Extrapolators and Contrarians: Forecast Bias and Individual Investor Stock Trading*

Steffen Andersen
Danmarks Nationalbank, Copenhagen Business School, and CEPR
sta@nationalbanken.dk

Stephen G. Dimmock
National University of Singapore
dimmock@nus.edu.sg

Kasper Meisner Nielsen
Copenhagen Business School
kmn.fi@cbs.dk

Kim Peijnenburg
Tilburg University and CEPR
k.peijnenburg@tilburguniversity.edu

December 2025

Abstract

We test whether forecast bias affects individual investors' stock trading by combining bias measures from laboratory experiments with administrative trade data. Forecast bias is positively associated with past excess returns of purchased stocks: Compared to contrarians, extrapolators purchase stocks with higher past returns. Forecast bias is negatively associated with capital gains of sold stocks. Forecast bias also explains investor heterogeneity in the relation between market returns and net flows. Taken together, forecast bias provides a unifying mechanism through which different salient performance measures — past stock returns, capital gains, and past market returns — shape corresponding purchase, sale, and net flow decisions.

JEL Classifications: G5, G11, G41, D84, D81

Keywords: Extrapolation, Contrarian bias, Forecast bias, Expectations, Household finance, Experimental finance, Individual investors, Individual investor trading

* We thank Claes Bäckman, Milo Bianchi, Pedro Bordalo, Stefano Cassella, Patrick Coen, Zhi Da, Carey Deck, Andreas Fuster, Nicola Gennaioli, Francisco Gomes, Tobin Hanspal, Glenn Harrison, Lawrence Jin, Augustin Landier, Steffen Meyer, Gianpaolo Parise, Cameron Peng, Sugata Ray, Mirco Rubin, Mark Schneider, Michael Shin, Sean Shin, Paolo Sodini, Theresa Spickers, David Thesmar, Raman Uppal, Zhiqiang Ye, Terry Zhang, and conference/seminar participants at Aarhus University, Academy of Behavioral Finance and Economics, Australia National University Summer Camp, Bank of England Workshop, Copenhagen Business School, CEPR Household Finance Conference, ESSEC Business School, FIRS Conference, Frankfurt School of Management, Goethe University, Helsinki Finance Summit, IESEG Business School, London Business School, Research in Behavioral Finance Conference, SFS Cavalcade Asia-Pacific, Tilburg University, Toulouse School of Economics, University of Alabama, University of Geneva, University of Hawaii, University of Zurich, Workshop on Micro Data Meet Macro Models, and WHU – Otto Beisheim School of Management for helpful comments and suggestions. We also thank Akshaya Muthu and Phoon Ruei for excellent research assistance. We received ethics approval from the Copenhagen Business School Ethics Council. Andersen and Peijnenburg gratefully acknowledge support from the Center for Big Data in Finance (Grant no. DNRF167). Andersen gratefully acknowledges support from the European Research Council through ERC grant "RDRECON" (ID:639383). Dimmock gratefully acknowledges support from the Singapore Ministry of Education research grant A-0003873-00-00. Nielsen gratefully acknowledges support from the Danmarks Frie Forskningsfond (grant ID 10.46540/2033-00136B) and the Danish Finance Institute. The views expressed herein are those of the authors, and do not necessarily reflect the official views of Danmarks Nationalbank.

Expectations play a key role in both behavioral and rational models of investment decisions. A growing literature uses surveys to measure investors' expectations of stock market returns, and finds that recent past returns strongly affect investors, resulting in biased expectations. This literature also documents substantial heterogeneity across investors in how past returns affect their expectations.[1] Numerous studies further show that investors' expectations of market returns predict their allocations to risky asset classes.[2]

While these prior studies examine allocations to asset classes, there is little evidence of how biases in investors' expectations affect the selection of individual securities within an asset class. We elicit individual-level measures of forecast bias in a laboratory experiment and link them to administrative records of our subjects' stock trading decisions. Using these data, we provide the first direct evidence that forecast bias affects individual stock selection decisions for purchases and sales, as well as net flows into stocks.

We invite a representative sample of investors to participate in a laboratory experiment designed to elicit their forecast biases. Our experiment closely follows Afrouzi et al. (2023) and asks subjects to forecast a stochastic process. The subjects are eligible to win monetary prizes based on their forecast accuracy. At the start, each subject observes 40 realizations of the process and is then asked to forecast the next realization. After making this initial forecast, they observe the next realization of the process and then forecast the next realization. This

[1] De Bondt (1993), Fisher and Statman (2000), Greenwood and Shleifer (2014), and Adam, Matveev, and Nagel (2021) show that investors' stock market return expectations are strongly related to past returns, and argue that investors' update their beliefs in a biased manner. Dominitz and Manski (2011), Heiss et al. (2022), von Gaudecker and Wogrolly (2022), and Atmaz, Cassella, Gulen, and Ruan (2024) document large heterogeneity across investors' in how past returns affect stock market expectations.

[2] Vissing-Jorgensen (2003), Dominitz and Manski (2007), Malmendier and Nagel (2011), Amromin and Sharpe (2014), Merkle and Weber (2014), Hoffman, Post, and Pennings (2015), Giglio, Maggiori, Stroebel, and Utkus (2021), Beutel and Weber (2023), and Laudenbach, Weber, Weber, and Wohlfart (2024) show that investors' stock market expectations predict asset allocation decisions.

continues for a total of 40 rounds. Using these forecasts, we construct an individual-level measure of bias in belief formation, *Forecast Bias*.

Our aim is to elicit a general measure of our subjects' forecast bias and relate it to their individual stock trading decisions. Conceptually, variation in forecasts can arise from across-subject variation in information or from across-subject variation in information processing. Our laboratory experiment allows us to control and standardize the information provided, allowing us to measure variation in how subjects process information. By focusing on measuring information processing, rather than information itself, we obtain a single parameter that is widely applicable across different securities, time periods, and investment decisions.

Our estimates of forecast bias are consistent with earlier experimental studies (e.g., Dominitz and Manski, 2011; Afrouzi et al., 2023). On average, subjects have positive forecast bias (extrapolators): forecasts are too high following high realizations and too low following low realizations. There is, however, substantial heterogeneity across subjects. Many subjects have negative forecast bias (contrarians) and many have near-zero bias. We use this cross-sectional variation in lab-measured forecast bias to explain cross-sectional differences in real-world trading behavior.

We link the experimental results with 11 years of administrative register data on stock trading, income, wealth, and demographics from before and after the experiment. The trade-level data covers 2011-2021 and comprises all trades of every Danish investor, including our subject pool.

The combination of a laboratory experiment with administrative data offers several advantages. The laboratory experiment setting allows us to cleanly measure the forecast bias parameter, while controlling the underlying data-generating process and the information available to the subjects. The administrative data provide a representative sample of investors with complete and accurate records of their stock trading and holdings. Integrating experimental and administrative field data enables a deeper understanding of the heterogeneity in belief formation and its relation to stock selection decisions.

Theory shows that forecast bias increases sensitivity to recent past performance when forming expectations of future performance (Barberis, Greenwood, Jin, and Shleifer, 2015, 2018). In particular, extrapolation bias makes stocks that recently performed well more attractive while contrarian bias makes these stocks less attractive. Thus, theory predicts that higher extrapolation (contrarian) bias results in buying stocks with relatively high (low) past performance. Following similar logic, higher extrapolation (contrarian) bias results in selling stocks with relatively low (high) past performance.

Using our subject-specific measure, *Forecast Bias*, we test how subjects' biases interact with past stock performance to affect stock selection decisions. In these tests, we focus solely on days when a subject trades, and for each investor-day we construct the set of individual stocks the investor could plausibly trade. Our key regression specification explains the actual stock traded using an interaction between the investor's *Forecast Bias* and that individual stock's past performance, while including investor-day fixed effects. This approach allows us to examine how *Forecast Bias* alters the effect of an individual stock's past performance on trading decisions. The investor-day fixed effects control for aggregate market performance and remove the direct effect of time-variant (and time-invariant) investor and portfolio characteristics, such as past experiences, risk preferences, or wealth effects. By isolating the interaction between *Forecast Bias* and each individual stock's past performance relative to other stocks available to trade, we can identify how a subject's forecast bias shapes their stock selection among the set of feasible investment opportunities.

First, we test whether *Forecast Bias* affects stock purchase decisions. The results show that individuals' biases are related to the past returns of the stocks they buy: extrapolators tend to buy stocks with higher past returns and contrarians tend to buy stocks with lower past returns.[3] Supplementary tests reveal that a one-standard deviation increase in *Forecast Bias* implies purchasing stocks with a 3.0 percentage point higher past one-year return.

Second, we test whether *Forecast Bias* affects stock sale decisions. The results show that *Forecast Bias* is negatively associated with the capital gains of stocks that are sold and unrelated to the past one-year return. This result suggests that for sales, capital gains is the most salient measure of past performance used in forming expectations. A one standard deviation increase in *Forecast Bias* implies a 6.1 percentage point reduction in the probability that a subject sells a stock with a one standard deviation higher capital gain, relative to the baseline probability.

As we rely on individual heterogeneity in forecast bias, omitted variable bias is a potential concern. However, any omitted variable would need to explain both a monotonic relation between *Forecast Bias* and the past performance of purchases and sales, as well as the U-shaped relation between Forecast Bias and forecast error. For instance, interpretations based on cognitive abilities are unlikely, as both high and low values of *Forecast Bias* reflect deviations from rationality.

To address concerns about omitted variable bias, we conduct several additional tests. First, we confirm a monotonic relation between *Forecast Bias* and the past performance of

[3] Note that the relation between *Forecast Bias* and past returns does not necessarily indicate an investor mistake. In our lab experiment, the true data generating process is known and so forecast errors can be reliably labeled as bias in a normative sense. In contrast, the data generating process for stocks is unknown, and there is evidence for both momentum (Jegadeesh and Titman, 1993) and long-term reversals (De Bondt and Thaler, 1985). Accordingly, we do not take a position on the optimal trading strategy with respect to past returns, and instead adopt a positive stance, documenting the correlation between *Forecast Bias* and the past returns of purchased stocks.

purchases and sales. Second, our purchase and sales analyses are robust to including interaction terms of investor characteristics with past performance. These interactions are not accounted for in our baseline specification, as investor-day fixed effects absorb only the direct influence of investor characteristics on trading. Third, an assessment of omitted variable bias based on coefficient stability, as suggested by Oster (2019), indicates that omitted factors are unlikely to drive our estimated coefficients of interest.

We show that the results for purchases and sales are robust across several alternative specifications. First, we show that measurement error in *Forecast Bias* is limited and the results are similar when excluding subjects for whom measurement error is larger. Second, the results are robust when using alternative time-horizons for measuring past returns. Finally, to alleviate concerns about reverse causality, we show that the results are similar when analyzing only trades conducted after the subjects have participated in the experiment.

While our main set of results tests how forecast bias relates to cross-sectional security selection decisions, we also examine how forecast bias relates to net flows to stocks. We find that investors with higher forecast bias increase their allocations to stocks following high market returns over the past year relative to investors with lower forecast bias.

Finally, using the full universe of Danish investors, we document within-investor consistency in buying and selling behavior in relation to past performance. Investors who buy stocks with relatively high (low) past performance tend to sell stocks with relatively low (high) past performance. This pattern is consistent with forecast bias being a common mechanism affecting both purchases and sales, and supports the external validity of our results.

We contribute to the literature on individual investors' stock market expectations[4] and, in particular, to those studies that focus on the role of expectations in allocations to risky asset classes (e.g., see Vissing-Jorgensen, 2003; Malmendier and Nagel, 2011; Das, Kuhnen, and Nagel, 2020; Giglio, Maggiori, Stroebel, and Utkus, 2021; Charles, Frydman, and Kilic, 2024) and portfolio turnover (Liu, Peng, Xiong, and Xiong, 2022).[5] Our findings are also closely related to Laudenbach, Weber, Weber, and Wohlfart (2024), who show that beliefs about the historical autocorrelation of aggregate stock market returns relate to flows to the stock market during the COVID-19 crash, as well as Beutel and Weber (2023), who use an information experiment to show that beliefs affect risky asset allocations.

Our study differs from the prior literature in two significant ways. First, we directly measure forecast bias in a controlled laboratory setting instead of measuring beliefs (and inferring biases from stated beliefs). Second, this approach allows us to study how forecast bias affects individual stock selection decisions, whereas prior literature studies portfolio allocations to equities as an asset class. We provide the first direct evidence linking laboratory elicited biases in expectation formation to individual stock purchase and sale decisions. Our results have implications for theoretical models of extrapolation biases and investment decisions, highlighting the importance of considering investor heterogeneity in forecast biases.

Our study also contributes to the literature on how past performance affects individual investor decisions. Prior studies show that different past performance measures affect different types of investment decisions. The decision to purchase a stock is related to that stock's past

[4] Prior studies use survey data and show a positive relation between past returns and investors' stated expectations of future returns on the aggregate market (De Bondt, 1993; Fisher and Statman, 2000; Vissing-Jorgensen, 2003; Malmendier and Nagel, 2011; Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014) or investors' stated expectations of future returns on individual stocks (Da, Huang, and Jin, 2021). Across investors there is significant heterogeneity in how individual investors incorporate past returns into their expectations (Dominitz and Manski, 2011; von Gaudecker and Wogrolly, 2022; Laudenbach, Weber, Weber, and Wohlfart, 2024).

[5] In the Online Appendix, Liu, Peng, Xiong, and Xiong (2022) examine how extrapolation beliefs relate to past returns of purchases, but their tests pool market timing and cross-sectional security selection decisions.

return (Grinblatt and Keloharju, 2000; Barber and Odean, 2008). The decision to sell a stock is related to the investor's capital gain on that stock (Odean, 1998; Ben-David and Hirshleifer, 2012; Hartzmark, 2015). Decisions about net flows to stocks are related to past market returns (Greenwood and Shleifer, 2014). Our study contributes to this literature by showing that a single mechanism, forecast bias, affects how these different salient performance measures affect different types of investment decisions.

Finally, our study relates to work in asset pricing on extrapolation and contrarian biases. An extensive literature in asset pricing establishes stylized facts about stock returns and attributes these to forecast bias.[6] Although we do not directly study asset pricing topics in this study, our findings provide support for many of the forecast-based mechanisms posited in the literature. That is, we show a direct link between forecast bias and our subjects' stock selection decisions.

1. Eliciting Individuals' Forecast Bias

We conduct a laboratory experiment to measure our subjects' forecast bias. The experiment is designed to capture biases in how the subjects process information when forming expectations. This differs from much of the related literature, which uses survey measures of subjects' expectations of stock market returns to study investor decisions at the asset class level (e.g., Vissing-Jorgensen, 2003; Greenwood and Shleifer, 2014; Giglio, Maggiori, Stroebel, and Utkus, 2021). Our approach allows us to study investor security selection within an asset class by combining a simple measure of forecast bias with past stock returns of individual assets.

[6] For instance, the literature documents short-term momentum (Jegadeesh and Titman, 1993) and long-term reversal (De Bondt and Thaler, 1985; Lakonishok, Shleifer, and Vishny, 1994), which the authors attribute to investors' forecast biases. Similarly, several models explain cross-sectional return patterns by assuming investors suffer from forecast biases (e.g., Barberis, Shleifer, and Vishny, 1998; Hong and Stein, 1999; Cassella and Gulen, 2018; Bordalo, Gennaioli, La Porta, and Shleifer, 2019; Cassella, Chen, Gulen, and Petkova, 2022; Atmaz, Cassella, Gulen, and Ruan, 2024; Jin and Peng, 2024). See Barberis (2018) and Adam and Nagel (2022) for reviews of the literature on expectations and asset pricing.

Depending on each individual's level of *Forecast Bias*, past performance is interpreted differently, allowing us to examine security selection without the need to construct a time-series of each investor's expectations for every stock — a task that would be prohibitively complex.

1.1 The Elicitation Procedure

We develop an experimental module that includes a task to elicit individuals' forecast bias.[7] The task closely follows Afrouzi et al. (2023). In the forecast task, the subjects observe past values of an investment and then make forecasts about the future value.

The underlying data-generating function for the value of the investment is a first-order autoregressive (AR(1)) process with the first value set to 100:

$$x_{t+1} = 100 + 0.5 \cdot (x_t - 100) + \varepsilon_t. \tag{1}$$

The AR(1) coefficient is set to 0.5, the mean to 100, and the error term is drawn from a normal distribution with a standard deviation of 25.[8] Afrouzi et al. (2023) validate this method with a series of experiments, and show that forecast biases are similar with different parameter choices for the mean and standard deviation of the process, and across subject pools with different levels of sophistication (MIT students versus MTurk participants).

In our forecast task, subjects are not informed about the data generating process. Afrouzi et al. (2023) show that informing a sample of MIT engineering undergrads that the underlying data generating process is an AR(1) does not significantly alter their elicited biases, suggesting uncertainty about the true data generating process is unlikely to drive the observed forecast bias. The data generating process is not calibrated to stock prices and the forecasting horizon is generic, as we want to capture a general measure of forecast bias; however, prior research

---

[7] See Online Appendix A for the complete instructions of the forecasting experiment.

[8] Afrouzi et al. (2023) experiment with different values of the AR(1) coefficient in the range from 0 to 1 with 0.2 increments and find that overreaction is stronger for less persistent processes. We choose a single value for the AR(1) coefficient to ensure all responses are comparable.

suggests this design choice should not affect the results. For example, Landier, Ma, and Thesmar (2019) show that forecast biases are the same regardless of how the process is labelled.[9] Similarly, Frydman and Nave (2017) use a within-person design to show that subjects who exhibit extrapolative biases in a stock market experiment also exhibit similar biases in perceptual tasks.

Based on these findings, our experimental setting is likely to generate an interpersonal ranking between subjects' forecast bias that is independent of (the labelling of) the underlying process. Our tests of whether interpersonal variation in forecast bias relates to trading decisions assume that the interpersonal ranking between subjects' forecast bias is stable across different settings and autocorrelations – to the extent this assumption is incorrect, it will reduce the power of our tests and bias against finding significant results.

To begin, the subjects see 40 past realizations, and submit one- and two-period-ahead forecasts. Figure 1 provides a screenshot of the forecasting task.[10] The top panel shows the first 40 realizations as well as two "**x**"s, one blue and one orange, to indicate forecasts one period and two periods ahead, respectively. The subjects submit their forecasts for the next two periods by sliding the "**x**"s up or down to their desired value and clicking "Make forecast." Once the subject clicks "Make forecast," they observe the next realization, and are asked to make two new forecasts, as seen in the bottom panel of Figure 1. This step is repeated until each subject has submitted 40 rounds of forecasts. On average, the subjects take 9 minutes and 47 seconds to make the 40 rounds of forecasts, equivalent to four forecasts per minute, with only 27 (8) out of 959 subjects taking less than five (more than 20) minutes.

[9] Specifically, Landier, Ma, and Thesmar (2019) show that for a fixed stochastic process, subjects exhibit similar biases regardless of whether the process is labelled a "stable random process" or given an economic context (GDP growth, CPI, stock returns, or house price growth).

[10] The actual experiment is conducted in Danish. The caption "Værdien over tid" translates to "value over time." Online Appendix A contains an English translation of the complete instructions of the forecasting experiment.

To incentivize the subjects, in addition to the show-up fee, each subject has a 10% chance of being eligible to receive an incentive payment based on the accuracy of their forecasts. Each subject rolls a 10-sided dice to determine if they are eligible for the incentive payment, and if so, they roll a 4-sided and a 10-sided dice to randomly determine which of their 40 forecasts is selected to calculate their forecast accuracy. To ensure incentive compatibility and prevent risk aversion from affecting forecasts, we follow Hossain and Okui (2013) by letting the forecast accuracy affect the probability of winning a prize and not the amount of the prize. Thus, for the selected forecast, the subject's probability of winning a prize is: $100 - 5 \times \left| forecast_{i,t} - realization_{i,t} \right|$. If the forecast differs from the realized value by more than 20 in absolute terms, the probability of winning the prize is set to zero. The subject then rolls two 10-sided dice, and if the value from the roll is smaller than the winning probability, the subject receives 2,000 DKK (approximately €260).[11] Based on this procedure, 17 subjects received a prize from the forecasting task.

1.2 Measures of Forecast Bias

Using the forecasts elicited from the experiment, we construct our main measure of forecast bias at the individual level. We follow Afrouzi et al. (2023) and estimate the forecast bias that is implied by each subject's predictions using the following regression:

$$F_{i,t}(x_{i,t+1}) - E_{i,t}(x_{i,t+1}) = a_i + b_i \cdot (x_{i,t} - \bar{x}) + \varepsilon_{i,t} \tag{2}$$

where $F_{i,t}(x_{i,t+1})$ indicates subject i's forecast of next period's realization $x_{i,t+1}$ and $E_{i,t}(x_{i,t+1})$ is the rational forecast given the true data generating process.[12] Thus, the left-hand

[11] At the time of our experimental sessions, one DKK equals U.S. \$0.15-\$0.16 and €0.13.

[12] As Afrouzi et al. (2023) discuss, an advantage of measuring *Forecast Bias* using a laboratory experiment is that we know the true data generating process allowing us to reliably measure over- and under-reaction to information. Fuster, Laibson, and Mendel (2010) and Fuster, Hebert, and Laibson (2011) show that when the data generating process is unknown, it is difficult to separate forecast bias from uncertainty about the data generating process.

side of equation (2) is the subject's forecast error. The parameter b_i measures forecast bias.[13] A value of $b_i > 0$ indicates extrapolation bias: forecasts are too high (low) following high (low) realizations. A value of $b_i < 0$ indicates contrarian bias: forecasts are too low (high) following high (low) realizations.

Forecast Bias reflects systematic deviations from rational expectations, with individuals at both extremes exhibiting persistent departures from rationality. Consequently, any alternative interpretation of our findings must account for both (1) the U-shaped relation between *Forecast Bias* and forecast error and (2) the monotonic relation between our measure and past performance of purchases and sales, as documented in Section 3.3.

Each subject observes a unique series of realizations, which may differ by chance in persistence, mean, or variance. To account for these differences, we construct alternative measures of forecast bias. *Forecast Bias Residual* adjusts for subject-specific persistence and variance, while *Forecast Bias Limited Information* is estimated by comparing each subject's forecasts with the optimal within-sample forecast computed from that subject's prior realizations. *Forecast Bias Rank* is the percentile rank across subjects and is included to mitigate outliers. We also examine four additional measures that assume specific models of expectation formation: *Diagnostic Expectations* (Bordalo, Gennaioli, and Shleifer, 2018; Bordalo, Gennaioli, La Porta, and Shleifer, 2019), *Sticky Expectations* (Woodford, 2003), *Extrapolative Expectations* (Metzler, 1941), and *Adaptive Expectations* (Cagan, 1956).[14]

[13] Due to the small sample of forecasts, the OLS estimator of the persistence parameter of the AR(1) process is biased. The Kendall approximation that corrects for this bias is $b_i + \frac{1+3\vartheta}{T}$, which implies a bias of 0.06 for an AR(1) parameter of $\vartheta = 0.5$ and 40 forecasts. Our *Forecast Bias* measure thus underestimates the tendency to extrapolate. Note that the bias is consistent across subjects, and so does not affect our cross-sectional tests.

[14] The specifications of the alternative measures of forecast bias are the same as in Afrouzi et al. (2023), except we multiply Sticky Expectations by -1 so that it is directionally consistent with the other measures. See Online Appendix B for the exact specifications.

2. Data and Variables

Access to the data used in this study is provided by Statistics Denmark, the government agency with central authority for Danish statistics. We use the research infrastructure at Statistics Denmark to recruit subjects based on administrative register data and to conduct our laboratory experiment, as described later in this section. Statistics Denmark provides demographic, economic, and financial data, including stock holdings as well as trading records reported by banks and brokerage firms to the Danish Tax Authorities. The administrative registers are comprehensive and cover the entire Danish population.

2.1 Sample Recruitment and Lab Experiment

The starting point of our analysis is to recruit subjects for our experimental tasks. Statistics Denmark recruits the subjects using the following criteria provided by the authors. The initial population includes the 5,806,081 individuals residing in Denmark as of January 1, 2019. We then restrict the pool of eligible subjects in four steps. First, we exclude all individuals younger than 30 or older than 60, to remove students and retirees. Second, we exclude all individuals who do not reside within a 45-minute drive of the Statistics Denmark office in Copenhagen where the experiments are conducted. Third, we exclude individuals who are not homeowners for at least two years between 2014 to 2018. Finally, we exclude individuals who do not own at least 10,000 DKK in risky assets (stocks and mutual funds) in at least three of the years between 2014 and 2018.[15] After applying these criteria, the pool of eligible subjects contains 75,847 individuals. From the pool of eligible subjects, Statistics Denmark randomly invites 24,821 individuals to participate in our study.

[15] We exclude own company stock from this measure of risky assets.

In total, 959 subjects accept the invitation and participate in the experiment (3.9% participation rate). Online Appendix Table C.1 compares experiment participants and nonparticipants in terms of their demographic and economic characteristics, as well as their trading behaviors. Although some of the variables are significantly different, the magnitudes of the differences are small. Participants are slightly older (50.5 versus 49.5 years old), more educated (16.5 versus 16.0 years of education), more likely to be male (69% versus 56%), and less likely to be married (64% versus 69%) or have children (81% versus 86%). The differences are not significant for financial assets or housing wealth. For trading, sample participants trade slightly more and purchase stocks with slightly lower returns relative to the nonparticipants.

The experiment was conducted in-person, in sessions of around 15 subjects, which took place at Statistics Denmark in Copenhagen; Statistics Denmark obtained informed consent from all subjects. We conducted two sessions per day on 21 of the days between February 5, 2020 and March 11, 2020, at which time the experiment was suspended to comply with Covid protocols. The experiment was later resumed with an additional 12 days of two sessions per day between November 9 and 26, 2020.[16]

2.2 Measures of Forecast Bias

Figure 2 summarizes the measure of forecast bias for our subjects. We report the mean and quasi-median defined as the average value for the 45th through 55th percentiles. We report quasi-median instead of median, because our data agreement with Statistics Denmark prohibits reporting any statistics that are not based on at least three observations. Negative values of forecast bias imply contrarian bias relative to the data generating process: forecasts are too high (low) following low (high) realizations. Positive values imply extrapolation bias: forecasts are

[16] The average *Forecast Bias* is not significantly different pre- and post-Covid.

too high (low) following high (low) realizations. The mean (quasi-median) of the *Forecast Bias* parameter is 0.14 (0.28), which is significantly greater than the benchmark of zero (p-value < 0.0001).

The histogram in Figure 2 shows the distribution of *Forecast Bias*. Observations at zero indicate no bias, while observations to the left and right of zero indicate progressively greater contrarian bias and extrapolation bias, respectively. The heterogeneity is consistent with prior empirical studies such as Dominitz and Manski (2011), Heiss et al. (2022), von Gaudecker and Wogrolly (2022), and Laudenbach, Weber, Weber, and Wohlfart (2024). Although the point estimate of the bias is positive for 61.3% of subjects and negative for 38.7%, for 42.0% of the subjects the bias is not significantly different from zero at the 5% level. For simplicity of exposition, we refer to subjects with positive point estimates as extrapolators and those with negative as contrarians. We note, however, that our main analyses do not rely on these categorical classifications but rather on showing that the cross-sectional variation in lab-measured *Forecast Bias* is linked to cross-sectional variation in real-world trading behavior.

Online Appendix Table C.2 reports summary statistics of the alternative measures of forecast bias in Panel A: *Forecast Bias Residual*, *Forecast Bias Limited Info*, *Forecast Bias Rank, Diagnostic Expectations*, *Sticky Expectations, Extrapolative Expectations*, and *Adaptative Expectations*. Panel B shows that all measures are highly correlated and Panels C and D show that our results are robust to using these alternative measures.

As a simple check of reliability, we calculate two alternative measures identical to *Forecast Bias* except the first uses only the odd numbered rounds of forecasts and the second uses only the even numbered rounds (thus we have two measures per subject, each based on only 20 non-overlapping observations). Online Appendix Figure C.1 shows a close relation between *Forecast Bias Odd* and *Forecast Bias Even* using a scatter plot and the estimated regression line. *Forecast Bias Odd* and *Forecast Bias Even* have a correlation coefficient of 0.801. The Spearman-Brown reliability score between these two variables is 0.889, indicating high internal consistency. These findings provide reassurance that *Forecast Bias* is measured with limited measurement error at the individual level. Section 3.4 further addresses potential measurement error concerns.

Online Appendix Table C.3 shows summary statistics of the relation between *Forecast Bias* and individual characteristics. We define contrarians as subjects with negative *Forecast Bias* and extrapolators as subjects with positive *Forecast Bias*. Although some individual characteristics are statistically different, the magnitudes of the differences are small. Overall, *Forecast Bias* has little relation with economic and demographic characteristics, preferences, and proxies of cognitive abilities.[17]

Conceptually, optimism and overconfidence could cause subjects to have biased expectations about the mean of the stochastic process. But optimism and overconfidence are unlikely to affect our measure, because *Forecast Bias* does not capture a persistent upward or downward bias. Rather, our measure captures the directional response in forecasts to recent realizations of values (e.g., forecasts that are both consistently too high following high realizations and too low following low realizations).

2.3 Trading and Portfolio Data

We combine data from several administrative registers provided by Statistics Denmark. Income, wealth, and investments data come from the official records of the Danish Tax and Customs Administration (SKAT) for the years 2011 to 2021, and are comparable to the data

[17] Online Appendix Table C.3 shows statistically significant differences in proxies for quantitative reasoning ability for different levels of *Forecast Bias*, but the magnitudes of the differences are small.

from other Nordic countries.[18] Danish tax law requires third parties to report information on income, wealth, and trading directly to SKAT. For example, banks and brokerages report investment holdings and trades at the individual level. Thus, our trading data are reported directly from administrative sources and are not self-reported by individuals. The data contain information on individuals' stock holdings by ISIN number at the end of the year as well as daily records of all stock transactions, including both domestic and international stocks.[19] We supplement this information with demographics from the Civil Registration System and educational records from the Ministry of Education. We match the data at the individual level using the civil registration number, which is the Danish equivalent of the social security number in the United States.

A total of 680 of the 959 (71%) participants in the experiment purchase at least one stock between 2011 and 2021 and a total of 583 of the 959 (61%) individuals sell at least one stock between 2011 and 2021. Panel A of Table 1 summarizes the purchases and sales. On average, subjects make 44 purchases and 33 sales, for a total of 50,298 unique trades.[20] The distribution of trading activity is highly skewed, with the 52 most active traders making about half of total trades. To ensure that our results are not driven by a few extremely active investors, our empirical specifications weight subjects equally. The average purchase has a value of 59,293 DKK (€7,708) and the average sale is 82,819 DKK (€10,766). In aggregate, the value of trades in our sample is slightly greater than 3.4 billion DKK (€442 million).

We supplement the administrative data with return data from Refinitiv and Compustat Global. The Refinitiv data are matched using ISIN codes. The Compustat data are matched

[18] For example, Grinblatt and Keloharju (2000), Kaustia and Knüpfer (2008), and Knüpfer, Rantapuska, and Sarvimäki (2017, 2023) study data from Finland; Hvide and Östberg (2015) and Fagereng, Guiso, Malacrino, and Pistaferri (2020) study data from Norway; and Calvet, Campbell, and Sodini (2007, 2009) study data from Sweden.

[19] In our sample of trades, 56.6% of purchases and 59.7% of sales between 2011 and 2021 are of Danish stocks.

[20] We aggregate trades in the same stock within a day to get unique investor-stock-day purchases. Our sales variable includes both partial sales and full divestment of a position.

using the GVKEY to ISIN mapping files provided by Capital IQ. For benchmark returns, we use the WRDS World Index for Denmark.

For each security, we calculate returns using daily data for the year ending the day before a purchase or sale. We report averages including returns of both the traded stocks as well as the stocks included in the consideration set of individuals but ultimately not traded. The average lagged annual return in our purchase sample is 26.2%, but returns are positively skewed and the quasi-median is 8.9% (for comparison, the average lagged annual return for the Danish stock market is 20.0% for our sample). The average lagged annual return in our sales sample is 29.3%, and the quasi-median is 13.5%. We also calculate the capital gain since purchase for the stocks in our sample.[21] The average capital gain is 23.1% and the quasi-median is 3.4%.

3. Forecast Bias and Stock Trading Decisions

3.1 Forecast Bias and Stock Purchases

We hypothesize that *Forecast Bias* affects how investors react to stock performance. Specifically, extrapolators purchase stocks with higher past returns than those purchased by contrarians. We test how *Forecast Bias* and past returns interact to affect which stock is purchased by the investor. We examine purchase decisions at the investor-stock-day level, and limit our sample to include only days in which the investor makes at least one purchase (Section 4 considers the issue of whether to trade). These tests require data on both the stocks the investor purchases and those they do not.

Individual investors have limited time and attention and are unlikely to evaluate all of the thousands of available stocks. Instead, each investor will focus on a smaller set of stocks that have captured their attention via news coverage, financial advice, or conversations with

[21] For individuals who make multiple purchases of stocks over time, we use the weighted average capital gain per share. For partial sales when the subject has purchased shares in multiple tranches at different prices, we assume the subject sells from each tranche on a pro rata basis.

friends and family. This set of stocks forms the investor's consideration set. Ideally, for each purchase we would know this consideration set. However, as the consideration set is unobserved, we construct a proxy set that includes the purchased stock and all stocks meeting the following requirements: (1) The stock is traded at least once by one of our subjects during the sample period and (2) At least one investor in Denmark (considering all investors in the country, not just those in our sample) purchased the stock during the month. This creates a comprehensive proxy consideration set, containing the overwhelming majority of plausibly considered stocks, but is likely overinclusive containing far more stocks than the investor actually considers: for an average investor-purchase day, there are 1.2 purchased stocks and 1,434 stocks in the proxy consideration set.[22]

Using this sample, we estimate the following linear probability model:

$$Purchase_{i,j,t} = \beta_1 \cdot ForecastBias_i \cdot Perform_{j,t} + \beta_2 \cdot Perform_{j,t} + \boldsymbol{\delta X_{i,j,t}} + \theta_{i,t} + \varepsilon_{i,j,t} \qquad (3)$$

where $Purchase_{i,j,t}$ is an indicator variable equal to 100 if subject i buys stock j on date t. $Perform_{j,t}$ is the lagged annual return of stock j as of the end of the prior trading day, winsorized at the 1st and 99th percentiles to ensure the results are not driven by outliers. $\boldsymbol{X_{i,j,t}}$ is a matrix of control variables, and $\theta_{i,t}$ is an investor-day fixed effect. The unit of observation is investor-stock-day. Because trading activity is highly skewed, we use weights to ensure that each person in our sample contributes equally to the estimation. Thus, our estimates reflect the effects of *Forecast Bias* for an average person in our sample, corresponding to the economic concept we are interested in – how people's *Forecast Bias* affects their trading choices (see

[22] Online Appendix Table C.4 shows that our results are robust to using different consideration sets. In particular, we vary the first requirement: (1) The stock is traded at least once by more than one of our subjects in our sample during the sample period, (2) The stock is traded at least once by more than three of our subjects during the sample period, and (3) The stock is traded at least once by more than five of our subjects during the sample period. These alternative consideration sets contain 815, 426, and 307 unique stocks, respectively. As the number of stocks in the consideration set is reduced, the fraction of the dependent variable equal to one increases, thereby mechanically increasing the coefficient of interest.

Solon, Haider, and Wooldridge, 2015). The standard errors reported below the coefficients are clustered at the investor level.

Conditioning on the investor's decision to trade on a given day, the investor-day fixed effect absorbs all sources of variation common to that investor and the market on that day, including individual and portfolio characteristics. As a result, the specification exploits only the variation across stocks, and the investor's reaction to that variation. The investor-day fixed effect helps to address concerns about omitted variable bias. First, it removes the overall market performance, and thus the regressions examine the likelihood an investor chooses to purchase a stock given its performance relative to other stocks available at that point in time. Thus, we examine the relation between forecast bias and security selection, without confounding this relation with market timing. Second, it removes the direct effect of any investor and portfolio characteristic at that point in time, such as wealth, risk-aversion, or past experiences. One caveat with this approach is that the indirect effect from interactions between investor characteristics and past performance is not absorbed, a concern we address in Section 3.3.

We include several variables to control for the likelihood of individuals to purchase certain stocks, irrespective of past return. $Held\ before_{i,j,t}$ and $Current\ holding_{i,j,t}$ are indicator variables equal to one if subject i has ever held stock j and if subject i currently holds stock j, respectively. In our sample, 55.3% of purchases are of stocks that the investor either currently holds or has previously held, and 41.9% are of stocks currently held. $Portfolio\ weight_{i,j,t}$ is the weight of a current holding in the portfolio (and zero for stocks not currently held). $Stock\ purchase\ fraction_{j,m-1}$ is constructed from the full sample of all Danish investors. It is defined as the percentage of all investor-stock purchases in the preceding month, m-1, that were in stock j. This variable measures the popularity of stock j among Danish investors.[23]

As discussed earlier, we hypothesize that the coefficient on the interaction term, $ForecastBias_i \cdot Perform_{j,t}$, will be positive. Extrapolators, who have positive *Forecast Bias*, tend to buy stocks with higher past returns[24] than contrarians, who have negative *Forecast Bias*. Consistent with the predictions of theory, the coefficient in column (1) of Table 2 is positive and significant.[25] Relative to the baseline probability of 0.08%, subjects with a one standard deviation higher level of forecast bias have a 4.4% higher probability of buying a stock with a one standard deviation higher past return. Column (2) of Table 2 reports coefficients from a conditional logit regression that conditions out investor-day effects. Aside from the different estimation method, the specification is the same as in column (1). The results are similar to those in column (1).

Online Appendix Table C.6 presents coefficient estimates allowing us to quantify the economic magnitude of the relation in returns rather than probabilities. Specifically, we limit the sample to include only investor-stock-day purchases, use the stock's excess return as the dependent variable, and *Forecast Bias* as the independent variable. The coefficient estimate in

[23] Including this control variable causes the coefficient on the performance measure, the lagged 12-month returns to become negative. The coefficient on $Stock\ purchase\ fraction_{j,m-1}$ is positive and significant.

[24] Our main specification uses lagged one-year returns as the performance measure. The past year is a natural evaluation period as brokerages and financial media often report returns over the past year, and this is a commonly used period in the literature on past returns and individual investors' decisions (e.g., see Barber and Odean, 2002; Laudenbach, Weber, Weber, and Wohlfart, 2024). As a robustness test, we evaluate lagged returns over three, six, and 36-month time-periods. The results in Online Appendix Table C.5 show that the coefficient on the interaction term is significant for the six and 36-month periods.

[25] We conduct several robustness tests. First, Panel C of Online Appendix Table C.2 shows that our results are robust when using alternative measures of *Forecast Bias*. Further details regarding these alternative measures are provided in Online Appendix B. Second, Online Appendix Table C.7 shows robustness results including stock, stock-year, stock-month, or stock-day fixed effects. Third, as investors with *Forecast Bias* close to zero may be effectively unbiased, we re-estimate our baseline specification in equation (3) on a subsample that excludes those with *Forecast Bias* not different from zero at the 5% and 10% significance level (Panel A Online Appendix Table C.8). Our results are unaffected.

column (2) implies that a one standard deviation increase in *Forecast Bias* is associated with buying stocks that had 3.0 percentage points higher excess returns over the past year.

Table 3 reports an alternative approach examining the relation between *Forecast Bias* and stock purchase decisions. We partition the sample into subjects with positive *Forecast Bias* and subjects with negative *Forecast Bias*, which we label extrapolators and contrarians, respectively. We then identify the effect of *Forecast Bias* by comparing how past performance affects the trading of individuals with positive versus negative *Forecast Bias*.

We estimate the following linear probability model:

$$Purchase_{i,j,t} = \beta_1 \cdot Lagged\ annual\ return_{j,t} + \boldsymbol{\delta X_{i,j,t}} + \theta_{i,t} + \varepsilon_{i,j,t}, \tag{4}$$

where the control variables and fixed effects are the same as in the baseline specifications (column (1) of Table 3). The inclusion of investor-day fixed effects implies that we are comparing the relative return performance of stocks on a given day, rather than their absolute returns.

Column (1) of Table 3 shows the relation between stock purchases and past returns for all investors in our sample. The literature finds mixed results on the relation between past returns and purchases by individual investors. Some studies find individual investors are net buyers of stocks with relatively low past returns (e.g., Nofsinger and Sias, 1999; Grinblatt and Keloharju, 2000, 2001; Kogan, Makarov, Niessner, and Schoar, 2024). Other studies find they buy stocks with high past returns (e.g., Odean, 1998; Da, Huang, and Jin, 2021). Our results, in column (1), show that the likelihood of our subjects purchasing a stock declines when the stock's past return is higher.

Columns (2) and (3) show the results for the extrapolator (*Forecast Bias*>0) and contrarian (*Forecast Bias*$\leq$0) subsamples. Our main hypothesis for purchases is that the coefficient on past returns will be larger for extrapolators than contrarians, consistent with a positive relation between *Forecast Bias* and the past returns of purchased stocks. The results show that the relation between past returns and purchases is more negative for contrarians than for extrapolators. More importantly, the difference between the coefficients is statistically significant at the 1% level (p-value = 0.002). Overall, the results are consistent with our main finding that investors with higher *Forecast Bias* purchase stocks with higher past returns.

3.2 *Forecast Bias and Stock Sales*

Conceptually, the relation between *Forecast Bias* and stock sales mirrors that of purchases: Extrapolators prefer to sell stocks with lower past performance than those sold by contrarians. There is, however, an important distinction between sales and purchases. While purchases can be selected from the entire universe of stocks, sales are almost exclusively selected from the more limited set of existing holdings.[26] The mean (quasi-median) number of individual stocks held at the time of a sale is 11.5 (8.5). This enables investors to easily compare performance across all their holdings, providing a clearly defined consideration set. Thus, for the sale regressions we limit the sample to include only stocks held by the investor as of the end of the previous day.

For sales, we estimate a linear probability model similar to that in Eq. (5) for purchases:

$$Sale_{i,j,t} = \beta_1 \cdot ForecastBias_i \cdot Perform_{i,j,t} + \beta_2 \cdot Perform_{i,j,t} + \boldsymbol{\delta X_{i,j,t}} + \theta_{i,t} + \theta_l + \varepsilon_{i,j,t} \quad (5)$$

where $Sale_{i,j,t}$ is an indicator variable equal to 100 if subject i sells stock j on date t and $Perform_{i,j,t}$ is a measure of investor i's performance on stock j as of the end of the prior trading day (e.g., capital gain since purchase), winsorized at the 1st and 99th percentiles to ensure

[26] Although short sales can be made from the entire universe of stocks, shorting is extremely rare for individual investors (e.g., Barber and Odean, 2008).

the results are not driven by outliers.[27] $\boldsymbol{X_{i,j,t}}$ is a matrix of control variables, $\theta_{i,t}$ is an investor-day fixed effect and θ_l is a fixed effect for the length of the holding in months. We limit the sample to include only days when the subject sells at least one stock. The unit of observation is investor-stock-day. Because trading activity is highly skewed, we estimate weighted regressions such that each subject receives equal weight. The standard errors are clustered at the investor level.

Including the investor-day fixed effect means the specification examines the likelihood that an investor chooses to sell a stock *relative* to that of other stocks held in the portfolio.[28] We include a length of holding period fixed effect as prior studies show a strong relation between length of holding and sales (e.g., see Ben-David and Hirshleifer, 2012; Hartzmark, 2015). We include two control variables. $Portfolio\ weight_{i,j,t}$ is the weight of stock j in investor i's portfolio. $Stock\ sales\ fraction_{j,m-1}$ is constructed from the full sample of all Danish investors, and is defined as the percentage of all investor-stock sell trades in the preceding month, *m-1*, that were in stock j.[29]

Table 4 reports regression results for sales and considers two different performance measures. Column (1) follows the purchase specification and uses lagged annual returns. Columns (2) and (3) use capital gain since purchase. *Forecast Bias* measures how investors incorporate past performance into expectations. For sales, investors have a performance measure that is not available for purchases – their capital gain on a stock – and a large literature

[27] Because we require holding-level capital gains, the sample for the sales regressions begins in 2013 as we cannot observe purchase prices before 2011. For positions that were initiated prior to January 1, 2011, we calculate capital gains using the price as of the end of December 2010.

[28] Panel D of Online Appendix Table C.2 demonstrates that our results are robust when using alternative measures of *Forecast Bias*. Further details regarding these alternative measures are provided in Online Appendix B. Online Appendix Table C.7 shows robustness results including stock, stock-year, stock-month, or stock-day fixed effects.

[29] The purchase regressions also include $Held\ before_{i,j,t}$ and $Current\ holding_{i,j,t}$ as controls. We do not include these variables in the sale regressions because the sample includes only stocks the investor currently holds.

shows capital gains are strongly related to sales decisions (e.g., Odean, 1998 Ben-David and Hirshleifer, 2012; Hartzmark, 2015). Capital gains are salient and typically easy to observe in brokerage accounts,[30] making it likely investors pay attention to them and incorporate them into their forecasting process (Frydman and Rangel, 2014; Frydman and Wang, 2020).

In column (1), the performance measure is lagged annual returns.[31] In this specification, the interaction between performance and *Forecast Bias* is not significant. In columns (2) and (3), the performance measure is capital gains since purchase. Column (2) reports a linear probability model and column (3) reports a conditional logit regression. In both columns, the coefficient on the interaction term between *Forecast Bias* and capital gain is negative and significant. Comparing across positions held in their portfolio, investors with higher *Forecast Bias* are more likely to sell stocks with lower capital gains. The result in column (2) implies that subjects with a one standard deviation higher level of *Forecast Bias* have a 3.8% lower probability of selling a stock with a one standard deviation higher capital gain, relative to the baseline sales probability of 23.2%. The contrast between the results for the two performance measures is consistent with the personally experienced performance measure, capital gains, being more salient for sales decisions.

As we did for purchases, we partition the sample into extrapolators and contrarians based on the sign of *Forecast Bias* and compare the relation between past performance and trading decisions for these two groups. We estimate the following linear probability model:

$$Sale_{i,j,t} = \beta_1 \cdot Capital\ gains_{i,j,t} + \boldsymbol{\delta X_{i,j,t}} + \theta_{i,t} + \theta_l + \varepsilon_{i,j,t}, \tag{6}$$

[30] The interfaces of three major brokerage accounts in Denmark, accessible to the authors, prominently display the percentage capital gains since purchase for stocks held in their portfolio, not lagged returns.

[31] Column (1) includes the stocks' lagged returns over the past year as an independent variable and this causes the loss of observations for which we do not observe a full year of returns.

where the control variables and fixed effects are the same as in the baseline specifications (column (2) of Table 4). Column (1) of Table 5 shows that, in the full sample, the relation between capital gains and sales decisions is not significant. The results in columns (2) and (3) show that the relation between capital gains and sales is negative for extrapolators and insignificantly positive for contrarians. The coefficients in the two subsamples are significantly different at the 5% level (p-value = 0.029). This result supports our baseline finding that investors with higher *Forecast Bias* tend to sell stocks with lower capital gains.

Our findings complement the literature on the relation between capital gains and investor selling decisions. Much of this literature centers on the disposition effect: the empirical pattern that, on average, investors are more likely to sell winners than losers, and explanations for that pattern. It is important to establish that *Forecast Bias* is distinct from previously documented explanations for the disposition effect. We address this in two ways. First, the leading explanations for the disposition effect are preference-based[32] (e.g., realization utility as in Barberis and Xiong, 2012). Our elicitation method is designed to avoid contamination from preferences (see Hossain and Okui, 2013). Second, Online Appendix Table C.9 reports additional tests addressing this issue. Column (1) shows our results are robust to including additional controls for the square and cube of capital gains, an indicator for positive capital gains (capital gains > 0), and indicators for the highest- and lowest-performing stocks within each investor's portfolio (Ben-David and Hirshleifer, 2012; Hartzmark 2015). Column (2) adds an interaction between an indicator for positive capital gains and *Forecast Bias*, because a discontinuity at zero capital gains is a central prediction of the disposition effect; the coefficient on this interaction term is not significant, indicating that the relation between *Forecast Bias* and

---

[32] However, there are some belief-based explanations of the disposition effect (Ben-David and Hirshleifer, 2012; Andersen, Hanspal, Martínez-Correa, and Nielsen, 2021).

sales is not driven by the sign of the capital gain, and that the pattern we document is distinct from the disposition effect.

3.3 *Omitted Variable Bias*

Forecast Bias is designed to capture how individuals process information to form expectations. Although we do not know exactly what information our subjects process (e.g., past returns, media coverage, etc.), past performance should capture the general direction and magnitude of this information. Our results are consistent with this interpretation: compared to contrarians, extrapolators purchase (sell) stocks with higher (lower) past performance.

As we rely on individual heterogeneity in forecast bias, omitted variable bias is a potential concern. Any alternative interpretation, based on omitted variable bias, must account for two key patterns: (1) the relation between *Forecast Bias* and the past performance of purchases and sales is monotonic, and (2) the relation between *Forecast Bias* and forecast errors is U-shaped. By design, *Forecast Bias* quantifies directional errors relative to the rational benchmark; negative values correspond to excessive contrarianism and positive values to excessive extrapolation. Hence, individuals with both positive and negative values of *Forecast Bias* systematically deviate from the rational benchmark, but in opposite directions. Thus, any alternative interpretation must explain not only that individuals who make forecast errors trade based on past returns, but also that the direction of the forecast errors observed in the lab align with the direction of the relation between past performance and trading decisions.

We test whether there is a monotonic relation between *Forecast Bias* and the lagged performance of stock purchases and sales. We divide our subjects into five groups based on their *Forecast Bias*. We then interact an indicator variable for each bin with lagged performance, and estimate regressions analogous to our baseline specifications. Specifically,

$$Purchase_{i,j,t} = \sum_{b=1}^{5} \beta_b \cdot \mathbb{I}_{i,b} \cdot Lagged\ 12\ month\ return_{j,t} + \boldsymbol{\delta X_{i,j,t}} + \theta_{i,t} + \delta_{j,t} + \varepsilon_{i,j,t} \quad (7)$$

$$Sale_{i,j,t} = \sum_{b=1}^{5} \beta_b \cdot \mathbb{I}_{i,b} \cdot Capital\ gains_{i,j,t} + \boldsymbol{\delta X_{i,j,t}} + \theta_{i,t} + \delta_{j,t} + \theta_l + \varepsilon_{i,j,t} \quad (8)$$

where $\mathbb{I}_b$ is an indicator variable equal to one if subject *i*'s *Forecast Bias* falls in bin *b*. We generate five bins of *Forecast Bias* using cut point of [-1, -0.25, 0.25, 1] as these are symmetric around zero and ensure sufficient observations in each bin. Figure 3 shows a monotonic relation: Moving from the lowest *Forecast Bias* bin to the highest, subjects purchase (sell) stocks with progressively higher (lower) performance. This evidence raises the bar for omitted variable explanations, as any alternative interpretation of *Forecast Bias* must reconcile the monotonic relation between *Forecast Bias* and past performance with the non-monotonic relation between *Forecast Bias* and deviations from the rational benchmark.

We perform additional robustness tests to further alleviate concerns regarding omitted variable bias. To begin, Online Appendix Table C.3 presents summary statistics of characteristics for extrapolators and contrarians. Although some differences are statistically significant, their magnitudes are small, suggesting that extrapolators and contrarians have similar characteristics.

Online Appendix Table C.10 replicates our main purchase specification (Table 2, column 1) and sales specification (Table 4, column 2) but further includes interactions between past performance and the following investor characteristics: education, financial literacy, numeracy, trust, optimism, and overconfidence.[33] The investor-day fixed effects remove the direct effect of all investor characteristics that do not interact with past returns, and the new

[33] This set of control variables follows Dimmock et al. (2016, 2021). These control variables have missing values, which we impute using median imputation. For these seven variables, on average about 1% of the observations are missing.

interaction terms capture any within person-day variation in trading behavior related to the correlation between individual characteristics and past performance. The results in Online Appendix Table C.10 show that the coefficient estimates are stable across specifications and only three out of 28 coefficients on interaction terms are significant.

Following the recommendations of Oster (2019) for evaluating omitted variable bias, we compare the within-R^2s and coefficient estimates of models without controls to the full models reported in column (8) of Panels A and B in Online Appendix Table C.10. The results suggest that the true effects are close to the estimated effect.[34] Specifically, the bias-adjusted coefficient for the purchase analysis is 0.0045 which is slightly above our estimated coefficient of 0.0043 in the full model. The bias adjusted coefficient for the sales analysis is -0.634 which is slightly closer to zero than our estimated coefficient of -0.670. For both panels, the results show that the bias-adjustments are small and do not materially affect our inferences.

Overall, the robustness tests in this section mitigate concerns about omitted variable bias. While some degree of confounding effects remains possible, it is unlikely that omitted variable bias fully explains our findings.

3.4 Measurement Error

Forecast Bias is measured with error; however, Section 2.2 shows that this error is limited. To further address measurement-error concerns, we conduct several robustness analyses. For each subject, *Forecast Bias* is estimated from a regression using 40 forecasts, to

[34] The model without controls includes only past performance and its interaction with *Forecast Bias*. The model with controls adds all standard controls as well as the interaction terms from Online Appendix Table C.10. The method of Oster (2019) requires two assumptions: (i) an upper bound on the achievable R^2 if all unobservables were included as controls and (ii) a ratio, δ, that compares the relative contribution of observable and unobservable factors in explaining outcome variation. Following the recommendation in Oster (2019), we set the upper bound on the achievable R^2 to 1.3 times the R^2 value of the full model and we set δ=1. We use within-R^2 to account for fixed effects, as it captures the variation explained by the included regressors while excluding that absorbed by the fixed effects. Relying on the overall R^2 would inflate the total explanatory power of the model, overstating the variation already accounted for, and in turn underestimating the potential bias from omitted variables.

obtain both the point estimate and its standard error. The standard error serves as a subject-level proxy for estimation noise. Online Appendix Table C.11 columns (1) in Panels A and B re-estimate the baseline purchase and sales specification after excluding subjects whose standard errors are in the top decile. The results remain robust.

We further address concerns about the possible attenuation of our coefficients due to measurement error in *Forecast Bias* using Obviously Related Instrumental Variables (ORIV) (Gillen, Snowberg, and Yariv, 2019). Each subject's forecasts are split into odd- and even-numbered rounds, and we estimate two proxies based on these non-overlapping samples. *Forecast Bias Odd* and *Forecast Bias Even* are highly correlated (see Online Appendix Figure C.1). ORIV then instruments the interaction *Forecast Bias Odd* × *Performance Measure* with *Forecast Bias Even* × *Performance Measure*, addressing attenuation bias arising from classical errors-in-variables.

Online Appendix Table C.11 columns (2) in Panels A and B show that the ORIV estimates are similar to our baseline estimates in column (2) Table 2 and column (2) Table 4. For purchases, the ORIV point estimate is essentially unchanged relative to the baseline specification. For sales, the ORIV point estimate is moderately larger than the baseline, consistent with some attenuation from measurement error in the original estimates. Overall, the ORIV evidence suggests that measurement error does not change the main findings and, to the extent it is present, it modestly biases coefficients toward zero.

3.5 Forecast Bias and Post-Experiment Trading

Our prior tests relate *Forecast Bias* to purchase and sales decisions made both before and after the experiment. This raises a potential concern of reverse causality; for example, if learning from pre-experiment trading influenced the subjects' responses in the elicitation procedure. As a robustness test, Table 6 reports results in which we estimate our main analyses on the subset of trades made after the lab experiments. Except for the change in sample, the regressions for purchases in column (1) and sales in column (2) are identical to the baseline specifications. In both columns, the results are similar to the main results.

4. *Forecast Bias*, Aggregate Stock Market Returns, and Net Flows into Stocks

The primary focus of this study is to test how forecast bias affects individuals' *cross-sectional* stock selection decision conditional on trading. In this section, we step back from stock selection and instead test how forecast bias relates to net flows to stocks. This aligns with much of the literature that examines surveys of investors' expectations, and focuses on the *time-series* of beliefs about stock market index returns (e.g., Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014; Adam, Matveev, and Nagel, 2021; Giglio, Maggiori, Stroebel, and Utkus, 2021). Most closely related to our study, Laudenbach, Weber, Weber, and Wohlfart (2024) use survey measures of investor-level beliefs about historical stock index autocorrelations to explain investors' net flows into stocks. We expand the extant literature by testing whether bias in information processing affects net flows to stocks. Specifically, we examine how forecast bias interacts with past market index returns and with each investor's own excess returns to affect net flows. Note that our previous results on buys and sells do not imply a mechanical relation between *Forecast Bias* and flows: the earlier results could arise from within-portfolio reallocations, which can occur independently from net flows. Therefore, we separately investigate the relation between *Forecast Bias* and net flows.

Column (1) of Table 7 reports regression results in which the dependent variable is net flows into stocks. The unit of observation is investor-month, and the sample includes all months in which an investor owns stocks, even if the investor does not trade (i.e., the sample does not condition on trading). The dependent variable, net flows, is defined as the difference between the value of stock purchases and sales in a month, divided by the value of stocks owned at the beginning of the month (multiplied by 100). We regress this variable on *Forecast Bias* interacted with lagged market and own-portfolio excess returns:

$$\begin{aligned} Net\ Flows_{i,t} = {} & \beta_1 \cdot MarketRet_t \cdot ForecastBias_i \\ & + \beta_2 \cdot ExcessRet_{i,t} \cdot ForecastBias_i \\ & + \beta_3 \cdot ExcessRet_{i,t} + \boldsymbol{\delta X_{i,t}} + \theta_i + \theta_t + \varepsilon_{i,t} \end{aligned} \tag{9}$$

where $MarketRet_t$ is the return on the Danish stock market index over the prior 12 months and $ExcessRet_{i,t}$ is return on investor *i*'s stock portfolio in excess of the market index over the prior 12 months.[35] The specification includes investor fixed effects, which subsume the direct effect of *Forecast Bias* and control for the investors' savings rates, general trading tendencies, etc. The control variables, $\boldsymbol{X_{i,t}}$, include the value of beginning of month stock holdings, financial assets, housing assets, income, education, children, and marital status. The specification also includes year-month fixed effects, which control for overall market returns, the state of the economy, etc.

The results in column (1) show a positive and significant coefficient on the interaction between *Forecast Bias* and lagged market index returns. Relative to contrarians, extrapolators' net flows increase when the market does well. The interaction between *Forecast Bias* and investors' own lagged excess returns on their stock portfolios is not significant. The results for these two coefficients show what type of returns interact with forecast bias – the returns of the underlying asset class are important but the investor's own excess returns are not. This provides evidence that forecast bias is distinct from overconfidence or biased self-attribution – individuals may increase allocations to stocks following high excess returns – but any effect from the investor's own excess returns does not interact with forecast bias.

[35] Because we require the past 12 months of the subjects' investment returns, these tests are estimated over the period 2012-2021, with the 2011 data used only to calculate the subjects' lagged investment returns.

Although the coefficient on the interaction term between *Forecast Bias* and lagged market index returns is significant, the implied economic magnitude is small. The coefficient estimate implies that, following a lagged market index return one standard deviation above the mean, a one standard deviation increase in *Forecast Bias* is associated with a net flow into stocks of 12 basis points. The small economic magnitude found in this unconditional regression is consistent with the literature. Giglio, Maggiori, Stroebel, and Utkus (2021) find that beliefs have little explanatory power for the timing of trades, but that, conditional on trade occurring, beliefs explain the direction of trade.[36] Accordingly, we separate the decision to trade from the action taken conditional upon trading.

Columns (2), (3), and (4) of Table 7 examine the decision of *when* to actively adjust the amount allocated to stocks, ignoring the size of the adjustment. In column (2), the dependent variable is *Trade Month*, an indicator equal to 100 for months when the absolute value of the investor's net flow is greater than 1%.[37] In column (3), the dependent variable is *Buy Month*, an indicator equal to 100 for months when the investor's net flow is greater than 1%. In column (4), the dependent variable is *Sell Month*, an indicator equal to 100 for months when the investor's net flow is less than -1%. None of the interaction terms in these columns are significant; *Forecast Bias* lacks the ability to predict when investors will trade.[38] This is similar to Giglio, Maggiori, Stroebel, and Utkus (2021), who find that changes in beliefs have little ability to predict when investors will trade.

Column (5) turns to the relation between net flows and the interaction between *Forecast Bias* and lagged returns *conditional* upon trading in that month. This regression is identical to

---

[36] Andries, Bianchi, Huynh, and Pouget (2025) show that signal precision affects forecast bias and the magnitude of the passthrough to investment decisions in an experimental asset market.

[37] We define *Trade month* as absolute net flows greater than 1% because portfolio rebalancing might lead to small absolute net flows if the values of sales and purchases are not identical. We do not want to classify portfolio rebalancing as an active decision about net flows.

[38] Online Appendix Table C.12 shows the results are similar for a conditional logit model.

that in column (1), except we restrict the sample to include only months in which the investor's absolute net flow is greater than 1%. The results are directionally similar to the unconditional results, but the implied economic magnitude is 7.2 times larger. The set of results in Table 7 shows that the relation between *Forecast Bias* and past returns is driven by the intensive margin of trading – actions taken conditional upon trading – and not by the decision to trade.

Taken together, our results on stock purchases, sales, and net flows provide evidence of a single underlying mechanism that affects the relation between past performance and investor trading decisions. Prior studies find that different types of salient performance measures affect different types of investment choices. The decision to purchase a stock is linked to that stock's historical returns (Grinblatt and Keloharju, 2000; Barber and Odean, 2008); the decision to sell a stock is linked to the investor's capital gains on that stock (Odean, 1998; Ben-David and Hirshleifer, 2012; Hartzmark, 2015); and decisions about net flows to stocks are linked with past market returns (Greenwood and Shleifer, 2014). Consistent with the literature, we find that lagged individual stock returns, capital gains, and overall market returns affect stock purchases, stock sales, and net flows into the market, respectively. Furthermore, our results show that forecast bias is a mechanism through which these different types of salient performance measures influence the corresponding investment decisions.

5. Behavioral Consistency Between Purchases and Sales

The previous results show that *Forecast Bias* affects both purchase and sale decisions. Frydman and Camerer (2016) posit that if a common mechanism drives both purchase and sale behaviors there will be within-investor consistency across both domains, and provide experimental evidence supporting this hypothesis. Following this idea, we test investor consistency in purchases and sales. We restrict the sample to investors with sufficient trading activity to reduce noise in these ratios, requiring at least two purchases, two sales, and ten total trades. For each investor, we compute the ratio of extrapolation purchases to total purchases and the ratio of extrapolation sales to total sales over the sample period. An extrapolation purchase is defined as a stock purchase following a positive excess return over the prior 12 months, and an extrapolation sale as a stock sale following a negative capital gain.

Figure 4 Panel A plots the extrapolation purchases ratio against the extrapolation sales ratio, showing 25 bins and the line of best fit. The figure shows a strong positive relation between buying and selling behaviors in our sample. Specifically, an increase of 0.2 in the extrapolation purchase ratio is associated with an increase of 0.027 in the extrapolation sale ratio, suggesting that extrapolation trading in one domain predicts extrapolation trading in the other.

A key feature of the behavioral consistency test is that it does not rely on our experimental measure, allowing us to examine trading consistency in the full universe of Danish investors. Figure 4 Panel B shows the relation between extrapolation purchases and extrapolation sales for the 316,770 Danish investors with sufficient trades during our sample period. The results are similar to those in our experimental sample, indicating significant behavioral consistency in the general population and supporting the external validity of our findings.

6. Conclusion

Our study is the first to show a relation between individual-level measures of forecast bias and cross-sectional stock trading decisions. We elicit a measure of forecast bias using a laboratory experiment for a sample of investors in Denmark. On average, individuals have positive forecast bias (extrapolators), though there is substantial variation: many individuals have negative forecast bias (contrarians) and many have near-zero bias. We link our measure of forecast bias to administrative register data on stock trades from 2011-2021 and examine how it affects individuals' stock selection decisions.

We find that forecast bias is positively related to the past excess returns of stocks purchased by individual investors: extrapolators tend to purchase stocks with higher past annual excess returns compared to contrarians. For sales decisions, forecast bias is negatively related to investors' capital gains since purchase of stocks they sell. Beyond security selection decisions, we find that investors with higher forecast bias increase their allocations to stocks following higher annual market returns relative to investors with lower forecast bias. In additional tests we find evidence of behavioral consistency between purchase and sale decisions in our sample as well as in the general population of investors: Individuals with a high ratio of extrapolator purchases also have a high ratio of extrapolator sales. Overall, our results show that heterogeneity in forecast bias – errors in how investors incorporate past performance into expectations – explains across-investor variation in how past performance affect investors' decisions about trading individual stocks and net flows to stocks.

Our results demonstrate widespread, yet heterogeneous, biases in investors' forecasts, which significantly affect their stock trading decisions. These findings have broader implications for other domains where forecasting is important, such as the effects of inflation expectations on consumption or interest rate expectations on mortgage choices. Heterogeneity in forecast bias can provide a potential rationale for the observed heterogeneity in beliefs and responses, particularly in domains where seemingly similar individuals form different beliefs and react differently to the same data.

References

Adam, K., D. Matveev, and S. Nagel. 2021. Do survey expectations of stock returns reflect risk adjustments? *Journal of Monetary Economics* 117, 723-740.

Adam, K., and S. Nagel. 2022. Expectations data in asset pricing. In *Handbook of Economic Expectations*, R. Bachman, G. Topa, and W. van der Klaauw Editors, Elsevier.

Afrouzi, H., S.Y. Kwon, A. Landier, Y. Ma, and D. Thesmar. 2023. Overreaction in expectations: Evidence and theory. *Quarterly Journal of Economics* 138, 1713-1764.

Amromin, G., and S.A. Sharpe. 2014. From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science* 60, 845-866.

Andersen, S., T. Hanspal, J. Martínez-Correa, and K.M. Nielsen. 2021. Beliefs and the disposition effect. Working paper, Copenhagen Business School.

Andries, M., M. Bianchi, K. Huynh, and S. Pouget. 2025. Return predictability, expectations, and investment: Experimental evidence. *Review of Financial Studies* 38, 1687-1729.

Atmaz, A., S. Cassella, H. Gulen, and F. Ruan. 2024. Contrarians, extrapolators, and stock market momentum and reversal. *Management Science* 70, 5949-5984.

Barber, B.M., and T. Odean. 2002. Online investors: Do the slow die first? *Review of Financial Studies* 15, 455-487.

Barber, B.M., and T. Odean. 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies* 21, 786-818.

Barberis, N.C. 2018. Psychology-based models of asset prices and trading volume. In *Handbook of Behavioral Economics*, D. Bernheim, S. DellaVigna, and D. Laibson Editors, Elsevier.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer. 2015. X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics* 115, 1-24.

Barberis, N., R. Greenwood, L. Jin, and A. Shleifer. 2018. Extrapolation and bubbles. *Journal of Financial Economics* 129, 203-227.

Barberis, N., and W. Xiong. 2012. Realization utility. *Journal of Financial Economics* 104, 251-271.

Barberis, N., A. Shleifer, and R. Vishny. 1998. A model of investor sentiment. *Journal of Financial Economics* 49, 307-343.

Ben-David, I., and D. Hirshleifer. 2012. Are investors really reluctant to realize their losses? Trading responses to past returns and the disposition effect. *Review of Financial Studies* 25, 2485-2532.

Beutel, J., and M. Weber. 2023. Beliefs and portfolios: Causal evidence. Working paper, University of Chicago.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer. 2019. Diagnostic expectations and stock returns. *Journal of Finance* 74, 2839-2874.

Bordalo, P., N. Gennaioli, and A. Shleifer. 2018. Diagnostic expectations and credit cycles. *Journal of Finance* 73, 199-227.

Cagan, P. 1956. The monetary dynamics of hyperinflation. In *Studies in the Quantity Theory of Money*, M. Friedman Editor, University of Chicago Press.

Calvet, L.E., J.Y. Campbell, and P. Sodini. 2007. Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy* 115, 707-747.

Calvet, L.E., J.Y. Campbell, and P. Sodini. 2009. Fight or flight? Portfolio rebalancing by individual investors. *Quarterly Journal of Economics* 124, 301-348.

Cassella, S., Z. Chen, H. Gulen, and R. Petkova. 2022. Extrapolators at the gate: Market-wide misvaluation and the value premium. Working paper, Tilburg University.

Cassella, S., and H. Gulen. 2018. Extrapolation bias and the predictability of stock returns by price-scaled variables. *Review of Financial Studies* 31, 4345-4397.

Charles, C., C. Frydman, and M. Kilic. 2024. Insensitive Investors. *Journal of Finance* 79, 2473-2503.

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

Das, S., C.M. Kuhnen, and S. Nagel. 2020. Socioeconomic Status and Macroeconomic Expectations. *Review of Financial Studies* 33, 395-432.

De Bondt, W.F. 1993. Betting on trends: Intuitive forecasts of financial risk and return. *International Journal of Forecasting* 9, 355-371.

De Bondt, W.F. and R. Thaler. 1985. Does the stock market overreact? *Journal of Finance* 40, 793-805.

Dimmock, S.G., R. Kouwenberg, O.S. Mitchell, and K. Peijnenburg. 2016. Ambiguity aversion and household portfolio choice puzzles: Empirical evidence. *Journal of Financial Economics* 119, 559-577.

Dimmock, S.G., R. Kouwenberg, O.S. Mitchell, and K. Peijnenburg. 2021. Household portfolio underdiversification and probability weighting: Evidence from the field. *Review of Financial Studies* 34, 4524-4563.

Dominitz, J., and C.F. Manski. 2007. Expected equity returns and portfolio choice: Evidence from the Health and Retirement Study. *Journal of the European Economic Association* 52, 369-379.

Dominitz, J., and C.F. Manski. 2011. Measuring and interpreting expectations of equity returns. *Journal of Applied Econometrics* 26, 352-370.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri. 2020. Heterogeneity and persistence in returns to wealth. *Econometrica* 88, 115-170.

Fisher, K.L., and M. Statman. 2000. Investor sentiment and stock returns. *Financial Analysts Journal* 56, 16-23.

Frydman, C., and C. Camerer. 2016. Neural evidence of regret and its implications for investor behavior. *Review of Financial Studies* 29, 3108-3139.

Frydman, C., and G. Nave. 2017. Extrapolative beliefs in perceptual and economic decisions: Evidence of a common mechanism. *Management Science* 63, 2340-2352.

Frydman, C., and A. Rangel. 2014. Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price. *Journal of Economic Behavior & Organization* 107, 541-552.

Frydman, C., and B. Wang. 2020. The impact of salience on investor behavior: Evidence from a natural experiment. *Journal of Finance* 75, 229-276.

Fuster, A., B. Hebert, and D. Laibson. 2011. Natural expectations, macroeconomic dynamics, and asset pricing. *NBER Macroeconomics Annual* 26, 1-48.

Fuster, A., D. Laibson, and B. Mendel. 2010. Natural expectations and macroeconomic fluctuations. *Journal of Economic Perspectives* 24, 67-84.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus. 2021. Five facts about beliefs and portfolios. *American Economic Review* 111, 1481-1522.

Gillen, B., E. Snowberg, and L. Yariv. 2019. Experimenting with measurement error: Techniques with applications to the Caltech cohort study. *Journal of Political Economy* 127, 1826-1863.

Greenwood, R., and A. Shleifer. 2014. Expectations of returns and expected returns. *Review of Financial Studies* 27, 714-746.

Grinblatt, M., and M. Keloharju. 2000. The investment behavior and performance of various investor types: A study of Finland's unique data set. *Journal of Financial Economics* 55, 43-67.

Grinblatt, M., and M. Keloharju. 2001. What makes investors trade? *Journal of Finance* 66, 589-616.

Hartzmark, S.M. 2015. The worst, the best, ignoring all the rest: The rank effect and trading behavior. *Review of Financial Studies* 28, 1024-1059.

Heiss, F., M. Hurd, M. van Rooij, T. Rossman, and J. Winter. 2022. Dynamics and heterogeneity of subjective stock market expectations. *Journal of Econometrics* 231, 213-231.

Hoffman, A.O.I., T. Post, and J.M.E. Pennings. 2015. How investor perceptions drive actual trading and risk-taking behavior. *Journal of Behavioral Finance* 16, 94-103.

Hong, H., and J.C. Stein. 1999. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance* 54, 2143-2184.

Hossain, T., and R. Okui. 2013. The binarized scoring rule. *Review of Economic Studies* 80, 984-1001.

Hvide, H., and P. Östberg. 2015. Social interaction at work. *Journal of Financial Economics* 117, 628-652.

Jegadeesh, N., and S. Titman. 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance* 48, 65-91.

Jin, L.J., and C. Peng. 2024. Law of small numbers in financial markets: Theory and evidence. Working paper, Cornell University.

Kaustia, M., S. Knüpfer. 2008. Do investors overweight personal experience? Evidence from IPO subscriptions. *Journal of Finance* 63, 2679-2702.

Knüpfer, S., E.H. Rantapuska, and M. Sarvimäki. 2017. Formative experiences and portfolio choice: Evidence from the Finnish Great Depression. *Journal of Finance* 72, 133-166.

Knüpfer, S., E.H. Rantapuska, and M. Sarvimäki. 2023. Social interaction in the family: Evidence from investors' security holdings. *Review of Finance* 27, 1297-1327.

Kogan, S., I. Makarov, M. Niessner, and A. Schoar. 2024. Are cryptos different? Evidence from retail trading. *Journal of Financial Economics* 159, 103897.

Lakonishok, J., A. Shleifer, and R. Vishny. 1994. Contrarian investment, extrapolation, and risk. *Journal of Finance* 49, 1541-1578.

Landier, A., Y. Ma, and D. Thesmar. 2019. Biases in expectations: Experimental evidence. Working paper, HEC Paris.

Laudenbach, C., A. Weber, R. Weber, and J. Wohlfart. 2024. Beliefs about the stock market and investment choices: Evidence from a field experiment. *Review of Financial Studies*, forthcoming.

Liu, H., C. Peng, W.A. Xiong, and W. Xiong. 2022. Taming the bias zoo. *Journal of Financial Economics* 143, 716-741.

Malmendier, U., and S. Nagel. 2011. Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics* 126, 373-416.

Merkle, C., and M. Weber. 2014. Do investors put their money where their mouth is? Stock market expectations and investing behavior? *Journal of Banking and Finance* 46, 372-386.

Metzler, L.A. 1941. The nature and stability of inventory cycles. *Review of Economics and Statistics* 23, 113-129.

Nofsinger, J.R., and R.W. Sias. 1999. Herding and feedback trading by institutional and individual investors. *Journal of Finance* 54, 2263-2295.

Odean, T. 1998. Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775-1798.

Oster, E. 2019. Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics* 37, 187-204.

Solon, G., S.J. Haider, and J.M. Wooldridge. 2015. What are we weighting for? *Journal of Human Resources* 50, 301-316.

Vissing-Jorgensen, A. 2003. Perspectives on behavioral finance: Does "irrationality" disappear with wealth? Evidence from expectations and actions. *NBER Macroeconomics Annual* 18, 139-194.

von Gaudecker, H-M., and A. Wogrolly. 2022. Heterogeneity in households' stock market beliefs: Levels, dynamics, and epistemic uncertainty. *Journal of Econometrics* 231, 232-247.

Woodford, M. 2003. Imperfect common knowledge and the effects of monetary policy. In *Knowledge, Information, and Expectations in Modern Macroeconomics*, P. Aghion, R. Frydman, J.E. Stiglitz, and M. Woodford Editors, Princeton University Press.

Figure 1: Elicitation of forecast bias

This figure shows an example of the forecasting task. The upper panel shows an example of the first round of the prediction task. The subject observes 40 past realizations of the process (green dots with numbers showing exact values). The subject is asked to make forecasts for the next two rounds by sliding the blue and orange "**x**" up and down, and then clicking the "Make forecast" button. The next realization of the process is then revealed, as seen in the lower panel, and the subject is asked to make two new predictions. This process continues for a total of 40 rounds.

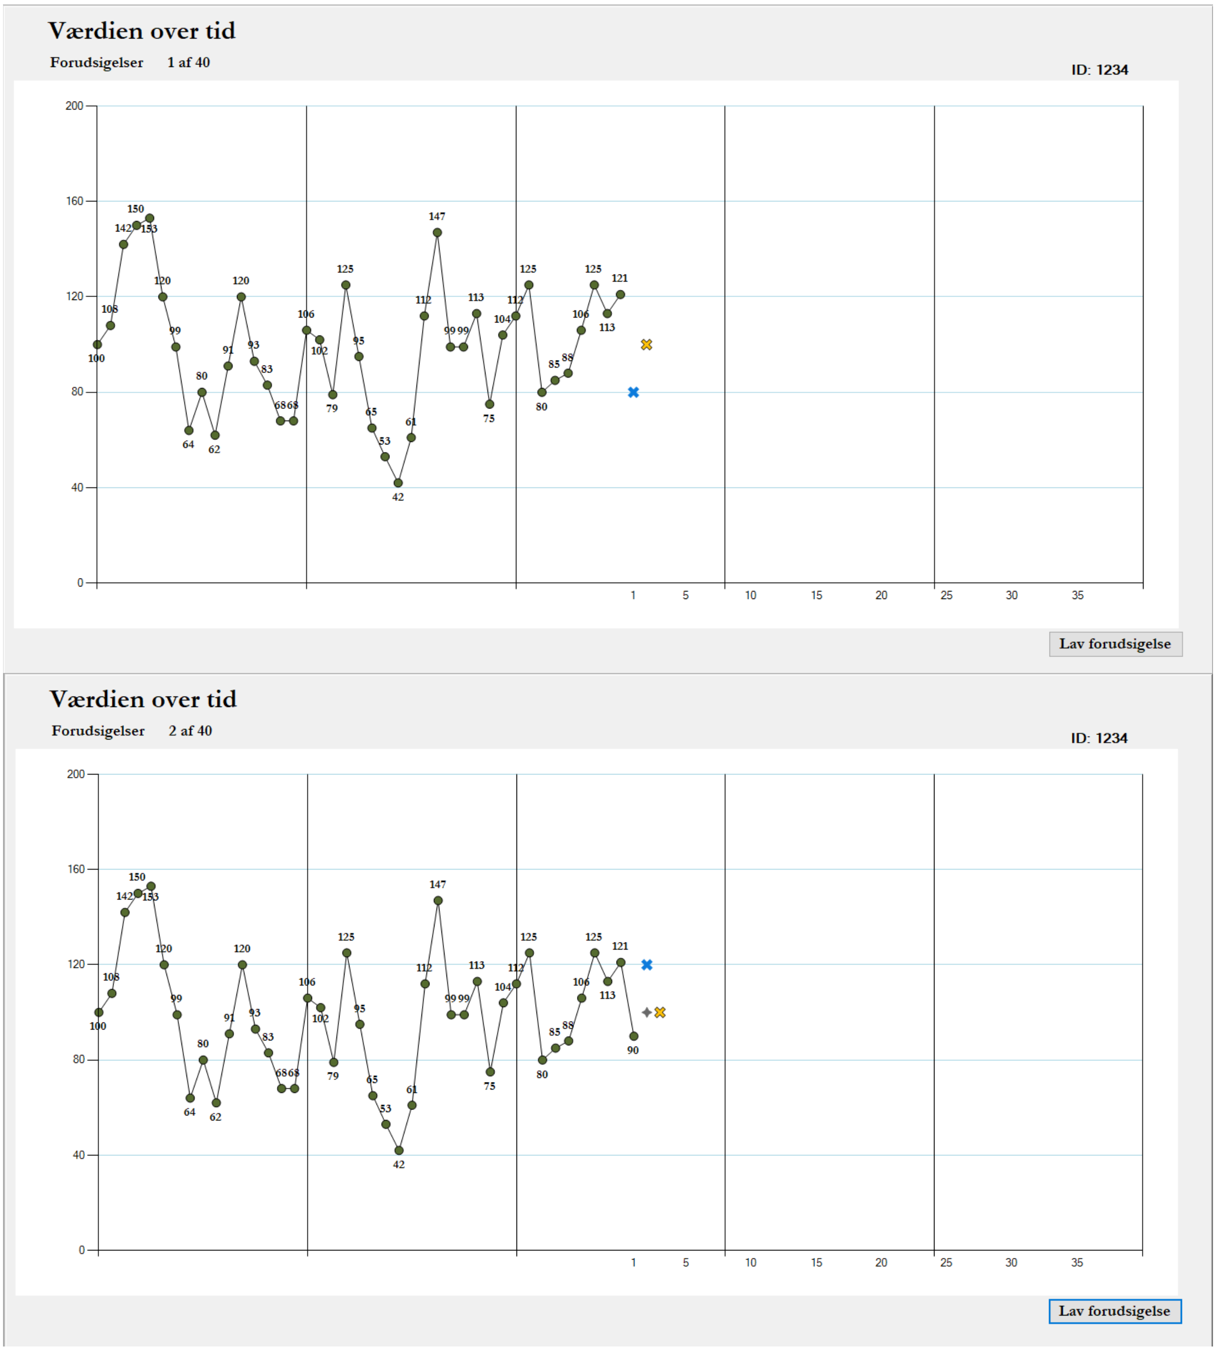

Figure 2: Histogram of *Forecast Bias*

This figure shows a histogram of the distribution of *Forecast Bias*. A value of zero implies no bias, a value greater than zero implies extrapolation bias (i.e., forecast is biased in the direction of recent realizations), and a value below zero implies contrarian bias (i.e., forecast is biased in the opposite direction of recent realizations). We truncate the tails to avoid reporting bins with fewer than three observations, in accordance with our data agreement with Statistics Denmark. The mean is 0.14, the standard deviation is standardized to 1, and the quasi-median is 0.28 ("quasi-median" is the average value for subjects between the 45th and 55th percentiles, and is used to avoid reporting non-aggregated values).

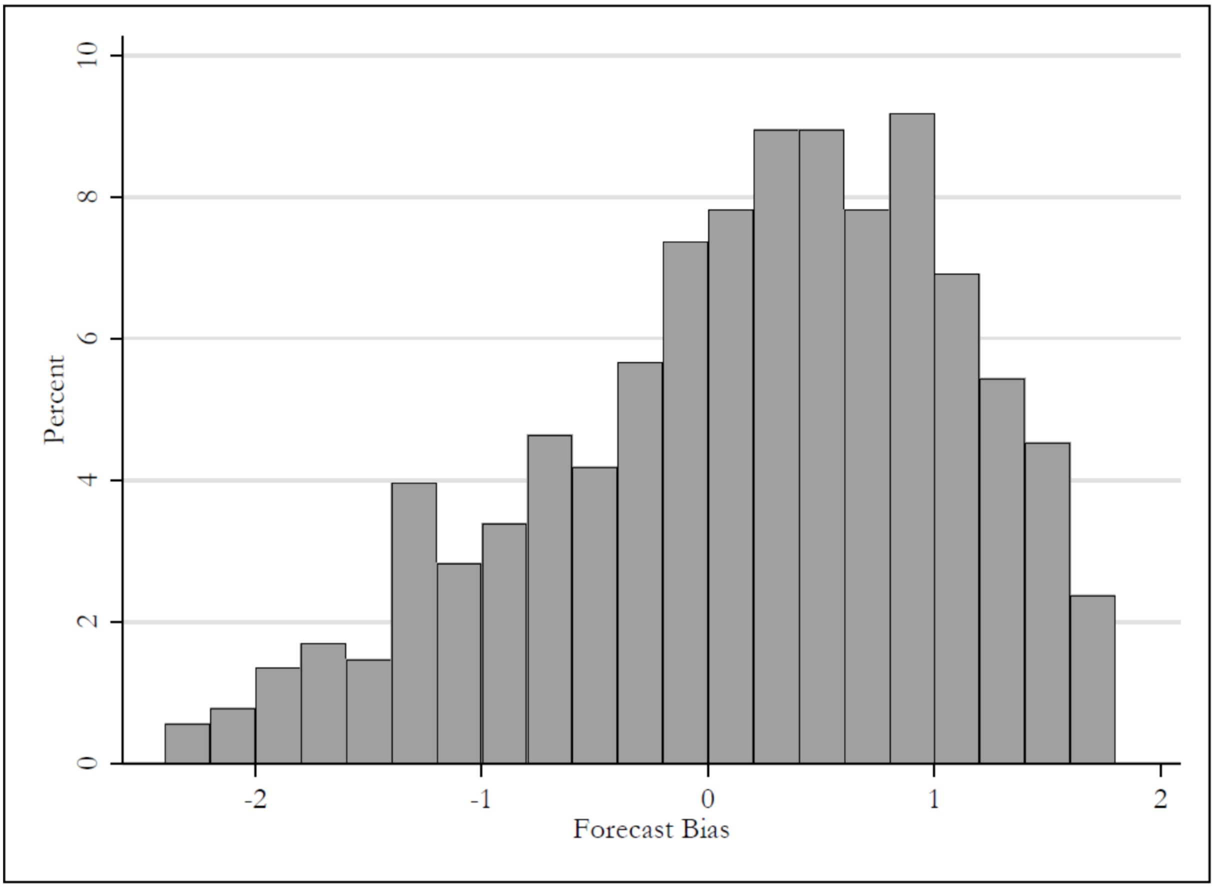

Figure 3: Relation between *Forecast Bias* and performance of traded stocks

This figure shows results of regressions of the relation between *Forecast Bias* and stock purchases and sales. *Forecast Bias* is partitioned into five bins: $Forecast\ Bias < -1$; $-1 \leq Forecast\ Bias < -0.25$; $-0.25 \leq Forecast\ Bias < 0.25$; $0.25 \leq Forecast\ Bias < 1$; and $Forecast\ Bias \geq 1$. An indicator variable for each bin of *Forecast Bias* is interacted with the performance measure. The key independent variables are $\mathbb{I}$(*Forecast Bias bin*) *x Performance measure*. Panel A reports the coefficients of a weighted-least squares regression in which the dependent variable equals 100 if the stock is purchased and zero otherwise. The performance measure is *Lagged annual return*. The return is winsorized at the 1st and 99th percentiles. The unit of observation is investor-stock-day over the period 2011-2021. Panel B reports the coefficients of a weighted-least squares regression in which the dependent variable equals 100 if the stock is sold and zero otherwise. The performance measure is *Capital gain*. The capital gain is winsorized at the 1st and 99th percentiles. The unit of observation is investor-stock-day over the period 2013-2021. Panel A includes an indicator for a stock previously held by the subject, an indicator for a stock currently owned by the subject, the portfolio weight of a current holding (zero for stocks not currently held), and the fraction of all investor-stock purchases for the full Danish population in the past month that were in the stock. Panel B includes monthly holding length fixed effects, as well as controls for the portfolio weight of a current holding, and the fraction of all investor-stock sales for the full Danish population in the past month that were in the stock. In both panels, the observations are weighted such that each investor has equal weight in the regressions. *Forecast Bias* is adjusted to have a standard deviation of one. Standard errors are clustered at the individual-level. The dots indicate the coefficient point estimates, and the vertical bars show the 95% confidence intervals.

Panel A: Purchases

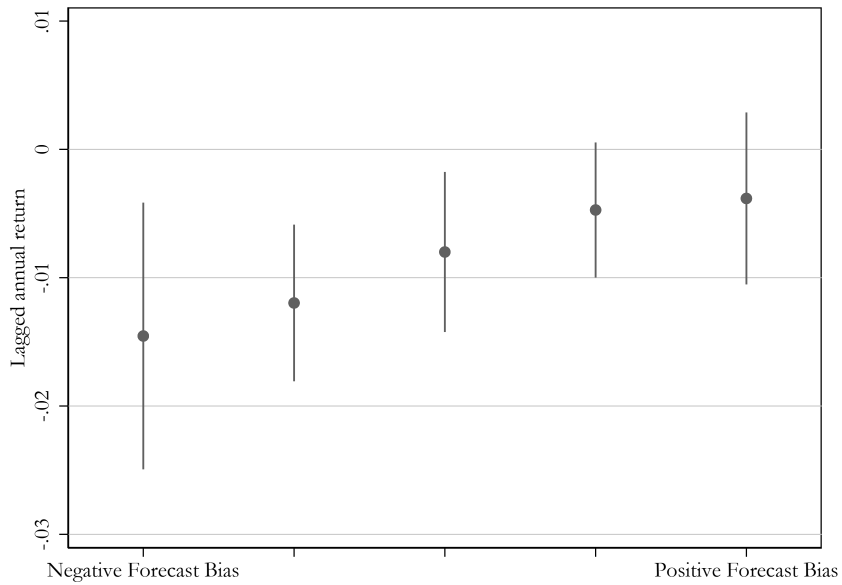

Panel B: Sales

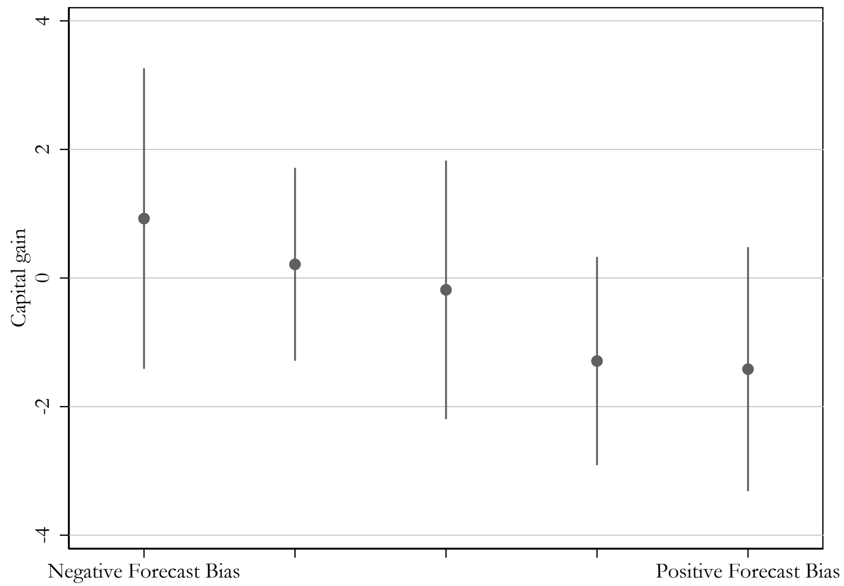

Figure 4: Trading based on past performance for stock purchases and sales

This figure shows results of regressions of the relation between the ratio of extrapolation purchases and the ratio of extrapolation sales. An extrapolation purchase is defined as a stock purchase following a positive excess return over the prior 12 months, and an extrapolation sale as a stock sale following a negative capital gain. The ratio of extrapolation purchases (sales) is the number of extrapolation purchases (sales) divided by total purchases (sales) calculated over the period 2013 to 2021. The sample is restricted to investors with at least two purchases, two sales, and ten total trades. Panel A includes all subjects in our experiment conditional on meeting the restriction. Panel B includes all Danish investors conditional on meeting the restrictions. The number of observations in Panels A and B are 417 and 316,770, respectively. The solid line is the linear regression line, and the dots present 25 bins of observations.

Panel A: Experiment sample

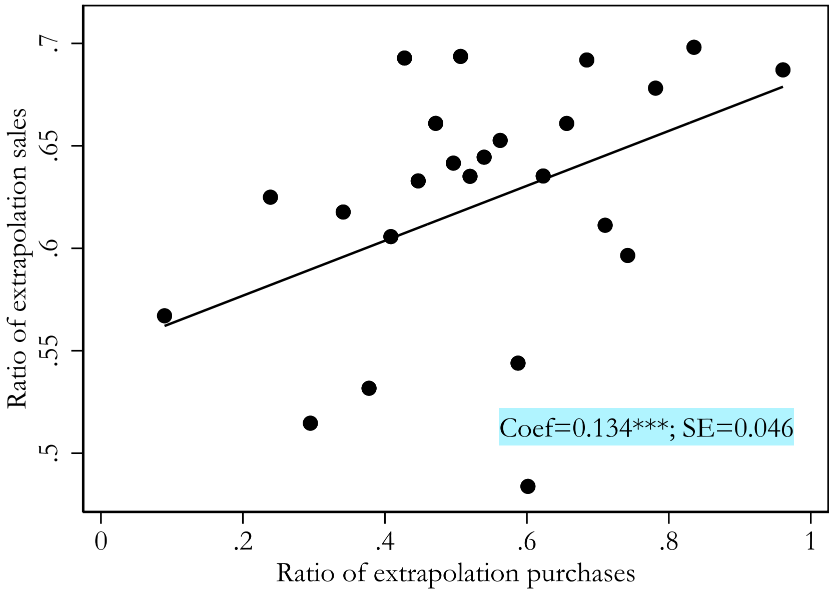

Panel B: All Danish investors

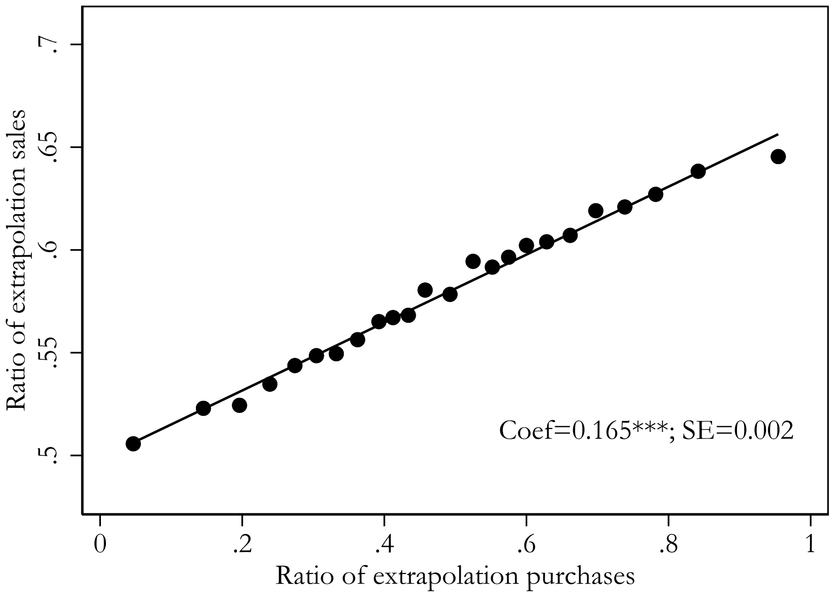

Table 1: Summary statistics

This table reports summary statistics. Appendix Table A1 defines variables. For each variable, we report the mean and standard deviation. The column "quasi-median" reports the average value of the variable for subjects between the 45th and 55th percentile (this is done because our data agreement prohibits reporting non-aggregated values). Panels A and B report summary statistics for the trading and stock characteristics, and individual characteristics, respectively. The summary statistics for the individual characteristics are for the year the experiment is conducted, 2020.

Panel A: Trading and stock characteristics

	Mean	Std. Dev.	Quasi-median
Number of buys	43.57	99.85	14.18
Value of buy	59,299	134,950	25,318
Lagged annual ret. (buys)	26.20%	84.57%	8.86%
Lagged annual excess ret. (buys)	6.41%	82.85%	-9.52%
Number of sales	33.10	75.92	9.94
Value of sale	82,819	624,121	32,815
Prior annual ret. (sales)	29.27%	78.83%	13.47%
Capital gain	23.09%	99.46%	3.37%
Held before (%)	1.81%	13.34%	0%
Current holding (%)	0.21%	4.53%	0%
Portfolio weight (buys)	0.02%	0.90%	0%
Portfolio weight (sales)	8.67%	13.32%	4.14%
Stock purchase fraction (%)	0.06%	0.41%	0.00%
Stock sales fraction (%)	1.07%	1.83%	0.27%
Net flows	0.28	10.63	0
Trade month	13.82%	34.51%	0%
Buy month	7.73%	26.70%	0%
Sell month	6.09%	23.91%	0%
Conditional net flows	2.15	28.37	2.14

Panel B: Individual characteristics

	Mean	Std. dev.	Quasi-median
Age	50.56	7.87	52.52
Male	0.69	0.46	1
Married	0.64	0.48	1
Children	0.81	0.39	1
Education	16.45	2.20	16.98
Financial assets (000's)	2,432.81	16,112.57	855.07
Income (000's)	771.00	635.61	646.26
Housing assets (000's)	1,920.04	1,916.59	1633.44
Post-Covid experiment	0.34	0.47	0
Risk aversion	0.49	0.16	0.49
Financial literacy	3.40	0.80	4
Numeracy	2.83	0.43	3
Optimism	4.59	7.96	5
Overconfidence	0.19	0.91	0
Trust	4.23	1.55	5

Table 2: Forecast bias and stock purchases

This table reports results of regressions of the relation between stock purchases and *Forecast Bias*. Column (1) reports the coefficients of a weighted-least squares regression in which the dependent variable equals 100 if the stock is purchased and zero otherwise, and includes an investor-day fixed effect. Column (2) reports the coefficients of a conditional logit regression in which the dependent variable equals 1 if the stock is purchased and zero otherwise, and conditions out investor-day effects. The key independent variable is *Forecast Bias x Lagged annual return*. The return is winsorized at the 1st and 99th percentiles. The unit of observation is investor-stock-day over the period 2011-2021. In both columns, the observations are weighted such that each investor has equal weight in the regressions. *Forecast Bias* is adjusted to have a standard deviation of one. Both columns include an indicator for a stock previously held by the subject, an indicator for a stock currently owned by the subject, the portfolio weight of a current holding (zero for stocks not currently held), and the fraction of all investor-stock purchases for the full Danish population in the past month that were in the stock. Standard errors are clustered at the individual-level and appear in parentheses. The symbols *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Weighted-least squares	Conditional logit
	(1)	(2)
Forecast Bias × Lagged annual return	0.004***	0.061**
	(0.001)	(0.026)
Lagged annual return	-0.008***	-0.019
	(0.002)	(0.026)
Investor-Day Fixed Effect	Yes	Yes
Control variables	Yes	Yes
N	33,675,400	33,675,400

Table 3: Past performance and stock purchases for extrapolators and contrarians

This table reports results of weighted-least squares regressions of the relation between lagged annual return and purchases. The dependent variable equals 100 if the stock is purchased and zero otherwise, and includes an investor-day fixed effect. Column (1) presents results for the full sample, column (2) presents results for investors with positive *Forecast Bias,* and column (3) presents results for investors with negative *Forecast Bias*. Returns are winsorized at the 1st and 99th percentiles. The unit of observation is investor-stock-day over the period 2011-2021. The observations are weighted such that each investor has equal weight in the regressions. All columns include an indicator for a stock previously held by the subject, an indicator for a stock currently owned by the subject, the portfolio weight of a current holding (zero for stocks not currently held), and the fraction of all investor-stock purchases for the full Danish population in the past month that were in the stock. Standard errors are clustered at the individual-level and appear in parentheses. The symbols *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Full sample	Extrapolators	Contrarians
	(1)	(2)	(3)
Lagged annual return	-0.008***	-0.004**	-0.013***
	(0.002)	(0.002)	(0.003)
Investor-Day Fixed Effect	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
N	33,675,400	19,545,733	14,129,667

Table 4: Forecast bias and stock sales

This table reports results of regressions of the relation between stock sales and *Forecast Bias*. Columns (1) and (2) report the coefficients of a weighted-least squares regression in which the dependent variable equals 100 if the stock is sold and zero otherwise, and includes an investor-day fixed effect. Column (3) reports the coefficients of a conditional logit regression in which the dependent variable equals 1 if the stock is sold and zero otherwise, and conditions out investor-day effects. The key independent variables are *Forecast Bias x Performance measure*, where the performance measure is *Lagged annual return* in column (1) and *Capital gain* in columns (2) and (3). Returns and capital gains are winsorized at the 1st and 99th percentiles. The unit of observation is investor-stock-day over the period 2013-2021. In both columns, the observations are weighted such that each investor has equal weight in the regressions. *Forecast Bias* is adjusted to have a standard deviation of one. All columns include monthly holding length fixed effects, as well as controls for the portfolio weight of a current holding, and the fraction of all investor-stock sales for the full Danish population in the past month that were in the stock. Standard errors are clustered at the individual-level and appear in parentheses. The symbols *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Weighted-least squares		Conditional logit
Performance measure	Lag. annual return	Capital gain	Capital gain
	(1)	(2)	(3)
Forecast Bias × Perform.	-0.432	-0.893**	-0.063**
	(0.695)	(0.401)	(0.027)
Performance measure	3.759***	-0.427	-0.012
	(0.885)	(0.448)	(0.026)
Investor-Day Fixed Effect	Yes	Yes	Yes
Length of Holding Fixed Effect	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
N	162,508	176,451	165,947

Table 5: Past performance and stock sales for extrapolators and contrarians

This table reports results of weighted-least squares regressions of the relation between capital gain and sales. The dependent variable equals 100 if the stock is sold and zero otherwise, and includes an investor-day fixed effect. Column (1) presents results for the full sample, column (2) presents results for investors with positive *Forecast Bias,* and column (3) presents results for investors with negative *Forecast Bias*. Capital gains are winsorized at the 1st and 99th percentiles. The unit of observation is investor-stock-day over the period 2013-2021. The observations are weighted such that each investor has equal weight in the regressions. All columns include monthly holding length fixed effects, as well as controls for the portfolio weight of a current holding, and the fraction of all investor-stock sales for the full Danish population in the past month that were in the stock. Standard errors are clustered at the individual-level and appear in parentheses. The symbols *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Full sample	Extrapolators	Contrarians
	(1)	(2)	(3)
Capital gain	-0.579	-1.445**	0.985
	(0.458)	(0.597)	(0.618)
Investor-Day Fixed Effect	Yes	Yes	Yes
Length of Holding Fixed Effect	Yes	Yes	Yes
Control variables	Yes	Yes	Yes
N	176,451	120,766	55,685

Table 6: Forecast bias and trading post-experiment

This table reports the coefficients of weighted-least squares regressions using only trading observations post-experiment for purchases (column 1) and sales (column 2). The specification for purchases in column (1) is the same as in column (1) of Table 2. The specification for sales in column (2) is the same as in column (2) of Table 4. The performance measure for purchases (column 1) is lagged annual return and the performance measure for sales (column 2) is capital gain. In both columns, the sample includes only trade-days occurring after the subject participated in the laboratory experiment. Standard errors are clustered at the individual-level and appear in parentheses. The symbols *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Purchases	Sales
	(1)	(2)
Forecast Bias × Performance	0.003**	-1.046**
	(0.001)	(0.482)
Performance measure	-0.007***	-0.396
	(0.002)	(0.587)
Investor-Day Fixed Effect	Yes	Yes
Length of Holding Fixed Effect	No	Yes
Control variables	Yes	Yes
N	12,416,384	72,771

Table 7: Net flows, trading, and conditional net flows

This table reports the results of weighted-least squares regressions of monthly net flows, trading activity, and conditional net flows during the period 2012-2021. In columns (1) and (5), the dependent variable is net flows, which is defined as the value of purchases less the value of sales divided by beginning of month portfolio value, and is winsorized at the 1st and 99th percentiles. In columns (2), (3), and (4), the dependent variables are indicators equal to 100 if, respectively, the absolute value of the subject's net flows is greater than 1%, the value of net flows is greater than 1%, and the value of net flows is less than -1%. In columns (1) through (4), the sample includes all investor-months in which the subject owns individual stocks. In column (5), the sample includes only investor-months in which the absolute value of the subject's net flow is greater than 1%. In all columns, the observations are weighted such that each investor has equal weight in the regressions. Lagged annual market return is the return on the Danish stock market index over the prior year. Lagged annual excess return is the subject's stock return over the prior year less the lag market return. All columns include controls for the value of beginning of month stock holdings, financial assets, housing assets, income, education, children, and marital status, as well as individual and year-month fixed effects. Appendix Table A1 defines the variables. Standard errors are clustered at the individual-level and appear in parentheses. The symbols *, **, and *** denote significance at the 10%, 5%, and 1% level, respectively.

	Net flows	Trade month	Buy month	Sell month	Conditional net flows
	(1)	(2)	(3)	(4)	(5)
Forecast Bias × Lag market return	0.854**	-0.739	-0.341	-0.398	6.121***
	(0.346)	(1.057)	(0.782)	(0.705)	(2.151)
Forecast Bias × Lag excess return	0.099	0.547	0.173	0.374	0.457
	(0.194)	(0.594)	(0.432)	(0.409)	(0.900)
Lag excess return	0.080	3.035***	1.656***	1.379***	-0.955
	(0.237)	(0.696)	(0.544)	(0.455)	(1.124)
Investor Fixed Effect	Yes	Yes	Yes	Yes	Yes
Year-Month Fixed Effect	Yes	Yes	Yes	Yes	Yes
Control variables	Yes	Yes	Yes	Yes	Yes
N	87,140	87,140	87,140	87,140	11,947

Appendix Table A1: Variable definitions

Variable name	Definition
Forecast Bias	The forecast bias parameter estimated as in equation (2). For ease of interpretation, we divide the parameter estimate by its standard deviation.
Lagged annual return	The return on the purchased stock over the prior year ending the day before purchase
Lagged annual market return	The return on the value-weighted Danish stock market over the prior year ending the day before purchase
Lagged annual excess return	The difference between the lagged annual return of the stock and the lagged annual market return
Capital gain	The percentage change in the value of the position relative to its purchase price
Held before	Indicator if the subject held the stock previously
Current holding	Indicator if the subject holds the stock currently
Portfolio weight	A stock's weight in the portfolio
Stocks purchase fraction	Prior month purchases of a stock as a fraction of prior month total stock sales in Denmark (multiplied by 100)
Stocks sales fraction	Prior month sales of stock as a fraction of prior month total stock sales in Denmark (multiplied by 100)
Net flows	Value of stock purchases in a month less the value of stock sales divided by the beginning-of-month value of stocks owned (multiplied by 100)
Conditional net flows	Net flows conditional on trading that month
Trade month	Indicator variable equal to 100 in months in which the absolute value of the investor’s net flow into stocks is greater than 1% of the beginning-of-month portfolio value
Buy month	Indicator variable equal to 100 in months in which the value of the investor’s net flow into stocks is greater than 1% of the beginning-of-month portfolio value
Sell month	Indicator variable equal to 100 in months in which the value of the investor’s net flow into stocks is less than -1% of the beginning-of-month portfolio value
Income	The sum of labor income, social transfers, pension income, income from investments, and other personal income, reported in Danish kroner (DKK)
Financial assets	The sum of stocks, bonds, and bank deposits (DKK)
Housing assets	The value of the subjects’ home (DKK)
Age	Age in years
Education	Years of formal education
Male	Indicator for male
Married	Indicator if subject is currently married
Children	Indicator for whether the subject has children
Risk aversion	Fraction of paired lottery choice questions for which the subject chose the safer option (details in Online Appendix B)

Financial literacy	Number of the four financial literacy questions answered correctly (details in Online Appendix B)
Numeracy	Number of the three numeracy questions answered correctly
Optimism	Subjects' stated life expectancy less objective life expectancy from actuarial tables (details in Online Appendix B)
Overconfidence	The sum of financial literacy and numeracy questions the subject believes they answered correctly less the number they actually answered correctly (details in Online Appendix B)
Trust	Likert scale where zero indicates "Most people can be trusted" and six indicate "One has to be very careful with other people" (details in Online Appendix B)
Post-Covid experiment	Indicator for subjects whose experimental session was in November 2020