Discussion Paper Series – CRC TR 224

Discussion Paper No. 684
Project C 01

Experience Effects on Wall Street vs. Main Street: Field and Lab Evidence of Context Dependence

Benjamin Christoffersen[1]
Arvid Hoffmann[2]
Zwetelina Iliewa[3]
Lena Jaroszek[4]

April 2025

[1]Churney APS, Email : boennecd@gmail.com
[2]University of Adelaide, Email : arvid.homann@adelaide.edu.au
[3]University of Bonn & CESifo, Email : zwetelina.iliewa@gmail.com
[4]Copenhagen Business School, Email : lja.@cbs.dk

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Experience Effects on Wall Street vs. Main Street: Field and Lab Evidence of Context Dependence

BENJAMIN CHRISTOFFERSEN, ARVID HOFFMANN, ZWETELINA ILIEWA, LENA JAROSZEK*

This draft: April 15, 2025

Abstract

We examine how and why context influences experiential learning, comparing professional- and private-context stock market experiences. We find opposing patterns: In professional contexts, experiential learning exhibits a *primacy bias*, where sticky early experiences cause an underreaction to subsequent experiences. In contrast, in private contexts, a *recency bias* causes beliefs to fluctuate excessively over time. To identify the causal effect of context, we leverage (i) panel data on the dynamics of context-related experiences and expectations of finance professionals and (ii) experimental data on investment choices. Our experimental design allows us to identify the cognitive mechanisms underlying the documented context dependence of experience effects.

JEL classification: D83, D84, G02, G17.

---

* Benjamin Christoffersen: Churney APS (boennecd@gmail.com). Arvid O. I. Hoffmann: Adelaide Business School, University of Adelaide (arvid.hoffmann@adelaide.edu.au). Zwetelina Iliewa: University of Bonn, CESifo (zwetelina.iliewa@gmail.com). Lena M. Jaroszek: Copenhagen Business School (lja.fi@cbs.dk). We would like to thank Tobias Berg, Peter Bossaerts, Alexander Hillert, Markku Kaustia, Theresa Kuchler, Camelia Kuhnen, Stefan Nagel, Alexandra Niessen-Ruenzi, Stefan Ruenzi, Paul Smeets, Michael Ungeheuer, Martin Weber, Michael Weber as well as seminar and conference participants at the University of Aarhus, University of Bonn, University of Maastricht, University of Mannheim, University of Melbourne, Research in Behavioral Finance Conference, CESifo Venice Summer Institute on Expectation Formation, and the Women in Finance Workshop for their helpful suggestions. We are grateful to the Centre for European Economic Research (ZEW) for providing the field data on finance professionals. Zwetelina Iliewa would like to acknowledge that this project was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2126/1- 390838866 and through CRC / TR 224 (Project C01). All remaining errors are our own. Approval for the experimental studies in this paper was obtained from the Institutional Review Board (IRB) at the German Association for Experimental Economic Research under project No. JmqZYPwU.

People often acquire new information through a sequence of experiences. The way in which they process these experiences may vary across different contexts. For example, a person going for a walk may appreciate that the weather is good or that it has improved compared to the previous day. In contrast, a meteorologist not only appreciates the positive short-term trend but also evaluates whether or not their forecast was accurate. Similarly, a stock market analyst may become concerned if market movements deviate from their forecasts for too many days in a row, unlike a private investor who does not generate formal predictions and therefore may not experience this sense of error as long as the market movements are profitable. In professional contexts, people generally engage in additional activities that are absent in private contexts. How does this set of professional activities change the way people *process* experience sequences and, consequently, learn from them?

We study how and why the context — private versus professional — affects experiential learning. Using both field evidence from finance professionals' stock market forecasts and experimental data on investment choices, we document a robust yet contrasting pattern. While experiential learning occurs in both private and professional contexts, it exhibits a *primacy bias* — a persistent tendency to overreact to early-career experiences and underreact to subsequent experiences — only in *professional* settings. Among the various professional activities that could lead to primacy bias, we identify *forecasting* as the most relevant — where people's experiences are accompanied by simultaneously making explicit predictions and receiving timely feedback.[1] We experimentally identify the differences in the dynamic processing of experience sequences, with and without simultaneous forecasting, that drive the context dependence.

The occurrence of experiential learning has been demonstrated repeatedly before and it deviates from neoclassical models of rational expectations (for a summary see Malmendier, 2021; Malmendier and Wachter, 2024). Positive theoretical models, which we discuss in detail later, incorporate various cognitive mechanisms to predict both recency bias and primacy bias, both under realistic assumptions. However, field evidence in the contexts that have previously been studied points overwhelmingly to recency bias, while evidence of primacy bias in the field is scarce. This gap may

[1] This is distinct from an unconditional recency bias in extrapolative behavior, independent of one's individual window of experience (see, e.g., Greenwood and Shleifer, 2014). Any form of experiential learning, including those with a primacy bias, inherently leads to an unconditional recency bias. This is because the overweighted window of experience always includes the most recent subset of the data sequence, regardless of how individual data points within that subset are weighted relative to each other.

be due to the excessive emphasis on private-context over professional-context experiences in prior research.[2] Examining the conditions for the occurrence of primacy bias is important because of its implications. While a recency bias causes beliefs and choices to fluctuate excessively over time, a primacy bias implies a slowly-resolving overreaction to early-career experiences causing learning of the true parameters to be very slow. For example, the primacy-bias estimates in our field setting imply that the first-year career experience alone receives 67% weight after two years of tenure, 50% weight after four years, 28% weight after 21 years (which is the median tenure in our field setting) and 20% weight after 55 years. As a result, beliefs and choices may be systematically biased for a long time depending on how extreme and unlikely these first-year career experiences are. For example, based on our primacy-bias estimates in the field, finance professionals who experience a stock market return above 20% or below -20% in their first year — which is the case for about half of the professionals in our setting — are unlikely to fully overcome the resulting systematic optimism (or pessimism) bias before retirement.

An ideal experiment would isolate the causal impact of context on experiential learning by randomly assigning individuals to either a private or professional context while they undergo the same sequence of experiences. Ideally, this would involve randomly assigning individuals to work in or outside the finance industry or randomly instructing them to perform relevant professional activities, assuming these activities were well-defined. In this study, we approximate this ideal design in two ways. First, in our *field setting* we use a novel dataset of finance professionals that allows us to exploit within-subject variation of the context. Second, in our *lab setting* we employ a between-subject design to randomly assign subjects to various structured representations of potential private-context and professional-context activities.

We begin by analyzing context-dependent experiential learning using a field dataset of 330 finance professionals surveyed over 178 months. Importantly, we observe how the same subject experiences a data sequence of the same underlying stock market index in either their private or professional context. To measure experienced returns in the private and professional context, we use

[2] Prior research has demonstrated the occurrence of experience effects among professionals from their private-context experiences (Malmendier et al., 2021, 2011) and professional experiences (Schoar and Zuo, 2017; Dittmar and Duchin, 2016). However, previous studies do not provide any estimates of tenure-dependent differences in the weighting of past experiences in the professional context.

data on subjects' birth year and career start in the finance industry. Importantly, experiences in the private and professional context are empirically distinct due to the large variation in career start ages among the professionals in our sample. Furthermore, the longitudinal panel structure of our dataset allows us to identify the type of experiential learning — the relative over- or underweighting of various episodes of the experienced sequences — from the within-subject dynamics of expectations over time. Following Malmendier and Nagel (2011), we calculate the weighted average stock market return that each professional has experienced since their career start (i.e. professional-context experience) and throughout their lifetime (i.e. private-context experience).

Our findings reveal strong context dependence in experiential learning. Finance professionals extrapolate from their experience in both the private and the professional contexts.[3] A decline in the weighted average of experienced professional-context stock market returns predicts a roughly 1 percentage point reduction in expected future returns.[4] While private-context experiences also influence expectations, their effect is weaker: a 1 percentage point decrease in private-context returns predicts an expected return reduction of approximately 0.25 to 0.5 percentage points. This suggests that prior studies that focused solely on private-context experiences may have underestimated the overall impact of experiential learning on expectations. Second, regarding the dynamics of learning, we analyze how different episodes of the past are weighted relative to each other the private and the professional context. For private-context experiences, finance professionals show a recency bias in line with the literature on retail investors (Malmendier and Nagel, 2011). Specifically, finance professionals overweigh their recent private-context stock-market experiences when forming stock-market expectations. In contrast, in the professional context, the first impressions finance professionals happen to make at the beginning of their careers leave a particularly strong mark on their expectations. We rule out that our main effects are driven by informativeness of

[3] Note that the documented experience effect is inconsistent with associative memory, which suggests that individuals recall past information based on its similarity to new information (see, e.g., Wachter and Kahana, 2024; Enke et al., 2024a). Assuming that similarity depends on the context in which information is learned (private vs. professional), our finding of cross-context effect of private experiences in professional forecasts is inconsistent with associative learning.

[4] Our results do not directly compare to previous studies, since each data set warrants slightly different analytical approaches (e.g., fixed effects regression vs. probit model). Also, we would not expect experiences in different domains – e.g., stock markets vs. inflation – to leave the same mark on the respective expectations. Neither would we expect finance professionals to extrapolate from their private-context experiences to the same extent as retail investors.

experienced returns.

The financial setting is unique because experiences in both the private and professional context are reasonably exogenous — per Efficient Market Hypothesis both the birth year and career start decision are orthogonal to subsequent experience sequences of the stock market.[5] However, to facilitate identification and help pin down the exact professional activity that drives the context dependence, as well as the mechanisms behind it, we conduct an online experiment with retail investors (N=741). Participants experience a random sequence of returns from the distribution of a risky asset and are consequently asked to make an investment decision after the experience stage is over. Using a between-subject design, we randomly vary the *context* — the activity that the participants perform during the experience stage. Our design thus allows us to rule out context-dependent informativeness of returns, risk preferences and wealth effects as potential confounding factors. In the *Forecast Value* treatment, subjects are asked to play the role of professional forecasters to provide forecasts of future value. Alternatively, a *Forecast Sentiment* treatment asks participants to forecast the consensus forecast of the other investors in the experiment. At last, participants in the *Observe & Recall* treatment are prompted to recall the most recent experience without being asked to actively engage in any type of forecasting, thus closely resembling experiential learning in the private context. All experience-stage activities are incentivized to ensure that difference between the private-context and professional-context treatments are not driven by skin in the game. Furthermore, the experience-stage returns are drawn randomly for each subject, hence they are orthogonal to the context treatments. This experimental design allows for a clean identification of the context dependence and isolation of potential confounding factors.

We find that a primacy bias emerges only in the *Forecast Value* treatment. Notably, participants in this treatment undergo a structured representation of the occupation of professional forecasters — conducting regular incentivized forecasts of the future value and receiving immediate feedback as the actual value is subsequently revealed. Our experimental design allows us to distinguish between potential cognitive mechanisms. We find that the emergence of the context dependence in experiential learning is most consistent with differences in the dynamic processing of sequential

[5] In addition, we analyze survey non-response behavior and withdrawal from the survey or the industry, and we do not find any evidence that they are related to professionals' expectations, as we discuss in greater detail in Appendix Online Appendix B.

experiences. In particular, experiential learning in the *private* context (i.e. Observe & Recall treatment) is driven by *salience* and *attention*, consistent with salience-channeled attention as a micro-foundation of overreaction to new information (see Ba et al., 2024; Bordalo et al., 2024). In contrast, neither salience nor attention play any statistically significant role in the professional context. This is consistent with the growing body of literature documenting a stronger impact of salience-channeled attention on the financial decisions of retail investors than professional investors (see, e.g., Barber and Odean, 2007; DellaVigna and Pollet, 2009).[6] In this paper we show that susceptibility to exogenous salience and attention variations is not necessarily a function of financial expertise per se or a higher cognitive capacity of professionals. Instead, our findings suggest that differences in the susceptibility to salience-channeled attention may vary within-subject depending on the context of information processing.

In terms of the mechanism driving the primacy bias, we find that the relative weighting of experiences in the *Forecast Value* treatment is driven by the forecast error attached to the experiences. Specifically, experiences which occur as expected are overweighted, while surprising experiences are underweighted independent of the sign of the surprise. The impact of the forecast error is a distinguishing characteristic of model-free reinforcement learning. Barberis and Jin (2023) formally develop a hybrid theoretical framework of dynamic information processing as an interplay of both model-based and model-free learning algorithms. The framework predicts that only model-free reinforcement learning could result in a primacy bias. The intuition for this lies in the inherent inefficiency of model-free reinforcement learning — on average it underreacts to new information because it does not learn from the entire available data but only learns about the value of the outcomes of its previous actions. Coupled with context features that limit the decision-makers' tendency to explore different options, they may strive to quickly find a satisfactory action early on and then stick to that action for a very long time, ultimately underreacting to new information. This is exactly the dynamic implied by the estimated weighting function in the professional context *Forecast Value* treatment. This mechanism implies inefficiency in data processing caused by inherent features of a widespread professional activity — financial forecasting.[7] Specifically, subjects'

[6] It should be noted that effects of attention and salience are also evident in studies of finance professionals (see, e.g., Lu et al., 2016; Fedyk, 2024).

[7] Our evidence is consistent with Carlos Alos-Ferrer (2023), who show that incentives induce model-free reinforcement

experiences are coupled with simultaneous *error* cues — a necessary condition for reinforcement learning — only in the professional context *Forecast Value* treatment. The notion of inefficiency is intuitively at odds with professional expertise. However, it is consistent with a growing body of literature showcasing the susceptibility of finance professionals to behavioral biases and learning mistakes (see, e.g., Ben-David et al., 2013; Akepanidtaworn et al., 2023; Glaser et al., 2019; Linnainmaa et al., 2021). Against this backdrop, it is important to recognize that true expertise in financial forecasting — particularly in stock market forecasting — may be unattainable, as argued by Kahneman and Klein (2009). This is due to fundamental characteristics of the learning environment that likely hinder the learning process and, thus, the development of true expertise.

Related Literature. Our paper contributes to the strand of literature on experience effects. Prior studies have documented robust experience effect in the beliefs and choices of retail investors (see, e.g., Malmendier and Nagel, 2011, 2016; Kuchler and Zafar, 2019) and professionals (Schoar and Zuo, 2017; Malmendier et al., 2021). On the empirical side, experience effects can be rationalized with information search costs (Kuchler and Zafar, 2019; D'Acunto et al., 2021), memory encoding and retrieval (Laudenbach et al., 2019; Das et al., 2020; Gödker et al., forthcoming). On the theoretical side, experience effects are widely explained by context-dependent memory (Gilboa and Schmeidler, 1995; Mullainathan, 2002; Bordalo et al., 2023; Wachter and Kahana, 2024), which gives rise to recency effects (see, e.g., Bordalo et al., 2019). The theory is supported by experimental work on context-dependent expectations among finance professionals (Cohn et al., 2015; Guiso et al., 2018), retail investors (Laudenbach et al., 2021), and online experimental subjects (Afrouzi et al., 2023; Enke et al., 2024b; Graeber et al., 2024). Our contributions to this strand of literature are twofold. First, we suggest that different cognitive mechanisms may drive experience effects in different contexts. Second, in addition to previously proposed mechanisms related to memory and information retrieval, we highlight cognitive mechanisms related to the information processing Ba et al. (2024); Barberis and Jin (2023).

We also add to a growing body of literature on asymmetric and context-dependent learning. Along these lines, prior studies have suggested differences in learning in the gain versus loss domain (Kuhnen, 2015), in booms versus bursts (Ke, 2024), depending on ownership of the stock

learning. In our experiment, both the professional-context and the private-context treatments are incentivized but the incentives structures are inherently different.

(Hartzmark et al., 2020), for buy versus sell decisions (Akepanidtaworn et al., 2023), for beliefs versus choices (Barberis and Jin, 2023), from numerical information versus qualitative information (Graeber et al., 2024), with our without incentives (Carlos Alos-Ferrer, 2023). We contribute to this literature by adding the context of information processing — private versus professional. Importantly, we identify the type of incentives and the frequency of feedback as aspects of the professional environment which interfere with the efficiency of the information processing.

Context dependence in information processing, as evidenced in this study, provides important insights for modeling. Notably, the co-exitence of primacy bias and recency bias in the field cannot be accounted for by model-based learning alone, as shown by Barberis and Jin (2023), emphasizing the need for dual-system or hybrid models of learning (see, e.g., Thaler and Shefrin, 1981; Camerer et al., 2004; Ba et al., 2024).

The rest of the paper proceeds as follows: Section I introduces our field setting and methodology, while Section II outlines our findings from the field. Section III outlines our experimental design and results, while Section IV discusses potential mechanisms. In Section V we illustrate the implications of primacy and recency bias for the dynamics of learning over time. Finally, Section VI concludes.

I. Field Data and Methodology

A. Survey of Professionals' Expectations

We use data from the German ZEW Financial Market Survey, conducted monthly since December 1991 among roughly 300 financial professionals. Participants provide semi-annual forecasts on macroeconomic and financial variables.[8] The survey results influence highly liquid financial markets (see, e.g., Entorf et al., 2012) and rank among the top three indicators for institutional investors in Germany.[9]

Most survey participants—economists, analysts, and financial advisors at banks, insurers, and employees in financial departments of large corporations — regularly produce stock market expectations as part of their job. The expectations of these professionals serve as input for the investment

[8] Recent survey data and information on the survey can be found at `http://www.zew.de/economicsentiment`.

[9] More than 95% of institutional investors interested in Germany subscribe to the announcement of the survey responses of the professionals in our panel. For details on Bloomberg attention data, see Online Appendix A.

decisions of others. Hence, biases in the professionals' expectations potentially pose a negative externality to investors who use them.[10]

When working with survey data, there might be concerns about strategic response behavior or response fatigue among subjects. To avoid traditional reasons for strategic response behavior, such as rational herding arising from rank-dependent remuneration (see, e.g., Lamont, 2002; Hirshleifer and Hong Teoh, 2003), the ZEW publishes the results strictly anonymously, and there are no monetary incentives for survey participation. To reduce response fatigue, participants answer only relevant, non-mandatory questions. We provide evidence against endogenous attrition, where subjects would leave the panel or not answer in some survey waves for reasons related to their stock market expectations in Online Appendix B.

B. Expectations

We examine professionals' six-month ahead DAX expectations since its introduction to the survey in February 2003 until December 2017. Specifically, the numerical question on DAX expectations reads:

"I expect the DAX in 6 months to be at ... points."

We convert the responses to the above survey question into expected stock market returns using the DAX opening level on the day of the response.[11]

Table I provides descriptive statistics on the survey. Our monthly panel covers 330 individual finance professionals. The panel structure is unbalanced, with the average professional submitting 71.2 responses over the entire sample period of 178 months. This long panel structure allows us to identify the effects of experienced returns out of the dynamics of expectations within-subject. The professionals in the sample have many years of work experience: Their average tenure is 21

---

[10] Among our professionals, 55% have external clients using their forecast, and for 79% the forecasts support the institutional trading strategy, according to a supplementary survey from June 2013 among 122 of the survey participants.

[11] While asking for expected prices instead of expected returns induces more pessimistic expectations and amplifies the recency bias (Glaser et al., 2019), both effects are orthogonal to the context-dependence in experiences we analyze here, as the questionnaire design is the same for all subjects. The type of recency bias documented in the study concerns extrapolation from commonly used charts of past performance and constitutes a time-fixed effect which we control for in our analyses. We provide more details on data handling in Online Appendix A.

years, and at the tenth percentile the tenure amounts to 9 years. Professionals who have left the panel during the sample period are included in our analysis to avoid survivorship bias. The average expected semi-annual return is 2.85%. This average is close to the historical average semi-annual DAX return of 2.94%.

[Insert Table I approximately here.]

C. Experiences

We consider for each professional a personally experienced period of stock market returns, specifically the returns of the German stock market index DAX. We measure experienced returns in two different contexts – the private and the professional context. The professionals' private-context experiences follow Malmendier and Nagel (2011) and stand in contrast to their professional-context experiences defined by their time in the finance industry related to Dittmar and Duchin (2016).

C.1. Professional-Context Experiences

Our dataset contains information on the professionals' birth year and the year of their career start in the finance industry. The sample displays a large heterogeneity in professionals' age at their career start.[12] This heterogeneity allows us to distinguish between return experiences in the private context and the professional context. Specifically, our panel covers a broad range of cohorts with diverse early career experiences. Importantly, a significant proportion of the sample has experienced a crisis (e.g., 1973-74 stock market crash, the Black Monday stock market crash of 1987, the early 1990s recession, and the collapse of the dot-com bubble 2000-2003) as well as booms of similar magnitude (e.g., 1993, 1997, etc.).[13]

We rule out endogenous selection into the sample. Specifically, we provide evidence that selection into the finance industry does not plausibly correlate with subsequent returns in and efficient market in Online Appendix A.

We provide evidence against endogenous sample attrition in Online Appendix B. While neither

---

[12] We provide the distribution of career start ages in Online Appendix A, Figure A1.

[13] We provide an illustration of market returns at different cohorts' career start in Online Appendix A, Figure A2 along with a detailed discussion to rule out endogenous selection into the sample.

sample selection nor sample attrition can be present in our experimental study, its results confirm our field study, providing further evidence against endogenous selection in or out of our field data.

C.2. Private-Context Experiences

Analyzing experienced returns in the private context requires assumptions about the age at which subjects become aware of the stock market and start perceiving its development. Previous studies assume that subjects start experiencing the stock market development at birth, potentially indirectly through its impact on their parents. Instead of deliberately fixing the age at which subjects start being aware of the stock market, we estimate the most plausible starting age as an additional parameter from our data. It is possible to distinguish the private context from the professional context due to the variation of age at career start in our sample.

D. Model Specifications

We estimate non-linear models given by:

$$y_{it} = \beta f\left(s_i, t; \vec{\theta}\right) + \vec{\zeta}^{\top} \vec{x}_{it} + \epsilon_{it} \tag{1}$$

where y_{it} is the expected DAX return of subject i at time t converted from semi-annual to average monthly returns. β and $\vec{\zeta}$ are slope coefficients. $\vec{x}_{it}$ refers to covariates in the model (including an intercept and subject-fixed effects), and ϵ_{it} are iid noise terms. The function f aggregates the weighted average stock return over the individually experienced period of subject i. We aggregate experienced returns using two weighting functions: a non-monotonic weighting function suggested by Prelec (1998), and a monotonic weighting function following Malmendier and Nagel (2011).

$$f\left(s_i, t; \vec{\theta}\right) = \sum_{j=1}^{t-s_i} \bar{w}_j(t - s_i; \vec{\theta}) R_{s_i+j-1} \qquad t > s_i \tag{2}$$

The individually experienced period of subject i begins at a starting point s_i and ends at time t. In the function f, $\vec{\theta}$ comprises of non-linear parameters determining the weights of experienced returns. For the monotonic weighting function, $\vec{\theta}$ comprises of only one parameter, λ, the *recency parameter*, which determines the difference between the primary episodes of the experienced period and the most recent episodes. For the non-linear weighting function, $\vec{\theta}$ includes two parameters: δ, the *primacy parameter*, determines the relevance of primary relative to recent episodes analogously to the parameter of the monotonic function. γ, the *hump-shape parameter*, determines the relative relevance of intermediate episodes compared to the two extreme periods (primary and recent episodes) jointly.[14]

We estimate $\vec{\theta}$, β, and $\vec{\zeta}$ in Equation (1) jointly using maximum likelihood following Golub and Pereyra (1973). We restrict the parameters δ and γ to be positive by maximizing the log of the parameters. The Wald-tests we provide in Section II are on the log scale. Two-way clustered standard errors are computed following Zeileis (2006), with a cluster on the individual and time and with the standard cluster bias adjustment (Cameron et al., 2011).

We include month-year fixed effects and subject-fixed effects to circumvent potential endogeneity arising from a correlation of experienced returns with time-level variables and subject-level variables. Specifically, subject-fixed effects cover personal traits such as risk aversion, optimism, overconfidence as well as potential selection effects that result from an interaction between personal traits and market history prior to the career start (e.g., if optimists were more likely to start a career in finance after a bust than pessimists). It should be noted that the fixed effects also capture two extreme scenarios of the relevance of experienced returns, hence our model setup identifies a lower bound of the potential effect of experienced returns: Firstly, month-year fixed effects cover recent returns, which are experienced by every professional in the sample. Hence, extreme cases where the professional bases his predictions entirely on returns realized over the past several months are not identifiable as they are subsumed by the time-fixed effects. Secondly, subject-fixed effects cover returns that are observed at the start of the career and are never forgotten. Hence, our analysis captures the remaining effect of experienced returns which varies within-subject over time because of (i) short-tenure professionals putting more weight on new observations than long-tenure professionals, and (ii) the memory of experienced returns fading over time.

All regressions include day-of-week fixed effects that capture possible calendar effects, which are not covered by the month-year fixed effects. Such additional calendar effects arise because, in each monthly survey, professionals do not respond on one day, but instead, the responses are

[14] For a more detailed description of the weighting function, see Online Appendix A.

spread over typically 11 trading days. To take out any potential effect of macro experiences, we provide two approaches: (i) we take out the impact of inflation expectations (as previously documented by D'Acunto et al. (2021); Malmendier and Nagel (2016); Malmendier et al. (2021)), using excess returns over the risk-free rate; (ii) we control for macro expectations assuming that any other macro experiences affect macro expectations directly and stock market expectations indirectly through macro expectations. We measure excess returns as weighted average monthly DAX returns in excess of the money-market rate.[15] We control for our finance professionals' macro expectations with similar measures as used by Greenwood and Shleifer (2014) and Amromin and Sharpe (2014).[16]

II. Experience Effects by Context: Evidence from the Field

In this section, we examine how experienced returns can explain the dynamics of expectations of different cohorts of finance professionals. We allow for different dynamics in the professionals' private context and their professional context: In Subsection II.A, we focus on the experienced returns in the professional context. In Subsection II.B, we analyze the effect of experienced returns in the professionals' private context. In both contexts, we examine which episodes of the respective experienced period are particularly important - the primary episodes, the recent episodes, or both. The focus on different episodes of the past may differ in the two contexts (i.e., private and professional), and learning about the importance of primacy and recency might help to understand why experienced returns affect expected returns, which is the subject of the next section.

A. Experienced Returns in the Professional Context

This subsection focuses on the returns finance professionals have experienced since the beginning of their finance careers, i.e., in their professional context. We uncover the effect of experienced returns in the professional context on the dynamics of professionals' stock market expectations. We formalize the analysis according to the models introduced in Section I.D.

[15] For 1959–2012, we use the Detusche Bundesbank's money market rate (code BBK01.SU0104). For data on a risk-free rate after May 2012, we use the 1-month Euribor, as available on Datastream.

[16] For definitions of the variables see Table AI.

[Insert Table II approximately here.]

The significant β estimates in Table II support that experienced returns affect professionals' expected returns. This result is robust across model specifications and holds for both the non-monotonic weighting function in Panel A and the monotonic weighting function in Panel B. According to the coefficient estimate β in the fully specified model in Column (ii), a decrease in the weighted average experienced returns by 1 percentage point predicts a lower expected return of around 1 percentage point. This interpretation applies to both weighting functions. With subject-fixed effects and time-fixed effects, the coefficient estimate for β captures the effect of within-subject variation of experienced returns resulting from (i) "fading memory" of the primary observations from the beginning of the career and (ii) decreasing "attention" to any single observation, including the most recent observations, as tenure increases. To illustrate the effect of "fading memory" of the primary observation, consider a professional who started his career during a financial crisis similar to 2008 when the DAX decreased by roughly 40%. Independent of the market development after the first year, the memory of the financial crisis would fade over time as the weight of the first year is reciprocal to the professional's tenure. The shape of the weighting function determines how slowly the memory of the financial crisis would fade. Following the coefficient estimate in Column (ii), a financial crisis comparable to 2008 can be expected to lead to overpessimism compared to the base rate and a Bayesian benchmark with limited information. The estimate implies that this overpessimism might take over 45 years to dissolve fully; hence, for most professionals, the bias would not be dissolved before retirement.

All model specifications include time-fixed effects, which means that experienced returns impact expected returns over and above the impact of any public information available to all professionals. Such information includes but is not limited to the entire history of stock market returns.

Furthermore, the effect of experienced returns is robust to including subject-specific macroeconomic expectations and assessments. Macroeconomic experiences may confound stock market experiences. Stock market expectations, however, should only be indirectly affected by macroeconomic experiences. Macroeconomic experiences should primarily affect macroeconomic expectations, which in turn can affect stock market expectations. Therefore, we control for macroeconomic expectations to isolate the effect of experiencing the stock market from a potential effect of macroeconomic experiences, which might be reflected in experienced stock market returns. We find an effect of macroeconomic expectations on stock market expectations, which is strong and consistent with previous studies.[17] The effect of stock market experiences is qualitatively unchanged after controlling for macroeconomic expectations and, hence, not driven by macroeconomic experiences.

Our results in Table II are consistent with previous experimental evidence that information learned from experience, observation, or simulation prevails over information learned from data or description (see, e.g., Hertwig et al., 2004; Kaufmann et al., 2013). Our findings validate that the experimental results apply to the field for sophisticated and experienced investors such as the professionals in our sample.

The empirical evidence in Table II is inconsistent with alternative explanations from selection effects. Specifically, the patterns we observe are at odds with selection towards (i) optimists and (ii) contrarians. First, in a scenario where professionals with a bullish career start in the finance industry become overly optimistic, such optimists will presumably endure adverse market environments and stay in the finance industry. If optimists were consequently over-represented in the sample, one would expect a high weight on the experiences at career start (under-reaction to recent episodes), but at the same time mean-reverting expectations after market downturns. While our evidence is in line with under-reaction to recent episodes, instead of mean-reverting expectations, we find extrapolative expectations. Taken together, we infer that subjects with good or bad early market experiences are equally likely to remain in the sample. Second, in a scenario where subjects extrapolate excessively from the market environment at career start, they would be prone to losing their jobs due to low forecast quality. The professionals remaining in the sample would be more likely to show mean-reverting expectations. However, our evidence is inconsistent with selection towards contrarian expectations because we obtain a positive beta, which implies that the sample subjects show extrapolative expectations.

Figure 1 illustrates the estimated shape of the weighting function for a professional with 10 years (120 months) of tenure in finance. The subfigures correspond to the different model specifications outlined in Table II. Subfigure a. is for the non-monotonic weighting function in Table

[17] Although it contradicts traditional asset-pricing theory, the empirical link between macroeconomic expectations and expected returns is consistent with previous literature on expectation formation of investors (see, e.g., Greenwood and Shleifer, 2014; Amromin and Sharpe, 2014).

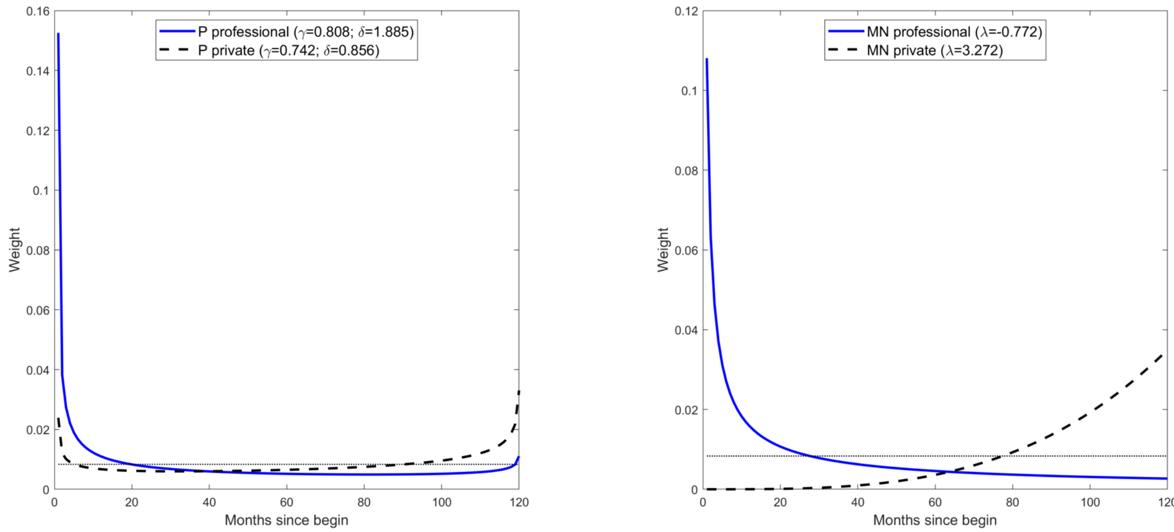

a. Non-monotonic weighting function

b. Monotonic weighting function

Figure 1. Weighting functions by context - professional versus private. This figure compares the weighting functions resulting from the different model specifications displayed in Table II, Columns (ii) and (iv) in Panel A and B for a subject with a tenure of 10 years (120 months). Panel A illustrates the estimates based on the non-monotonic weighting function by Prelec (1998) and Panel B displayes the estimates for the monotonic weighting function by Malmendier and Nagel (2011).

II, Column (ii). Subfigure b. is for the monotonic weighting function in Table II, Column (iv). Each subfigure shows the estimated weighting function relative to the Bayesian benchmark. For the Bayesian benchmark, we assume that subjects have limited information (i.e., no access to information before their career start) and a stationary data-generating process – yielding an equally weighted function. Figure 1 shows that for both weighting functions, compared to the Bayesian benchmark, professionals put higher weight on the primary episodes in the first months of their careers. This primacy effect is captured by the parameter λ in the monotonic weighting function. According to the λ estimates in Table II, Panel B, the primacy effect is statistically significant at the 1% level in all model specifications (columns (i)-(iii)). A qualitatively similar result obtains for the non-monotonic weighting function in Panel A (columns (i)-(iii)). Even though the non-monotonic function allows for various functional forms, it confirms the evidence favoring a primacy effect. To the best of our knowledge, we are the first to document a primacy effect in how finance professionals extrapolate from their experiences.

A primacy effect contradicts previous experimental evidence of recency and salience effects (e.g., Langer et al., 2005; Ariely and Zauberman, 2000) as well as previous empirical findings of a recency effect in the expectations and financial risk-taking of retail investors (e.g., Malmendier and Nagel, 2011; Kuchler and Zafar, 2019). The difference may indicate that subjects learn differently from experiences in the private context and from experiences in the professional context.

In particular, the difference can stem from contrasting private-context experiences, possibly over the subject's entire lifetime, with experiences from a defining moment in life. Extrapolation from private-context experiences, where the early episodes may be less formative than the recent episodes because of subjects' relatively low awareness and motivation for information acquisition during childhood, favors a recency effect. Extrapolation from experiences that begin at a defining moment in life, when subjects are particularly aware and attentive to the experience, forwards a primacy effect. The beginning of one's career is arguably an example of a defining moment that requires attention and awareness.[18] The primacy effect in our findings is consistent with the effect of early-career economic background on managerial style (Schoar and Zuo, 2017), as well as the effect of growing up in a low-socioeconomic-status environment on economic expectations (Das et al., 2020).

Furthermore, the difference can be due to differences between economic and financial time-series characteristics: economic experiences based on persistent economic time-series can give rise to a recency effect as suggested by Malmendier and Nagel (2016). In contrast, extrapolation from (stationary) financial time series might facilitate a primacy effect.

B. Experienced Returns in the Private Context

This subsection focuses on the returns subjects have experienced since a certain age – which is what we call the private context. Here again, we formalize the analysis of the effect of experienced returns in the private context on the dynamics of professionals' stock market expectations. We treat it as an empirical question to determine the age since when subjects begin experiencing returns and obtain the starting age of the experienced period with the best fit as a result of the maximum-likelihood estimation.

Our estimates in Table II, Panel B, column (iv)-(vi), show that, in the monotonic weighting

---

[18] An opposite effect of distraction by defining private events, in particular marriage, has been shown by Lu et al. (2016).

function model, professionals extrapolate from their experienced returns since the age of 23. Furthermore, in their private context, professionals exhibit a recency effect as indicated by a positive λ estimate (also see Figure 1). The estimates are statistically significant at the 10% level.The recency effect in the extrapolation from returns experienced in a private context is consistent with previous findings on the risk-taking of retail investors (Malmendier and Nagel, 2011), and findings on housing and inflation expectations of consumers (Kuchler and Zafar, 2019). Hence, we confirm these previous results for (i) a direct measure of expected stock market returns and (ii) professionals. In contrast to previous results that capture a joint effect of cross-cohort variation and dynamics of expectations within a cohort, our identification relies exclusively on the dynamics of expectations.

C. Robustness

One channel through which personal experiences affect beliefs and choices is when personally experienced returns constitute a valuable information source. For example, experienced returns can be information if data access to objective information is restricted or costly (see, e.g., Kuchler and Zafar, 2019). However, in our setting, experienced returns have no marginal value as an information source. This is because we analyze subjects who have abundant information at their disposal, including information on the complete history of stock market returns, which is more data than the subset they could have personally experienced. By estimating models with month-year fixed effects, we control among other things for the effect of the long-term average DAX return, which constitutes the objective base rate. Hence, our results show that experiencing the past has an effect over and above the objective information on the past, which is available to everybody. The effect of experienced returns, which we trace in this study, is an effect of non-informative observations. When exhaustive descriptive information is available and personal observations are non-informative, one would not expect subjects to resort to the subsample of experienced observations, because they cover only a short period of the past, leading to a small sample bias and under- or overweighting of the probability of rare events.

Even though experienced returns are non-informative in the long run, in a small sample they could have a positive correlation with realized returns by chance. Following this line of reasoning, it can be hypothesized that professionals may have come to appreciate their experienced returns because they might have coincidentally been an accurate predictor of future returns over the relatively short period of their careers. To test this hypothesis, we analyze whether the professionals in our sample have observed a positive correlation between their experienced returns and actual realized returns over the course of their careers.

[Insert Table IV approximately here.]

Specifically, in Table IV we calculate the aggregate experienced return for all professionals in our panel in a given year and examine its correlation with the realized return. For stability, we require that at least 50 subjects from our pool of professionals have already started their finance career for a year to be considered. The results do not provide any evidence that the correlation over time is positive. In the long run, the correlation is negative and statistically significant at the 1% level. This finding is consistent with the theoretical prediction of an asset-pricing model with agents who extrapolate from personal experiences (Nagel and Xu, 2022).

Another way in which past returns might affect investment behavior is through risk preferences. Guiso et al. (2018) elicit the risk aversion of retail investors and report a significant decrease after the 2008-2009 financial crisis. Differences in risk aversion may be linked through experienced returns, e.g. through wealth effects, if our proxy for experienced returns is correlated with the income and/or portfolio returns of the professionals. At the same time, the professionals' risk aversion may affect their expected returns (i) through a required risk premium, assuming that professionals understand expected returns as required returns, or (ii) through a risk-adjustment, assuming that professionals understand expected returns as risk-neutral returns (Cochrane, 2011). With respect to the latter, Adam et al. (2021) show that survey expectations are inconsistent with the risk-adjustment hypothesis. With respect to the former, previous evidence suggests that professionals do not think about the required risk premium when thinking about expected returns (Kaustia et al., 2009). To account for the potential effect of risk preferences, as well as other potential confounding factors, we design an experimental study, which is destribed in detail in the next section.

III. Context-Dependent Experience Effects in the Lab

Results from the field, outlined in Section II, suggest a context dependence of experience effects — finance professionals extrapolate from the *most recent* returns experienced in the *private context* but from the *initial* returns experienced in the *professional context*. While the field setting enables us to identify a context-dependence, it limits our ability to test for potential explanations. An open question remains: How do individuals experience stock market developments differently in private versus professional contexts, and which of these differences are relevant for the shift between “recency bias” and “primacy bias”. We designed an experimental study to identify and explain how the context in which experiences are accumulated causally influences their effect on subsequent choices.

A. Experimental Design

The aim of the experiment is to (i) establish a *causal* link between the context of the experience and the shift between “recency” and “primacy”, and (ii) to test for alternative *cognitive foundations* of the context-dependence. People face an incentivized allocation decision between a risky and a zero-interest risk-free asset. Importantly, people learn about the properties of the risky asset by sequentially drawing (i.e. “experiencing”) 10 random observations from its distribution (see, e.g., Hertwig et al., 2004; Bohren et al., 2024, for an overview of experiential learning in experiments). The random sequences consist of independently and identically distributed draws from a normal distribution with parameters matching the sample distribution of the MSCI World index ($\sim N(0.11, 0.15^2)$ *iid*). The random sequences are drawn individually for each subject. To enhance external validity, we present the sequence of random draws in the experiencing stage using a value format, as prior evidence suggests that retail investors and finance professionals predominantly rely on values or prices rather than returns (see Glaser et al., 2019).

In a between-subjects design, we manipulate the *context* of the experience — that is, the activity participants engage in while *processing* each random draw. These activities are incentivized in all context treatments. The experiment consists of three context treatments, that are outlined in Figure 2 and described below.

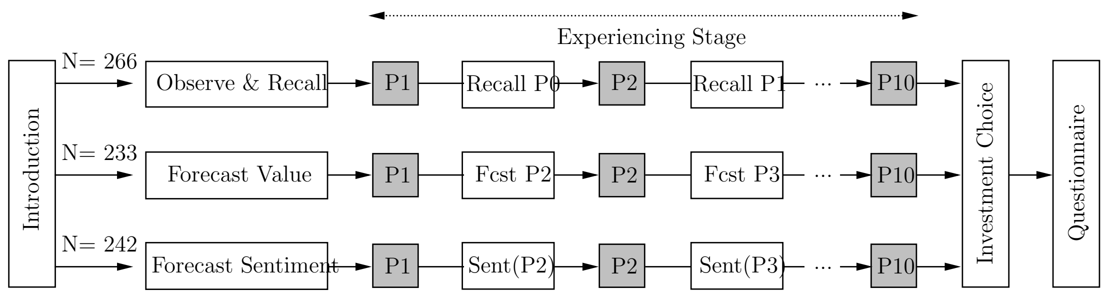

Figure 2. Experimental design and sample size.

In the "*Forecast Value*" treatment, we ask subjects to state their expectation for the next draw after observing each draw. This treatment is representative of the professional environment of professional forecasters. It also aims to capture any activity which requires a decision to be made after each new piece of information. This context treatment is motivated by recent neurophysiological evidence suggesting that the formation of expectations is initiated in the process of decision-making and does not happen inherently in the process of information acquisition (Nursimulu et al., 2016). We assume that the frequency of expectation formation — whether explicitly through forecasting or implicitly as part of decision-making — is higher in the professional context than in the private context, making it more aligned with the frequency of information updates. Put differently, a professional investor who frequently makes forecasts or decisions processes stock market experiences differently from an investor who, while attentive and with skin-in-the-game, remains mostly passive.[19] Another inherent feature of step-by-step expectation formation is the immediacy of feedback, particularly hit-or-miss feedback. As a recent study by Carlos Alos-Ferrer (2023) demonstrates, hit-or-miss feedback facilitates model-free reinforcement learning — a learning algorithm in which successful choices are likely to be repeated, while unsuccessful ones prompt exploration of alternatives. Barberis and Jin (2023) theoretically demonstrate that, unlike Bayesian learning, model-free learning can generate a primacy bias. At the very least, it predicts a weaker recency bias than model-based Bayesian learning.

[19] It is important to note that a high frequency of expectation formation and decision-making is not exclusive to the professional context (see, e.g., Heimer and Imas, 2022; Heimer et al., 2025, for examples of high-frequency trading by retail investors in the "private" context). However, since retail participation in high-frequency trading strategies surged only toward the end of our sample period, we argue that it is unlikely to have significantly influenced the private-context experiences of the subjects in our field dataset.

Another key distinction between the "professional" and "private" contexts is the degree to which individuals form second-order expectations about other investors' beliefs and choices. Specifically, second-order beliefs are more likely to emerge in the "professional" context, where the prevalence of rankings and tournament incentives in the finance industry plays a significant role (see, e.g., Kirchler et al., 2018). Consistent with this, Camerer et al. (2004) suggests that prompting individuals to consider others' responses enhances their k-level thinking, with professional stock market portfolio managers exhibiting more advanced k-level reasoning than non-professionals. Thus, we hypothesize that second-order belief formation is more prevalent in the "professional" than in the "private" context.[20] To test whether second-order belief formation contributed to the recency-primacy-shift in experience effects, we ask subjects in the "*Forecast Sentiment*" treatment to predict the average forecast of others during the experience stage.

The "*Observe & Recall*" treatment most closely captures the experience of the stock market in a private context. The main challenge in designing this treatment is to isolate potential confounding factors. One factor that has been extensively studied in prior theoretical and empirical literature on experience effects is memory (see Malmendier and Wachter, 2024, for an overview). While memory processes related to the encoding and retrieval of experiences mostly predict a recency bias in the weighting of past experiences, they have the potential, under certain conditions, to also generate a primacy bias (see Section IV for a discussion). In the experiment, our focus is on other potential factors of the recency-primacy shift, which is why we need to make sure that the effect of memory is possibly small and symmetric in all treatments. Moreover, we need to ensure that incentives during the experiencing stage do not confound the between-subject treatment effects. To this end, subjects in the private-context "Observe & Recall" treatment, like the participants in both other treatments, perform an incentivized task while processing the sequence of experiences. Specifically, we ask them to actively recall the experience immediately preceding the most recent one. This activity mirrors how a reasonably attentive non-investor engages with the stock market

[20] Given the composition of our panel of finance professionals, which is by construction dominated by analysts, we consider this a useful generalization. However, it is important to note, that thinking of others is neither inherent to nor exclusive for the "professional" context. Specifically, a qualitative study by Andre et al. (2024) measures the mental models of finance professionals and retail investors. The study shows that (i) within the group of finance professionals, financial advisors often think about other investors' reactions and fund managers largely do not, and (ii) some retail investors also think about other investors' reactions.

in the real world — learning about recent stock market performance and rehearsing the information over a couple subsequent days through social interactions or media channels. The assumption of reasonably attentive non-investor is particularly fitting for the early-life private-context experiences of individuals who later become finance professionals such as the subjects in our field dataset. By mirroring these naturalistic informational processes, the treatment also provides a more ecologically valid framework for examining how past experiences shape decision-making in a private context.

The experiment (N=741, 49% male, average age of 42 years) was pre-registered[21], programmed in Qualtrics and administered among investors[22] on Prolific in January 2025. We use further sample selection criteria based on location (US) and approval rate (> 90%). Participants received approximately $23 per hour for their participation in the experiment, which is twice the typical amount offered on Prolific.[23] Before the beginning of the experiment, subjects underwent a bot check, as well as attention and comprehension checks (see Figure D1 in Online Appendix C). Only subjects who passed all checks were allowed to proceed with the study. Fewer than 5% of participants exited the experiment after being randomized into treatments, which assuages concerns about selective attrition.

It is essential to verify that the between-subject context treatments are effective, meaning that participants do not respond randomly but instead engage in the respective activities as intended during the experiencing stage. Effort is most easily verified in the "Observe & Recall" treatment, where the recall question has a definitive correct answer. Among participants, 89.1% made no more than two mistakes (mostly towards the end, possibly due to fatigue effects), and 60.1% answered all recall questions correctly.[24] We expect both professional-context treatments to be more effortful than the "Observe & Recall" treatment, as they require participants to engage in an expectation-formation task. Unlike the information-processing task present in all context

[21] The pre-registration report is available at https://aspredicted.org/ftpm-fkv9.pdf

[22] An *investor* is defined as anyone who responds affirmatively to the question "Have you ever made investments (either personal or through your employment) in the common stock or shares of a company?"

[23] Online platforms, such as Prolific, are increasingly used in economics, finance and the broader social sciences to recruit subjects for experiments. Studies have shown that laboratory results broadly replicate on these online platforms (e.g., Horton et al., 2011; Snowberg and Yariv, 2021).

[24] The analyses of the "Observe & Recall" treatment presented below are robust to exclusion of the participants who have made at least one mistake in the recall-question sequence.

treatments, expectation formation does not occur automatically and demands additional cognitive effort (see Nursimulu et al., 2016). Consistent with this hypothesis, Figure 3 shows that subjects in both professional-context treatments (i) spend more time per period during the experiencing stage and (ii) rate the task as more mentally demanding, as measured by perceived complexity following Maynard and Hakel (1997) (see Online Appendix C for details and wording). The difference between the "Observe & Recall" treatment and each professional-context treatment is significant at the 5% level, while no significant difference is observed between the two professional-context treatments.

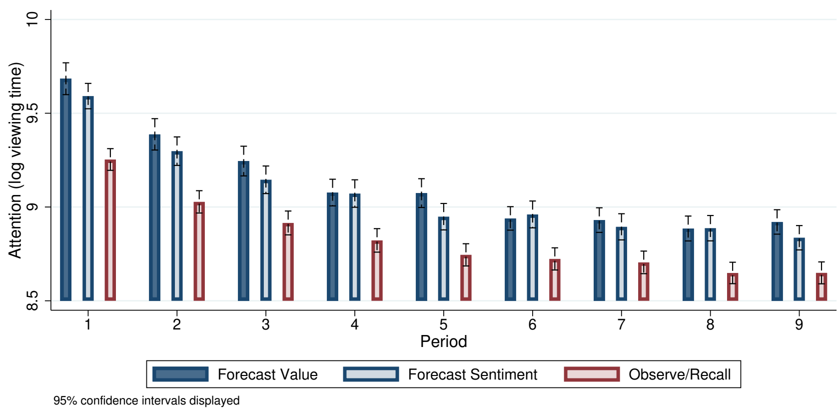

a. Dynamics of Attention by Treatment

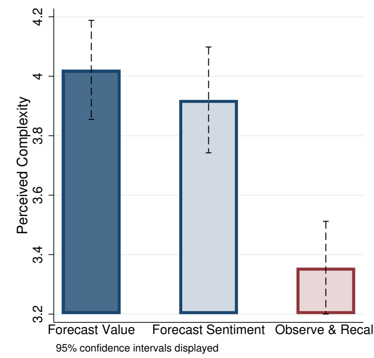

b. Perceived Complexity by Treatment

Figure 3. Measures of effort by experimental task. This figure compares the effort exerted by the subjects in each experimental treatment. First, in subfigure a, we plot the attention to the individual experienced returns period by period. Attention corresponds to the logarithmized viewing time in each period. Logarithmizing is appropriate as the distribution of viewing time is positively skewed. Second, in subfigure b, we show the perceived complexity of the investment task using the measure of Maynard and Hakel (1997).

B. Context-Dependent Experience Effects in the Lab: Results

Our results, outlined in Table V, indicate that participants extrapolate from past returns across all between-subject treatments. However, subjects in the "Forecast Value" treatment, in particular, assign roughly twice as much weight to experiences from the first half of the sample (cumulative weight = 67.9%) compared to those from the second half (cumulative weight = 32.1%). Notably, the most recent period is significantly underweighted (assigned weight = 1.2% vs. an equal-weighting benchmark of 10%, p-value < 0.001), as is the second-most recent period (assigned weight = 3.5%, p-value = 0.021). Table V and Figure 4 display estimates of the non-monotonic weighting function by Prelec (1998). The table shows that the estimates are robust to the inclusion of risk preferences and other demographic control variables. Taken together, these findings suggest that, in the “Forecast Value” treatment, participants exhibit a systematic and robust primacy bias.

[Insert Table V approximately here.]

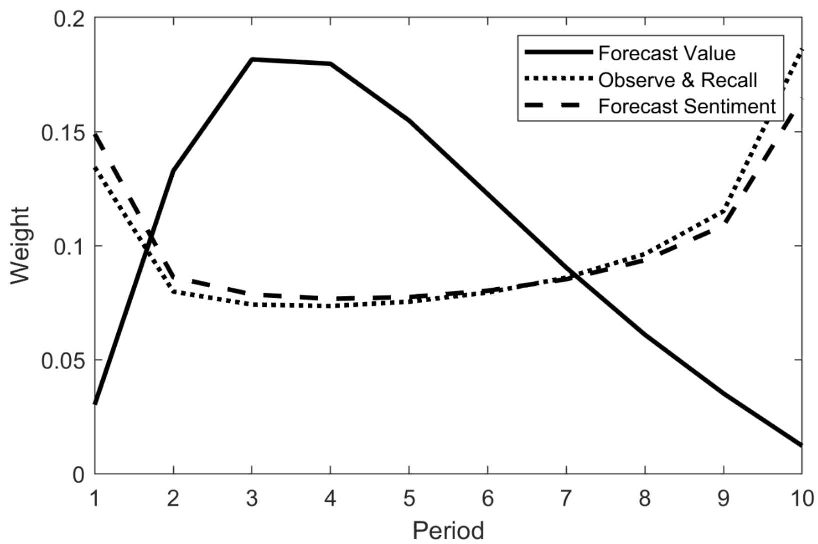

Figure 4. Weighting functions by experimental treatments. This figure illustrates the weighting functions by context treatment, as reported in Table V, Columns (i), (iii) and (iv).

In contrast, participants in the ”Observe & Recall” treatment extrapolate equally from both recent and initial observations. The cumulative weight of the second half of the sample is 56.3%. The estimated weighting function is statistically indistinguishable from an equal-weighting function at the 10% significance level. This finding is consistent with previous experimental studies on passive sequential observations, which are most closely linked to the design of the “Observe & Recall” treatment (see, e.g., Ungemach et al., 2009).[25] Similarly, the weighting function in the “Forecast Sentiment” treatment is not distinguishable from equal weighting at the 10% level. Even though both the “Forecast Value” and “Forecast Sentiment” treatment simulate activities, which

[25] It should be noted that our experimental design of the “Observe & Recall” treatment deviates from the traditional experiential learning paradigm in two important aspects, which may mitigate potential recency bias. First, the length of the sequence of experiences is exogenous. Hertwig et al. (2004) shows that this design choice makes the reliance on salience cues more likely, as we show in Section IV.A. Second, the participants are asked to actively “rehearse” past experiences, which may result in their re-encoding in memory and consequently reduce the recency bias (see Mullainathan, 2002; Wachter and Kahana, 2024). Both design choices are necessary to preserve the internal validity of the treatment comparison.

are more characteristic of the professional than the private domain, a primacy bias in experience effects, similar to the one in the field data, emerges only in the "Forecast Value" treatment.

Is it reasonable to expect that the activities in the "Forecast Value" treatment reflect those performed by finance professionals in their professional domain? A one-time survey question administered in June 2013 to 194 respondents from the ZEW panel indicates that 58% regularly engage in the same activity as in the "Forecast Value" treatment — regularly and explicitly forecasting stock market values. This supports the validity of the assumption. However, it remains unclear why this particular activity, unlike other professional or private-context tasks, leads to a primacy bias. To fill this gap, Section IV delves into the cognitive mechanisms underlying the primacy bias.

IV. Cognitive Mechanisms of the Context-Dependence

Experiential learning is characterized by *sequential* observations and *dynamic* updating of beliefs. In what follows, we explore potential drivers of context- dependence related to both processes.

The first set of potential explanations comes from the fact that past experiences unfold as a sequence of events that are committed to memory. Consequently, beliefs and choices are shaped by the encoding and retrieval of the memories of these events rather than the objective data. As a consequence, memory processes are proposed to be major drivers of biases in experiential learning (see Malmendier, 2021; Malmendier and Wachter, 2024, for a review of memory-based explanations). To the extent that memory drives experience effects, a *recency* bias is implied by one of the fundamental laws of human memory, the Law of Recency (see Kahneman et al., 1993; Kahana, 2012, for an overview). Furthermore, a recency bias may result from imprecise memory due to memory constraints (Azeredo da Silveira et al., 2024). In contrast, potential explanations of a systematic *primacy* bias are few and far between. Distant memories could be retrieved when the decision context shares more features with the context of distant memories than the context of recent memories, thus creating stronger associations with distant memories (e.g., Tversky and Kahneman, 1973; Bordalo et al., 2023; Enke et al., 2024b). Alternatively, distant memories may interfere with the encoding of new memories (see, e.g., Das et al., 2020; Laudenbach et al., forthcoming). Last but not least, distant memories may play an increasingly stronger role over time if they are reencoded, or “rehearsed”, each time they are retrieved (Wachter and Kahana, 2024; Mullainathan, 2002). Taken together, memory-based explanations formulate a set of conditions for the retrieval of distant memories over recent memories. However, such explanations do not incorporate an unconditional prioritization of the retrieval of distant memories over recent memories, which does not depend on other features of the distant experience other than its timing.

The second set of explanations focuses on the *dynamic* updating of beliefs, given perfect memory about past experiences. Even with precise memory, salience (see, e.g., Tversky and Kahneman, 1973) and salience-channeled attention may distort the mental representation of the data (see, e.g., Griffin and Tversky, 1992; Bordalo et al., 2020, 2024). A study by Ba et al. (2024) identifies two main complementary cognitive mechanisms of belief updating. One mechanism — salience-channeled limited attention — distorts the mental representation of the data, resulting in overreaction to new information, consistent with a recency bias. This mechanism is in line with prior literature explaining the recency bias in experience effects with salience and limited attention (see, e.g., Malmendier and Nagel, 2011, 2016). The other mechanism — a cognitive imprecision distorting the processing of complex data in the absence of salience cues — results in underreaction to news, as implied by the primacy bias. Importantly, the conditions for over- and underreaction in complex environments are factors that may vary exogenously with features of the decision context.

In addition, the recency-primacy shift in the weighting of past experiences may occur as a result of the application of different dynamic algorithms of learning. A hybrid model of learning by Barberis and Jin (2023) accounts for an interplay of model-free reinforcement learning and model-based learning. The study predicts both recency and primacy bias in the weighting of past experiences and specifies conditions under which each bias occurs. A primacy bias, as in the professional context, can arise under certain conditions, but only as a result of model-free reinforcement learning. Model-free reinforcement learning is governed by past experiences — if a subject has experienced a high reward of a given action in the past, they are more likely to repeat that action in the future (see, e.g., Sutton and Barto, 1998, for an overview). A distinguishing feature of model-free reinforcement learning is that subsequent choices depend on the past reward prediction error — the experienced reward of a past action relative to the anticipated reward. Importantly, whether or not reinforcement learning is at play may vary exogenously with features of the decision environment (Carlos Alos-Ferrer, 2023).

In the following, we discuss several tests of the above-mentioned cognitive mechanisms. In summary, we show that *attention and salience* explain the weighting of past experiences in the *private* context but not in the professional context. In contrast, the weighting function in the professional context depends on the prediction errors of past experiences and, thus, displays the traces of *model-free reinforcement learning*.

A. Salience and Attention in Private Context

First, we explore attention and salience as potential drivers of the context-dependence in experience effects. Attention varies systematically over time in our lab experiment and with tenure in our field setting. Junior finance professionals reportedly work excessively long hours compared to professionals in senior positions.[26] Similarly, participants in the "Forecast Value" treatment exhibit a similar time-trend of attention, spending more time early in the experiencing stage (see Figure D1, Panel A). Moreover, attention is often assumed to be driven by salience (see, e.g., Barber and Odean, 2007; Engelberg and Parsons, 2011), even for highly-attentive finance professionals (see Fedyk, 2024). In the following, we investigate whether these factors also contribute to context dependence in the weighting of past experiences.

We test for the role of attention for the professional-context experience effect in our *field setting* by examining variations in the effect around the career start. To this end, we assume that prospective finance professionals start paying more attention to financial markets before their actual career start, already at the time of their job market entry. Accordingly, if the primacy bias is driven by attention, we expect to observe the professional-context experience effect already in the years before their career start in the finance industry. Table III shows results when shifting the starting points of experiences to years prior to the actual career start in the finance industry (s). The table shows that the experience effects are smaller and not statistically significant when extending the experience window to include the job-market years before the actual entry in the finance industry. This result indicates that features unique to the professional occupation rather than excess attention to financial markets are what drives experience effects characterized by primacy bias.

In addition, our lab setting allows for a direct measure of attention. We use a proxy of attention

[26] For anecdotal evidence on excess working hours of junior finance professionals, see, e.g.: `https://www.wsj.com/finance/banking/bank-america-jpmorgan-overtime-work-hours-f9f204a7`

given by the viewing time of each observation during the experiencing stage. To test for attention as a driver of the primacy bias, we replace the weighting function from our main model specification with a weighting function which ranks experienced returns based on the subject's attention to it, from the return that attracts the least attention to the return that attracts the most attention. Table VI, Column (ii) shows that the estimated weighting function, which is illustrated in Figure 5, Panel b, is flat — the weights subjects put on the observations they pay most attention to is the same as the weight they put on observations that they pay the least attention to.

[Insert Table VI approximately here.]

Analogously, we can test for salience as a potential driver of the overweighting of primary returns. In our experiment, the salience of returns does not vary systematically, as returns are i.i.d. However, by design, cumulative returns are more likely to become salient in later periods of the experiencing stage. Table VI, Column (iii) and Figure 5, Panel c, illustrate the estimated salience-ranked weighting function of experienced returns. The corresponding within-subject salience ranks are calculated based on absolute returns. The results show no significance of the salience-based weighting of experienced returns at the 10% level.

Previous studies suggest that investment decisions in the private context are more likely governed by salience-channeled attention because retail investors face more binding attention limitations than institutional investors (see, e.g., Barber and Odean, 2007; DellaVigna and Pollet, 2009). Consistent with this hypothesis, we find that, unlike the professional context, both attention and salience explain the weighting of past experiences in the private context (see Table VI, Columns (v) and (vi)). Interestingly, the difference between the private and professional contexts in our experiment cannot be explained by different incentives, availability to more resources or different time constraints, as suggested in previous studies (see Barber and Odean, 2007).

Our experimental findings on the impact of attention and salience in the professional context together with our field evidence challenge the notion that these factors drive the overweighting of primary returns. Absent salience-channeled attention, Ba et al. (2024) suggest that people form more unbiased mental representation of the decision environment. Instead, the study suggests that potential distortions in belief updating are rather overtaken by biases in the processing of information. In cases the information is complex, such as the information in both our field and lab setting, these distortions predict an underreaction, consistent with a primacy bias.

B. Reinforcement Learning in Professional Context

To trace potential reliance on model-free reinforcement learning in the professional context, we use a direct measure of the prediction error of past experiences in the "Forecast Value" treatment. For each subject we rank the experienced returns according to the subject's corresponding prediction error — the difference between the actual return and the subject's respective incentivized forecast. Table VI, Column (i) displays the parameter estimates and Figure 5, Panel a illustrates the corresponding weighting function. We find that participants significantly underweight past returns that they failed to predict (both extremely overestimated and extremely underestimated returns) and overweight past returns when their forecasts were relatively accurate. This result contradicts Bayesian learning, which relies solely on the sequence of experienced returns rather than prediction errors (Barberis and Jin, 2023).[27] The finding is consistent with an experimental study by Carlos Alos-Ferrer (2023), showing that the reliance on model-free reinforcement learning varies within-subject and is induced by hit-or-miss incentives, such as the experience-stage incentives in the "Forecast Value" treatment.

To provide an intuition why reinforcement learning is related to greater overweighting of primary returns relative to model-based (Bayesian) learning, it is useful to think of the different types of weighting functions in terms of the dynamics of beliefs that they imply. Section V illustrates the expected dynamics of beliefs conditional on different parameter combinations of the weighting function of past returns. Taken together, a slower speed of convergence is observed for weighting functions that are characteristic of reinforcement learning, whether due to a weaker recency bias or the presence of a primacy bias. This is because model-free reinforcement learning is inherently inefficient — it updates beliefs only based on direct experience rather than drawing from the full range of available data. Furthermore, environmental factors that limit random exploration of new actions can reinforce early choices to the point that the decision-maker sticks with an early action for a substantial period of time, hence exhibiting a primacy bias. Barberis and Jin (2023) demonstrate

[27] While it could be argued that latent subjective beliefs are inherently driven solely by model-based approaches grounded in probabilistic reasoning, incentivized forecasts can result from a decision-making process similar to the allocation decision. As such forecasts may be governed both by model-based and model-free reinforcement learning.

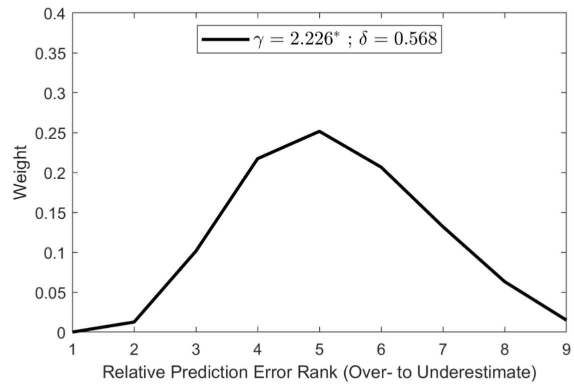

a. Reinforcement Learning

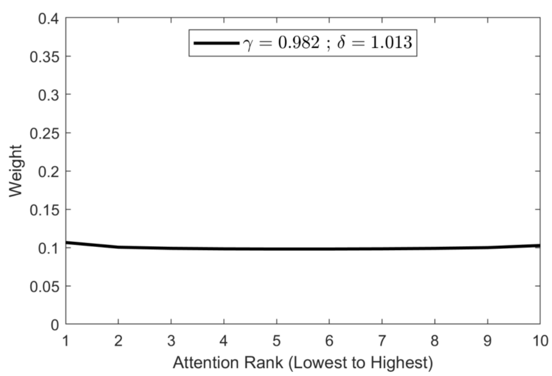

b. Attention

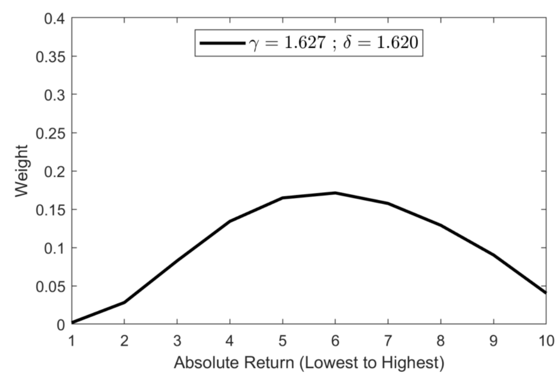

c. Salience

Figure 5. Professional context: comparison of potential cognitive mechanisms. This figure plots the estimates of alternative weighting functions of experienced past returns in the "Forecast Value" treatment. Unlike the main regressions, the returns are not sorted along the time dimension. The parameter estimates γ and δ of the Prelec (1998) function are provided in the legend. To test for the impact of *reinforcement learning*, Panel (a) plots the weighting of returns depending on their prediction-error rank (1 indicates most extreme overestimation/negative surprise; 10 indicates most extreme underestimation/positive surprise). The prediction error is computed as the relative surprise (in percent) relative to the subject's forecast. To test for the effect of *attention* on the weighting of experienced returns, Panel (b) sorts experienced returns according to their corresponding viewing time (1 indicates lowest/fastest; 10 indicates highest/slowest). To test for the effect of *salience*, in Panel (c) we estimate the weighting of returns based on their salience rank. The salience definition with the highest explanatory power defines salience as the absolute return (1 indicates return closest to 0 — positive or negative; 10 indicates most extreme return). Significance levels of the parameter estimates are indicated as *** (p<0.01), ** (p<0.05), * (p<0.10).

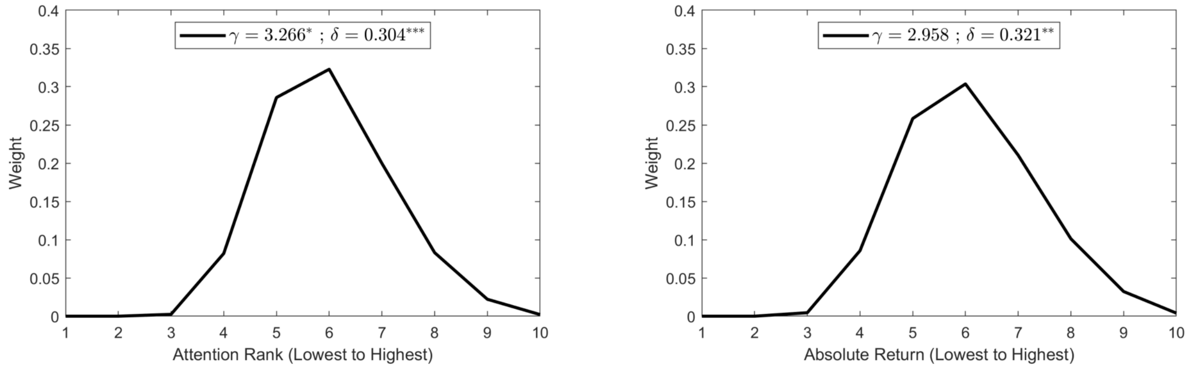

a. Attention b. Salience

Figure 6. Private context: comparison of potential cognitive mechanisms. This figure plots the estimates of alternative weighting functions of experienced past returns in the "Observe & Recall" treatment. Unlike the main regressions, the returns are not sorted along the time dimension. The parameter estimates γ and δ of the Prelec (1998) function are provided in the legend. To test for the effect of *attention* on the weighting of experienced returns, Panel (a) sorts experienced returns according to their corresponding viewing time (1 indicates lowest/fastest; 10 indicates highest/slowest). To test for the effect of *salience*, in Panel (b) we estimate the weighting of returns based on their salience rank. The salience definition with the highest explanatory power defines salience as the absolute return (1 indicates return closest to 0 — positive or negative; 10 indicates most extreme return). Significance levels of the parameter estimates are indicated as *** (p<0.01), ** (p<0.05), * (p<0.10).

that exploration levels below a certain threshold result in a functional form of the weighting function similar to our estimate in the "Forecast Value" treatment. Among the discovered factors that limit exploration, are high frequency of feedback and high cognitive load (see, e.g., Brown et al., 2022). These factors notably distinguish the "Forecast Value" treatment from the "Observe & Recall" treatment, where no feedback is provided and cognitive load is perceived to be lower. In the field setting, additional factors that are characteristic of the professional environment may further contribute to higher cognitive load and the primacy bias in the field (e.g., a larger choice set than in the private context, see Brown et al., 2022).

V. Implications: Overreaction vs. Underreaction to News

Why is it important to know when a primacy bias or a recency bias occur? To understand the implications of primacy and recency bias, it is useful to think about the two biases in terms of the dynamics of beliefs that they imply.

Learning with a *primacy bias* implies coining one's beliefs to a few observations experienced early on and sticking with these beliefs for a long period of time, hence *underreacting* to new information. A primacy bias implies that the individual learns very *slowly* — their beliefs take a long time to converge to the base rate. How slowly exactly is illustrated in Figure 7. The figure shows the expected dynamics of weighted-average experienced returns of an individual who has experienced a crisis at the very beginning of their experiencing stage and exhibits the parameter estimates in our field study.[28] The figure shows that the primacy bias results in a systematic underestimation of the true base rate. The scope of the underestimation decreases gradually with time and would eventually be fully resolved given a long enough horizon. However, as the figure shows, a primacy bias as pronounced as in our field data, is unlikely to be fully resolved 45 years into one's career, which is when most finance professionals would have already exited the finance industry.

In contrast, learning with a *recency bias* implies overreacting to new information and *underweighting* distant past experiences to the point of almost completely ignoring them. In contrast to the primacy bias, a recency bias implies that the individual learns and un-learns very *fast*. In expectation, their beliefs align with the true base rate shortly after an exogenous crisis experience. However, they fluctuate excessively around it. Figure A2, Panel b, shows that based on the estimates in our setting, the recency bias in the private context causes around five times more belief fluctuation than the professional-context primacy bias in the months following the crisis experience. Similarly, a J-shaped weighting function, displayed in Panel d, implies around three times more belief fluctuations.

Both overreaction to new information, resulting from a recency bias, and the underreaction, driven by a primacy bias, imply that beliefs are inaccurate for a long time. If the individuals learned and made decisions *in isolation*, the temporary inaccuracies, caused by either one of the two biases, would vanish with experience ultimately leading to correct asymptotic beliefs. However, in our field setting, learning in isolation is an implausible assumption for most finance professionals. By the nature of their occupation, they make their forecasts observable to other investors, while the source of their inaccuracy remains largely private information. The latter hinders other in-

[28] The dynamics are similar for a stock-market boom of similar magnitude. Figure A2 shows that a substantial part of the professionals in our panel experience either a large positive or a large negative return in the first year of their tenure.

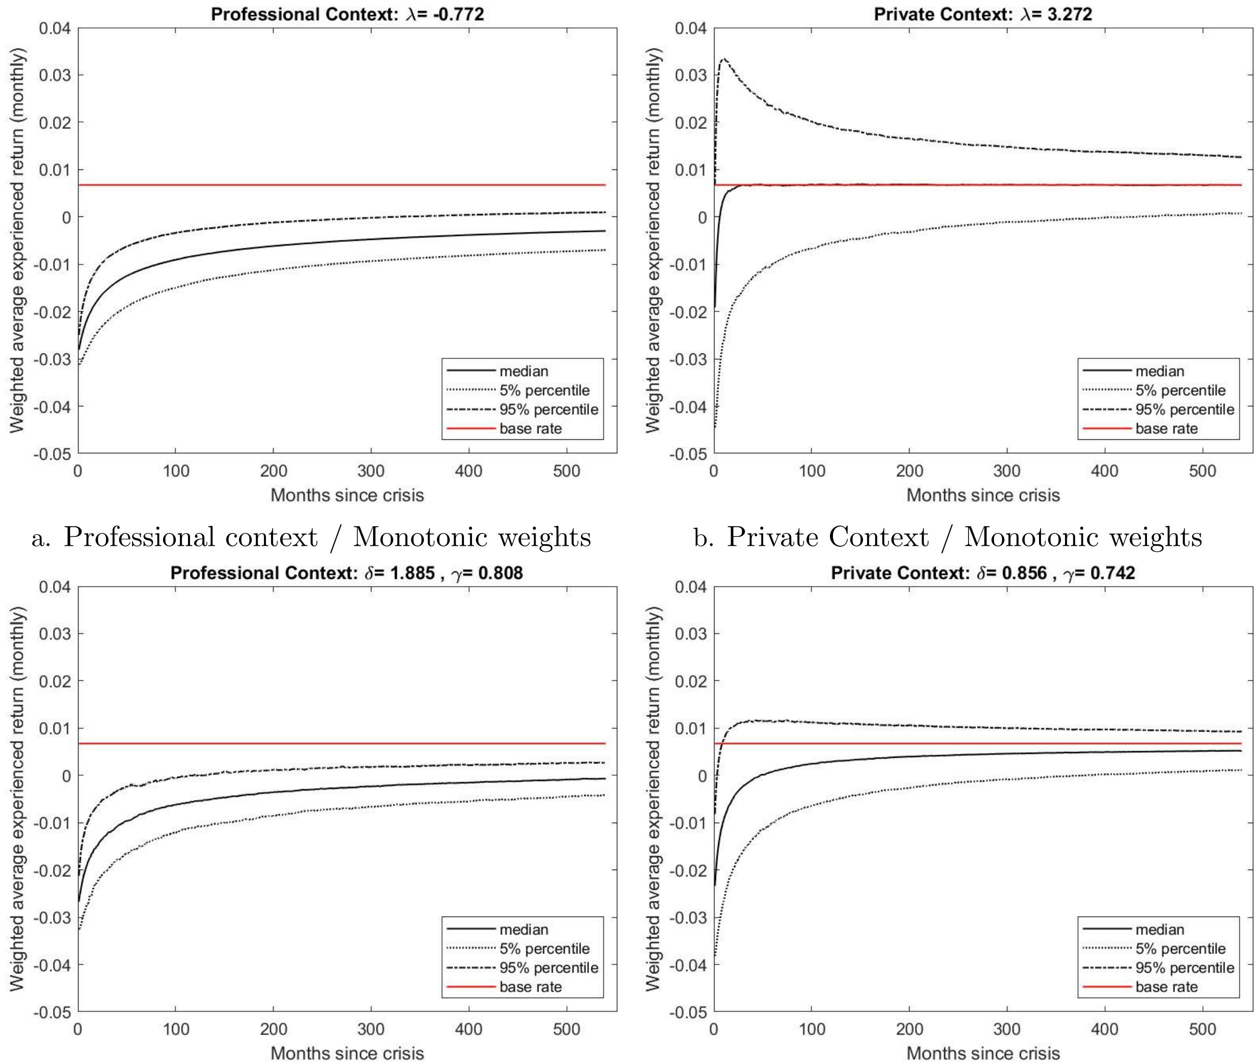

a. Professional context / Monotonic weights   b. Private Context / Monotonic weights

c. Professional context / Non-monotonic weights   d. Private context / Non-monotonic weights

Figure 7. Belief dynamics following a stock market crash experience. This figure illustrates the expected dynamics of weighted average experienced returns using the weighting parameters estimated in our field setting. At the beginning of the experiencing stage, we introduce an exogenous stock market crash of -30% p.a., similar to the financial market crash in 2008. The subsequent dynamics of the distribution of weighted average experienced returns are based on 1000 simulations of 45-years long stock-market return sequences. For the purposes of the simulations, we assume that stock market returns $R_i \sim N(0.084, 0.19^2)$ iid, following the historical sample distribution of the german DAX total return index. Panels (a) and (b) display the expected dynamics based on the estimated weighting parameter, λ, of monotonic weighting function by Malmendier and Nagel (2011). Panels (c) and (d) display the corresponding context-dependent dynamics based on the weighting parameters δ and γ of the non-monotonic weighting function by Prelec (1998).

vestors from accounting for any potential bias, when updating their own beliefs after observing the professionals' forecasts and recommendations (see, e.g., Hirshleifer et al., 2009, for social learning in financial markets). Bohren and Hauser (2021) analyze the dynamics of beliefs in *social learning* environments, like the real-life environment of the professionals in our field study. The study shows that, unlike individual learning in isolation, in social learning environments underreaction to news — such as that induced by primacy bias — can result in incorrect asymptotic learning.

VI. Conclusion

In this paper, we show how and why well-versed finance professionals from a market-moving survey in Germany extrapolate from the past stock market development that they have experienced over the course of their career. We distinguish between two domains of the past - a personal and a professional domain - and show that professionals extrapolate differently from the two domains. Whereas in the personal domain they tend to overweight the most recent episodes of the past, in the professional domain it is the primary episodes that are most formative.

Our results of a behavioral biases of financial professionals in the professional context have important implications for household finance as they may pose negative externalities on households' financial mistakes. Recently, European regulators have taken on the challenge of investor protection by restricting conflicts of interest in financial advice. The update to the Market in Financial Instrument Directions (MiFID II), which came into force in 2018, attempts to strengthen the trust in financial advisors by restricting conflicted sell-side advice. In the presence of behavioral biases of financial advisors, such regulatory measures can only be partially effective at best. Automated robo-advisors offered by FinTech companies may offer investors an advantage in this regard.

REFERENCES

Adam, Klaus, Dmitry Matveev, and Stefan Nagel, 2021, Do survey expectations of stock returns reflect risk-adjustment?, *Journal of Monetary Economics* 117, 723–740.

Afrouzi, Hassan, Spencer Y. Kwon, Augustin Landier, Yueran Ma, and David Thesmar, 2023, Overreaction in expectations: Evidence and theory, *The Quarterly Journal of Economics* 138, 1713–1764.

Akepanidtaworn, Klakow, Rick Di Mascio, Alex Imas, and Lawrence D.W. Schmidt, 2023, Selling fast and buying slow: Heuristics and trading performance of institutional investors, *The Journal of Finance* 78, 3055–3098.

Amromin, Gene, and Steven A. Sharpe, 2014, From the horse's mouth: Economic conditions and investor epxectations of risk and return, *Management Science* 60, 845–866.

Andersen, Steffen, Tobin Hanspal, and Kasper Meisner Nielsen, 2019, Once bitten, twice shy: The power of personal experiences in risk taking, *Journal of Financial Economics* 132, 97–117.

Andre, Peter, Philipp Schirmer, and Johannes Wohlfahrt, 2024, Mental models of the stock market, Available at SSRN: https://ssrn.com/abstract=4589777.

Ariely, Dan, and Gal Zauberman, 2000, On the making of the experience: The effect of breaking and combining experiences on their overall evaluation, *Journal of Behavioral Decision Making* 13, 219–232.

Azeredo da Silveira, Rava, Yeji Sung, and Michael Woodford, 2024, Optimally imprecise memory and biased forecasts, *American Economic Review* 114, 3075–3118.

Ba, Cuimin, J. Aislinn Bohren, and Alex Imas, 2024, Over- and underreaction to information, Available at SSRN: `https://ssrn.com/abstract=4274617`.

Barber, Brad M., and Terrance Odean, 2007, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *The Review of Financial Studies* 21, 785–818.

Barberis, Nicholas, and Lawrence Jin, 2023, Model-free and model-based learning as joint drivers of investor behavior, NBER Working Paper No. w31081.

Baucells, Manel, Martin Weber, and Frank Welfens, 2011, Reference-point formation and updating, *Management Science* 57, 506–519.

Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey, 2013, Managerial miscalibration, *Quarterly Journal of Economics* 128, 1547–1584.

Bohren, J. Aislinn, Josh Harscher, Alex Imas, Michael Ungeheuer, and Martin Weber, 2024, A cognitive foundation for perceiving uncertainty, NBER Working Paper No. w32149.

Bohren, J. Aislinn, and Daniel N. Hauser, 2021, Learning with heterogeneous misspecified model: Characterization and robustness, *Econometrics* 89, 3025–3077.

Bordalo, Pedro, John J. Conlon, Nicola Gennaioli, Spencer Y. Kwon, and Andrei Shleifer, 2023, Memory and probability, *The Quarterly Journal of Economics* 138, 265–311.

Bordalo, Pedro, Nicola Gennaioli, Giacomo Lanzani, and Andrei Shleifer, 2024, A cognitive theory of reasoning and choice, NBER Working Paper No. w33466.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2019, Diagnostic expectations and stock returns, *The Journal of Finance* 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, attention, and choice, *The Quarterly Journal of Economics* 135, 1399–1442.

Brown, Vanessa M., Michael N. Hallquist, Michael J. Frank, and Alexandre Y. Dombrovski, 2022, Humans adaptively resolve the explore-exploit dilemma under cognitive constraints: Evidence from a multi-armed bandit task, *Cognition* 229, 105233.

Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong, 2004, A cognitive hierarchy model of games, *Quarterly Journal of Economics* 119, 861–898.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2011, Robust inference with multiway clustering, *Journal of Business & Economic Statistics* 29, 238–249.

Carlos Alos-Ferrer, Michele Garagnani, 2023, Part-time bayesians: Incentives and behavioral heterogeneity in belief updating, *Management Science* 69, 5523–5542.

Cochrane, John, 2011, Presidential address: Discount rates, *Journal of Finance* 66, 1047–1108.

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michel André Maréchal, 2015, Evidence for countercyclical risk aversion: An experiment with financial professionals, *American Economic Review* 105, 860–885.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Micahel Weber, 2021, Exposure to grocery prices and inflation expectations, *Journal of Political Economy* 129.

Das, Sreyoshi, Camelia M. Kuhnen, and Stefan Nagel, 2020, Socioeconomic status and macroeconomic expectations, *The Review of Financial Studies* 33, 395–432.

DellaVigna, Stefano, and Joshua M. Pollet, 2009, Investor inattention and friday earnings announcements, *The Journal of Finance* 64, 709–749.

Dittmar, Amy, and Ran Duchin, 2016, Looking in the rearview mirror: The effect of managers' professional experience on corporate financial policy, *Review of Financial Studies* 29, 565–602.

Engelberg, Joseph E., and Christopher A. Parsons, 2011, The causal impact of media in financial markets, *The Journal of Finance* 66, 67–97.

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann, 2024a, Associative memory, beliefs and market interactions, *Journal of Financial Economics* 157, 103853.

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann, 2024b, Associative memory, beliefs and market interactions, *Journal of Financial Economics* 157, 103853.

Entorf, Horst, Anne Gross, and Christian Steiner, 2012, Business cycle forecasts and their implications for high frequency stock market returns, *Journal of Forecasting* 31, 1–14.

Fedyk, Anastassia, 2024, Front-page news: The effect of news positioning on financial markets, *The Journal of Finance* 79, 5–33.

Gilboa, Itzhak, and David Schmeidler, 1995, Case-based decision theory, *The Quarterly Journal of Economics* 110, 605–639.

Glaser, Markus, Zwetelina Iliewa, and Martin Weber, 2019, Thinking about prices versus thinking about returns in financial markets, *The Journal of Finance* 74, 2997–3039.

Gödker, Katrin, Peiran Jiao, and Paul Smeets, forthcoming, Investor memory, *Review of Financial Studies* .

Golub, Gene H., and Victor Pereyra, 1973, The differentiation of pseudo-inverses and nonlinear least squares problems whose variables separate, *SIAM Journal on Numerical Analysis* 10, 413–432.

Graeber, Thomas, Christopher Roth, and Florian Zimmermann, 2024, Stories, statistics, and memory, *The Quarterly Journal of Economics* qjae020.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *Review of Financial Studies* 27, 714–746.

Griffin, D., and A. Tversky, 1992, The weighing of evidence and the determinants of confidence, *Cognitive Psychology* 24, 411–435.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2018, Time varying risk aversion, *Journal of Financial Economics* 128, 403–421.

Hartzmark, Samuel, Samuel Hirshman, and Alex Imas, 2020, Ownership, learning and beliefs, *Working Paper* .

Heimer, Rawley, Zwetelina Iliewa, Alex Imas, and Martin Weber, 2025, Dynamic inconsistency in risky choice: Evidence from the lab and field, *American Economic Review* 115, 330–363.

Heimer, Rawley, and Alex Imas, 2022, Biased by choice: How financial constraints can reduce financial mistakes, *Review of Financial Studies* 35, 1643–1681.

Hertwig, Ralph, Greg Barron, Elke U. Weber, and Ido Erev, 2004, Decisions from experience and the effect of rare events in risky choice, *Psychological Science* 15, 534–539.

Hirshleifer, David, and Siew Hong Teoh, 2003, Herd behaviour and cascading in capital markets: A review and synthesis, *European Financial Management* 9, 25–66.

Hirshleifer, David, Sonya Seongyeon Lim, and Siew Hong Teoh, 2009, Driven to distraction: Extraneous events and underreaction to earnings news, *The Journal of Finance* 64, 2289–2325.

Horton, John J., David G. Rand, and Richard J. Zeckhauser, 2011, The online laboratory: Conducting experiments in a real labor market, *Experimental Economics* 14, 399–425.

Kahana, Michael Jacob, 2012, *Foundations of human memory* (OUP USA).

Kahneman, Daniel, Barbara L. Fredrickson, Charles A. Schreiber, and Donald A. Redelmeier, 1993, When more pain is preferred to less: Adding a better end, *Psychological Science* 4, 401–405.

Kahneman, Daniel, and G. Klein, 2009, Conditions for intuitive expertise: A failure to disagree, *American Psychologist* 64, 515–526.

Kaufmann, Christine, Martin Weber, and Emily Haisley, 2013, The role of experience sampling and graphical displays on one's investment risk appetite, *Management Science* 59, 323–340.

Kaustia, Markku, Heidi Laukkanen, and Vesa Puttonen, 2009, Should good stocks have high prices or high returns?, *Financial Analysts Journal* 65.

Ke, Shikun, 2024, Analysts' belief formation in their own words, SSRN Working Paper No. 5025830.

Kirchler, Michael, Florian Lindner, and Utz Weitzel, 2018, Rankings and risk-taking in the finance industry, *The Journal of Finance* 73, 2271–2302.

Kuchler, Theresa, and Basit Zafar, 2019, Personal experiences and expectations about aggregate outcomes, *The Journal of Finance* 74, 2491–2542.

Kuhnen, Camelia, 2015, Asymmetric learning from financial information, *The Journal of Finance* 70, 2029–2062.

Lamont, Owen A., 2002, Macroeconomic forecasts and microeconomic forecasters, *Journal of Economic Behavior & Organization* 48, 265–280.

Langer, Thomas, Rakesh Sarin, and Martin Weber, 2005, The retrospective evaluation of payment sequences: Duration neglect and peak-and-end effects, *Journal of Economic Behavior & Organization* 58, 157–175.

Laudenbach, Christine, Benjamin Loos, Jenny Pirschel, and Johannes Wohlfart, 2021, The trading response of individual investors to local bankruptcies, *Journal of Financial Econommics* 142, 928–953.

Laudenbach, Christine, Ulrike Malmendier, and Alexandra Niessen-Ruenzi, 2019, Emotional tagging and belief formation: The long-lasting effects of experiencing communism, *American Economic Association Papers and Proceedings* 109, 567–571.

Laudenbach, Christine, Ulrike Malmendier, and Alexandra Niessen-Ruenzi, forthcoming, The long-lasting effects of experiencing communism on attitudes towards financial markets, *The Journal of Finance* .

Linnainmaa, Juhani T., Brian T. Melzer, and Alessandro Previtero, 2021, The misguided beliefs of financial advisors, *The Journal of Finance* 76, 587–621.

Lu, Yan, Sugata Ray, and Melvyn Teo, 2016, Limited attention, marital events and hedge funds, *Journal of Financial Economics* 122, 607–624.

Malmendier, U., G. Tate, and J. Yan, 2011, Overconfidence and early-life experiences: The effect of managerial traits on corporate financial policies, *Journal of Finance* 66, 1687–1733.

Malmendier, Ulrike, 2021, Experience effects in finance: Foundations, applications, and future directions, *Review of Finance* 25, 1339–1363.

Malmendier, Ulrike, and Stefan Nagel, 2011, Depression babies: Do macroeconomic experiences affect risk taking?, *The Quarterly Journal of Economics* 126, 373–416.

Malmendier, Ulrike, and Stefan Nagel, 2016, Learning from inflation experiences, *The Quarterly Journal of Economics* 131, 53–87.

Malmendier, Ulrike, Stefan Nagel, and Zhen Yan, 2021, The making of hawks and doves, *Journal of Monetary Economics* 117, 19–42.

Malmendier, Ulrike, and Jessica A. Wachter, 2024, Memory of past experiences and economic decisions, in *The Oxford Handbook of Human Memory, Two Volume Pack: Foundations and Applications* (Oxford University Press).

Maynard, Douglas C., and Milton D. Hakel, 1997, Effects of objective and subjective task complexity on performance, *Human Performance* 10, 303–330.

Mullainathan, Sendhil, 2002, A memory-based model of bounded rationality, *The Quarterly Journal of Economics* 117, 735–774.

Nagel, Stefan, and Zhengyang Xu, 2022, Asset pricing with fading memory, *The Review of Financial Studies* 35, 2190–2245.

Nursimulu, Anjali, Ulrike Toepel, Peter L. Bossaerts, and Micah Murray, 2016, Spatiotemporal brain signatures of risk and reward, Available at SSRN: https://ssrn.com/abstract=2831371.

Prelec, Drazen, 1998, The probability weighting function, *Econometrica* 66, 497–527.

Ridout, Martin S., and Peter J. Diggle, 1991, Testing for random dropouts in repeated measurement data, *Biometrics* 47, 1617–1621.

Schoar, Antoinette, and Luo Zuo, 2017, Shaped by booms and busts: How the economy impacts ceo careers and management styles, *The Review of Financial Studies* 30, 1425–1456.

Simonsohn, Uri, Niklas Karlsson, George Loewenstein, and Dan Ariely, 2008, The tree of experience in the forest of information: Overweighing experience relative to observed information, *Games and Economic Behavior* 62, 263–286.

Snowberg, Erik, and Leeat Yariv, 2021, Testing the waters: Behavior across participant pools, *American Economic Review* 111, 687–719.

Sutton, Richard S., and Andrew G. Barto, 1998, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA).

Thaler, Richard H., and H. M. Shefrin, 1981, An economic theory of self-control, *Journal of Political Economy* 89, 392–406.

Tversky, Amos, and Daniel Kahneman, 1973, Availability: A heuristic for judging frequency and probability, *Cognitive Psychology* 5, 207–232.

Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and uncertainty* 5, 297–323.

Ungemach, Christoph, Nick Chater, and Neil Stewart, 2009, Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)?, *Psychological Science* 20, 473–479.

Wachter, Jessica A., and Michael Jacob Kahana, 2024, A retrieved-context theory of financial decisions, *The Quarterly Journal of Economics* 139, 1095–1147.

Zeileis, Achim, 2006, Object-oriented computation of sandwich estimators, *Journal of Statistical Software, Articles* 16, 1–16.

Table I
Descriptive Statistics

The panel data of forecasts runs from February 2003 to December 2017. Our sample includes all subjects who provided information on their career start in the finance industry or their birth year within the scope of the ZEW survey. Expected returns are semi-annual returns calculated from six months ahead expected DAX index level and the DAX level on the day of the response. Market data are from Datastream.

	N	Mean	Std.	10^{th} pct	Median	90^{th} pct
Birth year	323	1964.50	8.80	1952	1965	1976
Year of career start	330	1988.90	9.50	1975	1990	1999
Age at career start	323	24.30	4.10	19	25	29
# of survey responses per forecaster	330	71.20	54.30	10	57	159
# of participants per survey-wave	178	133.30	21.20	108	132	163
Tenure	23507	21.04	9.13	9	21	34
Expected semi-annual return	23507	2.85%	9.43%	-6.47%	3.18%	11.72%

Table II

Expected Returns and Experienced Returns Professional vs. Private Context

This table presents regressions of expected returns on experienced (excess) returns. Experienced returns are weighted average monthly DAX returns (in excess of the money-market rate). Columns (i) to (iii) show results for experienced returns in the professional context since the professional's career start in the finance industry. For private-context experiences in columns (iv) to (vi), we aggregate experienced returns since the "best starting age." The best starting ages result from the maximum-likelihood estimation. Expected returns are DAX return expectations winsorized at the 1% and the 99% level and converted to average monthly returns. All regressions include calendar effects (survey wave and day-of-week FE) and subject-level fixed effects. Columns (ii), (iii), (v), and (vi) include the professionals' economic expectations and their assessment of the current economic situation. *Panel A* provides results for the non-monotonic weighting function, *Panel B* shows results for the monotonic weighting function, which are both described in Section I.D. Standard errors are double-clustered. t-statistics or F-statistics are provided in parentheses. Significance levels are indicated as *** (p<0.01), ** (p<0.05), * (p<0.10).

Panel A: $y_{it} = \beta f(\gamma, \delta) + \vec{\zeta}^{\top} \vec{x}_{it} + \epsilon_{it}$

	Professional Context			Private Context		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
β	1.103***	1.058***	1.048***	0.525**	0.488	1.343**
	(2.852)	(2.771)	(2.731)	(2.100)	(1.512)	(2.160)
Hump-shape parameter (γ)	0.840	0.808	0.809	0.730	0.742	0.696
Primacy parameter (δ)	1.837	1.885	1.879	1.234	0.856	1.505
F-test ($\gamma = 1 \wedge \delta = 1$)	(4.242)**	(4.433)**	(4.270)**	(1.156)	(1.604)	(1.457)
Economic expectations		0.004***	0.004***		0.004***	0.004***
		(12.127)	(12.124)		(11.980)	(11.954)
Economic assessment		0.001***	0.001***		0.001***	0.001***
		(3.873)	(3.871)		(3.842)	(3.823)
Returns	raw	raw	excess	raw	raw	excess
Subject FE	yes	yes	yes	yes	yes	yes
Calendar FE	yes	yes	yes	yes	yes	yes
Adj. R^2	0.346	0.382	0.380	0.346	0.383	0.381
N	23507	23507	23507	22949	22949	22949
Estimated starting age				26	25	20

Panel B: $y_{it} = \beta f(\lambda) + \vec{\zeta}^{\top} \vec{x}_{it} + \epsilon_{it}$

	Professional Context			Private Context		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
β	1.078***	1.049**	1.024**	0.268*	0.244*	0.247*
	(2.700)	(2.353)	(2.324)	(1.868)	(1.916)	(1.897)
Recency parameter (λ)	-0.703***	-0.772***	-0.763***	1.243	3.272*	3.103*
	(-3.151)	(-3.258)	(-3.111)	(1.098)	(1.702)	(1.654)
Economic expectations		0.004***	0.004***		0.004***	0.004***
		(12.108)	(12.106)		(11.895)	(11.892)
Economic assessment		0.001***	0.001***		0.001***	0.001***
		(3.909)	(3.909)		(3.949)	(3.949)
Returns	raw	raw	excess	raw	raw	excess
Subject FE	yes	yes	yes	yes	yes	yes
Calendar FE	yes	yes	yes	yes	yes	yes
Adj. R^2	0.345	0.381	0.380	0.346	0.382	0.381
N	23507	23507	23507	22949	22949	22949
Estimated starting age				25	23	23

Table III
Expected Returns and Experienced Returns Around the Career Start

This table presents regressions of expected returns on experienced returns calculated using different starting points. Experienced returns are weighted average monthly DAX returns since the individual professional's career start in the finance industry, s, as well as since the year $s-2$, and delayed up to the year $s+2$. Expected returns are DAX return expectations winsorized at the 1% and the 99% level, and converted to average monthly returns. All regressions include subject-level controls and calendar effects (survey wave and day-of-week FE), as well as the professionals' economic expectations and their assessment of the current economic situation (in line with column (ii) in the main Table II. *Panel A* provides results for the non-monotonic weighting function, *Panel B* shows results for the monotonic weighting function, which are both described in Section I.D. Significance levels are indicated as *** (p<0.01), ** (p<0.05), * (p<0.10).

	Panel A: $y_{it} = \beta f(\gamma, \delta) + \vec{\zeta}^{\top} \vec{x}_{it} + \epsilon_{it}$				
	$s-2$	$s-1$	s	$s+1$	$s+2$
β	0.235	1.050**	1.058***	0.788***	0.461*
t-stat	(0.878)	(2.141)	(2.771)	(3.040)	(1.659)
Adj. R^2	0.380	0.382	0.382	0.381	0.381
N	23389	23437	23507	23504	23262
Economic exp. & assess.	yes	yes	yes	yes	yes
Subject FE	yes	yes	yes	yes	yes
Calendar FE	yes	yes	yes	yes	yes

	Panel B: $y_{it} = \beta f(\lambda) + \vec{\zeta}^{\top} \vec{x}_{it} + \epsilon_{it}$				
	$s-2$	$s-1$	s	$s+1$	$s+2$
β	0.385	0.779**	1.049**	0.602**	0.454*
t-stat	(0.838)	(2.386)	(2.353)	(2.494)	(1.955)
Adj. R^2	0.380	0.381	0.381	0.380	0.381
N	23389	23437	23507	23504	23262
Economic exp. & assess.	yes	yes	yes	yes	yes
Subject FE	yes	yes	yes	yes	yes
Calendar FE	yes	yes	yes	yes	yes

Table IV
Predictability of Future Returns with Experienced Returns

This table test the predictability of future returns with experienced returns. The dependent variable is the average monthly realized DAX return within the next six months. We focus on realized returns over a semi-annual horizon to ensure consistency with the forecast horizon of the ZEW survey (which is six months). Experienced returns are the weighted average monthly returns since the beginning of the professional's finance career weighted with $\lambda = -0.772$ (i.e., primacy in professional context), or since the age of 23 weighted with $\lambda = 3.272$ (i.e., recency in private context). The weighting parameters correspond to the estimates from Table II, Panel B, columns (ii) and (v) respectively for the different contexts. For each month, we calculate the average experienced return among all professionals who have already started their career (professionals are weighted equally). We exclude the periods before 1982 because we have less than 75 professionals who had started their career at this point. If a professional leaves the ZEW panel, we assume that he has left the finance industry. This is a plausible assumption, since the ZEW follows-up with professionals who change positions/employers and keeps them in the survey pool for as long as their occupation is in the field of finance. We report Newey-West standard errors to account for autocorrelation in the error term for up to 5 lags. t-statistics are provided in parentheses. Significance levels are indicated as *** (p<0.01), ** (p<0.05), * (p<0.10).

	(i)	(ii)	(iii)	(iv)	(v)	(vi)
(Sub)sample	02/1982-12/2017	02/1982-05/2000	06/2000-12/2017	02/1982-12/2017	02/1982-05/2000	06/2000-12/2017
Experienced return (professional context)	-1.715*	-0.660	-7.897***			
(Recency parameter $\hat{\lambda} = -0.772$)	(-1.693)	(-0.657)	(-3.029)			
Experienced return (private context)				-0.434	-0.638	-1.286*
(Recency parameter $\hat{\lambda} = 3.272$)				(-1.277)	(-1.596)	(-1.796)
Adj.R^2	0.022	0.002	0.151	0.012	0.033	0.055
N	417	208	209	417	208	209

Table V
Experienced Returns and Investment Choices in the Lab

This table presents regressions of investment choices on experienced returns by treatment. The dependent variable is the portfolio share of the risky asset. Experienced returns are the weighted average of the 10 sequentially displayed, randomly drawn observations from the distribution of the risky asset. We estimate the non-monotonic weighting function by Prelec (1998). Demographic control variables are included if indicated: age, gender, statistical skills (self-reported). t-statistics or F-statistics are provided in parentheses. Significance levels are indicated as *** (p<0.01), ** (p<0.05), * (p<0.10).

$$y_{it} = \beta f(\gamma, \delta) + \vec{\zeta}^\top \vec{x}_{it} + \epsilon_{it}$$

	Forecast Value		Observe&Recall		Forecast Sentiment	
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
β	1.669***	1.450***	1.172***	1.097***	0.700*	0.619*
	(4.476)	(4.322)	(3.376)	(3.237)	(1.865)	(1.844)
Hump-shape parameter (γ)	1.834*	1.912*	0.738	0.748	0.765	1.475
F-test ($\gamma = 1$)	(2.850)	(2.800)	(0.970)	(0.760)	(0.430)	(0.300)
Primacy parameter (δ)	1.319	1.290	0.922	0.897	0.994	0.332*
F-test ($\delta = 1$)	(0.610)	(0.470)	(0.070)	(0.140)	(0.000)	(3.060)
Risk preferences		0.035***		0.037***		0.020***
		(4.966)		(4.800)		(2.555)
Demographics	no	yes	no	yes	no	yes
Adj. R^2	0.080	0.178	0.028	0.127	0.003	0.029
N	233	233	266	266	242	242

Table VI
Cognitive Mechanism of Context-Dependence: Alternative Weighting Functions of Experienced Returns

This table presents non-linear regressions of investment choices in the lab on weighted average experienced returns using the Prelec (1998) weighting function in the "Forecast Value" and the "Observe & Recall" treatments. Unlike the main regressions, the returns are not sorted along the time dimension, but along different return rankings. The resulting weighting functions are illustrated in Figure 5. To test for the impact of *reinforcement learning*, in Column (i) experienced returns are sorted according to their prediction-error rank (1 indicates most extreme overestimation/negative surprise; 10 indicates most extreme underestimation/positive surprise). The prediction error is computed as the relative surprise (in percent) relative to the subject's forecast. To test for the effect of *attention* on the weighting of experienced returns, in Columns (ii) and (v) we sort experienced returns according to their corresponding viewing time (1 indicates lowest/fastest; 10 indicates highest/slowest). To test for the effect of *salience*, in Columns (iii) and (vi) we sort experienced returns according to their salience rank. The salience definition with the highest explanatory power defines salience as the absolute return (1 indicates closest to 0 — positive or negative; 10 indicates most extreme return). Significance levels are indicated as *** (p<0.01), ** (p<0.05), * (p<0.10).

$$y_{it} = \beta f(\gamma, \delta) + \vec{\zeta}^{\top} \vec{x}_{it} + \epsilon_{it}$$

	Forecast Value			Observe & Recall		
	Predict. Err (i)	Attention (ii)	Salience (iii)	Predict. Err (iv)	Attention (v)	Salience (vi)
β	1.535***	1.653***	1.484***	-	0.952***	0.978***
	(4.590)	(4.160)	(3.160)	-	(3.660)	(3.530)
Hump-shape parameter (γ)	2.226*	0.982	1.627	-	3.266*	2.958
F-test ($\gamma = 1$)	(2.730)	(0.010)	(0.249)	-	(3.230)	(1.810)
Primacy parameter (δ)	0.568	1.013	0.620	-	0.304***	0.321**
F-test ($\delta = 1$)	(2.510)	(0.000)	(0.348)	-	(15.490)	(5.120)
Adj. R^2	0.084	0.060	0.069	-	0.048	0.042
N	224	233	233	-	266	266

Online Appendix for “Experience Effects on Wall Street vs. Main Street: Field and Lab Evidence of Context Dependence”

Benjamin Christoffersen

Arvid Hoffmann

Zwetelina Iliewa

Lena Jaroszek

The contents of this Online Appendix are as follows:

Online Appendix A provides additional details about the field data.

Online Appendix B provides tests of potential endogenous attrition in the field data.

Online Appendix C contains the experimental instructions.

Online Appendix A. Field Data

Online Appendix AI. Survey Expectations

As dependent variable in our analyses, we examine professionals' six-month ahead expectations regarding the main German stock market index DAX. The question on the expected DAX returns was included in the survey in February 2003. We gather data until December 2017. Specifically, the numerical question on DAX expectations reads:

"I expect the DAX in 6 months to be at ... points."

We convert the responses to the above survey question into expected stock market returns using the DAX opening level on the day of the response.[29] If the response is submitted during a bank holiday or during the weekend, we use the preceding trading day's closing level. Daily market data is obtained from Thomson Reuters' Datastream. To prevent extreme outliers from affecting our results, we winsorize expected returns at the 1% and 99% levels. To have all data on the same frequency, we convert semi-annual return expectations to average monthly returns.

Online Appendix AII. Experiences

Our measure of stock market experiences follows Malmendier and Nagel (2011), in that information on subjects' past investments is unavailable. Instead, the measure relies on the period the subjects have personally experienced. For our data as well, we cannot know which kind of first-hand experiences our professionals have from their private investments or their professional activity. For example, in their professional context, they might not trade at all. If they trade, they are not constrained to long-only investments in stocks but may also take short positions or positions in other asset classes. However, all surveyed professionals follow the DAX development closely; otherwise they would neither be invited to the survey nor feel proficient in providing their

[29] It is common practice in real-world surveys to elicit expected price/index levels and convert them to expected returns (see, e.g., Livingston Survey of the Federal Reserve Bank of Philadelphia). However, asking for expected prices instead of expected returns induces more pessimistic expectations and amplifies the recency bias, as Glaser et al. (2019) experimentally show based on the subject pool of the ZEW Financial Market Survey. For the study at hand, as the questionnaire design is the same for all subjects, we conjecture that both effects are orthogonal to the experienced returns. The type of recency bias documented in the study concerns extrapolation from commonly used charts of past performance and constitutes a time-fixed effect which we control for in this study.

forecasts. Since previous research suggests that less direct involvement weakens the impact of experiential learning (Andersen et al., 2019; Simonsohn et al., 2008), our results likely understate the effects of experienced returns on expectations.

Table AI below provides the definitions of the variables that we draw from our survey of finance professionals.

Table AI
List of Variables

Expected returns	Expected returns are DAX return expectations elicited on a monthly basis. The survey question is given as follows: *"I expect the DAX in 6 months to be at ... points."* DAX level forecasts are converted into expected returns based on the DAX 30 performance index opening value on the day of the response submission. If the response was submitted on a bank holiday or a weekend we use the last available DAX level (closing level on the last trading day). Market data is downloaded from Datastream. Expected returns are winsorized at 1% and 99% levels. The 6 month horizon return expectations are converted to average monthly returns.
experienced return, $f\left(s_i, t; \vec{\theta}\right)$	experienced returns are the weighted average monthly DAX returns. We consider experienced returns in two different contexts: private context (i.e. since particular age, which we estimate to best fit the data) and professional context (i.e. since the beginning of the career in finance). We use two weighting functions for the weighted experienced returns. The first is a monotonic weighting function following Malmendier and Nagel (2011). The second is a non-monotonic weighting function suggested by Prelec (1998).
Economic expectations	Responses to the question: "In the mid-term the economic situation in Germany will: ... " *improve* = 1, *stay the same* = 0, or *worsen* = −1
Economic assessment	Responses to the question: "I assess the current economic situation in Germany as: ... " *good* = 1, *normal* = 0 or *bad* = −1

Online Appendix AIII. Professionals' Careers and Selection into the Sample

Figure A1 shows the distribution of finance professionals' age at their career start in the finance industry. The average professional in our sample was born in 1965 and entered the finance industry at the age of 24. The distribution of age at career start peaks twice - around the ages of 20 and 26. Career starts around age 20 are plausible for subjects who completed an apprenticeship, whereas a career starting around the age of 26 suggests a university education. The large variation of age at career start allows us to empirically distinguish between the private and the professional context.

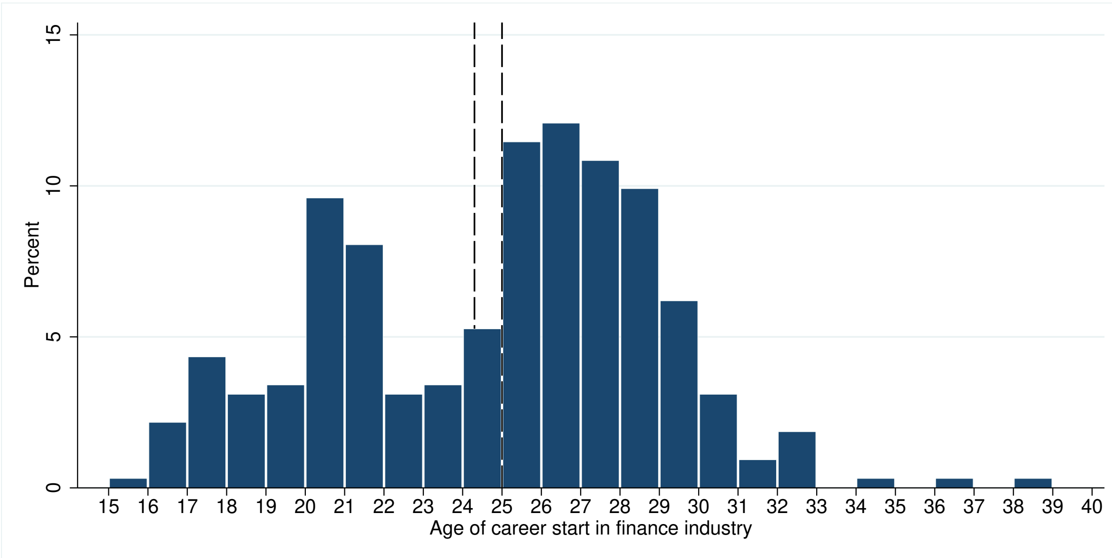

Figure A1. Sample distribution of age of professionals' career start in finance. This figure displays the distribution of ages at career start among the 323 professionals for whom both the year of birth and the year of career start is know. We use the difference between the year of career start and the year of birth as a proxy for age at career start. The reference lines indicate the sample average (24.3) and median (25).

The data on career starts come from several sources: The first source is the ZEW survey registration form, which requires all professionals who joined the panel in 2010 or later to report their birth year and the year in which they started their career in the finance industry. The other two data sources are supplementary questionnaires from 2003 and 2006 soliciting the same information. We exclude observations with conflicting information, as this suggests that the timing of the career start is ill-defined. We further exclude one professional, who reportedly entered the finance industry before the age of 14, which is likely incorrect. Since professionals report the year, but not the month of their career start, we disregard experiences from the reported career start year and gather data from January 1st of the year following the reported career start year. This approach avoids any endogeneity that would arise if stock market returns that were experienced just before the career start were correlated with unobserved factors related to the subject's decision to enter the finance industry. For the empirical analyses, we require all subjects to have at least two years of experience in the finance industry.

Figure A2 shows market returns at different cohorts' career start in the finance industry. The figure plots the number of career starts over time, where the green (red) bars indicate positive (negative) market returns in the year of the career start. To provide evidence against endogenous selection into the finance industry, i.e., into the sample, we examine the correlation between the number of career starts and preceeding DAX returns.

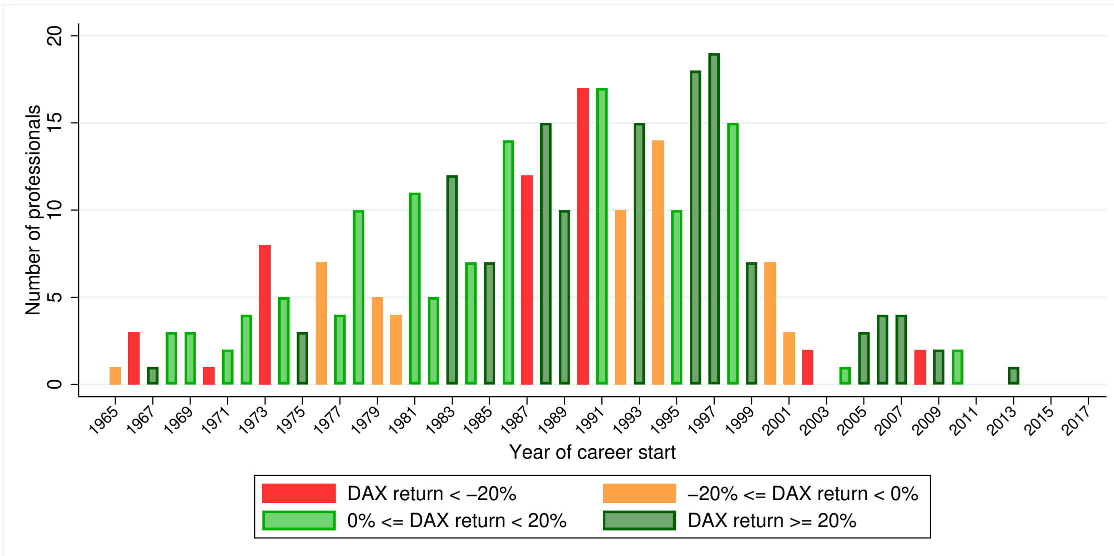

Figure A2. Timing of professionals' career start in finance. The bars show how many professionals in the ZEW survey started their career in the finance industry in a given year. The bars' color indicates the DAX returns in the respective years. The correlation of the number career starts and the contemporaneous DAX return is 0.19 (p-value of .192). The correlation of the number of career starts and the average DAX return over the last five years is 0.33 (p-value of .031).

Figure A2 reflects that our panel covers a broad range of professional cohorts with diverse early career experience. Importantly, many subjects at their career start have experienced a crisis (e.g., the 1973-74 stock market crash, the Black Monday stock market crash of 1987, the early 1990s recession, or the collapse of the dot-com bubble of 2000-2003) or a boom of similar magnitude (e.g., the boom years 1993, or 1997). The number of professionals from the ZEW panel starting their careers increases until 1997 and drops thereafter. This increase is in line with the historical growth of the finance industry during this period. The drop after 1997 can be attributed to the ZEW survey targeting professionals with a long tenure in the finance industry. The figure shows a positive correlation between the number of career starts and the average DAX return over the three-year preceding horizon. The positive correlation is around 0.2 but not statistically significant at the 10% level. A positive correlation might indicate that the demand or the supply of finance positions is pro-cyclical. It does not mean, however, that the decision to start a career in finance is related to the subsequently experienced returns. This form of potential endogeneity is precluded in efficient financial markets where past returns are no predictor of future returns.

Online Appendix AIV. External Validity and Attention to the ZEW Index Announcements

Figure A3 shows that the age structure in the ZEW survey resembles the age structure in the overall German finance industry.

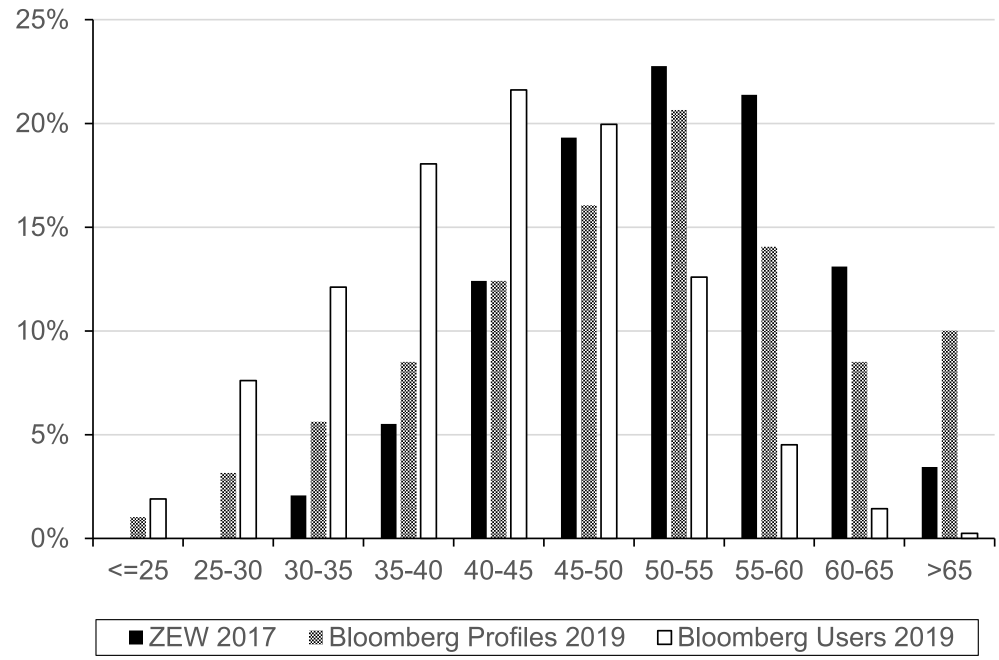

Figure A3. Participants' age-cohort distribution. The figure compares the ZEW survey professionals' age cohort distributions to those of German finance professionals for whom Bloomberg provides biographical data. Professionals' age cohorts are depicted on the x-axis. The y-axis shows the share of professionals in the respective age cohort in the three samples. Patterned bars show the distribution of all 1458 German finance professionals for whom Bloomberg provides biographical data. Framed bars show the distribution of 421 German finance professionals who have a biography and furthermore use Bloomberg as a financial trading platform. Solid black bars show the distribution of 145 finance professionals with birth year information in the ZEW survey.

In Figure A3, the data for the age distribution in the German finance industry stems from biographic profiles on Bloomberg. Specifically, we compare the ZEW survey professionals' five-year age-cohort distributions as of the information from 2017 to those of German finance professionals for whom Bloomberg provides biographical data as of 2019. Bloomberg provides biographical data for 1458 German finance professionals overall (patterned bars). A sub-sample of 421 professionals have a biography and furthermore use Bloomberg as a financial trading platform (framed bars). This latter group of professionals who also use the platform for trading are younger and without requirements on their tenure. The distribution of 145 finance professionals with age information in the ZEW survey (solid bars) compares better to the overall sample of Bloomberg users (patterned bars) because the ZEW survey specifically targets professionals with long tenure in the finance industry. Comparing ZEW and Bloomberg profile data, the patterns in retirement or job switches out of the finance industry appear very similar.

The scheduled announcements of the ZEW survey results attract the attention of institutional investors, as indicated by Bloomberg attention data. Bloomberg measures the attention of institutional investors (Bloomberg users) by tracking their alert subscriptions for the press release of the survey results. Figure A4 shows that among the institutional investors who are interested in the German market, 96.7% have set up an alert for the ZEW survey results.

[Insert Figure A4 approximately here.]

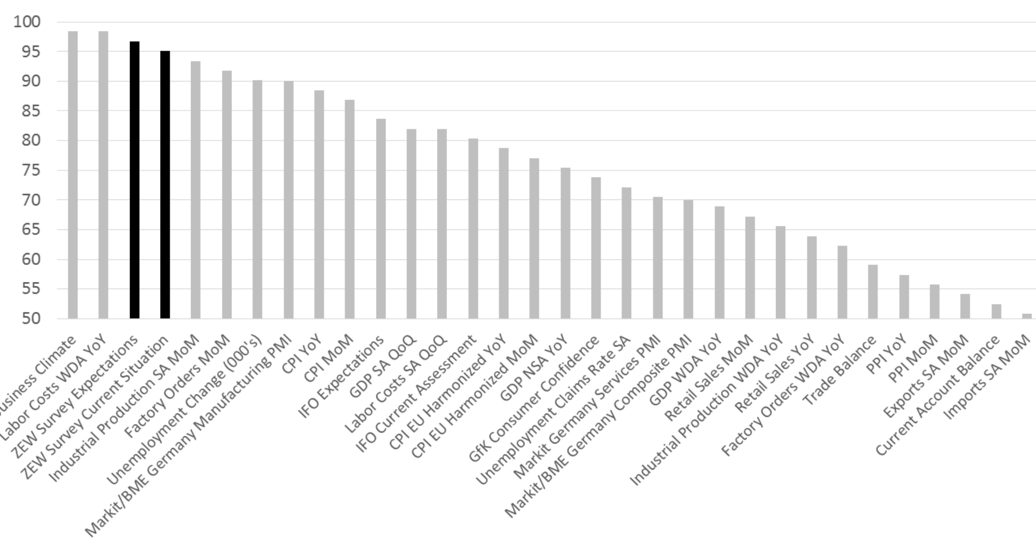

Figure A4. Institutional investors' attention to the ZEW Financial Market Survey. This figure displays international institutional investors' attention to the ZEW Financial Market Survey in comparison to other scheduled-release indicators for Germany. Attention is measured by the Bloomberg relevance index. It measures the percentage of alert subscriptions for the respective indicator among all Bloomberg users who have subscribed to any alerts regarding Germany. Displayed are scheduled-release indicators with a relevance index above 50. Bloomberg data is as of January 21st, 2019.

Online Appendix AV. Weighting Functions

A non-monotonic weighting function, e.g., U-shaped (J-shaped), would arise if both early and recent experiences are more formative than intermediary episodes (or if the weight on early or recent experiences is much more pronounced relative to the remaining horizon). To allow comparisons to the earlier experience effect literature, we also use the monotonic weighting function suggested by Malmendier and Nagel (2011). It determines whether early or recent experiences have a higher weight in forming expectations.

We consider for each professional a personally experienced period of stock market returns, specifically the returns of the German stock market index DAX. We aggregate the individual time-series of past stock market returns for each professional into their experienced returns using two weighting functions: a non-monotonic weighting function suggested by Prelec (1998), and a monotonic weighting function following Malmendier and Nagel (2011). In particular, experienced returns, $f\left(s_i, t; \vec{\theta}\right)$, are a weighted average of the returns over the respective experienced period. Hence, f is a function that depends on the time t, the monthly stock market returns R_t, and the starting time s_i of subject i. Furthermore, f depends on the non-linear parameters in $\vec{\theta}$.

The non-monotonic weighting function we use was suggested by Prelec (1998) for rank-dependent weights and was applied by Baucells et al. (2011) for time-dependent weights:

$$\begin{aligned}
f\left(s_i, t; \vec{\theta}\right) &= \sum_{j=1}^{t-s_i} \bar{w}_j(t-s_i; \vec{\theta}) R_{s_i+j-1} && t > s_i \\
\bar{w}_j(n; \vec{\theta}) &= \tilde{w}_j(n; \vec{\theta}) - \tilde{w}_{j-1}(n; \vec{\theta}) && 1 \leq j \leq n \\
\tilde{w}_j(n; \vec{\theta}) &= \begin{cases} 0 & j = 0 \\ \exp\left(-\left(-\log(j/n)\right)^{\gamma} / \delta\right) & 0 < j < n \\ 1 & j = n \end{cases}
\end{aligned} \tag{3}$$

where R_t is the stock market return in month t and $\vec{\theta}$ is defined by two non-linear parameters: δ, the *primacy parameter*, determines the relevance of primary relative to recent episodes analogously to the parameter of the monotonic function. γ, the *hump-shape parameter*, determines the relative relevance of intermediate episodes compared to the two extreme periods (primary and recent episodes) jointly.[30]

For the monotonic weighting function, we follow Malmendier and Nagel (2011):

$$
\begin{aligned}
f\left(s_i, t; \vec{\theta}\right) &= \sum_{j=1}^{t-s_i} w_j(t - s_i; \vec{\theta}) R_{s_i+j-1} && t > s_i \\
w_j(n; \vec{\theta}) &= \frac{j^{\lambda}}{\sum_{k=1}^{n} k^{\lambda}} && 1 \leq j \leq n
\end{aligned} \tag{4}
$$

where R_t is the stock market return in month t, and $\vec{\theta}$ comprises of only one parameter λ, the *recency parameter*, which determines the difference between the primary episodes of the experienced period and the most recent episodes.

[30] Note that there are other weighting functions with only one parameter that allow U-shaped (J-shaped) weighting functions (e.g., the weighting function in Tversky and Kahneman (1992)). However, we favored the two-parameter function suggested by Prelec (1998), because it allows for an overestimation of early experienced market returns relative to recently experienced returns.

Online Appendix B. Endogenous Attrition

For a formal analysis of endogenous attrition, we write n for the number of professionals with an expectation about the future stock return at each time t which we denote by $Y_{it} \in \mathbb{R}$. The ZEW survey runs from time $1, \ldots, T$. The professionals participate in the survey in some interval $\{t_i^{(L)}, t_i^{(L)} + 1, \ldots, t_i^{(U)}\}$. However, a professional may choose to not reply to the survey and we let $D_{it} \in \{0, 1\}$ be an indicator which is zero when we do not observe the professional's expectation. We further consider the covariates in the model (such as the professional's experience, economic assessments, or fixed effects) which we denote by $\vec{x}_{it}$.

The question is whether the professional's survey expectation Y_{it} is missing completely at random, missing at random, or not missing at random. One may suspect that people may not reply to the survey e.g., when they are busy due to poor performance of their firm in adverse time periods and that such subjects may have a more pessimistic expectation of the stock index. Such effects may be unrelated to the covariates in our model and thus problematic. On the other hand, some may not reply e.g., as they are on holiday which should be unrelated to their expectations.

We will test whether subjects who do not reply are not different in their previous observed responses with a test like the parametric approach suggested by Ridout and Diggle (1991). We use this test because it has the advantage that it easily allows to account for other covariates. We will cover the method in the sequel.

Each subject i has c_i uninterrupted sequences of observed responses to the DAX expectation question (a consecutive period where $D_{it} = 1$) and a dropout for a period between each sequence (a consecutive period where $D_{it} = 0$). Further, we let m_{ik} be the length of the kth uninterrupted sequence of observed expectations by subject i. Then we define the sets

$$R_j = \{(i, k) : m_{ik} \geq j\}$$

$$r_j = \{(i, k) : m_{ik} = j\} \subseteq R_j$$

as respectively the set uninterrupted sequences of observed expectations with j or more observed expectations and the set with exactly j observed expectations. The question we pose is whether there is a partial association between the observed expected stock return and whether we observe the next expectation by comparing the $R_j \setminus r_j$ observations to the r_j observations. To do so, we estimate a conditional logistic regression of the form

$$\operatorname{logit}\left(E(D_{i,t+1} \mid \vec{x}_{it})\right) = \alpha_{J(i,t)} + \vec{\gamma}^\top \vec{x}_{it} + \beta y_{it} \tag{5}$$

where $J(i,t) : \{1,\ldots,n\} \times \{1,\ldots,T\} \to \{1,\ldots,T\}$ is a map from the time t observed expectation of subject i to the length of the sequence of consecutive observed expectations of subject i at time t. E.g., if we have observed 3 consecutive expectations from subject i at time t then $J(i,t) = 3$. The hypothesis we are interested in is $H_0 : \beta = 0$. Evidence against this hypothesis suggests that expectations are not missing at random. We will also allow for a more general association with the model

$$\operatorname{logit}\left(E(D_{i,t+1} \mid \vec{x}_{it})\right) = \alpha_{J(i,t)} + \vec{\gamma}^\top \vec{x}_{it} + g(y_{it}; \vec{\beta}) \tag{6}$$

were g is a natural cubic spline using 5 degrees of freedom. The null hypothesis in this case is $H_0 : \vec{\beta} = \vec{0}$. As in our main regression, we winsorize the expected stock return. The covariates that we will use are the same as in our main regression. Further, we also include a so-called tensor product spline between the career start year and the experience length in months at time t of subject i. This way we will be able to estimate an effect similar to the weighted past experienced returns which we use in our main regression, since the weighted past experienced return is a function of the experience length and the career start year.[31]We use natural cubic spline with 3 degrees of freedom for each of the two marginal splines in the tensor product spline.

The likelihood ratio tests are shown in Table BI. We only have weak evidence of a partial association between the last observed expectations and whether we observe the next in spline models (i) and (ii) in Panel B in Table BI. This is mainly driven by a higher tendency to reply when expectations in the previous period were high as shown in Figure B1. However, we cannot reject that there is no association after including the other controls. Thus, these results provide no evidence that expectations are not missing at random.

[31] We do not include this tensor product spline in the conditional logistic regression models with subject fixed effects for numerical reasons.

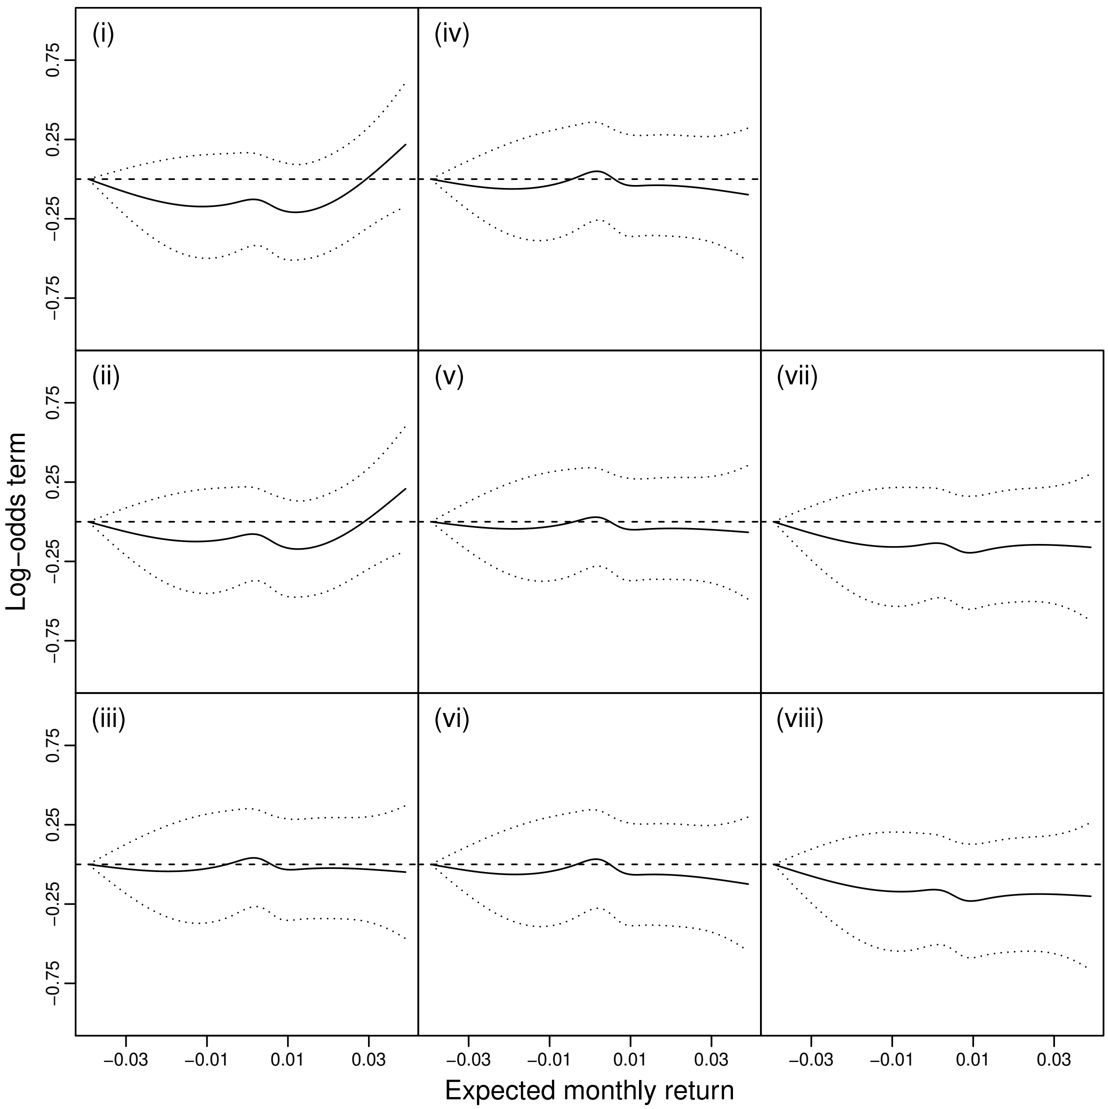

Figure B1. Estimated splines. This figure illustrates the estimated splines from the model in Equation (6). The x-axis shows the winsorized expected monthly return of the previous month and the y-axis shows the estimated partial effect on the log-odds scale of whether we observe the next expectations. Dotted lines are 95 pct. confidence bounds. The labels in the top left corresponds to Table BI. The narrow confidence bounds on the left-hand side are due to the boundary constraints of the natural cubic spline.

Table BI
Potential Endogenous Attrition

Panel A shows the likelihood ratio tests for $H_0 : \beta = 0$ in the models with a linear association shown in Equation (5) and Panel B shows the likelihood ratio tests for $H_0 : \vec{\beta} = \vec{0}$ for the models shown in Equation (6). Respondent level controls in columns (iv) and (vi) are described in Table AI. Columns (v), (vi) and (viii) include the subjects' economic expectations and their assessment of the current economic situation. Columns (vii) and (viii) include subject-level fixed effects. Columns (iii) to (viii) include calendar effects (survey wave and day-of-week FE). Columns (ii) to (viii) include a spline for the career start year, and columns (ii) to (vi) include a tensor-product spline for the present experience length and career start.

	Panel A							
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
Test statistic	0.33	0.05	0.05	0.14	0.20	0.49	0.44	1.09
Df	1	1	1	1	1	1	1	1
p-value	0.5637	0.8212	0.8285	0.7102	0.6549	0.4858	0.5051	0.2964
Subject-level controls	no	no	no	yes	no	yes	no	no
Economic assessments	no	no	no	no	yes	yes	no	yes
Subject FE	no	no	no	no	no	no	yes	yes
Calendar FE	no	no	yes	yes	yes	yes	yes	yes
Career start	no	yes	yes	yes	yes	yes	yes	yes
Experience $\times$ career start	no	yes	yes	yes	yes	yes	no	no

	Panel B							
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
Test statistic	12.29**	9.83*	2.05	3.54	2.16	3.77	1.66	2.42
Df	5	5	5	5	5	5	5	5
p-value	0.0310	0.0803	0.8422	0.6176	0.8270	0.5837	0.8943	0.7890
Subject-level controls	no	no	no	yes	no	yes	no	no
Economic assessments	no	no	no	no	yes	yes	no	yes
Subject FE	no	no	no	no	no	no	yes	yes
Calendar FE	no	no	yes	yes	yes	yes	yes	yes
Career start	no	yes	yes	yes	yes	yes	yes	yes
Experience $\times$ career start	no	yes	yes	yes	yes	yes	no	no

Online Appendix C. Experimental Instructions

Online Appendix CI. Attention Checks (All Treatments)

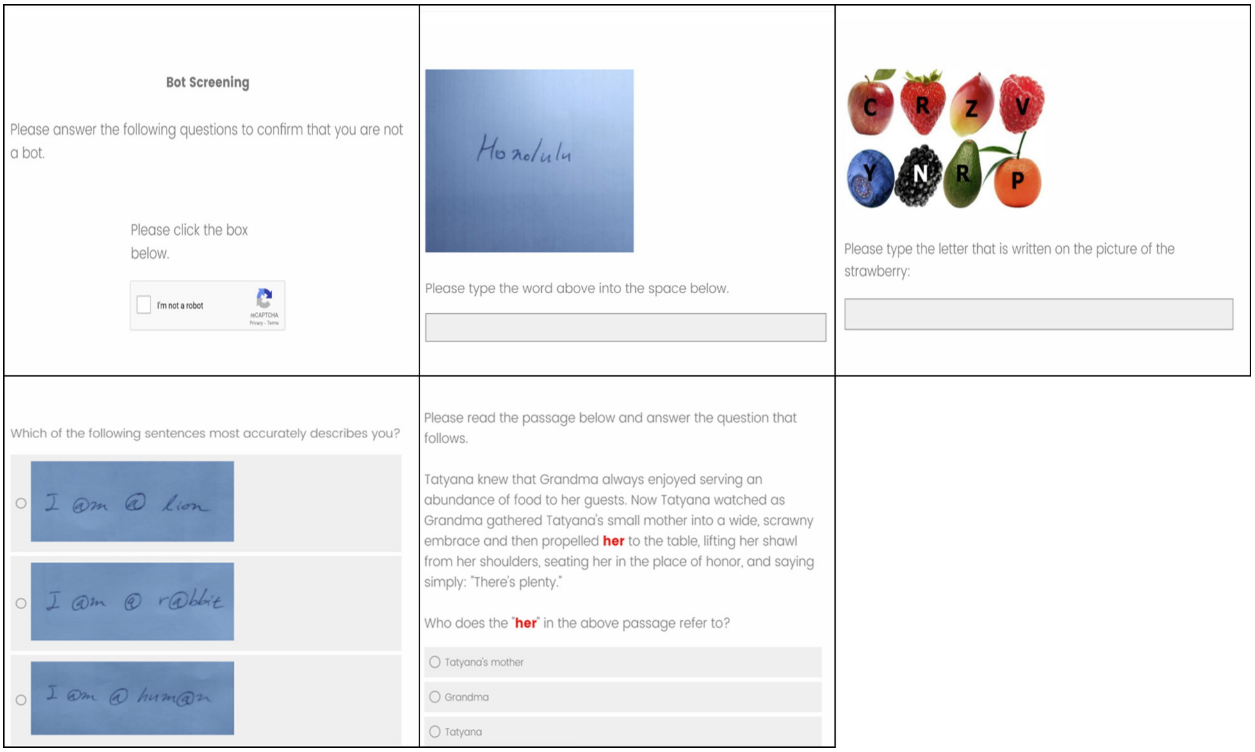

Figure D1. Screenshots of bot, attention and comprehension checks To participate in the experiment subjects need to pass all checks.

Online Appendix CII. Welcome Screen (All Treatments)

Dear Participant,

Thank you for taking the time to participate in our study. The aim of this study is to gain insights about financial forecasts and investment choices. The study takes approximately 6 minutes to complete. It consists of 2 parts:

- In Part 1 you will learn about the properties of a financial asset and will subsequently make an investment choice. - Part 2 is a questionnaire

You will receive a fix payment of $1.2 for participating. Additionally, you will earn a bonus.

To receive your remuneration you need to complete the entire study. At the end of the study, you will receive a completion code you need to submit on the Prolific platform.

Click "Next" to proceed.

Online Appendix CIII. Demographics (All Treatments)

Before we start with the experiment please answer the following questions. [Note: Your participation does not depend on your answers.]

- Your age
- Your gender
- Your highest level of education
- How would you rate your statistical knowledge? Please choose a category between 1 ("very bad") and 6 ("very good").

Online Appendix CIV. Past Performance in "Forecast Value" Treatment

`Screen: Instructions`

In this task, you will review the past performance of a financial asset over the last 10 periods. Each observation represents the simulated historical value of a $100 invested in this financial asset 10 periods ago. At the end of the task, you will decide whether or not to invest in the asset for the next period.

The future value of the financial asset will be simulated subsequently. Its characteristics will remain unchanged, and your investment decision will not influence its future value. Your investment bonus will depend on your investment decision and the future value of the financial asset. More details will be provided at the beginning of the investment task.

You will also have the chance to earn a forecasting bonus. As you observe the past performance, we will ask you to predict the value of the financial asset for the next period. At the end of the task, one of your forecasts will be selected randomly. If your forecast is close enough to the actual value, you will earn a bonus of 25 cents, otherwise your bonus will be 0.

Click "Start" to see the past performance of a $100 investment in the financial asset period by period.

`Screen: Period: -10 (ten periods ago)`

Value: 100

`Screen: Period: -9 (nine periods ago)`

Value: [the value is randomly drawn for each subject.]

What do you expect the value to be in the next period?

`Screen: Period: -8 (eight periods ago)`

Value: ...

What do you expect the value to be in the next period?

...

`Screen: Period 0 (today)`

Value: ...

Online Appendix CV. Past Performance in "Forecast Sentiment" Treatment

`Screen: Instructions`

In this task, you will review the past performance of a financial asset over the last 10 periods. Each observation represents the simulated historical value of a $100 invested in this financial asset 10 periods ago. At the end of the task, you will decide whether or not to invest in the asset for the next period.

The future value of the financial asset will be simulated subsequently. Its characteristics will remain unchanged, and your investment decision will not influence its future value. Your investment bonus will depend on your investment decision and the future value of the financial asset. More details will be provided at the beginning of the investment task.

You will also have the chance to earn a forecasting bonus. After this study, we will conduct a survey of 100 Prolific participants with some financial market experience. These participants will be shown the same asset performance you are about to see and asked to predict the asset's next value in each period. Your task is to guess what the average forecast of these participants will be as accurately as possible. At the end of the task, one of your guesses will be randomly selected. If your selected guess is close enough to the average forecast from that survey, you will earn a bonus of 25 cents, otherwise your bonus will be 0.

Click "Start" to see the past performance of a $100 investment in the financial asset period by period.

`Screen: Period: -10 (ten periods ago)`

Value: 100

`Screen: Period: -9 (nine periods ago)`

Value: [the value is randomly drawn for each subject.]

What is your guess for the average forecast of the next value?

`Screen: Period: -8 (eight periods ago)`

Value: ...

What is your guess for the average forecast of the next value?

...

`Screen: Period: 0 (today)`

Value: ...

Online Appendix CVI. Past Performance in "Observe & Recall" Treatment

`Screen: Instructions`

In this task, you will review the past performance of a financial asset over the last 10 periods. Each observation represents the simulated historical value of a $100 invested in this financial asset 10 periods ago. At the end of the task, you will decide whether or not to invest in the asset for the next period.

The future value of the financial asset will be simulated subsequently. Its characteristics will remain unchanged, and your investment decision will not influence its future value. Your investment bonus will depend on your investment decision and the future value of the financial asset. More details will be provided at the beginning of the investment task.

You will also have the chance to earn an attention bonus. As you observe the past performance period by period, we will ask you to recall and report the previous value of the financial asset from the last period. At the end of the task, one of your responses will be selected randomly. If your response is correct, you will earn a bonus of 25 cents, otherwise your bonus will be 0.

Click "Start" to see the past performance of a $100 investment in the financial asset period by period.

`Screen: Period: -10 (ten periods ago)`

Value: 100

`Screen: Period: -9 (nine periods ago)`

Value: ...

What was the value of the financial asset in the previous period?

`Screen: Period: -8 (eight periods ago)`

Value: ...

What was the value of the financial asset in the previous period?

...

`Screen:Period: 0 (today)`

Value: ...

What was the value of the financial asset in the previous period?

Online Appendix CVII. Investment Choice (All Treatments)

The current value of the financial asset is [VALUE].

Suppose you will receive $100 as an additional payment. You can choose to invest a certain amount of the $100 in the asset for the next period. You can invest any amount between $0 and $100. The remaining amount will stay in your current account, where it will not earn interest or incur any fees.

The future value of the financial asset will be simulated subsequently. Its characteristics will remain unchanged, and your investment decision will not influence its future value.

Your bonus payment will be 1% of the future value of your total assets. This includes:

- The future value of your investment in the financial asset (based on the amount you invested and the simulated future value of the asset).
- The amount you kept in your current account.

For example, let's denote the amount of your investment in the financial asset with $X and the actual future value of the financial asset one period from now with $F. Hence, you will purchase (X/[VALUE]) shares of the financial asset, which will be worth (X/[VALUE]]*F) one period from now. Any amount you do not invest (i.e. (100-X)) will stay in your current account.

Please indicate the amount (in $) you wish to invest in the financial asset.

[Slider from 0 to 100, no default]

Online Appendix CVIII. Questionnaire (All Treatments)

`Screen: Questionnaire 1/2`

Please answer the following questions about your perception of the task you just performed. Please choose a category between 1 ("Totally disagree") and 7 ("Totally agree").

- I found this to be a complex task.
- This task was mentally demanding.
- This task required a lot of thought and problem-solving.
- I found this to be a challenging task.

`Screen: Questionnaire 2/2`

Please answer the following questions:

- How do you see yourself: are you generally a person who is fully prepared to take risks or do you try to avoid taking risks? Please choose a category between 1 ("not at all willing to take risks") and 11 ("very willing to take risks").
- People may behave differently in different domains: How do you see your willingness to take risks in financial matters? Please choose a category between 1 ("not at all willing to take risks") and 11 ("very willing to take risks").