Network for Studies on Pensions, Aging and Retirement

Investor Memory

Katrin Gödker, Peiran Jiao, Paul Smeets

NETSPAR ACADEMIC SERIES

DP 03/2024-017

Investor Memory

Katrin Gödker, Peiran Jiao, and Paul Smeets*

March 11, 2024

ABSTRACT

We provide experimental evidence of a positive memory bias which affects individuals' beliefs, decisions to reinvest, and overconfidence in the stock market. Individuals over-remember positive investment outcomes of chosen assets and under-remember negative ones. Based on their memories, subjects form overly optimistic beliefs about their investment, reinvest too much, and become overconfident about their investment ability relative to others. We further provide evidence on motivation driving the memory bias. This positive memory bias offers a cognitive microfoundation for why gains weight more than losses when people learn from experiences. This helps reconciling various stylized facts in investor beliefs and behavior.

JEL classification: D01, G4

Keywords: Selective Memory, Beliefs, Motivated Reasoning, Investor Behavior, Experimental Finance

*Katrin Gödker: Bocconi University, CESifo, and IGIER, katrin.goedker@unibocconi.it; Peiran Jiao: School of Business and Economics, Maastricht University and Nuffield College, University of Oxford, p.jiao@maastrichtuniversity.nl; Paul Smeets: University of Amsterdam, p.m.a.smeets@uva.nl.

We thank Kai Barron, Roland Bénabou, Sabine Bernard, Pedro Bordalo, Luca Braghieri, Colin Camerer, Lauren Cohen, Sebastian Ebert, Nicola Gennaioli, Thomas Graeber, Robin Greenwood, Glenn Harrison, Alex Imas, Mats Köster, Florian Nagler, Salvatore Nunnari, Terrance Odean, Sébastien Pouget, Julien Sauvagnat, Nicolas Serrano Velarde, Theresa Spickers, Philipp Strack, and Florian Zimmermann for very helpful comments and suggestions and Bin Dong and JD O'Hea for great research assistance. We thank seminar participants at Caltech, DICE at Heinrich Heine University Duesseldorf, Erasmus School of Economics, Frankfurt Goethe University, Frankfurt School of Finance and Management, HEC Paris, Imperial College London, Lund University, Max Planck Institute for Research on Collective Goods, Stockholm Business School, Toulouse School of Economics, University of Alicante, University of Amsterdam, University of Copenhagen, University of Innsbruck, University of Mannheim, University of St. Gallen, Vrije Universiteit (VU) Amsterdam, and participants at AFA 2020 in San Diego, Belief Based Utility (BBU) Conference 2022 in Amsterdam, Boulder Summer Conference on Consumer Financial Decision Making, briq Institute Workshop on Beliefs 2021, CRC Workshop on Memory, Narratives and Belief Formation in Riederau 2021, Early Career Behavioral Economist (ECBE) Conference 2022 at briq institute, Early Career Women in Finance Conference (ECWFC) at Stanford 2020, and Helsinki Finance Summit on Investor Behavior 2019 for helpful discussions. Part of this paper was written while Katrin Gödker was visiting the Haas School of Business (UC Berkeley). Katrin Gödker thanks the Graduate School of Business, Economics, and Social Sciences at Hamburg University for financial support. Paul Smeets is supported by a VENI grant from the Netherlands Organisation for Scientific Research (NWO) under grant number 016.Veni.175.019. Peiran Jiao is supported by Horizon 2020 Marie Curie Individual Fellowship under project code 797396.

I. Introduction

How does memory influence individuals' investment decisions? Substantial evidence from psychology suggests that people view their past through rosy glasses (see Adler and Pansky (2020) for a review). They tend to recall their past experiences with a bias towards the positive. Particularly when emotionally significant, positive memories are long-lasting, are more readily recalled, and exert a greater influence on behavior compared to negative ones (e.g., Williams, Ford, and Kensinger (2022)). This paper's key contribution lies in presenting empirical evidence for such a systematic positive memory bias within the investment context and documenting its impact on individuals' beliefs and investment choices.

In three experiments, we document positive recall for investment outcomes. People over-remember the gains and under-remember the losses from their prior choices. After one week, subjects who chose to invest in a risky stock remember 23% more positive outcomes and 10% fewer negative outcomes than actually occurred. This positive memory bias offers a cognitive microfoundation for a well-documented pattern, namely that gains weight more than losses when people learn from experiences. This helps reconciling various stylized facts in investor beliefs and behavior. For example, investors tend to learn more from past successes than from past failures, leading to investor overconfidence (Gervais and Odean, 2001; Barber, Lee, Liu, Odean, and Zhang, 2020). Moreover, investors tend to repeat actions that generated gains in the past much more than they avoid actions that generated losses in the past, which has been documented in IPO subscriptions and stock repurchase behavior (Kaustia and Knüpfer, 2008; Strahilevitz, Odean, and Barber, 2011). In contrast to the widely acknowledged notion that prospective "losses loom larger than gains" in choices under risk (Kahneman and Tversky, 1979), individuals seem to react more strongly to gains than losses in response to *past* outcomes. We provide direct evidence for memory processes being a microfoundation for these central aspects in investor beliefs and behavior by showing that a positive memory bias causes individuals' to (i) form overly optimistic beliefs about the quality of their investments, (ii) reinvest too much, and (iii) become overconfident in their ability to invest.

Our experiments are designed to test for the existence of selective recall of investment outcomes and its causal effect on investor beliefs, reinvestment decisions, and overconfidence. Subjects choose between two investments with known probability to generate positive and negative outcomes. Subjects observe a series of investment outcomes, and we then elicit their memory of these outcomes. We exogenously manipulate the time span between observation of outcomes and memory elicitation to identify the memory effect. Subjects were asked to recall their outcomes either immediately or one week after the observation. We compare subjects' recall in the immediate vs. one week delay conditions. This allows us to identify the effect of memory processes, ruling out alternative factors that might also influence recall but are related to information acquisition or processing, such as attention (Barber and Odean, 2008; DellaVigna and Pollet, 2009), salience (Bordalo, Gennaioli, and Shleifer, 2012, 2013), or information ignorance at time of acquisition to maintain positive emotions (Kuhnen and Knutson, 2011).

To study how subjects' memory bias relates to their beliefs, decisions to reinvest, and overconfidence, we (i) measure their beliefs about the quality of their investment, (ii) ask them to make an additional investment decision, and (iii) elicit their confidence about their ability to select good investments relative to others. Our experimental design has three key features. First, by using experimental variations of time spans we can isolate a causal memory effect. Second, the lab allows us to elicit individual memory and beliefs in a controlled environment and in a direct and incentivized way. Third, to explore behavioral deviations from standard theory, we can compare subjects' elicited beliefs and choices to normative benchmarks.

We have four main findings (from Experiment 1 and 3). First, subjects exhibit a systematic memory bias. They recall more gains and fewer losses from their investments than actually occurred as well as compared to a control group. Second, the memory bias is related to subjects' belief formation. One week after observing the investment outcomes, subjects form too optimistic beliefs about the quality of their investments compared to an objective benchmark. Their beliefs are on average 8.16%-points too optimistic compared to the objective Bayesian posterior. Further, we show that the larger subjects' memory bias, the more optimistic their beliefs about the quality of their investments are. Third, the memory bias affects subjects' reinvestment decisions. One week after observing the investment outcomes, subjects are more likely to reinvest in the stock even when doing so leads to a lower expected return. Specifically, 54% of the subjects keep investing in the stock, although from a Bayesian perspective investing in an alternative asset was optimal. In addition, subjects with a larger memory bias are more likely to reinvest in such a suboptimal way. Fourth, the memory bias increases subjects' overconfidence. We elicit subjects' confidence in their investment ability relative to others. For this, we let subjects bet on the optimality of their own investment choice in a parimutuel betting market (Thaler and Ziemba, 1988; Enke, Graeber, and Oprea, 2023) after they observed the outcomes of their investment choice. Subjects exhibit more confidence in their ability to invest when they must rely on their memory from last week, even when they observe outcomes that suggest that they made a suboptimal investment choice (i.e., they exhibit overconfidence). Further, subjects' overconfidence increases with their size of the memory bias.

The novel contribution of these results is that we isolate memory as a channel through which experienced outcomes affect investor beliefs and behavior. We investigate the complete chain from previously experienced investment outcomes to memories of those outcomes, beliefs, subsequent investment choices, and overconfidence at the individual level.

Next, we explore the underlying mechanism of selective recall in the investment context. We implement two additional conditions in Experiment 2, which aim at exploring motivation as the underlying mechanism of the positive memory bias. Following existing research in psychology and behavioral economics, we study the role of motivated memory suppression and the role of active choice. First, theory (Bénabou and Tirole, 2002) and experimental evidence (Zimmermann, 2020) suggests that people suppress memory for motivational reasons. Following Zimmermann (2020), we identify this mechanism by exogenously manipulating the monetary benefit of reporting accurate memory in the memory elicitation task. Second, the active decision to do something and its consequences can be of high importance for one's ego or self-image and thus could give rise to motivated memory (Cioffi and Garner, 1996; Allison and Messick, 1988; Fazio, Chen, McDonel, and Sherman, 1982). In psychological experiments, Mather, Shafir, and Johnson (2000, 2003) document that people have the motivation to recall features of chosen options in a more positive light than those of assigned options. They call this choice-supportive memory. To isolate the role of active choice in our setting, we follow this design and compare subjects' recall in a condition, in which subjects receive a random endowment of a risky stock to our baseline condition in which subjects actively choose their investment.

In line with motivation being an underlying mechanism of the memory bias, we find that subjects suppress memories of negative outcomes rather than actually forgetting them and that their active choice plays an important role for the memory bias. In fact, subjects only exhibit a memory bias if they made an active investment choice, not when they were randomly endowed with an investment.

Our findings offer novel insights into well-established facts concerning investor beliefs and behavior. First, our study contributes to the literature on reinforcement learning in finance. This strand of literature explores the impact of experienced investment outcomes on future investment actions. For this, it is important to understand how investors learn, i.e., form and update beliefs, based on past investment outcomes. What reinforces investor actions might not be the objective outcomes from the past, but investors' own recollections of those outcomes. We identify selective recall as a crucial mechanism shaping how experienced outcomes impact investor beliefs and actions over time. In a seminal paper on reinforcement learning in finance, Kaustia and Knüpfer (2008) document that investors' likelihood to repeat an IPO subscription after previously generating relatively high returns is significantly larger than their likelihood to avoid an IPO subscription after previously generating relatively low returns. Thus, investors seem to be more sensitive to prior gains than to prior losses. This observation can be explained with our finding that investors tend to over-remember gains and to under-remember losses.

This complements existing work showing that individuals' learning about assets is determined by the way how the brain's valuation centers encode new information about assets (Kuhnen, Rudorf, and Weber, 2017). We document that the effect of prior outcomes on beliefs and investment decisions changes significantly and systematically over time due to memory processes in the brain. Because of the positive memory bias, previous gains from own investments receive asymmetrically and increasingly larger weights when time passes. That is, the distinction between how individuals respond to past gains versus losses depends on whether people make decisions immediately after experiencing the outcomes or after some delay.

Moreover, we examine how individuals decide whether to reinvest in a previously chosen stock, drawing a direct connection to the existing finance literature on stock repurchases. We link selective recall of own past realizations from a stock to the decision to repurchase and can explain a similarly observed asymmetric reaction to past gains versus losses. Existing research, such as Strahilevitz et al. (2011) and Du, Niessen-Ruenzi, and Odean (2019), has shown that both private and professional investors are more inclined to repurchase a stock if they had previously sold it for a gain rather than a loss. Importantly, Strahilevitz et al. (2011) show that investors' likelihood to repurchase a stock previously sold for a gain is considerably larger than their likelihood to avoid repurchasing a stock previously sold for an equal-sized loss. This can again be explained by investors over-remembering gains and under-remembering losses. We show that selective recall increases the relative number of recalled positive outcomes of stocks and decreases the relative number of recalled negative outcomes of stocks, which can explain this gain-loss asymmetry in stock repurchases, and in turn an increase in net repurchases of stocks. We further document that when time passes, individuals with a positive memory bias reinvest even if this reduces their expected return.

Second, we establish a direct link between selective memory and investor overconfidence, a well-documented phenomenon in behavioral finance where investors consistently perceive their own investment abilities to be superior, both absolutely and relative to others. Our research demonstrates that selective memory amplifies investor overconfidence. This has important implications for investor learning over time. A recent field study documents that such a positive memory bias generates persistent overconfidence of managers in settings with repeated feedback (Huffman, Raymond, and Shvets, 2022). Thus, our documented memory bias could explain why investors stay overconfident over time despite receiving repeated feedback in the stock market.

Further, in various settings – ranging from theoretical frameworks to experimental and actual markets – overconfident investors have been shown to trade more than is in their own best interest and contribute to price volatility (Odean, 1998; Deaves, Lüders, and Luo, 2009; Barber and Odean, 2001). We provide empirical evidence that selective recall increases investor overconfidence and, in turn, can contribute to high trading volume and price volatility.

Third, our findings directly connect to the fundamental question of whether individual trading is rational (Nicolosi, Peng, and Zhu, 2009; Seru, Shumway, and Stoffman, 2010; Barber et al., 2020). Rational models claim that "trading to learn" explains widespread speculative trading and excessive trading. The study by Barber et al. (2020) tests the proposition that individual investors rationally speculate as day traders to learn about their trading ability. Inconsistent with rational learning models, they find that most of the day trading volume in Taiwan is generated by traders with a history of losses. In other words, unsuccessful investors keep trading. This observation can be explained by investors' tendency to over-remember gains and to under-remember losses. Essentially, if investors forget their losses, they cannot learn from their failures.

Recent work emphasizes the crucial role of memory in how individuals form beliefs in economics (Afrouzi, Kwon, Landier, Ma, and Thesmar, 2023; Bénabou and Tirole, 2002; Bordalo, Conlon, Gennaioli, Kwon, and Shleifer, 2023; Fudenberg, Lanzani, and Strack, forthcoming). The basic intuition is that human memory has limitations, and with those limitations come trade-offs, namely prioritizing which information to hold on to, and which to let go of. Economic theory formalizes this limitation through selective recall. Selective recall can be categorized into two broad classes of recall processes: associativeness and motivation, also called positive memory bias (Fudenberg et al., forthcoming). Under associativeness, it is easier to recall experiences that are similar to features of the current decision problem, for example, past experiences that generated a similar mood (Matt, Vázquez, and Campbell, 1992). Associativeness can distort beliefs if it leads individuals to over- or underweight data relative to its true informativeness (Mullainathan, 2002). The simplest version of associativeness is similarity weighting (Bordalo et al., 2023). It can explain why individuals' beliefs in some situations overreact and in other situations underreact to available data, including the well-established representativeness heuristic (Bordalo, Coffman, Gennaioli, Schwerter, and Shleifer, 2021). Enke, Schwerter, and Zimmermann (forthcoming) provide experimental evidence for this. They show that subjects in their experiment remember past signals that get cued by the current context much more than those that get not cued. Subjects' beliefs systematically overreact to the cued signals. Wachter and Kahana (forthcoming) apply a version of associative memory to financial decision making. They explain three empirical patterns observed in financial markets: the role of early life experience in shaping investment choices, the rise of financial crises, and the impact of fear on asset allocation.

Under motivation, individuals are more likely to recall experiences that are more preferable, because they induce a larger utility. Motivation distorts beliefs by over-weighting experiences with high utility and under-weighting experiences with low utility (Bénabou and Tirole, 2002). Zimmermann (2020) finds that subjects who received poor scores on an IQ test are more likely to state that they "cannot recall" their test results, even though that answer is payoff dominated in the experiment. Chew, Huang, and Zhao (2020) even find that people selectively forget some negative events and create memories of fictitious positive ones. Motivated recall generates straightforwardly overly positive beliefs, including overconfidence (Bénabou and Tirole, 2002). We extend this strand of literature to the domain of financial decision making. So far, it is unclear whether motivated memory extends to the domain of financial decision making. Previous experimental studies document that in the domain of financial decision making, individuals form beliefs consistent with Bayesian updating (Barron, 2021) or even exhibit stronger updating in response to losses than to gains (Kuhnen, 2015). Yet, these studies do not allow for memory distortions. We show that when individuals have time to form memories about their investment outcomes, they form overly optimistic beliefs about the quality of their investments and become overconfident about their relative investment ability. This provides empirical evidence for memory being an important factor driving motivated belief formation and updating in the domain of financial decisions.

Our paper also contributes to the broader economic literature on motivated beliefs. This line of work argues that people form and update beliefs in order to maintain a positive self-view (Bénabou and Tirole, 2002; Köszegi, 2006). Recent experimental work provides evidence for overly optimistic belief formation about ego-relevant information on intelligence (Zimmermann, 2020), beauty (Eil and Rao, 2011), generosity (Saucet and Villeval, 2019; Di Tella, Perez-Truglia, Babino, and Sigman, 2015; Carlson, Maréchal, Oud, Fehr, and Crockett, 2020) and strategic interactions (Schwardmann and Van der Weele, 2019). A key question in the literature on motivated reasoning remains how such self-serving beliefs translate into behavior. Decades of psychology research points out that positivity of memory and beliefs is a powerful mechanism to maintain individuals' mental well-being (Taylor and Brown, 1988). However, if individuals make decisions based on biased memory, future behavior can be mistaken, as documented in our experiments. We show that individuals are more likely to make unprofitable reinvestment decisions because they have overly positive recollections of their past investment returns.

Moreover, our findings have implications for the extensive work on the two modes of learning – learning from description and learning from experience – and the resulting "description-experience gap" in choice under risk Hertwig, Barron, Weber, and Erev (2004); Hertwig and Erev (2009); de Palma, Abdellaoui, Attanasi, Ben-Akiva, Erev, Fehr-Duda, Fok, Fox, Hertwig, Picard, et al. (2014); Wulff, Mergenthaler-Canseco, and Hertwig (2018). The general idea of this strand of literature is that people can learn about the probabilistic outcomes of their actions in two ways: One involves consulting descriptions detailing the outcomes and probabilities associated with an action (such as reading about the potential price impact of a company's new CEO). The other entails personally experiencing the probabilistic outcomes of an action (for instance, experiencing price changes of a purchased stock). Our findings suggest that positive memory plays a crucial role when individuals learn from experience. This can be helpful for policy: De-biasing a positive memory bias – and resulting investor overconfidence or unprofitable re-purchase behavior – promises to be more effective in forms of reminders (with the aim to improve learning from experience) than informational campaign (with the aim to improve learning from description).

We further document the relevance of active choices for a positive memory bias. Individuals only suppress negative returns in their memory when they actively make their own investment decisions. This finding complements related work showing that endowment of a risky asset impacts how people react to information when updating their beliefs about the asset and how they make financial decisions – in the laboratory as well as in the field (Anagol, Balasubramaniam, and Ramadorai, 2018, 2021; Hartzmark, Hirshman, and Imas, 2021). Our results show that beyond ownership, active choices result in different memory processes, beliefs and subsequent behavior than passive ownership, i.e., delegation in financial markets.

II. Experimental Design

To investigate a memory bias for investment outcomes and its effect on beliefs and choices we use an experimental setup with (i) an investment decision that generates outcomes, (ii) exogenous variation in investment outcomes (positive and negative outcomes), (iii) direct memory elicitation, (iv) an experimental manipulation to isolate memory effects, and (v) subsequent incentivized belief and choice elicitation. In this section, we outline the features of our three experiments in detail.[1]

A. Basic Setting

All three experiments follow the same basic setup. First, we let subjects choose between two investment assets. After this initial investment decision, subjects observe the investment outcomes over the course of 12 periods. The outcome of each of the 12 periods is sequentially presented on a screen for 2 seconds. After subjects observe the outcomes, we ask them to recall their investment outcomes. This memory task is incentivized and is not announced beforehand.

Importantly, to clearly identify the effect memory processes have on subjects' recollection, we manipulate the time span between the observation phase and the memory elicitation, following a between-subject design. We randomly assign each subject to one and only one experimental condition for the duration of the whole experiment. In the *Delay* condition, subjects perform the memory task one week after the observation and in the *Immediate* condition, subjects perform the memory task immediately after they observed the investment outcomes. We compare subjects' responses and choices in the *Delay* group one week after the observation of outcomes vs. in the *Immediate* group immediately after the observation of outcomes. Comparisons between the *Delay* and *Immediate* treatments allow us to isolate memory effects from other factors such as attention (Barber and Odean, 2008; DellaVigna and Pollet, 2009), salience (Bordalo et al., 2012, 2013), or information ignorance at time of acquisition to maintain positive emotions (Kuhnen and Knutson, 2011), which might also influence subjects' recollection of outcomes, but are related to subjects' information acquisition and processing. Attention and salience should influence subjects' information acquisition and processing equally

[1] The experiment instructions are provided in Appendix A.

in both conditions, so any difference in participants' responses in the *Delay* vs. *Immediate* condition can be attributed to the time delay and the memory processes during the delay. Subjects in the *Delay* treatment observe investment outcomes in week t and perform the memory elicitation task in week $t+1$. In contrast, subjects in the *Immediate* condition perform all experimental tasks in the course of one session.

Subsequently, we elicit individual beliefs and choices in different forms across the three experiments to identify memory effects on (i) beliefs about the quality of investments (i.e., expectations about future investment outcomes), (ii) reinvestment decisions, and (iii) confidence in own investment ability, measured by betting behavior. In addition, we implement two additional treatment variations to explore the underlying mechanism of the memory bias.

All experiments were organized in two sessions. Subjects had to sign up for and participate in two experimental sessions, with one week in between irrespective of which treatment condition they were randomly assigned to, to avoid selection effects. Further, in order to reduce attrition, (i) the first and the second session took place on the same day of the week and at the same time (only one week later) for each subject, and (ii) subjects were reminded via email or Prolific message about the second session. Attrition was low at 4% in Experiment 1, 8.5% in Experiment 2 and 2.6% in Experiment 3.[2]

The experiments and their procedures were ethically approved by the Hamburg University Experimental Laboratory Committee (Experiment 1), the Office for Protection of Human Subjects (OPHS) at UC Berkeley, protocol number 2021-06-14375, and the Ethics Committee of the Centre for Experimental Social Sciences (CESS) at Oxford University, protocol number VE_0009 (Experiment 2), as well as Ethical Review Committee Inner City Faculties (ERCIC) of Maastricht University (Experiment 3). We obtained subjects' consent before participating in the experiment. Experiment 1 is programmed and conducted with z-Tree (Fischbacher, 2007) and Experiment 2 and 3 with o-Tree (Chen, Schonger, and Wickens, 2016). Experiment 3 is pre-registered at AsPredicted under ID 153791.

B. Experiment 1

In Experiment 1, subjects choose to invest either in a stock with risky outcomes (positive and negative outcomes) or in a bond with known safe outcomes (cf. Kuhnen, 2015). The risky stock may be good or bad, with equal probability. That is, the stock is either more likely to generate positive outcomes or more likely to generate negative outcomes. Good stocks have positive outcomes with a 60% probability and negative outcomes with a 40% probability each period. Bad stocks have positive outcomes with a 40% probability

[2]Attrition is not correlated with a specific experimental condition and the memory elicitation task in the second session of the *Delay* treatment was not announced beforehand.

and negative outcomes with a 60% probability.[3]

The positive outcomes of the stock are either 11, 13, or 15 EUR with equal probability and the negative outcomes are either -5, -3, or -1 EUR with equal probability. Thus, given that the outcome in a period is positive, it is randomly drawn from {11, 13, 15} with equal probability. Given that the outcome in a period is negative, it is randomly drawn from {-5, -3, -1} with equal probability. The determination of the outcomes is independent across periods. The bond has a certain outcome of 3.10 EUR each period and subjects start with an initial endowment of 60 EUR.[4]

To measure subjects' memory bias, we let subjects observe the stock outcomes and then elicit their memory of these outcomes. Subjects see the outcomes of the stock, irrespective of whether they chose to invest in the stock or bond. After subjects observe the outcomes, we ask them to recall how many positive and negative outcomes they observed and, more specifically, how often the stock paid 11, 13, 15, -5, -3, and -1 EUR, respectively.

Further, we aim at relating subjects' memory bias to their beliefs. We therefore elicit subjects' beliefs about the stock's quality, i.e., the chance of paying a positive outcome in the future, after they observed the investment outcomes. Note, subjects can form and update beliefs about the quality of the stock based on what they observe, because each stock has a specific, but constant probability of paying a positive versus a negative outcome in our setting. Initially, subjects do not know the quality of the stock. They start with a prior that the stock is either good or bad with equal probability. After observing the outcomes, subjects can make informed inferences about the stock's probability of being the good stock. A fully rational (Bayesian) subject counts the number of positive outcomes in the course of the 12 periods. The value of the objective Bayesian posterior after observing t outcomes can be calculated as:

$$\mu_t^G(h_t) = \frac{1}{1 + \frac{1-\mu_0^G}{\mu_0^G} * (\frac{\theta}{1-\theta})^{t-2n_t^+}} \tag{1}$$

where μ_0^G is 50% and indicates the prior that the stock is good; θ is 60%, the probability that a good stock generates the positive outcome in each period; t is the total number of observed outcomes; h_t represents the history of outcomes for t observations; n_t^+ represents the number of positive outcomes. This posterior serves as the benchmark for objectively correct beliefs in our experimental setting.

Finally, to investigate the consequence of a memory bias for subsequent reinvestment behavior, we ask subjects to make a second investment decision after observing the outcomes. They choose between investing in the stock they have observed or a different bond for further 12 periods (Table VIII). Risk-neutral subjects

[3]To make sure that subjects understand the distributions, we included a phase of experience sampling in Experiment 1 (for both the distribution of the good stock's and the bad stock's outcomes) which did not influence subjects payout.

[4]See Table VIII in the Appendix for an overview of subjects' investment options and choices.

should invest in the stock if it has a higher expected payoff than the bond. Note that for the post-observation round of investment decisions, the bond pays 5.10 EUR per period. If the stock is good, the expected outcome of the stock is higher (6.60 EUR per period) than the expected outcome of the bond. Yet, if the stock is bad, the expected outcome is lower (3.40 EUR per period). Given these expected outcomes, a risk-neutral Bayesian subject should always invest in the stock if there were more than 6 positive outcomes, which leads to a Bayesian posterior about the stock being good of 69.2% or greater.[5] Risk-averse subjects should require a higher posterior belief about the stock being good in order to choose the stock. Since we are particularly interested in subjects' decision to invest in the stock despite an objectively low Bayesian posterior probability (Section V), our results should hold for a range of reasonable risk attitude parameters.[6]

As indicated in Table I, in the *Delay* condition, we elicit subjects' subsequent beliefs and choices one week after the observation phase and in the *Immediate* conditions immediately after the observation phase. To avoid order effects, the order in which subjects perform the experimental tasks, i.e., the memory elicitation, belief elicitation, and the second investment task, is random. In addition, we tested whether subjects' memory bias, stated beliefs, and investment decisions after the observation are influenced by the task order and we do not find any significant impact on the memory bias, subjective beliefs, and suboptimal reinvesting (T-test, $p = 0.314$, $p = 0.890$, and $p = 0.428$, respectively).

To rule out timing effects, we create two *Immediate* conditions that are exactly the same except that they vary whether subjects perform the tasks in week t (*Immediate1* condition) or in week $t+1$ (*Immediate2* condition). Subjects were randomly assigned to either the *Immediate1* condition or *Immediate2* condition. In our main analyses, we pool the data of the *Immediate1* condition and the *Immediate2* condition and control for the session the subject participated in. In Section VII.D we provide robustness checks for the timing.

Further, a *NoRecall* condition, in which subjects do not perform any memory elicitation task, serves as a control condition. This condition allows us to avoid spill-over effects from the mere fact of asking subjects to recall investment outcomes (the memory elicitation) on subjects' stated beliefs and subsequent investment decisions. We do not find a significant difference in subjects' beliefs or investment decisions between the *Delay* and *NoRecall* condition (T-test, $p = 0.714$ and $p = 0.343$, respectively), suggesting that asking subjects to recall the stock's outcomes does not drive our results. Please refer to Section VII.E for a detailed description of the analyses.

All experimental tasks were incentivized. Subjects were paid a show-up fee for participating in the study.

[5] In general, a risk-neutral subject should invest in the stock if the observed outcomes lead to a posterior belief of the stock being good above 53.1% and otherwise invest in the bond.

[6] Typically, people are risk-averse even at low payoff levels (Holt and Laury, 2002, 2005). However, note that the small stakes in the laboratory setting might lead subjects to behave in a risk-neutral manner (Rabin, 2000).

We further randomly drew three participants from each session (with maximum 30 participants per session) who were paid a potential bonus based on their performance in one of the tasks. In the first investment choice, subjects could earn an initial endowment of 60 EUR plus either 37.20 EUR from investing in the bond or accumulated outcomes over 12 periods from investing in the stock. In the belief elicitation task, subjects were paid according to the accuracy of their probability estimates. We paid them 120 EUR for a probability estimate within +/-5 percent of the objective Bayesian value. In the memory elicitation task, subjects could earn 12 EUR for each correct answer provided, which could add up to 120 EUR if they answered all memory questions correctly. In the second investment choice, subjects could earn an initial endowment of 60 EUR plus either 61.20 EUR from investing in the bond or accumulated outcomes over 12 periods from investing in the stock.

The experiment was followed by a short questionnaire with background and control questions. We elicited subjects' general risk preferences (Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner, 2011), financial literacy, stock market participation, understanding of the risk-return relationship of investments, IQ, and mood (Watson, Clark, and Tellegen, 1988). Further, subjects were asked to indicate their age, gender, and highest level of education.

A total of 229 subjects participated in both sessions of the laboratory experiment, mostly business and economics students from the University of Hamburg. On average, subjects earned 16.90 EUR. For each subject, both sessions took about 45 minutes each.

C. Experiment 2

In Experiment 2, subjects choose between two equally risky stocks with positive and negative outcomes.[7] Before subjects choose between the two risky stocks, we show them the outcomes of each of the stocks from three previous periods. These three outcomes give subjects noisy information about the quality of the two stocks from which they can choose from.[8] The positive outcomes of the stock are either 2, 4, or 6 USD/GBP with equal probability and the negative outcomes are either -6, -4, or -2 USD/GBP with equal probability.[9]

We implement two additional between-subjects conditions in Experiment 2 which aim at exploring the underlying motivational mechanism of the investor memory bias, namely the role of motivated memory suppression and the role of active choice.[10] In each condition we additionally randomly assign subjects to an *Immediate* or *Delay* treatment group. Further, we want to rule out anticipatory utility affecting subjects'

[7]A choice between two stocks rules out any potential bias due to subjects' endogenous selection into investing in the stock.

[8]A fully rational (Bayesian) subject counts the number of positive outcomes of the two stocks and selects the stock with the most positive outcomes. In case both stocks generated the same amount of positive outcomes in previous periods, a Bayesian subject selects randomly.

[9]We switch to symmetric payoffs in order to rule out any difference in memory bias for positive and negative outcomes that is attributable to the size of the outcomes, for instance, because larger values are more salient.

[10]Table IX in Appendix E summarizes our treatment conditions as well as the number of subjects per treatment condition.

memory of outcomes from the first investment task.[11] Thus, in our *Baseline* condition, all subjects receive the payoff from the first investment decision soon after the session in week t and definitely before the session in week $t + 1$.

Memory suppression. Theory suggests that people suppress memory for motivational reasons (Bénabou and Tirole, 2002). These reasons can include the wish to maintain a positive self-image. Importantly, memory suppression is different from storing or coding the information wrongly in memory in the first place, in which case the true information could never be retrieved. Thus, in situations where people suppress memory for motivational reasons, sufficiently high incentives can induce people to access that information in memory nonetheless (Bénabou and Tirole, 2002). Zimmermann (2020) documents the existence of such motivational memory suppression for feedback about own intelligence.

To shed light on this motivational channel for the memory bias, we conduct a high stakes version of the *Baseline* condition. Following Zimmermann (2020), we manipulate the monetary benefit of reporting accurate memory in a *HighStake* condition. Compared to the *Baseline* condition, we increase the financial incentive for a correct answer in the memory elicitation task from 8 GBP/USD to 50 GBP/USD. A comparison between the *Baseline* condition and the *HighStake* condition tests whether the experienced outcomes are correctly encoded in subjects' memory but simply suppressed in the retrieval process. If so, high incentives would induce subjects to report memory more accurately.

This experimental manipulation is guided by the model laid out in Bénabou and Tirole (2002), in which agents trade off utility from accurate memory vs. utility from motivated memory. To be precise, we manipulate the utility of holding motivated memory in our experiment. In our experiment, subjects might have the tendency to hold motivated memory of investment outcomes (i.e., over-remember gains and under-remember losses), because that gives them additional utility, i.e., they like to over-remember gains and under-remember losses for motivational reasons. Thus, based on Bénabou and Tirole (2002), any increase in the incentive to report memory accurately while keeping the utility from motivated memory constant across two treatment groups should then reduce the elicited bias in recall.

Note that outside the lab, the utility from accurate memory as well as the utility from motivated memory is considerably higher. First, investment mistakes are expensive for individual investors (Beshears, Choi, Laibson, and Madrian, 2018). Second, the benefit from perceiving oneself as a good investor in the real world is larger than the benefit from perceiving oneself as a good investor in a lab experiment. Thus, we argue that the difference between 8 GBP/USD and 50 GBP/USD is designed to successfully shift the tradeoff between utility from accurate memory vs. utility from motivated memory (Bénabou and Tirole, 2002) in

[11] Theoretically, subjects who care about expected future utility flows exhibit higher utility now if they overestimate the probability that their investments pay off well in the future (Brunnermeier and Parker, 2005).

an experiment like ours. Yet, we expect that financial incentives to shift the tradeoff between utility from accurate memory vs. utility from motivated memory in the field, would need to be much larger.

Active choice. The fact whether an investment choice is made actively or not might affect how the investment outcomes are remembered. First, we expected that choosing actively increases the ego-relevance of investment outcomes. The active decision to do something and its consequences can be of high importance for one's self-image (Cioffi and Garner, 1996; Allison and Messick, 1988; Fazio et al., 1982). Outcomes from active choices not only mean changes in wealth, but also represent successes or failures of the investor. In psychological experiments, Mather et al. (2000, 2003) document that people have the motivation to recall features of chosen options in a more positive light than those of assigned options. They call this choice-supportive memory. To isolate the role of active choice in our setting, we follow this design and compare subjects' recall in a condition, in which subjects receive a random endowment of a risky stock to our baseline condition in which subjects actively choose their investment. Hence, we expect a stronger memory bias in an active rather than in a passive choice setting. To isolate this effect, we implement a *NoChoice* condition, in which subjects do not make any choice in the first investment task, but receive a randomly assigned risky stock from the two available options. To keep the level of information constant across the *Baseline* condition and the *NoChoice* condition, subjects in the *NoChoice* condition see the outcomes from three previous periods of the assigned risky stock. By comparing the *NoChoice* condition with the *Baseline* condition, in which subjects make an active investment decision, we can investigate the effect of active choice on the memory bias.

To relate subjects' memory bias to their beliefs, we elicit subjects' beliefs about the stock's quality, i.e., the chance of paying a positive outcome in the future, exactly as in Experiment 1. Further, subjects are asked to make a reinvestment decision, similar to Experiment 1. They choose between investing in the stock they have observed and investing in a new randomly drawn stock (from the same underlying processes, with a prior of being the good stock of 50%) for further 12 periods. Risk-neutral subjects should invest in the observed stock if it has a higher expected payoff than the new randomly assigned stock.

All experimental tasks were incentivized. Subjects were paid a show-up fee for participating in the study. In addition, each subject was paid a potential bonus based on their performance in one of the tasks. For each subject, the computer randomly decided which task determined his or her payment. In the first investment choice, subjects could earn the accumulated outcomes over 12 periods from investing. In the belief elicitation task, subjects were paid according to the accuracy of their probability estimates. We paid them 8 USD/GBP for a probability estimate within +/-5 percent of the objective Bayesian value. In the memory elicitation task, subjects could earn either 8 USD/GBP (*Baseline*) or 50 USD/GBP (*HighStake*) if they answered all memory questions correctly. In the second investment choice, subjects could earn the accumulated outcomes over 12 periods from investing.

The experiment was followed by a short questionnaire with background and control questions. We elicited subjects' self-deceptive enhancement, that is, the tendency for motivated self-enhancement (Paulhus, 1991), proneness to regret (Schwartz, Ward, Monterosso, Lyubomirsky, White, and Lehman, 2002), general risk preferences (Dohmen et al., 2011), financial literacy, as well as subjects' identification with the investment task. Further, subjects were asked to indicate their age, gender, and highest level of education.

A total of 498 subjects completed both sessions of the experiment. Experiment 2 was conducted online as a virtual lab experiment with the students subject pools at the Centre for Experimental Social Sciences (CESS) at Nuffield College, University of Oxford, and at the Experimental Social Science Laboratory (Xlab) of the University of California, Berkeley. The number of subjects who participated in Experiment 2 by lab can be found in Appendix E.[12] On average, subjects earned 15.50 GBP at CESS / 23 USD at Xlab. For each subject, the two sessions took about 30 minutes each.

D. *Experiment 3*

Experiment 3 follows the set-up of Experiment 2. Subjects choose between two stocks with risky outcomes (positive and negative outcomes). Before subjects choose between the two risky stocks, we show them the outcomes of each of the stocks from three previous periods. These three outcomes give subjects noisy information about the quality of the two stocks from which they can choose from. The positive outcomes of the stock are either 2, 4, or 6 GBP with equal probability and the negative outcomes are either -6, -4, or -2 GBP with equal probability.

In Experiment 3 we implement an additional *Reminder* condition to isolate memory effects. Subjects observe investment outcomes in week t and perform the second choice elicitation in week $t+1$. In contrast to the *Delay* condition, subjects do not have to rely on their memory while making their choice in week $t+1$, because we remind them of their investment outcomes by showing them the record of outcomes, on the same screen as their experimental choice task in week $t+1$. In the *Reminder* condition, full information on the investment outcomes is available to subjects. Comparing subjects' choices in the *Delay* condition to those in the *Reminder* condition allows us to isolate memory effects on decisions while keeping the timing constant (week $t+1$). This controls for potential factors in week t, right after the first investment choice is made, that might impact subjects' confidence and subsequent behavior, such as sentiment (Baker and Wurgler, 2006, 2007) or mood (Hirshleifer, Jiang, and DiGiovanni, 2020).

The aim of Experiment 3 is to study the consequence of biased memory for subjects' confidence in their

[12] We conducted one additional treatment with 154 subjects (77 in *Delay* and 77 in *Immediate*) that is not part of this paper, but is available upon request.

investment ability. For this, we let subjects bet on the optimality of their own investment choice in a 10-person parimutuel betting market (Thaler and Ziemba, 1988; Enke et al., 2023) after observing the ouctomes of their investment choice. Please find the detailed instructions in Appendix C. We define an investment choice, i.e., choosing one stock over the other, as optimal when it is expected payoff-maximizing.[13] In such betting markets, bettors use their private information to place their bet (Thaler and Ziemba, 1988). This betting market setting allows us to study subjects' belief about their ability to select good stocks (i) in an incentivized way, (ii) relative to other subjects in a market, and (iii) on the intensive margin. Subjects were informed that a cohort of 9 other subjects in the study completed the same investment choice between two risky stocks as they did and that the 10 subjects would be grouped together into a betting market to bet on the quality of their investment choice. Each subject is endowed with 100 points (ECUs). Their task is to decide how many of those 100 points (if any) to bet on the proposition that her own choice was optimal. This decision was implemented using a simple slider that ranged from 0 to 100, with no default value.

Following Enke et al. (2023), the payoff is calculated as follows. If a subject's investment decision was not optimal, all points bet are lost and the subject only keeps the remaining points she did not bet. If the subject's investment decision was optimal, the subject's payoff from betting is given by:

$$payoff_i = \frac{b_i}{\frac{\sum_{i=1}^{10} x_i b_i}{\sum_{i=1}^{10} b_i}} + (100 - b_i) \tag{2}$$

with b_i denoting points bet by subject i and x_i an indicator that equals 1 if the subject's investment choice was optimal. A subject whose investment choice was optimal is guaranteed to earn back at least what she bet, and the bonus is higher the more she bets, and the more points bet by other subjects whose investment decisions were not optimal.

Importantly, our measure has the advantage of having an objective benchmark of optimality required to define overconfidence. We can test how memory causally impacts subjects' overconfidence, i.e., subjects' betting on the optimality of their investment choice despite having observed investment outcomes that suggest that they made a suboptimal investment choice ("suboptimal betting"). Note, given the payoff scheme, subjects should bet 0 points on their choice in case they believe that they did not make the optimal choice in order to maximize their payoff.

The other experimental tasks were incentivized as well. Subjects were paid a show-up fee for participating in the study. In addition, each subject was paid a potential bonus based on their performance in one of the tasks. For each subject, the computer randomly decided which task determined his or her payment. In the first investment choice, subjects could earn the accumulated outcomes over 12 periods from investing. In

[13] In cases in which both stocks have the same expected payoff, both choices are defined as optimal.

the memory elicitation task, subjects could earn 8 GBP if they answered all memory questions correctly.

The experiment was followed by a short questionnaire with background and control questions. We elicited subjects' age, gender, and highest level of education.

A total of 487 subjects completed both sessions of the experiment. Experiment 3 was conducted online. Subjects were recruited via Prolific with restriction on being located in the United Kingdom, holding at least a high school diploma, and having an approval rate of 99-100 on Prolific. On average, subjects earned 10.47 GBP and the two sessions took in total 8 minutes.

E. Summary

Table I summarizes the treatment conditions which identify the memory bias and its effects on beliefs and choices, across the experiments. As outlined above, in order to explore the underlying mechanism of the memory bias, we implement two additional variations of the setup in Experiment 2 (*HighStake* and *NoChoice*), yet we still rely on the treatment conditions *Delay* and *Immediate* to identify the memory bias in each of these variations (see Appendix E).

In the main results section, we report findings from Experiment 1. In addition, Section VI reports findings from additional treatment conditions in Experiment 2 to explore the underlying mechanism of the memory bias, and Section V.B outlines our findings from additional treatment conditions in Experiment 3 to study memory effects on investor confidence. Note that our Experiments 2 and 3 successfully replicated our main results in Experiment 1. The replication results are reported in Appendix M.

Table I. Overview of Treatment Conditions Across Three Experiments

This table provides an overview of the treatment and control conditions of the experiments.

Treatment	**Observation of outcomes**	**Memory elicitation**	**Reminder**	**Belief and/or choice elicitation**	**Experiment**
Delay	Week t	Week t+1	No	Week t+1	1, 2, 3
Immediate(1)	Week t	Week t	No	Week t	1, 2
Immediate2	Week t+1	Week t+1	No	Week t+1	1
NoRecall	Week t	No	No	Week t+1	1
Reminder	Week t	No	Week t+1	Week t+1	3

III. Evidence for a Systematic Memory Bias

Based on our direct memory elicitation, subjects' memory bias is estimated at the individual level by taking the difference between each subject's recalled number of positive and negative outcomes in the memory elicitation task and the actual observed number of outcomes.[14] If subjects had perfect memory, their memory bias would be zero. A memory bias above zero means that subjects recollect a higher number of outcomes than actually observed and a memory bias below zero means that subjects remember a lower number of outcomes than observed.

Result 1. *After one week, subjects exhibit a systematic memory bias: Subjects recall more gains and fewer losses from their investments (i) than actually occurred and (ii) compared to a control group (Immediate condition).*

Figure 1 displays the systematic memory bias measured in Experiment 1. After one week, subjects in the *Delay* condition recall significantly more positive outcomes (left-hand side) and fewer negative outcomes (right-hand side) from their investments than actually occurred.

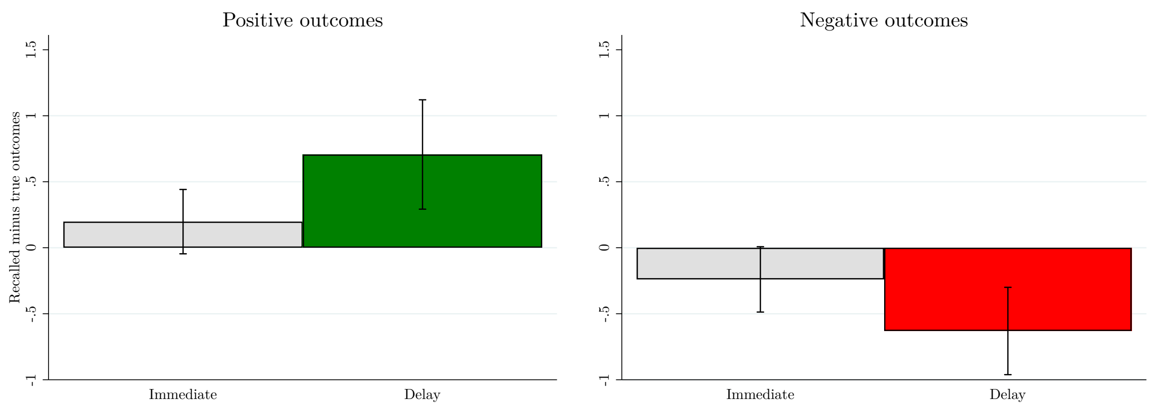

Figure 1. Subjective Memory Bias. This figure displays mean values of subjects' recollection of experienced positive investment outcomes (left-hand side) and negative investment outcomes (right-hand side) from Experiment 1. The bars represent the mean values by treatment. Error bars indicate 95% confidence intervals.

Table II provides an overview of subjects' memory bias in Experiment 1. It reports results from T-tests against the null hypothesis that the memory bias is zero. We first focus on the *Delay* condition, in which people had time to form memory in the week between observing the stock's outcomes and recollecting them. Column 1 illustrates that subjects have biased memory of the stock's outcomes in the *Delay* condition. If they invested in the stock, they remember significantly more positive outcomes and significantly less negative

[14] When eliciting subjects' memory, we reminded them that they observed in total 12 outcomes, but we did not force them to report a total number of 12 recalled outcomes. In our analyses throughout the paper, we include both memory bias variables to account for subjects' deviations from the total number of 12 outcomes.

outcomes than actually observed. On average, subjects significantly over-remember the absolute number of positive outcomes by 0.9 and significantly under-remember the absolute number of negative outcomes by 0.7. We calculate the percentage of over-remembered and under-remembered outcomes compared to actually occurred outcomes at the subject level.[15] On average, after one week, subjects remember 23% more positive outcomes and 10% less negative outcomes than actually occurred.

Table II. Subjective Memory Bias Across Treatments

This table displays subjects' memory bias in the *Delay* and *Immediate* condition in Experiment 1. Memory bias is estimated at the individual level by subtracting the actual individually observed number of positive (negative) outcomes from subject's recalled number of the stock's positive (negative) outcomes. The table reports mean values and T-test results against the null hypothesis that the memory bias is zero and T-test results of the difference in means between our conditions *Delay* and *Immediate* for subjects who invested in the stock. Note that the sample is restricted to subjects who invested in the stock.

Observed outcomes	***Immediate***		***Delay***		**Difference**
	Mean (N=78)	T-test	Mean (N=74)	T-test	T-test
Positive outcomes	0.27	p = 0.058	0.89	p = 0.001	p = 0.018
Negative outcomes	-0.28	p = 0.037	-0.73	p = 0.000	p = 0.038

Importantly, we find that the positive memory bias is also significantly larger compared to what subjects recall in the control group immediately after observing their outcomes (*Immediate* condition). When individuals have one week time to form memory, the bias significantly increases. Thus, our results suggest a memory effect on subjects' recollection beyond factors that might be related to subjects' experience, information acquisition, or processing capacities. The comparison of subjects' recollection between the *Immediate* and *Delay* condition isolates the effect of memory from these alternative explanations.

Further results in Table X in the Appendix illustrate that while subjects who invested in the stock show a significant memory bias for investment outcomes, the small sample of subjects who did not invest do not display a significant memory bias, which is in line with a motivated memory bias. We discuss this underlying mechanism in more detail in Section VI.[16] Our findings are robust to different estimations of the memory bias. Please refer to Section VII.A for robustness of these findings to using estimations based on the relative fraction of the recalled number of positive and negative outcomes as well as the difference between the

[15]For instance, the percentage of over-remembered number of positive outcomes is the difference between the subject's recalled number of positive outcomes and the actual number of positive outcomes observed divided by the actual number of positive outcomes.

[16]Note that we designed our experiment to study investor memory, i.e., how individuals remember stock outcomes when invested. Thus, the sample size of subjects who did not invest in the stock is small and we restrict our further analyses to the sub-sample of subjects who invested in the stock. We address potential selection bias due to subjects selecting into investing in a stock in our Experiment 2 and 3 by letting them choose between two stocks with the same risk characteristics. The results hold.

number of recalled positive and negative outcomes. Further, the results of Experiment 2 and 3 replicate the documented memory bias (Appendix M).

IV. Memory-based Beliefs

In this section, we provide evidence that subjects' memory bias is associated with overly optimistic subjective beliefs about the stock quality, i.e., the expectations about future investment outcomes of that stock.

Result 2. *After one week, subjects form too optimistic beliefs about their investments compared to an objective benchmark.*

Figure 2 gives a graphical presentation of subjects' elicited beliefs in Experiment 1. The figure displays subjects' believed probability that their chosen stock is the good stock relative to the objective Bayesian probability. The x-axis indicates the value of each possible objective Bayesian posterior belief, i.e., the objective probability that the stock is good, and the y-axis represents the average belief indicated by the subject, i.e., the subjective probability that the stock is good. If subjects indicated the objectively correct probability, their subjective posteriors would line up along the 45° reference line. The figure suggests that subjective beliefs deviate from the objective posteriors in a systematic way and are associated with subjects' memory. In the *Delay* condition (right-hand side) subjective beliefs are overly optimistic regarding the likelihood that the stock is good. In numbers, subjects' beliefs in the *Delay* condition are on average 8.16 %-points too optimistic compared to the objective Bayesian posteriors (T-test, $p = 0.004$, see Appendix I for detailed results).

Result 3. *Subjects with a larger memory bias form more optimistic beliefs about their investment.*

Our results further document that the size of subjects' elicited memory bias is significantly correlated with more optimistic beliefs about the stock quality. Figure 2 displays that the belief of subjects with an above average memory bias (solid line) is further away from Bayesian objective posteriors relative to the belief of subjects with a below average memory bias (dashed line).

In accordance with this graphical display, we find that subjects' elicited memory bias positively correlates with subjective beliefs. The regression models in Table III indicate that subjects' memory bias for positive and negative outcomes is significantly correlated with their beliefs. We use subjects' belief that the stock is the good stock (columns 1 and 2) and their belief distortion compared to the Bayesian posterior (columns 3 and 4) as dependent variables. Subjects' belief distortion is the difference between the posterior log-likelihood ratios of subjects' elicited probabilities and the objective Bayesian probabilities. The elicited memory bias

Figure 2. Subjective Beliefs

This figure displays the average subjective posterior that the stock is the good stock, as a function of the correct objective Bayesian probability. The sample is limited to subjects who invested in the stock (first choice) in Experiment 1. If subjective estimates were Bayesian, they would line up on the 45° line. The left panel presents subjective beliefs from the *Delay* condition and the right panel presents subjective beliefs from the *Immediate* condition. Subjects' probability estimates for each level of the objectively correct Bayesian posterior are shown on solid lines for subjects with a high memory bias, and on dashed lines for subjects with a low memory bias. *High Memory Bias* is a dummy variable equal to 1 if the subject has a memory bias that is larger than the mean memory bias in its treatment condition. *Low Memory Bias* is a dummy variable equal to 1 if the subject has a memory bias that is smaller than the mean memory bias in its treatment condition.

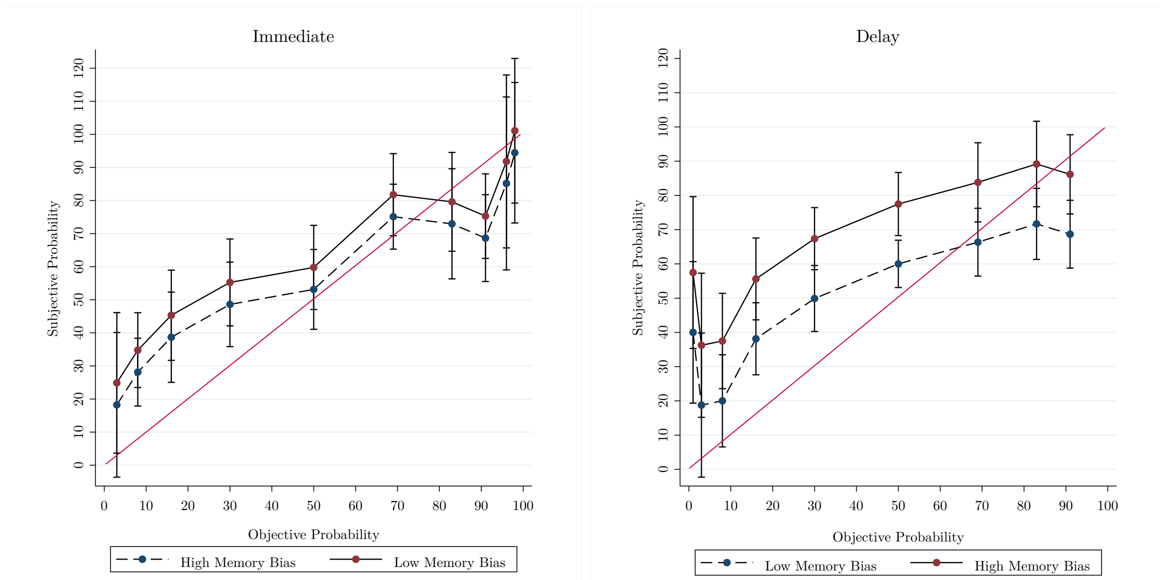

for positive and negative outcomes serve as independent variables. In our first two models, we control for the correct Bayesian probability that the stock is good, given the information seen by the subject. In all models we control for session fixed effects. Column 1 shows that subjects' beliefs are on average 6.20% higher for each positive outcome they over-remember. That is, subjects who recollect a significantly higher number of positive outcomes than observed form more optimistic beliefs about the stock. We find the opposite direction for negative investment outcomes (column 2). Subjects' beliefs are on average 7.96% lower for every loss they under-remember. This means that subjects who remember a lower number of losses than observed form also significantly more optimistic beliefs about the stock.[17] Similarly, results in column 3 and 4 document that subjects' elicited memory bias is significantly correlated with their actual deviation from the objective Bayesian probability. Subjects who remember a higher number of positive outcomes (or a lower number of negative outcomes) than observed, form overly optimistic beliefs compared to the Bayesian benchmark.

Please refer to Section VII.A for robustness of these findings. We show that subjects' belief distortion in the *Delay* condition does not rely on a specific memory elicitation method. Further, we replicate the

[17] Note that the magnitude of the regression coefficients for the objective probability as well as the constant are similar to previous results reported by Kuhnen (2015).

Table III. Memory-Based Beliefs

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable is the subjective posterior belief that the stock is the good stock (from 1 to 100), *Subjective Probability*, or subjects' belief distortion measured by the difference between the posterior log-likelihood ratios of subjects' elicited probabilities and the objective Bayesian probabilities, *Belief Distortion* from Experiment 1. *Memory Bias (for Pos. Outcomes)* represents subject's memory bias for observed positive outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of positive stock outcomes from subject's recalled number of the stock's positive outcomes. *Memory Bias (for Neg. Outcomes)* represents subject's memory bias for observed negative outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of negative stock outcomes from subject's recalled number of the stock's negative outcomes. *Objective Probability* is the value of the objective Bayesian probability that the stock is the good stock (from 1 to 100). *Session* is a dummy variable representing the different sessions of the experiment. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Subjective Probability	(2) Subjective Probability	(3) Belief Distortion	(4) Belief Distortion
Memory Bias (for Pos. Outcomes)	6.196***		0.424***	
	(0.92)		(0.05)	
Memory Bias (for Neg. Outcomes)		-7.958***		-0.525***
		(1.02)		(0.06)
Objective Probability	0.684***	0.701***		
	(0.05)	(0.05)		
Constant	16.751***	15.204***	-0.185	-0.223
	(5.16)	(5.00)	(0.27)	(0.26)
Session	Yes	Yes	Yes	Yes
N	188	188	182	182
R^2	0.52	0.55	0.31	0.35

documented memory effect on subjects' beliefs in Experiment 2 (Appendix M).

The results are consistent with a model of memory-based beliefs in which motivational concerns form the basis for how information is remembered. In the Appendix, we formalize such a model in which the memory of investment outcomes is systematically biased in a motivated way and thereby distorts beliefs. We follow the idea that selective recall might be a potential mechanism for motivated biased beliefs (Bénabou and Tirole, 2002, 2004; Chew et al., 2020; Zimmermann, 2020). In the model, the memory bias stems from quasi-Bayesian belief updating with a probability of under-remembering specific previously observed signals. This probability depends on whether the signals are consistent with the decision maker's motivation. Please refer to Appendix F for the model.

V. Consequences for Investor Behavior

In this section, we show that subjects' biased memory of past investment outcomes is related to future investment behavior.

A. Reinvestment Behavior

In previous empirical work, it has been shown that investors engage in reinforcement learning, meaning that they (i) repeat actions that generated previous gains and (ii) avoid actions that generated previous losses. We find that subjects in our experiment behave in this way (see Table XVII in the Appendix). The objective Bayesian posterior (calculated based on all experienced outcomes of the subject) is significantly positively correlated with subjects' probability to reinvest. In other words, experienced positive outcomes increase subjects' probability to reinvest, and experienced negative outcomes decrease their probability to reinvest. In the following, we present how memory impacts how experienced outcomes affect subsequent actions.

Result 4. *After one week, subjects are (i) more likely to reinvest and (ii) more likely to reinvest suboptimally compared to a control group (Immediate condition).*

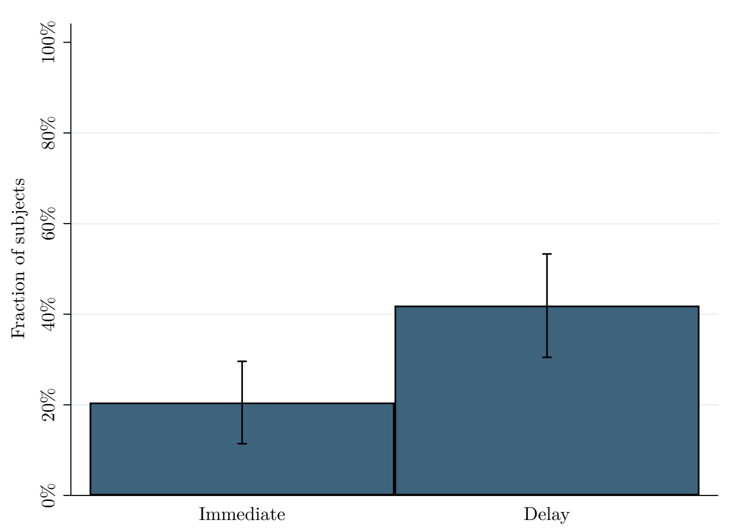

Figure 3. Suboptimal Reinvestment. This figure displays the fraction of subjects reinvesting in their previously chosen stock despite having observed less than 6 positive outcomes in Experiment 1. The bars represent the fractions by treatment. Error bars indicate 95% confidence intervals.

In the *Immediate* condition, subjects were asked to decide whether they would like to reinvest immediately after the observation and in the *Delay* condition with a one week delay. Figure 3 shows our key finding from Experiment 1. It displays the fraction of subjects who chose to reinvest in the stock after observing the investment outcomes although this was a suboptimal choice from a Bayesian perspective, assuming risk neutrality. Given the individual outcomes subjects observed, they invested in the stock, although the stock's expected outcome was lower than the bond's outcome. The results show that 41.9% of the subjects in the *Delay* condition reinvest in the stock despite its lower expected outcome, while in the *Immediate* condition only half of this fraction (20.5% of the subjects) reinvest in such a manner. The difference between the two treatments is significant at the 1% level (see Table XV in the Appendix). Thus, subjects seem to be more likely to avoid suboptimal decisions when immediately deciding whether to reinvest or not (*Immediate* condition), whereas subjects who have to rely on their memory from last week (*Delay* condition) are likely to reinvest, although the decision is suboptimal. Note that we assume risk-neutral subjects. If instead some subjects were risk averse, the fraction of subjects investing suboptimally would be even larger. Hence, our estimates are rather conservative.

Table IV outlines our results from Experiment 1 in more detail. The table reports results from Logit regressions with a dummy variable equal to one for subjects who chose to reinvest in the stock (*Reinvestment*) and a dummy variable equal to one for subjects who chose to reinvest in the stock with a lower expected outcome than the bond after observing the outcomes as dependent variable (*Reinvestment (Suboptimal)*). Odds ratios from the Logit regressions are reported. We use a treatment dummy as well as subjects' memory bias for positive and negative outcomes as independent variables and control for session fixed effects. Subjects in the *Delay* condition are significantly more likely to reinvest in the stock (column 1). The important question is whether this effect is associated with the quality of subjects' choices. Therefore, columns 2-4 report results for subjects who made a suboptimal choice from a Bayesian perspective, assuming risk neutrality. Subjects in the *Delay* condition are significantly more likely to reinvest suboptimally.

Result 5. *Subjects with a larger memory bias are more likely to reinvest suboptimally.*

Further, the results show that subjects' suboptimal investment decisions are correlated with the magnitude of their elicited individual memory bias. Column 3 of Table IV indicates that subjects' probability to reinvest in the stock with a lower expected outcome than the bond increases by 59.9% with each positive outcome they over-remember. In other words, subjects who recollect a higher number of positive outcomes than actually observed have a significantly higher probability to reinvest in the stock with a lower expected outcome. We find the opposite direction for negative investment outcomes (column 4). Subjects' probability to reinvest in the stock with the lower expected outcome is on average 38.3% lower for every loss they over-remember. Note that the average subject remembers a lower number of negative outcomes than actually occurred (Table II). In this regard, each under-remembered loss increases the probability to reinvest in the stock with a lower expected outcome. Together, these results document that subjects' memory bias is related to suboptimal reinvestment choices. We replicate this finding in Experiment 2 (see Appendix M).

Table IV. Subjective Memory and Reinvestment Decisions

This table contains the odds ratios and standard errors (in parentheses) of Logit regressions in which the dependent variable is a dummy variable which is equal to one if the subject invested in the stock with a lower expected outcome than the bond after the observation phase (second choice) in Experiment 1. *Treatment (Delay)* is a dummy variable representing our treatment with 1 = Delay condition and 0 = Immediate condition. *Memory Bias (for Pos. Outcomes)* represents subject's memory bias for observed positive outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of positive stock outcomes from subject's recalled number of the stock's positive outcomes. *Memory Bias (for Neg. Outcomes)* represents subject's memory bias for observed negative outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of negative stock outcomes from subject's recalled number of the stock's negative outcomes. *Session* is a dummy variable representing the different sessions of the experiment. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Reinvestment	(2) Reinvestment (Suboptimal)	(3) Reinvestment (Suboptimal)	(4) Reinvestment (Suboptimal)
Treatment (*Delay*)	2.213**	3.542***		
	(0.78)	(1.44)		
Memory Bias (for Pos. Outcomes)			1.599***	
			(0.22)	
Memory Bias (for Neg. Outcomes)				0.617***
				(0.10)
Constant	0.453	0.043***	0.077**	0.076**
	(0.27)	(0.05)	(0.08)	(0.08)
Session	Yes	Yes	Yes	Yes
N	152	152	152	152
Pseudo R^2	0.05	0.11	0.13	0.11

B. *Investor Confidence*

So far, we have shown that subjects form too optimistic beliefs about the quality of their stock based on memory and reinvest too often in their chosen stock as a result. In this section, we show that biased memory also affects individuals' confidence in making the optimal choice, that is in their ability to select good stocks relative to others in a market setting. Importantly, we document that biased memory affects individuals' *over*confidence.

Result 6. *After one week, subjects exhibit (i) more confidence and (ii) more overconfidence in their ability to select good stocks relative to others when they must rely on their memory from last week (Delay condition) compared to the control group (Reminder condition).*

Figure 4 shows our key finding from Experiment 3. It displays subjects' overconfidence, i.e., subjects' betting on the optimality of their investment choices although the decision was suboptimal from a Bayesian perspective ("suboptimal betting") across treatments. In contrast to the *Delay* condition, subjects in the *Reminder* condition do not have to rely on their memory while they make their subsequent choice in week $t+1$, because we remind them of their investment outcomes on their screen. The results show that on average subjects bet 46.3 points in the *Delay* condition despite the choice being suboptimal, while in the *Reminder* condition subjects only bet 31.5 points. Subjects bet on average 47% more given suboptimal decisions when they must rely on their memory from last week (*Delay* condition) compared to the control group (*Reminder* condition). This difference is significant at the 1%-level (see Table XVI in the Appendix).

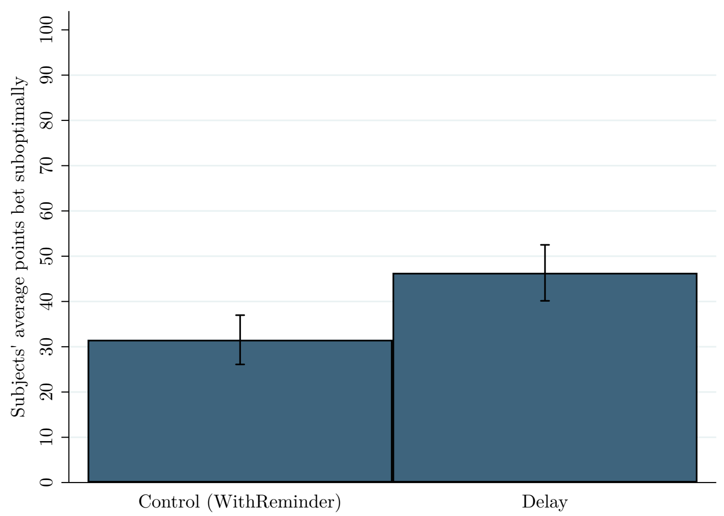

Figure 4. Overconfidence: Suboptimal Betting. This figure displays suboptimal betting by subjects in Experiment 3. It shows the average points bet on the optimality of their investment choices by subjects despite having made a suboptimal choice, because they invested in the stock with the lower probability of being the good stock ("suboptimal betting"). The bars represent the mean values by treatment. Error bars indicate 95% confidence intervals.

Table V outlines our results from Experiment 3 in more detail. The table reports results from OLS regressions with a variable indicating subjects' number of points bet on the optimality of their investment choice (*Points Bet*) and a variable indicating subjects' number of points bet on the optimality of their investment choice despite having made a suboptimal choice, because they invested in the stock with the lower probability of being the good stock (*Points Bet (Suboptimal)*). We use a treatment dummy as well as subjects' memory bias for positive outcomes as independent variables. Note, in Experiment 3 we forced subjects' responses in the memory elicitation task to always add up to 12 outcomes (i.e., the true number of outcomes observed). Thus, analyses using subjects' memory bias for negative outcomes yield the same finding as using positive outcomes.[18] Subjects in the *Delay* condition bet significantly more points (column 1). Again, the important question is whether this effect is associated with the quality of subjects' choices. Therefore, columns 2 and 3 report results for subjects who made a suboptimal investment choice from a Bayesian perspective, i.e., it reports results for subjects' overconfidence. Subjects in the *Delay* condition bet on average 14.8 points more on suboptimal decisions.

[18] We do not control for session fixed effects, since we collected all the data of Experiment 3 within one online session on the same date.

Table V. Subjective Memory and Confidence

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable indicates subjects' number of points bet on the optimality of their investment choice (*Points Bet*) and a variable indicating subjects' number of points bet on the optimality of their investment choice despite having observed less than 6 positive outcomes (*Points Bet (Suboptimal)*) in Experiment 3. *Treatment (Delay)* is a dummy variable representing our treatment with 1 = Delay condition and 0 = Reminder condition. *Memory Bias (for Pos. Outcomes)* represents subject's memory bias for observed positive outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of positive stock outcomes from subject's recalled number of the stock's positive outcomes. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Points Bet	(2) Points Bet (Suboptimal)	(3) Points Bet (Suboptimal)
Treatment (Delay)	9.381***	14.788***	
	(2.85)	(4.18)	
Memory Bias (for Pos. Outcomes)			2.790**
			(1.25)
Constant	47.714***	31.537***	41.990***
	(1.96)	(2.75)	(3.63)
N	487	191	83
R^2	0.02	0.06	0.06

Result 7. *Subjects with a larger memory bias exhibit more overconfidence in their ability to select good stocks relative to others.*

Further, the results show that subjects' overconfidence measured by their betting given suboptimal choices (and therefore also suboptimal betting) is correlated with the magnitude of their elicited individual memory bias. Column 3 indicates that with each positive outcome subjects over-remember, they bet 2.8 more points suboptimally.[19] In other words, subjects who recollect a higher number of positive outcomes than actually observed bet a significantly higher amount of points suboptimally.

VI. Exploring the Underlying Mechanism of the Memory Bias

In this section, we explore motivation as underlying mechanism of the memory bias. We report experimental evidence that supports the idea of a motivated memory bias for investment outcomes.

Result 8. *In line with an underlying motivational mechanism of the memory bias, we find that (i) subjects suppress memories of negative outcomes rather than actually forgetting them and (ii) subjects' memory bias depends on whether one made an active investment choice or not.*

Table VI shows our main findings from Experiment 2. It illustrates subjects' mean memory bias for positive outcomes in the *Delay* treatment group and the treatment effect (difference in mean memory bias

[19]Note that we could only elicit memory in the *Delay* condition. Thus, the observations in column 3 are restricted to subjects in the Delay condition who bet points suboptimally (N = 83).

between *Delay* and *Immediate* group) across the three experimental conditions. Note, in Experiment 2 we forced subjects' responses in the memory elicitation task to always add up to 12 outcomes (i.e., the true number of outcomes observed). Thus, analyses using subjects' memory bias for negative outcomes yield the same finding as using positive outcomes. Table XI in the Appendix shows subjects' memory bias in Experiment 2, separately for positive and negative outcomes. The table reports mean values and T-tests results against the null hypothesis that the memory bias is zero (column 1) and T-test results of the difference in means between *Delay* and *Immediate* (column 2). The memory bias is estimated at the individual level by taking the difference between each subject's recalled number of positive outcomes in the memory elicitation task and the actual observed number of positive outcomes.

Table VI. Subjective Memory Bias Across Experimental Treatments

This table displays subjects' mean memory bias for positive outcomes (compared to the true number of outcomes) in the *Delay* group, separately for subjects in the *Baseline*, *HighStakes*, and *NoChoice* condition as well as the treatment effect (the difference in memory bias for positive outcomes in the *Delay* group vs. *Immediate* group)in Experiment 2. Memory bias is estimated at the individual level by subtracting the actual individually observed number of positive outcomes from subject's recalled number of the stock's positive outcomes. The table reports mean values and T-test results against the null hypothesis that the memory bias is zero as well as the T-test results of the difference in means between *Delay* and *Immediate*.

Treatment	**Mean Memory Bias (for Pos. Outcomes) in *Delay* condition**	**Treatment Effect**	**Mechanism**
Baseline (N= 185)	0.36 (p = 0.023)	0.48 (p = 0.010)	-
HighStakes (N = 154)	0.11 (p = 0.591)	0.32 (p = 0.151)	Motivated memory suppression
NoChoice (N = 159)	0.18 (p = 0.398)	0.49 (p = 0.051)	Active choice vs. passive choice

A. *Memory Suppression*

We have two key results in support of motivated memory suppression. First, as reported in Table VI, we find that subjects show a significant positive bias in memory in the *Baseline* condition. In contrast, subjects in the *HighStakes* condition do not exhibit a systematic memory bias one week after they observed the investment outcomes. Further, we do not find a significant treatment effect between the *Delay* and *Immediate* groups when subjects face high recall incentives. That is, with high financial incentives for memory accuracy we observe no significant memory bias.

Second, we expect the effect of high incentives for accurate memory to be more effective for subjects with high motivational reasons to suppress memory otherwise (Bénabou and Tirole, 2002; Zimmermann, 2020). Thus, we explore whether the memory suppression is specifically related to subjects' self-image concerns. We elicit subjects' tendency for self-deceptive enhancement, that is, the tendency for motivated self-enhancement (Paulhus, 1991). Self-enhancement is a type of motivation that works to make people feel good about themselves and to maintain self-esteem. We expect that the effect of a high monetary benefit of reporting accurate memory reduces the memory bias only, or at least more, for subjects scoring high on this self-enhancement scale.

Table VII. The Effect of High Incentives on Subjective Memory Bias

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable is subject's memory bias for observed positive outcomes *Memory Bias (for Pos. Outcomes)* in Experiment 2. The two regression models restrict the sample to subjects who score above-average on the self-deceptive enhancement (SDE) scale (1) and to subjects who score below-average on the self-deceptive enhancement (SDE) scale (2). *Treatment (HighStakes)* is a dummy variable that equals one if the subject took part in the *HighStake* condition and zero if the subject took part in the *Baseline* condition. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Memory Bias (for Pos. Out.) - high SDE	(2) Memory Bias (for Pos. Out.) - low SDE
Treatment (*HighStakes*)	-0.469**	0.115
	(0.19)	(0.22)
Constant	0.182	0.058
	(0.12)	(0.16)
N	172	167
R^2	0.04	0.00

Table VII provides evidence for testing this hypothesis. We use subjects' memory bias for positive outcomes as dependent variable. We split the sample based on subjects' responses to a validated survey scale, which measures the tendency to deceive oneself for self-enhancement (Paulhus, 1991). A dummy variable that equals one if the subject took part in the *HighStake* condition and zero if the subject took part in the *Baseline* condition serves as independent variable. In our first model, we restrict the sample to subjects who score above-average on the self-deceptive enhancement (SDE) scale and in our second model, we restrict the sample to subjects who score below-average on the SDE scale. We find that higher incentives significantly decrease the memory bias of subjects who score above-average on the SDE scale, but not the memory bias of subjects who score below-average on the SDE scale.

B. Active Choice

In order to explore the role of active choice for subjects' memory bias, we compare subjects' memory bias in the *Delay* treatment groups in the *Baseline* and *NoChoice* condition in Table VI. We find that in contrast to subjects in the *Baseline* condition, subjects in the *NoChoice* condition do not exhibit a systematic memory bias one week after they observed the investment outcomes. That is, the memory bias is alleviated when subjects do not actively choose their stock.

Interestingly, the recall bias in the *Immediate* group of the *NoChoice* condition is significantly negative (and significantly different from the memory bias in the *Delay* group). Immediately after the observation of outcomes, subjects report significantly less positive outcomes than occurred (see Table XII in the Appendix). We argue that this is in line with the notion that, in contrast to people's memory, people's attention to the bad rather than attention to the good is dominant when motivational reasons are limited.[20]

VII. Robustness of Results

Our results are robust to different measures of memory bias, to between-subject differences in risk preferences and financial knowledge, and to the order and timing of the experimental tasks. Remarkably, better financial knowledge does not alleviate the memory bias. Additionally, we show that the mere fact that we asked subjects to recall outcomes did not drive subjective beliefs or subsequent investment decisions. In addition, the results from Experiment 1 are replicated in Experiment 2 and Experiment 3. Moreover, we rule out anticipated regret as an alternative underlying mechanism driving the memory bias.

A. *Different Measures of Memory Bias*

The main findings in Table II are robust to different measures of subjects' memory bias. We test two alternative measures, a memory bias measure based on the fraction of recalled positive or negative outcomes among all observed outcomes (Table XVIII) and a memory bias of the absolute difference between recalled positive and negative outcomes (Table XIX). Our results are qualitatively unchanged. See Appendix L for more details.

Based on our model in (Appendix F), we also estimate subjects' bias in recall without relying on any memory elicitation. The regression specification allows us to directly estimate the probabilities of under-remembering positive and negative outcomes from subjects' reported subjective posterior beliefs that the stock is the good stock in Experiment 1. Consistent with our documented memory bias, our regression results suggest that after one week, subjects in the *Delay* condition have a probability of under-remembering negative outcomes of almost 60% ($p < 0.001$), which is higher than the probability in the *Immediate* condition. For details see Appendix H.

[20]Research suggests that people's attention to the bad dominates their attention to the good (Baumeister, Bratslavsky, Finkenauer, and Vohs, 2001). For example, psychology research provides evidence that subjects in the lab pay more attention to photographs depicting negative content compared to photographs depicting positive content (Fiske, 1980) and negative social information is more attention-grabbing than positive one (Pratto and John, 1991).

B. Subjects' Risk Preferences

In Experiment 1, investing in the stock could be driven by subjects' risk preferences. Yet, we show that our findings are robust to differences in subjects' self-reported risk preferences (Dohmen et al., 2011). First, subjects' memory bias is not correlated with self-reported risk preferences (Table XX). Second, our main findings are robust to controlling for individual risk preferences (Table XXI). See Appendix L for more details. Note, that in Experiment 2 and Experiment 3, risk preferences cannot impact our results.

C. Subjects' Financial Literacy

Our findings are robust to differences in subjects' level of financial literacy. To measure subjects' financial literacy, we asked them to calculate the expected value of an investment allocated between a stock and a savings account at the end of the session (Kuhnen, 2015). First, Table XXII reports the robustness of our findings in Table II of Section III, by separately testing the significance of memory bias for subjects with different levels of financial literacy. We find that higher financial knowledge does not alleviate the memory bias. Second, subjects' memory bias is not correlated with their financial knowledge (Table XXIII. Third, Table XXI shows the robustness of our main findings of Section IV and V to differences in subjects' financial literacy. The results are qualitatively unchanged. See Appendix L for more details.

D. Order and Timing of Experimental Tasks

To avoid order effects the order in which subjects perform the experimental tasks after observing the investment outcomes, i.e., the memory elicitation, belief elicitation, and the second investment task, is random. In addition, we tested whether subjects' memory bias, stated beliefs, and investment decisions after the observation are influenced by the task order and we do not find any significant impact on the memory bias, subjective beliefs, and suboptimal reinvesting (T-test, $p = 0.314$, $p = 0.890$, and $p = 0.428$, respectively). Further, subjects in the *Immediate* condition perform all experimental tasks in the course of one session. To rule out timing effects, we vary whether subjects perform the tasks in week t (*Immediate1* condition) or in week $t + 1$ (*Immediate2* condition). In half of the experimental sessions, subjects perform the tasks in week t and in the other half of the sessions in week $t + 1$. Our experimental findings are robust to the differences in timing of the tasks (see Table XXIV in Appendix L).

E. Timing of Memory Elicitation

We included a *NoRecall* condition, in which subjects do not perform a memory elicitation task as a control condition. This control condition allows us to identify potential effects of simply asking subjects to recall investment outcomes. Table XXV in Appendix L indicates no significant difference in subjects' beliefs or investment decisions between the *Delay* and *NoRecall* condition (T-test, $p = 0.714$ and $p = 0.343$, respectively). This holds for the sub-sample of subjects who invested in the stock with a lower expected return ($p = 0.383$). Thus, simply asking subjects to recall investment outcomes seems to have no effect on their subjective beliefs or choices afterwards. We find similar results for the sub-sample of subjects who performed the memory elicitation task *after* the belief elicitation task and second investment task.

F. *Replication Results from Experiment 2 and 3*

Experiment 2 replicates all findings from Experiment 1. Subjects exhibit a systematic bias in memory for their investment outcomes. They remember significantly more positive and less negative outcomes than actually occurred and compared to a control group. Further, the memory bias is related to overly optimistic subjective beliefs about the investment and suboptimal reinvesting decisions. In addition, the design of Experiment 3 allows to test for the replicability of the positive memory bias in a general population sample and we indeed replicate the memory bias in the general population sample as well. See Appendix M for more details.

G. *Alternative Underlying Mechanism of the Memory Bias*

In contrast to memory suppression and the role of active choice, which both are consistent with a motivational bias linked to self-image concerns, another motivational mechanism of our documented memory bias could also encompass minimizing anticipated regret (Loomes and Sugden, 1982). Therefore, we test for the role of subjects' anticipated regret. We elicit subjects' proneness to regret on a 5-item scale and create an index variable (Schwartz et al., 2002). We find that subjects' memory bias is not significantly correlated with subjects' proneness to regret (see Appendix N).

VIII. Conclusion

In three experiments, we study how investors remember their investment outcomes, how they form beliefs based on these memories, and how these translate into decisions. We document that individuals hold too positive memories of their investment outcomes. They over-remember positive investment outcomes of their chosen assets and under-remember negative ones. We further show that individuals do not adjust their beliefs or behavior to account for the fallibility of their memory. We provide direct evidence that the positive memory bias leads individuals' to (i) form overly optimistic beliefs about the quality of their investments,

(ii) reinvest too much, and (iii) become overconfident in their ability to invest. The memory bias offers a cognitive microfoundation for a well-documented pattern in investor beliefs and behavior: gains weight more than losses when investors learn from experience.

We explicitly designed experiments in which subjects have to rely on their recollection of investment outcomes and we directly elicit subjects' memory of these outcomes at different points in time. This allows us to isolate biased memory as a microfoundation of subjective beliefs and choices. People often rely on their subjective memory rather than on records of historical data or information. A neoclassical explanation for this is that information acquisition can be costly due to search frictions. On the other hand, people might avoid objective information actively because they are boundedly rational. For example, it has been shown that people frequently refuse information, because of the so-called ostrich effect, regret aversion, or the motive to maintain positive feelings (Caplin and Leahy, 2001; Karlsson, Loewenstein, and Seppi, 2009; Golman, Hagmann, and Loewenstein, 2017). When people do not look up objective information, they might make investment decisions that are biased by their subjective memory.

Moreover, we provide evidence in support of the notion that motivation drives the documented positive memory bias. We find that subjects in our experiments suppress memories of negative outcomes rather than actually forgetting them and that their active choice plays an important role for the memory bias.

The memory bias documented in this paper opens several avenues for further research. First, we document positivity in selective recall. This complements prior experimental work showing evidence for selective memory based on associativeness (Bordalo et al., 2021; Enke et al., forthcoming). Theoretical and empirical investigations of the interaction between selective recall processes based on associativeness and motivation (i.e., positivity) seem promising avenues for future research. Second, we designed our experiment in a way that it rules out the effects of extreme outcomes and documents memory bias for day-to-day fluctuations. An interesting avenue for future research is to explore the memory bias with regard to extreme experiences. Psychologists have shown that people can exhibit enhanced memory for extreme emotional events (Phelps, LaBar, Anderson, O'Connor, Fulbright, and Spencer, 1998; Hamann, Ely, Grafton, and Kilts, 1999; Hamann, 2001). Economists have shown that experiences of extreme events have long-lasting effects on risk-taking (Malmendier and Nagel, 2011). Building a bridge between these two streams of literature seems promising. Third, we study selective recall of experienced investment outcomes. It would be interesting to contrast our findings to how investors recall "uninvolved information," for instance, how investors recall outcomes when they are not invested in the stock market. This could provide novel insights into reinforcement learning in financial markets. Fourth, our experiments studied the effect of biased memory one week after the initial investment decision and observation of outcomes. Future work can explore the degree and dynamics of biased memory over longer time horizons. Fifth, motivational reasons and the fundamental tendency to maintain and enhance self-esteem (Bénabou and Tirole, 2002) are used to explain how investment ideas spread in financial markets (Han, Hirshleifer, and Walden, 2018; Shiller, 2017). It has been suggested that investors have a self-enhancing bias in conversations, i.e. they talk more about good than about bad returns, which influences the popularity and thus pricing of particular investment strategies. Future work can explore the interplay between memory bias and social interaction in financial markets. Sixth, we conduct laboratory and online experiments that provide tight identification of biased memory and an objectively correct benchmark for subjects' beliefs and choices. A promising avenue for future research is to explore the effects of a positive memory bias in the field, similar to recent work investigating associative memory of investors in the field (Jiang, Liu, Peng, and Yan, 2022; Charles, 2022).

REFERENCES

Adler, Orly, and Ainat Pansky, 2020, A "Rosy View" of the Past: Positive Memory Biases, in *Cognitive Biases in Health and Psychiatric Disorders*, 139–171 (Elsevier).

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.

Allison, Scott T., and David M. Messick, 1988, The Feature-Positive Effect, Attitude Strength, and Degree of Perceived Consensus, *Personality and Social Psychology Bulletin* 14, 231–241.

Anagol, Santosh, Vimal Balasubramaniam, and Tarun Ramadorai, 2018, Endowment Effects in the Field: Evidence from India's IPO Lotteries, *The Review of Economic Studies* 85, 1971–2004.

Anagol, Santosh, Vimal Balasubramaniam, and Tarun Ramadorai, 2021, Learning from Noise: Evidence from India's IPO lotteries, *Journal of Financial Economics* 140, 965–986.

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor Sentiment and the Cross-section of Stock Returns, *The Journal of Finance* 61, 1645–1680.

Baker, Malcolm, and Jeffrey Wurgler, 2007, Investor Sentiment in the Stock Market, *Journal of Economic Perspectives* 21, 129–151.

Barber, Brad M., Yi-Tsung Lee, Yu-Jane Liu, Terrance Odean, and Ke Zhang, 2020, Learning, Fast or Slow, *Review of Asset Pricing Studies* 10, 61–93.

Barber, Brad M., and Terrance Odean, 2001, Boys Will be Boys: Gender, Overconfidence, and Common Stock Investment, *The Quarterly Journal of Economics* 116, 261–292.

Barber, Brad M., and Terrance Odean, 2008, All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors, *Review of Financial Studies* 21, 785–818.

Barron, Kai, 2021, Belief Updating: Does the 'Good-News, Bad-News' Asymmetry Extend to Purely Financial Domains?, *Experimental Economics* 24, 31–58.

Baumeister, Roy F., Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D Vohs, 2001, Bad is Stronger than Good, *Review of General Psychology* 5, 323–370.

Bénabou, Roland, and Jean Tirole, 2002, Self-Confidence and Personal Motivation, *Quarterly Journal of Economics* 117, 871–915.

Bénabou, Roland, and Jean Tirole, 2004, Willpower and Personal Rules, *Journal of Political Economy* 112, 848–886.

Bénabou, Roland, and Jean Tirole, 2016, Mindful Economics: The Production, Consumption, and Value of Beliefs, *Journal of Economic Perspectives* 30, 141–64.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, 2018, Behavioral Household Finance, in *Handbook of Behavioral Economics: Applications and Foundations 1*, 177–276 (Elsevier).

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, Frederik Schwerter, and Andrei Shleifer, 2021, Memory and Representativeness, *Psychological Review* 128, 71.

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, and Andrei Shleifer, 2012, Salience Theory of Choice under Risk, *The Quarterly Journal of Economics* 127, 1243–1285.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2013, Salience and Asset Prices, *American Economic Review* 103, 623–628.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2020, Memory, Attention, and Choice, *The Quartely Journal of Economics* 135.

Brunnermeier, Markus K., and Jonathan A. Parker, 2005, Optimal Expectations, *American Economic Review* 95, 1092–1118.

Caplin, Andrew, and John Leahy, 2001, Psychological Expected Utility Theory and Anticipatory Feelings, *Quarterly Journal of Economics* 116, 55–79.

Carlson, Ryan W., Michel Maréchal, Bastiaan Oud, Ernst Fehr, and Molly J. Crockett, 2020, Motivated Misremembering of Selfish Decisions, *Nature Communications* 11, 2100.

Charles, Constantin, 2022, Memory Moves Markets, Working Paper, Available at SSRN 4019728.

Chen, Daniel L., Martin Schonger, and Chris Wickens, 2016, oTree—An Open-Source Platform for Laboratory, Online, and Field Experiments, *Journal of Behavioral and Experimental Finance* 9, 88–97.

Chew, Soo H., Wei Huang, and Xiaojian Zhao, 2020, Motivated False Memory, *Journal of Political Economy* 128, 3913–3939.

Cioffi, Delia, and Randy Garner, 1996, On Doing the Decision: Effects of Active versus Passive Choice on Commitment and Self-Perception, *Personality and Social Psychology Bulletin* 22, 133–147.

de Palma, André, Mohammed Abdellaoui, Giuseppe Attanasi, Moshe Ben-Akiva, Ido Erev, Helga Fehr-Duda, Dennis Fok, Craig R. Fox, Ralph Hertwig, Nathalie Picard, et al., 2014, Beware of Black Swans: Taking Stock of the Description–Experience Gap in Decision Under Uncertainty, *Marketing Letters* 25, 269–280.

Deaves, Richard, Erik Lüders, and Guo Ying Luo, 2009, An Experimental Test of the Impact of Overconfidence and Gender on Trading Activity, *Review of Finance* 13, 555–575.

DellaVigna, Stefano, and Joshua M. Pollet, 2009, Investor Inattention and Friday Earnings Announcements, *Journal of Finance* 64, 709–749.

Di Tella, Rafael, Ricardo Perez-Truglia, Andres Babino, and Mariano Sigman, 2015, Conveniently Upset: Avoiding Altruism by Distorting Beliefs About Others' Altruism, *American Economic Review* 105, 3416–3442.

Dohmen, Thomas, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G. Wagner, 2011, ndividual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences, *Journal of the European Economic Association* 9, 522–550.

Du, Mengqiao, Alexandra Niessen-Ruenzi, and Terrance Odean, 2019, Stock Repurchasing Bias of Mutual Funds, Working Paper, Available at SSRN 3247066.

Eil, David, and Justin M. Rao, 2011, The Good News - Bad News Effect: Asymmetric Processing of Objective Information About Yourself, *American Economic Journal: Microeconomics* 3, 114–38.

Enke, Benjamin, Thomas Graeber, and Ryan Oprea, 2023, Confidence, Self-selection, and Bias in the Aggregate, *American Economic Review* 113, 1933–1966.

Enke, Benjamin, Frederik Schwerter, and Florian Zimmermann, forthcoming, Associative Memory, Beliefs and Market Interactions, in *Journal of Financial Economics*.

Fazio, Russell H., Jeaw-Mei Chen, Elizabeth C. McDonel, and Steven J. Sherman, 1982, Attitude Accessibility, Attitude-Behavior Consistency, and the Strength of the Object-Evaluation Association, *Journal of Experimental Social Psychology* 18, 339–357.

Fischbacher, Urs, 2007, z-Tree: Zurich Toolbox for Ready-Made Economic Experiments, *Experimental Economics* 10, 171–178.

Fiske, Susan T., 1980, Attention and Weight in Person Perception: The Impact of Negative and Extreme Behavior, *Journal of Personality and Social Psychology* 38, 889–906.

Fudenberg, Drew, Giacomo Lanzani, and Philipp Strack, forthcoming, Selective Memory Equilibrium, in *Journal of Political Economy*.

Gervais, Simon, and Terrance Odean, 2001, Learning to be Overconfident, *Review of Financial Studies* 14, 1–27.

Golman, Russell, David Hagmann, and George Loewenstein, 2017, Information Avoidance, *Journal of Economic Literature* 55, 96–135.

Hamann, Stephan B., 2001, Cognitive and Neural Mechanisms of Emotional Memory, *Trends in Cognitive Sciences* 5, 394–400.

Hamann, Stephan B., Timothy D. Ely, Scott T. Grafton, and Clinton D. Kilts, 1999, Amygdala Activity Related to Enhanced Memory for Pleasant and Aversive Stimuli, *Nature Neuroscience* 2, 289–294.

Han, Bing, David Hirshleifer, and Johan Walden, 2018, Social Transmission Bias and Investor Behavior, *Journal of Financial and Quantitative Analysis* 1–42.

Hartzmark, Samuel M., Samuel D. Hirshman, and Alex Imas, 2021, Ownership, Learning, and Beliefs, *Quarterly Journal of Economics* 136, 1665–1717.

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.

Hertwig, Ralph, and Ido Erev, 2009, The Description-Experience Gap in Risky Choice, *Trends in Cognitive Sciences* 13, 517–523.

Hirshleifer, David, Danling Jiang, and Yuting Meng DiGiovanni, 2020, Mood Beta and Seasonalities in Stock Returns, *Journal of Financial Economics* 137, 272–295.

Holt, Charles A., and Susan K. Laury, 2002, Risk Aversion and Incentive Effects, *American Economic Review* 92, 1644–1655.

Holt, Charles A., and Susan K. Laury, 2005, Risk Aversion and Incentive Effects: New Data without Order Effects, *American Economic Review* 95, 902–912.

Huffman, David, Collin Raymond, and Julia Shvets, 2022, Persistent Overconfidence and Biased Memory: Evidence from Managers, *American Economic Review* 112, 3141–3175.

Jiang, Zhengyang, Hongqi Liu, Cameron Peng, and Hongjun Yan, 2022, Investor Memory and Biased Beliefs: Evidence from the Field, Working Paper, Available at SSRN.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect Theory: An Analysis of Decision under Risk, *Econometrica* 47, 263–292.

Karlsson, Niklas, George Loewenstein, and Duane Seppi, 2009, The Ostrich Effect: Selective Attention to Information, *Journal of Risk and Uncertainty* 38, 95–115.

Kaustia, Markku, and Samuli Knüpfer, 2008, Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions, *Journal of Finance* 63, 2679–2702.

Köszegi, Botond, 2006, Ego Utility, Overconfidence, and Task Choice, *Journal of the European Economic Association* 4, 673–707.

Kuhnen, Camelia M., 2015, Asymmetric Learning from Financial Information, *Journal of Finance* 70, 2029–2062.

Kuhnen, Camelia M, and Brian Knutson, 2011, The Influence of Affect on Beliefs, Preferences, and Financial Decisions, *Journal of Financial and Quantitative Analysis* 46, 605–626.

Kuhnen, Camelia M., Sarah Rudorf, and Bernd Weber, 2017, The Effect of Prior Choices on Expectations and Subsequent Portfolio Decisions, Working Paper No. w23438, National Bureau of Economic Research.

Loomes, Graham, and Robert Sugden, 1982, Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty, *The Economic Journal* 92, 805–824.

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

Mather, Mara, Eldar Shafir, and Marcia K. Johnson, 2000, Misremembrance of Options Past: Source Monitoring and Choice, *Psychological Science* 11, 132–138.

Mather, Mara, Eldar Shafir, and Marcia K. Johnson, 2003, Remembering Chosen and Assigned Options, *Memory & Cognition* 31, 422–433.

Matt, Georg E., Carmelo Vázquez, and W. Keith Campbell, 1992, Mood-congruent Recall of Affectively Toned Stimuli: A Meta-analytic Review, *Clinical Psychology Review* 12, 227–255.

Mobius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat, 2014, Managing Self-Confidence: Theory and Experimental Evidence, Working paper, National Bureau of Economic Research.

Mullainathan, Sendhil, 2002, A Memory-Based Model of Bounded Rationality, *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.

Nicolosi, Gina, Liang Peng, and Ning Zhu, 2009, Do Individual Investors Learn From Their Trading Experience?, *Journal of Financial Markets* 12, 317–336.

Odean, Terrance, 1998, Volume, Volatility, Price, and Profit When All Traders Are Above Average, *Journal of Finance* 53, 1887–1934.

Paulhus, Delroy L., 1991, *Measures of Personality and Social Psychological Attitudes*, chapter Measurement and Control of Response Bias, 17–59 (Academic Press).

Phelps, Elizabeth A., Kevin S. LaBar, Adam K. Anderson, Kevin J. O'Connor, Robert K. Fulbright, and Dennis D. Spencer, 1998, Specifying the Contributions of the Human Amygdala to Emotional Memory: A Case Study, *Neurocase* 4, 527–540.

Pratto, Felicia, and Oliver P. John, 1991, Automatic Vigilance: The Attention-Grabbing Power of Negative Social Information, *Journal of Personality and Social Psychology* 61, 380–391.

Rabin, Matthew, 2000, Risk Aversion and Expected-Utility Theory: A Calibration Theorem, *Econometrica* 68, 1281–1292.

Rabin, Matthew, 2002, Inference by Believers in the Law of Small Numbers, *Quarterly Journal of Economics* 117, 775–816.

Rabin, Matthew, and Joel L. Schrag, 1999, First Impressions Matter: A Model of Confirmatory Bias, *Quarterly Journal of Economics* 114, 37–82.

Saucet, Charlotte, and Marie Claire Villeval, 2019, Motivated Memory in Dictator Games, *Games and Economic Behavior* 117, 250–275.

Schwardmann, Peter, and Joel Van der Weele, 2019, Deception and Self-deception, *Nature Human Behaviour* 3, 1055–1061.

Schwartz, Barry, Andrew Ward, John Monterosso, Sonja Lyubomirsky, Katherine White, and Darrin R. Lehman, 2002, Maximizing versus Satisficing: Happiness is a Matter of Choice, *Journal of Personality and Social Psychology* 83, 1178–1197.

Seru, Amit, Tyler Shumway, and Noah Stoffman, 2010, Learning by Trading, *The Review of Financial Studies* 23, 705–739.

Shiller, Robert J., 2017, Narrative Economics, *American Economic Review* 107, 967–1004.

Strahilevitz, Michal A., Terrance Odean, and Brad M. Barber, 2011, Once Burned, Twice Shy: How Naive Learning, Counterfactuals, and Regret Affect the Repurchase of Stocks Previously Sold, *Journal of Marketing Research* 48, 102–120.

Taylor, Shelley E., and Jonathon D. Brown, 1988, Illusion and Well-being: A Social Psychological Perspective on Mental Health, *Psychological Bulletin* 103, 193–210.

Thaler, Richard H., and William T. Ziemba, 1988, Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries, *Journal of Economic Perspectives* 2, 161–174.

Wachter, Jessica A., and Michael Jacob Kahana, forthcoming, A Retrieved-context Theory of Financial Decisions, in *The Quartely Journal of Economics*.

Watson, David, Lee Anna Clark, and Auke Tellegen, 1988, Development and Validation of Brief Measures of Positive and Negative Affect: The PANAS Scales, *Journal of Personality and Social Psychology* 54, 1063.

Williams, Samantha E., Jaclyn H. Ford, and Elizabeth A. Kensinger, 2022, The Power of Negative and Positive Episodic Memories, *Cognitive, Affective, & Behavioral Neuroscience* 22, 869–903.

Wulff, Dirk U., Max Mergenthaler-Canseco, and Ralph Hertwig, 2018, A Meta-analytic Review of Two Modes of Learning and the Description-Experience Gap, *Psychological Bulletin* 144, 140–176.

Zimmermann, Florian, 2020, The Dynamics of Motivated Beliefs, *American Economic Review* 110, 337–361.

Appendices

A. Experimental Instructions

A. Introduction

Welcome to our financial decision making study

For the duration of the study, we ask you to follow a few rules. Should there be questions, please raise your hand and an experimenter will answer your question privately. We ask you not to communicate with each other or use a calculator during the study.

We also ask you to turn off your cell phones and other devices, or at least to put them on silent, and to pack them away with your bag or belongings. We do not want you or other participants to be disturbed or distracted. If you do not adhere to these rules, this will lead to an automatic exclusion from the study and from payment.

The study consists of 3 stages and will last approximately 1.5 hours. You will perform different tasks of the study at different points in time, which last about 45 min today and 45 min next week.

After the study, you will receive a payout for your participation. The actual amount will depend on your decisions in the experiment and luck.

Everyone will earn 8 EUR for participating in this study. In addition, the computer will randomly pick three out of the present participants who get paid his or her earnings from one of the study's tasks.

Please press 'proceed' to continue with the general instructions.

-Next Page-

B. General Instructions

In this study you complete investment tasks, related to two securities: a risky security (i.e., a stock with risky payoffs) and a riskless security (i.e., a bond with a known payoff), and will provide estimates as to how good an investment in the risky security is.

Please click 'proceed' to continue with the detailed instructions for the tasks. Take your time to read the instructions carefully. Note that you cannot go back to previous pages. Please let us know if you have any questions.

-Next Page-

C. Stage 1 of the Study

Investment task 1

First, you will decide to invest in one of two securities for 12 periods: a risky security (i.e., a stock with risky payoffs) and a riskless security (i.e., a bond with a known payoff).

Either way, you start with an endowment of 60 EUR. In addition to this endowment, you will get payoffs from investing.

If you choose to invest in the bond, you get a payoff of 3.10 EUR for sure in each period.

If you choose to invest in the stock, you will receive a dividend in every period, which can be either positive or negative. A positive dividend is either 11, 13, or 15 EUR with equal probability. A negative dividend is either -1, -3, or -5 EUR with equal probability.

The stock can either be good or bad, and this will determine the likelihood of its dividend being positive or negative. If the stock is good then the probability of receiving a positive dividend is 60% and the probability of receiving a negative dividend is 40%. If the stock is bad then the probability of receiving a positive dividend is 40% and the probability of receiving a negative dividend is 60%.

If you decide to invest in the stock, you will have the possibility to choose between two stocks, Stock BLUE and Stock YELLOW. One of the stocks is good and one is bad. You will not know which type of stock you chose. You may be facing the good stock, or the bad stock, with equal probability.

The dividends of the stock are independent from period to period, but come from the same distribution. In other words, once you decided for a stock (BLUE or YELLOW) and it is a good stock, then in each period the odds of the dividend being positive are 60%, and the odds of it being negative are 40%. If the chosen stock is a bad stock then the probability of receiving the positive dividend is 40% and the probability of receiving the negative dividend is 60% in each period.

If you decide to invest in a stock you accumulate the dividends paid by the stock over 12 periods and if you invest in the riskless security you accumulate the known payoff over 12 periods, i.e. you earn 37.20 EUR for sure.

At the end of the task, you will be told how much you have accumulated. Your task earnings will be your accumulated payoffs plus your initial endowment of 60 EUR.

-Next Page-

Stock evaluation task

You will then see the dividends of the stock, no matter if you chose to invest in the stock or the bond. You will see either the dividends paid by the stock you chose, or – if you decided to invest in the bond – by one of the stocks randomly picked by the computer.

After that, we will ask you to tell us two things:

(1) what you think is the probability that the stock is the good one (the answer must be a number between 0 and 100);

(2) how much you trust your ability to come up with the correct probability estimate that the stock is good. In other words, we want to know how confident you are that the probability you estimated is correct.

There is always an objective, correct, probability that the stock is good, which depends on the history of dividends paid by the stock already. For instance, at the beginning of the task, the probability that the stock is good is exactly 50%, and there is no doubt about this value.

As you observe the dividends of the stock, you will update your belief whether or not the stock is good. It may be that after a series of good dividends, you think the probability of the stock being good is 75%.

However, how much you trust your ability to calculate this probability could vary. Sometimes you may not be too confident in the probability estimate you calculated and sometimes you may be highly confident in this estimate. For instance, at the very beginning of the task, the probability of the stock being good is 50% and you should be highly confident in this number because nothing else has happened since then.

If you provide us with a probability estimate that is within 5% of the correct value (e.g., correct probability is 80% and you say 84%, or 75%) you will earn 120 EUR in this task.

-Next Page-

Investment task 2

Further, you will again decide to invest in the stock you observed (Stock BLUE/Stock YELLOW) or in another bond for the next 12 periods.

Either way, you start with an endowment of 60 EUR. In addition to this endowment, you will get payoffs from investing.

If you choose to invest in the bond, you get a payoff of 5.10 EUR for sure in each period.

Again, if you choose to invest in the stock, you will receive a dividend in every period, which can be either positive or negative. A positive dividend is either 11, 13, or 15 EUR with equal probability. A negative dividend is either -1, -3, or -5 EUR with equal probability.

The dividends of the stock are independent from period to period, but **come from the same distribution as in Investment Task 1**. In other words, once you decided for a stock (BLUE or YELLOW) and it is a good stock, then in each period the odds of the dividend being positive are 60%, and the odds of it being negative are 40%. If the chosen stock is a bad stock then the probability of receiving the positive dividend is 40% and the probability of receiving the negative dividend is 60% in each period.

As in the Investment Task 1, if you decide to invest in the stock you accumulate the dividends of the stock over 12 periods and if you invest in the riskless security you accumulate the known payoff over 12 periods, i.e. you earn 61.20 EUR for sure.

At the end of the task, you will be told how much you have accumulated. Your task earnings will be your accumulated payoffs plus your initial endowment of 60 EUR.

-Next Page-

D. Stage 2 of the study

We will then ask you to complete three IQ test questions. The more questions you answer correctly, the more you earn in this task. For each correct answer you earn 40 EUR. For example, if you answer all three questions correctly, you earn 3 x 40 EUR i.e., 120 EUR. If you don't answer any question correctly, you will earn nothing.

-Next Page-

Your final payment at the end of the study

Your final payment will be:

You will get paid 8 EUR for participating in our study regardless of your task earnings.

In addition, your earnings in one of the experimental tasks can determine your payment. We will randomly draw three participants out of each session (with maximum 30 participants) who will get paid one of her or his task earnings. The computer will randomly decide which of the above-described tasks will determine the participants' payment. Remember, your task earnings depend on your decisions and answers:

Investment Task 1: Your initial endowment of 60 EUR and either 37.20 EUR from investing in the bond or accumulated dividends over 12 periods from investing the stock.

Stock Evaluation Task: Either 120 EUR if you provide us with a probability estimate that is within 5% of the correct value or nothing.

Investment Task 2: Your initial endowment of 60 EUR and either 61.20 EUR from investing in the bond or accumulated dividends over 12 periods from investing the stock.

IQ Test: Between 120 EUR and nothing; dependent on how many questions you answer correctly (40 EUR for each correct answer).

-Next Page-

E. Stage 3: Post-questionnaire

At the end of the experiment, we will ask you some personal questions. Note that all answers will be treated confidentially and will be analyzed anonymously.

B. Memory Elicitation

The memory elicitation task was not announced beforehand. We randomized the order of the specific recall questions, i.e. whether we first asked to recall positive or negative outcomes.

A. *Instructions*

Before we proceed to the next part, we ask you to complete a Recall Task, related to information you observed [last week]. Similar to the other experimental tasks, your answers can determine your final payment. This Recall Task can as well be selected by the computer for additional payment, which three of you will receive.

We will ask you 10 questions. The more questions you answer correctly, the more you earn in this task. For each correct answer, you earn 12 EUR. For example, if you answer all ten questions correctly, you earn 10 x 12 EUR i.e., 120 EUR. If you don't answer any question correctly, you will earn nothing.

-Next Page-

B. *Questions for positive and negative outcomes*

First, your task consists of recalling the stock dividends from last week. You observed 12 dividends of the stock.

How many positive dividends (11, 13, or 15) did you observe? [ ]

... and how many negative dividends (-1, -3, or -5) did you observe? [ ]

-Next Page-

C. *Questions for specific outcomes*

How often did the stock pay a dividend of -1 EUR? [ ]

How often did the stock pay a dividend of -3 EUR? [ ]

How often did the stock pay a dividend of -5 EUR? [ ]

-Next Page-

How often did the stock pay a dividend of 11 EUR? [ ]

How often did the stock pay a dividend of 13 EUR? [ ]

How often did the stock pay a dividend of 15 EUR? [ ]

C. Betting Market Instructions of Experiment 3

Now, Task 2 builds on your decision in Part 1 of the study (Investment Task 1, last week). This works as follows. All participants of this study participated in the same task as you did in Part 1, i.e., they chose one among two stocks. 9 randomly drawn participants and you (thus, 10 participants in total) take part in a **betting task** that relates to their decision in Part 1 of the study (Investment Task 1, last week).

Each participant is given a budget of 100 points for betting. You can decide how many points to **bet that you chose the better stock** among the two in Part 1 of the study (Investment Task 1, last week).

Every point you don't bet, you keep. The other 9 participants will also bet.

Your **bonus payment** is determined based on how much everyone bets and whether you/they chose the better stock. This works as follows:

- If you did not choose the better stock, every point you bet will be lost. You only keep the points you did not bet.
- If you chose the better stock, every point you bet will yield a positive profit for you. In this case, the bonus from this task is given by:
  Bonus = (Number of points you bet * Number of points bet by all participants / Number of points bet by participants who chose the better stock) + Number of points you did not bet.

While this may sound complicated, what it means is relatively simple: if you chose the better stock, you are guaranteed to earn back at least what you bet, but probably more (the more you bet).

Your potential bonus payment is 1/10 of your points in GBP.

Please place your bet. How many points do you want to bet that you chose the better stock among the two in Part 1 of the study (Investment Task 1, last week)?

Note, the better stock is the stock that is more likely to be the good stock. In case both stocks were equally likely to be the good stock, we count your decision as choosing the better stock regardless.

Click the bar below to reveal the slider. Drag it to indicate your answer.

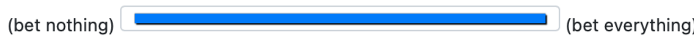

D. Subjects' Investment Options in Experiment 1

This table provides an overview of subjects' investment options during Experiment 1. Subjects made two investment decisions, a first choice before the observation of outcomes and a second choice after the observation of outcomes.

Table VIII. Overview of subjects' investment options

	Investment Option	**Risk about Asset Type**	**Asset Type**	**Possible Outcome(s)**	**Probability of Outcome(s)**	**Expected Outcome**
First Choice (Before Observation)	Stock	50% probability	Good Stock	11, 13, 15 EUR	60%	6.60 EUR
				-5, -3, -1 EUR	40%	
		50% probability	Bad Stock	11, 13, 15 EUR	40%	3.40 EUR
				-5, -3, -1 EUR	60%	
	Bond	No	-	3.10 EUR	100%	3.10 EUR
Second Choice (After Observation)	Stock	Based on Subjective posterior	Good Stock	11, 13, 15 EUR	60%	6.60 EUR
				-5, -3, -1 EUR	40%	
		Based on Subjective posterior	Bad Stock	-5, -3, -1 EUR	60%	3.40 EUR
				11, 13, 15 EUR	40%	
	Bond	No	-	5.10 EUR	100%	5.10 EUR

E. Additional Experimental Conditions to Identify the Underlying Mechanism (Experiment 2)

This table provides an overview of the additional experimental conditions and the number of subjects participated in Experiment 2 across laboratories.

Table IX. Number of subjects across labs and experimental conditions of Experiment 2.

Treatment	Groups	CESS	Xlab	Total
Baseline	*Immediate*	68	23	91
	Delay	64	30	94
HighStake	*Immediate*	42	35	77
	Delay	38	38	76
NoChoice	*Immediate*	34	40	75
	Delay	46	38	84

F. A Model of Memory-based Belief Distortion

Economic theory has begun to emphasize the role of selective recall in how people deal with ego-threatening information (Bénabou and Tirole, 2002, 2004) and to formalize consequences of memory limitations for economic choice (Bordalo, Gennaioli, and Shleifer, 2020; Mullainathan, 2002; Nagel and Xu, 2022). In this section, we present a simple behavioral mechanism for a positive memory bias and its effect on beliefs. In our model, the agent uses observed signals to update beliefs about an underlying state in a quasi-Bayesian way (e.g. Rabin, 2002). She behaves as a Bayesian updater, but places wrong weights on signals (e.g. Rabin and Schrag, 1999; Mobius, Niederle, Niehaus, and Rosenblat, 2014). In our setting, the biased signal weighting happens when the agent retrieves historical data, such as earnings surprises, returns, or dividends, from memory. The weights are biased in a self-serving way (e.g. Köszegi, 2006):[21] the agent relatively under-remembers signals that are inconsistent with her motivation, e.g., maintaining a positive self-image, and becomes over-optimistic about more preferred states.

A. The Model

Suppose in each period t, a stock generates an outcome $d_t \in \mathcal{D}$, where $\mathcal{D}$ is a finite ordered set of outcomes. Outcomes in all periods are *i.i.d.* In a financial market, these outcomes can be thought of as earnings surprises, returns, or dividends. Which outcome is generated in each period is determined by the stock's underlying type. The agent's task is to observe the outcomes and to make inference about the stock's type. For simplicity and for the convenience in our subsequent experimental testing, we use binary outcomes and a binary type space. The outcomes are either positive or negative: $d_t \in \{0, 1\}$, where 0 represents a negative outcome and 1 represents a positive outcome.[22] The stock is of one of 2 different types good or bad, represented by G and B respectively. Each type corresponds to an underlying distribution from which outcomes are drawn. The probability of a positive outcome is θ_G for a good stock, and θ_B for a bad stock, with $0 < \theta_B < \theta_G < 1$. In other words, the good type has an outcome distribution that first-order stochastically dominates the bad type. Let μ_0^G and μ_0^B represent the prior belief about the good and bad type respectively.

The true history of outcomes up to period t is represented by $h_t = (d_1, ..., d_t)$. The agent observes the history of outcomes h_t, stores the outcomes in memory as the history of remembered outcomes $h_t^R = (d_1^R, ..., d_t^R)$, and forms beliefs about the stock's type and future outcomes based on h_t^R. Up to period t, the total number of positive outcomes occurred is n_t^+, the total number of negative outcomes occurred is n_t^-.

[21] See Bénabou and Tirole (2016) for a review of this literature.
[22] This is without loss of generality, as long as the set of outcomes is ordered.

The total number of positive and negative outcomes sum up to t.[23] Let the number of remembered positive and negative outcomes be denoted as $n_t^{+,R}$ and $n_t^{-,R}$, respectively.

Assume that the agent always remembers the total number of periods correctly, thus $n_t^{+,R} + n_t^{-,R} = t$.[24] However, $n_t^{+,R}$ and $n_t^{-,R}$ may not be correct representations of n_t^{+} and n_t^{-}. If the agent invests in the stock, she may relatively under-remember negative outcomes compared to positive outcomes.[25] This memory bias could emerge because obtaining negative outcomes might suggest that the initial investment decision was wrong and positive outcomes might justify the initial investment decision. In other words, negative outcomes might be inconsistent with the agent's motivation, e.g., maintaining a positive self-image, and positive outcomes might align with the agent's motivation. In a similar vein, if the agent does not invest in the stock, she may relatively under-remember positive outcomes compared to negative outcomes. However, the effect does not have to be symmetric. For instance, the decision to invest could be of higher importance for the agent's motivation, e.g., the maintenance of a positive self-image, compared to the decision not to invest. Many studies document the power of action to alter self-views and provide evidence for an asymmetry between the influence of action and non-action on self-perception (Cioffi and Garner, 1996; Allison and Messick, 1988; Fazio et al., 1982).

Suppose R_t represents the memory bias in period t, given the agent's initial choice to invest in the stock and dependent on whether the outcome is positive or negative:

$$R_t = \delta_t [I_t q^- + (1 - I_t) q^+], \tag{A.1}$$

where $\delta_t \in [0, 1]$ captures the degree of memory bias depending on whether an active choice (to invest or not to invest in a given stock) is made in period t. If the agent makes an active choice, $\delta_t = 1$ so that there is full memory bias; if the agent does not make an active choice, for instance if the stock position is inherited, $\delta_t \in [0, 1)$, so that the memory bias is lower and it is possible that no memory bias emerges ($\delta_t = 0$). I_t is an indicator function that is equal to 1 if the agent invests in the stock in period t, and 0 otherwise. Additionally, $q^- \in [0, 1]$ is the probability of under-remembering a negative outcome and $q^+ \in [0, 1]$ is the probability of under-remembering a positive outcome. The assumption here is that if the agent actively invests, she faces a probability q^- of forgetting each negative outcome, but she does not mis-remember positive outcomes. Thus, she relatively under-remembers negative compared to positive outcomes. In case the agent invests (does not invest), if $q^- = 0$ ($q^+ = 0$) there is no memory bias, and if $q^- = 1$ ($q^+ = 1$) the memory bias

[23] Under the current notation, we also conveniently have $n_t^+ = \sum_0^t d_t$, and $n_t^- = t - \sum_0^t d_t$.

[24] This means if the agent relatively over-remembers positive outcomes, she must relatively under-remember negative outcomes. Of course the agent could also under-remember all outcomes in general, but this memory decay is not what we would like to pursue.

[25] Relative mis-remembering a kind of outcome relative to the other is the phenomenon of interest here. Uniform mis-remembering of both positive and negative outcomes is not explored here.

is extreme, i.e., the outcomes are forgotten completely. It is possible that $q^- = q^+$, or $q^- \neq q^+$. That is, the agent may or may not have the same bias if she invests or if she does not. Further, it is possible that either q^- or q^+ is 0, which means the agent has a memory bias only if she invests in the stock, or only if she does not invest in the stock. We will demonstrate this point in our experimental results. And finally, when $q^- = q^+ = 0$, the model reverts to rational Bayesian updating.

For simplicity, in what follows we focus on the case where an active choice is made, or $\delta_t = 1$.[26] If the agent invests in the stock in period 1 to t, then in period t, the remembered number of negative outcomes is expected to be $n_t^{-,R} = (1 - q^-)n_t^- \leq n_t^-$, and the remembered number of positive outcomes is expected to be $n_t^{+,R} = t - n_t^{-,R} \geq n_t^+$. If the agent does not invest in the stock, then in period t, the remembered number of positive outcomes is expected to be $n_t^{+,R} = (1 - q^+)n_t^+ \leq n_t^+$, and the remembered number of negative outcomes is expected to be $n_t^{-,R} = t - n_t^{+,R} \geq n_t^-$. Note that equality holds if there is no outcome that contradicts with the agent's motivation. This is the case when $n_t^- = 0$ if the agent invests, and $n_t^+ = 0$ if the agent does not invest.

An unbiased agent uses the true history h_t to update her belief about the stock's type. The posterior belief that the stock's type is G of an unbiased agent is represented by:

$$\mu_t^{G,Bayesian}(h_t) = \frac{P(h_t|\theta_G)\mu_0^G}{\sum_{j=G,B} P(h_t|\theta_j)\mu_0^j}.$$

The likelihood ratio of type G relative to type B is given by:

$$\Lambda^{Bayesian}(h_t) = \frac{\theta_G^{n_t^+}(1 - \theta_G)^{n_t^-}\mu_0^G}{\theta_B^{n_t^+}(1 - \theta_B)^{n_t^-}\mu_0^B}. \tag{A.2}$$

A biased agent uses the history of remembered outcomes h_t^R to update her belief about the stock. The posterior belief that the stock's type is G of a biased agent is represented by:

$$\mu_t^{G,Biased}(h_t^R) = \frac{P(h_t^R|\theta_G)\mu_0^G}{\sum_{j=G,B} P(h_t^R|\theta_j)\mu_0^j}.$$

The likelihood ratio of type G relative to type B for a biased agent is

$$\Lambda^{Biased}(h_t^R) = \frac{\theta_G^{n_t^{+,R}}(1 - \theta_G)^{n_t^{-,R}}\mu_0^G}{\theta_B^{n_t^{+,R}}(1 - \theta_B)^{n_t^{-,R}}\mu_0^B}, \tag{A.3}$$

If the agent invests in the stock $n_t^{+,R} \geq n_t^+$, and thus $\Lambda^{Biased}(h_t^R) \geq \Lambda^{Bayesian}(h_t)$: the biased agent overestimates the good type relative to the bad type compared to a Bayesian agent. Conversely, if the agent

[26] In our experimental results we will demonstrate the difference in memory bias between an active choice and a passive endowment.

does not invest in the stock $n_t^{+,R} \leq n_t^+$, and thus $\Lambda^{Biased}(h_t^R) \leq \Lambda^{Bayesian}(h_t)$. Again, the biased agent updates her beliefs like a Bayesian agent if there is actually no outcome that contradicts with her motivation.

In general, after a biased agent invests in the stock, her memory of observed outcomes will be biased. She relatively under-remembers less preferred outcomes, and thus when updating beliefs, she becomes overly optimistic about the underlying type of the stock. The opposite might happen if the biased agent does not invest in the stock. However, as stated before, the memory bias might not be symmetric for the cases in which the agent invests and does not invest. Note that when the probability of under-remembering (q^- or q^+) is zero, Equation (A.3) and Equation (A.2) coincide.

The agent with a motivated memory bias can also be seen as placing biased weights on past outcomes when retrieving them from memory. If she invests in the stock, she relatively overweights positive outcomes, and overestimates the probability that the stock is of the good type. By contrast, if she does not invest in the stock, she relatively overweights negative outcomes, and overestimates the probability that the stock is of the bad type.

B. *Signal Weighting*

Without loss of generality, suppose that $\theta_G = \theta$, and $\theta_B = 1 - \theta$, where $\theta \in (0.5, 1)$. Additionally, the two types are equally likely, $\mu_0^G = \mu_0^B = \mu_0 = 0.5$.[27] This further simplifies matters and helps us to derive simple representations of biased signal weightings, which are testable in our experiment.

Rewriting Equation (A.2), an unbiased agent's posterior likelihood ratio of G relative to B is

$$\Lambda^{Bayesian}(h_t) = \frac{\theta^{n_t^+}(1-\theta)^{n_t^-}}{\theta^{n_t^-}(1-\theta)^{n_t^+}}$$

For a biased agent, if she invests in the stock, the posterior likelihood ratio of G relative to B is

$$\Lambda^{Biased,INV}(h_t^R) = \frac{\theta^{t-(1-q^-)n_t^-}(1-\theta)^{(1-q^-)n_t^-}}{\theta^{(1-q^-)n_t^-}(1-\theta)^{t-(1-q^-)n_t^-}} \geq \Lambda^{Bayesian}(h_t)$$

If the biased agent does not invest in the stock, the posterior likelihood ratio of G relative to B is

$$\Lambda^{Biased,NOT}(h_t^R) = \frac{\theta^{(1-q^+)n_t^+}(1-\theta)^{t-(1-q^+)n_t^+}}{\theta^{t-(1-q^+)n_t^+}(1-\theta)^{(1-q^+)n_t^+}} \leq \Lambda^{Bayesian}(h_t)$$

A Bayesian agent uses the correct number of occurred positive and negative outcomes to update her beliefs, placing equal weight on positive and negative outcomes. As in Equation (B.1), the weight should be exactly equal to $\ln(\frac{\theta}{1-\theta})$, which represents the informativeness of positive relative to negative signals.

[27] In the experiment, we also have binary signals generated from the underlying distribution, either positive or negative outcomes. We used three positive and three negative outcomes, respectively, that are equally likely. However, this is irrelevant for a Bayesian updater when forming beliefs. Beliefs should be solely formed based on whether the signal is positive or negative.

$$\ln \Lambda^{Bayesian}(h_t) = \ln(\frac{\theta}{1-\theta})n_t^+ - \ln(\frac{\theta}{1-\theta})n_t^-. \quad \text{(B.1)}$$

However, the agent with a motivated memory bias has the following log likelihood ratios

$$\ln \Lambda^{Biased,INV}(h_t^R) = \ln(\frac{\theta}{1-\theta})(n_t^+ + q^- n_t^-) - \ln(\frac{\theta}{1-\theta})(1-q^-)n_t^-;$$

$$\ln \Lambda^{Biased,NOT}(h_t^R) = \ln(\frac{\theta}{1-\theta})(1-q^+)n_t^+ - \ln(\frac{\theta}{1-\theta})(n_t^- + q^+ n_t^+).$$

It can be shown that the following equations characterize the difference between the log likelihood ratios of a biased agent and a Bayesian agent.

$$\ln \Lambda^{Biased,INV}(h_t^R) - \ln \Lambda^{Bayesian}(h_t) = \ln(\frac{\theta}{1-\theta})2q^- n_t^-, \quad \text{(B.2)}$$

$$\ln \Lambda^{Biased,NOT}(h_t^R) - \ln \Lambda^{Bayesian}(h_t) = -\ln(\frac{\theta}{1-\theta})2q^+ n_t^+. \quad \text{(B.3)}$$

Hence, the belief distortion, measured as the log-likelihood deviation from the Bayesian belief, is positively correlated with θ, q^- (or q^+), and n_t^- (or n_t^+). The magnitude of the belief distortion is non-zero, when $q^- n_t^- \neq 0$ (or $q^+ n_t^+ \neq 0$), as $\theta \in (0.5, 1)$.

Equations (B.2) and (B.3) yield the following propositions.

Proposition 1: *Given $q^- n_t^- \neq 0$ (or $q^+ n_t^+ \neq 0$), the magnitude of the belief distortion is positively correlated with the contrast between the underlying processes (θ).*

Proposition 2: *Given $q^- \neq 0$ (or $q^+ \neq 0$), the magnitude of the belief distortion is positively correlated with the number of signals inconsistent with the agent's motivation (n^- or n^+).*

Proposition 3: *Given $n_t^- \neq 0$ (or $n_t^+ \neq 0$), the magnitude of the belief distortion is positively correlated with the magnitude of the memory bias (q^- and q^+).*

C. Comparative Statics

Our model proposes interesting comparative statics. The model predicts that holding θ, q^- and q^+ constant, the belief distortion compared to the Bayesian posterior is larger if there are more outcomes that are contradictory to the biased agent's motivation (Proposition 2). These are the negative outcomes if the agent invests and the positive outcomes when the agent does not invest. Further, the model predicts that holding θ, n^- and n^+ constant, the belief distortion relative to the Bayesian posterior is larger for agents with a larger memory bias (Proposition 3).

In line with the model propositions, we show that the belief distortion of subjects who invested in the stock is larger if more negative outcomes occurred, as indicated in Proposition 2. Figure 5 displays this finding for subjects who invested in the stock. The bars represent the mean value of subjects' belief distortion for different numbers of observed negative outcomes (n^-). The red line indicates the prediction for a linear regression. The figure illustrates a significant positive trend ($p < 0.001$), in line with our Proposition 2. The magnitude of the belief distortion of subjects who invested in the stock is positive as soon as subjects observed more than four negative outcomes. Further, the belief distortion increases significantly with a growing number of observed negative outcomes.

Figure 5. Belief Distortion and Observed Outcomes

This figure displays mean values of subjects' belief distortion measured by the difference between the posterior log-likelihood ratios of subjects' elicited probabilities and the objective Bayesian probabilities (*Belief Distortion*). The sample is limited to subjects who invested in the stock (first choice). The bars represent the mean values for different numbers of individually observed negative outcomes by subjects (n^-). Error bars indicate 95% confidence intervals. The red line represents the prediction for a linear regression of *Belief Distortion* on a Dummy variable for different numbers of n^-.

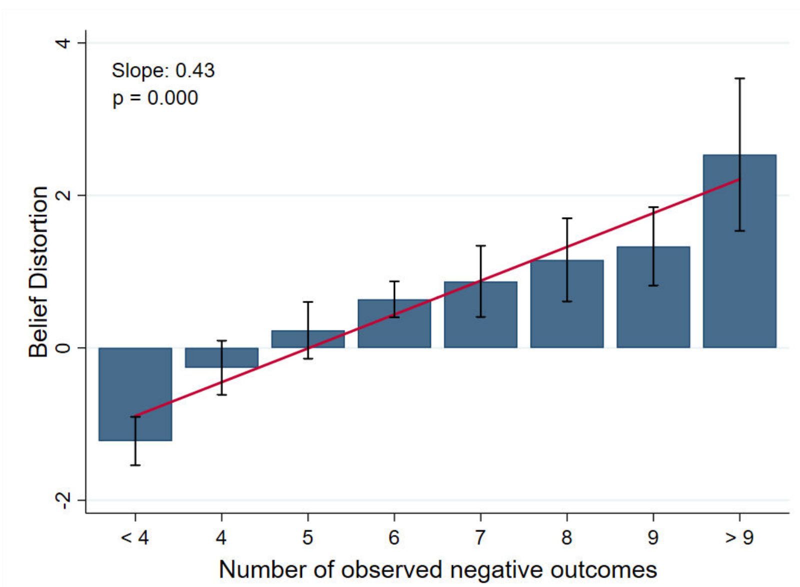

G. Subjective Memory Bias Across Different Specifications

In this Appendix sections, we provide additional results on subjects' memory bias across different specifications and treatments. First, Table X shows subjects' memory bias in Experiment 1, including subjects who did not invest in the stock.

Table X. Subjective Memory Bias Across Treatments (including subjects who did not invest in the stock)

This table displays subjects' memory bias in the *Delay* and *Immediate* condition, separately for subjects who invested in the stock and subjects who did not invest in the stock, i.e., invested in the bond. Memory bias is estimated at the individual level by subtracting the actual individually observed number of positive (negative) outcomes from subject's recalled number of the stock's positive (negative) outcomes. The table reports mean values and T-test results against the null hypothesis that the memory bias is zero (columns 1-8) and T-test results of the difference in means between our conditions *Delay* and *Immediate* for subjects who invested in the stock (column 9).

Observed outcomes	***Delay***				***Immediate***				Difference (if invested)
	Invested (N = 74)	T-test	**Not invested** (N = 18)	T-test	**Invested** (N = 78)	T-test	**Not invested** (N = 18)	T-test	T-test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Positive outcomes	0.89	p = 0.001	-0.06	p = 0.918	0.27	p = 0.058	-0.11	p = 0.668	p = 0.018
Negative outcomes	-0.73	p = 0.000	-0.22	p = 0.664	-0.28	p = 0.037	-0.06	p = 0.875	p = 0.038

Second, Table XI shows subjects' memory bias in Experiment 2, separately for positive and negative outcomes. Note, in Experiment 2 we forced subjects' responses in the memory elicitation task to always add up to 12 outcomes (i.e., the true number of outcomes observed). Thus, subjects' memory bias for positive and negative outcomes yield the same finding, just with opposite signs.

Table XI. Subjective Memory Bias in Experiment 2 (for pos. and neg. outcomes separately)

This table displays subjects' mean memory bias for positive and negative outcomes (compared to the true number of outcomes observed) in the *Delay* group as well as the treatment effect (the difference in memory bias for positive outcomes in the *Delay* group vs. *Immediate* group) in Experiment 2. Memory bias is estimated at the individual level by subtracting the actual individually observed number of positive/negative outcomes from subject's recalled number of the stock's positive/negative outcomes. The table reports mean values and T-test results against the null hypothesis that the memory bias is zero as well as the T-test results of the difference in means between *Delay* and *Immediate*. The sample is restricted to subjects in *Baseline*.

Baseline condition (N = 185)	**Mean Memory Bias in Delay condition**	**Treatment Effect (Delay vs. Immediate)**
For Pos. Outcomes	0.36 (p = 0.023)	0.48 (p = 0.010
For Neg. Outcomes	-0.36 (p = 0.023)	-0.48 (p = 0.010)

Third, Table XII shows subjects' memory bias in Experiment 2, separately for subjects in the active choice setting (*Baseline* condition) and subjects in the passive choice setting (*NoChoice* condition).

Table XII. Memory Bias in Active vs. Passive Choice Settings

This table displays subjects' memory bias in the *Delay* and *Immediate* group, separately for subjects in the *Baseline* and *NoChoice* condition. Memory bias is estimated at the individual level by subtracting the actual individually observed number of positive outcomes from subject's recalled number of the stock's positive outcomes. The table reports mean values and T-test results against the null hypothesis that the memory bias is zero as well as the T-test results of the difference in means between *Delay* and *Immediate*.

Condition	**Mean bias in *Delay* (vs. true number of outcomes)**	**Mean bias in *Immediate* (vs. true number of outcomes)**	**Treatment effect (difference between *Delay* and *Immediate*)**
Baseline (N = 94)	0.36 (p = 0.023)	-0.12 (p = 0.212)	0.48 (p = 0.010)
NoChoice (N = 84)	0.18 (p = 0.398)	-0.31 (p = 0.008)	0.49 (p = 0.051)

H. Model-based Measure of the Memory Bias

In accordance with this model, we test for subjects' memory bias based on how subjects weight observed outcomes when forming beliefs about the stock. Equations (B.2) and (B.3) provide the foundation for our regression analysis with regard to memory-based belief distortion. With subjects' belief distortion as the dependent variable, $\ln(\frac{\theta}{1-\theta})2n_t^-$ or $-\ln(\frac{\theta}{1-\theta})2n_t^+$ as the independent variable, we can directly estimate the value of q^- and q^+, i.e., the memory bias, on average across subjects. In the model, q^- is the probability of under-remembering a negative outcome if invested in the stock and q^+ is the probability of under-remembering a positive outcome if not invested. If q^- or q^+ equals zero, the probability of under-remembering an outcome is zero and if it equals one, the probability of under-remembering an outcome is 100%.

Table XIII. Estimated Memory Bias Based on Relative Weighting of Observed Outcomes

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable is subjects' belief distortion, *Belief Distortion* (*Belief Dist.*), measured by the difference between the posterior log-likelihood ratios of subjects' elicited probabilities and the objective Bayesian probabilities. As independent variables, we use the combination of the informativeness of evidence (the log likelihood ratio $\ln \frac{\theta}{1-\theta}$) and the strength of evidence (n^- or n^+) in our experimental situation, according to Equations (B.2) and (B.3). *Session* is a dummy variable representing the different sessions of the experiment. Regression coefficients represent parameters q^- and q^+. We present the regression models for our treatments *Delay* (*Del.*) and *Immediate* (*Imm.*) as well as for subjects who invested and who did not invest separately. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Belief Dist. (*Del.*, inv.)	(2) Belief Dist. (*Del.*, not inv.)	(3) Belief Dist. (*Imm.*, inv.)	(4) Belief Dist. (*Imm.*, not inv.)
$ln(\frac{\theta}{1-\theta})2n_t{}^-$	0.599***		0.473***	
	(0.07)		(0.07)	
$-ln(\frac{\theta}{1-\theta})2n_t{}^+$		0.341		0.147
		(0.27)		(0.36)
Constant	-3.189***	0.686	-1.880***	1.830
	(0.54)	(1.81)	(0.50)	(2.02)
Session	Yes	Yes	Yes	Yes
N	73	18	73	19
R^2	0.55	0.58	0.46	0.40

We find that subjects who invested in the stock relatively underweight negative compared to positive outcomes, resulting in overly optimistic beliefs. Regressions in Table XIII estimate q^- and q^+. The table reports results from regressions with subject's belief distortion as dependent variable. As independent variable, we use a combination of the informativeness of evidence (the log likelihood ratio $\ln \frac{\theta}{1-\theta}$) and the strength of evidence (n^- or n^+) in our experimental situation, according to Equations (B.2) and (B.3) and control for session fixed effects. Thus, the regression coefficients represent our estimations for q^- and q^+ from our model based on subjects' relative under-weighting of negative and positive outcomes n^- and n^+ when forming beliefs about the stock. The regressions are reported separately for subjects who invested and for those who did not invest as well as across our two treatments *Delay* and *Immediate*.

Our results show that subjects who invested in the stock (columns 1 and 3) have a significantly positive probability of under-remembering negative outcomes, and this probability is higher in the *Delay* condition (59.9%, $p < 0.001$) than in the *Immediate* condition (47.3%, $p < 0.001$). Subjects who did not invest in the stock (columns 2 and 4), do not display a significant probability to under-remember positive outcomes. This is consistent with the proposed behavioral mechanism in our model as well as our previous results based on subjects' directly elicited memory bias (Table II). Thus, we find that subjects' belief distortion stems from relatively under-remembering negative outcomes when forming beliefs about the stock. This finding does not rely on a specific memory elicitation method. This finding and the observation that subjects' elicited memory bias is significantly correlated with their belief distortion (Table III) is supportive of our Proposition 3.

I. Difference between Subjective and Objective Posterior

Table XIV reports mean values and T-tests for differences in means between subjects' reported subjective and objective (Bayesian) posteriors in the *Delay* condition of Experiment 1.

Table XIV. Difference between Subjective and Objective Posterior in Experiment 1

Delay treatment (N = 92)	**Mean**	**Difference**
Subjective posterior	58.01	8.16
Objective posterior (Bayesian)	49.85	(p = 0.004)

J. T-Tests for Treatment Effects in Reinvestment and Confidence

In this section we report mean values of subjects' reinvestments and confidence levels as well as T-tests for differences in group means between our conditions *Delay* and *Immediate*.

Table XV. Memory-Based Reinvestment Decisions

This table displays subjects' investment decisions before (first choice) and after (second choice) observing the stock's outcomes in the *Delay* and *Immediate* condition of Experiment 1. *Investment in stock* is a dummy variable equal to one if the subjects chose to invest in the stock.*Reinvestment* is a dummy variable equal to one if the subjects chose to reinvest in the stock. The table reports mean values and T-tests for differences in group means between our conditions *Delay* and *Immediate*. *Reinvestment (suboptimal)* is a dummy variable equal to one if the subject reinvested in the stock with a lower expected outcome than the bond after the observation phase (second choice).

	First choice		Second choice			
	Investment in stock	N	Reinvestment	N	Reinvestment (suboptimal)	N
	(1)		(2)		(3)	
Immediate	80.41	97	46.15	78	20.51	78
Delay	80.43	92	64.86	74	41.89	74
Difference (T-test)	p = 0.997		p = 0.020		p = 0.008	

The results of Table XV show that in Experiment 1, 41.9% of the subjects in the *Delay* condition reinvest in the stock despite its lower expected outcome, while in the *Immediate* condition only half of this fraction (20.5% of the subjects) reinvest in such a manner.

The results of Table XVI show that on average subjects in Experiment 3 bet 46.3 points in the *Delay* condition despite the choice being suboptimal, while in the *Reminder* condition subjects only bet 31.5 points. Subjects bet on average 47% more given suboptimal decisions when they must rely on their memory from last week (*Delay* condition) compared to the control group (*Reminder* condition).

Table XVI. Memory-Based Betting

This table displays how many points subjects bet in the *Delay* and *Reminder* condition. *Points bet* is the number of points bet on the optimality of own investment choice by subjects. *Points bet (suboptimal)* is the number of points bet suboptimally on the optimality of own investment choice by subjects. The table reports mean values and T-tests for differences in group means between our conditions *Delay* and *Reminder*.

	Second choice			
	Points bet	N	Points bet (suboptimal)	N
	(1)		(2)	
Control (*Reminder*)	47.71	255	31.54	108
Delay	57.09	232	46.33	83
Difference (T-test)	$p = 0.001$		$p < 0.001$	

K. Evidence for Reinforcement Learning

In previous empirical work, it has been shown that investors engage in reinforcement learning, meaning that they (i) repeat actions that generated previous gains and (ii) avoid actions that generated previous losses. We find that subjects in our experiment behave in this way (see Table XVII). The objective Bayesian posterior (calculated based on all experienced outcomes of the subject) is significantly positively correlated with subjects' probability to reinvest. In other words, experienced positive outcomes increase subjects' probability to reinvest, and experienced negative outcomes decrease their probability to reinvest.

Table XVII. Subjective Memory and Reinvestment Decisions

This table contains the odds ratios and standard errors (in parentheses) of Logit regressions in which the dependent variable is a dummy variable which is equal to one if the subject invested in the stock after the observation phase (second choice) in Experiment 1. *Treatment (Delay)* is a dummy variable representing our treatment with 1 = Delay condition and 0 = Immediate condition. *Objective Probability* is the value of the objective Bayesian probability that the stock is the good stock (from 1 to 100). *Session* is a dummy variable representing the different sessions of the experiment. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Reinvestment	(2) Reinvestment
Objective Probability	1.013**	1.015**
	(0.01)	(0.01)
Treatment (*Delay*)		2.336**
		(0.85)
Session	Yes	Yes
N	184	152
Pseudo R^2	0.03	0.08

L. Robustness Tests

In this Appendix section, we provide the detailed analyses for Section VII in the main text. Table XVIII and Table XIX show robustness of the memory bias in Experiment 1 to different measures of the bias, namely based on recalled fractions and the absolute difference between recalled positive and negative outcomes.

Table XVIII. Subjective Memory Bias based on the Recalled Fraction

	Delay				*Immediate*				Difference (if invested)
	Invested (N = 74)	T-test	Not invested (N = 18)	T-test	Invested (N = 78)	T-test	Not invested (N = 18)	T-test	T-test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Positive outcomes	0.07	p = 0.000	0.01	p = 0.846	0.02	p = 0.042	-0.00	p = 0.953	p = 0.015
Negative outcomes	-0.07	p = 0.000	-0.01	p = 0.846	-0.02	p = 0.042	0.00	p = 0.953	p = 0.015

Notes: This table displays subjects' memory bias in the *Delay* and *Immediate* condition, seperately for subjects who invested in the stock and subjects who did not invest in the stock, i.e., invested in the bond. Memory bias is estimated at the individual level by subtracting the actual individually observed fraction of positive (negative) outcomes relative to the 12 observed outcomes from subject's recalled fraction of the stock's positive (negative) outcomes relative to the 12 observed outcomes. The table reports mean values and T-tests against the null hypothesis that the memory bias is zero (columns 1-8) and of differences in group means between our conditions *Delay* and *Immediate* (column 9).

Table XIX. Subjective Memory Bias based on the Recalled Absolute Difference

	Delay				*Immediate*				Difference (if invested)
	Invested (N = 74)	T-test	Not invested (N = 18)	T-test	Invested (N = 78)	T-test	Not invested (N = 18)	T-test	T-test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Difference between positive and negative outcomes	1.62	p = 0.000	0.16	p = 0.871	0.55	p = 0.045	-0.06	p = 0.927	p = 0.013

Notes: This table displays subjects' memory bias in the *Delay* and *Immediate* condition, separately for subjects who invested in the stock and subjects who did not invest in the stock, i.e., invested in the bond. Memory bias is estimated at the individual level by subtracting the actual difference between observed positive and negative outcomes from subject's recalled difference between the stock's positive and negative outcomes. The table reports mean values and T-tests for differences in group means between our conditions *Delay* and *Immediate* (column 9).

Table XX shows that subjects' risk preferences do not correlate with the memory bias in Experiment 1, separately for positive and negative outcomes.

Table XX. Subjective Memory Bias and Individual Risk Preferences

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variables are subject's memory bias for observed positive and negative outcomes. Both variables, *Memory Bias (for Pos. Outcomes)* and *Memory Bias (for Neg. Outcomes)*, are estimated at the individual level by subtracting the actual individually observed number of positive (negative) stock outcomes from subject's recalled number of the stock's positive (negative) outcomes. *Risk Tolerance* represents subjects' self-stated general risk preferences on a 10-point scale from 0 (lowest) to 10 (highest) (Dohmen et al., 2011). *Session* is a dummy variable representing the different sessions of the experiment. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Memory Bias (for Pos. Outcomes)	(2) Memory Bias (for Neg. Outcomes)
Risk Tolerance	-0.003	0.030
	(0.05)	(0.04)
Constant	-0.136	-0.082
	(0.43)	(0.38)
Session	Yes	Yes
N	187	187
R^2	0.06	0.03

Table XXI. Robustness of Main Results to Individual Risk Preferences and Financial Literacy

In columns 1 and 2 this table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable is subjects' belief distortion measured by the difference between the posterior log-likelihood ratios of subjects' elicited probabilities and the objective Bayesian probabilities, *Belief Distortion*. In columns 3 and 4 this tables contains the odds ratios and standard errors (in parentheses) of Logit regressions in which the dependent variable is a dummy variable which is equal to one if the subject invested in the stock with a lower expected outcome than the bond after the observation phase (second choice), *Investment (Subopt.)*. *Memory Bias (for Pos. Outcomes)* represents subject's memory bias for observed positive outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of positive stock outcomes from subject's recalled number of the stock's positive outcomes. *Memory Bias (for Neg. Outcomes)* represents subject's memory bias for observed negative outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of negative stock outcomes from subject's recalled number of the stock's negative outcomes. *Risk Tolerance* represents subjects' self-stated general risk preferences on a 10-point scale from 0 (lowest) to 10 (highest) (Dohmen et al., 2011). *Financial Literacy* is a dummy variable equal to one for subjects who indicated the correct answer to a financial literacy question used in Kuhnen (2015). *Session* is a dummy variable representing the different sessions of the experiment. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Belief Distortion	(2) Belief Distortion	(3) Investment (Subopt.)	(4) Investment (Subopt.)
Memory Bias (for Pos. Outcomes)	0.427***		1.607***	
	(0.05)		(0.21)	
Memory Bias (for Neg. Outcomes)		-0.528***		0.635***
		(0.06)		(0.09)
Risk Tolerance	-0.019	-0.003	1.103	1.111
	(0.04)	(0.04)	(0.08)	(0.08)
Financial Literacy	-0.155	-0.166	1.449	1.500
	(0.21)	(0.21)	(0.62)	(0.63)
Constant	-0.046	-0.149	0.026***	0.022***
	(0.32)	(0.31)	(0.03)	(0.03)
Session	Yes	Yes	Yes	Yes
N	181	181	187	187
R^2	0.31	0.35		
Pseudo R^2			0.13	0.11

Table XXI shows that the key findings of Experiment 1 are robust to controlling for subjects' self-reported risk preferences as well as to measured financial literacy. Further, Table XXII reports the robustness of our findings in Table II of Section III, by separately testing the significance of memory bias for subjects with different levels of financial literacy.

Table XXII. Subjective Memory Bias and Individual Financial Literacy

	Financial Literacy Question Correct				*Financial Literacy Question Incorrect*			
	Delay (N = 23)	T-test	Immediate (N = 15)	T-test	Delay (N = 51)	T-test	Immediate (N = 62)	T-test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Positive outcomes	1.04	p = 0.013	0.60	p = 0.108	0.82	p = 0.004	0.17	p = 0.252
Negative outcomes	-0.78	p = 0.005	-0.60	p = 0.108	-0.71	p = 0.002	-0.19	p = 0.182

Notes: This table displays subjects' memory bias in the *Delay* and *Immediate* condition, separated by subjects' financial literacy measured in the post-questionnaire Kuhnen (2015). Memory bias is estimated at the individual level by subtracting the actual individually observed number of positive (negative) outcomes from subject's recalled number of the stock's positive (negative) outcomes. The table reports mean values and T-tests against the null hypothesis that the memory bias is zero.

Moreover, we find that subjects' memory bias is not correlated with their financial knowledge (Table XXIII).

Table XXIII. Subjective Memory Bias and Individual Financial Literacy

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variables are subject's memory bias for observed positive and negative outcomes. Both variables, *Memory Bias (for Pos. Outcomes)* and *Memory Bias (for Neg. Outcomes)*, are estimated at the individual level by subtracting the actual individually observed number of positive (negative) stock outcomes from subject's recalled number of the stock's positive (negative) outcomes. *Financial Literacy* is a dummy variable equal to one for subjects who indicated the correct answer to a financial literacy question used in Kuhnen (2015). *Session* is a dummy variable representing the different sessions of the experiment. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Memory Bias (for Pos. Outcomes)	(2) Memory Bias (for Neg. Outcomes)
Financial Literacy	0.248	-0.217
	(0.29)	(0.25)
Constant	-0.237	0.126
	(0.38)	(0.34)
Session	Yes	Yes
N	187	187
R^2	0.06	0.03

Table XXIV shows that our experimental findings from Experiment 1 are robust to the differences in timing of the experimental tasks. We control for the order of tasks, which was randomized across subjects.

Table XXIV. Robustness of Main Results to Timing in Immediate Condition

In columns 1 and 2 this table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable is subjects' belief distortion measured by the difference between the posterior log-likelihood ratios of subjects' elicited probabilities and the objective Bayesian probabilities, *Belief Distortion*. In columns 3 and 4 this tables contains the odds ratios and standard errors (in parentheses) of Logit regressions in which the dependent variable is a dummy variable which is equal to one if the subject invested in the stock with a lower expected outcome than the bond after the observation phase (second choice), *Investment (Subopt.)*. *Memory Bias (for Pos. Outcomes)* represents subject's memory bias for observed positive outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of positive stock outcomes from subject's recalled number of the stock's positive outcomes. *Memory Bias (for Neg. Outcomes)* represents subject's memory bias for observed negative outcomes. This variable is estimated at the individual level by subtracting the actual individually observed number of negative stock outcomes from subject's recalled number of the stock's negative outcomes. *Timing* is a dummy variable equal to one for sessions in which subjects in the *Immediate* condition perform the tasks in week t and zero for sessions in which subjects in the *Immediate* condition perform the tasks in week $t+1$. *Session* is a dummy variable representing the different sessions of the experiment. *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Belief Distortion	(2) Belief Distortion	(3) Investment (Subopt.)	(4) Investment (Subopt.)
Memory Bias (for Pos. Outcomes)	0.424***		1.563***	
	(0.05)		(0.19)	
Memory Bias (for Neg. Outcomes)		-0.525***		0.661***
		(0.06)		(0.09)
Timing	-0.401	-0.544	0.140*	0.128*
	(0.36)	(0.35)	(0.16)	(0.14)
Constant	0.216	0.322	0.358**	0.353**
	(0.24)	(0.23)	(0.15)	(0.15)
Session	Yes	Yes	Yes	Yes
N	182	182	188	188
R^2	0.31	0.35		
Pseudo R^2			0.12	0.10

In addition, we included a *NoRecall* condition, in which subjects do not perform a memory elicitation task as a control condition. This control condition allows us to identify potential effects of simply asking subjects to recall investment outcomes. Table XXV indicates no significant difference in subjects' beliefs or investment decisions between the *Delay* and *NoRecall* condition (T-test, $p = 0.714$ and $p = 0.343$, respectively).

Table XXV. Comparison of Subjective Beliefs and Choices between *Delay* and *NoRecall*

	Delay		*NoRecall*		Differences
	Mean	N	Mean	N	T-test
	(1)		(2)		(3)
Subjective Probability	58.011	92	59.725	40	p = 0.714
Investment in stock (second choice)	0.565	92	0.475	40	p = 0.343
Investment in stock (second choice, suboptimal)	0.544	92	0.435	40	p = 0.383

Notes: This table displays subjective beliefs as well as subjects' investment choice after the observation phase (second choice) in the *Delay* and *NoRecall* condition. Subjective beliefs are subject's indicated probability that the observed stock is the good stock (1 to 100). Investment in stock is a dummy variable which is equal to one if the subject invested in the stock after the observation phase (second choice); Investment in stock (suboptimal) is a dummy variable which is equal to one if the subject invested in the stock with a lower expected outcome than the bond after the observation phase (second choice). The table reports mean values and T-tests for differences in group means between our conditions *Delay* and *NoRecall* (column 3).

M. Replication of the Findings of Experiment 1

Subjects exhibit a systematic bias in memory for their investment outcomes. They remember significantly more positive and less negative outcomes than actually occurred. Further, the memory bias is related to overly optimistic subjective beliefs about the investment and suboptimal reinvesting decisions.

Table XXVI summarizes our findings from Experiment 2. The regression models include the full sample of subjects in Experiment 2. Column 1 displays our treatment effect on subjects' memory bias. We use subjects' memory bias for positive outcomes as dependent variable. The memory bias is estimated at the individual level by taking the difference between each subject's recalled number of positive outcomes in the memory elicitation task and the actual observed number of positive outcomes. Note, in Experiment 2 we forced subjects' responses in the memory elicitation task to always add up to 12 outcomes (i.e., the true number of outcomes observed). Thus, analyses using subjects' memory bias for negative outcomes yield the same finding as using positive outcomes. A treatment dummy variable serves as explanatory variable.

We find that after one week, subjects in the *Delay* group exhibit a significantly larger memory bias for positive outcomes than subjects in the *Immediate* group. That is, across all experimental conditions, subjects who were asked to recall their investment outcomes after one week, falsely remember on average 0.43 more positive outcomes than subjects who were asked immediately after the observation of investment outcomes.

Table XXVI. Subjective Memory Bias, Beliefs, and Investment Decisions from Experiment 2

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable is subjects' memory bias for observed positive outcomes *Memory Bias (for Pos. Outcomes)*, or the subjective posterior belief that the stock is the good stock (from 1 to 100), *Subjective Probability*. In addition, the table contains the odds ratios and standard errors (in parentheses) of Logit regressions restricted to the sample of subjects for whom the previous stock has a lower expected outcome compared to the alternative (new) stock in the second investment task and the dependent variable is a dummy variable which is equal to one if the subject reinvested in the stock. *Objective Probability* is the value of the objective Bayesian probability that the stock is the good stock (from 1 to 100). *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Memory Bias (for Pos. Out.)	(2) Subjective Probability	(3) Re-Investment (Subopt.)
Treatment	0.433***		
	(0.13)		
Memory Bias (for Pos. Outcomes)		2.880***	0.162*
		(0.51)	(0.10)
Objective Probability		0.340***	
		(0.02)	
Constant	-0.206**	40.030***	-0.927***
	(0.09)	(1.54)	(0.17)
N	498	498	204
R^2	0.02	0.29	
Pseudo R^2			0.01

Column 2 shows that subjects' memory bias is positively related to their subjective beliefs about the investment (for a comparison to Experiment 1, see Table III, column 1). We use subjects' reported probability that the stock is the good stock (between 0 and 100) as dependent variable and subjects' memory bias for positive outcomes as explanatory variable. We control for the correct Bayesian posterior that the stock is a good stock given the information seen by the subject. The regression result shows that subjects' beliefs are on average 2.9% higher for each positive outcome they over-remember.

Column 3 shows that subjects' memory bias is positively correlated with suboptimal decisions to reinvest (for a comparison to Experiment 1, see Table IV, column 1). Note, in Experiment 2 subjects are asked to choose between reinvesting in the stock they have observed and investing in a new randomly drawn stock for further 12 periods. Risk-neutral subjects should invest in the observed stock if it has a higher expected payoff than the new randomly assigned stock. The expected payoff of the observed stock can be evaluated using one's posterior belief, and the expected value of the new stock can be calculated from the given underlying distribution. The table reports the result of a Logit regression. We restrict the sample to subjects for whom the previous stock has a lower expected outcome compared to the alternative (new) stock in the second investment task. The dependent variable is a dummy variable equal to one for subjects who chose to invest in the previous (suboptimal) stock and zero when choosing the new stock. Subjects' memory bias for positive outcomes is used as the explanatory variable. We find that subjects who recollect a higher number of positive outcomes than actually occurred have a marginally significantly higher probability to reinvest in the stock suboptimally. In other words, subjects' probability to reinvest suboptimally increases by 16.2% with each positive outcome they over-remember. In total, 35% of subjects invest in the asset despite its lower expected outcome.

Table XXVII. Subjective Memory Bias in Experiment 3 (for pos. and neg. outcomes separately)

This table displays subjects' mean memory bias for positive and negative outcomes (compared to the true number of outcomes observed) in the *Delay* condition in Experiment 3. Memory bias is estimated at the individual level by subtracting the actual individually observed number of positive/negative outcomes from subject's recalled number of the stock's positive/negative outcomes. The table reports mean values and one-sided T-test results against the null hypothesis that the memory bias is zero.

Delay treatment (N = 233)	**Mean Memory Bias in *Delay* condition**	**One-sided T-test**
For Pos. Outcomes	0.26	p = 0.061
For Neg. Outcomes	-0.26	p = 0.061

Further, Table XXVII shows that we replicate subjects' memory bias in Experiment 3. The memory bias is reported separately for positive and negative outcomes. Note, in Experiment 3 we forced subjects' responses in the memory elicitation task to always add up to 12 outcomes (i.e., the true number of outcomes observed). Thus, subjects' memory bias for positive and negative outcomes yield the same finding, just with opposite signs.

N. The Role of Anticipated Regret

Table XXVIII reports the relation between subjects' memory bias and proneness to regret. We use subjects' memory bias for positive outcomes as dependent variable. An index variable based on subjects' responses to the five items measuring proneness to regret serves as independent variable. We find that subjects' memory bias for positive outcomes is not significantly correlated with subjects' proneness to regret. The index variable is not significantly correlated with the memory bias when exploring the total sample (column 1) nor when restricting the sample to subjects who participated in each of the three experimental conditions separately (columns 2-4).

Table XXVIII. Memory Bias and Subjects' Anticipated Regret

This table contains the coefficients and standard errors (in parentheses) of OLS regressions in which the dependent variable is subject's memory bias for observed positive outcomes *Memory Bias (for Pos. Outcomes)*. The three regression models in column 2, 3, and 4 restrict the sample to subjects who participated in each of the experimental conditions respectively. *Regret* is an index variable based on subjects' responses to five items measuring proneness to regret (Schwartz et al., 2002). *, **, and *** denote significance at the 10%, the 5%, and the 1% level, respectively.

	(1) Memory Bias - total sample	(2) Memory Bias - Baseline	(3) Memory Bias - HighStakes	(4) Memory Bias - NoChoice
Regret	-0.015	-0.067	-0.024	0.074
	(0.08)	(0.13)	(0.14)	(0.17)
Constant	0.067	0.359	0.038	-0.308
	(0.29)	(0.44)	(0.49)	(0.60)
N	490	178	154	158
R^2	0.00	0.00	0.00	0.00