

**Beliefs about the Stock Market
and Investment Choices:
Evidence from a Survey and
a Field Experiment**

Christine Laudenbach, Annika Weber, Rüdiger Weber, Johannes Wohlfart

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Beliefs about the Stock Market and Investment Choices: Evidence from a Survey and a Field Experiment

Abstract

We survey retail investors at an online bank to study how beliefs about the autocorrelation of aggregate stock returns shape investment decisions measured in administrative account data. Individuals' beliefs exhibit substantial heterogeneity and predict trading responses to market movements. We inform half of our respondents that, historically, the autocorrelation was close to zero, which causes them to update their perceived current autocorrelation and return expectations. The treatment shifts respondents' equity purchases during the Covid-19 crash months later in the direction implied by the intervention. Our results provide causal evidence on the drivers of disagreement and trade in asset markets.

JEL-Codes: D830, D840, D910, E710, G110, G120, G410, G500.

Keywords: expectation formation, individual investors, trading, field experiment.

Christine Laudenbach
Department of Finance, SAFE
Goethe University Frankfurt / Germany
laudenbach@safe-frankfurt.de

Annika Weber
House of Finance
Goethe University Frankfurt / Germany
annika.weber@hof.uni-frankfurt.de

Rüdiger Weber
WU Vienna and VGSF
Vienna / Austria
ruweber@wu.ac.at

Johannes Wohlfart
Department of Economics and CEBI
University of Copenhagen / Denmark
johannes.wohlfart@econ.ku.dk

1 Introduction

Expectations about asset returns are central in macroeconomics and finance. They shape portfolio choices and saving behavior, influence asset prices, and ultimately guide the allocation of scarce capital resources in the economy. In the context of aggregate stock returns, households' expectations often deviate from what is implied by standard theories. For instance, there are important deviations from rational expectations (Adam, Marcet and Beutel, 2017; Malmendier and Nagel, 2011) and a substantial amount of disagreement across households, which is reflected in heterogeneity in portfolio decisions (Giglio, Maggiori, Stroebel and Utkus, 2021a). Based on these findings, a recent literature incorporates more realistic belief formation mechanisms into macro-finance models, which has important implications for model predictions about both aggregate and individual outcomes (Adam and Nagel, 2022). However, to date there is limited *causal* evidence on (i) what is driving disagreement about expected returns and (ii) how return expectations affect investors' trading decisions.

In this paper, we study these two questions using a field experiment with retail investors. We propose that heterogeneous beliefs about the predictiveness of specific state variables for future returns are an important driver of disagreement about expected returns. Heterogeneous beliefs about predictiveness could arise for different reasons, such as investors' reliance on different subjective models (Andre, Pizzinelli, Roth and Wohlfart, 2022a), the use of different heuristics (Barberis, Greenwood, Jin and Shleifer, 2015), or investors forming their forecasting rules in light of different experiences (Malmendier and Nagel, 2011). In particular, we focus on beliefs about return predictability based on recent realized returns, that is, the perceived autocorrelation of aggregate stock returns. Heterogeneity in these beliefs could lead to differences in how information about new return realizations is processed, resulting in disagreement about future expected returns and trading. Our focus is motivated by previous literature suggesting that recent return realizations are central to households' beliefs about future returns (Dominitz and Manski, 2011; Greenwood and Shleifer, 2014; Vissing-Jorgensen, 2003). However, we also explore the idea that individuals hold divergent views on the predictiveness of other state variables, such as valuation ratios, which might contribute to disagreement about expected returns.

We conduct a survey with about 2,000 stockholders that are clients of a major German online bank. In our survey, we elicit respondents' beliefs about the historical autocorrelation of stock returns using a new,

individual-level measure. Specifically, we first ask respondents to think of six bins of historical annual return realizations of the German stock market index (DAX) during the last 50 years. For each return bin, respondents are asked to provide an estimate of the average return of the DAX over the subsequent 12 months when the return over the previous 12 months fell into the respective bin. Subsequently, a random half of the respondents are informed about the actual historical conditional mean return over the following year for each of the six bins. Actual conditional mean returns in the six bins vary only narrowly around the unconditional historical average return of the DAX of 8.5%, illustrating that, historically, the predictive power of recent 12-month returns for future 12-month returns has been very low.¹ In both the main survey and a four-week follow-up survey, we then measure our respondents' posterior beliefs about the autocorrelation of returns and the expected return over the 12 months after the survey.

Our information intervention generates exogenous variation in beliefs about the autocorrelation of aggregate returns. This allows us to examine whether heterogeneity in these beliefs is a causal driver of disagreement in return expectations. Moreover, by linking the treatment variation with administrative account data on trading decisions before and after the intervention, we obtain causal evidence on the role of individuals' subjective beliefs in their investment choices. Due to the randomized nature of our intervention, our experimental evidence is immune to concerns related to omitted variables or reverse causality, which could confound correlational estimates of the relationship between beliefs and economic decisions.

We document four main sets of results. First, we provide descriptive evidence on our respondents' prior beliefs. There is a large degree of heterogeneity in beliefs about the historical autocorrelation of stock returns. Respondents disagree strongly about how differences in realized returns between scenarios are associated with different year-ahead returns. More than half of the respondents believe in a negative autocorrelation – i.e., in mean reversion – and about one fourth believe in a positive autocorrelation – i.e., in persistence of returns. Respondents' beliefs about the autocorrelation of returns are predictive of their expectations about the 12-month-ahead return at the time of the survey. This suggests that part of the disagreement in return expectations in the stock market is due to investors processing information about

¹This is in line with other evidence suggesting that, empirically, the autocorrelation of aggregate returns is close to zero at the annual horizon (Fama and French, 1988; Huang, Li, Wang and Zhou, 2020; Nagel and Xu, 2022b). By contrast, other variables such as valuation ratios have been shown to be predictive of future aggregate returns (Cochrane, 2008, 2011).

the same return realization differently depending on their heterogeneous beliefs about predictiveness.

The finding of a prevalence of beliefs in mean reversion contrasts with a number of earlier studies documenting extrapolative belief formation in the context of stock returns (Amromin and Sharpe, 2013; Giglio, Maggiori, Stroebel and Utkus, 2021b; Greenwood and Shleifer, 2014; Vissing-Jorgensen, 2003). In contrast to these studies, which exploit time series variation in realized returns and expectations, we measure investors' beliefs by asking them directly to think about the past autocorrelation. Our elicitation method may yield different results than these earlier studies for at least two reasons: First, even investors who believe in a negative autocorrelation of returns may hold extrapolative expectations in specific episodes, e.g., because the episode is accompanied by a specific narrative (Andre, Haaland, Roth and Wohlfart, 2022b). For instance, the studies by Greenwood and Shleifer (2014) and Vissing-Jorgensen (2003) largely focus on the build-up and burst of the dot-com bubble in the early 2000s, which may have featured the narrative that "this time is different". Second, our mode of elicitation holds constant beliefs about the return in the previous period, which in other studies are unobserved and may deviate from actual return realizations.

More generally, our first result highlights that beliefs about the autocorrelation of returns are highly heterogeneous. The exact fractions of investors believing in mean reversion or persistence could vary across different groups of economic agents and over time. Indeed, using a longer sample period than previous literature, Nagel and Xu (2022b) document extrapolative updating behavior in response to realized returns in a US household sample but contrarian updating behavior among US financial professionals. Other studies find that beliefs in persistence, beliefs in mean reversion, and beliefs that returns have negligible autocorrelation all seem to be fairly common among households (Dominitz and Manski, 2011; Heiss, Hurd, van Rooij, Rossmann and Winter, 2022; Luo, Ravina, Sammon and Viceira, 2022; von Gaudecker and Wogroly, 2022). Using auxiliary data collections, we show that beliefs in a negative autocorrelation are less prevalent in samples from the general population than among retail investors. First, this confirms that our direct elicitation method is in principle able to pick up both types of beliefs. Second, this underscores that beliefs about the statistical properties of returns are highly heterogeneous across different types of economic agents, where our retail investors seem to be closer on the spectrum to, e.g., the professionals in Nagel and Xu (2022b) than to general population samples.

Second, we examine the effect of the information intervention on individuals' beliefs. The treatment strongly reduces respondents' beliefs in predictability of future returns based on recent returns, as measured by their agreement with a set of verbal statements. Our treatment completely closes the gaps in these beliefs between those perceiving a positive or negative historical relationship between past and future returns and those perceiving an autocorrelation close to zero before the intervention. Moreover, the information induces respondents to update their expectations about the return of the German stock market over the 12 months after the survey in line with the change in their perceived autocorrelation, reducing expectation dispersion. Most of the experimentally induced changes in beliefs persist in a follow-up survey conducted four weeks after the intervention, mitigating concerns related to experimenter demand effects or numerical anchoring (Cavallo, Cruces and Perez-Truglia, 2017; Haaland, Roth and Wohlfart, 2023). Taken together, these findings highlight that investors' beliefs about the dynamics of returns are elastic and can be changed through factual information. Moreover, these results imply that heterogeneity in the perceived autocorrelation of returns is a causal driver of disagreement in stock return expectations.

Third, we study whether respondents' perceived autocorrelation of returns matches the timing of their equity transactions at the online bank over the five years before the intervention. For each month of our sample period, we predict each respondent's return expectation based on the realized return over the preceding 12 months and the respondent's belief about the autocorrelation. Respondents whose predicted return expectations over a given month increase exhibit significantly higher net purchases of equity during that month. These trading adjustments mostly reflect increases in gross purchases, while gross sales are less responsive to belief changes. Moreover, the effects on net buying are fully driven by the intensive margin of trading: predicted belief changes are correlated with net purchases of equity conditional on trading, but are unrelated to whether or not an investor trades. These patterns are in line with recent evidence on the relationship between trading and return expectations measured in repeated surveys (Giglio et al., 2021a). These findings validate our survey measure and suggest that beliefs about the autocorrelation of returns play a key role in retail investors' decisions.

Fourth, we study whether respondents adjust their trading behavior at the online bank in response to the information. Four to five months after our intervention, realized returns dropped from about 17% to about -23% during the Covid-19 crash. Our treatment significantly shifts respondents' net purchases of

equity during the downturn in the direction of the belief changes caused by the intervention. These effects are mostly driven by changes in the gross purchase volume, while the gross sales volume is less affected, in line with our correlational evidence. Different to the correlational results, the treatment affects trading both through the extensive and the intensive margin. This may reflect the exceptional size and speed of the Covid-19 downturn, which may have increased investor attention. Our treatment more than offsets differences in trading responses to the crash across individuals with different priors in the control group. We detect no systematic effects on respondents' trading decisions during the time before the crash – a period without major stock market movements. These findings indicate that retail investors' beliefs about the dynamics of returns causally affect their trading responses to stock market fluctuations.

Lastly, we present additional evidence suggesting that memory and personal experiences (Bordalo, Gennaioli and Shleifer, 2022; Enke, Schwerter and Zimmermann, 2022; Malmendier and Nagel, 2011) seem to play an important role in shaping investors' perceived autocorrelation. Specifically, when respondents think about the past autocorrelation of returns, they recall memories of salient stock market events, many of which involved return reversals. Respondents who recall episodes of return reversals are in turn more likely to believe in a negative autocorrelation of returns.

Taken together, our findings highlight substantial heterogeneity in beliefs about the predictiveness of recent returns for future returns. This heterogeneity causally contributes to the previously documented disagreement in stock return expectations across investors (Giglio et al., 2021a). Retail investors' beliefs about the dynamic properties of returns also causally affect their investment choices. Heterogeneity in these beliefs is thus potentially an important driver of trade in asset markets. Our findings lend support to a class of models in which trade arises due to differences in how investors evaluate the same piece of information (Banerjee and Kremer, 2010; Dumas, Kurshev and Uppal, 2009; Harris and Raviv, 1993; Scheinkman and Xiong, 2003). In particular, our results support theories where one class of investors extrapolates recent returns, while another class believes in a negative relationship between recent and future returns (Barberis et al., 2015; Barberis, Greenwood, Jin and Shleifer, 2018; Cutler, Poterba and Summers, 1991; De Long, Shleifer, Summers and Waldmann, 1990). In line with such models, our findings suggest that extrapolators' higher equity demand following high return realizations is at least partly accommodated by retail investors believing in a negative autocorrelation. While we

focus on recent return realizations, our findings likely extend to other state variables that individuals may use to form their return expectations. In the appendix we report results from an additional data collection highlighting that many of our findings on expectation formation carry over to beliefs about the predictiveness of valuation ratios for future returns. At a more general level, our findings are in line with a recent literature documenting heterogeneity in individuals' perceptions of macroeconomic relationships in different contexts (Andrade, Crump, Eusepi and Moench, 2016; Andre et al., 2022a,b; Armona, Fuster and Zafar, 2019), pointing to the importance of accounting for such heterogeneity in empirical and theoretical research in macroeconomics and finance.

A number of papers have used survey data to study the link between stock market beliefs and investor behavior (Ameriks, Kézdi, Lee and Shapiro, 2020; Amromin and Sharpe, 2013; Arrondel, Calvo-Pardo, Giannitsarou and Haliassos, 2022; Beutel and Weber, 2023; Das, Kuhnen and Nagel, 2020; Dominitz and Manski, 2007; Drerup, Enke and Von Gaudecker, 2017; Kézdi and Willis, 2011; Malmendier and Nagel, 2011). Only few studies have linked survey data on beliefs with administrative account data on investment decisions (Andersen, Hanspal, Martinez-Correa and Nielsen, 2021; Hoffmann, Post and Pennings, 2015; Merkle and Weber, 2014). For instance, Giglio et al. (2021a) show that the return expectations of Vanguard clients measured in repeated surveys correlate with their investment decisions as measured in administrative data, but the relationship is an order of magnitude smaller than what is implied by standard models. We confirm this finding using both correlational and experimental variation in beliefs. Giglio et al. (2021b) study the joint dynamics of stock return expectations and trading decisions of Vanguard clients during the Covid-19 crash. Meeuwis, Parker, Schoar and Simester (2022) study relative changes in investment decisions among households with different party affiliation around the 2016 presidential election, which are most likely driven by changes in beliefs. We add to this literature (i) by showing that not only investors' return expectations but also their subjective beliefs about the underlying data-generating process are meaningfully linked to their investment choices, highlighting the need to better understand investors' mental models of financial markets; and (ii) by exploiting experimental variation to study the causal effect of beliefs on investment decisions in the field.

Our paper also contributes to a literature that uses information provision experiments to study macroeconomic expectation formation of households (Armantier, Nelson, Topa, van der Klaauw and

Zafar, 2016; Binder and Rodrigue, 2018; Cavallo et al., 2017; Coibion, Gorodnichenko and Weber, 2022; D’Acunto, Fuster and Weber, 2022; Hanspal, Weber and Wohlfart, 2021; Roth and Wohlfart, 2020).² Complementary to our paper, Beutel and Weber (2023) examine how households’ return expectations and investment behavior respond to information about the realization of specific variables, e.g., the return over the past 12 months. Respondents’ updating about future returns then provides indirect evidence on the average respondent’s belief about the underlying data generating process. By contrast, we directly measure each respondent’s belief about the relationship between past and future stock returns and shift this belief by providing information.³ This allows us to provide causal evidence on how beliefs about the data generating process affect expectation formation and decisions. In contrast to our result of a widespread belief in mean reversion, Beutel and Weber (2023) find that respondents on average upward revise their expectations when learning about a high recent return, indirectly pointing to a belief in return persistence. This difference in findings could be due to their focus on a broader sample from the general population. Indeed, our auxiliary data collections point to a higher prevalence of beliefs in persistence in the general population than among retail investors. While Beutel and Weber (2023) take a more comprehensive view than our paper – focusing on a broader set of variables and studying how investors acquire information – they measure investment choices in an investment game included in their survey. To the best of our knowledge, ours is the first study to link an information experiment with administrative data on actual trading decisions.⁴ From a methodological perspective, our results highlight that simple information interventions can change economic decisions months later, demonstrating the relevance of information provision experiments as a method to study belief formation and the effects of beliefs on real-world behaviors. In follow-up work, Haaland and Næss (2023) use a similar setup as ours to demonstrate that mis-perceived returns to active investing shape the decisions of investors on a social trading platform.

Our paper also relates to recent work from laboratory experiments (Charles, Frydman and Kilic, 2022;

²See Haaland et al. (2023) for a review of the literature using information provision experiments.

³Armona et al. (2019) use a similar approach as Beutel and Weber (2023) to show that US households under-estimate the long-run mean reversion in local house prices.

⁴A few other papers link information experiments shifting macroeconomic expectations with non-self-reported data on decisions, e.g., Coibion et al. (2022), Coibion, Georgarakos, Gorodnichenko and Weber (2023) and Galashin, Kanz and Perez-Truglia (2022) for inflation expectations and spending, and Bottan and Perez-Truglia (2022) in the context of home selling decisions.

Kuhnen, Rudolf and Weber, 2017). For instance, Andries, Bianchi, Huynh and Pouget (2020) document that naive extrapolation of recent returns is most common when individuals believe that they have no informative signal about future returns available, but that in such cases individuals strongly discount their return expectations when making investment decisions. Consistent with their findings, in our sample of relatively experienced retail investors, who may plausibly believe to have ample relevant information available, beliefs in mean reversion are more prevalent than extrapolative behavior.

2 Experimental design and data

2.1 Main survey

Our main experiment consists of three stages: (i) a baseline stage eliciting respondents' prior beliefs; (ii) a treatment stage in which respondents receive information; and (iii) a final stage eliciting posterior beliefs and a set of background characteristics. Appendix B provides the survey instructions translated to English.

Baseline stage: Prior beliefs We first elicit respondents' point beliefs about the return of the German stock market index (DAX) over the 12 months before the survey and over the 12 months after the survey. To measure respondents' prior beliefs about the autocorrelation of returns, we then present participants with six different intervals of 12-month return realizations of the DAX, which are mutually exclusive and collectively exhaustive.⁵ For each interval, starting from the lowest one, we instruct respondents to think of all points in time over the past 50 years at which the return of the DAX over the preceding 12 months had fallen into the respective interval, and ask them to estimate the average return of the DAX over the subsequent 12 months for these cases. For each return interval, the prediction is elicited on a separate screen. A graph displays respondents' estimates for the current and for all previous scenarios in real time as blue bars. Figure 1 Panel A displays an example survey screen after forecasts for all six scenarios have been submitted. If a respondent believes that, historically, the return of the stock market was predictable by the realized return over the previous year, this should show up as non-constant estimates entered across the six bins. For instance, an upward sloping pattern in the displayed graph

⁵The intervals are "less than -20%", "between -20% and -10%", "between -10% and 0%", "between 0% and 10%", "between 10% and 20%", and "above 20%".

would indicate a belief in a positive autocorrelation, while a downward sloping pattern would indicate a belief in a negative autocorrelation.⁶ An advantage of this way of eliciting beliefs is that it requires no knowledge of statistical concepts beyond averages from respondents.

Treatment stage: Non-informativeness of recent return realizations In the second stage, a random half of the respondents are assigned to the treatment group, while the other half are assigned to the control group. Respondents in the treatment group receive information on the actual historical average realizations of 12-month-ahead returns for each of the six bins of returns over the preceding 12 months. The actual mean realized returns displayed to respondents vary narrowly between 7.4% and 9.5% across the six intervals, and there is no clear monotonic relationship between past 12-month returns and returns over the next 12 months. This illustrates that, historically, recent realized returns of the DAX have not been informative for future returns at the one-year horizon.⁷

The treatment information is communicated as follows: respondents are again shown the graph displaying their prior estimates, illustrated as six blue bars. Respondents have to repeatedly click on a button and learn about the actual historical mean return realization over the next 12 months one-by-one for each interval of previous returns. Actual historical values are displayed as orange bars next to participants' priors. In addition, for each case, a written sentence is displayed above the graph that reminds participants of their corresponding prior and informs them about the actual historical average return realization. Figure 1 Panel B displays an example screen of the treatment graph once all six actual realizations are displayed. On the next screen, respondents are again shown the complete graph with both their own estimates and the actual historical values. In addition, we provide them with a short text summarizing the content of the treatment.⁸

⁶In line with previous literature (Amromin and Sharpe, 2013; Dominitz and Manski, 2011; Giglio et al., 2021a), we focus on nominal stock returns to make our survey questions easy to understand for participants. One concern could be that longer-run trends in the risk-free rate or inflation could give rise to a positive autocorrelation of nominal stock returns. However, as can be seen in Figure 1, this is not the case for the German stock market, and also does not seem to hold more generally (Fama and French, 1988).

⁷We use data on total returns on the German stock index DAX provided by Thomson Reuters/Eikon Datastream (item DAXINDX). The DAX is a performance index for the largest German stocks with index levels capturing returns from both dividends and capital gains. For the time before 1988, when the DAX was established, Thomson Reuters computes a synthetic DAX time series. Returns are calculated on a weekly basis by comparing average index levels 52 weeks before and after the respective week. One of the values, a next-year average return of 9.58% for previous returns between -10% and 0%, was accidentally rounded to 9.5% instead of 9.6% in the information displayed to respondents.

⁸Among others, this summary contains a sentence stating that “*any return predictability would be quickly exploited and removed by large institutional investors*”. Given that time-varying risk aversion or changes in risk perceptions could in principle

Control group respondents receive no information on how informative returns over the previous 12 months were for returns over the next 12 months, but are merely informed about the unconditional historical average annual return on the DAX over the last 50 years of 8.5%. We provide this information to the control group because respondents' beliefs about the unconditional average historical return could also be affected by the treatment. By comparing respondents in the treatment and the control group we can therefore identify the effect of a change in beliefs about how future returns correlate with past realizations, holding constant beliefs about the unconditional average return.

Final stage: Posterior beliefs After the information treatment, we elicit participants' agreement on seven-point categorical scales reaching from "Strongly disagree" to "Strongly agree" with verbal statements describing different patterns of autocorrelation of aggregate stock returns. For instance, one of these statements reads "*When the stock market has recently increased it is more likely that stock returns will be positive over the following time than when the stock market has recently decreased.*", which is designed to capture a belief in persistence. All of these statements aim to measure respondents' beliefs about the *current* autocorrelation of returns. We abstain from a quantitative elicitation of this object (i) to avoid survey fatigue among respondents and (ii) to circumvent concerns related to numerical anchoring on the treatment information about the historical autocorrelation. We then re-elicite respondents' expectations about the 12-month-ahead return of the DAX, both as a point forecast (as for the prior) and as a subjective probability distribution. At the end of the survey, participants answer a series of questions on their financial behavior and background characteristics.

Discussion of the experimental design With our experiment we aim to study how beliefs about the predictiveness of a specific state variable for future returns affect expectation formation and trading. We focus on the recent realized return as predictor variable for three main reasons. First, recent returns seem to play a central role in households' return expectations and trading behavior. For instance, previous literature documents a tendency to extrapolate recent returns (Greenwood and Shleifer, 2014; Vissing-

give rise to a form of return predictability that would not be exploited, one could consider this sentence as problematic. However, the statement appears in the specific context of the autocorrelation of returns, so it likely does not appear as a general negation of any type of predictability in the stock market. Moreover, the sentence is less salient than other aspects of the information treatment, such as the dynamic figure, and therefore likely does not play a major role in driving treatment effects. Consistent with this intuition, Appendix A.7 describes a complementary experiment in which the treatment does not contain a statement of this kind, which yields qualitatively similar results as our main experiment.

Jorgensen, 2003) or beliefs in mean reversion among subsets of investors (Dominitz and Manski, 2011; Heiss et al., 2022). Similarly, households' and experts' average return expectations strongly co-move with recent return realizations, but are largely uncorrelated with business cycle variables or aggregate asset valuation measures (Nagel and Xu, 2022b). Second, in additional surveys conducted with clients of the same bank, respondents report higher levels of familiarity with returns than with other financial concepts, such as volatility or valuation ratios (Appendix Figure A.4), and higher levels of confidence in their beliefs about the recent return (64% are “confident” or “very confident”) than in their beliefs about the current price-dividend ratio (33% are “confident” or “very confident”).⁹ Focusing on the recent return as predictor variable should thus facilitate a relatively straightforward elicitation of beliefs, providing us with meaningful belief data. Lastly, focusing on the autocorrelation of annual returns offers an empirical benchmark of an autocorrelation close to zero – which has been documented for the US (Fama and French, 1988; Nagel and Xu, 2022b) and which also applies in our German setting (see Appendix A.2).¹⁰

In our experiment, we elicit our respondents' beliefs about the autocorrelation of returns during the past 50 years. This allows us to compare beliefs to a factual benchmark and to shock these beliefs by providing information on the muted relationship between recent and future returns over this time period. The resulting variation in investors' beliefs about return predictability allows us to study the causal effect of these beliefs on investors' return expectations and trading decisions. Naturally, the current, forward-looking relationship between recent and future returns may differ to some extent from its historical counterpart. Our empirical strategy does not require our respondents to fully adopt the view that there is a stable lack of predictiveness of recent return realizations for future returns, and allows respondents to believe that the relationship is time-varying. For our intervention to generate exogenous variation in beliefs we merely require respondents to consider the information on the historical empirical autocorrelation to be somewhat relevant for the dynamics of stock returns at the time of our survey and afterwards.

⁹Appendix Table A.1 provides an overview of our different data collections.

¹⁰The empirical benchmark is less clear for other state variables. For instance, while the price-dividend ratio is a predictor of future returns in the US, the price-dividend ratio of the German stock market mostly predicts variation in future dividend growth rather than returns (Appendix A.2 and Rangvid, Schmeling and Schrimpf (2014)). Nevertheless, in Appendix A.7 we use a complementary experiment to show that many of our take-aways carry over to beliefs about return predictability based on valuation ratios.

2.2 Follow-up survey

We invite respondents to a follow-up survey approximately four weeks after they completed the main survey. We choose a four-week gap to trade off between testing for persistence of treatment effects and maximizing the re-contact rate and therefore statistical power in the follow-up. The follow-up survey starts by re-eliciting respondents' beliefs about past and future 12-month returns and the historical autocorrelation.

Our prior before running the experiment was that changes in beliefs would fade quickly, as it is suggested by previous literature (Haaland et al., 2023). To achieve a more persistent first-stage effect on respondents' beliefs, which we could then use to study the causal effect of beliefs on trading decisions, we therefore decided to use the follow-up survey to again present respondents in the treatment group with the information. We do so *after* the block of questions measuring persistence of changes in beliefs caused by the initial intervention, and use the same treatment design as in the main survey. Participants in the control group are again informed about the unconditional historical average return of the DAX. After that, we elicit respondents' agreement with two additional verbal statements describing patterns of autocorrelation of stock returns.

2.3 Background and survey administration

We administered the survey to clients of a German online bank in September and October 2019 in a relatively stable market environment. The return of the DAX over the 12 months before the sample period averaged 1.1% and the DAX increased steadily after our survey until the Covid-19 crash in February 2020 (see Appendix Figure A.3).

The bank is among the top five online banks in terms of market share in Germany as measured by the number of clients. The bank provides full bank services offering savings and credit products in addition to its brokerage entity, and is hence used as principal bank by many clients. In a different data collection based on the same sample selection procedure at the same bank, 71% of respondents state that the portfolio they hold with the bank is their main investment account. Clients at the bank trade financial securities online in a self-directed manner. The broker does not offer any financial advice to these clients. This is important as an intermediary would likely reduce the direct impact of an individual's subjective

expectations on her choices.

We sent e-mail invitations to 14,000 individuals randomly selected from the bank's client pool. To eliminate deserted accounts, we only invited clients that had conducted at least one equity transaction with the broker over the previous 12 months. We offered invitees 10 EUR for completing the main survey and 5 EUR for the shorter follow-up. All payoffs were paid in the form of online shopping vouchers and distributed via e-mail.

Overall, 2,083 individuals completed the main survey. This corresponds to a relatively high response rate of 14.9%. 80.9% of respondents agreed to be invited for a second survey when asked at the end of the main survey. 987 investors ultimately completed the follow-up, corresponding to a re-contact rate of 58.5% among those who got invited to the follow-up. At the median, recontacted respondents completed the follow-up 26 days after the main survey. The mean (median) response time was 22.1 minutes (17.8 minutes) for the main survey and 14.2 minutes (10.2 minutes) for the follow-up. Appendix Figure A.1 displays the distribution of response times for both surveys.

We conducted several auxiliary data collections with additional samples, which we introduce when relevant throughout the paper. Appendix Table A.1 provides an overview.

2.4 Data

Sample definition We take two steps to screen out participants who likely did not take the survey seriously or just quickly "clicked" through the questions to obtain the shopping voucher. First, we follow a similar procedure as Armona et al. (2019) and drop respondents who in the main survey report prior or posterior point expectations about the return of the DAX over the 12 months after the survey lower than -20% or higher than 20%, roughly corresponding to the first and 98th percentiles of the response distributions. This step also ensures that our OLS estimates are not driven by outliers. Second, we drop participants who take less than 8 minutes or more than 60 minutes to complete the main survey. These steps leave us with 1,961 respondents in our baseline sample for the main survey, out of which 903 respondents form the follow-up sample.¹¹ Our results are robust to varying the cutoffs for distributions of point forecasts or for response time used to define the sample.

¹¹The follow-up sample excludes participants that are not part of the baseline sample used in the main survey and those who report expectations in the follow-up survey outside the interval $[-20\%; 20\%]$.

Sample characteristics Columns 2-6 of Table 1 display summary statistics of our sample. Column 1 shows population benchmarks from stockholders in the 2017 wave of the Bundesbank's Panel of Household Finances (PHF). Our respondents are predominantly male (84% vs 51% in the population). The average age is 45.2 years (50.6 years in the population). Sample participants are relatively highly educated, with 54% holding a university degree (36% in the population). Our respondents report an average net monthly household income of 3,914 EUR and a net household wealth of 300,488 EUR, fairly similar to the population of stockholders.

Administrative account data We obtain data on our respondents' month-end holdings and daily executed purchases and sales of securities from December 2014 until March 2020. Investors in our sample on average hold financial wealth of 55,272 EUR with the sample bank, of which 39,405 EUR are invested in equity (including direct holdings and holdings through mutual funds). Throughout our analysis, we focus on transactions in equity. The average number of equity trades per month is 1.73. Our sample is therefore comparable to the one used in Giglio et al. (2021a) in terms of trading activity.¹²

Selection into the survey We also use the administrative account data to examine which clients of the online bank select themselves into our survey. We do not have access to administrative account data for all the 14,000 clients who were invited to participate. However, we have access to a sample of 3,701 clients that were randomly selected from the bank's client pool based on the same criteria as the 14,000 invited clients (at least one equity transaction with the broker over the previous 12 months). Appendix Table A.3 compares our main survey sample with this random sample. As it is common in surveys, participation in our experiment is correlated with investors' characteristics. Investors in our sample are less wealthy and trade less often compared to the average client at the bank. In addition, our respondents are somewhat younger and more likely to be male and to be employed. They are very similar to the random sample in terms of their equity share held at the bank and in terms of their risk attitude as measured by the bank. Appendix Table A.14 highlights that there are no significant differences in investors' average trading responses to stock market fluctuations between the two samples. In Appendix A.1.4 we show that the tendency to participate in our survey and the composition of our sample do not systematically vary with

¹²To exclude deserted accounts, we set observations with an end-of-month-level of financial wealth lower than 500 EUR to missing in all our analyses of trading behavior using the administrative account data.

the realized return of the DAX over the previous 12 months (Appendix Figure A.2 and Table A.4).

Integrity of the randomization Our sample is well-balanced between the treatment and the control group for a set of key demographic and financial characteristics as well as a set of pre-treatment beliefs (Table 1 Columns 7 and 8). There are a few exceptions, such as slight imbalances by education and age. To rule out any concerns, we include a set of control variables in all our estimations.

3 Prior beliefs and updating

3.1 Prior beliefs

Prior perceived autocorrelation We start by describing our respondents' prior beliefs about the historical autocorrelation of annual returns of the German stock market. Figure 2 Panel A shows respondents' mean estimates of the historical conditional average 12-month-ahead returns of the DAX for the six past-return intervals in our belief elicitation task. On average, respondents believe that high returns tend to be followed by low returns and vice versa, consistent with a belief in mean reversion. Specifically, respondents on average estimate a mean 12-month-ahead return of 13.5% for cases in which realized returns were in the lowest bin (less than -20%), and a mean 12-month-ahead return of 3.8% for instances in which previous returns were in the highest bin (more than 20%). Over the intermediate intervals, respondents' average estimates of the historical mean 12-month-ahead return monotonically decrease in the level of the previous 12-month return. Averaging over the six bins, respondents perceive a historical return of the DAX of 8.4%, almost identical to the actual unconditional historical average of 8.5%.¹³

The means conceal substantial heterogeneity in respondents' beliefs. Figure 2 Panel B displays box plots illustrating the distributions of respondents' estimates of the historical conditional mean 12-month-ahead returns for the six past-return intervals. Disagreement is highest for the two most extreme return bins. For instance, the interquartile range is three times as high for the highest bin (15 percentage points (pp)) than for moderately positive returns between 0% and 10% (5 pp).

We next study the perceived autocorrelation of returns at the individual level. To do this, we calculate

¹³If we weight the six bins by their historical relative frequency, the average perceived unconditional return is slightly lower at 7.2%.

the individual-specific difference between a respondent's average estimated 12-month-ahead return across the three intervals of positive previous-year returns and the respondent's average estimate for the three negative previous return intervals, which we label the "gain-loss difference". As shown in Figure 2 Panel C, the majority of our respondents (70.5%) believe that returns over the following year were systematically higher when returns in the previous year were negative than when they were positive. Thus, a belief in a negative autocorrelation seems to be most common among the investors in our sample. We classify respondents as "mean reverters" if the gain-loss difference is lower than -4 pp (52.5% of our sample), as "neutral" if this difference lies between -4 pp and 4 pp (31.9%), and as "extrapolators" if it exceeds 4 pp (15.6%). None of our findings are sensitive to the exact choice of cutoffs.

How are the different types distributed across the population? In Appendix Table A.6 we show results of regressions of different belief measures on a set of covariates. For instance, higher financial literacy, higher investment experience, higher financial wealth, and higher attention to the DAX are associated with a significantly stronger tendency to believe in mean reversion.

Perceived autocorrelation and return expectations Similarly as in other studies, respondents to our survey exhibit substantial disagreement about the expected return over the 12 months after the survey (Appendix Figure A.5 Panel A), which could partially reflect heterogeneous beliefs about return predictability based on state variables such as the recent return. Do investors' beliefs about the historical autocorrelation of returns predict their return expectations at the time of our survey? We first select the past-return interval covering respondents' perceived return over the 12 months before the survey. Respondents on average believe that the return was 5.1% (median: 5%), and there is strong heterogeneity in respondents' return perceptions (Figure A.5 Panel B). We then study the correlation between the respondent's belief about the historical mean 12-month-ahead return in the scenario containing her return perception and the respondent's actual expected return for the 12 months after the survey. The strong heterogeneity in perceptions of past realized returns is consistent with recent evidence by Beutel and Weber (2023) from a German representative household survey. Appendix Figure A.6 and Appendix Table A.7, which are based on auxiliary data collections, show similar amounts of disagreement yet smaller average misperceptions about realized returns among retail investors from our online bank

compared to general population samples from Germany and the US.¹⁴

Figure 3 Panel A displays a binned scatterplot of the relationship between actual return expectations and return expectations predicted from the perceived historical autocorrelation, partialing out a set of controls that is used throughout the paper.¹⁵ A one pp higher perceived return in the relevant historical scenario is associated with a 0.134 pp higher expected return at the time of the survey, and the relationship is highly statistically significant ($p < 0.01$). We obtain a similar effect size when selecting the relevant return bin based on the actual realized return over the 12 months before the survey instead of the respondent's perceived return (Panel B). The effects decrease somewhat in size but remain statistically significant when we control for the respondent's mean belief about the historical 12-month-ahead return for the five bins into which the previous return did *not* fall (Panels C and D). This suggests that whether the currently realized return is higher or lower than at other times matters for respondents' expectations about returns going forward. Appendix Figure A.7 shows that our estimates barely change when calculating the predicted historical year-ahead return based on fitted linear or quadratic functional forms instead of simply using the respondent's belief about the conditional mean in the relevant bin.

There are several potential reasons for the less than one-to-one pass-through from predicted to actual return expectations. First, individuals likely do not exclusively base their return expectations on previous return realizations but may also consider other state variables, and a fraction of respondents may not consider previous returns at all. Second, respondents may think that the historical autocorrelation of returns in the last 50 years is only partially informative of the autocorrelation of returns at the time of our survey, e.g., due to changes in the economy. Lastly, there could be some attenuation bias due to measurement error in respondents' beliefs about recent returns and the historical autocorrelation. Taken together, our first main result is the following:

Result 1. There is substantial heterogeneity in retail investors' beliefs about the autocorrelation of aggregate returns. A majority of investors believe that, historically, high returns tended to be followed by low returns and vice versa. Respondents' perceived historical autocorrelation is predictive of their

¹⁴For instance, the difference between the 90th and the 10th percentile is 31 pp both in an additional descriptive investor survey and in a general population survey from Germany, both run in August 2022. Yet, average absolute misperceptions of the realized return are 19 pp in the general population sample and only 13.5 pp in the additional investor survey.

¹⁵We include a set of demographics, survey measures of investor behavior such as trading experience and risk tolerance, measures of the respondent's holdings with the bank such as the equity share, and controls for technical issues such as taking the survey on a mobile phone. The exact definition of the control variables is provided in Appendix A.1.2.

expected returns over the 12 months after the survey.

These findings suggest that heterogeneity in investors' beliefs about the predictiveness of recent returns contributes to the disagreement in stock return expectations across investors that has been documented by previous literature (Giglio et al., 2021a). Specifically, investors who base their expectations on their perceived autocorrelation will process the same incoming information about return realizations differently and update their expectations about future returns in different directions.

Comparison to previous findings Our findings differ from previous literature, which has often documented a prevalence of extrapolative belief formation about stock returns (Amromin and Sharpe, 2013; Giglio et al., 2021b; Greenwood and Shleifer, 2014; Malmendier and Nagel, 2011; Vissing-Jorgensen, 2003). Part of this difference in findings could be due to the fact that most other studies rely on time-series variation in realized returns and return expectations. First, when experiencing actual stock returns, individuals may not exclusively form expectations in light of their perceived autocorrelation. Episode-specific factors and narratives are likely to influence investors' expectation formation, and could sometimes make extrapolative behavior more common (e.g., because investors think that "this time is different"). Second, changes in expectations over time capture both individuals' beliefs about statistical properties of returns and their information set regarding realized returns, while our measure isolates the first component and fixes beliefs about the realized return.

To shed light on other potential sources of the differences in findings, we examine how beliefs about the autocorrelation vary across different countries and across different types of agents. For this purpose, we turn to auxiliary surveys with another sample of retail investors from the same bank and with general population samples from Germany and the US conducted in July and August 2022.¹⁶ Appendix Figure A.8 Panels A and B highlight that, similarly as in our main survey, respondents to our additional investor survey tend to perceive a negative relationship between realized and future returns. By contrast, as shown in Panels C-F, respondents to our German and US general population surveys on average perceive a positive relationship. These patterns highlight that in principle our elicitation method is able to pick up beliefs in different forms of autocorrelation. Moreover, these findings corroborate our result that beliefs about the autocorrelation of returns are highly heterogeneous, where an agent's degree of exposure

¹⁶Appendix A.1.1 and Table A.1 provide an overview of our additional data collections.

to the stock market seems to play an important role.¹⁷ Part of the differences in findings to previous literature may thus arise from the general high degree of variability of beliefs about the dynamics of returns across different types of agents and over time. By contrast, differences across countries seem to matter less, suggesting that our focus on a sample from Germany is not driving the differences between our findings and other studies.

3.2 Updating of beliefs

Manipulation check We next examine whether the information treatment changes respondents' beliefs about the autocorrelation of aggregate stock returns. After the treatment, we ask respondents to rate their agreement on 7-point scales with three verbal statements capturing different beliefs about how informative recent past returns are for 12-month-ahead returns. While our quantitative prior elicitation focuses on historical scenarios, the qualitative post-treatment measures are designed to capture beliefs about the *current* autocorrelation. Table 2 reports OLS estimates of the effect of the treatment on respondents' agreement with these statements (z-scored using the mean and standard deviation in the sample), including the same set of controls as previously. We also report specifications in which the treatment indicator is interacted with dummies indicating prior beliefs in mean reversion or in persistence for the historical return scenarios.

First, the treatment significantly increases agreement with the statement “*With an investment in stocks one can expect a positive return, independently of how the stock market has developed in the recent past.*” by 9.2% of a standard deviation in the full sample (Column 1, $p < 0.05$). This effect is driven by those believing in mean reversion before the intervention (Column 2, $p < 0.05$), while it is smaller and insignificant among those classified as neutrals and close to zero among prior extrapolators. In line with the patterns in the treatment effects, mean reverters in the control group agree less with the statement than neutrals, although this difference is not statistically significant ($p = 0.172$), and there is no difference between neutrals and extrapolators. Thus, the type of return predictability perceived by extrapolators does not seem to be well captured by the first statement, which could explain the muted treatment effect

¹⁷Consistent with this idea, Nagel and Xu (2022b) show that households' return expectations are positively related to recent returns, while the return expectations of professionals in the Livingston survey co-move negatively with recent returns. Greenwood and Shleifer (2014) find a smaller response to realized returns among CFOs than among households. Older literature points to a weaker tendency to extrapolate past stock price trends among experts than among laypeople (De Bondt, 1993; Fisher and Statman, 2000).

for this group.

Second, the treatment does not significantly change respondents' agreement with the statement "*When the stock market has recently increased it makes no sense to buy stocks*" in the full sample (Column 3). However, those who initially believe in mean reversion reduce their agreement by 15.5% of a standard deviation in response to the information (Column 4, $p < 0.01$). This implies that the treatment fully offsets the higher baseline agreement with this statement among mean reverters compared to neutrals, as indicated by the difference in the control group (Column 4, $p < 0.05$).

Third, treated respondents agree 14.7% of a standard deviation less with the statement "*When the stock market has recently increased it is more likely that stock returns will be positive over the following time than when the stock market has recently decreased.*" (Column 5, $p < 0.01$). In the control group, extrapolators agree more with this statement than neutrals (Column 6, $p < 0.01$). In line with this, the information reduces agreement significantly more for extrapolators than for neutrals or mean reverters (Column 6, p-values of these differences < 0.05).

Taken together, these results indicate that the treatment substantially reduces beliefs in return predictability based on recent returns among respondents, and it does so differentially and in the expected directions across groups with different priors.¹⁸

Updating of return expectations We next turn to respondents' updating of their expectations about the return over the 12 months after the survey in response to the information. Depending on (i) respondents' prior beliefs about the return over the 12 months *before* the survey and (ii) respondents' prior beliefs about the historical autocorrelation, our treatment implies an information shock that should be relevant for respondents' expectations about the return over the 12 months after the survey.

We define a respondent's perception gap as follows: First, out of the six intervals of realized returns we select the one into which the respondent's perceived realized return over the 12 months before the survey falls, interval(Perceived ret 12m before survey_{*i*}). Second, we calculate the difference between the actual historical conditional mean 12-month-ahead return for the relevant interval and the respondent's

¹⁸Appendix Table A.8 demonstrates robustness of the estimations by group in Columns 2, 4 and 6 of Table 2 to using different cutoffs of the perceived gain-loss difference to define mean reverters, neutrals and extrapolators. Appendix Table A.9 demonstrates similar patterns of average and heterogeneous treatment effects using two additional statements that were included after treated respondents were shown the information for a second time in the follow-up survey.

corresponding prior:

$$\begin{aligned}
& \text{Perception gap}_i \\
& = \text{Actual hist 12m ahead ret}[\text{interval}(\text{Perceived ret 12m before survey}_i)] \\
& - \text{Prior perceived hist 12m ahead ret}_i[\text{interval}(\text{Perceived ret 12m before survey}_i)] \tag{1}
\end{aligned}$$

If respondents form their return expectations at least partially based on their beliefs about the historical predictiveness of last period’s return, a larger perception gap should lead to a stronger updating of expectations about the return over the 12 months after the survey among respondents in the treatment group. We estimate specifications of the following form:

$$\begin{aligned}
\text{Updating}_i = & \alpha_0 + \alpha_1 \text{Perception gap}_i \times \text{Treatment}_i \\
& + \alpha_2 \text{Perception gap}_i + \alpha_3 \text{Treatment}_i + \Pi \mathbf{X}_i + \varepsilon_i \tag{2}
\end{aligned}$$

where Updating_i is the difference between a respondent’s posterior and prior beliefs about the return over the 12 months after the survey.¹⁹ Our main coefficient of interest is α_1 , which captures the extent to which treated respondents update their prior towards the information, i.e., the “learning rate”. These treatment effects proportional to the distance between the information and a respondent’s prior are commonly viewed as reflecting treatment effects operating through genuine changes in beliefs. α_2 captures differential changes in expectations across respondents with different perception gaps independently of the treatment, e.g., due to respondents with extreme priors being more likely to reconsider their forecasts when reporting them for a second time. α_3 captures any treatment effects that are independent of the respondents’ priors, e.g., updating working through mechanisms such as emotional responses, salience or priming. In practice, given that priors are measured with error and given that learning may be non-linear in the perception gap, this coefficient could also capture some updating due to genuine changes in beliefs. See Haaland et al. (2023) for a more extensive discussion of these issues. \mathbf{X}_i includes the same control variables as used previously.

¹⁹As explained in Section 2.4, we drop participants who report prior or posterior expectations lower than -20% or higher than 20%, which should reduce the influence of outliers to a large extent. To account for the few remaining outliers, we winsorize both the perception gap and the updating variable at -20% and 20%. None of our findings are sensitive to the exact choice of cutoffs.

Table 3 presents the results. Column 1 shows a simple OLS estimation. In Column 2 we instrument the perception gap and the interaction term, which are calculated based on the respondents' subjective return perception over the last 12 months, with versions of the gap and the interaction term that are based on the actual realized return over the 12 months before taking the survey (which varies over the survey period), in order to mitigate attenuation bias due to measurement error in subjective beliefs. Columns 3 and 4 show OLS and IV estimations in which the posterior is based on the mean of the respondent-level subjective distribution over 12-month-ahead returns instead of the point belief. Across specifications, we estimate coefficients between 0.09 and 0.14 on the interaction term ($p < 0.01$ or $p < 0.05$). The change in beliefs about return predictability thus causes our respondents to update their 12-month-ahead return expectations in the expected direction. Respondents adjust their return expectations only partially towards the information, consistent with the less than one-to-one relationship between respondents' prior return expectation and the prediction implied by their perceived autocorrelation (Figure 3 Panel A). This suggests that investors consider also other factors than realized returns when forming return expectations or that they view the historical autocorrelation as only partially informative for the relevant autocorrelation at the time of the survey. That said, our estimated learning rates are within the range of estimates from previous information provision experiments on macroeconomic expectation formation (Haaland et al., 2023). We also find some updating in response to the treatment that is independent of a respondent's perception gap, which could be due to salience effects or priming. These effects are only present when calculating updating based on a respondent's point forecast and not when calculating it based on the mean over the subjective distribution, perhaps because the point forecast is elicited closer to the treatment. Appendix Table A.10 shows that the estimated effects of the interaction terms hardly change when we additionally control for interactions of the treatment dummy with a set of respondent characteristics. This suggests that our estimated learning rates are not confounded by the fact that the perception gap is not randomly assigned and might vary systematically across groups. Appendix Table A.11 shows that the estimated learning rates are very similar when calculating the perception gap based on fitted linear or quadratic functional forms instead of using the respondent's belief about the conditional mean in the relevant bin and the corresponding signal.

If differences in the perceived autocorrelation of returns contribute to disagreement in stock return

expectations, we should observe a decline in the dispersion of return expectations in response to our intervention. Figure 4 displays the distribution of prior and posterior 12-month-ahead return expectations separately for prior and posterior beliefs within each treatment arm, including vertical lines for different percentiles. The difference between the 90th and the 10th percentile of the distribution of return expectations is 15 pp both in the control and in the treatment group when focusing on respondents' priors (Panels A and B). Yet, it is substantially reduced to 10.6 pp in posterior beliefs among treated respondents (Panel D), while it remains at 15 pp in the posteriors of control group respondents (Panel C). Appendix Table A.12 confirms this pattern for the cross-sectional standard deviation of return expectations: Levene's tests indicate that the standard deviation does not differ significantly across treatment arms – and is in fact somewhat higher in the treatment group – when based on respondents' priors ($p = 0.241$). However, it is significantly lower in the treatment group than in the control group when focusing on posteriors ($p < 0.01$). Thus, greater agreement about the dynamic properties of returns causes a decline in the dispersion of return expectations. Taken together, our second main result is the following:

Result 2. Our information intervention removes notions of return predictability based on recent returns among our respondents. Moreover, respondents significantly adjust their expectations about the return over the 12 months after the survey in response to the information, which results in a reduction in disagreement in expectations.

These results highlight that heterogeneous beliefs about the dynamic properties of returns are a causal driver of disagreement in stock return expectations. In Appendix A.5 we demonstrate that changes in beliefs persist in the four-week follow-up survey. Cavallo et al. (2017) and Coibion et al. (2022) use placebo treatments to show that changes in reported beliefs due to numerical anchoring vanish in follow-up surveys. Moreover, experimenter demand effects should become less important the more time has passed since the experimental intervention, as subjects become less likely to remember exact details of the manipulation (Haaland et al., 2023). Our follow-up evidence thus mitigates concerns related to numerical anchoring or experimenter demand effects.²⁰

²⁰Such concerns are further mitigated by the fact that our treatment shifts respondents' actual decisions as measured in the administrative account data, as described in Section 4.2.

4 Beliefs and investment choices

4.1 Correlational evidence on beliefs and investment choices

In this section we examine how investors' prior beliefs about the autocorrelation of returns are related to their portfolio decisions as measured in the administrative account data. For this analysis, we make use of the entire history of transactions and security holdings of survey respondents with the broker since December 2014 – a period of almost five years preceding the survey. We include control group respondents until the end of January 2020, while investors in the treatment group are only part of the sample until including August 2019 – the month before the survey period.²¹

Throughout our analysis of trading behavior, we mostly focus on the relative net buying volume – the difference between the amount of equity purchased and the amount sold within a given month, divided by an investor's overall financial wealth held at the bank at the end of the previous month:

$$\text{Relative net buying volume}_{i,t} = \frac{\text{Amount purchased}_{i,t} - \text{Amount sold}_{i,t}}{\text{Financial wealth}_{i,t-1}} \quad (3)$$

If there are no inflows to or outflows from a respondent's accounts at the bank, the relative net buying volume is approximately equal to the active change in the equity share – the change in exposure to equity due to active trading, abstracting from changes that are purely due to price changes.²² We also construct a dummy variable indicating whether a respondent is a net buyer of equity in a given month, which we define as exhibiting a relative net buying volume of at least 1 pp. In some of our specifications, we split the relative net buying volume into the relative gross buying and the relative gross selling volume.

We start by examining how the average monthly trading activity of our respondents varies with the return of the DAX over the preceding 12 months.²³ Columns 1 and 2 of Appendix Table A.14 highlight that both the end-of-previous-month level of the realized past 12-month return as well as the change

²¹We obtain very similar results if we also exclude control group respondents beginning in September 2019, or if we include them until the end of March 2020, which includes the Covid-19 stock market crash. We analyze trading responses to the crash separately in Section 4.2.

²²As we do not observe inflows to or outflows from respondents' accounts at the bank, we cannot directly calculate the active change in the equity share. Our results remain very similar when we use the uncorrected change in the equity share, which includes price changes. To limit the impact of outliers, we winsorize the relative net buying volume at values close to the bottom and top percentiles both in our month-level dataset used in this section and in the week-level dataset used in Section 4.2. None of our results are sensitive to the exact choice of cutoffs.

²³Appendix Figure A.3 depicts the return of the DAX over the previous 12 months for each month in our sample period.

in the realized return over the current month are significantly negatively related to a respondent's net buying of stocks during that month, conditional on a set of macroeconomic and individual-level control variables. Thus, consistent with a large fraction of our respondents believing in a negative autocorrelation, recent return realizations are not associated with higher net buying of equity in our sample – different to previous evidence, e.g., based on fund flows in the US (Greenwood and Shleifer, 2014).

We next examine whether respondents' perceived autocorrelation of returns is predictive of their trading activity. We use the respondents' beliefs elicited across the six historical realized return scenarios and the actual return over the previous twelve months to predict a respondent's expected return at any given point in our sample period. The resulting predicted return expectation, $E_{i,t}(\text{hist. } R_{12m} | \text{previous } R_{12m})$, can be thought of as the return an investor would expect if return expectations were fully based on the perceived historical autocorrelation of returns. We then regress different measures of a respondent's trading activity during month t , $Y_{i,t}$, on the predicted return expectation at the end of the previous month and the change in the predicted return expectation over the current month:

$$Y_{i,t} = \alpha_0 + \alpha_1 \Delta E_{i,t}(\text{hist. } R_{12m} | \text{previous } R_{12m}) + \alpha_2 E_{i,t-1}(\text{hist. } R_{12m} | \text{previous } R_{12m}) + \Pi \mathbf{X}_{i,t-1} + \mu_t + \varepsilon_{i,t} \quad (4)$$

The vector $\mathbf{X}_{i,t-1}$ includes a set of controls measured at the end of the previous month, such as a respondent's equity share, dummies indicating extreme equity shares of 0% or 100% or log financial wealth held with the bank, as well as our baseline set of time-invariant controls measured in the survey as described in Appendix A.1.2.²⁴ We also control for month-year fixed effects, μ_t , which account for all observed and unobserved macroeconomic variables, including effects of realized returns on trading that are common across investors with a different perceived autocorrelation (such as effects working through portfolio rebalancing). Specification 4 is similar to the one used in Giglio et al. (2021a) to study the effect of changes in return expectations reported in different survey waves on trading, facilitating a comparison of our findings to theirs.

Table 4 Panel A shows the results for the full sample. A one pp higher predicted expected return is

²⁴For our analysis of trading decisions, we exclude variables related to trading and portfolio decisions at the time of the survey from our baseline set of controls.

associated with a 0.05 pp increase in the probability to be a net buyer of stocks during a given month (Column 1, $p < 0.1$) and a 0.015 pp higher relative net buying volume (Column 2, $p < 0.05$). This is mostly driven by a 0.011 pp higher relative gross buying volume (Column 3, $p < 0.05$), while the gross selling volume is less elastic to changes in predicted return expectations (Column 4). The lagged level of an investor's predicted return expectations is positively but mostly insignificantly related to an investor's net buying of stocks. Giglio et al. (2021a) document that changes in investors' return expectations do not predict the timing of an investor's trading decisions, but shape the direction and magnitude of trades conditional on trading. We therefore study whether changes in return expectations due to a respondent's perceived autocorrelation of returns affect trading through the extensive or the intensive margin. The absolute change in a respondent's predicted return expectation is not predictive of whether a respondent conducts any trade, defined as exhibiting a relative net buying volume of at least 1 pp in absolute terms (Column 5). However, conditional on trading, increases in respondents' return expectations are associated with a statistically insignificantly higher tendency to be a net buyer of stocks (Column 6, $p = 0.127$) and a higher relative net buying volume (Column 7, $p < 0.1$). Naturally, the effects are more noisily measured when focusing on the smaller sample of investors who trade in the current month.

Panel B repeats all estimations on an "active" sample of observations for which the corresponding investor traded at least once over the preceding three months. The magnitude and statistical significance of the association between predicted return expectations and a respondent's trading activity are considerably higher in this restricted sample.

Taken together, respondents' perceived autocorrelation of returns as elicited in our survey significantly predicts trading responses to market movements: investors buy more equity when they adjust their return expectations upward due to a change in realized returns. Moreover, the changes in return expectations implied by a respondent's belief about autocorrelation affect trading in a similar way as return expectations measured directly in repeated surveys (Giglio et al., 2021a) – namely through the intensive rather than the extensive margin of trading. In Appendix A.6 we provide a detailed discussion of the size of our estimates. Our estimated effects are an order of magnitude smaller than the effects documented by Giglio et al. (2021a), which are themselves smaller than theory benchmarks from frictionless models. One likely reason for the smaller effect sizes in our context is that we predict respondents' return expectations based

on their perceived autocorrelation, while Giglio et al. (2021a) elicit respondents' return expectations directly. Taken together, our third main result is the following:

Result 3. Changes in predicted return expectations based on a respondent's perceived autocorrelation predict investors' trading decisions.

These findings validate our main survey measure and provide correlational evidence that heterogeneous beliefs about the dynamic properties of returns are reflected in differences in investment choices across investors.

4.2 Experimental evidence on beliefs and investment choices

Next, we investigate whether changes in beliefs about the autocorrelation of returns in response to the treatment affect investors' trading behavior. For this analysis, we use transaction data until March 2020 – five to six months after the intervention. This period provides a unique setup to study effects of beliefs about the autocorrelation of returns on trading decisions, as it includes the stock market crash that was triggered by the Covid-19 pandemic. Specifically, the German stock market dropped by about 30% between mid-February and mid-March 2020, which was reflected in a change in past 12-month realized returns from about 17% to about -23% . Depending on a respondent's prior beliefs, we would expect our intervention to lead to changes in trading decisions during the crash. We focus on the same outcomes as in our correlational analysis of beliefs and trading decisions in Section 4.1. To precisely pin down the beginning and the end of the crash period in the middle of February and the middle of March, we transform our monthly dataset to a weekly panel starting in June 2019 – three months before the survey period – and ending in mid-March 2020 – the trough of the stock market at the end of the crash.

Specification We compare the development of trading decisions from before to after our intervention between the treatment and the control group. Given that our treatment should shift trading decisions into different directions depending on investors' prior beliefs, we also include interaction terms with respondents' priors. Specifically, we estimate “triple-difference”-type specifications of the following form:

$$\begin{aligned}
Y_{i,t} = & \alpha + \beta_1 \text{Crash}_t \times \text{Treatment}_i \\
& + \beta_2 \text{Crash}_t \times \text{Treatment}_i \times \mathbb{1}(\text{Predicted exp. adjustment during crash} > 0)_i \\
& + \beta_3 \text{Crash}_t \times \mathbb{1}(\text{Predicted exp. adjustment during crash} > 0)_i \\
& + \gamma_1 \text{Post-survey pre-crash}_t \times \text{Treatment}_i \\
& + \gamma_2 \text{Post-survey pre-crash}_t \times \text{Treatment}_i \times \mathbb{1}(\text{Predicted exp. adjustment during crash} > 0)_i \\
& + \gamma_3 \text{Post-survey pre-crash}_t \times \mathbb{1}(\text{Predicted exp. adjustment during crash} > 0)_i \\
& + \delta \log(\text{Fin. wealth})_{i,t-1} + \mu_i + \mu_t + \varepsilon_{i,t}
\end{aligned} \tag{5}$$

where $Y_{i,t}$ is a measure of respondent i 's trading activity during week t . Crash_t is an indicator for the weeks from 17th February until 13th March 2020 and $\text{Post-survey pre-crash}_t$ is a dummy for all weeks from 16th September 2019 – the week the first respondents took the survey – until 16th February 2020. The period from 3rd June 2019 until 15th September 2019 is the omitted pre-survey base period. $\mathbb{1}(\text{Predicted exp. adjustment during crash} > 0)$ is a dummy variable indicating whether a respondent should upward adjust his or her return expectation in response to the crash if expectations were fully based on prior beliefs about the historical autocorrelation. Thus, it takes value one for respondents who estimate a higher conditional mean 12-month-ahead return for the lowest bin of previous returns (less than -20%) than for the second-to-highest bin (between 10% and 20%). All specifications include individual fixed effects, μ_i , which control for persistent differences in trading behavior across investors (e.g., investors with different prior beliefs). The individual fixed effects also account for any potential imbalances between treatment and control group – according to both observed and unobserved variables – and increase power. We also include week-year fixed effects, μ_t , in all our specifications, which control for common shocks hitting investors in both the treatment and the control group at any given point in time. Lastly, we control for lagged log financial wealth, $\log(\text{Fin. wealth})_{i,t-1}$.²⁵

The coefficient β_1 captures treatment effects on trading responses to the crash among those who should (weakly) negatively update their expectations over the course of the crash in the absence of the treatment. Conversely, the sum $\beta_1 + \beta_2$ captures treatment effects on those who should positively update their

²⁵Our results hardly change if we exclude financial wealth, which may be affected by the treatment and therefore be a “bad control”.

expectations during the crash according to their prior beliefs. The coefficient β_2 captures the difference in treatment effects between respondents whose priors imply expectation adjustments into different directions. We consider β_2 our main coefficient of interest for the following reason: differential effects of information interventions depending on prior differences in beliefs are commonly viewed as being due to genuine belief changes (Haaland et al., 2023) – in our case changes in return expectations. By contrast, the coefficients β_1 and $\beta_1 + \beta_2$ also capture potential side-effects of the treatment operating through other channels, e.g., effects due to changes in respondents’ general understanding of the functioning of the stock market or effects shifting their general willingness to trade. β_3 captures the difference in trading reactions to the crash between investors with different priors in the control group. Lastly, the coefficients γ_1 , γ_2 and γ_3 capture analogous effects for the post-survey pre-crash period, during which the stock market was fairly stable and experienced a steady but slow increase in realized returns from about 1% to about 17%. These coefficients are not of direct interest for our analysis, but the inclusion of the corresponding variables allows us to control for any differential treatment effects on respondents with different priors already during the period before the crash. This is important, as otherwise our estimates of β_1 , β_2 or β_3 could capture these effects. In our main table we do not display the estimates of γ_1 , γ_2 and γ_3 for parsimony.

Results Table 5 presents the results for the full sample. The treatment increases the tendency to be a net buyer of equity by 2.52 pp (Column 1, $p = 0.22$) and the relative net buying volume by 0.30 pp (Column 2, $p < 0.05$) during the crash among respondents whose prior beliefs imply a downward adjustment of expectations over the crash period in the absence of the treatment (β_1). Among respondents whose prior beliefs imply an upward adjustment of expectations, the implied treatment effects ($\beta_1 + \beta_2$) are -2.39 pp for the tendency to be a net buyer (Column 1, $p < 0.01$) and -0.07 pp for the relative net buying volume (Column 2, $p = 0.420$). Our main coefficients of interest, the differences in treatment effects between the two groups, β_2 , are -4.92 pp for the tendency to be a net buyer (Column 1, $p < 0.05$) and -0.361 pp for the relative net buying volume (Column 2, $p < 0.05$). Thus, the treatment has significantly more negative effects on respondents whose prior beliefs would point to an upward adjustment of return expectations during the crash than among respondents whose priors would imply a downward adjustment. This suggests that the belief changes caused by our intervention significantly shift investors’ trading

responses to the downturn in the expected directions.

The effects on net buying are mostly driven by a higher gross buying volume (Column 3), while the gross selling volume (Column 4) seems to be less elastic to belief changes. These patterns are consistent with our correlational evidence on beliefs and trading decisions presented in Section 4.1 above. When conditioning on trading, the differential treatment effects on net buying activity across investors with different prior beliefs increase several times in magnitude (Columns 6 and 7, $p < 0.05$ and $p < 0.01$). At the same time, the differential treatment effect on the tendency to trade also has a considerable magnitude of -4.37 pp, though it is statistically insignificant (Column 5, $p = 0.159$). This suggests that belief changes caused by the intervention affect investors' decisions *both* through the extensive and the intensive margin of trading – in contrast to our correlational evidence. One potential driver of the differential findings could be the exceptional speed and size of the drop in stock prices during the Covid-19 crash, which may have drawn investors' attention to their stockholdings. Appendix Table A.15 restricts the sample to “active” observations for which the corresponding investor has traded at least once over the preceding three months. In line with our correlational evidence, our estimated treatment effects are considerably larger and more statistically significant in this sample.

Magnitudes We use two different approaches to interpret the magnitudes of the treatment effects. First, we compare the differential treatment effects on respondents whose priors would point to expectation adjustments during the crash into different directions, β_2 , to differences in trading reactions between investors with different priors in the control group, β_3 . This comparison indicates how much of the difference in trading reactions to the crash between individuals with different priors in the control group is “closed” through our intervention. In the control group, investors whose priors imply a positive expectation adjustment increase their relative net buying volume more strongly over the crash period than those whose priors imply a negative adjustment, although these differences are imprecisely estimated. Across outcomes, our estimates of the difference in treatment effects by prior beliefs, β_2 , is larger in absolute magnitude than the control group gap in trading between investors with different priors, β_3 . This suggests that our treatment effects are large relative to trading behavior in the control group.

Second, we conduct a back-of-the-envelope calculation to put the treatment effects on trading responses in relation to the implied treatment effects on return expectation adjustments over the crash

period. We estimate a differential treatment effect on the response of the relative net buying volume to the crash between those whose priors imply an upward adjustment of expectations over the crash period and those whose priors imply a downward adjustment, β_2 , of -0.361 (Column 2). Given that β_2 measures a reaction in the weekly net buying volume, we have to multiply it by four to obtain the approximate effect on the overall relative net buying volume over the four-week-long crash period, giving us an effect of -1.444 . Over the crash period, realized 12-month returns decreased from about 17% to about -23% . In the absence of the treatment, this should trigger an increase in return expectations by, on average, 10.60 pp among those whose priors imply an upward adjustment, and an average decrease by 8.55 pp among those whose priors imply a (weak) downward adjustment.²⁶ Thus, the treatment closes a gap in predicted return expectation adjustments of $10.6 \text{ pp} - (-8.55 \text{ pp}) = 19.15 \text{ pp}$ between these two groups with different prior beliefs. Given the less than one-to-one pass-through from return expectations based on the perceived historical autocorrelation to actual return expectations, this likely overstates how much of the gap in actual return expectation adjustments is closed by the treatment. Using our estimate of the relationship between actual and predicted return expectations of 0.169 (Panel B of Figure 3), we obtain a differential treatment effect on actual return expectation adjustments during the crash of approximately $-19.15 \text{ pp} \times 0.169 = -3.24 \text{ pp}$. If we divide the differential treatment effect on trading responses by the differential treatment effect on return expectation adjustments, we obtain an elasticity of $-1.444/(-3.24) = 0.446$.²⁷ This is about five times as large as our correlational estimate presented in Table 4: if we adjust the correlational estimate for the limited pass-through from predicted to actual return expectations, we obtain an elasticity of $0.015/0.169 = 0.089$. Potential reasons for the larger experimental estimates include (i) a reduced impact of measurement error or (ii) the exceptional size and speed of the Covid-19 crash, which may have increased the elasticity of trading decisions to beliefs, e.g., by increasing investors' attention. Nevertheless, our experimental estimate of the elasticity of trading to belief changes is still substantially smaller than benchmarks from frictionless models, in line with correlational evidence both from our own setting and from Giglio et al. (2021a). Appendix A.6 provides a discussion of benchmarks and of magnitudes of the correlational estimates.

²⁶This assumes that respondents hold accurate beliefs about the change in realized returns over the crash period and that expectations are fully based on the prior perceived historical autocorrelation.

²⁷If respondents do not perfectly recall the treatment information months after taking the survey, the convergence in return expectation adjustments between individuals with different priors in response to the treatment could be further attenuated. This would lead to a somewhat larger implied elasticity of trading to changes in return expectations.

Our final main result is the following:

Result 4. Changes in beliefs about the autocorrelation of aggregate returns induced by the experimental intervention shift investors' trading responses to the Covid-19 crash in the expected directions. Thus, investors' beliefs about the autocorrelation of returns causally shape their trading decisions.

Additional evidence on treatment effects on trading We expect our treatment to differentially shift respondents' trading behavior depending on the development of the stock market and depending on respondents' prior perceived autocorrelation. This implies that treatment effects on equity purchases over longer periods that include both ups and downs of the stock market – e.g., when combining post-survey pre-crash and crash period – are difficult to interpret. Similarly, average treatment effects may be muted when the treatment shifts trading decisions into different directions depending on respondents' prior beliefs. For completeness, we nevertheless report specifications estimating average treatment effects pooling across respondents with different priors (i) for the full post-survey period and (ii) separately for its two sub-periods. We focus on the relative net buying volume – our main outcome of interest – and use the same weekly panel dataset as before. Column 1 of Table 6 presents results from a difference-in-difference specification regressing the relative net buying volume on the interaction of a treatment dummy with an indicator for the entire post-survey period (including both the post-survey pre-crash period and the crash period) as well as individual and week-year fixed effects. This estimation highlights that there is no significant average effect on trading when pooling across all post-survey periods and across respondents with different prior beliefs. Column 2 shows a similar regression including separate interactions of the treatment indicator with dummy variables for the two post-survey subperiods. This estimation highlights a muted average treatment effect on trading during the crash, reflecting that the treatment shifts trading into different directions for respondents with different prior beliefs. We also find no significant average treatment effect during the post-survey pre-crash period, which is confirmed in a specification that drops the crash period from the sample (Column 3).

We next explore heterogeneity in treatment effects during the post-survey pre-crash period depending on respondents' prior beliefs. As in Column 3, we drop the crash period from the sample and regress the relative net buying volume on the interaction of a treatment indicator with a dummy for the post-survey pre-crash period as well as individual and time fixed effects. However, we now add interaction terms

with dummy variables indicating specific predicted expectation adjustments based on respondents' priors. Column 4 presents a placebo analysis that includes interaction terms with a dummy variable indicating whether the respondent's priors would imply an upward adjustment of return expectations over the *crash period* or not. Column 5 instead analyzes heterogeneity of treatment effects on trading before the crash according to the expectation adjustment over the *post-survey pre-crash period* implied by respondents' priors – i.e., the expectation adjustment that should actually be relevant for respondents' decisions in this period. Neither estimation provides evidence of systematic heterogeneity in treatment effects according to respondents' priors. Overall, our treatment seems to have had no strong effects on trading behavior during the period after the survey but before the crash. This might reflect that the stock market did not undergo major fluctuations during this period: rolling realized 12-month returns increased steadily but slowly from 1% to 17%.²⁸

5 Additional evidence: Sources of beliefs

Although our focus lies in understanding how heterogeneity in the perceived autocorrelation translates into disagreement in return expectations and trading, we also provide some evidence on potential sources of investors' beliefs about the autocorrelation. We draw both on our main survey and on two additional surveys that we conducted in July and August 2022: a descriptive survey with another sample of retail investors at the same bank and an additional experimental survey with retail investors focusing on beliefs about return predictability by valuation ratios. Appendix A.1.1 and Table A.1 provide an overview of our additional data collections.

Beliefs about return predictability by valuation ratios One possible explanation for why many of our respondents believe in a negative autocorrelation is that they are familiar with results from research in finance pointing to return predictability based on the price-dividend ratio (Campbell and Shiller, 1988; Fama and French, 1989). According to the Campbell and Shiller (1988) decomposition, variation in

²⁸The average predicted return expectation adjustment for respondents that should positively update based on their priors in response to this change in realized return is 4.49 pp, while it is -2.62 pp among those who should negatively update. Thus, the treatment closes a gap in expectation adjustments of 7.11 pp over this period, which decreases to $7.11 \text{ pp} \times 0.169 = 1.20$ pp when accounting for the incomplete pass-through from expectations predicted from respondents' perceived historical autocorrelation to actual return expectations (see Panel B of Figure 3) – only a third of the implied effect on expectations during the crash. This may have been too small to trigger adjustments in respondents' decisions. See also Figure 2 for the much more moderate variation in prior perceived year-ahead returns for scenarios of intermediate realized returns compared to scenarios of extreme realized returns.

the price-dividend ratio necessarily reflects changes in expected future cash flow growth or changes in expected future returns. Thus, observing high recent returns could make investors conclude that the price-dividend ratio is currently high, which in turn would lead them to predict lower future returns. This is an unlikely explanation for our findings, for at least three reasons.

First, past-12-months returns are only weakly correlated with current valuation ratios in the German stock market, consistent with 12-month returns mostly capturing higher frequency variation and not longer-term swings in valuation ratios (see Appendix A.2 including Table A.5 Panel A). Second, variation in the price-dividend ratio is mostly driven by variation in future cash flow growth in the German context (Rangvid et al., 2014), which is reflected in only weak predictability of future returns based on current valuation ratios (Table A.5 Panel B). Investors' *perceived* relationship between future returns and the current PD ratio is even weaker than the actual relationship, as we demonstrate in our additional experimental survey (see Appendix A.7). Lastly, in our additional descriptive survey we elicit investors' beliefs about the autocorrelation of returns and on the next survey screen ask them which of the factors on a list presented to them contributed to their estimates. Only few respondents indicate that their knowledge of academic research in finance (8.7%) or their memory of the development of other financial variables such as valuation ratios (11.5%) played a role in their estimates (Appendix Figure A.9 Panel A). Moreover, these groups do not exhibit a significantly higher tendency to believe in mean reversion (Table 7 Columns 1 and 2). Similarly, believing in mean reversion is unrelated to respondents' self-reported knowledge of financial market theories (Column 3). These patterns are consistent with other evidence showing that recent returns and current valuation ratios affect investors' return expectations in a largely orthogonal manner (Nagel and Xu, 2022b). Together, these points suggest that believing in a negative autocorrelation of returns is not driven by familiarity with the theory-based link between current valuation ratios and future returns.

Memory and experiences An alternative driver of beliefs is investors' memory. When investors think about the dynamic properties of stock returns, they may selectively retrieve memories of specific episodes that involved persistence or mean reversion. To shed light on this possibility, we ask respondents to our additional descriptive survey to explain in open text which past episodes – if any – they thought of in particular when reporting their perceived autocorrelation. We hand-code the resulting text data and

identify nine episodes in the last decades that are mentioned by respondents and that unambiguously refer to specific stock market developments. All of these episodes are linked to salient global events and involved major market turbulences. 55.8% of investors mention at least one of them. As shown in Figure A.9 Panel B, the most frequently mentioned episodes are the Covid-19 crash in 2020 (27.9%), the Global Financial Crisis 2007-9 (35.1%), and the burst of the “dot-com bubble” in the early 2000s (18.3%). Table A.13 highlights that investors particularly retrieve those historical episodes that occurred while they were already participating in the stock market or while they had already reached adulthood. Investors’ database of past stock market developments thus seems to be strongly shaped by their own experiences (Malmendier and Nagel, 2011; Nagel and Xu, 2022a).

While seven of these historical episodes – including the Global Financial Crisis – were associated with reversals of returns, returns did not exhibit a clear tendency to revert during the two remaining episodes – including the early 2000s recession associated with the end of the “dot-com bubble”.²⁹ As shown in Table 7 Column 4, retrieving at least one reversal episode is associated with a 20.4 pp higher tendency to believe in mean reversion ($p < 0.01$), conditional on controls. The effect size is substantial in light of an overall fraction of mean reverters of 49% in our additional descriptive investor survey. By contrast, retrieving at least one non-reversal episode is unrelated to whether a respondent believes in mean reversion (Column 5). These patterns remain unchanged when we include proxies for knowledge of research results from finance (Column 6). Thus, thinking about the dynamic properties of returns seems to trigger recall of salient past stock market episodes, which in turn shape investors’ perceived autocorrelation. These findings are consistent with growing evidence on the role of associative memory in belief formation and financial decisions (Andre et al., 2022a; Bordalo, Conlon, Gennaioli, Kwon and Shleifer, 2023; Bordalo et al., 2022; Charles, 2022; Enke et al., 2022).

6 Conclusion

We study beliefs about the dynamic properties of aggregate stock returns and the causal effect of these beliefs on investment decisions using a survey with German retail investors that embeds an information

²⁹We classify the historical episodes as follows: For each week during an episode, we compute the return over the past and the subsequent 12 months. Whenever the two returns have the opposite sign, we classify this as a reversal observation. When during a specific episode, such as the Covid-19 crash, there are more than 50% of weeks classified as reversal observations, we classify the episode as a reversal episode.

provision experiment. There is substantial heterogeneity in beliefs about the predictability of future returns based on recent return realizations. Prior to our intervention, a majority of investors believe that high returns tend to be followed by low returns and vice versa. Our respondents significantly update their beliefs in response to information about the low historical degree of predictability of the stock market based on recent returns. Beliefs about the autocorrelation of returns predict investors' trading decisions over the five years before the intervention. The treatment shifts respondents' trading reactions to the Covid-19 stock market crash four to five months after the survey in the directions of the belief changes caused by the intervention.

Our findings provide causal evidence on the role of subjective beliefs in trading decisions, underlining the importance of incorporating the empirical evidence on investors' subjective beliefs into models of investment behavior and asset prices. Our paper highlights that investors disagree strongly about the statistical properties of aggregate stock returns, which calls for more research on agents' mental models of financial markets. Future research could supplement our method with open-ended survey questions to better understand the origins of such models and how they vary with investor characteristics such as financial sophistication, experiences and risk-bearing capacity.

Our method can be flexibly applied to study beliefs about relationships between other financial variables, as we show for the relationship between valuation ratios and future returns in the appendix. For instance, future research could apply our elicitation method to measure beliefs about the autocorrelation of returns for individual stocks, for returns of varying horizons, or based on scenarios of hypothetical future returns rather than historical realized returns. While our paper focuses on beliefs about statistical properties of returns, future work could explore other sources of extrapolative and contrarian belief formation such as the narratives that accompany specific market episodes.

References

- Adam, Klaus, Albert Marcet, and Johannes Beutel**, "Stock Price Booms and Expected Capital Gains," *American Economic Review*, 2017, 107 (8), 2352–2408.
- **and Stefan Nagel**, "Expectations Data in Asset Pricing," *Handbook of Economic Expectations*, 2022.
- Ameriks, John, Gábor Kézdi, Minjoon Lee, and Matthew D Shapiro**, "Heterogeneity in Expectations, Risk Tolerance, and Household Stock Shares: The Attenuation Puzzle," *Journal of Business & Economic Statistics*, 2020, 38 (3), 633–646.

- Amromin, Gene and Steven A Sharpe**, “From the Horse’s Mouth: Economic Conditions and Investor Expectations of Risk and Return,” *Management Science*, 2013, 60 (4), 845–866.
- Andersen, Steffen, Tobin Hanspal, Jimmy Martinez-Correa, and Kasper Meisner Nielsen**, “Beliefs and the Disposition Effect,” *Working Paper*, 2021.
- Andrade, Philippe, Richard K Crump, Stefano Eusepi, and Emanuel Moench**, “Fundamental Disagreement,” *Journal of Monetary Economics*, 2016, 83, 106–128.
- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart**, “Subjective Models of the Macroeconomy: Evidence from Experts and Representative Samples,” *Review of Economic Studies*, 2022, 89 (6), 2958–2991.
- , **Ingar Haaland, Christopher Roth, and Johannes Wohlfart**, “Narratives about the Macroeconomy,” *CEPR Discussion Paper No. DP16758*, 2022.
- Andries, Marianne, Milo Bianchi, Karen K Huynh, and Sébastien Pouget**, “Return Predictability, Expectations, and Investment: Experimental Evidence,” *Working Paper*, 2020.
- Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar**, “The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment,” *Review of Economics and Statistics*, 2016, 98 (3), 503–523.
- Armona, Luis, Andreas Fuster, and Basit Zafar**, “Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment,” *Review of Economic Studies*, 2019, 86 (4), 1371–1410.
- Arrondel, Luc, Hector Calvo-Pardo, Chryssi Giannitsarou, and Michael Haliassos**, “Informative Social Interactions,” *Journal of Economic Behavior and Organization*, 2022, 203, 246–263.
- Banerjee, Snehal and Ilan Kremer**, “Disagreement and Learning: Dynamic Patterns of Trade,” *Journal of Finance*, 2010, 65 (4), 1269–1302.
- Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer**, “X-CAPM: An Extrapolative Capital Asset Pricing Model,” *Journal of Financial Economics*, 2015, 115 (1), 1–24.
- , – , – , and – , “Extrapolation and Bubbles,” *Journal of Financial Economics*, 2018, 129 (2), 203–227.
- Beutel, Johannes and Michael Weber**, “Beliefs and Portfolios: Causal Evidence,” *Working Paper*, 2023.
- Binder, Carola and Alex Rodrigue**, “Household Informedness and Long-Run Inflation Expectations: Experimental Evidence,” *Southern Economic Journal*, 2018, 85 (2), 580–598.
- Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Yongwook Kwon, and Andrei Shleifer**, “Memory and Probability,” *The Quarterly Journal of Economics*, 2023, 138 (1), 265–311.
- , **Nicola Gennaioli, and Andrei Shleifer**, “Overreaction and Diagnostic Expectations in Macroeconomics,” *Journal of Economic Perspectives*, 2022, 36 (3), 223–44.
- Bottan, Nicolas L and Ricardo Perez-Truglia**, “Betting on the House: Subjective Expectations and Market Choices,” *Working Paper*, 2022.
- Campbell, John Y and Robert J Shiller**, “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1988, 1 (3), 195–228.

- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia**, “Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments,” *American Economic Journal: Macroeconomics*, 2017, 9 (3), 1–35.
- Charles, Constantin**, “Memory and Trading,” *Working Paper*, 2022.
- , **Cary Frydman, and Mete Kilic**, “Insensitive Investors,” *Available at SSRN 3971082*, 2022.
- Cochrane, John H**, “The Dog that Did not Bark: A Defense of Return Predictability,” *Review of Financial Studies*, 2008, 21 (4), 1533–1575.
- , “Presidential Address: Discount Rates,” *The Journal of Finance*, 2011, 66 (4), 1047–1108.
- Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber**, “Forward Guidance and Household Expectations,” *Journal of the European Economic Association*, 2023.
- , **Yuriy Gorodnichenko, and Michael Weber**, “Monetary Policy Communications and Their Effects on Household Inflation Expectations,” *Journal of Political Economy*, 2022, 130 (6), 1537–1584.
- Cutler, David M, James M Poterba, and Lawrence H Summers**, “Speculative Dynamics,” *Review of Economic Studies*, 1991, 58 (3), 529–546.
- D’Acunto, Francesco, Andreas Fuster, and Michael Weber**, “A Diverse Fed Can Reach Underrepresented Groups,” *Working Paper*, 2022.
- Das, Sreyoshi, Camelia M Kuhnen, and Stefan Nagel**, “Socioeconomic Status and Macroeconomic Expectations,” *Review of Financial Studies*, 2020, 33 (1), 395–432.
- De Bondt, Werner PM**, “Betting on Trends: Intuitive Forecasts of Financial Risk and Return,” *International Journal of Forecasting*, 1993, 9 (3), 355–371.
- Dominitz, Jeff and Charles F Manski**, “Expected Equity Returns and Portfolio Choice: Evidence from the Health and Retirement Study,” *Journal of the European Economic Association*, 2007, 5 (2-3), 369–379.
- **and** —, “Measuring and Interpreting Expectations of Equity Returns,” *Journal of Applied Econometrics*, 2011, 26 (3), 352–370.
- Drerup, Tilman, Benjamin Enke, and Hans-Martin Von Gaudecker**, “The Precision of Subjective Data and the Explanatory Power of Economic Models,” *Journal of Econometrics*, 2017, 200 (2), 378–389.
- Dumas, Bernard, Alexander Kurshev, and Raman Uppal**, “Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility,” *Journal of Finance*, 2009, 64 (2), 579–629.
- Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann**, “Associative Memory and Belief Formation,” *NBER Working Paper*, 2022.
- Fama, Eugene F and Kenneth R French**, “Permanent and Temporary Components of Stock Prices,” *Journal of Political Economy*, 1988, 96 (2), 246–273.
- **and** —, “Business Conditions and Expected Returns on Stocks and Bonds,” *Journal of Financial Economics*, 1989, 25 (1), 23–49.
- Fisher, Kenneth L and Meir Statman**, “Investor Sentiment and Stock Returns,” *Financial Analysts Journal*, 2000, 56 (2), 16–23.

- Galashin, Mikhail, Martin Kanz, and Ricardo Perez-Truglia**, “Macroeconomic Expectations and Credit Card Spending,” *Working Paper*, 2022.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus**, “Five Facts About Beliefs and Portfolios,” *American Economic Review*, 2021, *111* (5), 1481–1522.
- , —, —, —, and —, “The Joint Dynamics of Investor Beliefs and Trading During the COVID-19 Crash,” *Proceedings of the National Academy of Sciences*, 2021, *118* (4).
- Greenwood, Robin and Andrei Shleifer**, “Expectations of Returns and Expected Returns,” *Review of Financial Studies*, 2014, *27* (3), 714–746.
- Haaland, Ingar and Ole-Andreas Elvik Næss**, “Misperceived Returns to Active Investing,” *CESifo Working Paper*, 2023.
- , **Christopher Roth, and Johannes Wohlfart**, “Designing Information Provision Experiments,” *Journal of Economic Literature*, 2023, *61* (1), 1–38.
- Hanspal, Tobin, Annika Weber, and Johannes Wohlfart**, “Exposure to the COVID-19 Stock Market Crash and its Effect on Household Expectations,” *Review of Economics and Statistics*, 2021, *103* (5), 994–1010.
- Harris, Milton and Artur Raviv**, “Differences of Opinion Make a Horse Race,” *Review of Financial Studies*, 1993, *6* (3), 473–506.
- Heiss, Florian, Michael Hurd, Maarten van Rooij, Tobias Rossmann, and Joachim Winter**, “Dynamics and Heterogeneity of Subjective Stock Market Expectations,” *Journal of Econometrics*, 2022, *231* (1), 213–231.
- Hoffmann, Arvid, Thomas Post, and Joost Pennings**, “How Investor Perceptions Drive Actual Trading and Risk-taking Behavior,” *Journal of Behavioral Finance*, 2015, *16* (1), 94–103.
- Huang, Dashan, Jianguan Li, Liyao Wang, and Guofu Zhou**, “Time Series Momentum: Is It There?,” *Journal of Financial Economics*, 2020, *135* (3), 774 – 794.
- Kézdi, Gábor and Robert J Willis**, “Household Stock Market Beliefs and Learning,” *Working Paper*, 2011.
- Kuhnen, Camelia M, Sarah Rudolf, and Bernd Weber**, “The Effect of Prior Choices on Expectations and Subsequent Portfolio Decisions,” *Working Paper*, 2017.
- Long, J Bradford De, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann**, “Positive Feedback Investment Strategies and Destabilizing Rational Speculation,” *Journal of Finance*, 1990, *45* (2), 379–395.
- Luo, Cheng, Enrichetta Ravina, Marco Sammon, and Luis M Viceira**, “Retail Investors’ Contrarian Behavior Around News and the Momentum Effect,” *Working Paper*, 2022.
- Malmendier, Ulrike and Stefan Nagel**, “Depression Babies: Do Macroeconomic Experiences Affect Risk-taking?,” *The Quarterly Journal of Economics*, 2011, *126* (1), 373–416.
- Meeuwis, Maarten, Jonathan A Parker, Antoinette Schoar, and Duncan I Simester**, “Belief Disagreement and Portfolio Choice,” *The Journal of Finance*, 2022, *77* (6), 3191–3247.
- Merkle, Christoph and Martin Weber**, “Do Investors Put Their Money Where Their Mouth Is? Stock Market Expectations and Investing Behavior,” *Journal of Banking and Finance*, 2014, *46*, 372–386.

Nagel, Stefan and Zhengyang Xu, “Asset Pricing with Fading Memory,” *Review of Financial Studies*, 2022, 35 (5), 2190–2245.

— and —, “Dynamics of Subjective Risk Premia,” *Journal of Financial Economics*, 2022.

Rangvid, Jesper, Maik Schmeling, and Andreas Schrimpf, “Dividend Predictability Around the World,” *Journal of Financial and Quantitative Analysis*, 2014, 49 (5-6), 1255–1277.

Roth, Christopher and Johannes Wohlfart, “How Do Expectations About the Macroeconomy Affect Personal Expectations and Behavior?,” *Review of Economics and Statistics*, 2020, 102 (4), 731–748.

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

Vissing-Jorgensen, Annette, “Perspectives on Behavioral Finance: Does “Irrationality” Disappear with Wealth? Evidence from Expectations and Actions,” *NBER Macroeconomics Annual*, 2003, 18, 139–194.

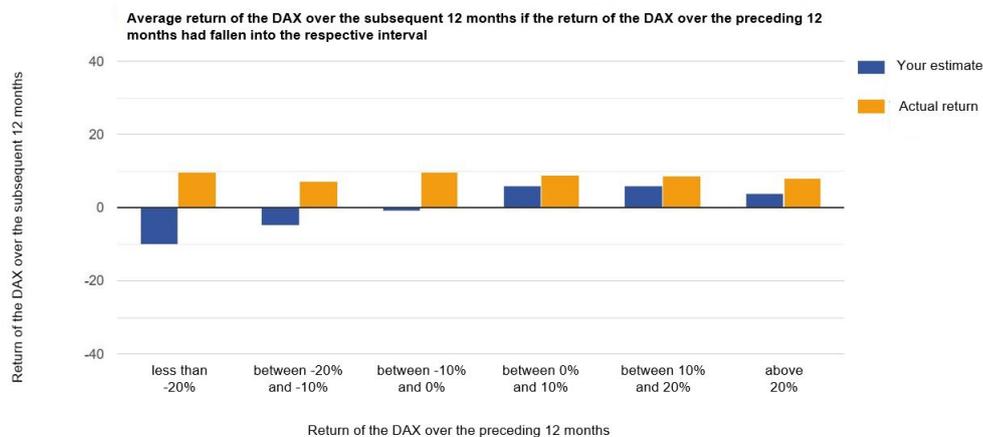
von Gaudecker, Hans-Martin and Axel Wogroly, “Heterogeneity in Households’ Stock Market Beliefs: Levels, Dynamics, and Epistemic Uncertainty,” *Journal of Econometrics*, 2022, 231 (1), 232–247.

Main figures

Figure 1: Screenshots of dynamic figures in the survey for prior elicitation and information treatment



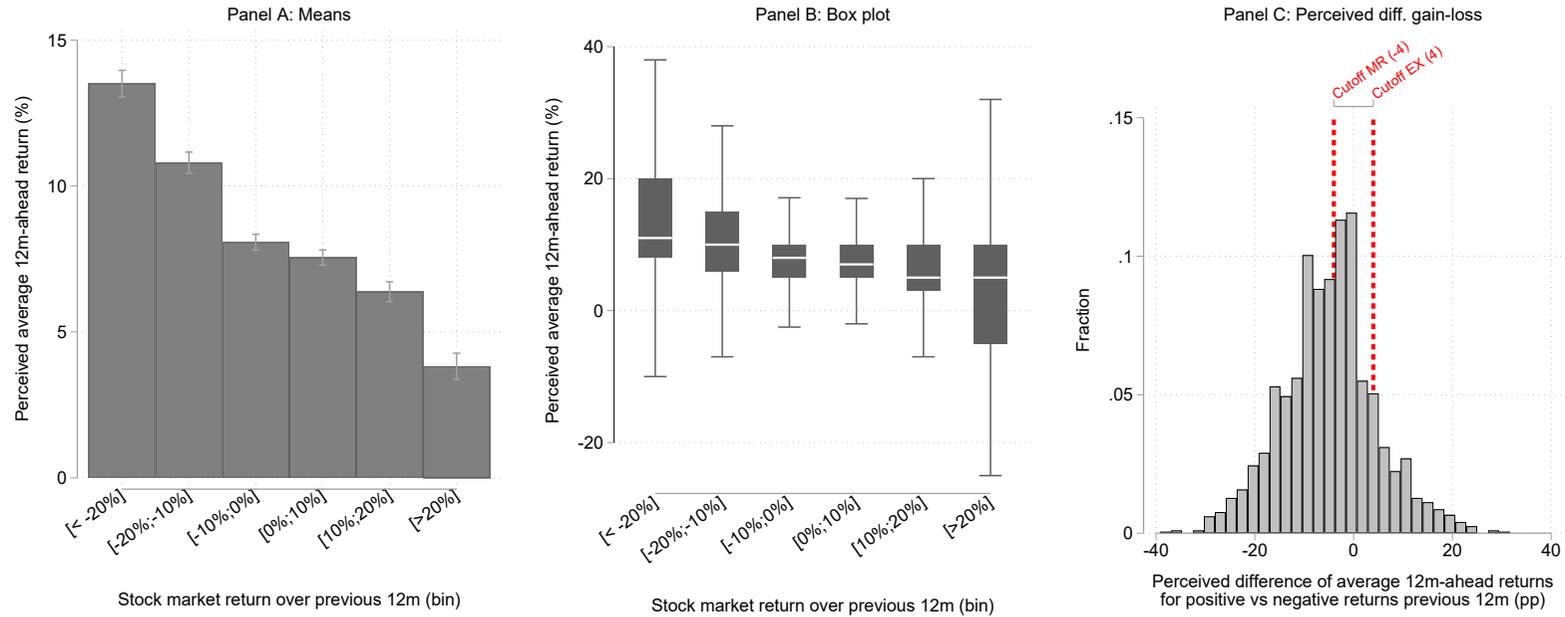
(a) Panel A: Priors



(b) Panel B: Information treatment

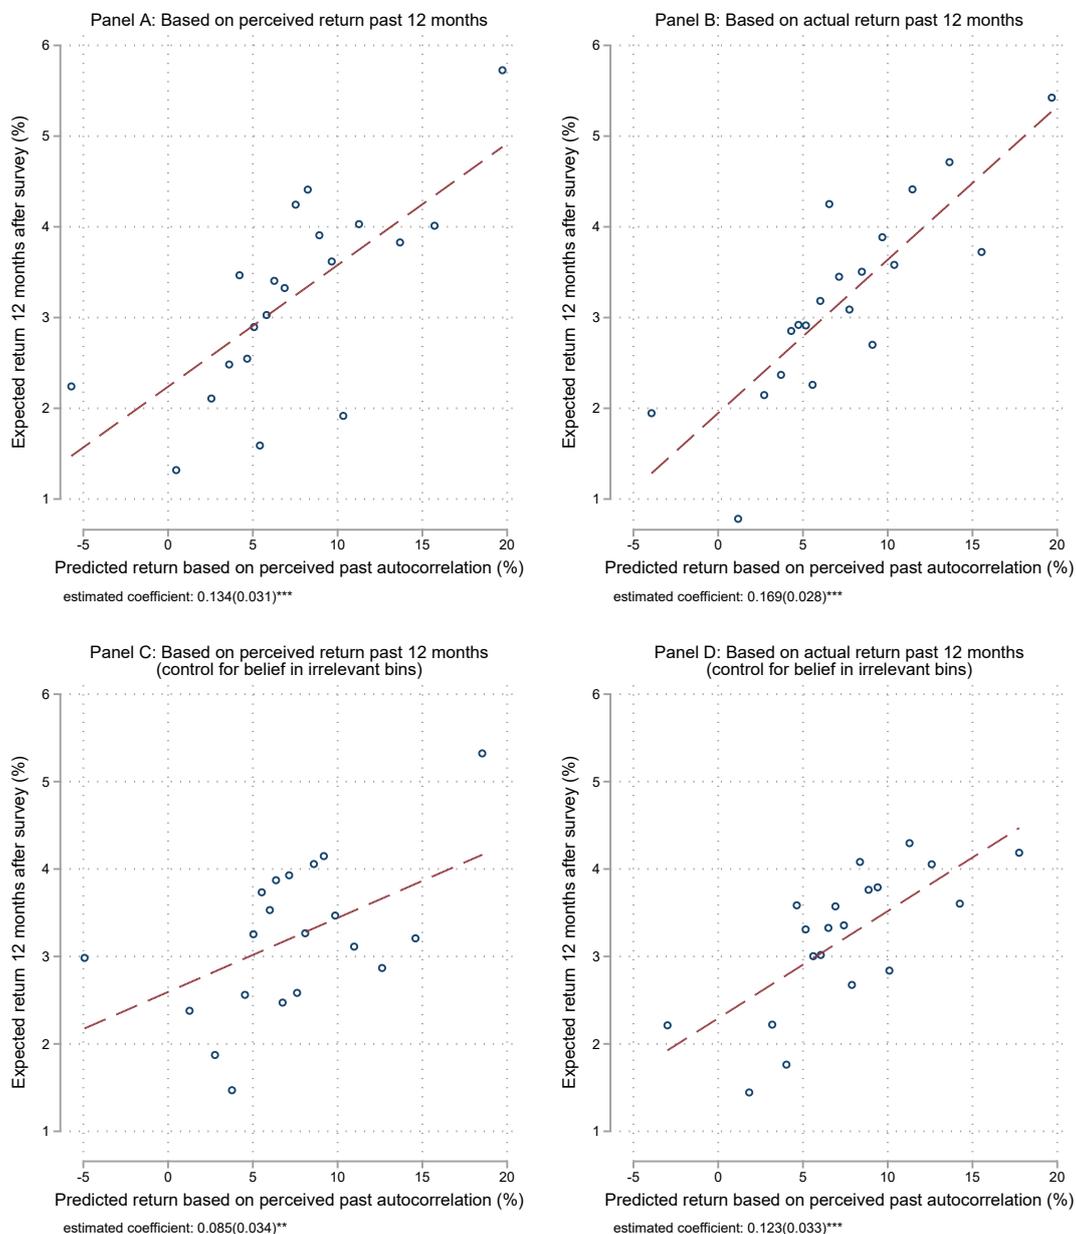
Notes: This figure presents examples from screenshots of the dynamic figures included in the main survey for the elicitation of priors and for the information treatment (translated from German to English). Panel A illustrates the figure used in the elicitation of prior beliefs about the historical autocorrelation of aggregate stock returns. For each of the six intervals on the horizontal axis, starting with the lowest one on the left, respondents are instructed to think of all points in time over the past 50 years at which the return of the DAX over the preceding 12 months had fallen into the respective interval, and ask them to estimate the conditional average return of the DAX over the subsequent 12 months. Each interval is asked about on a separate screen. On each screen, the graph displays the respondent's entry for the current interval as well as his or her estimates for previous intervals (blue bars). Panel B displays the figure on the information treatment screen shown to respondents in the treatment group. The orange bars illustrate the actual historical conditional mean 12-month-ahead returns in the six past-return intervals, respectively. Initially, the screen only shows the participants' entries previously made in all six scenarios. Participants are instructed to repeatedly click on a button to receive information on the actual values interval-by-interval. In addition, for each bin, we display a sentence above the figure comparing the respondent's prior with the actual value for the respective bin.

Figure 2: Prior beliefs about the autocorrelation of aggregate returns



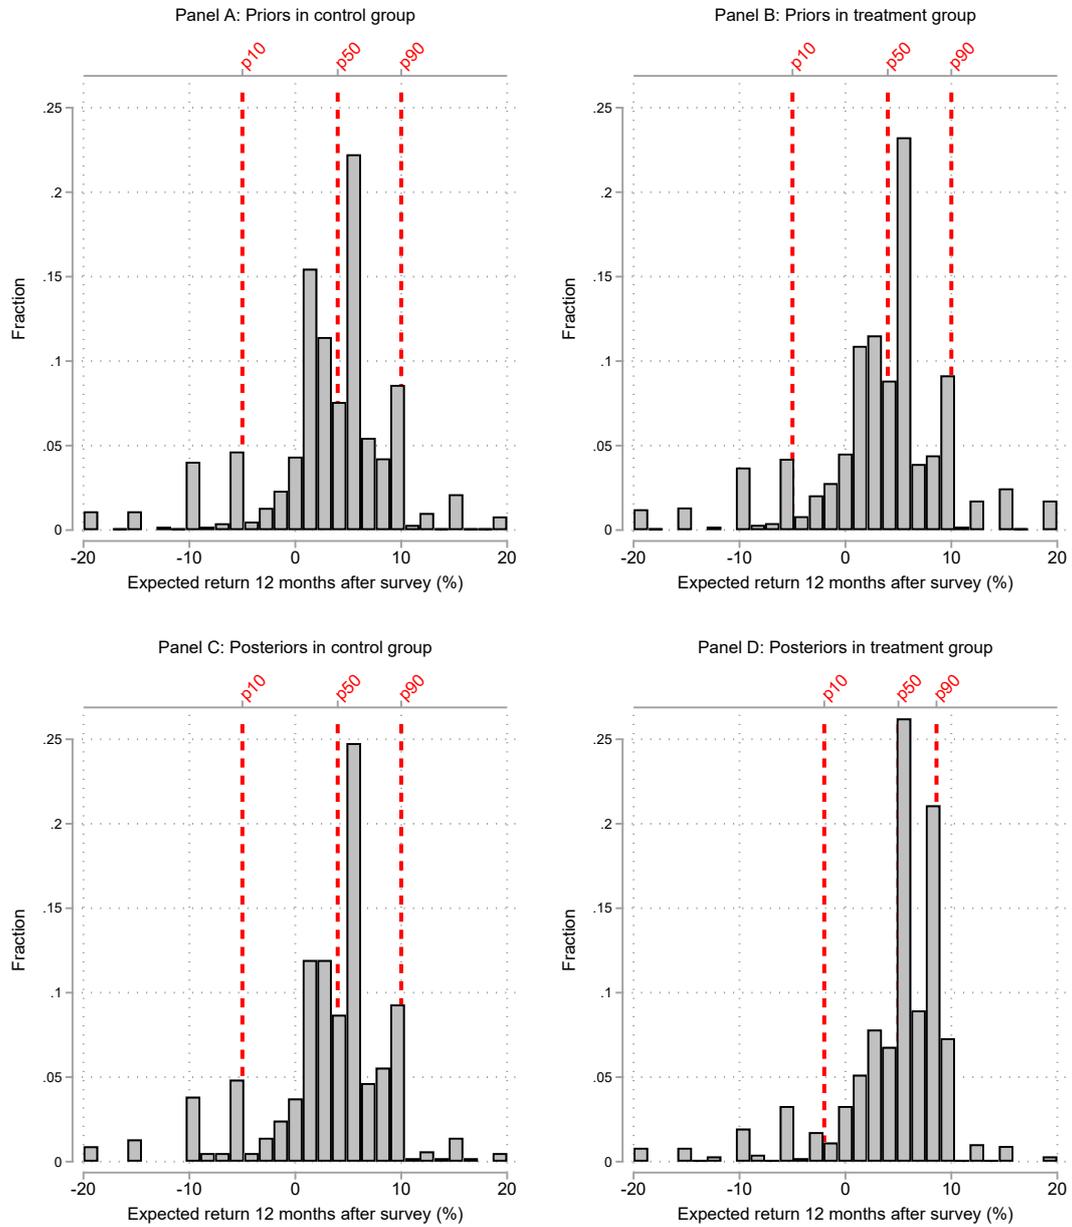
Notes: This figure summarizes prior beliefs about the autocorrelation of returns of the German stock market in the last 50 years among respondents to our main survey. Panel A shows the sample means of respondents' beliefs about average 12-month-ahead stock returns for six intervals of realized returns over the previous 12 months. Panel B displays box plots of respondents' prior beliefs about average 12-month-ahead stock returns for the six intervals of realized returns over the previous 12 months, including median, 25th and 75th percentile for each interval. Panel C shows a histogram of respondents' perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios, including the cutoffs we use to define mean reverters, neutrals and extrapolators.

Figure 3: Binned scatter plot of prior expected 12-month-ahead return vs predicted return expectation based on perceived historical autocorrelation



Notes: This figure shows binned scatter plots of respondents' prior expected return over the 12 months after the survey against respondents' perceived average historical 12-month-ahead return in the relevant interval of realized returns, which is selected based on the respondent's perceived return over the 12 months before the survey (Panels A and C) or based on the actual realized return over the 12 months before the respondent took the survey (Panels B and D), among respondents to our main survey. The binned scatter plots partial out the baseline set of controls described in Appendix A.1.2. Panels C and D additionally control for the respondent's average belief about the historical 12-month-ahead return for the five bins into which the perceived or actual return over the 12 months before the respondent took the survey did *not* fall. Robust standard errors are displayed below the different panels.

Figure 4: Prior and posterior beliefs about the 12-month-ahead return by treatment arm



Notes: This figure displays histograms of beliefs about the return of the German stock market over the 12 months after the survey, both based on priors (Panels A and B) and based on posteriors (Panels C and D). Panels A and C display beliefs for the control group, while Panels B and D display beliefs for the treatment group. The vertical dashed lines indicate the 10th percentile, the median, and the 90th percentile of the distributions.

Main tables

Table 1: Summary statistics and balance check

	PHF	Online brokerage sample							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	2017 Mean	Mean	Median	SD	p25	p75	Treatment Group: Mean	Control Group: Mean	p-value (7) = (8)
Female	0.49	0.16	0.00	0.37	0.00	0.00	0.17	0.16	0.511
Age	50.55	45.24	45.00	14.15	34.00	55.00	45.84	44.66	0.067
University	0.36	0.54	1.00	0.50	0.00	1.00	0.52	0.56	0.079
Employed	0.65	0.77	1.00	0.42	1.00	1.00	0.75	0.78	0.150
Household net income	3,808	3,914	4,000	2,769	2,000	5,250	3,927	3,902	0.837
Household net wealth	361,783	300,488	125,000	458,044	12,500	375,000	307,809	293,294	0.483
Total financial wealth at bank		55,272	22,082	98,312	5,581	65,752	55,073	55,468	0.929
Portfolio value at bank		43,970	14,872	87,671	3,726	47,620	43,438	44,489	0.795
Equity holdings at bank		39,405	13,381	78,678	3,437	42,458	38,318	40,467	0.553
Average monthly equity trades		1.73	0.67	3.29	0.00	2.00	1.75	1.71	0.812
Risk tolerance (1-7)		4.56	5.00	1.17	4.00	5.00	4.54	4.58	0.403
Trading experience (years)		14.13	15.00	10.87	4.00	20.00	14.38	13.88	0.309
Financial literacy score (0-3)		1.82	2.00	0.78	1.00	2.00	1.81	1.83	0.543
Follow DAX developments (1-7)		4.76	5.00	1.81	3.00	6.00	4.78	4.75	0.737
Investment horizon ≥ 5 years		0.49	0.00	0.50	0.00	1.00	0.48	0.50	0.373
Perceived return last 12 months		5.09	5.00	6.07	2.00	8.00	4.99	5.19	0.475
Confident in perceived return		0.64	1.00	0.48	0.00	1.00	0.64	0.64	0.855
Expected return next 12 months		3.21	4.00	6.28	1.50	6.00	3.32	3.09	0.423
Confident in expected return		0.54	1.00	0.50	0.00	1.00	0.54	0.53	0.494
Perceived mean hist. ret. intervals		8.36	7.83	4.62	5.17	11.17	8.39	8.32	0.739
Perceived diff. gain-loss historical		-4.88	-4.67	9.64	-10.67	0.50	-4.68	-5.08	0.360
Extrapolator (diff. ≥ 4)		0.16	0.00	0.36	0.00	0.00	0.16	0.15	0.725
Mean-reverter (diff. < -4)		0.53	1.00	0.50	0.00	1.00	0.53	0.52	0.789
In follow-up sample		0.46	0.00	0.50	0.00	1.00	0.46	0.46	0.886
Observations		1,961					972	989	

Notes: This table shows summary statistics for the sample of retail investors at the online bank that responded to our main survey (Columns 2-6), as well as benchmarks from the German population of individuals participating in the stock market as measured in the 2017 wave of the Bundesbank's Panel of Household Finance (Column 1). Columns 7-9 provide a check of balance of means between treatment and control group. Variables on income, wealth and wealth components are expressed in euro terms. Financial wealth at the bank, portfolio value at the bank, and equity holdings at the bank are measured in the month prior to the survey. Average monthly equity trades are measured over the three months preceding the survey. All belief variables reported in the table refer to respondents' priors elicited before the information treatment.

Table 2: Manipulation check

	Positive return irrespective of previous return (z)		No sense to buy after high return (z)		Positive return more likely after high return (z)	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.092** (0.044)		-0.054 (0.044)		-0.147*** (0.045)	
Treatment × Extrapolator (diff. ≥ 4) (a)		-0.037 (0.112)		0.021 (0.114)		-0.375*** (0.115)
Treatment × Neutral ($-4 \leq \text{diff.} < 4$)		0.078 (0.079)		0.075 (0.080)		-0.084 (0.081)
Treatment × Mean-reverter (diff. < -4) (b)		0.140** (0.059)		-0.155*** (0.060)		-0.114* (0.062)
Extrapolator (diff. ≥ 4)	-0.035 (0.069)	0.022 (0.097)	-0.018 (0.071)	0.008 (0.098)	0.143** (0.072)	0.288*** (0.102)
Mean-reverter (diff. < -4)	-0.066 (0.050)	-0.097 (0.071)	0.046 (0.051)	0.160** (0.070)	-0.127** (0.053)	-0.113 (0.072)
p-value (a=b)		0.161		0.174		0.047
Observations	1,961	1,961	1,961	1,961	1,961	1,961
R-squared	0.10	0.10	0.08	0.08	0.04	0.04

Notes: This table shows estimations of the effect of the information treatment on posterior agreement with verbal statements describing beliefs about the autocorrelation of aggregate returns among respondents to our main survey. Agreement with the statements is elicited on 7-point categorical scales and is z-scored using the means and standard deviations in the sample. The statements are: “With an investment in stocks one can expect a positive return, independently of how the stock market has developed in the recent past.” (Columns 1-2); “When the stock market has recently increased it makes no sense to buy stocks.” (Columns 3-4); “When the stock market has recently increased it is more likely that stock returns will be positive over the following time than when the stock market has recently decreased.” (Columns 5-6). Columns 2, 4 and 6 show heterogeneous treatment effects for prior extrapolators (perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios at least 4 pp), neutrals (difference at least -4 pp and less than 4 pp), and mean reverters (difference less than -4 pp). All estimations include the baseline set of controls described in Appendix A.1.2. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table 3: Updating of 12-month-ahead return expectations

	Updating (point belief)		Updating (mean distr.)	
	(1) OLS	(2) IV	(3) OLS	(4) IV
Treatment × Perception gap	0.086** (0.038)	0.138*** (0.051)	0.115*** (0.044)	0.142** (0.060)
Perception gap	-0.004 (0.025)	-0.019 (0.033)	0.022 (0.028)	0.044 (0.038)
Treatment	1.077*** (0.212)	1.007*** (0.219)	0.047 (0.263)	0.019 (0.266)
First stage F-stat		1020.48		1020.48
Observations	1,961	1,961	1,961	1,961
R-squared	0.05	0.04	0.04	0.04

Notes: This table examines changes in expectations about aggregate stock returns over the 12 months after the survey in response to the information among respondents to our main survey based on estimations of specification 2. The outcomes are the difference between posterior and prior point expectations about the 12-month-ahead return (Columns 1-2) and the difference between the mean of the respondent-level posterior distribution over 12-month-ahead returns and the prior point expectation (Columns 3-4). The perception gap is based on the respondent's prior belief about the historical autocorrelation of aggregate returns. It is the difference between the actual conditional mean 12-month-ahead return and the respondent's corresponding prior for the relevant scenario of realized returns over the previous 12 months, which is selected based on respondent's perceived return over the 12 months before the main survey. In Columns 2 and 4 the perception gap is instrumented with a version in which the relevant return interval is selected based on the actual realized return of the DAX over the 12 months before the survey. All estimations include the baseline set of controls described in Appendix A.1.2. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table 4: Perceived autocorrelation and trading behavior

	Probability net buy (%)	Relative net buying volume (%)	Relative gross buying volume (%)	Relative gross selling volume (%)	Probability trade (%)	Probability net buy (%)	Relative net buying volume (%)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Full sample							
Δ Perceived conditional hist. return (%)	0.050* (0.029)	0.015** (0.007)	0.011** (0.006)	-0.005 (0.005)		0.096 (0.062)	0.042* (0.024)
Lagged perceived conditional hist. return (%)	0.025 (0.042)	0.010* (0.006)	0.010 (0.007)	0.001 (0.007)	0.006 (0.043)	0.049 (0.068)	0.030 (0.020)
Δ Perceived conditional hist. return (%)					0.011 (0.079)		
Time FE Specification	Yes	Yes	Yes	Yes	Yes	Yes Cond. on trading	Yes Cond. on trading
Observations	74,569	74,569	74,569	74,569	74,569	22,131	22,131
Number of investors	1,871	1,871	1,871	1,871	1,871	1,782	1,782
R-squared	0.14	0.03	0.07	0.04	0.15	0.08	0.06
Panel B: Active sample							
Δ Perceived conditional hist. return (%)	0.097** (0.044)	0.025*** (0.009)	0.022*** (0.007)	-0.003 (0.007)		0.113* (0.063)	0.048** (0.023)
Lagged perceived conditional hist. return (%)	0.044 (0.062)	0.016* (0.008)	0.022** (0.011)	0.008 (0.011)	0.027 (0.059)	0.033 (0.075)	0.036* (0.019)
Δ Perceived conditional hist. return (%)					0.042 (0.116)		
Time FE Specification	Yes	Yes	Yes	Yes	Yes	Yes Cond. on trading	Yes Cond. on trading
Observations	46,056	46,056	46,056	46,056	46,056	18,855	18,855
Number of investors	1,841	1,841	1,841	1,841	1,841	1,639	1,639
R-squared	0.11	0.03	0.08	0.05	0.11	0.08	0.05

Notes: This table examines the association between the perceived autocorrelation of returns and equity trading decisions among respondents to our main survey based on investor-month level estimations. The “relative net buying volume” is defined as the percent ratio of net purchases of equity over a given month relative to end-of-previous-month financial wealth. The “relative gross buying volume” and the “relative gross selling volume” capture gross purchases and sales relative to end-of-previous-month financial wealth. “Probability net buy” and “probability trade” are dummy variables, multiplied by 100, indicating whether a respondent has a relative net buying volume of at least 1 pp or of at least 1 pp in absolute terms, respectively. The “perceived conditional hist. return” is the return an investor would expect if his or her return expectations were exclusively based on the investor’s prior perceived historical autocorrelation of returns, assuming accurate beliefs about the current realized return over the previous 12 months. We include the change in this return prediction from the end of the previous month to the end of the current month as well as the level of this return prediction at the end of the previous month. The transaction data span the period from December 2014 until August 2019 (directly before the survey period) for the treatment group and until including January 2020 for the control group. Panel A reports results for the full sample, while Panel B focuses on observations for which the corresponding investor conducted at least one equity trading transaction over the preceding three months. All specifications control for month-year fixed effects and lagged log financial wealth with the bank, the lagged equity share, dummies indicating a lagged equity share of 0% or of 100% as well as all non-trading related variables from the baseline set of controls measured at the time of the survey described in Appendix A.1.2. Standard errors are two-way clustered by investor and trading month and are presented in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table 5: Treatment effects on trading responses to the Covid-19 crash: Full sample

	Probability net buy (%)	Relative net buying volume (%)	Relative gross buying volume (%)	Relative gross selling volume (%)	Probability trade (%)	Probability net buy (%)	Relative net buying volume (%)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Crash × Treatment [β_1]	2.521 (2.042)	0.296** (0.118)	0.238* (0.135)	0.003 (0.126)	1.664 (2.557)	8.603** (3.919)	1.576** (0.609)
Crash × Treatment × 1 (Predicted exp. adjustment > 0) [β_2]	-4.915** (2.314)	-0.361** (0.143)	-0.333* (0.178)	-0.030 (0.139)	-4.365 (3.044)	-11.309** (4.403)	-1.911*** (0.585)
Crash × 1 (Predicted exp. adjustment > 0) [β_3]	2.642 (1.786)	0.125 (0.113)	0.153 (0.117)	0.072 (0.081)	2.931* (1.738)	5.291 (3.725)	1.090** (0.469)
p-value($\beta_1 + \beta_2 = 0$)	0.007	0.420	0.241	0.541	0.031	0.355	0.444
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Specification						Cond. on trading	Cond. on trading
Observations	71,223	71,223	71,223	71,223	71,223	8,612	8,612
Number of investors	1,843	1,843	1,843	1,843	1,843	1,288	1,288
R-squared	0.16	0.07	0.14	0.07	0.16	0.27	0.21

Notes: This table presents treatment effects of our intervention on trading responses to the Covid-19 crash among respondents to our main survey based on investor-week level estimations of specification 5. The “relative net buying volume” is defined as the percent ratio of net purchases of equity over a given month relative to end-of-previous-month financial wealth. The “relative gross buying volume” and the “relative gross selling volume” capture gross purchases and sales relative to end-of-previous-month financial wealth. “Probability net buy” and “probability trade” are dummy variables, multiplied by 100, indicating whether a respondent has a relative net buying volume of at least 1 pp or of at least 1 pp in absolute terms, respectively. “1 (Predicted exp. adjustment > 0)” indicates whether an investor would upward adjust his or her return expectations in response to the crash if return expectations were exclusively based on the investor’s prior perceived historical autocorrelation of returns, assuming accurate beliefs about the change in the realized return over the crash period. “Crash” takes value one for the weeks from 17th February until 13th March 2020 and zero otherwise. The sample spans the time from 3rd June 2019 until 13th March 2020. The specifications also include interactions of a dummy indicating the “Post-survey pre-crash” from 16th September 2019 until 16th February 2020 with a treatment dummy, with the “1 (Predicted exp. adjustment > 0)” dummy, and with both of these dummies. The estimations are based on the full sample, not restricting by the degree of trading activity. All specifications control for individual fixed effects, week-year fixed effects and lagged log financial wealth with the bank. Standard errors are two-way clustered by investor and trading week and are presented in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table 6: Additional results on treatment effects on trading behavior

	Relative net buying volume (%)				
	(1)	(2)	(3)	(4)	(5)
Post-survey × Treatment	-0.030 (0.032)				
Crash × Treatment		0.014 (0.068)			
Post-survey pre-crash × Treatment		-0.037 (0.031)	-0.042 (0.031)	0.007 (0.064)	-0.071 (0.042)
Post-survey pre-crash × Treatment × ‡ (Pred. exp. adjustment during crash > 0)				-0.062 (0.059)	
Post-survey pre-crash × ‡ (Pred. exp. adjustment during crash > 0)				0.035 (0.045)	
Post-survey pre-crash × Treatment × ‡ (Pred. exp. adjustment during post-survey pre-crash > 0)					0.116 (0.072)
Post-survey pre-crash × ‡ (Pred. exp. adjustment during post-survey pre-crash > 0)					-0.057 (0.053)
Individual FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Sample period	Full	Full	Excl. crash	Excl. crash	Excl. crash
Observations	71,223	71,223	64,336	64,336	63,432
Number of investors	1,843	1,843	1,842	1,842	1,813
R-squared	0.07	0.07	0.07	0.07	0.07

Notes: This table presents additional results on treatment effects of our intervention on trading behavior. All estimations are based on investor-week level datasets. The “relative net buying volume” is defined as the percent ratio of net purchases of equity over a given month relative to end-of-previous-month financial wealth. “Post-survey” takes value one for the weeks from 16th September 2019 until 13th March 2020 and zero otherwise. “Crash” takes value one for the weeks from 17th February until 13th March 2020. “Post-survey pre-crash” takes value one for the weeks from 16th September 2019 until 16th February 2020. “‡ (Pred. exp. adjustment during crash > 0)” indicates whether an investor would upward adjust his or her return expectations in response to the crash if return expectations were exclusively based on the investor’s prior perceived historical autocorrelation of returns, assuming accurate beliefs about the change in the realized return over the crash period. “‡ (Pred. exp. adjustment during post-survey pre-crash > 0)” is a similarly defined indicator for the period reaching from before a respondent took the survey until just before the crash. The samples span the time from 3rd June 2019 until 13th March 2020 (Columns 1-2) and the time from 3rd June 2019 until 16th February 2020 (Columns 3-5). All specifications control for individual fixed effects, week-year fixed effects and lagged log financial wealth with the bank. Standard errors are two-way clustered by investor and trading week and are presented in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table 7: Sources of beliefs: Memory and return predictability by valuation ratios

	Mean-reverter (diff. < -4)					
	(1)	(2)	(3)	(4)	(5)	(6)
Knowledge of research contributed to estimates	3.118 (13.834)					4.295 (14.404)
Other financial variables contributed to estimates		-18.712* (10.619)				-18.991* (10.489)
High knowledge of financial market theories			-0.158 (9.200)			-0.108 (9.830)
Thought of specific reversal episode				20.393*** (6.759)		20.445*** (6.964)
Thought of specific non-reversal episode					7.901 (8.539)	1.489 (8.595)
Observations	208	208	208	208	208	208
R-squared	0.13	0.15	0.13	0.17	0.14	0.19

Notes: This table examines memory and beliefs about return predictability by valuation ratios as two potential sources of beliefs in mean reversion among German retail investors participating in an additional descriptive survey. The outcome is a dummy variable, multiplied by 100, indicating whether the respondent is classified as a mean reverter (perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios lower than -4 pp). The main independent variables are the following: dummy variables indicating whether the respondent reports that knowledge of results from finance research or memory of past developments of financial variables other than the return (such as valuation ratios) contributed to her estimate of the historical autocorrelation (Columns 1 and 2); a dummy indicating whether the respondent reports a relatively high level of knowledge of financial market theories (Column 3); a dummy indicating whether the respondent thought of at least one specific past reversal episode when estimating the historical autocorrelation as measured in an open-ended question (Column 4); and a dummy indicating whether the respondent thought of at least one specific past non-reversal episode (Column 5). We classify the historical episodes as explained in the main text, giving us the following classification: reversal episodes: the Covid-19 downturn (2020), the downturn associated with Brexit (2018), the euro crisis (2011), the Global Financial Crisis (2007-9), the downturn associated with the Iraq War (2003), the downturn associated with the Gulf War (1990), and the downturn following the Black Monday (1987); non-reversal episodes: the downturn associated with the terror attacks on 9/11 (2001) and the burst of the dot-com bubble (2000-3). All estimations include a parsimonious set of controls (gender, age, employment status, education, household income, financial wealth, investment experience, financial literacy, and attention to DAX developments). Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

For online publication only:

Beliefs About the Stock Market and Investment Choices: Evidence from a Survey and a Field Experiment

Overview of the online appendix

Section A.1 provides additional details on the surveys and data.

Section A.2 provides background on returns of the German stock market.

Section A.3 contains additional evidence on beliefs.

Section A.4 contains additional evidence on trading behavior.

Section A.5 demonstrates persistence of treatment effects in a four-week follow-up.

Section A.6 provides additional details for the interpretation of effect sizes.

Section A.7 presents an additional experiment on beliefs about the PD ratio.

Section B.1 provides the instructions of the main survey.

Section B.2 provides the instructions of the four-week follow-up survey.

A Supplementary details and analyses

A.1 Additional details on surveys and data

A.1.1 Details on other surveys

Next to our main and follow-up surveys conducted between September and November 2019, which are described in detail in Section 2, we ran several additional surveys. Table A.1 provides a complete overview of all our data collections. The additional experimental investor survey run in July and August 2022 is described in detail in Appendix A.7. In the current section we briefly describe our other data collections. The instructions for the additional surveys can be found under <https://drive.google.com/file/d/1IhxTWR5pXyVSR0V6tPfr4y06p7UwygmQ/view?usp=sharing>.

Additional descriptive investor survey In July and August 2022 we ran an additional descriptive survey with retail investors from the same online bank we used in our main experiment. We invited 3,000 investors to participate in this survey, who were selected in the same way as invitees to our original survey. Out of these, 227 respondents eventually completed our survey, corresponding to a response rate of 7.6%. Due to changes in the data protection policies at the bank that came into effect after we had run our main survey, we were not allowed to link these new survey data to administrative account data. The survey starts with a question on investors' familiarity with different financial variables. We then elicit respondents' perceived autocorrelation of aggregate stock returns using the same elicitation format as in our main survey, which is followed by questions on respondents' thoughts when reporting their perceived autocorrelation. The survey ends with a measure of respondents' knowledge of research results in finance and questions on a range of background characteristics. As in our main survey, we drop respondents with extreme response times, which results in a sample of 208 respondents. Table A.2 Column 2 provides a range of basic summary statistics for this sample. The composition of the sample is very similar to the sample from our main analysis, the main difference being that respondents are somewhat wealthier and older on average.

Descriptive general population survey Germany We also ran a survey on a general population sample from Germany in collaboration with the survey company Dynata, which is widely used in the social sciences (Haaland et al., 2023). The content and structure of the survey are very similar to the additional

descriptive investor survey described above. In total, 504 individuals completed our survey, out of which 490 individuals form the final sample after dropping those with extreme response times. Table A.2 Column 4 provides a range of basic summary statistics. The sample is broadly representative of the adult German population, as can be seen from comparison with benchmarks taken from the 2017 wave of the Bundesbank’s Panel of Household Finances (PHF) displayed in Column 3.

Descriptive general population survey US In addition, we ran a survey on a general population sample from the US in collaboration with the survey company Lucid, which is widely used in the social sciences (Haaland et al., 2023). The content and structure of the survey are very similar to the additional descriptive investor survey described above. 508 individuals completed our survey. After dropping those with extreme response times, our final sample consists of 493 individuals. Table A.2 Column 6 provides summary statistics and shows that the sample is broadly representative of the adult US population (see the benchmarks taken from the 2019 American Community Survey in Column 5).

A.1.2 Details on control variables

To account for small imbalances across treatment arms in our main experiment (see Table 1) and to increase power, we include a set of control variables in all our estimations. Our baseline set of control variables used for our main experiment is the following: demographics: a dummy for being female, a dummy for above-median age, dummies for being employed and for holding a university degree, the logs of the respondent’s household’s net income and net wealth¹; survey measures of investor behavior: dummies for different levels of trading experience, financial literacy, attention to the DAX, investment horizon, sources of financial information and risk tolerance; administrative measures of holdings with the bank at the time of the survey: the log of total financial wealth held with the bank, the equity share and dummies for holding an equity share of 0% or 100%, the share of other securities, dummies for number of equity trades over the previous three months and length of relationship with the bank; technical controls: dummies for passing an attention screener, self-reported survey difficulty, use of external information in the response, experiencing a technical issue and taking the survey on a mobile phone. None of our results are sensitive to the exact set and construction of control variables included. We use a very similar set of

¹We elicit net wealth and income using survey questions with categorical response options, and construct continuous variables based on the mid-points of the corresponding bins. The lowest response categories are “no net wealth” or “no income”, for which we assign the value zero. We construct the logs of the variables after adding the value one.

control variables in the estimations based on the additional experimental investor survey described in Appendix A.7, the main difference being that all variables are constructed based on survey measures, as no administrative account data are available for any of our additional data collections. We use a more parsimonious set of control variables in our estimations based on the additional descriptive investor survey described in Appendix A.1.1 above due to its smaller sample size.

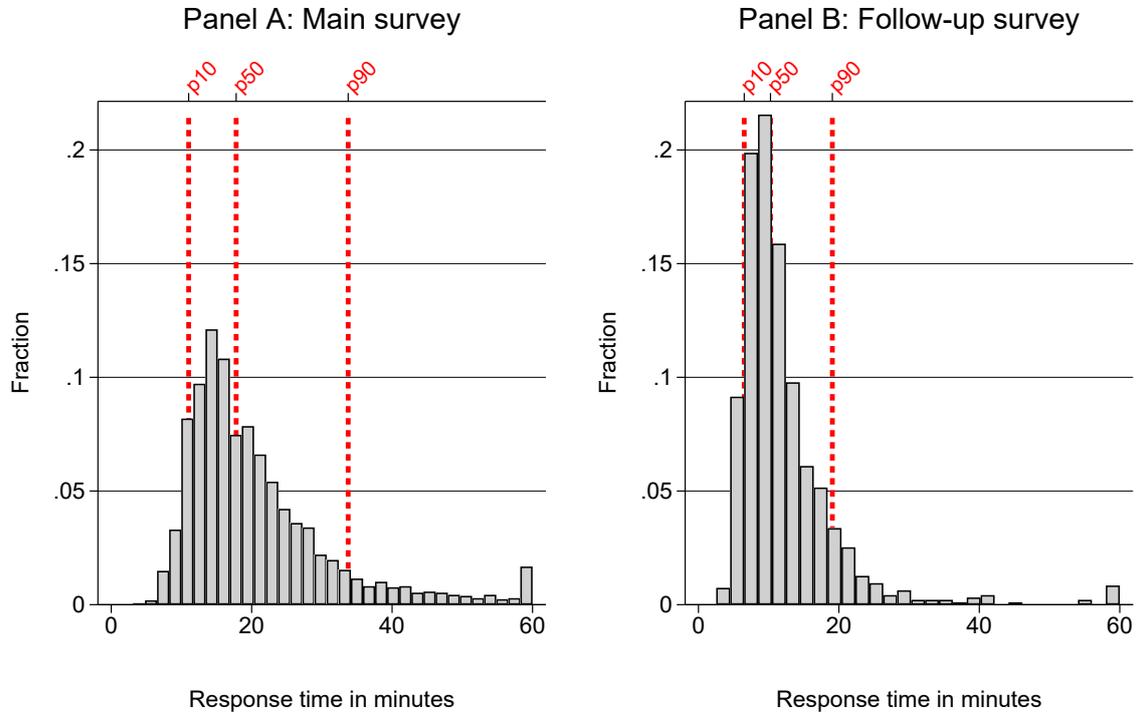
A.1.3 Distributions of response times

Figure A.1 displays the distribution of the response time for the main survey (Panel A) and the follow-up survey (Panel B) among all respondents completing the respective survey, i.e., before restricting the sample (among others by response time).

A.1.4 Selection into the survey

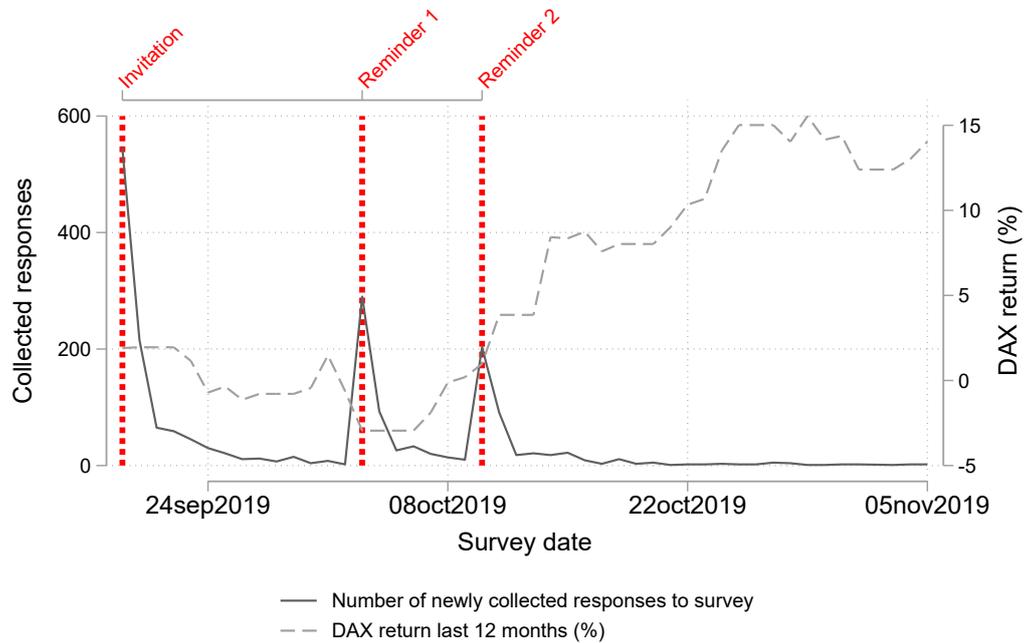
Table A.3 compares our main survey sample with a random sample of 3,701 investors drawn from the client pool of the same bank using the same criteria as for the main survey sample. Our sample over-represents less wealthy investors and those with lower trading frequency. We also study whether the realized return of the DAX affects participation in our survey. Figure A.2 highlights that participation in the survey peaks on days on which respondents receive the invitation or receive one of the two reminders we sent out two and three weeks after the initial invitation, and quickly drops in the respective subsequent days. There is no noticeable relationship with the realized return of the DAX, suggesting that the realized return does not affect the propensity to participate in the survey. Table A.4 highlights that also the composition of the sample in terms of observable respondent characteristics does not systematically vary with the realized return of the DAX.

Figure A.1: Response times to the main and the follow-up survey



Notes: This figure displays histograms of the response time for the main survey (Panel A) and the follow-up survey (Panel B) among all respondents completing the survey, before restricting the sample. The vertical dashed lines indicate the 10th percentile, the median, and the 90th percentile of the response time for each survey. The response time is winsorized at 60 minutes.

Figure A.2: Selection into the survey: Realized return and number of survey responses



Notes: This figure examines whether the realized return over the 12 months before a given day in the main survey period is systematically associated with the tendency to participate in the survey. The y-axis on the left indicates the number of collected survey responses on a given day. The y-axis on the right indicates the realized 12-months return of the DAX as of the end of the previous day. The vertical dashed lines indicate the day on which the invitation email was sent out (19th September) and the days on which the respondents received a first reminder (3rd October 2019) or a second reminder (10th October 2019). The sample is our main survey sample.

Table A.1: Overview of data collections

Data collection	Sample	Treatments	Main variables
Main survey (September-November 2019): Autocorrelation experiment	Retail investors at German online bank ($n = 1,961$)	Information on empirical autocorrelation of returns and control	Beliefs about autocorrelation, return expectations, trading decisions (account data)
Follow-up survey (September-November 2019)	Retail investors at German online bank from main survey ($n = 903$)	Repeated information on empirical autocorrelation of returns and control	Beliefs about autocorrelation, return expectations, trading decisions (account data)
Additional descriptive investor survey (July-August 2022)	Retail investors at German online bank ($n = 208$)	None	Beliefs about autocorrelation, sources of beliefs, and familiarity with financial variables
Additional experimental investor survey (July-August 2022): PD experiment	Retail investors at German online bank ($n = 693$)	Information on empirical correlation between valuation ratios and subsequent returns and control	Beliefs about link between valuation ratios and subsequent returns, return expectations
Descriptive general population survey Germany (July-August 2022)	Online panel in collaboration with Dynata ($n = 490$)	None	Beliefs about autocorrelation, return expectations
Descriptive general population survey US (July-August 2022)	Online panel in collaboration with Lucid ($n = 493$)	None	Beliefs about autocorrelation, return expectations

Table A.2: Summary statistics other surveys

	Means					
	(1) PHF 2017 stock- holders	(2) Descriptive investor survey 2022	(3) PHF 2017 all	(4) German representative survey 2022	(5) ACS 2019 all	(6) US representative survey 2022
Female	0.49	0.15	0.51	0.50	0.51	0.54
Age	50.55	47.73	50.12	53.25	47.78	48.77
University	0.36	0.55	0.19	0.22	0.31	0.38
Employed	0.65	0.72	0.60	0.55	0.62	0.50
Household income	3,808	4,301	2,790	2,925	8,050	6,944
Household net wealth	361,783	444,172	171,161	156,042		278,942
Stockowner	1.00	1.00	0.17	0.41		0.60
Average monthly equity trades		0.69		0.19		0.34
Risk tolerance (1-7)		4.38		2.64		3.44
Trading experience (years)		14.76		5.50		7.40
Financial literacy score (0-3)		1.82		1.06		1.34
Follow stock market developments (1-7)		4.37		2.92		3.43
Investment horizon ≥ 5 years		0.53		0.44		0.30
Perceived return last 12 months		-5.01		4.11		9.05
Confident in perceived return		0.64		0.38		0.43
Expected return next 12 months		6.68		11.47		32.40
Confident in expected return		0.45		0.36		0.43
Perceived mean hist. ret. intervals		9.46		8.30		11.95
Perceived diff. gain-loss historical		-4.31		1.55		1.87
Extrapolator (diff. ≥ 4)		0.16		0.30		0.34
Mean-reverter (diff. < -4)		0.49		0.18		0.19
Observations		208		490		493

Notes: This table shows summary statistics for the additional descriptive retail investor survey (Column 2), for the descriptive survey on a representative sample from Germany (Column 4), and for the descriptive survey on a representative sample from the US (Column 6). All three surveys were run in July and August 2022. The table also includes benchmarks from the German population of individuals participating in the stock market as measured in the 2017 wave of the Bundesbank's Panel of Household Finance (PHF, Column 1), the overall German population as measured in the PHF (Column 3), and the overall US population as measured in the 2019 wave of the American Community Survey (Column 5). Variables on income and wealth are expressed in euro terms (Columns 1-4) or in dollar terms (Columns 5-6). "Household income" indicates monthly household income after taxes for the German samples (Columns 1-4) and monthly household income before taxes for the US samples (Columns 5-6), reflecting differences in the availability of benchmark data.

Table A.3: Selection into the survey: Comparison with random sample

	(1) Survey Sample: Mean	(2) Survey Sample: SD	(3) Random Sample: Mean	(4) Random Sample: SD	(5) p-value (1) = (3)
Female	0.16	0.37	0.22	0.41	0.000
Age	45.24	14.15	52.02	15.25	0.000
Employed	0.77	0.42	0.61	0.49	0.000
Risk attitude (1-5)	4.25	1.18	4.37	1.18	0.000
Total financial wealth at bank	55,272	98,312	82,216	142,817	0.000
Portfolio value at bank	43,970	87,671	63,144	117,574	0.000
Equity Share	0.73	0.45	0.70	0.30	0.415
Average monthly trades	1.88	3.50	3.55	10.88	0.000
Average monthly equity trades	1.73	3.29	2.65	7.27	0.000
Observations	1,961	3,701	1,961	3,701	

Notes: This table shows summary statistics for our main survey sample (Columns 1-2) and a sample randomly drawn from the bank's client pool (Columns 3-4). Column 5 provides the p-values for a test for differences in means between the survey and the random sample. Gender, age, employment status and risk attitude are based on data provided by the bank. Variables on wealth and portfolio holdings are expressed in euro terms. For the survey sample, financial wealth at the bank, portfolio value at the bank, and equity holdings at the bank are measured in the month prior to the survey. Average monthly overall trades and equity trades are measured over the three months preceding the survey. For the random sample, we measure financial wealth at the bank, portfolio value at the bank, and equity holdings at the bank in August 2019 and report the average monthly equity trades in June, July and August 2019.

Table A.4: Selection into the survey: Realized return and composition of respondents

	DAX return last 12 months (%)		
	(1)	(2)	(3)
Female	-0.030 (0.190)	-0.017 (0.170)	0.021 (0.170)
Age	-0.001 (0.006)	-0.006 (0.005)	-0.004 (0.005)
Employed	0.001 (0.171)	-0.049 (0.151)	-0.051 (0.149)
Risk tolerance (1-7)	-0.013 (0.060)	0.009 (0.053)	0.016 (0.052)
Log(Total financial wealth with bank)	-0.132 (0.164)	-0.147 (0.145)	-0.135 (0.144)
Log(Portfolio value at bank)	0.071 (0.154)	0.070 (0.137)	0.066 (0.136)
Equity share in total financial wealth	-0.484 (0.375)	-0.041 (0.322)	0.023 (0.322)
Average monthly trades	-0.050 (0.097)	0.049 (0.084)	0.039 (0.084)
Average monthly equity trades	0.059 (0.104)	-0.045 (0.090)	-0.035 (0.090)
Time since last contacted		0.463*** (0.032)	0.193*** (0.057)
(Time since last contacted) ²			0.019*** (0.003)
p(Coeffs. on predictor vars. jointly equal zero)	0.871	0.264	0.424
Observations	1,961	1,961	1,961
R-squared	0.00	0.22	0.25

Notes: This table examines whether the realized return over the 12 months before a given day in the main survey period is associated with selection of respondents with systematically different characteristics into the survey. It displays regressions of the realized return over the 12 months until the day just before a respondent took the survey on a set of respondent characteristics. Financial wealth at the bank, portfolio value at the bank, and equity holdings at the bank are measured in the month prior to the survey. Average monthly overall trades and equity trades are measured over the three months preceding the survey. The variable “Time since last contacted” indicates the number of days since the invitation email was sent out (19th September) or the time since receiving a first reminder (3rd October 2019) or a second reminder (10th October 2019), depending on which one is closest to the date the respondent took the survey. The table includes the p-value of an F-test of the Null hypothesis that the coefficients of the included predictors except the controls for time since the respondent was last contacted are jointly equal to zero. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

A.2 Background on the German stock market

In this appendix we provide some background on the dynamics of returns and return predictability in the German stock market.

Figure A.3 presents realized (overlapping) past-12-month returns of the DAX for all months in the period from December 2014 until March 2020 – the time for which administrative account data for our main investor sample are available.

Table A.5 Panel A provides evidence on (i) the autocorrelation of aggregate returns and (ii) the correlation of realized past-12-month returns with the current price-dividend ratio for the overall German stock market over the period 1969-2021. We focus on the full German stock market instead of the DAX for this exercise to have a more meaningful measure of the price-dividend ratio than would be available for the DAX with its low number of constituent firms. The table highlights that, historically, 12-month-ahead returns do not vary systematically with realized returns over the previous 12 months. Moreover, realized past-12-month returns are only weakly correlated with the current price-dividend ratio, highlighting that realized 12-month returns reflect short-term fluctuations rather than (the often longer-term) swings in valuation ratios. Thus, observing a high realized 12-month return provides only a weak signal about the level of the current price-dividend ratio.

Panel B analyzes the predictability of future returns based on the current price-dividend ratio over the period 1969-2021. The correlation coefficients shown in Panel B.1 give a first indication that the price-dividend ratio is only weakly negatively related to 12-month-ahead returns while the association of the price-dividend ratio with 12-month-ahead dividend growth is fairly large and positive. Since the price-dividend ratio must – under rational expectations as well as ex-post – predict future returns or future cash flows (Campbell and Shiller, 1988), we can decompose the variance of the price-dividend ratio into variance explained by future returns and variance explained by future dividend growth. We compute the shares as in Cochrane (2008) using regressions of log returns, r , log dividend growth, Δd , and the log

price-dividend ratio, pd , in t on the log price-dividend ratio in $t - 1$:

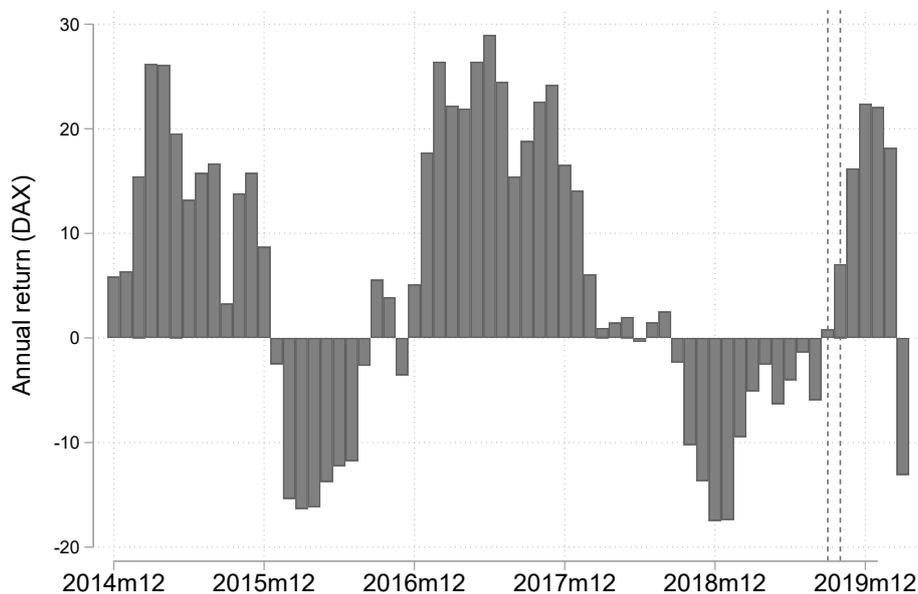
$$r_t = \mu_r + \beta_r pd_{t-1}$$

$$\Delta d_t = \mu_d + \beta_d pd_{t-1}$$

$$pd_t = \alpha + \rho pd_{t-1}$$

The linearized return identity $r_t \approx k_0 + k_1 pd_t - pd_{t-1} + \Delta d_t$ (where $k_1 = \overline{PD}/(1 + \overline{PD})$ with \overline{PD} denoting the mean price-dividend ratio) implies that $\beta_r \approx k_1 \rho - 1 + \beta_d \Leftrightarrow \frac{\beta_r}{k_1 \rho - 1} - \frac{\beta_d}{k_1 \rho - 1} \approx 1$, where the absolute values of the two fractions on the left-hand side can be understood as shares of the variation in the price-dividend ratio that can be explained by variation in either returns or dividend growth. The results are shown in Panel B.2. Most of the variation in the price-dividend ratio (59%) can be explained by variation in future cash flow growth, confirming earlier results for the German context (Rangvid et al., 2014). Running direct regressions of cumulative weighted future returns and dividend growth on the log price-dividend ratio (see Cochrane, 2008) yields similar results.

Figure A.3: 12-month returns of the German Stock index (DAX) over the sample period



Notes: This figure shows the return of the DAX over the previous 12 months for each month (i.e., overlapping periods of 12 months) in the period for which account data for the respondents to our main survey are available. The dashed lines mark the time span in which investors responded to the survey.

Table A.5: Price-dividend ratios and returns in the German stock market

Panel A: Correlation of returns and PD ratio		
Correlation of past 12-month return and PD ratio	$Corr(PD_t, R_t)$	0.0925 (0.1411)
Correlation of past 12-month return and future 12-month return	$Corr(R_t, R_{t-1})$	-0.1271 (0.1068)
Panel B: Predictability by the PD ratio		
Panel B.1: Correlation		
Correlation of PD ratio and future 12-month return	$Corr(PD_t, R_{t+1})$	-0.0513 (0.139)
Correlation of PD ratio and future 12-month dividend growth	$Corr(PD_t, \Delta D_{t+1})$	0.2070 (0.145)
Panel B.2: Variance decomposition		
PD Ratio Variance decomposition %	returns	dividend growth
	41.60	58.64

Notes: This table shows statistics on the joint dynamics of the price-dividend (PD) ratio, returns and dividend growth in the German stock market. All data are taken from Datastream for a sample from 1969 to 2021. Price-dividend ratios and dividends are computed from returns with and without dividends as in Cochrane (2008). Panel A shows the correlation between the PD ratio at the end of year t with returns in year t as well as the correlation of returns in t and $t - 1$. Numbers in parentheses are bootstrapped standard-errors. Panel B.1 shows correlations between the PD ratio and future returns and dividend growth, respectively. Panel B.2 shows the variance decomposition of the PD ratio into variance explained by future returns and future dividend growth, as discussed in Section A.2 above.

A.3 Additional evidence: Beliefs

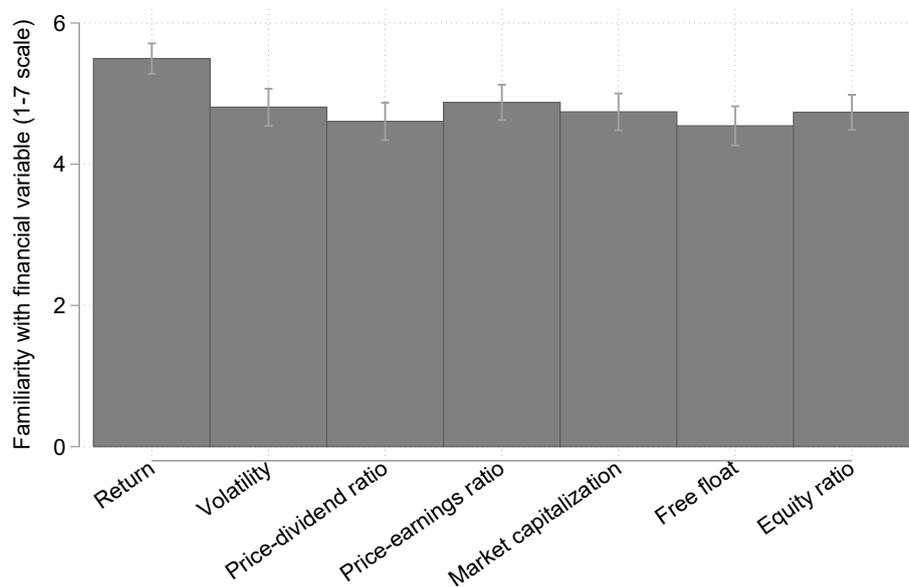
In this appendix we present additional evidence on respondents' beliefs.

Figure A.4 displays the average degree of familiarity with different financial indicators among respondents to an additional descriptive investor survey. Figure A.5 presents the distributions of prior beliefs about the returns over the 12 months before and the 12 months after the survey among respondents to the main survey. Figure A.6 displays the distributions of beliefs about the return of the aggregate stock market over the 12 months before the survey among respondents to the main survey and among additional samples (an additional descriptive survey among German investors as well as general population surveys from Germany and the US). Figure A.7 demonstrates robustness of the relationship between actual prior 12-month ahead return expectations and predicted expectations based on the respondent's prior perceived autocorrelation of returns shown in Figure 3 to calculating the predicted return based on fitted linear or quadratic functional forms instead of simply using the respondent's belief about the conditional mean in the relevant bin. Figure A.8 displays distributions of beliefs about the autocorrelation of aggregate stock returns across different samples (an additional descriptive survey among German investors as well as general population surveys from Germany and the US). Figure A.9 displays factors contributing to respondents' estimates of the historical autocorrelation of returns among respondents to an additional descriptive investor survey.

Table A.6 displays multivariate regressions of prior beliefs about the autocorrelation of returns on a set of co-variates. Table A.7 displays statistics on beliefs about the return of the aggregate stock market over the 12 months before the survey among respondents to the main survey and among additional samples (an additional descriptive survey among German investors as well as general population surveys from Germany and the US). Table A.8 demonstrates the robustness of the findings presented in Table 2 to varying the cutoffs used to define extrapolators, neutrals and mean reverters. Table A.9 shows average and heterogeneous treatment effects on additional qualitative measures of respondents' perceived autocorrelation included in the four-week follow-up survey. Table A.10 demonstrates robustness of the treatment effects on updating of return expectations displayed in Table 3 to controlling for interaction terms of the treatment indicator with a set of covariates. Table A.11 demonstrates robustness of the treatment effects on updating of return expectations displayed in Table 3 to calculating the perception gap

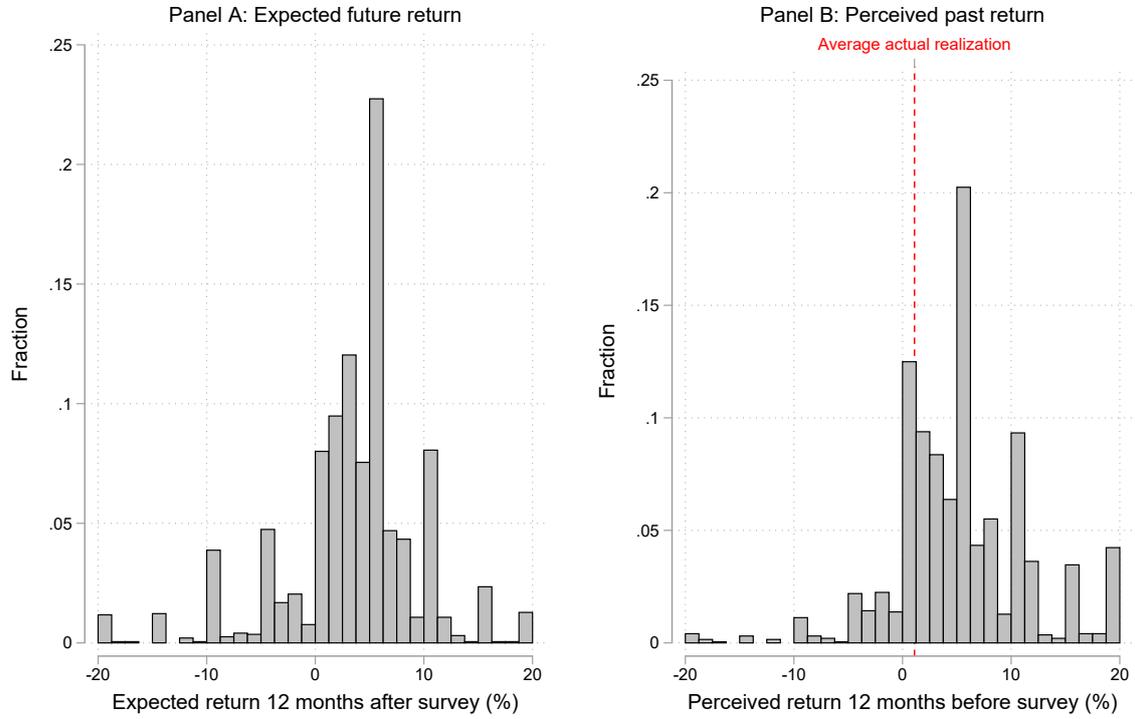
based on fitted linear or quadratic functional forms instead of simply using the respondent's belief about the conditional mean in the relevant bin and the corresponding signal. Table A.12 displays measures of disagreement in expectations within each treatment arm both before and after the intervention. Table A.13 examines how thoughts of specific historical episodes when estimating the historical autocorrelation are related to respondents' lifetime experiences among respondents to an additional descriptive investor survey.

Figure A.4: Familiarity with financial variables



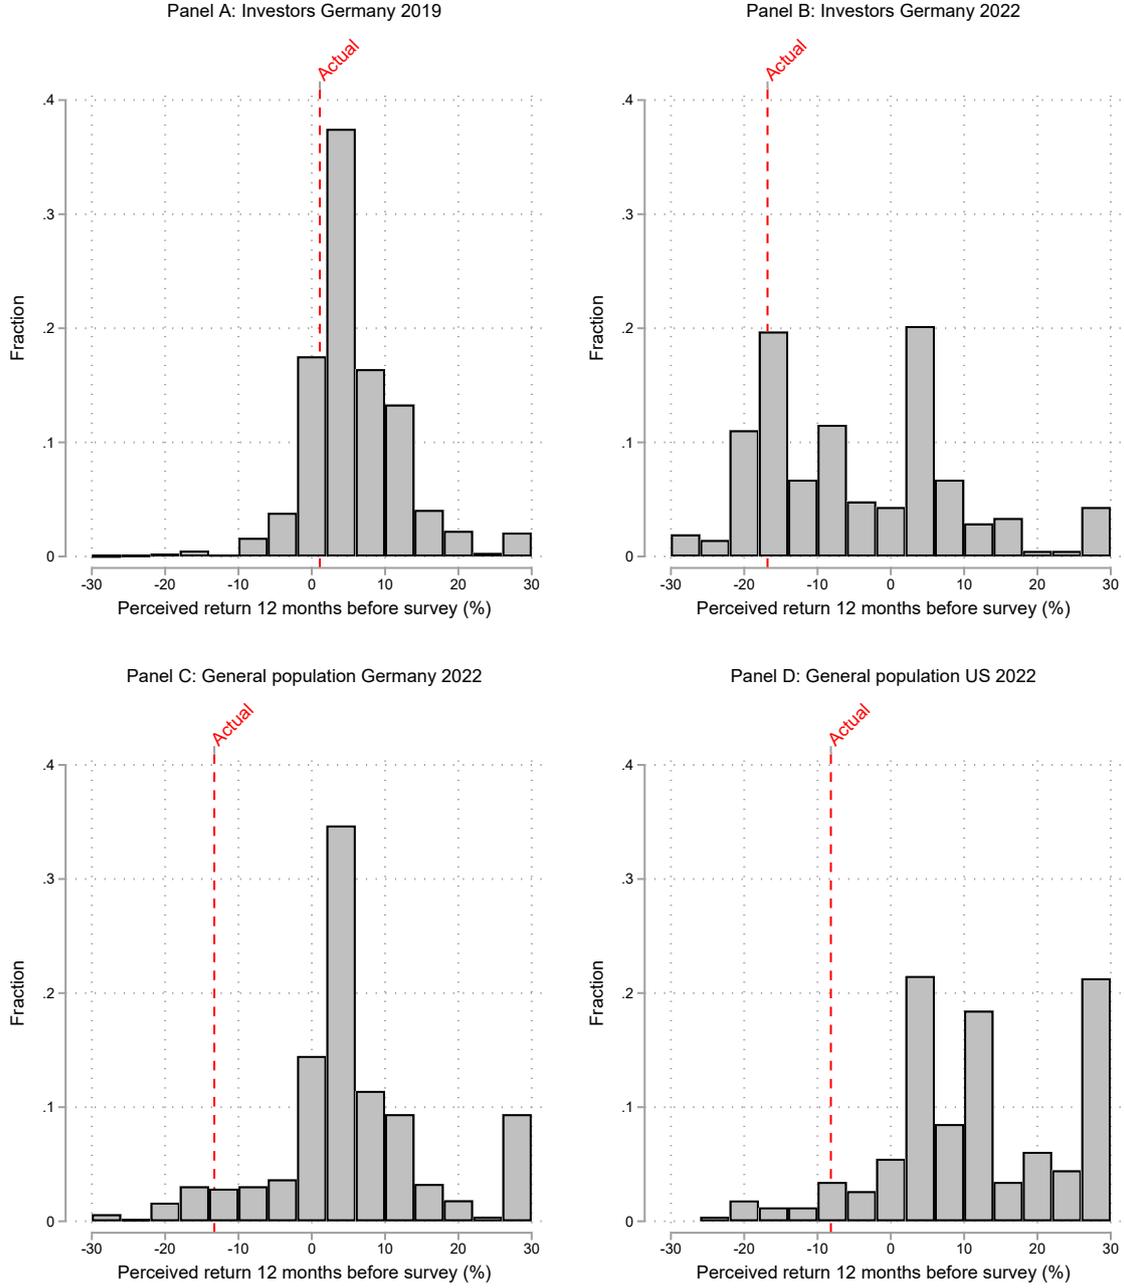
Notes: This figure displays the average degrees of familiarity with different financial concepts as measured on 7-point categorical scales in an additional descriptive survey of German retail investors.

Figure A.5: Prior beliefs about the return 12 months after and before the survey



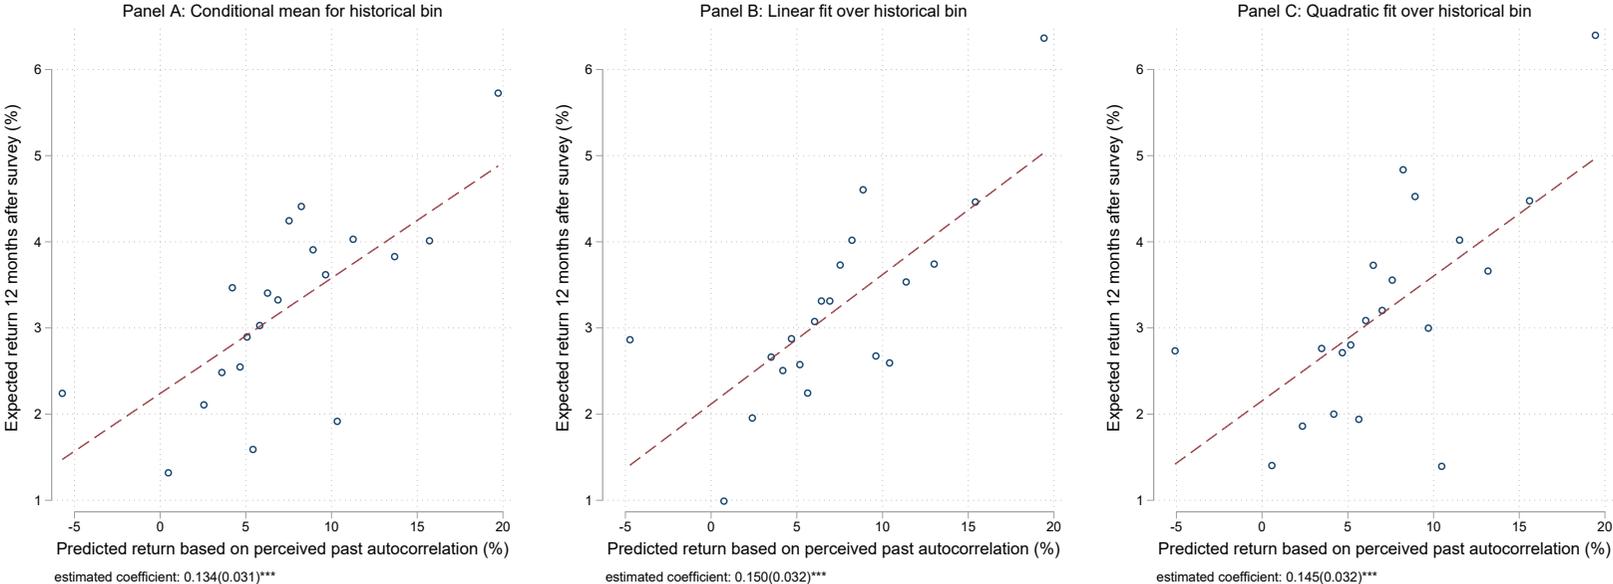
Notes: This figure displays histograms of prior beliefs about the return of the German stock market over the 12 months after (Panel A) and the 12 months before the survey (Panel B) among respondents to our main survey. Our sample focuses on respondents with a prior expected return over the next 12 months between -20% and 20%. The perceived return over the last 12 months is winsorized at -20% and 20%. The dashed red line in Panel B shows the average actual return realization over the sample period of 1.1%.

Figure A.6: Beliefs about the return 12 months before the survey across types of agents and countries



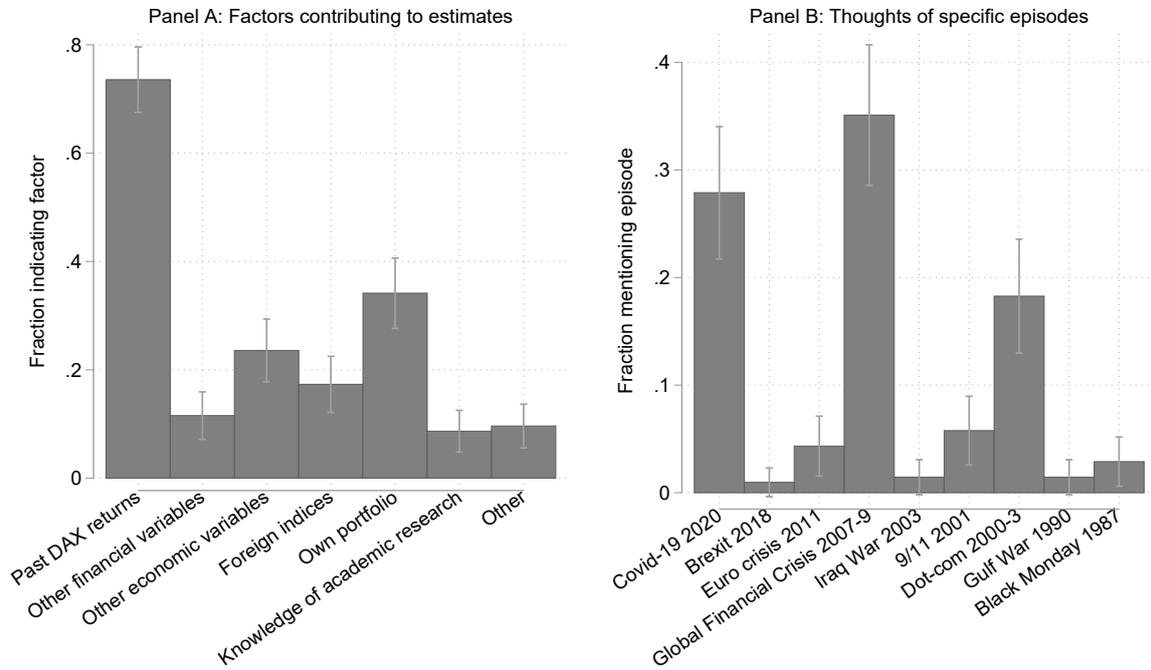
Notes: This figure displays histograms of beliefs about the return of the aggregate stock market over the 12 months before the survey among German retail investors participating in our main survey (Panel A), German retail investors participating in an additional descriptive survey (Panel B), respondents to a general population survey from Germany (Panel C), and respondents to a general population survey from the US (Panel D). The main survey was run between September and November 2019. All other surveys were run in July and August 2022. The perceived return over the last 12 months is winsorized at -30% and 30%. The dashed red lines show the average actual return realization over the sample period for each data collection.

Figure A.7: Binned scatter plot of prior expected 12-month-ahead return vs predicted return expectation based on perceived historical autocorrelation: Robustness to functional form



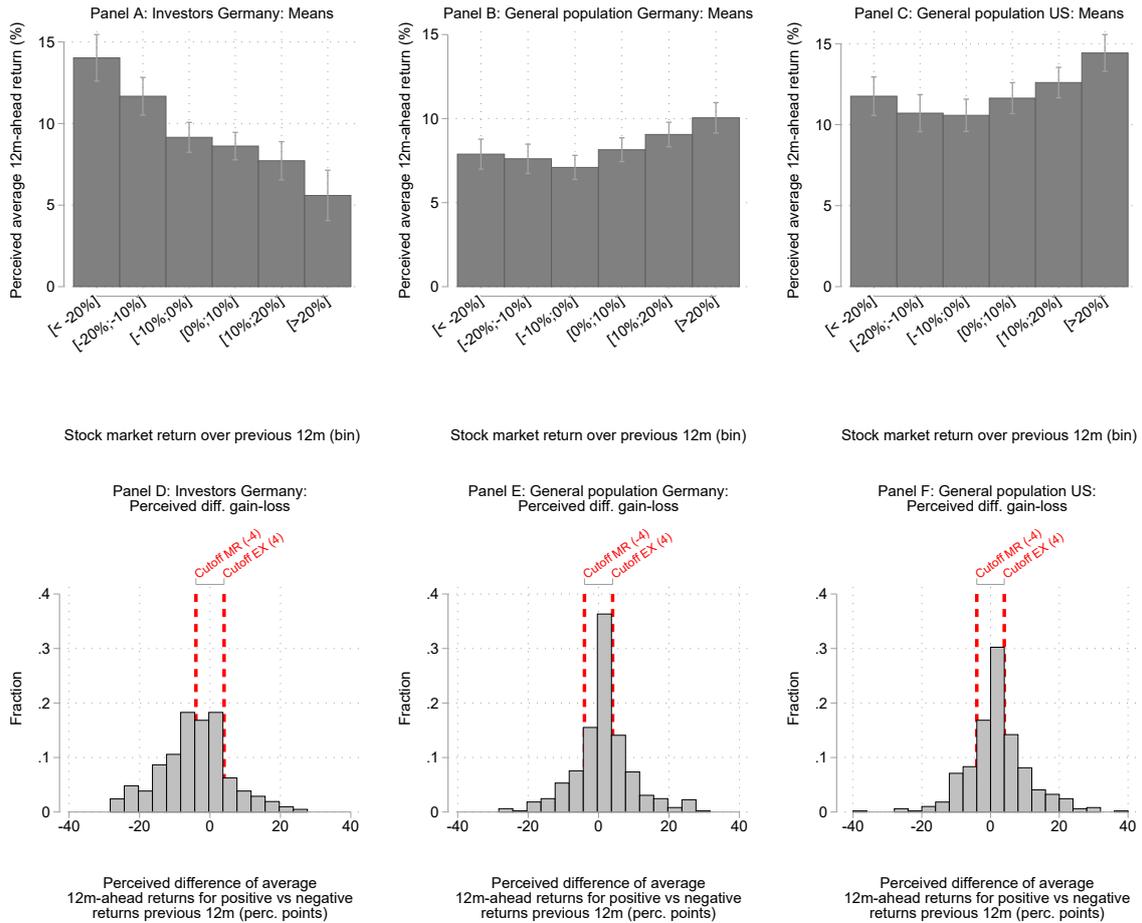
Notes: This figure shows binned scatter plots of respondents' prior expected return over the 12 months after the survey against the predicted return based on a respondent's perceived historical autocorrelation of returns and the respondent's perceived return over the 12 months before the survey, among respondents to our main survey. Panel A calculates the predicted return as the respondent's reported conditional mean year-ahead return for the relevant historical bin, reproducing Panel A of Figure 3. Panel B instead calculates the prediction fitting linear functions between the different historical bins (i.e., lines connecting the points characterized by the historical conditional mean *past 12-months* return and the respondent's corresponding predicted conditional mean *12-months-ahead* return between any two neighboring bins). Panel C instead calculates the prediction fitting two quadratic functions (i) for the three return intervals on the left and (ii) for the three intervals on the right as well as a linear function connecting the two middle bins. Both alternative ways of calculating the predicted return (Panels B and C) simply assign the respondent's predicted conditional mean year-ahead return in the two extreme bins in the few cases where the perceived return over the 12 months before the survey is lower than -29.7% or higher than 33.8% (the conditional mean past 12-month returns for the extreme bins). The binned scatter plots partial out the baseline set of controls described in Appendix A.1.2. Robust standard errors are displayed below the different panels.

Figure A.9: Factors contributing to estimates of the historical autocorrelation



Notes: This figure summarizes the factors contributing to respondents' estimates of the historical autocorrelation of returns of the DAX among German retail investors participating in an additional descriptive survey. Panel A displays the fractions of respondents indicating different factors that contributed to their estimates based on a structured survey question. The different factors refer to respondents' memory of the past development of the returns of the DAX, of other financial variables (such as valuation ratios), of other economic variables (such as GDP growth), of the returns of foreign stock market indices, or of their own portfolio, and their knowledge of academic research on financial markets. Panel B displays the fractions of respondents mentioning specific past episodes they thought about when estimating the historical autocorrelation based on an open-ended survey question.

Figure A.8: Beliefs about the autocorrelation of aggregate returns across types of agents and countries



Notes: This figure summarizes beliefs about the autocorrelation of aggregate stock returns in the last 50 years among German retail investors participating in an additional descriptive survey, respondents to a general population survey from Germany, and respondents to a general population survey from the US. The surveys were all run in July and August 2022. The belief elicitation focused on the DAX (Panels A, B, D and E) or on the overall US stock market (Panels C and F). Panels A-C show the sample means of respondents' beliefs about average 12-month-ahead stock returns for six intervals of realized returns over the previous 12 months. Panels D-F show histograms of respondents' perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios, including the cutoffs we use to define mean reverters, neutrals and extrapolators.

Table A.6: Correlates of beliefs

	Perceived diff. gain-loss	Extra- polator (diff. ≥ 4)	Neutral ($-4 \leq$ diff. < 4)	Mean- reverter (diff. < -4)
	(1)	(2)	(3)	(4)
Female	2.534*** (0.590)	9.721*** (2.545)	-1.343 (2.924)	-8.379*** (3.049)
Age	-0.006 (0.019)	-0.182** (0.077)	0.412*** (0.097)	-0.230** (0.099)
Employed	-0.542 (0.534)	-4.287** (2.122)	4.072 (2.537)	0.215 (2.723)
University	-1.556*** (0.440)	-1.616 (1.667)	-3.637* (2.175)	5.253** (2.297)
Log(Household income)	-0.120 (0.077)	-0.222 (0.313)	-0.619 (0.389)	0.841** (0.406)
Log(Fin. wealth with bank)	-0.208 (0.129)	-0.773 (0.520)	-1.067 (0.658)	1.840*** (0.685)
Invest. experience \geq Median	-1.299** (0.524)	-3.833* (2.052)	-5.563** (2.676)	9.396*** (2.784)
Full financial literacy score	-0.654 (0.513)	-2.225 (1.900)	-5.213** (2.521)	7.437*** (2.768)
Follow DAX \geq Median	-0.942** (0.469)	-0.384 (1.720)	-6.661*** (2.214)	7.046*** (2.402)
Mean dep. var.	-4.88	15.66	31.87	52.47
SD dep. var.	9.64	36.35	46.61	49.95
Observations	1,961	1,961	1,961	1,961
R-squared	0.04	0.03	0.02	0.04

Notes: This table shows multivariate regressions of beliefs on covariates among respondents to our main survey. The outcomes are the perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios (Column 1), and dummies, multiplied by 100, for extrapolators (for which this difference is at least 4 pp, Column 2), neutrals (difference at least -4 pp and less than 4 pp, Column 3), and mean reverters (difference less than -4 pp, Column 4). Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.7: Beliefs about the return 12 months before the survey across types of agents and countries

	Perceived return last 12 months				(5) p-value (2) = (3)
	(1) Main investor survey 2019	(2) Descriptive investor survey 2022	(3) German representative survey 2022	(4) US representative survey 2022	
Benchmark	1.06	-16.82	-13.29	-8.18	
Mean	5.30	-4.81	4.99	11.35	0.000
Mean absolute misperception	5.98	13.79	18.96	20.49	0.000
SD	6.79	13.48	11.29	12.96	0.000
p90 – p10	12.00	31.00	31.00	35.00	
Observations	1,961	208	490	493	

Notes: This table displays statistics on beliefs about the return of the aggregate stock market over the 12 months before the survey among German retail investors participating in our main survey (Column 1), German retail investors participating in an additional descriptive survey (Column 2), respondents to a general population survey from Germany (Column 3), and respondents to a general population survey from the US (Column 4). Column 5 presents p-values from tests for equality of means and from a Levene’s test for equality of the cross-sectional variance between the retail investor and the German general population samples from 2022. The benchmark is the actually realized return over the 12 months before a respondent took the survey. The main survey was run between September and November 2019. All other surveys were run in July and August 2022. The perceived return over the last 12 months is winsorized at -30% and 30%.

Table A.8: Manipulation check: Alternative type definitions

	Positive return irrespective of previous return (z)	No sense to buy after high return (z)	Positive return more likely after high return (z)
	(1)	(2)	(3)
Panel A: Baseline			
Treatment × Extrapolator (diff. ≥ 4) (a)	-0.037 (0.112)	0.021 (0.114)	-0.375*** (0.115)
Treatment × Neutral ($-4 \leq \text{diff.} < 4$)	0.078 (0.079)	0.075 (0.080)	-0.084 (0.081)
Treatment × Mean-reverter (diff. < -4) (b)	0.140** (0.059)	-0.155*** (0.060)	-0.114* (0.062)
p-value (a=b)	0.161	0.174	0.047
Panel B: Neutral narrow			
Treatment × Extrapolator (diff. ≥ 3) (a)	-0.067 (0.105)	0.076 (0.109)	-0.327*** (0.112)
Treatment × Neutral ($-3 \leq \text{diff.} < 3$)	0.128 (0.091)	-0.041 (0.091)	-0.076 (0.092)
Treatment × Mean-reverter (diff. < -3) (b)	0.126** (0.056)	-0.101* (0.057)	-0.117** (0.059)
p-value (a=b)	0.106	0.151	0.099
Panel C: Neutral broad			
Treatment × Extrapolator (diff. ≥ 5) (a)	0.004 (0.117)	0.024 (0.123)	-0.432*** (0.119)
Treatment × Neutral ($-5 \leq \text{diff.} < 5$)	0.066 (0.073)	0.034 (0.073)	-0.093 (0.074)
Treatment × Mean-reverter (diff. < -5) (b)	0.140** (0.062)	-0.151** (0.062)	-0.106 (0.066)
p-value (a=b)	0.303	0.206	0.017

Notes: This table shows estimations of the effect of the information treatment on posterior agreement with verbal statements describing beliefs about the autocorrelation of aggregate returns among respondents to our main survey for alternative definitions of belief types. Agreement with the statements is elicited on 7-point categorical scales, and is z-scored using the means and standard deviations in the sample. The statements are: “With an investment in stocks one can expect a positive return, independently of how the stock market has developed in the recent past.” (Column 1); “When the stock market has recently increased it makes no sense to buy stocks.” (Column 2); “When the stock market has recently increased it is more likely that stock returns will be positive over the following time than when the stock market has recently decreased.” (Column 3). The estimations show heterogeneous treatment effects for prior extrapolators, neutrals, and mean reverters. Panel A repeats the results based on the baseline type definition from Table 2. In Panel B, we use a more narrow cutoff to define neutrals (difference in estimated returns for the following year between positive and negative previous return scenarios of at least -3 pp and lower than 3 pp). In Panel C, we use a broader definition of neutrals (difference at least -5 pp and lower than 5 pp). All estimations include the baseline set of controls described in Appendix A.1.2 as well as non-interacted dummies for extrapolators and mean reverters using the relevant definition. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.9: Manipulation check: Additional measures included in the follow-up survey

	Above average return after negative return (z)		Negative return likely to continue next year (z)	
	(1)	(2)	(3)	(4)
Treatment	-0.326*** (0.067)		-0.452*** (0.065)	
Treatment × Extrapolator (diff. ≥ 4) (a)		-0.382** (0.176)		-0.916*** (0.189)
Treatment × Neutral ($-4 \leq \text{diff.} < 4$)		0.088 (0.123)		-0.274** (0.115)
Treatment × Mean-reverter (diff. < -4) (b)		-0.556*** (0.087)		-0.428*** (0.086)
Extrapolator (diff. ≥ 4)	-0.053 (0.109)	0.172 (0.131)	0.165 (0.115)	0.469*** (0.159)
Mean-reverter (diff. < -4)	0.031 (0.076)	0.340*** (0.098)	-0.076 (0.073)	-0.007 (0.100)
p-value (a=b)		0.379		0.019
Observations	903	903	903	903
R-squared	0.11	0.13	0.12	0.13

Notes: This table shows estimations of the effect of the information treatment on posterior agreement with additional verbal statements describing beliefs about the autocorrelation of aggregate returns included in the four-week follow-up survey after the repeated information treatment. Agreement with the statements is elicited on 7-point categorical scales and is z-scored using the means and standard deviations in the sample. The statements are: “When the stock market has fallen in the previous year one can expect above-average returns for the next year.” (Columns 1-2); “When the stock market has fallen over the previous 12 months there is a high probability that this trend will continue in the following 12 months.” (Columns 3-4). Columns 2 and 4 show heterogeneous treatment effects for prior extrapolators (perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios at least 4 pp), neutrals (difference at least -4 pp and less than 4 pp), and mean reverters (difference less than -4 pp). All estimations include the baseline set of controls described in Appendix A.1.2. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.10: Updating of 12-month-ahead return expectations: Robustness to additional interaction terms

	Updating (point belief)				Updating (mean distr.)			
	(1) OLS	(2) OLS	(3) IV	(4) IV	(5) OLS	(6) OLS	(7) IV	(8) IV
Treatment × Perception gap	0.086** (0.038)	0.081** (0.038)	0.138*** (0.051)	0.131*** (0.051)	0.115*** (0.044)	0.114*** (0.044)	0.142** (0.060)	0.144** (0.059)
Treatment × Female		0.325 (0.532)		0.330 (0.524)		0.480 (0.667)		0.472 (0.654)
Treatment × Age		0.099 (0.494)		0.105 (0.488)		1.084* (0.629)		1.051* (0.619)
Treatment × Employed		-0.477 (0.502)		-0.485 (0.495)		-0.291 (0.628)		-0.289 (0.620)
Treatment × University		1.117*** (0.414)		1.094*** (0.407)		0.903* (0.524)		0.893* (0.518)
Treatment × Log(Household income)		-0.023 (0.081)		-0.023 (0.080)		-0.174* (0.098)		-0.174* (0.097)
Treatment × Log(Fin. wealth with bank)		0.094 (0.130)		0.085 (0.129)		-0.143 (0.156)		-0.157 (0.154)
Treatment × Invest. experience ≥ Median		-0.029 (0.490)		-0.037 (0.485)		0.513 (0.629)		0.533 (0.622)
Treatment × Full financial literacy score		-0.725 (0.484)		-0.712 (0.478)		-0.400 (0.640)		-0.381 (0.633)
Treatment × Follow DAX ≥ Median		-0.184 (0.436)		-0.163 (0.430)		-0.377 (0.557)		-0.374 (0.549)
Perception gap	-0.004 (0.025)	-0.003 (0.026)	-0.019 (0.033)	-0.019 (0.033)	0.022 (0.028)	0.020 (0.028)	0.044 (0.038)	0.039 (0.038)
Treatment	1.077*** (0.212)	0.232 (1.496)	1.007*** (0.219)	0.261 (1.478)	0.047 (0.263)	1.840 (1.674)	0.019 (0.266)	1.947 (1.650)
First stage F-stat			1020.48	1032.79			1020.48	1032.79
Observations	1,961	1,961	1,961	1,961	1,961	1,961	1,961	1,961
R-squared	0.05	0.05	0.04	0.05	0.04	0.05	0.04	0.05

Notes: This table examines changes in expectations about aggregate stock returns over the 12 months after the survey in response to the information among respondents to our main survey based on estimations of specification 2, controlling for additional interaction terms of the treatment indicator with covariates. The outcomes are the difference between posterior and prior point expectations about the 12-month-ahead return (Columns 1-4) and the difference between the mean of the respondent-level posterior distribution over 12-month-ahead returns and the prior point expectation (Columns 5-8). The perception gap is based on the respondent's prior belief about the historical autocorrelation of aggregate returns. It is the difference between the actual conditional mean 12-month-ahead return and the respondent's corresponding prior for the relevant scenario of realized returns over the previous 12 months, which is selected based on respondent's perceived return over the 12 months before the main survey. In Columns 3, 4, 7 and 8 the perception gap is instrumented with a version in which the relevant return interval is selected based on the actual realized return of the DAX over the 12 months before the survey. All estimations include the baseline set of controls described in Appendix A.1.2, which also includes all variables that are interacted with the treatment dummy. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.11: Updating of 12-month-ahead return expectations: Robustness to functional form

	Updating (point belief)			Updating (mean distr.)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment × Perception gap	0.086** (0.038)	0.102** (0.042)	0.098** (0.041)	0.115*** (0.044)	0.126*** (0.048)	0.125*** (0.047)
Perception gap	-0.004 (0.025)	0.005 (0.028)	0.006 (0.028)	0.022 (0.028)	0.031 (0.030)	0.029 (0.030)
Treatment	1.077*** (0.212)	1.051*** (0.213)	1.057*** (0.212)	0.047 (0.263)	0.021 (0.264)	0.023 (0.264)
Perception gap calculated based on	Cond. mean	Linear fit	Quadr. fit	Cond. mean	Linear fit	Quadr. fit
Observations	1,961	1,961	1,961	1,961	1,961	1,961
R-squared	0.05	0.05	0.05	0.04	0.05	0.05

Notes: This table examines changes in expectations about aggregate stock returns over the 12 months after the survey in response to the information among respondents to our main survey based on estimations of specification 2, using different functional forms to calculate the perception gap. The outcomes are the difference between posterior and prior point expectations about the 12-month-ahead return (Columns 1-3) and the difference between the mean of the respondent-level posterior distribution over 12-month-ahead returns and the prior point expectation (Columns 4-6). The perception gap is based on the respondent's prior belief about the historical autocorrelation of aggregate returns and the signal implied by the treatment. In Columns 1 and 4, it is calculated as the difference between the actual conditional mean 12-month-ahead return and the respondent's corresponding prior for the relevant scenario of realized returns over the previous 12 months, which is selected based on respondent's perceived return over the 12 months before the main survey (replicating Columns 1 and 3 of Table 3). In Columns 2 and 5 the expectation predicted from the respondent's prior perceived autocorrelation and the relevant signal implied by the information treatment (which together are used to calculate the perception gap) are instead obtained by fitting linear functions between the different historical bins (i.e., lines connecting the points characterized by the historical conditional mean *past 12-months* return and the corresponding prior predicted or actual conditional mean *12-months-ahead* return between any two neighboring bins). Columns 3 and 6 instead calculate the implied prior and the implied signal by fitting quadratic functions (i) for the three return intervals on the left and (ii) for the three intervals on the right as well as a linear function connecting the two middle bins. Both alternative ways of calculating the perception gap (Columns 2 and 5 and Columns 3 and 6) simply assign the difference between the actual conditional mean return shown in the information and the respondent's corresponding prior in the two extreme bins in the few cases where the perceived return over the 12 months before the survey is lower than -29.7% or higher than 33.8% (the conditional mean past 12-month returns for the extreme bins). All estimations include the baseline set of controls described in Appendix A.1.2. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.12: Disagreement in priors and posteriors

	Expected return next 12 months		
	(1) Control Group	(2) Treatment Group	(3) p-value (1) = (2)
SD priors	6.10	6.45	0.241
SD posteriors	5.86	5.29	0.008
p90 – p10 priors	15.00	15.00	
p90 – p10 posteriors	15.00	10.60	
Observations	989	972	

Notes: This table displays the cross-sectional standard deviation and the difference between the 90th and the 10th percentile of expectations about aggregate stock returns over the 12 months after the survey, both based on priors and based on posteriors, separately for the control group (Column 1) and the treatment group (Column 2). Column 3 displays p-values of Levene's tests for equality of the standard deviation across treatment arms.

Table A.13: Sources of beliefs: Experiences and memory database

	Thought of specific episode		
	(1)	(2)	(3)
Started investing in stocks before the episode	0.067*** (0.018)		0.051** (0.020)
Turned 18 years before the episode		0.063*** (0.015)	0.038** (0.017)
Observations	1,872	1,872	1,872
R-squared	0.02	0.02	0.03

Notes: This table examines whether investors' memory databases are shaped by their own lifetime experiences among German retail investors participating in an additional descriptive survey. The dataset is at the investor \times episode level. There are 208 investors and the following nine historical episodes: the Covid-19 downturn (2020), the downturn associated with Brexit (2018), the euro crisis (2011), the Global Financial Crisis (2007-9), the downturn associated with the Iraq War (2003), the terror attacks on 9/11 (2001), the burst of the dot-com bubble (2000-3), the downturn associated with the Gulf War (1990), and the downturn following the Black Monday (1987). The outcome is a dummy variable taking value one if an investor thought of the specific event when estimating the historical autocorrelation. The main independent variables are dummies taking value one if the investor had started to invest in stocks before the respective episode (Column 1) or if the investor had turned 18 before the respective episode (Column 2). All estimations include a parsimonious set of controls (gender, employment status, education, household income, financial wealth, financial literacy, and attention to DAX developments). Standard errors clustered at the investor level are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

A.4 Additional evidence: Trading

In this appendix we present additional evidence on trading behavior among respondents to our main survey.

Table A.14 shows regressions of different measures of investors' net buying of stocks on the change in the realized 12-month DAX return from the end of the previous to the end of the current period and the lagged realized 12-month return among our main investor sample and a random sample of investors from the same subject pool.

Table A.15 replicates the analysis of treatment effects on trading responses to the Covid-19 crash reported in Table 5 and described in Section 4.2 on an "active sample" of observations for which the corresponding investor conducted at least one equity trading transaction over the preceding three months.

Table A.14: Realized returns and average trading decisions

	Full sample		Random sample	
	(1) Probability net buy (%)	(2) Relative net buying volume (%)	(3) Probability net buy (%)	(4) Relative net buying volume (%)
Δ Realized return	-0.127*	-0.031***	-0.125***	-0.027***
P-Value	(0.064)	(0.010)	[0.746]	[0.185]
Lagged realized return	-0.052*	-0.012***	-0.054**	-0.010**
P-Value	(0.029)	(0.004)	[0.673]	[0.850]
R-squared	.032	.013	.011	.011
Observations	74,569	74,569	180,039	180,039

Notes: This table examines the association between the development of the German stock market index DAX and trading activity at the online bank among respondents to our main survey (Columns 1-2) and among a random sample of investors at the online bank (Columns 3-4) based on investor-month level estimations. The "relative net buying volume" is defined as the percent ratio of net purchases of equity over a given month relative to end-of-previous-month financial wealth. "Probability net buy" is a dummy variable, multiplied by 100, indicating whether a respondent has a relative net buying volume of at least 1 pp. The change in the realized return is the difference between the realized return in the current and in the previous month. For the return calculation, we use the average return compared to 12 months earlier across the last seven days in the respective month. The transaction data span the period from December 2014 until August 2019 (directly before the survey period) for the treatment group and until including January 2020 for the control group and the random sample. The table also provides p-values testing for differences in trading reactions to changes in and lagged levels of realized returns across the two samples (comparing Columns 1 and 3 as well as Columns 2 and 4). All estimations control for lagged log financial wealth held with the bank, the lagged equity share, dummies indicating a lagged equity share of 0% or of 100% as well as a set of macroeconomic variables (year-on-year monthly inflation and quarterly, seasonally-adjusted GDP growth). Standard errors are two-way clustered by investor and trading month and are presented in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.15: Treatment effects on trading responses to the Covid-19 crash: Active sample

	Probability net buy (%)	Relative net buying volume (%)	Relative gross buying volume (%)	Relative gross selling volume (%)	Probability trade (%)	Probability net buy (%)	Relative net buying volume (%)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Crash \times Treatment [β_1]	4.650* (2.320)	0.423** (0.168)	0.436*** (0.158)	0.103 (0.180)	4.296 (2.556)	7.993* (4.024)	1.448** (0.679)
Crash \times Treatment \times $\mathbb{1}$ (Predicted exp. adjustment > 0) [β_2]	-8.502*** (2.894)	-0.639*** (0.210)	-0.615*** (0.181)	-0.055 (0.198)	-7.494** (2.996)	-11.841** (4.594)	-1.921*** (0.635)
Crash \times $\mathbb{1}$ (Predicted exp. adjustment > 0) [β_3]	2.915 (2.448)	0.200 (0.179)	0.183 (0.165)	0.042 (0.155)	2.798 (2.174)	6.492 (4.176)	1.118* (0.635)
p-value($\beta_1 + \beta_2 = 0$)	0.001	0.034	0.076	0.387	0.017	0.218	0.303
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Specification						Cond. on trading	Cond. on trading
Observations	50,125	50,125	50,125	50,125	50,125	7,672	7,672
Number of investors	1,710	1,710	1,710	1,710	1,710	1,077	1,077
R-squared	0.16	0.07	0.14	0.08	0.16	0.27	0.21

Notes: This table presents treatment effects of our intervention on trading responses to the Covid-19 crash among respondents to our main survey based on investor-week level estimations of specification 5. The “relative net buying volume” is defined as the percent ratio of net purchases of equity over a given month relative to end-of-previous-month financial wealth. The “relative gross buying volume” and the “relative gross selling volume” capture gross purchases and sales relative to end-of-previous-month financial wealth. “Probability net buy” and “probability trade” are dummy variables, multiplied by 100, indicating whether a respondent has a relative net buying volume of at least 1 pp or of at least 1 pp in absolute terms, respectively. “ $\mathbb{1}$ (Predicted exp. adjustment > 0)” indicates whether an investor would upward adjust his or her return expectations in response to the crash if return expectations were exclusively based on the investor’s prior perceived historical autocorrelation of returns, assuming accurate beliefs about the change in the realized return over the crash period. “Crash” takes value one for the weeks from 17th February until 13th March 2020 and zero otherwise. The sample spans the time from 3rd June 2019 until 13th March 2020. The specifications also include interactions of a dummy indicating the “Post-survey pre-crash” from 16th September 2019 until 16th February 2020 with a treatment dummy, with the “ $\mathbb{1}$ (Predicted exp. adjustment > 0)” dummy, and with both of these dummies. The estimations focus on observations for which the corresponding investor conducted at least one equity trading transaction over the preceding three months. All specifications control for individual fixed effects, week-year fixed effects and lagged log financial wealth with the bank. Standard errors are two-way clustered by investor and trading week and are presented in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

A.5 Persistence in four-week follow-up

How persistent are the changes in beliefs in response to the treatment information documented in Section 3.2? We address this question using data from the follow-up survey, in which respondents participated about four weeks after the main survey. We focus on responses in the follow-up that were given *before* the information was provided for a second time to respondents in the treatment group.

First, we examine respondents' beliefs about historical 12-month-ahead returns for the six intervals of realized returns over the previous 12 months. For each interval we regress the difference between a respondent's follow-up and prior beliefs on the gap between the information and the respondent's prior, a treatment dummy, and the interaction of the two. Table A.16 Columns 1-6 highlight estimated coefficients on the interaction term of about 0.25, indicating that treated respondents adjust their beliefs by about one fourth of the initial gap to the information. We find a similar effect size for the extent that treated respondents adjust their beliefs about the difference in year-ahead returns between previous positive and previous negative returns towards the information (Column 7).

Second, Table A.17 examines agreement with the three verbal manipulation check questions in the follow-up. Since these questions were included in both the main and the follow-up survey, we can quantify the persistence of initial treatment effects. For both the statement capturing beliefs about the absence of any form of predictability by recent returns and the statement capturing a belief in mean reversion, we find that treatment effects strongly persist, and, if anything, increase in size compared to the main survey. While the treatment effects on agreement with the statement capturing a belief in persistence do not persist, this finding should be interpreted in light of the very small group of prior extrapolators in the follow-up sample.

Third, Table A.18 examines the persistence of updating of expectations about the return over the 12 months after the survey based on specification 2. Columns 1 and 2 display estimates in the main survey restricted to those who later participate in the follow-up. The outcomes in Columns 3 and 4 are the difference between 12-month-ahead return expectations measured in the beginning of the follow-up survey and the prior elicited in the main survey. In both OLS and IV specifications the estimated effect sizes *increase* compared to the effect sizes in the main survey.

Taken together, these patterns highlight a strong persistence of treatment effects on respondents' beliefs. Previous studies using information provision experiments often find that treatment effects on respondents' beliefs persist at a reduced size in follow-up surveys (Armona et al., 2019; Cavallo et al., 2017; Coibion et al., 2022; Roth and Wohlfart, 2020). The higher persistence in our setting could be due to the fact that our information treatment aims to change respondents' beliefs about return predictability – and therefore the way they form return expectations – instead of providing them with information that might lose its relevance over time, such as, e.g., expert forecasts.

Our evidence on persistence mitigates two concerns. First, changes in return expectations could be driven by unconscious numerical anchoring on the information. Such anchoring is a short-lived phenomenon by definition, so the strong persistence of treatment effects in the follow-up suggests a limited role for numerical anchoring (Cavallo et al., 2017; Coibion et al., 2022). Second, experimenter demand effects – respondents guessing the experimental hypothesis and trying to conform with it – should be less important in the follow-up, where respondents are less likely to remember exact details of the experimental intervention (Haaland et al., 2023).²

²de Quidt, Haushofer and Roth (2018) show that demand effects seem to be of limited quantitative importance in online experiments.

Table A.16: Recall of treatment information in four-week follow-up

	Δ Estimated historical mean return next 12 months conditional on return previous 12 months in interval						Δ Perceived diff. gain-loss
	(1) $\leq -20\%$	(2) [-20%, -10%]	(3) [-10%, 0%]	(4) [0%, 10%]	(5) [10%, 20%]	(6) >20%	(7)
Treatment \times (Information - Prior)	0.277*** (0.054)	0.132** (0.063)	0.020 (0.076)	0.319*** (0.109)	0.260*** (0.060)	0.252*** (0.054)	0.317*** (0.057)
Treatment	-0.594 (0.571)	-0.986* (0.515)	-0.723* (0.401)	-0.766* (0.457)	0.384 (0.427)	0.888* (0.527)	-0.826 (0.572)
Information - Prior	0.495*** (0.041)	0.610*** (0.043)	0.759*** (0.052)	0.521*** (0.091)	0.567*** (0.046)	0.531*** (0.040)	0.420*** (0.043)
Observations	903	899	900	903	903	902	903
R-squared	0.45	0.43	0.41	0.30	0.45	0.48	0.40

Notes: This table examines whether respondents recall the treatment information in the four-week follow-up. The outcomes are differences between respondents' posterior beliefs measured in the four-week follow-up survey and prior beliefs measured in the main survey. The beliefs are the perceived historical average 12-month-ahead return when the return over the previous 12 months fell into one of six intervals (Columns 1-6), and the perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios (Column 7). Changes in beliefs are regressed on a treatment indicator, the difference between information and prior, and the difference between information and prior interacted with a treatment indicator (indicating whether respondents actually received the information). Changes in beliefs and differences between information and prior are winsorized at -20 and 20 pp. All estimations include the baseline set of controls described in Appendix A.1.2 as well as dummies for the time between main and follow-up survey. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.17: Manipulation check: Persistence in four-week follow-up

	Main survey		Main survey		Follow-up survey	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Positive return irrespective of previous return (z)						
Treatment	0.092**		0.107		0.113*	
	(0.044)		(0.065)		(0.068)	
Treatment ×		-0.037		-0.176		-0.021
Extrapolator (diff. ≥ 4)		(0.112)		(0.178)		(0.208)
Treatment ×		0.078		0.085		-0.031
Neutral (-4 ≤ diff. < 4)		(0.079)		(0.117)		(0.116)
Treatment ×		0.140**		0.198**		0.237***
Mean-reverter (diff. < -4)		(0.059)		(0.085)		(0.089)
Observations	1,961	1,961	903	903	903	903
Panel B: No sense to buy after high return (z)						
Treatment	-0.054		-0.050		-0.135**	
	(0.044)		(0.063)		(0.068)	
Treatment ×		0.021		-0.133		-0.259
Extrapolator (diff. ≥ 4)		(0.114)		(0.174)		(0.177)
Treatment ×		0.075		0.025		0.030
Neutral (-4 ≤ diff. < 4)		(0.080)		(0.112)		(0.118)
Treatment ×		-0.155***		-0.072		-0.199**
Mean-reverter (diff. < -4)		(0.060)		(0.084)		(0.094)
Observations	1,961	1,961	903	903	903	903
Panel C: Positive return more likely after high return (z)						
				0.031		
					(0.067)	
Treatment	-0.147***		-0.100			
	(0.045)		(0.069)			
Treatment ×		-0.375***		-0.610***		-0.026
Extrapolator (diff. ≥ 4)		(0.115)		(0.184)		(0.179)
Treatment ×		-0.084		-0.011		0.058
Neutral (-4 ≤ diff. < 4)		(0.081)		(0.119)		(0.120)
Treatment ×		-0.114*		-0.010		0.031
Mean-reverter (diff. < -4)		(0.062)		(0.093)		(0.091)
Observations	1,961	1,961	903	903	903	903

Notes: This table shows persistence of treatment effects on respondents' posterior agreement with verbal statements describing beliefs about the autocorrelation of aggregate returns in the four-week follow-up survey. Agreement with the statements is elicited on 7-point categorical scales, and is z-scored using the means and standard deviations in the sample. The statements are: "With an investment in stocks one can expect a positive return, independently of how the stock market has developed in the recent past." (Panel A); "When the stock market has recently increased it makes no sense to buy stocks." (Panel B); "When the stock market has recently increased it is more likely that stock returns will be positive over the following time than when the stock market has recently decreased." (Panel C). Columns 1-2 focus on responses in the main survey using the full sample. Columns 3-4 focus on responses in the main survey using those who later completed the follow-up. Columns 5-6 focus on responses in the four-week follow-up. Columns 2, 4 and 6 show heterogeneous treatment effects for prior extrapolators (perceived difference in average 12-month-ahead returns between the positive and the negative realized return scenarios at least 4 pp), neutrals (difference at least -4 pp and less than 4 pp), and mean reverters (difference less than -4 pp). All estimations include the baseline set of controls described in Appendix A.1.2. Columns 3-6 additionally control for dummies indicating the time between main and follow-up survey. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.18: Updating of 12-month-ahead return expectations measured in four-week follow-up

	Updating (point belief) main survey follow-up sample		Updating (point belief) follow-up survey	
	(1)	(2)	(3)	(4)
	OLS	IV	OLS	IV
Treatment × Perception gap main	0.118** (0.056)	0.134* (0.076)		
Perception gap main	-0.054 (0.035)	-0.020 (0.048)		
Treatment × Perception gap follow-up			0.136* (0.080)	0.276** (0.122)
Perception gap follow-up			-0.069 (0.056)	-0.129 (0.080)
Treatment	1.130*** (0.312)	1.109*** (0.316)	0.482 (0.426)	0.236 (0.436)
First stage F-stat		534.73		355.35
Observations	903	903	903	903
R-squared	0.07	0.07	0.05	0.04

Notes: This table examines changes in expectations about aggregate stock returns over the 12 months after the survey in response to the information based on estimations of specification 2. The outcomes are the difference between posterior and prior point expectations about the 12-month-ahead return, both measured in the main survey (Columns 1 and 2) and the difference between the posterior point expectation measured at the start of the follow-up survey (*before* the repeated information treatment) and the prior point expectation measured in the main survey (Columns 3-4). The perception gap is based on the respondent's prior belief about the historical autocorrelation of aggregate returns. It is the difference between the actual conditional mean 12-month-ahead return and the respondent's corresponding prior for the relevant scenario of realized returns over the previous 12 months, which is selected based on the respondent's perceived return over the 12 months before the main survey (Columns 1-2) or before the follow-up survey (Columns 3-4). In Columns 2 and 4 the perception gap is instrumented with a version in which the relevant return interval is selected based on the actual realized return of the DAX over the 12 months before the respective survey. All estimations are based on respondents who are part of the follow-up sample. All estimations include the baseline set of controls described in Appendix A.1.2 as well as dummies for the time between main and follow-up survey. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

A.6 Additional details on interpretation of magnitudes

In this appendix we provide more details on the interpretation of the magnitudes of our estimated effects of the pass-through of beliefs to investment choices. Specifically, we compare our estimated effects (i) to theory benchmarks and (ii) to estimates by Giglio et al. (2021a) from repeated surveys with Vanguard clients and matched administrative portfolio data from respondents' retirement accounts. Giglio et al. (2021a) estimate a relationship between an investor's equity share and her return expectations of between 0.7 and 1.2 depending on the specification. By contrast, frictionless benchmark models such as Merton (1969) would predict a relationship between 3 and 10 for plausible preference parameters. Our setting allows us to speak to the elasticity of investment choices to beliefs based on both correlational and experimental estimates.

Equity share To obtain estimates that are comparable to those in Giglio et al. (2021a), we start by running cross-sectional regressions of the equity share in a respondent's holdings at the bank on the return that the respondent would expect in that month if return expectations were fully based on the perceived autocorrelation. We include our baseline set of controls, log financial wealth and time fixed effects. As shown in Column 1 of Table A.19, a one pp higher predicted return expectation is associated with a 0.151 pp higher equity share ($p < 0.05$). Column 2 additionally controls for a respondent's average stated return expectation in the five bins into which the previous 12-month return did *not* fall. The estimated coefficient on the predicted return expectation increases to 0.160 pp and is highly statistical significant (Column 2, $p < 0.01$), while the mean perceived historical return across the irrelevant bins has no significant effect. Thus, whether a respondent believes that returns were higher following periods with a similar recent return as in the current period than at other times seems to matter for their portfolio allocation.

Our estimates of between 0.151 and 0.160 are smaller in magnitude than the estimates in Giglio et al. (2021a). One reason could be that we predict an investor's expectation in a particular month depending on the realized return over the previous 12 months and the respondent's belief about the historical autocorrelation. Investors' actual return expectations may deviate because (i) other variables than the recent return could influence investors' return expectations at any point in time, (ii) investors may think that the current autocorrelation differs from its historical counterpart, (iii) actually realized

DAX returns are an imperfect proxy for investors' perceived returns over the previous 12 months, or (iv) investors' beliefs about the historical autocorrelation themselves may have changed between the particular trading month and the survey period. The less than one-to-one relationship between actual return expectations and predicted return expectations based on an investor's perceived autocorrelation shown in Figure 3 Panel A is consistent with possibilities (i) and (ii).

To make the effect sizes comparable to theory benchmarks and the results in Giglio et al. (2021a), we first regress investors' 12-month-ahead return expectations at the time of our survey on a measure of investors' predicted 12-month-ahead return expectations based on the respondents' perceived historical autocorrelation of returns and the actual return realized over the 12 months before the respondent took the survey. As shown in Panel B of Figure 3, conditional on our baseline set of controls, we obtain an estimated relationship of 0.169, which is highly statistically significant. We then adjust our estimates of the elasticity of the equity share to beliefs ranging from 0.151 and 0.160 for the fact that we predict the respondent's expectations by dividing them by 0.169. The resulting adjusted elasticities range between 0.893 and 0.947 – remarkably close to the estimates in Giglio et al. (2021a). Thus, while differences in investors' return expectations due to different beliefs about the autocorrelation are reflected in portfolio decisions, the magnitude of the relationship is smaller than what is implied by frictionless benchmark models.

Trading In Table 4 we regress an investor's relative net buying volume during a given month on the change in the respondent's predicted return expectations over that month (calculated using the perceived autocorrelation and the change in the realized return). If there are no inflows to or outflows from a respondent's holdings at the bank, the relative net buying volume is equal to the active change in the equity share – i.e., the change in the equity share that is purely due to trading, abstracting from price changes. In the full sample, a one pp higher increase in the predicted return expectation is associated with a 0.015 pp higher relative net buying volume. This is an order of magnitude smaller than the cross-sectional estimates of the relationship between beliefs and the equity share displayed in Table A.19. This difference in magnitudes is in line with the evidence in Giglio et al. (2021a) and could reflect the impact of frictions. On top this, the impact of measurement error introduced by predicting a respondent's return expectation instead of eliciting it directly may be more severe when taking differences. As explained

in Section 4.2, our experimental estimates of the pass-through from return expectations to trading are about five times as large as the correlational ones. This is still substantially smaller than benchmarks from frictionless models.

Table A.19: Perceived autocorrelation and the equity share

	Equity share (%)	
	(1)	(2)
Perceived conditional historical return (%)	0.151** (0.065)	0.160*** (0.056)
Perceived mean historical return other bins (%)		-0.046 (0.135)
Time FE	Yes	Yes
Observations	78,691	78,691
Number of investors	1,909	1,909
R-squared	0.06	0.06

Notes: This table examines the association between the perceived autocorrelation of returns and the equity share held with the bank among respondents to our main survey based on investor-month level estimations. The “perceived conditional historical return” is the return an investor would expect if his or her return expectations were exclusively based on the investor’s prior perceived historical autocorrelation of returns, assuming accurate beliefs about the current realized return over the previous 12 months. The “perceived mean historical return other bins” is the average of a respondent’s historical 12-month-ahead return perceptions over the five bins into which the current realized return does *not* fall. The transaction data span the period from December 2014 until August 2019 (directly before the survey period) for the treatment group and until including January 2020 for the control group. The results are based on the full sample. All specifications control for month-year fixed effects and lagged log financial wealth with the bank as well as the baseline set of controls measured at the time of the survey described in Appendix A.1.2, excluding the variables relating to portfolio shares and trading activity. Standard errors are two-way clustered by investor and trading month and are presented in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

A.7 Experiment on beliefs about the price-dividend ratio

In our main experiment we demonstrate that retail investors hold highly heterogeneous beliefs about the autocorrelation of aggregate stock returns, which causally drives disagreement about expected returns and trading. While recent returns seem to be central to the formation of investors' expectations about future returns (Dominitz and Manski, 2011; Greenwood and Shleifer, 2014; Heiss et al., 2022; Vissing-Jorgensen, 2003), investors may also consider other state variables when forming return expectations. For instance, the price-dividend ratio (PD ratio) should be negatively related to future returns, even though movements in the PD ratio also reflect changes in expected future cash flow growth (Campbell and Shiller, 1988) and the relative importance of expected returns and cash flows in driving PD ratios differs across markets and time periods (Rangvid et al., 2014; Golez and Koudijs, 2018). In this appendix we use an additional experiment on a separate sample of retail investors to measure and shift respondents' beliefs about return predictability based on the PD ratio, and to study how these beliefs are linked to individuals' return expectations.

A.7.1 PD experiment: Design

We use a very similar design as in our main survey (described in Section 2.1), including three stages (i) measuring priors, (ii) providing a random subset of respondents with information, and (iii) eliciting posteriors. The main difference is that, instead of measuring respondents' perceived return over the last 12 months and their perceived historical autocorrelation, the experiment elicits respondents' perceived current PD ratio and their beliefs about historical return predictability based on the PD ratio. Moreover, due to advantages in terms of data availability for the empirical benchmark and the information treatment, we frame all belief elicitations around the entire German stock market instead of the DAX, which covers only a subset of listed firms. The instructions for the additional experiment can be found under <https://drive.google.com/file/d/1IhxTWR5pXyVSR0V6tPfr4y06p7UwygmQ/view?usp=sharing>.

Investors are less familiar with the PD ratio than with returns (Figure A.4). Before eliciting respondents' priors, we therefore provide them with a brief explanation of the PD ratio and anchor them on the range into which the ratio fell in the German stock market over the last 50 years (20 to 75, with an average of 40). This anchor should make it easier for our respondents to meaningfully report their

beliefs (Ansolabehere, Meredith and Snowberg, 2013). We use the same type of dynamic figure as in our main survey for the elicitation of beliefs about predictability and the information treatment. We elicit prior beliefs about conditional mean 12-month-ahead returns in the last 50 years for four intervals of the current PD ratio (“lower than 30”, “between 30 and 40”, “between 40 and 50”, “higher than 50”) and subsequently display to respondents in the treatment group the actual average realizations. These values indicate a lack of return predictability by the PD ratio across the three lower PD bins (13.4%, 13.3%, and 12.3%, respectively). Merely in the case of very high PD ratios, subsequent returns are systematically lower (6.7%).³ We also show treated respondents the following statement: *When the price-dividend ratio was very high, the return of the German stock market over the following 12 months was relatively low on average.* As in the main experiment, respondents in the control group are provided with the unconditional average return of the overall German stock market (11.1%). After the treatment, we elicit respondents’ agreement with two verbal statements describing a positive and a negative relationship between the current PD ratio and future returns, respectively, which we use as manipulation check.

A.7.2 PD experiment: Survey administration and sample

Survey administration We administered the survey in July and August 2022 to clients of the same online bank as used in our main experiment. Due to changes in the bank’s data protection policies that came into effect after we had conducted the main survey, we were not allowed to link the new survey data to administrative account data on the clients’ investment decisions. We therefore focus on investors’ belief formation in our analysis.

We sent email invitations to 9,000 individuals from the bank’s client pool, which were selected in the same way as done in our main survey. We offered invitees a 5 EUR reward in the form of an Amazon voucher for completing the survey. 772 individuals completed our survey, corresponding to a response rate of 8.6%, lower than in the main survey. Potential reasons for the lower response rate include the lower reward (5 EUR instead 10 EUR), changes in the bank’s client pool, and the timing of the survey during holiday season. The mean (median) response time to the survey was 18 (14.7) minutes.

³Thus, predictability of returns by the PD ratio is less pronounced in the German setting than in other markets and time periods, in line with movements in the PD ratio mostly reflecting changes in expected future cash flow growth in the German setting (1.6%, 8.1%, 12.2% and 15.2% 12-month-ahead dividend growth going from the lowest to the highest PD bucket). See also Appendix A.2 and Rangvid et al. (2014).

Sample characteristics We select our sample using the same criteria as in our main survey (described in Section 2.4), dropping individuals with extreme response times or prior or posterior return expectations.⁴ Again, our results do not hinge on the exact choices of cutoffs used to define the sample. Table A.20 Columns 2-6 display summary statistics of the resulting sample of 693 respondents. The composition of our sample is similar to the sample used in the main experiment, the main differences being that respondents to the additional survey are somewhat older and wealthier on average. Columns 7-9 highlight that the sample is mostly balanced across the treatment and the control group. To address any concern, we include a set of control variables in all estimations.

A.7.3 PD experiment: Prior beliefs

Prior perceived return predictability Figure A.10 Panel A highlights that respondents on average perceive a weak negative historical relationship between the current PD ratio and subsequent returns. For instance, while they on average perceive mean returns of 12.2% over the next year when the PD ratio was in the lowest bin (below 30), they believe subsequent returns were 9.7% when the PD ratio was in the highest bin (above 50). Thus, they perceive a flatter relationship than what is implied by actual historical data (13.4% and 6.7% for the lowest and the highest bin, respectively). The only weakly negative perceived relationship implies that the beliefs in mean reversion of returns documented in our main analysis are unlikely to be driven by beliefs about return predictability based on the PD ratio. Panel B highlights that there is a substantial amount of disagreement in each given bin, which is the lowest for moderate PD ratios between 40 and 50 and more pronounced for very low or very high ratios, similarly as for the perceived relationship between past and future returns measured in the main survey (see Section 3.1).

We next calculate for each respondent the difference between the perceived 12-month-ahead return for the highest PD scenario (above 50) and the average perceived 12-month-ahead return over the three scenarios with lower PD ratios. We focus on this difference because actual historical return realizations are lowest in the highest PD scenario (6.7% on average), while they are very similar and at a higher level across the three lower PD scenarios (13% on average). Panel C shows that there is a high level

⁴The only difference compared to the main survey is that we drop respondents with response time below 6 minutes instead of 8 minutes to account for the shorter survey length.

of heterogeneity in respondents' beliefs about this difference. A majority estimate the difference to be positive (40.8% of respondents) or less negative than the actual difference of -6.3 pp (37.1% of respondents) – i.e., most respondents do not fully account for how much lower historical returns were in cases of very high PD ratios compared to cases of low and moderate PD ratios. Similarly as in our main analysis, we classify respondents into “underestimators”, for whom the difference in perceived 12-month-ahead returns between the highest and lower PD bins is more negative than -10 pp (11.4% of our sample), “neutrals” (difference between -10 pp and -2 pp, 33.2% of our sample) and “overestimators” (difference more positive than -2 pp, 55.4% of our sample).

How do beliefs about return predictability by the PD ratio vary across groups? In Table A.21 we regress different belief measures on a set of co-variates. For instance, respondents with higher financial literacy are more likely to be in the “neutral” category and less likely to overestimate the difference in 12-month-ahead returns between high and low PD scenarios.

Perceived return predictability and return expectations We next examine whether beliefs about return predictability based on the price-dividend ratio are related to respondents' return expectations at the time of the survey. We first select the PD ratio interval covering the respondent's perceived PD ratio at the time of the survey. Respondents on average believe that the current PD ratio is 40.4, compared to an actual average PD ratio over the survey period of 27.6, and there is strong heterogeneity in respondents' perceived current PD ratio. We then study how respondents' expectations about the return over the 12 months after the survey are related to their perceived conditional average historical 12-month-ahead return from the relevant PD ratio interval. Figure A.11 Panel A shows a binned scatter plot of this relationship, partialling out a set of controls. A one pp higher perceived 12-month-ahead return in the relevant historical scenario is associated with a 0.26 pp higher return expectation at the time of the survey ($p < 0.01$). The relationship is similar when selecting the relevant PD ratio interval based on the actual instead of the respondent's perceived PD ratio (Panel B). When controlling for a respondent's mean belief about the 12-month-ahead return across the three intervals into which the (perceived or actual) current PD ratio did *not* fall, the effects decrease in size and significance but remain meaningful. Taken together, respondents' beliefs about historical return predictability based on the PD ratio are reflected in their return expectations at the time of the survey.

A.7.4 PD experiment: Updating of beliefs

Manipulation check We next examine whether the information indicating a negative relationship between actual historical return realizations and the PD ratio shifts respondents' beliefs about return predictability based on the PD ratio. To do this, we regress respondents' post-treatment agreement with different statements describing the relationship between the current PD ratio and future returns – measured on 7-point categorical scales, which we z-score using the mean and the standard deviation in the sample – on a treatment indicator and a set of controls. We also report specifications in which we interact the treatment indicator with dummies indicating prior underestimators or overestimators of the relationship or with a dummy for neutrals.

As shown in Table A.22 Column 1, the treatment increases respondents' agreement that “*On average, low returns follow high price-dividend ratios*” by 61.9% of a standard deviation ($p < 0.01$). In the control group, those who previously overestimated the relationship are significantly less likely to agree with this statement than neutrals (Column 2, $p < 0.01$). In line with this, the treatment increases agreement with the statement by 74.9% of a standard deviation among prior overestimators and only by 47.1% ($p < 0.01$) and by 44.6% ($p < 0.05$) among prior neutrals and underestimators, respectively (Column 2).

The treatment also significantly reduces respondents' agreement with the statement “*When, at a specific point in time, the price-dividend ratio is high, high returns can be expected over the following time*” by 48.2% of a standard deviation (Column 3, $p < 0.01$). These effects are fully driven by prior overestimators (Column 4, 68.9% of a standard deviation, $p < 0.01$) and neutrals (30% of a standard deviation, $p < 0.05$), while prior underestimators do not respond (0% of a standard deviation, $p = 0.826$). This implies that the treatment approximately offsets the control group differences between prior overestimators on the one hand and prior neutrals and underestimators on the other hand, as can be seen in Column 4.

Taken together, our intervention successfully shifts respondents' beliefs about return predictability based on the PD ratio towards the treatment information. The shifts in beliefs are proportional to respondents' prior mis-perceptions of the historical relationship between the PD ratio and subsequent returns.

Updating of return expectations If beliefs about return predictability based on the PD ratio are a causal driver of respondents' return expectations, respondents should update their return expectations in response to the treatment information. Similarly as done in our main analysis (Section 3.2), we define a respondent's perception gap as the difference between actual and perceived historical average 12-month-ahead returns in the PD interval that captures the respondent's perceived PD ratio at the time of the survey. We then regress a respondent's belief updating – i.e., the difference between a respondent's posterior and prior return expectation – on a treatment dummy, the perception gap, and the interaction between the two, as described in specification 2 in the main text. Similarly as in our main analysis, we also report specifications in which the perception gap and the interaction term are instrumented with versions that are based on the actual PD ratio at the time of the survey instead of a respondent's perception. This serves the purpose of mitigating measurement error in reported beliefs. We report specifications using the posterior point forecast or using the mean of a respondent's subjective probability distribution over future returns to construct the updating variable.

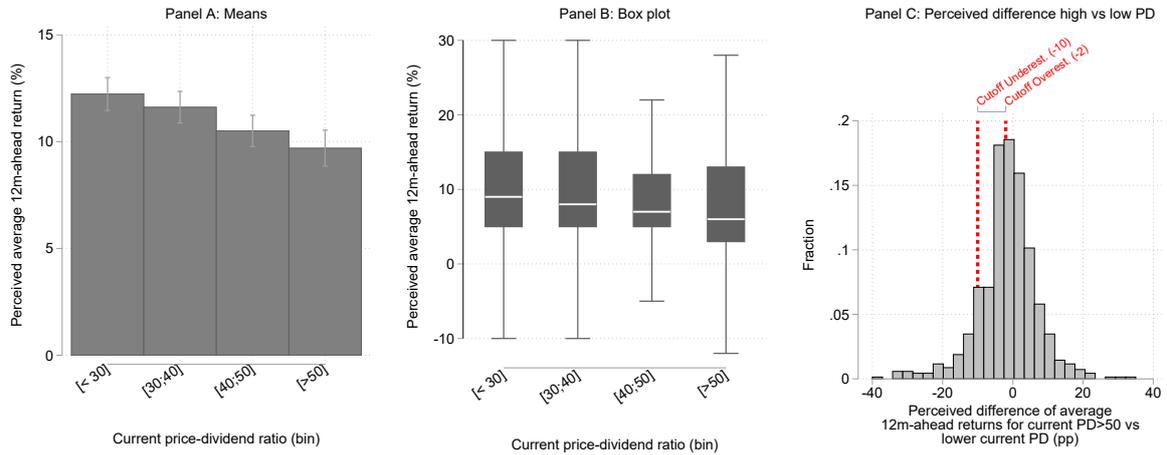
Table A.23 shows learning rates between 10.9% and 21.2% from the information, as indicated by the estimated coefficient on the interaction term, which are statistically significant at the 10% level or at more precise levels. Similarly as in our main experiment, our estimated learning rates are substantially smaller than one. This is in line with the ideas that (i) respondents do not exclusively rely on their beliefs about predictability by the PD ratio when forming return expectations and (ii) that respondents view the historical relationship between the PD ratio and subsequent returns as an imperfect proxy for the forward-looking relationship at the time of the survey. However, the significant updating of return expectations in response to the information highlights that beliefs about return predictability based on the PD ratio causally shape retail investors' return expectations.

A.7.5 PD experiment: Summary

Our additional experiment on beliefs about the PD ratio highlights that many of the take-aways from our main experiment do not exclusively apply to recent returns, but carry over to other state variables that respondents may use to form their return expectations. Specifically, investors' beliefs about the relationship between the PD ratio and subsequent returns are highly heterogeneous, they are responsive to information, and they are causally linked to respondents' return expectations. This corroborates that

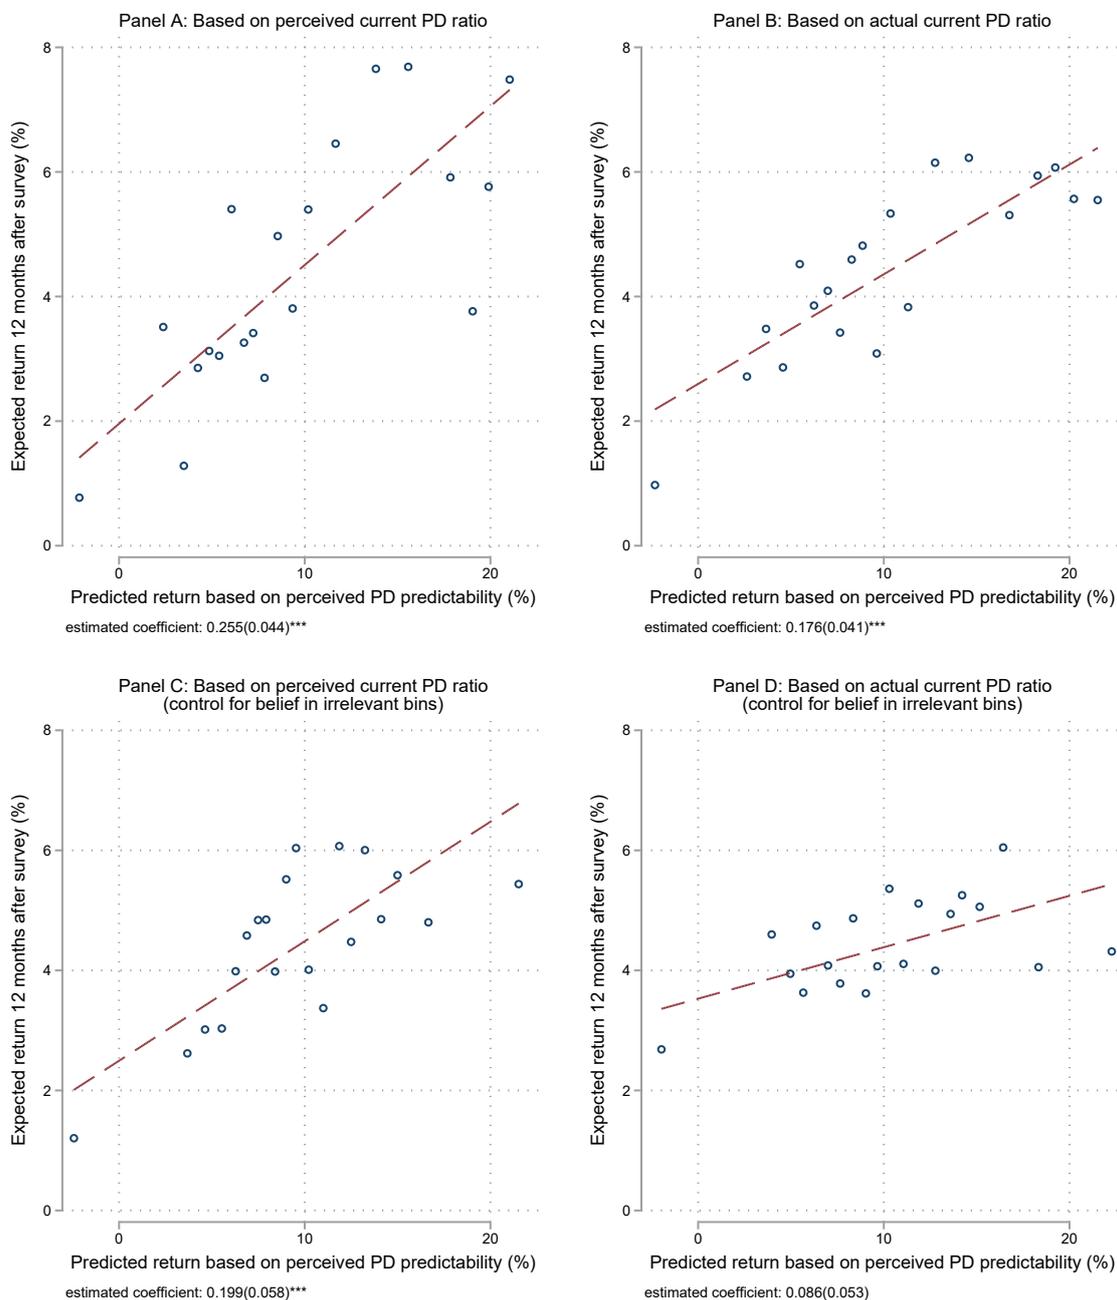
heterogeneous beliefs about the predictiveness of given state variables for subsequent returns are an important driver of disagreement about expected returns in the stock market.

Figure A.10: Prior beliefs about return predictability based on the price-dividend ratio



Notes: This figure summarizes prior beliefs about future return predictability based on the current price-dividend ratio of the German stock market in the last 50 years among retail investors participating in an additional experimental survey. Panel A shows the sample means of respondents' beliefs about average 12-month-ahead stock returns for four intervals of realized current price-dividend ratios. Panel B displays box plots of respondents' prior beliefs about average 12-month-ahead stock returns for the four intervals of realized current PD ratios, including median, 25th and 75th percentile for each interval. Panel C shows a histogram of respondents' perceived difference in average 12-month-ahead returns between cases where the PD ratio was higher than 50 (the highest bin) and cases where the PD ratio was at most 50 (the three lower bins), including the cutoffs we use to define underestimators and overestimators.

Figure A.11: Binned scatter plot of prior expected 12-month-ahead return vs predicted return expectation based on perceived historical predictability by the price-dividend ratio



Notes: This figure shows binned scatter plots of respondents' prior expected return over the 12 months after the survey against the respondents' perceived average historical 12-month-ahead return in the relevant interval of realized PD ratios, which is selected based on the respondent's perceived current PD ratio (Panels A and C) or based on the actual current PD ratio at the time when the respondent took the survey (Panels B and D). The sample consists of retail investors participating in an additional experimental survey. The binned scatter plots partial out a set of controls similar to the ones described in Appendix A.1.2. Panels C and D additionally control for the respondent's average belief about the historical 12-month-ahead return for the three bins into which the perceived or actual realized PD ratio at the time the respondent took the survey did *not* fall. Robust standard errors are displayed below the different panels.

Table A.20: Summary statistics and balance check (PD experiment)

	PHF	Online brokerage sample							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	2017 Mean	Mean	Median	SD	p25	p75	Treatment Group: Mean	Control Group: Mean	p-value (7) = (8)
Female	0.49	0.17	0.00	0.38	0.00	0.00	0.14	0.20	0.017
Age	50.55	48.02	48.00	14.70	36.00	59.00	47.52	48.54	0.361
University	0.36	0.55	1.00	0.50	0.00	1.00	0.56	0.55	0.736
Employed	0.65	0.73	1.00	0.44	0.00	1.00	0.73	0.73	0.840
Household net income	3,808	3,943	4,000	2,698	2,000	5,250	4,017	3,865	0.461
Household net wealth	361,783	413,384	175,000	563,650	25,000	625,000	438,515	386,682	0.227
Total financial wealth at bank		85,783	42,500	113,146	15,000	112,500	93,950	77,106	0.050
Portfolio value at bank		52,175	25,000	71,210	7,500	62,500	57,143	46,897	0.058
Equity holdings at bank		44,890	17,500	65,553	2,500	56,250	49,433	40,062	0.060
Trades equity at least once per quarter		0.73	1.00	0.44	0.00	1.00	0.74	0.73	0.693
Risk tolerance (1-7)		4.31	4.00	1.23	4.00	5.00	4.28	4.34	0.467
Trading experience (years)		14.24	11.00	11.53	4.00	25.00	14.01	14.49	0.586
Financial literacy score (0-3)		1.84	2.00	0.80	1.00	2.00	1.85	1.83	0.727
Follow stock market developments (1-7)		4.55	5.00	1.85	3.00	6.00	4.51	4.58	0.615
Investment horizon \geq 5 years		0.59	1.00	0.49	0.00	1.00	0.60	0.58	0.613
Perceived current price-dividend ratio		39.08	35.00	16.90	30.00	50.00	39.54	38.58	0.455
Confident in perceived PD ratio		0.33	0.00	0.47	0.00	1.00	0.33	0.32	0.739
Expected return next 12 months		4.41	5.00	6.29	2.00	7.00	4.42	4.40	0.961
Confident in expected return		0.59	1.00	0.49	0.00	1.00	0.63	0.56	0.081
Perceived mean hist. ret. intervals		11.02	8.25	8.61	5.25	13.50	11.00	11.03	0.958
Perceived diff. high vs low PD historical		-1.75	-1.50	9.03	-5.33	2.67	-2.17	-1.31	0.206
Overestimator (diff. \geq -2)		0.55	1.00	0.50	0.00	1.00	0.56	0.55	0.857
Underestimator (diff. $<$ -10)		0.11	0.00	0.32	0.00	0.00	0.12	0.11	0.942
Observations		693					357	336	

Notes: This table shows summary statistics for the sample of retail investors at the online bank that we use in an additional experiment conducted in July and August 2022 (Columns 2-6), as well as benchmarks from the German population of individuals participating in the stock market as measured in the 2017 wave of the Bundesbank's Panel of Household Finance (Column 1). Columns 7-9 provide a check of balance of means between treatment and control group. Variables on income, wealth and wealth components are expressed in euro terms. Financial wealth, portfolio value and equity holdings at the bank as well as trading frequency are elicited in the survey, as no match with administrative account data was feasible. All belief variables reported in the table refer to respondents' priors elicited before the information treatment.

Table A.21: Correlates of beliefs (PD experiment)

	Perceived diff. high vs low PD	Over- estimator (diff. ≥ -2)	Neutral ($-10 \leq$ diff. < -2)	Under- estimator (diff. < -10)
	(1)	(2)	(3)	(4)
Female	1.595* (0.856)	9.310* (5.046)	-6.815 (4.655)	-2.495 (3.016)
Age	0.001 (0.030)	0.264* (0.160)	-0.407*** (0.150)	0.142 (0.098)
Employed	0.873 (0.876)	5.528 (4.632)	-6.371 (4.382)	0.843 (2.907)
University	-0.882 (0.706)	-3.425 (3.793)	0.454 (3.585)	2.971 (2.426)
Log(Household income)	-0.155 (0.128)	-0.395 (0.743)	-0.040 (0.712)	0.435 (0.458)
Log(Fin. wealth with bank)	0.107 (0.145)	0.089 (0.725)	-0.019 (0.714)	-0.069 (0.473)
Invest. experience \geq Median	-1.497* (0.868)	-9.128** (4.581)	8.450* (4.312)	0.678 (3.087)
Full financial literacy score	-1.963** (0.777)	-13.269*** (4.554)	12.934*** (4.492)	0.335 (3.022)
Follow stock market \geq Median	-0.172 (0.792)	-6.734 (4.295)	8.541** (4.133)	-1.808 (2.682)
Mean dep. var.	-1.75	55.41	33.19	11.40
SD dep. var.	9.03	49.74	47.12	31.80
Observations	693	693	693	693
R-squared	0.03	0.04	0.04	0.01

Notes: This table shows multivariate regressions of respondents' beliefs on covariates. The sample consists of retail investors participating in an additional experimental survey. The outcomes are the perceived difference in historical average 12-month-ahead returns between the highest PD scenario (higher than 50) and lower PD scenarios (Column 1) and dummies, multiplied by 100, for being classified as overestimator of this difference (difference at least -2 pp, Column 2), as neutral (differences at least -10 pp and less than -2 pp, Column 3) or as underestimator (difference less than -10 pp, Column 4). Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.22: Manipulation check (PD experiment)

	Low returns follow high PD (z)		When PD high, high returns can be expected (z)	
	(1)	(2)	(3)	(4)
PD treatment	0.619*** (0.073)		-0.482*** (0.074)	
PD treatment × Overestimator (diff. ≥ -2) (a)		0.749*** (0.100)		-0.689*** (0.102)
PD treatment × Neutral ($-10 \leq \text{diff.} < -2$)		0.471*** (0.121)		-0.300** (0.117)
PD treatment × Underestimator (diff. < -10) (b)		0.446** (0.199)		-0.047 (0.214)
Overestimator (diff. ≥ -2)	-0.182** (0.080)	-0.326*** (0.112)	0.266*** (0.080)	0.469*** (0.112)
Underestimator (diff. < -10)	-0.062 (0.115)	-0.048 (0.163)	-0.049 (0.120)	-0.181 (0.191)
p-value (a=b)		0.171		0.007
Observations	693	693	693	693
R-squared	0.20	0.21	0.20	0.22

Notes: This table shows estimations of the effect of the PD information treatment on respondents' posterior agreement with verbal statements describing beliefs about aggregate return predictability based on the current price-dividend ratio. The sample consists of retail investors participating in an additional experimental survey. Agreement with the statements is elicited on 7-point categorical scales, and is z-scored using the means and standard deviations in the sample. The statements are: "On average, low returns follow high price-dividend ratios." (Columns 1-2); "When, at a specific point in time, the price-dividend ratio is high, high returns can be expected over the following time." (Columns 3-4). Columns 2 and 4 show heterogeneous treatment effects for prior overestimators (perceived difference in average 12-month-ahead returns between the highest ($PD > 50$) and lower current PD scenarios at least -2 pp), neutrals (difference at least -10 pp and less than -2 pp), and underestimators (difference less than -10 pp). All estimations include a set of controls similar to the ones described in Appendix A.1.2. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

Table A.23: Updating of 12-month-ahead return expectations (PD experiment)

	Updating (point belief)		Updating (mean distr.)	
	(1) OLS	(2) IV	(3) OLS	(4) IV
PD treatment × Perception gap	0.116*** (0.044)	0.109* (0.060)	0.099* (0.057)	0.212*** (0.079)
Perception gap	-0.007 (0.028)	-0.016 (0.041)	0.033 (0.041)	-0.087 (0.059)
PD treatment	0.985*** (0.370)	0.993*** (0.362)	0.380 (0.525)	0.299 (0.511)
First stage F-stat		365.15		365.15
Observations	693	693	693	693
R-squared	0.09	0.09	0.08	0.07

Notes: This table examines changes in expectations about aggregate stock returns over the 12 months after the survey in response to the PD information treatment based on estimations of specification 2. The sample consists of retail investors participating in an additional experimental survey. The outcomes are the difference between posterior and prior point expectations about the 12-month-ahead return (Columns 1-2) and the difference between the mean of the respondent-level posterior distribution over 12-month-ahead returns and the prior point expectation (Columns 3-4). The perception gap is based on the respondent's prior belief about the historical return predictability based on the price-dividend ratio. It is the difference between the actual conditional mean 12-month-ahead return and the respondent's corresponding prior for the relevant scenario of current price-dividend ratios, which is selected based on the respondent's perceived current PD ratio. In Columns 2 and 4 the perception gap is instrumented with a version in which the relevant PD ratio interval is selected based on the actual PD ratio at the time of the survey. All estimations include a set of controls similar to the ones described in Appendix A.1.2. Robust standard errors are in parentheses. * denotes significance at 10%-, ** at 5%-, and *** at 1%-level.

B Survey instructions

This appendix provides the survey instructions translated to English for both our main survey conducted between September and November 2019 and the four-week follow-up survey. The instructions of the other surveys we conducted, which are listed in Table A.1, can be found under <https://drive.google.com/file/d/1IhxTWR5pXyVSR0V6tPfr4y06p7UwygmQ/view?usp=sharing>.

[X] has been anonymized for journal submission.

B.1 Main survey

Welcome screen

Welcome to the survey from [X]!

Many thanks for answering our questions on the investment behavior of retail investors.

Completion of the survey takes *about 15 minutes*. Your participation is of course anonymous. Your responses will only be used for scientific research.

In return for completing the survey you will receive an *Amazon voucher of 10 euros*. You will receive more detailed information on this during the survey.

To receive a voucher, please leave your email address at the end of the survey. We will send you your voucher code within the next days by email. Your email address will be saved separately from your responses in the survey, and will be deleted after we have sent out the voucher. You can learn more under our information on data protection.

Hint: The survey contains graphics that cannot be optimally displayed on smartphones. We therefore ask you to complete the survey *using a computer or a tablet* if possible.

Do you have questions? Please contact us under [X]

Attention check

The next question is about the following problem. In questionnaires like ours there are sometimes participants who do not read the questions carefully and only quickly “click” through the questionnaire. This results in many random answers, which compromise the quality of research studies.

To show us that you read our questions carefully, please select “Very interested” and “Not at all interested” as your answer to the next question.

How interested are you in sports?

Very strongly interested - Strongly interested - Somewhat interested - Almost not interested - Not at all interested

Perceived recent stock market return

Let us think about the *last 12 months*.

What do you think, what was the return (in percent) of the DAX over the *last 12 months*?

The return is the percent change in value of an investment in the DAX over the last 12 months. A positive number indicates that the value of the DAX has increased, a negative number indicates that the value has decreased.

___ percent

According to your estimate, an investor who 12 months ago invested 100 euro in the DAX would own X euro today.

How certain are you about your response?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Prior expected future stock market return

Let us now think about the future. What do you think, what will the return (in percent) of the DAX be over the *next 12 months*?

The return is the percent change in value of an investment in the DAX over the next 12 months. A positive number indicates that the value of the DAX increases, a negative number indicates that the value decreases.

___ percent

According to your estimate, an investor who today invests 100 euro in the DAX would own X euro 12 months from now.

How certain are you about your response?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Transition to main belief elicitation

On the following pages we would like to ask you about your estimates of the return of the DAX under *six different scenarios*.

Please take a moment to read the questions carefully. High attention in responding to the questions is essential for the quality of the results of this study.

Hint: Each question will be shown to you only once, and you will not be able to go back to previous questions later on.

Prior belief about historical autocorrelation of stock returns (quantitative): Return below -20 percent

First think about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *less than -20 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

The blue bar in the figure below illustrates your response.

___ percent

Prior belief about historical autocorrelation of stock returns (quantitative): Return between -20 and -10 percent

Please think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between -20 and -10 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

__ percent

Prior belief about historical autocorrelation of stock returns (quantitative): Return between -10 and 0 percent

Please think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between -10 and 0 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

__ percent

Prior belief about historical autocorrelation of stock returns (quantitative): Return between 0 and 10 percent

Please think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between 0 and 10 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

__ percent

Prior belief about historical autocorrelation of stock returns (quantitative): Return between 10 and 20 percent

Think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between 10 and 20 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

__ percent

Prior belief about historical autocorrelation of stock returns (quantitative): Return above 20 percent

Think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *above 20 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

__ percent

Transition to information treatment [Treatment group only]

On the next screen we will provide you with information on the *actual average returns* of the DAX in the different cases.

Please take a moment to read the information carefully.

Hint: The information will be shown to you only once and you will not be able to go back to the information.

Information treatment screen 1 [Treatment group only]

The figure below shows you the *actual average returns* of the DAX over the *following 12 months*, depending on what the return was over the *preceding 12 months*.

The figure is based on the returns of the DAX over the last 50 years.

Through repeated clicking on the button below you will be shown the actual average returns in the different scenarios. Only when you have seen the actual average returns in all six scenarios will you be allowed to proceed with the survey.

Information treatment screens 1a-1f [Treatment group only]

The figure below shows you the *actual average returns* of the DAX over the *following 12 months*, depending on what the return was over the *preceding 12 months*.

The figure is based on the returns of the DAX over the last 50 years.

Through repeated clicking on the button below you will be shown the actual average returns in the different scenarios. Only when you have seen the actual average returns in all six scenarios will you be allowed to proceed with the survey.

When the return over the preceding 12 months was below -20%, the return over the following 12 months was 9.5% on average (your estimate: A%).

When the return over the preceding 12 months was between -20% and -10%, the return over the following 12 months was 7.4% on average (your estimate: B%).

When the return over the preceding 12 months was between -10% and 0%, the return over the following 12 months was 9.5% on average (your estimate: C%).

When the return over the preceding 12 months was between 0% and 10%, the return over the following 12 months was 8.8% on average (your estimate: D%).

When the return over the preceding 12 months was between 10% and 20%, the return over the following 12 months was 8.7% on average (your estimate: E%).

When the return over the preceding 12 months was above 20%, the return over the following 12 months was 8.1% on average (your estimate: F%).

Information treatment screen 2 [Treatment group only]

Independently of the interval in which the return over the preceding 12 months was, the return of the DAX over the following 12 months was on average always between 7.4% and 9.5%.

This means that *regardless of the return of the DAX* over a particular year the *best forecast* of the return over the following year is close to the long-run historical mean return of 8.5%.

High or low stock market returns over a particular year hence do not allow to make a prediction about stock market returns over the following year.

Imagine one could predict at which point stock prices would increase by more than on average. Large institutional investors would then buy securities in large amounts. This would put stock prices under upward pressure. The possibility to predict higher-than-average returns would vanish immediately.

Control group information screen [Control group only]

Think now about the development of the DAX in the last 50 years. The average annual return of the DAX over this time period was

8.5 percent per year.

Posterior beliefs about autocorrelation of stock returns (qualitative)

To what extent do you agree with the following statements?

“With an investment in stocks one can expect a positive return, independently of how the stock market has developed in the recent past.”

“When the stock market has recently increased it makes no sense to buy stocks.”

“When the stock market has recently increased it is more likely that stock returns will be positive over the following time than when the stock market has recently decreased.”

1 (strongly disagree) - 2 - 3 - 4 - 5 - 6 - 7 (strongly agree)

Posterior expected future stock market return

Let us again think about the *next 12 months*. What do you think, what will the return (in percent) of the DAX be over the *next 12 months*?

___ percent

According to your estimate, an investor who today invests 100 euro in the DAX would own X euro 12 months from now.

Please explain your response in 1-2 sentences.

—

Posterior expected future stock market return: Subjective distribution

In the following we show you *6 possible scenarios* on how the DAX might develop over the *coming 12 months*.

Please indicate how likely you consider each scenario to be.

To do this, assign a probability to each scenario. The probabilities across the six scenarios have to sum to 100 percent.

Scenario 1: A return greater than 20%: ___ percent

Scenario 2: A return between 10% and 20%: ___ percent

Scenario 3: A return between 0% and 10%: ___ percent

Scenario 4: A return between -10% and 0%: ___ percent

Scenario 5: A return between -20% and -10%: ___ percent

Scenario 6: A return less than -20 %: ___ percent

How certain are you about your response?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Advanced financial literacy test

At the end we would like to ask you a few general questions.

What happens with the price of a bond if interest rates increase.

Rises - Falls - Remains unchanged - I do not know.

Which of the following statements is correct? If someone buys a stock of company B, then ...

... he owns a share in this company. - ... he lends money to company B. - ... he is liable for the liabilities of company B. - No response is correct. - I do not know.

Is the following statement true or false? The value of a call option for a stock is - everything else unchanged - higher, the more volatile the stock is.

True - False - I do not know.

How many of these questions have you answered correctly?

—

Background questions I

To what extent do you agree with the following statement? *“I closely follow the development of the DAX.”*

1 (strongly disagree) - 2 - 3 - 4 - 5 - 6 - 7 (strongly agree)

What information sources do you typically use before securities purchases / sales?

General news (e.g. newspapers, TV) - Specialized press (e.g. investment magazines) - Own online research - Chart analysis - Securities rankings (e.g. daily top ten) - I mostly trade with securities I already have / had in my portfolio - Recommendation from family / friends / acquaintances - Recommendation from a financial advisor - Recent stock price development - Other: ___

How difficult have you found the questions in this survey?

1 (not difficult at all) - 2 - 3 - 4 - 5 - 6 - 7 (very difficult)

Have you looked up additional information to answer the survey (e.g. google)?

No - Yes, namely: ___

Background questions II

For how many years have you been investing in stocks or stock mutual funds?

Enter 0 if you have no experience with investment in stocks or stock mutual funds.

For __ years.

When you personally make saving or investment decisions, how would you generally describe your *attitude toward risk*?

1 (not at all willing to take risks) - 2 - 3 - 4 - 5 - 6 - 7 (very willing to take risks)

How often do you trade with stocks on average?

Daily - Weekly - Once or twice per month - Once or twice per quarter - Once or twice per year - Less than once a year

What is your typical *investment horizon* for securities investments?

Less than 3 months - 3-12 months - 1-3 years - 3-5 years - 5-10 years - Longer than 10 years

Background questions III

What is your current *employment status*?

Full-time employed (including apprenticeships) - Part-time employed - Temporary leave (e.g. parental leave) - In school, university or unpaid internship - Unemployed - Permanent leave - Retired - Housekeeper - Other: __

Which of the following categories best describes your household's monthly *available net income*?

In your response, please account for all income of your household (e.g. also income from letting or leasing and child allowance). By household we mean all family members living with you at your main residence, excluding renters and flat mates.

[Categories]

Into which of the following categories falls your household's *net wealth*?

The net wealth is the value of everything the household members own (e.g. real estate, vehicles, financial assets, insurances) minus all liabilities (e.g. credit, loans, mortgages).

[Categories]

Background questions IV

Please indicate your gender.

Female - Male

In which year were you born?

[Dropdown menu]

What is your highest *educational attainment*?

Secondary school qualification - Secondary school certificate - Higher education entrance qualification - Higher education degree - No school-leaving qualification - Other: ___

Technical questions

On which device have you filled out the questionnaire?

PC, laptop or tablet - Smartphone - Other: ___

Did you experience a *technical issue* during the survey?

Yes - No.

Feedback questions

Did you have difficulties understanding one or more questions in this survey?

Yes - No

Do you have any suggestions or criticism related to our survey? Please let us know here (optional):

—

Payment and invitation to follow-up survey

Many thanks!

As a thank you for your responses you receive and *Amazon voucher of 10 EUR*.

If you would like to receive the voucher, you simply have to confirm this below and in a next step provide a valid email address. The voucher codes will be sent by email within the next 2 weeks.

Yes, I would like to receive a voucher code by email. - No, I would not like to receive a voucher code by email.

Are you interested in participating in a *follow-up survey*?

We would be happy to invite you to it by email.

Of course, we also reward participation in follow-up surveys with a bonus.

Yes, please invite me to a follow-up survey. - No, please do not invite me to a follow-up survey.

Please enter a *valid email address*:

Of course we will not give your email address to the bank or to third parties. Contacting you by email will be exclusively done to send you the voucher and / or invite you to a follow-up survey. After completion of this study your email address will be immediately deleted.

—

Goodbye screen

Many thanks for your participation!

You receive your *Amazon voucher of 10 EUR* in return for your participation *within the coming 2 weeks* by email.

B.2 Follow-up survey

Welcome screen

Welcome to the survey from [X]!

Many thanks for again taking the time to respond to our questions.

Completion of the survey takes *about 10 minutes*. Your participation is of course anonymous. Your responses will only be used for scientific research.

In return for completing the survey you will receive an *Amazon voucher of 5 euros*. We will send you your voucher code within the next days by email. Your email address will be saved separately from your responses in the survey, and will be deleted after we have sent out the voucher. You can learn more under our information on data protection.

Hint: The survey contains graphics that cannot be optimally displayed on smartphones. We therefore ask you to complete the survey *using a computer or a tablet* if possible.

Do you have questions? Please contact us under [X]

Posterior expected future stock market return I

Let us now think about the future. What do you think, what will the return (in percent) of the DAX be over the *next 12 months*?

The return is the percent change in value of an investment in the DAX over the next 12 months. A positive number indicates that the value of the DAX increases, a negative number indicates that the value decreases.

___ percent

According to your estimate, an investor who today invests 100 euro in the DAX would own X euro 12 months from now.

How certain are you about your response?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Posterior expected future stock market return: Subjective distribution

In the following we show you *6 possible scenarios* on how the DAX might develop over the *coming 12 months*.

Please indicate how likely you consider each scenario to be.

To do this, assign a probability to each scenario. The probabilities across the six scenarios have to sum to 100 percent.

Scenario 1: A return greater than 20%: ___ percent

Scenario 2: A return between 10% and 20%: ___ percent

Scenario 3: A return between 0% and 10%: ___ percent

Scenario 4: A return between -10% and 0%: ___ percent

Scenario 5: A return between -20% and -10%: ___ percent

Scenario 6: A return less than -20 %: ___ percent

How certain are you about your response?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Posterior beliefs about autocorrelation of stock returns (qualitative) I

To what extent do you agree with the following statements?

“With an investment in stocks one can expect a positive return, independently of how the stock market has developed in the recent past.”

“When the stock market has recently increased it makes no sense to buy stocks.”

“When the stock market has recently increased it is more likely that stock returns will be positive over the following time than when the stock market has recently decreased.”

1 (strongly disagree) - 2 - 3 - 4 - 5 - 6 - 7 (strongly agree)

Transition to main belief elicitation

On the following pages we would like to ask you about your estimates of the return of the DAX under *six different scenarios*.

Please take a moment to read the questions carefully. High attention in responding to the questions is essential for the quality of the results of this study.

Hint: Each question will be shown to you only once, and you will not be able to go back to previous questions later on.

Posterior belief about historical autocorrelation of stock returns (quantitative): Return below -20 percent

First think about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *less than -20 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

The blue bar in the figure below illustrates your response.

___ percent

Posterior belief about historical autocorrelation of stock returns (quantitative): Return between -20 and -10 percent

Please think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between -20 and -10 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

___ percent

Posterior belief about historical autocorrelation of stock returns (quantitative): Return between -10 and 0 percent

Please think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between -10 and 0 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

___ percent

Posterior belief about historical autocorrelation of stock returns (quantitative): Return between 0 and 10 percent

Please think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between 0 and 10 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

___ percent

Posterior belief about historical autocorrelation of stock returns (quantitative): Return between 10 and 20 percent

Think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *between 10 and 20 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

___ percent

Posterior belief about historical autocorrelation of stock returns (quantitative): Return above 20 percent

Think now about *all points in time* in the last 50 years at which the return of the DAX over the *preceding 12 months* was *above 20 percent*.

What do you think, what was the return of the DAX in these cases on average over the *immediately following 12 months*?

___ percent

Posterior belief about historical autocorrelation of stock returns (quantitative): Confidence

How confident are you about your responses in the six scenarios?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Perceived recent stock market return

Let us think about the *last 12 months*.

What do you think, what was the return (in percent) of the DAX over the *last 12 months*?

The return is the percent change in value of an investment in the DAX over the last 12 months. A positive number indicates that the value of the DAX has increased, a negative number indicates that the value has decreased.

___ percent

According to your estimate, an investor who 12 months ago invested 100 euro in the DAX would own X euro today.

How certain are you about your response?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Transition to repeated information treatment [Treatment group only]

On the next screen we will provide you with information on the *actual average returns* of the DAX in the different cases.

Please take a moment to read the information carefully.

Hint: The information will be shown to you only once and you will not be able to go back to the information.

Repeated information treatment screen 1 [Treatment group only]

The figure below shows you the *actual average returns* of the DAX over the *following 12 months*, depending on what the return was over the *preceding 12 months*.

The figure is based on the returns of the DAX over the last 50 years.

Through repeated clicking on the button below you will be shown the actual average returns in the different scenarios. *Only when you have seen the actual average returns in all six scenarios will you be allowed to proceed with the survey.*

Repeated information treatment screens 1a-1f [Treatment group only]

The figure below shows you the *actual average returns* of the DAX over the *following 12 months*, depending on what the return was over the *preceding 12 months*.

The figure is based on the returns of the DAX over the last 50 years.

Through repeated clicking on the button below you will be shown the actual average returns in the different scenarios. *Only when you have seen the actual average returns in all six scenarios will you be allowed to proceed with the survey.*

When the return over the preceding 12 months was below -20%, the return over the following 12 months was 9.5% on average (your estimate: A%).

When the return over the preceding 12 months was between -20% and -10%, the return over the following 12 months was 7.4% on average (your estimate: B%).

When the return over the preceding 12 months was between -10% and 0%, the return over the following 12 months was 9.5% on average (your estimate: C%).

When the return over the preceding 12 months was between 0% and 10%, the return over the following 12 months was 8.8% on average (your estimate: $D\%$).

When the return over the preceding 12 months was between 10% and 20%, the return over the following 12 months was 8.7% on average (your estimate: $E\%$).

When the return over the preceding 12 months was above 20%, the return over the following 12 months was 8.1% on average (your estimate: $F\%$).

Repeated information treatment screen 2 [Treatment group only]

Independently of the interval in which the return over the preceding 12 months was, the return of the DAX over the following 12 months was on average always between 7.4% and 9.5%.

This means that *regardless of the return of the DAX* over a particular year the *best forecast* of the return over the following year is close to the long-run historical mean return of 8.5%.

High or low stock market returns over a particular year hence do not allow to make a prediction about stock market returns over the following year.

Imagine one could predict at which point stock prices would increase by more than on average. Large institutional investors would then buy securities in large amounts. This would put stock prices under upward pressure. The possibility to predict higher-than-average returns would vanish immediately.

Control group repeated information screen [Control group only]

Think now about the development of the DAX in the last 50 years. The average annual return of the DAX over this time period was

8.5 percent per year.

Posterior beliefs about autocorrelation of stock returns (qualitative) II

To what extent do you agree with the following statements?

“When the stock market has fallen in the previous year one can expect above-average returns for the next year.”

“When the stock market has fallen over the previous 12 months there is a high probability that this trend will continue in the following 12 months.”

1 (strongly disagree) - 2 - 3 - 4 - 5 - 6 - 7 (strongly agree)

Posterior expected future stock market return II

Let us now think again about the *next 12 months*. What do you think, what will the return (in percent) of the DAX be over the *next 12 months*?

The return is the percent change in value of an investment in the DAX over the next 12 months. A positive number indicates that the value of the DAX increases, a negative number indicates that the value decreases.

___ percent

According to your estimate, an investor who today invests 100 euro in the DAX would own X euro 12 months from now.

Beliefs about historical frequencies of return scenarios

Think of all *12-month periods* in the *last 50 years*. Please give an estimate. In how many percent of cases did the DAX achieve a return in a given interval.

To do this, assign a response in percent to each scenario. The responses across the six scenarios have to sum to 100 percent.

Scenario 1: A return greater than 20%: ___ percent

Scenario 2: A return between 10% and 20%: ___ percent

Scenario 3: A return between 0% and 10%: ___ percent

Scenario 4: A return between -10% and 0%: ___ percent

Scenario 5: A return between -20% and -10%: ___ percent

Scenario 6: A return less than -20 %: ___ percent

How certain are you about your responses?

1 (not at all certain) - 2 - 3 - 4 - 5 - 6 - 7 (very certain)

Background questions I

Have you learned about topics related to *economics or business* in *school*?

Yes - No

Have you completed a *university degree* with focus on *economics or business*?

Please indicate your highest university degree with corresponding focus.

No, I have *not* completed a *university degree* with focus on economics or business. - Yes, I have completed a *Bachelor degree* with focus on economics or business. - Yes, I have completed a *Master degree* with focus on economics or business. - Yes, I have completed a *doctorate* with focus on economics or business.

Are you or have you been working in the financial sector?

Yes - No.

Background questions II

Did you follow the German stock market during the last 4 weeks?

not at all - a little bit - closely - very closely

To what extent do you agree with the following statement: "I always follow the development of the DAX."

1 (strongly disagree) - 2 - 3 - 4 - 5 - 6 - 7 (strongly agree)

On which device have you filled out the questionnaire?

PC, laptop or tablet - Smartphone - Other: ___

Do you have any suggestions or criticism related to our survey? Please let us know here (optional):

—

Payment and invitation to future surveys

Many thanks!

As a thank you for your responses you receive an *Amazon voucher of 5 EUR*.

If you would like to receive the voucher, you simply have to confirm this below and in a next step provide a valid email address. The voucher codes will be sent by email within the next 2 weeks.

Yes, I would like to receive a voucher code by email. - No, I would not like to receive a voucher code by email.

Are you interested in participating in a *follow-up survey*?

We would be happy to invite you to it by email.

Of course, we also reward participation in follow-up surveys with a bonus.

Yes, please invite me to a follow-up survey. - No, please do not invite me to a follow-up survey.

Please enter a *valid email address*:

Of course we will not give your email address to the bank or to third parties. Contacting you by email will be exclusively done to send you the voucher and / or invite you to a follow-up survey. After completion of this study your email address will be immediately deleted.

—

Goodbye screen

Many thanks for your participation!

You receive your *Amazon voucher of 5 EUR* in return for your participation *within the coming 2 weeks* by email.

References for the appendix

- Ansolabehere, Stephen, Marc Meredith, and Erik Snowberg**, “Asking about Numbers: Why and how,” *Political Analysis*, 2013, 21 (1), 48–69.
- Armona, Luis, Andreas Fuster, and Basit Zafar**, “Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment,” *Review of Economic Studies*, 2019, 86 (4), 1371–1410.
- Campbell, John Y and Robert J Shiller**, “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1988, 1 (3), 195–228.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia**, “Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments,” *American Economic Journal: Macroeconomics*, 2017, 9 (3), 1–35.
- Cochrane, John H.**, “The Dog that Did not Bark: A Defense of Return Predictability,” *Review of Financial Studies*, 2008, 21 (4), 1533–1575.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber**, “Monetary Policy Communications and Their Effects on Household Inflation Expectations,” *Journal of Political Economy*, 2022, 130 (6), 1537–1584.
- de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth**, “Measuring and Bounding Experimenter Demand,” *American Economic Review*, 2018, 108, 3266–3302.
- Dominitz, Jeff and Charles F Manski**, “Measuring and Interpreting Expectations of Equity Returns,” *Journal of Applied Econometrics*, 2011, 26 (3), 352–370.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus**, “Five Facts About Beliefs and Portfolios,” *American Economic Review*, 2021, 111 (5), 1481–1522.
- Golez, Benjamin and Peter Koudijs**, “Four Centuries of Return Predictability,” *Journal of Financial Economics*, 2018, 127 (2), 248–263.
- Greenwood, Robin and Andrei Shleifer**, “Expectations of Returns and Expected Returns,” *Review of Financial Studies*, 2014, 27 (3), 714–746.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart**, “Designing Information Provision Experiments,” *Journal of Economic Literature*, 2023, 61 (1), 1–38.
- Heiss, Florian, Michael Hurd, Maarten van Rooij, Tobias Rossmann, and Joachim Winter**, “Dynamics and Heterogeneity of Subjective Stock Market Expectations,” *Journal of Econometrics*, 2022, 231 (1), 213–231.
- Merton, Robert C.**, “Lifetime Portfolio Selection Under Uncertainty: The Continuous-time Case,” *Review of Economics and Statistics*, 1969, pp. 247–257.
- Rangvid, Jesper, Maik Schmeling, and Andreas Schrimpf**, “Dividend Predictability Around the World,” *Journal of Financial and Quantitative Analysis*, 2014, 49 (5-6), 1255–1277.
- Roth, Christopher and Johannes Wohlfart**, “How Do Expectations About the Macroeconomy Affect Personal Expectations and Behavior?,” *Review of Economics and Statistics*, 2020, 102 (4), 731–748.
- Vissing-Jorgensen, Annette**, “Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions,” *NBER Macroeconomics Annual*, 2003, 18, 139–194.