to their beliefs compared to the other investors. This suggests how the models could be modified to match these new moments. Indeed, several factors can amplify the equilibrium price effects of these changes in expectations: a larger wealth share owned by the behavioral investors; a lower elasticity of the remaining investors to demand shocks, as studied for institutional investors by Gabaix and Koijen (2020); or other frictions, as in Adam et al. (2015).

III. Trading and the Pass-Through of Beliefs to Portfolios

Fact 1 highlights that low portfolio turnover reduces the measured sensitivity of portfolios to beliefs in the cross-section of survey respondents. In this section, we explore the relationship between trading activity and time-series variation in beliefs. We establish that active trades are not only infrequent, as is apparent from the summary statistics presented in Table 1, but also do not appear to be prompted by changes in beliefs. The way that belief changes translate into changes in portfolios is through the direction and magnitude of trading *conditional* on a trade occurring.

Before presenting the analysis, we briefly discuss how we measure trading; online Appendix Section A.7 provides additional details. We observe information on all transactions for clients with a retail account. These transactions include money being moved in and out of the Vanguard accounts, purchases and sales of securities, and purchases, sales, and exchanges of shares in mutual funds. We aggregate all trades by asset class: equity, fixed income, cash and cash-equivalents, and other investments. Since we observe beliefs only when an investor answers the survey, we also aggregate all trades that occur between two consecutive survey responses; these time windows differ across investors who respond to different survey waves.[24] This approach allows us to focus on changes in portfolio shares over a given time window that are induced by active trading, filtering out any changes resulting from market movements. We then regress the change in the equity share *due to trading* for individual i over time window w, given as $\Delta EquityShare_{i,w}$, on the expected 1-year stock return at the beginning of the window, $E_{i,w-}[R_{1y}]$, the change in this expectation during the window, $\Delta E_{i,w}[R_{1y}]$, and the equity share at the beginning of the window, $EquityShare_{i,w-}$:

$$\Delta EquityShare_{i,w} = \alpha + \beta E_{i,w-}[R_{1y}] + \gamma \Delta E_{i,w}[R_{1y}] + \delta EquityShare_{i,w-} + \phi X_{i,w} + \epsilon_{i,w}. \tag{3}$$

The vector $X_{i,w}$ includes a set of time-window-length fixed effects, as well as controls for age, gender, region of residence, wealth, wave fixed effects, and dummies for initial equity shares of 0 percent and 100 percent. Column 1 of Table 5 reports the main coefficients; in online Appendix Section A.7, we also report the coefficients on the control variables as well as estimates with the sample restricted to the pre-COVID-19 period. A 1 percentage point increase in expected returns at the beginning of the window predicts a 0.13 percentage point increase in the equity share due to trading over the following window; a 1 percentage point change in beliefs over the window predicts a 0.23 percentage point change in the equity share. While these sensitivities are statistically significant, they are smaller than what we

[24] For example, if an investor has answered waves 1, 2, and 5 of the survey, we would identify two time windows: the 2-month period between wave 1 and wave 2, and the 6-month period between wave 2 and wave 5. Each time window would appear as a separate observation in regression (3).

Table 5—Trading Analysis

	Δ Equity share (pp) (1)	Probability trade (2)	Probability trade (3)	Probability buy (4)	Δ Equity share (pp) (5)
Δ expected 1Y stock return (pp)	0.229			0.977	0.587
	(0.033)			(0.201)	(0.085)
Expected 1Y stock return (pp)	0.130	0.006		1.489	0.395
	(0.019)	(0.130)		(0.199)	(0.062)
Lagged equity share (pp)	−0.048	−0.121		−0.337	−0.161
	(0.004)	(0.021)		(0.031)	(0.011)
|Δ expected 1Y stock return (pp)|		0.258			
		(0.223)			
Extreme equity share dummies	Y	Y	Y	Y	Y
Time between wave dummies	Y	Y	Y	Y	Y
Other fixed effects and controls	Y	Y	N	Y	Y
Specification				Conditional on trading	Conditional on trading
R^2	0.031	0.380	0.364	0.483	0.154
Observations	22,439	22,439	23,441	6,606	6,606

Notes: Table shows results from regression (3). The unit of observation is a window between two consecutive survey responses by the same individual. The dependent variable in columns 1 and 5 is the change in the equity share due to active trading between the two survey waves. The dependent variable in columns 2 and 3 is an indicator for whether there was any active trading between the two survey waves, defined as an active change in the equity share of at least one percentage point. The dependent variable in column 4 is an indicator of whether the individual actively increased her portfolio share in equity by at least one percentage point during the window as a result of trading between the two survey waves. All columns control for the length of time between two consecutive answers, and for dummies capturing extreme start-of-period equity shares of 0 percent or 100 percent. All columns, except column 3, also control for the respondents' age, gender, region of residence, wealth, and the survey wave. Columns 4 and 5 condition the sample on windows with active trades. All results are obtained using ORIV. The R^2 is computed from OLS specifications. Standard errors are clustered at the respondent level.

obtained from the cross-sectional analysis in Section II.[25] Column 1 also shows that investors with high equity shares at the beginning of the window tend to actively reduce their equity exposures, potentially a sign of rebalancing of their positions.

The low sensitivity in column 1 could reflect two different mechanisms. First, it could simply result from the fact that individuals trade infrequently, so that the average sensitivity to beliefs appears low (extensive margin). Alternatively, it could reflect a low sensitivity of trading to beliefs even when investors trade actively (intensive margin). We next explore these explanations.

The Extensive Margin of Trading.—A large literature aims to explain trading volume in financial markets via a mix of changes in beliefs and overconfidence (e.g., Harrison and Kreps 1978; Hong and Stein 1999, 2007; Scheinkman and Xiong 2003). A related literature has studied the optimal frequency and size of trading in the presence of limited information and transaction costs (e.g., Duffie and Sun 1990; Gabaix and Laibson 2001; Abel, Eberly, and Panageas 2007; Alvarez, Guiso,

[25] Since these regressions are analyzing changes over time in portfolio choice, rather than levels as in Section II, 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.

and Lippi 2012). A natural question in our setup, therefore, is whether changes in beliefs are associated with trading activity.

Column 2 of Table 5 shows results from a regression similar to equation (3), except that the dependent variable is an indicator of whether the investor has actively traded during the time window (defined as an active change in the equity share of at least 1 percentage point), and the change in expected returns over the window is replaced with its absolute value. The ability of belief changes to predict trading is statistically and economically small, with both beginning-of-window expected returns and changes over the window having no effect on the probability of trade. While the R^2 of the regression appears high at 38 percent, column 3 of Table 5, which does not include controls for beliefs, portfolio shares, and demographics, displays a similarly high R^2 of 36 percent. The high R^2 in columns 2 and 3 are thus due to the window-length fixed effects: trading is mechanically more likely to occur over longer time windows. The incremental explanatory power of beliefs, portfolios, and demographics in predicting the extensive margin of trading is small.

The Intensive Margin of Trading.—In our next analysis, we condition on time windows during which individuals trade actively, and ask whether the direction and the magnitude of the trade can be explained by beliefs. We start by predicting the direction of trading. Column 4 of Table 5 reports the results of a regression similar to equation (3), except that the dependent variable is now an indicator of whether the investor has actively increased her equity share by at least one percentage point during the window. Beliefs predict the direction of trading conditional on a trade actually occurring: an investor who expects future returns at the beginning of the wave to be 1 percentage point higher is around 1.5 percentage points more likely to buy equities in a given window. The fact that beginning-of-wave beliefs affect subsequent trading activity is consistent with our earlier findings of infrequent trading: if trading occurs infrequently, one's current portfolio does not always capture one's current beliefs. A 1 percentage point increase in expected returns over the wave is also associated with a 1 percentage point higher probability of actively increasing the equity share (conditional on trading).

Column 5 of Table 5 explores the magnitude of trading conditional on a trade occurring. The dependent variable is the change in the equity share due to active trading: that is, the regression is the same as for column 1, but the results are conditional on trading taking place (again, measured by the equity share changing by at least 1 percent in any direction due to active trading). Conditional on trading, the sensitivity of trading to beliefs increases by a factor of three compared to the unconditional results: a 1 percentage point increase in investors' expected 1-year stock returns corresponds to a 0.4 percentage point increase in the equity share due to trading. When we condition on larger trades (at least a 5 percent change in the equity share) the magnitudes increase considerably with estimates around 0.92 percent for the change in beliefs and 0.77 percent for the level of beliefs.[26]

[26] We also investigate how the allocation of "new funds" across different asset classes depends on individuals' beliefs. We consider all cases in which, during a window, we see a net inflow of outside cash into the Vanguard account that is at least 20 percent of the existing 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. Since investors devote some time to deciding an allocation for funds when they first transfer them to their trading account, this represents

Online Appendix Table A.13 presents the same analysis as Table 5 but restricts the sample to end in February 2019, just before the COVID-19 crisis. It shows broadly similar patterns with a stronger quantitative relationship on the intensive margin between beliefs and trading. During the COVID-19 crisis, particularly during the stock market crash in March 2020, beliefs became substantially more pessimistic while portfolios change only by a small amount. Those who became more pessimistic do tend to trade out of equity conditional on trading, confirming the association of beliefs and trading even during a major turmoil.

Our analysis in this section confirms that trading patterns play a central role in reducing the pass-through of beliefs into asset demand. We summarize the key results in the following fact.

Fact 2: While belief dynamics 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).

As we discussed above, one promising path for future theoretical work hoping to generate the relatively low average sensitivity of portfolios to beliefs is to explicitly account for infrequent trading. The results in this section suggest that one parsimonious way to model such behavior is to introduce infrequent random trading, whereby an agent is selected at random based on a memory-less distribution to have the possibility of trading in a given period. This approach would be reminiscent of the Calvo (1983) adjustment model for firm pricing decisions, and would be consistent with both the fact that the only observable variable that predicts whether an individual trades is the window length, as well as with the low average pass-through of beliefs to portfolios established in Fact 1. Researchers who want to match the cross-sectional heterogeneity in trading frequency in addition to the low average trading frequency could explicitly model different arrival rates of trading opportunities for different individuals.[27]

IV. Variance Decomposition of Beliefs

Section I documented substantial heterogeneity in investors' beliefs. In this section, we further explore this heterogeneity by decomposing the panel variation of beliefs into three components: fixed individual characteristics, common variation in individual beliefs over time, and a residual component that captures both idiosyncratic individual time variation and measurement error.

To motivate this variance decomposition, panel A of Figure 3 shows the time-series of average 1-year expected returns in the GMSU-Vanguard survey. The average expected return displays meaningful time-series variation, with a range of over

a particularly informative window to observe how beliefs affect portfolio composition. We repeat the regression of column 5 of Table 5, but also condition 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 0.92 for belief changes and to 0.68 for belief levels.

[27] The correlation between age, wealth, and trading frequency established in online Appendix Section A.4, as well as the analysis presented in online Appendix Table A.6, can provide further guidance to researchers hoping to incorporate heterogeneous infrequent trading into richer life-cycle models.

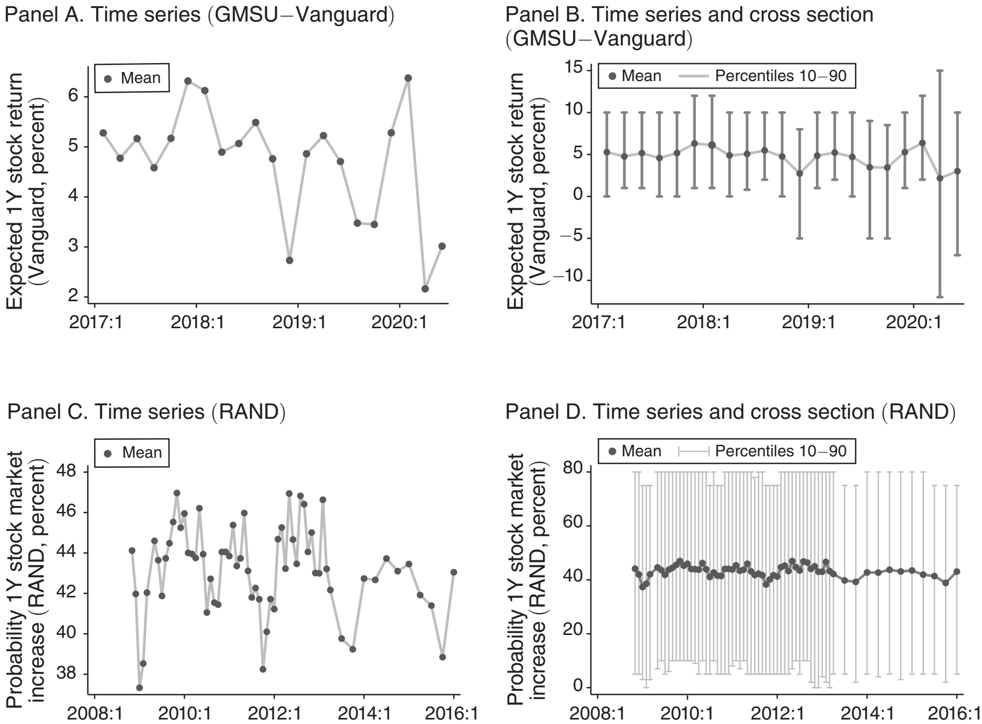

Figure 3. Time-Series and Cross-Sectional Variation: GMSU-Vanguard and RAND

Notes: Figure reports the time series of the average beliefs from the GMSU-Vanguard survey (top row, 1-year expected return question) and from the RAND survey (bottom row, probability of a 1-year stock market increase). The right panels in each row also plot the tenth and ninetieth percentiles of the survey answers in each wave.

4 percentage points over our sample period. The largest month-to-month change was in March 2020, when average expected returns fell by over 4 percentage points after the stock market crash induced by the COVID-19 crisis. Panel B of Figure 3 shows the same time series of average expected returns as in panel A, but also includes the tenth and ninetieth percentiles of the cross-sectional distribution of answers in each wave. The cross-sectional variation in expected returns dwarfs the time-series variation, not only in "normal times," but also during the COVID-19 crisis.

This pattern is not unique to our survey or our sample period. For example, the bottom row of Figure 3 shows similar plots for the RAND survey, which covers the period from November 2008 to January 2016 and thus includes part of the financial crisis and the following stock market recovery (see online Appendix Section A.3 for more details on the RAND survey). Unfortunately, the RAND survey does not directly elicit beliefs about expected returns, so we focus on beliefs about the probability of a stock market increase over the coming year. We find that the RAND survey also features cross-sectional dispersion in beliefs that is much larger than the time-series variation.

There are two potential interpretations consistent with the substantial cross-sectional dispersion in beliefs. At one extreme, individual responses might display substantial idiosyncratic variation over time, with the same individual

Table 6—Decomposing the Variation in Beliefs: Individual and Time Fixed Effects

	R^2 (percent) of panel regression			
	Time FE	Individual FE	Time + individual FE	Observations
Expected 1Y stock return (percent)	5.0	47.5	51.5	1,960
Expected 10Y stock return (percent p.a.)	0.5	45.0	45.5	1,964
Probability 1Y stock return $< -10\%$	2.7	51.5	53.6	2,011
SD expected 1Y stock return (percent)	0.5	56.7	57.2	2,011
Confidence (stock Qs)	1.4	60.6	62.0	1,988
Expected 3Y GDP growth (percent p.a.)	3.8	43.9	46.8	1,968
Expected 10Y GDP growth (percent p.a.)	0.6	39.7	40.2	1,952
Probability p.a. 3Y GDP growth $< 0\%$	5.1	45.4	49.3	2,010
SD expected p.a. 3Y GDP growth (percent)	1.0	56.5	57.3	2,010
Confidence (GDP Qs)	0.8	62.8	63.8	1,978
Expected 1Y return of 10Y bond (percent)	2.4	38.8	40.7	1,953
Confidence (bond Qs)	0.4	62.9	63.3	1,969

Notes: Table reports the R^2 values corresponding to the three regressions (4), (5), and (6), and the number of individual respondents' observations. We only include respondents who have responded to at least five waves. Each row corresponds to a different survey question that is used as the dependent variable.

reporting very different beliefs at different points in time. At the other extreme, the observed cross-sectional variation could be due to *persistent* heterogeneity in beliefs; that is, the *same* investors are always optimistic or always pessimistic. Since these interpretations have substantially different implications for theoretical models of asset pricing, we next exploit the panel dimension of our survey to determine their quantitative relevance.

The Dominance of Individual Fixed Effects.—We denote the belief expressed by individual i at time t as $B_{i,t}$. For the (unbalanced) panel of these beliefs, we then run the following regressions:

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

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

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

Equation (4) estimates a set of time (i.e., survey wave) fixed effects, χ_t, that absorb the common time-series variation of respondents' beliefs. Equation (5) estimates a set of individual fixed effects, ϕ_i, that absorb the average belief over time of each respondent. Equation (6) jointly estimates both individual and time fixed effects. In our baseline analysis, we estimate these regressions including all responses from individuals who have responded to at least five waves.

Table 6 reports the R^2 statistics of the three regressions for a subset of survey questions. Most of the panel variation in beliefs is absorbed by individual fixed effects. Consider, for example, the first row, which decomposes the panel variation in 1-year expected stock returns. Time fixed effects capture 5 percent of the total panel variation, whereas individual fixed effects capture 47.5 percent of the total variation. The remaining variation is a combination of idiosyncratic belief changes

Table 7—Decomposing the Variation in Beliefs: Robustness

	R^2 (total, percent)				Number of individuals			
	#Resp≥5	#Resp≥6	#Resp≥7	#Resp≥8	#Resp≥5	#Resp≥6	#Resp≥7	#Resp≥8
Expected 1Y stock return (percent)	47.5	46.5	46.6	45.9	1,960	1,361	974	712
Expected 10Y stock return (percent p.a.)	45.0	43.7	45.8	44.9	1,964	1,360	959	704
Probability 1Y stock return $< -10\%$	51.5	51.3	52.1	52.3	2,011	1,389	1,003	729
SD expected 1Y stock return (percent)	56.7	57.6	57.7	57.8	2,011	1,389	1,003	729
Confidence (stock Qs)	60.6	60.4	60.9	60.5	1,988	1,374	975	718
Expected 3Y GDP growth (percent p.a.)	43.9	43.2	42.8	39.6	1,968	1,371	978	715
Expected 10Y GDP growth (percent p.a.)	39.7	39.6	38.5	36.7	1,952	1,342	963	708
Probability p.a. 3Y GDP growth $< 0\%$	45.4	44.1	43.9	44.4	2,010	1,392	1,000	730
SD expected p.a. 3Y GDP growth (percent)	56.5	57.5	57.5	58.3	2,010	1,392	1,000	730
Confidence (GDP Qs)	62.8	62.7	62.2	62.1	1,978	1,364	978	721
Expected 1Y return of 10Y bond (percent)	38.8	37.3	35.6	34.7	1,953	1,342	968	705
Confidence (bond Qs)	62.9	62.9	62.6	63.1	1,969	1,363	981	703

Notes: The left panel reports the R^2 values corresponding to regression (5). The right panel reports the number of individuals that responded the required number of times. Across columns, we increase the minimum number of responses for an individual to be included in the sample from 5 to 8. Each row corresponds to a different survey question that is used as the dependent variable.

within individuals over time, as well as measurement error in beliefs.[28] This large difference in explanatory power of time fixed effects and individual fixed effects is common across all beliefs. The same patterns hold when we decompose the heterogeneity in individuals' confidence in their beliefs: most of the variation is across individuals rather than over time.[29]

There are two possible concerns with this analysis. First, the relatively short time period over which we observe survey responses might make the fixed effects appear more important than they truly are. Second, the large COVID-19 shock might overstate the importance of time-series variation, since shocks of that magnitude are historically rare.

On the concern that our analysis may be overfitting the fixed effects in sample, especially for investors that reply only a few times, we find that when we increase the minimum number of waves that an individual has to answer to be included in the analysis, results are very similar. In particular, Table 7 shows how the R^2 statistics of the individual fixed effects changes as we increase the minimum number of responses per individual.[30] We find at most a modest deterioration in the importance of individual fixed effects as we increase the minimum number of answers. This finding suggests that our results are not driven by overfitting the fixed effects for people who have responded only a few times.[31]

[28] These findings echo results in Dominitz and Manski (2011), who show that individuals' responses for the probability of a positive equity return over the coming twelve months were correlated across two waves of the Michigan Survey of Consumers. The authors also found substantial heterogeneity in this probability across individuals.

[29] In online Appendix Section A.10, we apply the same variance decomposition to portfolios and find that individual fixed effects explain 87 percent of the panel variation in equity shares during our sample period. When we relate the individual fixed effects extracted from beliefs to those extracted from portfolios, we recover a sensitivity very similar to our benchmark estimates in Section II.

[30] The table also reports the number of individuals who respond a certain number of times. The number of observations is of course greater since each individual has answered multiple times.

[31] Another possibility is that the fixed effects capture what in reality is a stationary but extremely *persistent* process of beliefs, even though there is no difference in the *permanent* component of beliefs. Since the economic

We also perform the variance decomposition presented above on the RAND survey, which also has a panel structure. As discussed above, the RAND survey ran for more than seven years and covered much of the Great Recession and subsequent recovery; overall, there are 61 survey waves. A total of 4,734 individuals participated in the survey, 3,166 of whom responded at least 10 times, 1,780 at least 30 times, and 1,032 at least 50 times. When we perform the same variance decomposition for the RAND survey, we find quantitatively similar results (see online Appendix Section A.9 for details). Indeed, across all questions in the RAND survey that relate to expected stock returns, time fixed effects explain around 1 percent of the panel variation, while individual fixed effects explain 50–60 percent of the variation. Importantly, the results are robust to increasing the minimum number of waves that an individual has to respond to in order to be included from three all the way to fifty. These results highlight that our findings are robust to different economic environments as well as to observing substantially more responses per individual.[32]

To explore the concern that our results might overstate the importance of time-series variation once the COVID-19 shock is included, online Appendix Table A.14 reports an analysis similar to Table 6, but restricting the sample to end in February 2020, before the COVID-19 crisis. Naturally, the importance of the time-series variation decreases, and the results are, if anything, strengthened: individual fixed effects are even more important to explain the panel variation in beliefs. While a reader might want to informally think of the period before and the period including the COVID-19 shock as upper and lower bounds for the relative importance of individual fixed effects and time series variation, the economic conclusions are very similar.

The importance of persistent cross-sectional dispersion in beliefs provides useful insights for the design of macro-finance models. In particular, much of the existing literature that builds on survey evidence of beliefs has focused on representative agent models disciplined by matching the time-series behavior of average beliefs (e.g., Barberis et al. 2015). This literature misses a crucial feature of the data: individual heterogeneity. Our results offer a new set of moments that can be used to enrich the models in this under-explored dimension. In this direction, models that explicitly feature heterogeneous agents with different beliefs, such as the model of Geanakoplos (2010), are likely to offer a fruitful starting point for further exploration.

interpretation of permanent versus extremely persistent differences in beliefs is not one that is crucial to most theories, we do not aim to definitively distinguish between these interpretations. Instead, we view our results as emphasizing that there is large and persistent cross-sectional dispersion of beliefs and the fixed-effects analysis is simply a transparent way to document this pattern. Nevertheless, one can try to statistically disentangle the two interpretations by estimating a panel model for beliefs that features both fixed effects and an AR(1) component and use a statistical test to distinguish between the two explanations. When we estimate this model using the Arellano and Bond (1991) estimator for dynamic panel data, we find that the autoregressive component is small in absolute value and statistically insignificant (for example, it is −0.03 for 1-year expected returns), providing suggestive evidence against the AR(1) interpretation of our results.

[32] While the pattern of persistent and large belief differences across retail investors appears consistent across surveys covering different time horizons and investor populations, it would be interesting to study the same relationship among institutional investors or professional forecasters. However, such analyses need to carefully account for the various incentives of the respondents, which is less of a concern in non-public surveys of retail investor beliefs. For example, Ottaviani and Sørensen (2006) discuss various aspects of professional forecasters' strategic behavior, highlighting the presence of incentives to herd (see also Graham 1999; Rangvid, Schmeling, and Schrimpf 2013).

Table 8—Beliefs Heterogeneity and Demographics

R^2	#Resp ≥ 1	#Resp ≥ 2	#Resp ≥ 3	#Resp ≥ 4	#Resp ≥ 5
Expected 1Y stock return (percent)	2.5	4.3	5.3	6.1	6.8
Expected 10Y stock return (percent p.a.)	1.9	3.5	4.4	4.8	5.6
Probability 1Y stock return $< -10\%$	3.0	4.4	6.3	6.5	6.9
SD expected 1Y stock return (percent)	7.1	8.3	10.8	10.2	9.4
Expected 3Y GDP growth (percent p.a.)	2.3	3.0	3.5	4.5	5.3
Expected 10Y GDP growth (percent p.a.)	3.0	3.7	3.8	4.0	4.4
Probability p.a. 3Y GDP growth $< 0\%$	5.0	6.4	8.5	8.5	7.9
SD expected p.a. 3Y GDP growth (percent)	4.8	4.8	5.3	5.1	4.3
Expected 1Y return of 10Y bond (percent)	2.2	3.0	3.0	3.7	4.1

Notes: 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.

Beliefs and Demographics.—Having established the importance of individual fixed effects in explaining the panel variation in beliefs, it is natural to ask whether observable characteristics can explain why some individuals are permanently optimistic and others are permanently pessimistic. We find that observable individual characteristics have little explanatory power for beliefs, even though some of these characteristics are related to beliefs in statistically significant ways. To establish this finding, we run the following regression:

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

where $\phi_{3,i}$ are the individual fixed effects estimated in regression (6), and X_i are the following individual characteristics: dummy variables for age groups, wealth quintiles, region of residence, gender, confidence, and quintiles for the number of days with Vanguard logins in an average month.[33] In addition, motivated by recent evidence that investors' past experiences influence their beliefs (e.g., Malmendier and Nagel 2011), we also include the average return on the equity and fixed income components of the investors' portfolios since 2011 as controls in X_i. Table 8 shows the R^2 statistics from these regressions, which capture the share of variation in the fixed effects that is explained by the demographics.

The observed characteristics have only small explanatory power, with values for the R^2 between 2 percent and 7 percent depending on the question (using our complete sample in the analysis). When we restrict the analysis to explaining fixed effects that are estimated on more observations, and which should therefore be more precise, there is only a modest increase in the R^2. We thus conclude that classical measurement error in beliefs cannot explain the low predictive power of demographics for beliefs.

Online Appendix Section A.8 reports the coefficients on the various demographic characteristics from regression (7). Despite the low overall explanatory power of demographics for beliefs, some of these characteristics have statistically significant relationships with beliefs. For example, we find that older individuals are

[33] 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.

more optimistic, while wealthier respondents are more pessimistic. In addition, we find that individuals who experienced higher past equity returns are more optimistic about future stock returns, while individuals who experienced higher past fixed income returns are more optimistic about future bond returns.

These results relate to the literature that explores the relationship between expectations and demographic characteristics and personal experiences. It is common in this literature to find strong statistical relationships but low explanatory power for expectations using variables such as wealth, gender, IQ, place of birth, current location, own past experience, or friends' past experiences (see, for example, Malmendier and Nagel 2011; Kuchler and Zafar 2019; Armona, Fuster, and Zafar 2016; Das, Kuhnen, and Nagel 2020; Bailey et al. 2017, 2018; Ben-David et al. 2018; Coibion, Gorodnichenko, and Kamdar 2018; D'Acunto et al. 2019). The low predictive power suggests that these individual fixed effects reflect complex combinations of individual characteristics and experiences, some of which economic research has yet to discover. We collect the findings from this section in the following fact.

Fact 3: Variation in individual beliefs is mostly characterized by heterogeneous individual fixed effects: between 40 percent and 60 percent of the panel variation in responses is due to individual fixed effects, and 5 percent is due to common time series variation. The remaining variation is accounted for by idiosyncratic individual variation over time and measurement error. Only a small part of the persistent heterogeneity in individual beliefs is explained by observable demographic characteristics.

V. Covariation in Expected Returns and Expected Cash Flows

Asset prices are determined by expectations of future returns and cash flows. It is therefore natural to investigate how expectations of returns and economic growth are related both across individuals and within each individual over time.

Figure 4 presents conditional binscatter plots of the relationship between short-run and long-run expectations of stock returns and GDP growth. Panel A shows that expectations about short-run and long-run stock returns are positively correlated, with an unconditional correlation coefficient of 0.30. Interestingly, even those respondents who expect negative returns over the next year expect long-run returns to be positive. Similarly, short-run and long-run GDP growth are positively correlated, with an unconditional correlation coefficient of 0.65 (see panel B). The bottom row of Figure 4 shows that expectations of stock returns and economic growth are positively correlated, both at the short horizon (panel C, unconditional correlation coefficient of 0.26) and at the long horizon (panel D, unconditional correlation coefficient of 0.27).

Table 9 presents these results in regression form, both with and without including individual fixed effects. While these regressions are restricted to linear specifications, and therefore miss some of the interesting nonlinearities in Figure 4, the findings confirm the strong link between expectations about different objects, both across horizons as well as across domains. Importantly, we see these patterns in both the cross-section and the time series (when we control for individual fixed effects).

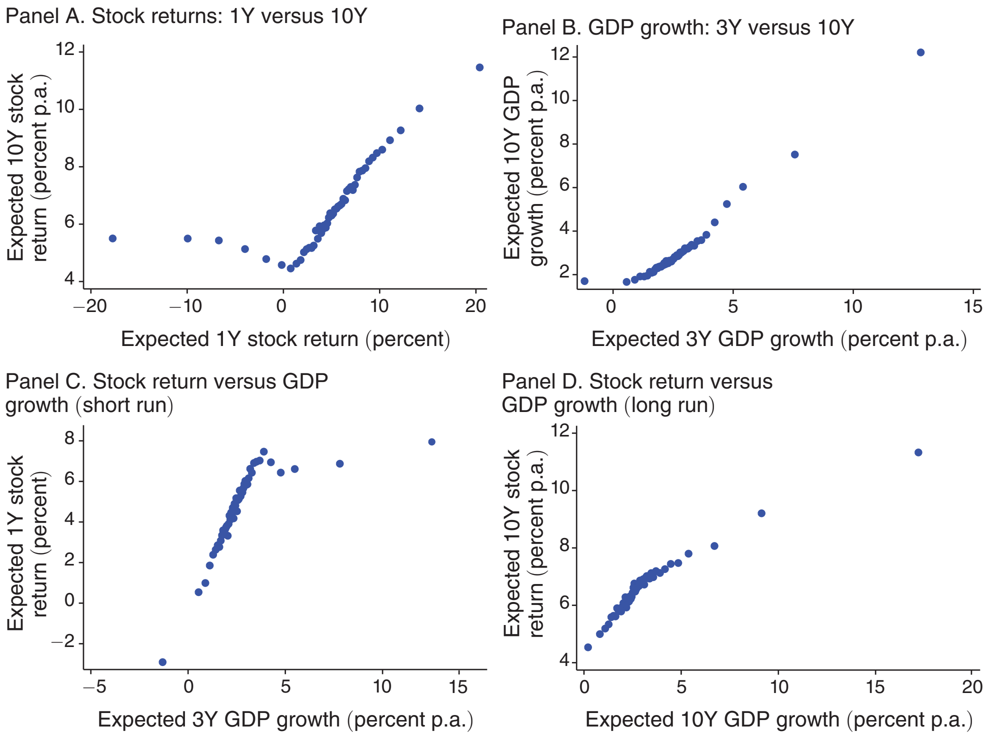

Figure 4. Relationships among Different Beliefs within the Same Survey

Note: Figure shows conditional binscatter plots across different answers from within the same survey response, conditional on the respondents' age, gender, region, wealth, and the survey wave.

In the cross-section, individuals who are more optimistic about stock returns tend to also be more optimistic about GDP growth. In the time series, we find that when an investor becomes more optimistic about stock returns, she also becomes more optimistic about GDP growth.[34] Online Appendix Section A.12 shows that these patterns also hold when excluding the period of the stock market crash in March 2020. We summarize these results in the following fact.

Fact 4: Higher expectations of stock returns are associated with higher expectations of GDP growth, and higher short-run expectations are associated with higher long-run expectations (for both stock returns and GDP growth), both across and within individuals.

The correlation between expected returns and cash flow growth is an informative moment for macro-finance models. To see why, it is useful to refer to the Campbell

[34] While Table 9 shows the results in a linear setting, a similar conclusion can be reached in a nonlinear setting as well, by building binscatter plots that relate the fixed effects of beliefs across domains (thus isolating the cross-sectional component) and by plotting the residual components after taking out fixed effects (thus focusing on the within-individual time variation). We report these plots in online Appendix Section A.11. The conclusions are identical to those in this section: the panel-correlation of beliefs across different domains and across different horizons reflects similar correlations in the persistent and transient components of beliefs.

Table 9—Correlation across Survey Responses

	Expected 10Y stock returns (percent p.a.)		Expected 10Y GDP growth (percent p.a.)	
	(1)	(2)	(3)	(4)
Panel A				
Expected 1Y stock return (percent)	0.198	0.100		
	(0.006)	(0.009)		
Expected 3Y GDP growth (percent p.a.)			0.824	0.640
			(0.016)	(0.039)
Controls	Y	Y	Y	Y
Individual fixed effects	N	Y	N	Y
R^2	0.111	0.711	0.442	0.810
Observations	42,978	42,978	42,751	42,751
Panel B				
Expected 3Y GDP growth (percent p.a.)	0.710	0.381		
	(0.023)	(0.044)		
Expected 10Y GDP growth (percent p.a.)			0.388	0.260
			(0.015)	(0.035)
Controls	Y	Y	Y	Y
Individual fixed effects	N	Y	N	Y
R^2	0.105	0.712	0.089	0.717
Observations	42,926	42,926	42,543	42,543

Notes: Table shows results from regressing answers to different expectation questions onto each other; panel A relates short-run and long-run beliefs within the same domain, while panel B relates beliefs across domains over similar time horizons. The unit of observation is a survey response. All specifications control for the respondents' age, gender, region of residence, wealth, and the survey wave. Columns 2 and 4 also control for respondent fixed effects. Standard errors are clustered at the respondent level.

and Shiller (1988) decomposition, which shows how prices, expected cash flows, and expected returns are linked:

$$pd_t \approx E_{i,t}\sum_{j=0}^{\infty}\rho^j \Delta d_{t+1+j} - E_{i,t}\sum_{j=0}^{\infty}\rho^j r_{t+1+j}.$$

In this equation, pd_t is the logarithm of the price-dividend ratio of an asset, Δd_{t+1} is the growth of cash flows between t and $t+1$, and r_{t+1} is the return of the asset between t and $t+1$. For expositional convenience, we assume that this equation holds approximately under each investor i's expectations $E_{i,t}$. If we take our GDP growth responses to proxy for cash flow growth, then we can use this equation to interpret our empirical results.

One immediate implication of Fact 4 is in the time-series dimension (see De la O and Myers 2021). As the Campbell-Shiller decomposition shows, cash flow expectations and expected returns have opposite effects on current valuations. All else equal, when investors become more optimistic about cash flows, asset prices rise; but if expected returns simultaneously increase, as they do in the data, prices will be lower through a discount-rate effect. Therefore, accounting for the joint variation of expected returns and expected cash flow growth is important to understand the movement of asset prices. For example, models that match survey variation in expected cash-flow growth, but ignore the correlated variation in expected returns, are likely to overstate the power of the variation in cash flow expectations for explaining time-series variation in asset prices (e.g., Bordalo et al. 2020).

The implications of our results for the cross-section (disagreement among investors at each point in time) are more subtle. While investors might disagree about future cash flows or about future expected returns, they all face the same current price pd_t. Assuming that the Campbell-Shiller identity holds for each individual, whether they are an optimist or a pessimist, this implies that expectations of cash flows and expectations of returns *need to be* positively correlated in the cross-section.[35] For example, consider two investors, one optimistic and one pessimistic about future cash flow growth. Given that they both face the same price, the optimistic investor *has to* expect higher returns than the pessimistic one.[36] However, the Campbell-Shiller decomposition is silent about the exact horizon at which this correlation will occur. For example, it does not tell us whether disagreement about short-term cash flows is matched by disagreement about short-term expected returns or by disagreement about long-term expected returns. This is where our empirical results add value to this decomposition: the results provide evidence on the correlations of cash flow and returns *at specific horizons*, thereby guiding the calibration of the term structure of disagreement in asset pricing models.

VI. Rare Disasters and Expected Returns

In the previous sections, we explored a number of moments of the belief distribution that have been of central interest to the asset pricing literature, such as average expected returns and average expected GDP growth rates. In addition, an important strand of the macro-finance literature has emphasized that expectations of rare but potentially catastrophic events, sometimes called rare disasters, can help explain expected returns, portfolio holdings, and asset prices (Rietz 1988, Barro 2006, Gabaix 2012).[37] To further understand these relationships, we exploit that our survey directly elicits expectations of disaster probabilities for both stock returns (i.e., 1-year stock returns of less then -30 percent) and GDP growth (i.e., annualized 3-year GDP growth of less than -3 percent).[38]

We first explore the relationship between individuals' expectations of the probabilities of stock market disasters and GDP disasters. The left panel of Figure 5 shows that expectations of the two types of disasters are positively related at the individual level (the slope of the regression line is 0.39); in unreported results, we find that this is also true within individuals over time. 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., beliefs about stock market disasters are not purely due to beliefs about discount rate variation).

[35] The equation is an identity and thus only imposes mild restrictions on expectations. Nonetheless, agents' expectations might violate the identity, e.g., because they have imperfect knowledge of the current price-dividend ratio.

[36] The Campbell-Shiller decomposition as written assumes that investors believe that all mispricing will be corrected eventually; if that was not the case, the equation would feature an extra term, containing the limit of pd_{t+n} for $n \to \infty$.

[37] Recently, Goetzmann, Kim, and Shiller (2018) have studied the determinants of beliefs about rare disasters.

[38] Online Appendix Section A.6 highlights the importance of subjective disaster probabilities for portfolio formation: holding fixed the mean, respondents with a higher perceived probability of stock market disasters also had lower equity shares.

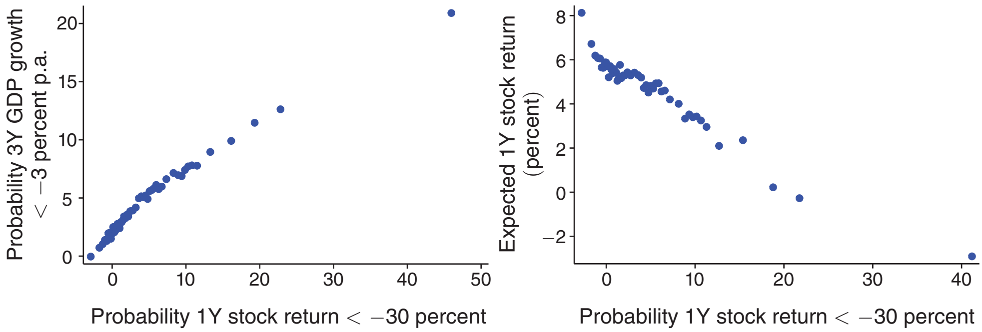

Figure 5. Stock Disasters, GDP Disasters, and Expected Returns

Notes: The left panel shows a conditional binned scatter plot of survey respondents' subjective probabilities that the 1-year stock returns are below −30 percent and their expectations that annualized average GDP growth over the next three years is below −3 percent. The right panel shows a conditional binned scatter plot of survey respondents' subjective probabilities that the 1-year stock returns are below −30 percent and their expected 1-year stock returns. Both panels condition on the respondents' age, gender, region, wealth, and the survey wave.

We also analyze the relationship between expected returns and expected disaster probabilities. The right panel of Figure 5 shows that individuals who report a higher subjective probability of a stock-market disaster also report lower expected stock returns. To explore this relationship more formally, we run the following regression:

$$E_{i,t}\big[R_{1y}\big] = \alpha + \beta \mathrm{Pr}_{i,t}\big[R_{1y} < -30\%\big] + \gamma X_{i,t} + \psi_t + \epsilon_{i,t}, \tag{8}$$

where the coefficient of interest is β. We additionally control for demographic characteristics such as age, gender, wealth, and region of residence, as well as survey-wave fixed effects.

The specification in column 1 of Table 10 corresponds to the right panel of Figure 5. The estimate of β implies that a 5 percentage point increase in an individual's subjective probability of a stock market disaster is associated with a 1 percentage point decline in her subjective expected returns. Column 2 shows that a similar negative relationship occurs when we consider the probability of less extreme outcomes, i.e., returns below −10 percent. Column 3 restricts the data to those answers that report the probability of a stock market return of less than −30 percent to be between 0.1 percent and 10 percent. We find that excluding extreme responses increases the magnitude of the sensitivity from −0.21 to −0.27. Column 4 shows that the results are not meaningfully affected by the order in which the buckets are presented to the respondent in the distribution question (high-to-low versus low-to-high). Column 5 includes individual fixed effects, and column 6 does the same but restricts the probabilities to be in the same range as in column 3. These latter columns show that the negative relationship between expected returns and subjective disaster probabilities also holds in the time series for each individual. Online Appendix Section A.12 shows that these patterns also hold when excluding the period of the stock market crash in March 2020. We collect the findings in this section in the following fact.

Table 10—Expected Stock Returns and Rare Disasters Beliefs

	Expected 1Y stock return (percent)					
	(1)	(2)	(3)	(4)	(5)	(6)
Probability 1Y stock return < −30%	−0.212 (0.008)		−0.270 (0.013)		−0.146 (0.013)	−0.153 (0.017)
Probability 1Y stock return < −10%		−0.152 (0.003)				
Probability 1Y stock return < −30% × low bucket first				−0.208 (0.010)		
Probability 1Y stock return < −30% × high bucket first				−0.217 (0.012)		
Controls	Y	Y	Y	Y	Y	Y
Individual fixed effects					Y	Y
Specification			Prob. $\in$ [0.1%, 10%]			Prob. $\in$ [0.1%, 10%]
R^2	0.117	0.244	0.055	0.438	0.722	0.764
Observations	43,492	43,492	23,466	43,492	43,492	22,670

Notes: Table shows results from regression (8). The unit of observation is a survey response; the dependent variable is the expected 1-year stock return. Columns 1 to 4 control for the respondents' age, gender, region of residence, wealth, and the survey wave. Columns 5 and 6 include individual fixed effects, survey wave fixed effects, and a dummy for the randomization order of the buckets in the distribution question. Columns 3 and 6 restrict the sample to individuals who report expected probabilities of a stock market disaster between 0.1 percent and 10 percent. Standard errors are clustered at the respondent level.

Fact 5: Higher subjective probabilities of stock market disasters are associated with lower expected stock market returns, both across and within individuals.

Our cross-sectional results in columns 1–4 of Table 10 map most closely to models in which agents disagree about the probability of disasters and are overconfident in their beliefs (i.e., they "agree to disagree"). For example, in the model of Chen, Joslin, and Tran (2012), agents differ in their subjective beliefs about the probability of cash flow disasters. Since all agents observe the current stock price, those agents who think that disasters are more likely also tend to expect lower returns. Our findings support the Chen, Joslin, and Tran (2012) model prediction that the optimists expect both high returns and a lower probability of disaster relative to the pessimists.

Our results on the within-individual time-series relationship between disaster beliefs and expected returns in columns 5 and 6 of Table 10 relate to the literature on time-varying rare-disasters with representative agents (e.g., Gabaix 2012, Wachter 2013). In these models, a representative agent with rational expectations prices assets in an economy affected by time-varying rare disasters. In equilibrium, expected returns and the disaster probability are positively related in the time series. The intuition is that a higher disaster probability induces individuals to demand higher compensation for holding the stock market, which increases equilibrium expected returns. The relationship in the data at the individual level appears with an opposite sign relative to the theory. Mapping our individual-level partial-equilibrium results into a general equilibrium representative-agent model is beyond the scope of this paper, but offers an interesting avenue to further develop the rare disaster paradigm (see Jin 2015 for a behavioral model in this direction).

VII. Conclusion

In this paper we analyzed a new survey of investor beliefs. We combined the survey responses with administrative data on respondents' portfolio holdings and trading activity to establish five facts about the relationship between investor beliefs and portfolios. These facts provide guidance on the construction of macro-finance models. In particular, we highlight three ingredients for new models that the future literature could develop (i) large and highly persistent heterogeneity in beliefs about both expected returns and cash flows, with the two beliefs positively related, (ii) a willingness to "agree to disagree" that allows for trading based on disagreement, and (iii) infrequent trading with an exogenous probability of trading that differs across agents. An interesting open question is how well such a model would perform in quantitatively matching asset prices in addition to the main features of beliefs and portfolios documented in this paper.

REFERENCES

Abel, Andrew B., Janice C. Eberly, and Stavros Panageas. 2007. "Optimal Inattention to the Stock Market." *American Economic Review* 97 (2): 244–49.

Adam, Klaus, Johannes Beutel, Albert Marcet, and Sebastian Merkel. 2015. "Can a Financial Transaction Tax Prevent Stock Price Booms?" *Journal of Monetary Economics* 76: S90–S109.

Adam, Klaus, Albert Marcet, and Johannes Beutel. 2017. "Stock Price Booms and Expected Capital Gains." *American Economic Review* 107 (8): 2352–408.

Adam, Klaus, Dmitry Matveev, and Stefan Nagel. 2018. "Do Survey Expectations of Stock Returns Reflect Risk-Adjustments?" NBER Working Paper 25122.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2018. "Perception of House Price Risk and Homeownership." NBER Working Paper 25090.

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): 2272–300.

Ameriks, John, Joseph S. Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti. 2015a. "Long-Term Care Utility and Late in Life Saving." NBER Working Paper 20973.

Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti. 2018. "The Long-Term-Care Insurance Puzzle: Modeling and Measurement." NBER Working Paper 22726.

Ameriks, John, Joseph S. Briggs, Andrew Caplin, Minjoon Lee, Matthew D. Shapiro, and Christopher Tonetti. 2017. "Older Americans Would Work Longer If Jobs Were Flexible." NBER Working Paper 24008.

Ameriks, John, Andrew Caplin, Minjoon Lee, Matthew D. Shapiro, and Christopher Tonetti. 2015b. "The Wealth of Wealthholders." NBER Working Paper 20972.

Ameriks, John, Gábor Kézdi, Minjoon Lee, and Matthew D. Shapiro. 2020. "Heterogeneity in Expectations, Risk Tolerance, and Household Stock Shares." *Journal of Business & Economic Statistics* 38 (3): 633–46.

Amromin, Gene, and Steven A. Sharpe. 2014. "From the Horse's Mouth: Economic Conditions and Investor Expectations of Risk and Return." *Management Science* 60 (4): 845–66.

Arellano, Manuel, and Stephen Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *Review of Economic Studies* 58 (2): 277–97.

Armona, Luis, Andreas Fuster, and Basit Zafar. 2016. "Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment." Federal Reserve Bank of New York Staff Report 798

Arnold, Marc, Matthias Pelster, and Marti G. Subrahmanyam. 2018. "Paying Attention." Unpublished.

Arrondel, Luc, Hector Calvo Pardo, and Derya Tas. 2014. "Subjective Return Expectations, Information and Stock Market Participation: Evidence from France." Unpublished.

Azeredo da Silveira, Rava, and Michael Woodford. 2019. "Noisy Memory and Over-Reaction to News." *AEA Papers and Proceedings* 109: 557–61.

Bacchetta, Philippe, Elmar Mertens, and Eric van Wincoop. 2009. "Predictability in Financial Markets: What Do Survey Expectations Tell Us?" *Journal of International Money and Finance* 28 (3): 406–26.

Bachmann, Rüdiger, Kai Carstensen, Stefan Lautenbacher, and Martin Schneider. 2018. "Uncertainty and Change: Survey Evidence of Firms' Subjective Beliefs." Unpublished.

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): 2224–76.

Bailey, Michael, Eduardo Dávila, Theresa Kuchler, and Johannes Stroebel. 2019. "House Price Beliefs and Mortgage Leverage Choice." *Review of Economic Studies* 86 (6): 2403–52.

Barber, Brad M., and Terrance Odean. 2000. "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors." *Journal of Finance* 55 (2): 773–806.

Barber, Brad M., and Terrance Odean. 2001. "Boys Will Be Boys: Gender, Overconfidence, and Common Stock Investment." *Quarterly Journal of Economics* 116 (1): 261–92.

Barber, Brad M., and Terrance Odean. 2013. "The Behavior of Individual Investors." In *Handbook of the Economics of Finance*, Vol. 2A, edited by George M. Constantinides, Milton Harris, and René M. Stulz, 1533–70. Amsterdam: Elsevier.

Barberis, Nicholas. 2018. "Psychology-Based Models of Asset Prices and Trading Volume." In *Handbook of Behavioral Economics: Foundations and Applications*, Vol. 1, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 79–176. Amsterdam: Elsevier.

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): 823–66.

Ben-David, Itzhak, Elyas Fermand, Camelia M. Kuhnen, and Geng Li. 2018. "Expectations Uncertainty and Household Economic Behavior." NBER Working Paper 25336.

Benjamin, Daniel J. 2019. "Errors in Probabilistic Reasoning and Judgment Biases." In *Handbook of Behavioral Economics: Applications and Foundations 1*, Vol. 2, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 69–186. Amsterdam: Elsevier.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian. 2009. "The Importance of Default Options for Retirement Saving Outcomes: Evidence from the United States." In *Social Security Policy in a Changing Environment*, edited by Jeffrey R. Brown, Jeffrey B. Liebman, and David A. Wise, 167–195. Chicago: University of Chicago Press.

Bhandari, Anmol, Jaroslav Borovička, and Paul Ho. 2016. "Identifying Ambiguity Shocks in Business Cycle Models Using Survey Data." NBER Working Paper 22225.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer. 2020. "Expectations of Fundamentals and Stock Market Puzzles." NBER Working Paper 27283.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2018. "Overreaction in Macroeconomic Expectations." NBER Working Paper 24932.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2016. "Understanding Booms and Busts in Housing Markets." *Journal of Political Economy* 124 (4): 1088–1147.

Caballero, Ricardo J., and Alp Simsek. 2017. "A Risk-Centric Model of Demand Recessions and Macroprudential Policy." NBER Working Paper 23614.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

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.

Case, Karl E., Robert J. Shiller, and Anne Thompson. 2012. "What Have They Been Thinking? Home Buyer Behavior in Hot and Cold Markets." NBER Working Paper 18400.

Chen, Hui, Scott Joslin, and Ngoc-Khanh Tran. 2012. "Rare Disasters and Risk Sharing with Heterogeneous Beliefs." *Review of Financial Studies* 25 (7): 2189–2224.

Cheng, Ing-Haw, Sahil Raina, and Wei Xiong. 2014. "Wall Street and the Housing Bubble." *American Economic Review* 104 (9): 2797–829.

Choi, James J., and Adriana Z. Robertson. 2020. "What Matters to Individual Investors? Evidence from the Horse's Mouth." *Journal of Finance* 75 (4): 1965–2020.

Clark, Jeffrey W., Stephen P. Utkus, and Jean A. Young. 2015. *Automatic Enrollment: The Power of the Default*. Valley Forge, PA: Vanguard.

Cochrane, John H. 2011. "Presidential Address: Discount Rates." *Journal of Finance* 66 (4): 1047–108.

Cochrane, John H. 2017. "Macro-Finance." *Review of Finance* 21 (3): 945–85.

Cogent Wealth Reports. 2018. "Investor Brand Builder 2018."

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1): 116–59.

Coibion, Olivier, Yuriy Gorodnichenko, and Rupal Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature* 56 (4): 1447–91.

Coval, Joshua, and Erik Stafford. 2007. "Asset Fire Sales (and Purchases) in Equity Markets." *Journal of Financial Economics* 86 (2): 479–512.

Cummins, Jason G., Kevin A. Hassett, and Stephen D. Oliner. 2006. "Investment Behavior, Observable Expectations, and Internal Funds." *American Economic Review* 96 (3): 796–810.

D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber. 2019. "IQ, Expectations, and Choice." Bank of Finland Research Discussion Paper 2/2019.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam. 1998. "Investor Psychology and Security Market Under- and Overreactions." *Journal of Finance* 53 (6): 1839–85.

Das, Sreyoshi, Camelia M. Kuhnen, and Stefan Nagel. 2020. "Socioeconomic Status and Macroeconomic Expectations." *Review of Financial Studies* 33 (1): 395–432.

De la O, Ricardo, and Sean Myers. 2021. "Subjective Cash Flow and Discount Rate Expectations." *Journal of Finance.* https://doi.org/10.1111/jofi.13016.

DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47 (2): 315–72.

DellaVigna, Stefano, and Joshua M. Pollet. 2009. "Investor Inattention and Friday Earnings Announcements." *Journal of Finance* 64 (2): 709–49.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990. "Noise Trader Risk in Financial Markets." *Journal of Political Economy* 98 (4): 703–38.

De Marco, Filippo, Marco Macchiavelli, and Rosen Valchev. 2018. "International Diversification and Information Heterogeneity." Unpublished.

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): 369–79.

Dominitz, Jeff, and Charles F. Manski. 2011. "Measuring and Interpreting Expectations of Equity Returns." *Journal of Applied Econometrics* 26 (3): 352–70.

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): 378–89.

Duffie, Darrell, and Tong-sheng Sun. 1990. "Transactions Costs and Portfolio Choice in a Discrete-Continuous-Time Setting." *Journal of Economic Dynamics and Control* 14 (1): 35–51.

Fuhrer, Jeffrey C. 2018. "Intrinsic Expectations Persistence: Evidence from Professional and Household Survey Expectations." Federal Reserve Bank of Boston Working Paper 18-9.

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 Ralph Koijen. 2020. "In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis." Unpublished.

Gabaix, Xavier, and David Laibson. 2001. "The 6D Bias and the Equity-Premium Puzzle." In *NBER Macroeconomics Annual 2001*, Vol. 16, edited by Ben S. Bernanke and Kenneth Rogoff, 257–312. Cambridge, MA: MIT Press.

García-Schmidt, Mariana, and Michael Woodford. 2019. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." *American Economic Review* 109 (1): 86–120.

Geanakoplos, John. 2010. "The Leverage Cycle." *NBER Macroeconomics Annual 2009*, Vol. 24, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 1–66. Chicago: University of Chicago Press.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. 2016. "Expectations and Investment." In *NBER Macroeconomics Annual 2015,* Vol. 30, edited by Martin Eichenbaum and Jonathan A. Parker, 379–431. Chicago: University of Chicago Press.

Gervais, Simon, and Terrance Odean. 2001. "Learning to Be Overconfident." *Review of Financial Studies* 14 (1): 1–27.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2020. "Inside the Mind of a Stock Market Crash." NBER Working Paper 27272.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021. "Replication Data for: Five Facts about Beliefs and Portfolios." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E123761V1.

Gillen, Ben, Erik Snowberg, and Leeat Yariv. 2019. "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study." *Journal of Political Economy* 127 (4): 1826–63.

Glaeser, Edward L., and Charles G. Nathanson. 2017. “An Extrapolative Model of House Price Dynamics.” *Journal of Financial Economics* 126 (1): 147–70.

Goetzmann, William, Dasol Kim, and Robert Shiller. 2018. “Availability, Media, Earthquakes, and Lottery Winners: Determinants of Crash Beliefs from Investor Surveys.” Unpublished.

Graham, John R. 1999. “Herding among Investment Newsletters: Theory and Evidence.” *Journal of Finance* 54 (1): 237–68.

Greenwood, Robin, and Andrei Shleifer. 2014. “Expectations of Returns and Expected Returns.” *Review of Financial Studies* 27 (3): 714–46.

Grinblatt, Mark, and Matti Keloharju. 2009. “Sensation Seeking, Overconfidence, and Trading Activity.” *Journal of Finance* 64 (2): 549–78.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2018. “Time Varying Risk Aversion.” *Journal of Financial Economics* 128 (3): 403–21.

Harrison, J. Michael, and David M. Kreps. 1978. “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations.” *Quarterly Journal of Economics* 92 (2): 323–36.

Hoffmann, Arvid O. I., and Thomas Post. 2016. “How Does Investor Confidence Lead to Trading? Linking Investor Return Experiences, Confidence, and Investment Beliefs.” *Journal of Behavioral and Experimental Finance* 12: 65–78.

Hong, Harrison, and Jeremy C. Stein. 1999. “A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets.” *Journal of Finance* 54 (6): 2143–84.

Hong, Harrison, and Jeremy C. Stein. 2007. “Disagreement and the Stock Market.” *Journal of Economic Perspectives* 21 (2): 109–28.

Hudomiet, Peter, Gabor Kezdi, and Robert J. Willis. 2011. “Stock Market Crash and Expectations of American Households.” *Journal of Applied Econometrics* 26 (3): 393–415.

Hurd, Michael, Maarten Van Rooij, and Joachim Winter. 2011. “Stock Market Expectations of Dutch Households.” *Journal of Applied Econometrics* 26 (3): 416–36.

Jin, Lawrence. 2015. “A Speculative Asset Pricing Model of Financial Instability.” Unpublished.

Kempf, Elisabeth, and Margarita Tsoutsoura. 2018. “Partisan Professionals: Evidence from Credit Rating Analysts.” NBER Working Paper 25292.

Kézdi, Gábor, and Robert J. Willis. 2009. “Stock Market Expectations and Portfolio Choice of American Households.” Unpublished.

Kézdi, Gábor, and Robert J. Willis. 2011. “Household Stock Market Beliefs and Learning.” NBER Working Paper 17614.

Kuchler, Theresa, and Basit Zafar. 2019. “Personal Experiences and Expectations about Aggregate Outcomes.” *Journal of Finance* 74 (5): 2491–542.

Kyle, Albert S., and F. Albert Wang. 1997. “Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test?” *Journal of Finance* 52 (5): 2073–90.

Landier, Augustin, Yueran Ma, and David Thesmar. 2017. “New Experimental Evidence on Expectations Formation.” CEPR Working Paper DP12527

Madrian, Brigitte C., and Dennis F. Shea. 2001. “The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior.” *Quarterly Journal of Economics* 116 (4): 1149–87.

Malmendier, Ulrike, and Stefan Nagel. 2011. “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?” *Quarterly Journal of Economics* 126 (1): 373–416.

Manski, Charles F. 2004. “Measuring Expectations.” *Econometrica* 72 (5): 1329–76.

Martin, Ian, and Dimitris Papadimitriou. 2019. “Sentiment and Speculation in a Market With Heterogeneous Beliefs.” CEPR Working Paper DP13857.

Meeuwis, Maarten, Jonathan A. Parker, Antoinette Schoar, and Duncan I. Simester. 2018. “Belief Disagreement and Portfolio Choice.” NBER Working Paper 25108.

Merkle, Christoph, and Martin Weber. 2014. “Do Investors Put Their Money Where Their Mouth Is? Stock Market Expectations and Investing Behavior.” *Journal of Banking and Finance* 46: 372–86.

Merton, Robert C. 1969. “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case.” *Review of Economics and Statistics* 51 (3): 247–57.

Mian, Atif, Amir Sufi, and Nasim Khoshkhou. 2015. “Government Economic Policy, Sentiments, and Consumption.” NBER Working Paper 21316.

Odean, Terrance. 1999. “Do Investors Trade Too Much?” *American Economic Review* 89 (5): 1279–98.

Ottaviani, Marco, and Peter Norman Sorensen. 2006. “The Strategy of Professional Forecasting.” *Journal of Financial Economics* 81 (2): 441–66.

Ouimet, Paige, and Geoffrey A. Tate. 2017. “Attention for the Inattentive: Positive Effects of Negative Financial Shocks.” Unpublished.

Peng, Lin, and Wei Xiong. 2006. “Investor Attention, Overconfidence and Category Learning.” *Journal of Financial Economics* 80 (3): 563–602.

Piazzesi, Monika, and Martin Schneider. 2009. "Momentum Traders in the Housing Market: Survey Evidence and a Search Model." *American Economic Review* 99 (2): 406–11.

Rangvid, Jesper, Maik Schmeling, and Andreas Schrimpf. 2013. "What Do Professional Forecasters' Stock Market Expectations Tell Us about Herding, Information Extraction and Beauty Contests?" *Journal of Empirical Finance* 20: 109–29.

Rietz, Thomas A. 1988. "The Equity Risk Premium: A Solution." *Journal of Monetary Economics* 22 (1): 117–31.

Scheinkman, Jose A., and Wei Xiong. 2003. "Overconfidence and Speculative Bubbles." *Journal of Political Economy* 111 (6): 1183–1220.

Vissing-Jorgensen, Annette. 2003. "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear with Wealth? Evidence from Expectations and Actions." In *NBER Macroeconomics Annual,* Vol. 18, edited by Mark Gertler and Kenneth Rogoff, 139–94. Cambridge, MA: MIT Press.

Wachter, Jessica A. 2013. "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance* 68 (3): 987–1035.