



Network for Studies on Pensions, Aging and Retirement

Investor Memory

Katrin Godker, Peiran Jiao, and Paul Smeets

NETSPAR ACADEMIC SERIES

DP 11/2021-036

Investor Memory

Katrin Gödker, Peiran Jiao, and Paul Smeets*

This version: November, 2021

Abstract

How does memory shape individuals' financial decisions? We provide experimental evidence of a self-serving memory bias. Individuals over-remember positive investment outcomes of their chosen investments and under-remember negative ones. In contrast, individuals who did not choose their investments or did not invest but merely observed outcomes do not have this bias. The memory bias affects individuals' beliefs and decisions to re-invest. After investing, subjects form overly optimistic beliefs about their investment and re-invest even when doing so leads to a lower expected return. Our findings contribute to the understanding of how people learn from their experiences in financial markets. More generally, the documented memory bias offers a consistent explanation for stylized facts about investor behavior as well as dynamic risk taking in many economic domains.

JEL Classification: D01, G4

Keywords: Memory, Selective Recall, Beliefs, Motivated Reasoning, Investor Behavior, Experimental Economics

*Katrin Gödker: School of Business and Economics, Maastricht University, k.godker@maastrichtuniversity.nl; Peiran Jiao: School of Business and Economics, Maastricht University and Nuffield College, University of Oxford, p.jiao@maastrichtuniversity.nl; Paul Smeets: School of Business and Economics, Maastricht University, pm.smeets@maastrichtuniversity.nl.

We thank Roland Bénabou, Pedro Bordalo, Colin Camerer, Lauren Cohen, Sebastian Ebert, Robin Greenwood, Glenn Harrison, Alex Imas, Mats Köster, Terrance Odean, Sébastien Pouget, Philipp Strack, and Florian Zimmermann for very helpful comments and suggestions. We thank seminar participants at Caltech, DICE at Heinrich Heine University Duesseldorf, Frankfurt Goethe University, Frankfurt School of Finance and Management, Max Planck Institute for Research on Collective Goods, Toulouse School of Economics, University of Innsbruck, University of St. Gallen, and participants at AFA 2020 in San Diego, Boulder Summer Conference on Consumer Financial Decision Making, brq Institute Workshop on Beliefs 2021, Early Career Women in Finance Conference at Stanford 2020 (ECWFC), and Helsinki Finance Summit on Investor Behavior 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. The experiments and their procedures were ethically approved by the University Hamburg Ethical Research Committee, by the Office for Protection of Human Subjects (OPHS) at UC Berkeley (protocol number 2021-06-14375) and by the Ethics Committee of the Centre for Experimental Social Sciences (CESS) at Oxford University (protocol number VE_0009).

1 Introduction

How does memory shape individuals' financial decisions? We all learn from experience, yet our memories of those experiences can be both selective and distorted (Schacter, 1999). While economic theory has begun to integrate psychology-based facts about memory formation and retrieval into behavioral models, empirical work on this question is sparse. The primary contribution of this paper is to provide evidence for a systematic memory bias in an important economic domain: financial decision-making. In two experiments, we document a self-serving memory bias for investment outcomes and its consequences for individuals' beliefs and choices. 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. The memory bias affects individuals' beliefs about their investments and decisions to re-invest. These findings are relevant for the understanding of how people learn from experience in financial markets and offer an explanation for well-documented phenomena in investor behavior, such as the repurchase effect and the persistent trading of unsuccessful investors.

We conducted two experiments in an investment context to test for the presence of a memory bias, its underlying mechanism, and its effect on beliefs and subsequent decisions. In our first experiment (Experiment 1), subjects choose to invest either in a risky asset or a risk-free asset. The risky asset is either a "good stock" or a "bad stock". A good stock is more likely to generate positive outcomes and a bad stock is more likely to generate negative outcomes. Subjects do not know whether their stock is good or bad. To identify a memory bias, subjects observe a series of investment outcomes and we then elicit their memory of these 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 isolate the effect of memory from alternative factors that might influence recall, but are related to information acquisition or processing, such as attention (Barber and Odean, 2008; DellaVigna and Pollet, 2009) or salience (Bordalo, Gennaioli, and Shleifer, 2012, 2013). To relate subjects' memory bias to their beliefs and future investment behavior, we elicit their beliefs about the stock's chance of being the good stock and ask them to make an additional investment decision. Our experimental design has two key features. First, we directly measure what people remember. Second, to explore behavioral deviations from standard theory, we compare subjects' elicited beliefs and choices to a normative Bayesian benchmark.

We have two main findings from Experiment 1 ($N = 229$). First, subjects over-remember investment gains and under-remember investment losses if they invested in the stock. In contrast, subjects

who decided not to invest in the stock (i.e., chose the risk-free asset) do not display this memory bias. Second, the memory bias is related to subjects' belief formation and future investment choices. One week after observing the investment outcomes, subjects who invested in the stock form overly optimistic beliefs about the stock and are more likely to re-invest 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 the risk-free asset is optimal.

The results are consistent with a model in which self-image concerns form the basis for how information is remembered. We formalize a model in which memory of investment outcomes is systematically biased in a self-serving way and thereby distorts beliefs. We follow the idea that selective recall might be a potential mechanism for self-servingly biased beliefs (Bénabou and Tirole, 2002, 2004; Chew, Huang, and Zhao, 2020; Zimmermann, 2020).¹ 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 positive self-image. The agent observes outcomes of an asset to update her belief about the quality of the asset. She is relatively less likely to recall outcomes that are inconsistent with her positive self-image, and thus underweights them when updating her belief. That is, if the agent actively invested in the asset, she is less likely to recall negative outcomes relative to positive ones and becomes over-optimistic about the quality of the asset.

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

In a second experiment, we successfully replicate our results from the first experiment in two different laboratories ($N = 498$). In addition, we implement several treatments to provide evidence on the underlying mechanism driving the memory bias. We have two key findings that are in line with the notion of a motivational mechanism involving self-image concerns (Bénabou and Tirole, 2002, 2004; Chew, Huang, and Zhao, 2020; Zimmermann, 2020). First, our results suggest that individuals suppress memories of negative outcomes rather than actually forgetting them. When we increase the monetary incentive for accurate recall from \$8 to \$50, the memory bias is reduced and no longer significant. This is consistent with our proposition that the memory bias is motivated.

¹Experienced events are generally stored in episodic memory, the database of past experiences (Kahana, 2012; Tulving, 1972; Tulving and Murray, 1985). Evidence in psychology suggests that episodic memory, in particular, tends to be biased in ways that maintain and enhance one's own positive self-image: People are more likely to remember personal successes than failures and to remember features of past options that are supportive of the choices they made (Mather and Johnson, 2000).

Cognitive limitations of memory capacity would predict similar results for the low and high stakes treatment. Instead, we find that individuals recall more accurately if the monetary benefits of doing so become relatively larger compared to the benefits of a positive self-image. Second, individuals only exhibit the memory bias if they actively chose the investment. When the stock is randomly allocated to them, the memory bias is reduced and no longer significant. Motivated recall thereby seems more pronounced when recalling the outcomes of their own choices rather than those related to passive ownership, which is consistent with the idea that ego-relevance is derived from having made an active choice and inconsistent with a theory of simple optimistically biased recall.

The experiments are designed to shed new light on well-established behavioral facts concerning investor behavior. First, our results extend the literature on risk taking in financial markets. The way people take risks in markets is crucial to the financial well-being and wealth of many households. One dominant finding is that individuals' level of risk taking in financial markets is positively related to the returns they have experienced (Malmendier and Nagel, 2011; Knüpfer, Rantapuska, and Sarvimäki, 2017; Guiso, Sapienza, and Zingales, 2018; Andersen, Hanspal, and Nielsen, 2019). The key implication of this experience hypothesis is that experiences of high returns tend to be correlated with higher subsequent financial risk taking and experiences of low returns tend to be correlated with lower financial risk taking. Yet, we find that positive and negative investment outcomes, such as returns, are remembered differently depending on whether people invested or not. For instance, people who invested can show high levels of risk taking despite experiencing low returns, because they form overly optimistic beliefs based on their biased memory. By contrast, those who did not decide to invest show unbiased memory and take less risk. Thus, we provide a cognitive foundation for heterogeneity of experience effects at the individual level.

Second, we observe subjects' decisions about whether or not to re-invest in a previously chosen stock. This directly connects our findings to the finance literature on stock repurchases. It has been documented that private as well as professional investors have a higher propensity to repurchase a stock if they had previously sold that stock for a gain rather than a loss (Strahilevitz, Odean, and Barber, 2011; Du, Niessen-Ruenzi, and Odean, 2019). Further, Frydman and Camerer (2016) use neural data collected from an experimental asset market to show that induced regret from a price increase after a stock sale is correlated with stock repurchase mistakes. We contribute to this strand of research by showing that positively biased memory of own past realizations from a stock is a source of stock repurchases. We further document that when time passes, individuals with a positive memory bias re-invest even if this reduces their expected payoff.

In addition, the memory mechanism we document could underlie investor overconfidence, which

is used to explain a number of stylized facts in finance, most prominently the high trading volume in stock markets. In particular, our results suggest that memory could be the source of a learning process that leads investors to become overconfident (Gervais and Odean, 2001). If investors under-remember their investment failures relative to successes, as indicated by our findings, and use their memories as a statistic to form beliefs about their ability to trade, they would become overconfident about their trading ability. This has important implications for investor learning over time. In line with recent evidence that selective memory generates persistent overconfidence of managers in settings with repeated feedback (Huffman, Raymond, and Shvets, 2019), the memory bias that we document could explain the persistent overconfidence of investors and also the high trading volumes displayed by unsuccessful investors documented, e.g., in (Barber, Lee, Liu, Odean, and Zhang, 2020). If traders forget their own losses, they cannot learn from their failures.

Our paper further contributes to the 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). In contrast to the well-documented observation that "losses loom larger than gains" in choices under risk (Kahneman and Tversky, 1979), individuals seem to fail to update fully in response to negative news when motivation is at play (Bénabou and Tirole, 2016). Recent experimental work provides evidence for overly optimistic belief formation about ego-relevant information on intelligence (Zimmermann, 2020), beauty (Eil and Rao, 2011), and generosity (Saucet and Villeval, 2019; Di Tella, Perez-Truglia, Babino, and Sigman, 2015; Carlson, Maréchal, Oud, Fehr, and Crockett, 2020).

We extend this to the domain of financial decision-making, where ego-relevance is derived from the subject's belief about their ability to make good investment choices. This allows us to explore self-serving belief formation and updating about the quality of one's choices rather than personal characteristics like intelligence, beauty, or generosity. We further document the relevance of active choices for a motivated 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 ownership of a good impacts how people react to information when updating their beliefs about the good (Hartzmark, Hirshman, and Imas, 2021). Our results show that beyond ownership, active choices result in different beliefs and subsequent behavior than passive ownership.

A key question in the literature on motivated reasoning remains how 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 risky investments because they have overly positive recollections of their investment returns.

More generally, economics and finance research has shown that prior outcomes affect subsequent risk taking (Thaler and Johnson, 1990; Imas, 2016; Suhonen and Saastamoinen, 2017; Kuhnen, Rudolf, and Weber, 2017). Our findings suggest that prior gains from own choices have a relatively larger impact on future risk taking compared to prior losses, because people forget their losses. Importantly, this effect increases significantly when time has passed. That is, the distinction between how individuals respond to prior gains versus losses depends on whether people make decisions immediately after experiencing the outcomes or after some delay. Such effects on individual risk taking are relevant in many economic domains, in particular for dynamic settings in which sequential risky choices are made. For example, the proposed self-serving memory bias might have implications for corporate investment decisions by CEOs, betting, or competitive behavior.

2 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, 2019). In this section, we present a simple behavioral mechanism for a self-serving 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):² the agent relatively under-remembers signals that are inconsistent with her positive self-image, and becomes over-optimistic about more preferred states.

2.1 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

²See Bénabou and Tirole (2016) for a review of this literature.

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.³ 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 $h_t = (d_1, \dots, d_t)$. The agent observes the history of outcomes h_t , stores the outcomes in memory as the history of remembered outcomes $h_t^R = (d_1^R, \dots, d_t^R)$, 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^- . The total number of positive and negative outcomes sum up to t .⁴ 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 $n_t^{+,R} + n_t^{-,R} = t$.⁵ 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.⁶ 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 positive self-image and positive outcomes might align with the agent's positive self-image. 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 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, Chen, McDonel, and Sherman, 1982).

Suppose R_t represents the memory bias in period t , given the agent's initial choice to invest in

³This is without loss of generality, as long as the set of outcomes is ordered.

⁴Under the current notation, we also conveniently have $n_t^+ = \sum_0^t d_t$, and $n_t^- = t - \sum_0^t d_t$.

⁵This 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.

⁶Relative 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.

the stock and dependent on whether the outcome is positive or negative:

$$R_t = \delta_t [I_t q^- + (1 - I_t) q^+], \quad (1)$$

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 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 $\delta_t = 1$.⁷ 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 positive self-image. 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:

⁷In our experimental results we will demonstrate the difference in memory bias between an active choice and a passive endowment.

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

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

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

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

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

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 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 positive self-image.

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 (3) and Equation (2) coincide.

The agent with a self-serving 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.

2.2 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, $\mu_0^G = \mu_0^B = \mu_0 = 0.5$.⁸ This further simplifies matters and helps us to derive simple representations of biased signal weightings, which are testable in our experiment.

Rewriting Equation (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 (4), the weight should be exactly equal to $\ln(\frac{\theta}{1-\theta})$, which represents the informativeness of positive relative to negative signals.

$$\ln \Lambda^{Bayesian}(h_t) = \ln\left(\frac{\theta}{1-\theta}\right)n_t^+ - \ln\left(\frac{\theta}{1-\theta}\right)n_t^- \quad (4)$$

However, the agent with a self-serving memory bias has the following log likelihood ratios

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

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

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

⁸In 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.

$$\ln \Lambda^{Biased,INV}(h_t^R) - \ln \Lambda^{Bayesian}(h_t) = \ln\left(\frac{\theta}{1-\theta}\right)2q^-n_t^-, \quad (5)$$

$$\ln \Lambda^{Biased,NOT}(h_t^R) - \ln \Lambda^{Bayesian}(h_t) = -\ln\left(\frac{\theta}{1-\theta}\right)2q^+n_t^+. \quad (6)$$

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 (5) and (6) 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 one's self-image (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^+).*

The model serves two important purposes for our subsequent experimental testing. In the experiment, we fix θ but vary n_t^- and n_t^+ . First, 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 true outcomes that are contradictory to the biased agent's self-image (Proposition 2). These are the negative outcomes if the agent invests and the positive outcomes when the agent does not invest. We will provide experimental evidence for this proposition.⁹ 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). We will also provide evidence for this proposition.

Second, Equations (5) and (6) provide the foundation for our regression analysis with regard to memory-based belief distortion. With subjects' belief distortion as the dependent variable, $\ln\left(\frac{\theta}{1-\theta}\right)2n_t^-$ or $-\ln\left(\frac{\theta}{1-\theta}\right)2n_t^+$ as the independent variable, we can directly estimate the value of q^- and q^+ , i.e., the memory bias, on average across subjects.

⁹Note that in our experiment, memory elicitation was a surprise task so we could not vary parameters within-subjects.

3 Evidence for the Investor Memory Bias

3.1 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 self-relevant outcomes, (ii) exogenous variation in investment outcomes (positive and negative outcomes), (iii) direct memory elicitation, and (iv) an experimental manipulation of the time span between observation of outcomes and tasks to isolate memory effects. In this section, we outline the features of our experiment (Experiment 1) in detail (Table 1 summarizes our treatment conditions).¹⁰

Table 1: Experimental Conditions

Treatment	First investment choice and observation of outcomes	Memory elicitation	Belief elicitation and second investment choice
<i>Delay</i>	Week t	Week t+1	Week t+1
<i>Immediate1</i>	Week t	Week t	Week t
<i>Immediate2</i>	Week t+1	Week t+1	Week t+1
<i>NoRecall</i>	Week t	No	Week t+1

Notes: This table provides an overview of the treatment and control conditions of the experiment with different time spans between tasks.

First, subjects make an investment decision. They 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). After that decision, investment outcomes are observed over the course of 12 periods. With equal probability, the stock may be good or bad. That is, the stock is either more likely to generate positive outcomes or more likely to generate negative outcomes. We choose $\theta = 0.6$. 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 and negative outcomes with a 60% probability.¹¹ The positive outcomes of the stock are either 11, 13, or 15 EUR and the negative outcomes are either -5, -3, or -1 EUR. More precisely, 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. Subjects start with an initial endowment of 60 EUR.¹² See Table 2 for an overview of subjects'

¹⁰The experiment instructions are provided in Appendix A.

¹¹To make sure that subjects understand the distributions, we included a phase of experience sampling (for both the distribution of the good stock's and the bad stock's outcomes) which did not influence subjects payout.

¹²The incentive structure is described in detail in Section 3.2.

investment options.

To measure subjects' memory bias, we let subjects observe the generated 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. 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 how many positive and negative outcomes they observed and, more specifically, how often the stock paid 11, 13, 15, -5, -3, and -1 EUR. This memory task is not announced beforehand.¹³

Importantly, to clearly identify the effect memory has on subjects' recollection, we manipulate the time span between the observation phase and the memory elicitation, following a between-subject design. We randomly assign subjects to one of three experimental conditions for 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 (Table 1). 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) or salience (Bordalo, Gennaioli, and Shleifer, 2012, 2013), which might influence subjects' recollection of outcomes, but are related to subjects' information acquisition and processing. Differences between the *Delay* and the *Immediate* conditions cannot be caused by subjects' information acquisition or processing, but solely by the fact that subjects have to remember the stock's outcomes from last week.

A *NoRecall* condition, in which subjects do not perform a memory elicitation task, serves as a control condition. This treatment allows us to identify potential effects of simply asking subjects to recall investment outcomes. Note that 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 simply asking subjects to recall the stock's outcomes does not drive our results. Please refer to Section 3.4.5 for a detailed description of the analyses.

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. This could happen either in week t or in week $t + 1$. 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 $t + 1$ (Table 1). In our main analyses, we pool the data of the *Immediate1* condition and the *Immediate2* condition and control for the

¹³The instructions for the memory elicitation are provided in Appendix B.

Table 2: 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%	
	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%	
	Bond	No	-	5.10 EUR	100%	5.10 EUR

Notes: This table provides an overview of subjects' investment options during the experiment.

session the subject participated in. In Section 3.4.4 we provide robustness checks for the timing.

Subjects had to sign up for and participate in two experimental sessions, with one week in between irrespective of whether they were randomly assigned to the *Delay* or *Immediate* treatments, to avoid selection effects. Further, in order to reduce attrition, (i) we made all payments from the experiment in the second session (in week $t + 1$), to maximize the incentive for subjects to show up to the second lab session, (ii) for each subject the first and the second lab session took place at the same day of the week and at the same time (only one week later), and (iii) subjects were reminded via email about the second lab session. Out of 239 subjects that participated in the first session, 229 participated in the second session, which is 96%.

To relate subjects' memory bias to their beliefs, we elicit subjects' beliefs about the stock's chance of paying a positive outcome after they observed the investment outcomes. 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 make informed inferences about the stock's probability of paying from the good distribution of outcomes. A fully rational (Bayesian) subject counts the number of positive outcomes (11, 13, or 15 EUR) in the course of the 12 periods. The value of the objective Bayesian posterior can be calculated as:

$$\mu_t^G(h_t) = \frac{1}{1 + \frac{1-\mu_0^G}{\mu_0^G} * \left(\frac{\theta}{1-\theta}\right)^{t-2n_t^+}} \quad (7)$$

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 observations; h_t represents the history of outcomes for t observations; n_t^+ represents the number of positive outcomes. This posterior serves as benchmark for objectively correct beliefs in our experimental setting.

To further investigate the consequence of a memory bias for subsequent investment behavior, we ask subjects to make a second investment decision after observing the outcomes. They choose to invest in the stock they have observed or another bond for further 12 periods (Table 2). Risk-neutral subjects should invest in the stock if it has a higher expected outcome 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 in our experimental setup, 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.¹⁴ Risk-averse subjects should require a higher posterior belief about the stock being a good stock 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 3.3.3), our results should hold for a range of reasonable risk attitude parameters.¹⁵

In the *Delay* condition, we elicit subjects' beliefs and second investment choice one week after the observation phase and in the *Immediate* condition immediately after the observation phase. Further, to avoid spill-over effects, the order in which subjects perform these tasks, i.e., the memory elicitation, belief elicitation, and the second investment task, is random.

3.2 Incentives and Procedures

The experimental sessions were organized in two parts. Subjects first made their decision to invest in the stock or bond, and observed the stock's outcomes, and afterwards participated in a memory elicitation, belief elicitation, and a second investment task.

Subjects were paid a show-up fee of 8 EUR for participating in the study.¹⁶ Further, we randomly drew three participants from each session (with maximum 30 participants per session) who were paid based on their performance in one of the tasks. For each drawn subject, the computer

¹⁴In 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.

¹⁵Typically, 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).

¹⁶In the first two sessions, subjects were paid 5 EUR for participating in the study. We increased this amount to 8 EUR after the first two sessions because the average duration of the experimental sessions was longer than expected.

randomly decided which task determined his or her payment.¹⁷ 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 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, as well as subjects' understanding of the risk-return relationship of investments. Further, subjects were asked to indicate their age, gender, and highest level of education. In addition, subjects were asked to solve three Raven matrices and we elicited their mood (Watson, Clark, and Tellegen, 1988).

A total of 229 subjects participated in the laboratory experiment, mostly business and economics students from one of the authors' universities.¹⁸ On average, subjects earned 16.90 EUR. For each subject, both sessions took about 45 minutes each. The experiment is programmed and conducted with z-Tree (Fischbacher, 2007) and the experimental sessions were organized and administrated with the software hroot (Bock, Baetge, and Nicklisch, 2014).

3.3 Results

The results of Experiment 1 document that subjects exhibit a positive memory bias for investment outcomes, which distorts their beliefs and drives subsequent investment decisions. The findings are consistent with our model in which image-concerns form the basis for how information is remembered.

3.3.1 Memory Bias

We find that (i) subjects' memory of investment outcomes is systematically biased; (ii) subjects' memory bias differs for positive and negative outcomes; and (iii) subjects' memory bias depends on

¹⁷It has been shown that paying a subset of participants is an effective payment scheme for economic experiments (Charness, Gneezy, and Halladay, 2016) and that random incentive systems do not bias risk-taking behavior in experiments (Starmer and Sugden, 1991; Cubitt, Starmer, and Sugden, 1998; Hey and Lee, 2005).

¹⁸The experiment and its procedure were ethically approved by the University Experimental Laboratory Committee. We obtained subjects' consent orally before participating in the experiment.

whether they invested or not. Together, these results provide evidence for a positive memory bias for own investment gains and losses.

Result 1. After investing, subjects recall significantly more positive and less negative outcomes than actually occurred. Subjects who did not invest recall the correct number of occurred positive and negative outcomes.

Based on our direct memory elicitation, 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. 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. Importantly, to isolate memory effects from other factors related to subjects' information acquisition or processing, we compare subjects' recollection in our two treatment conditions, *Immediate* and *Delay*.

Table 3 provides an overview of subjects' memory bias and reports results from *T*-tests against the null hypothesis that the memory bias is zero (column 2, column 4, column 6, and column 8). We first focus on the *Delay* condition, in which people had time to form inaccurate 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, as implied by a self-serving memory bias: If they invested in the stock, they remember significantly more positive outcomes and significantly less negative outcomes than actually observed. On average, subjects over-remember the absolute number of positive outcomes by 0.9 (*T*-test, $p < 0.001$) and under-remember the absolute number of negative outcomes by 0.7 (*T*-test, $p < 0.001$). We calculate the percentage of over-remembered and under-remembered outcomes compared to actually occurred outcomes at the subject level.¹⁹ On average, after one week, subjects remember 23% more positive outcomes and 10% less negative outcomes than actually occurred.

Moreover, the first columns (columns 1-4) illustrate that while subjects who invested in the stock show a significant memory bias for investment outcomes, subjects who did not invest do not display a memory bias. This is in line with studies suggesting that the decision to do something is of higher importance for one's self-image compared to the decision not to do something (Cioffi and Garner,

¹⁹For 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.

Table 3: Subjective Memory Bias

Observed outcomes	<i>Delay</i>				<i>Immediate</i>				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

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 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).

1996; Allison and Messick, 1988; Fazio, Chen, McDonel, and Sherman, 1982).

Further, 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. We indeed find that the memory bias is significantly larger in the *Delay* treatment than in the *Immediate* treatment, both for positive and for negative outcomes (*T*-test, $p < 0.05$).

Our findings are robust to different estimations of memory bias. Please refer to Section 3.4.1 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 number of recalled positive and negative outcomes.

3.3.2 Memory-Based Beliefs

In this section, we show that the memory bias is associated with overly optimistic subjective beliefs. First, our results document that subjects' elicited memory bias is significantly correlated with too optimistic beliefs about the stock. Second, in accordance with our model, we test for subjects' memory bias based on how they weight observed outcomes when forming beliefs about the stock. In line with the model propositions, we find that subjects who invested in the stock relatively underweight negative compared to positive outcomes, resulting in overly optimistic beliefs. This supports our finding of a memory bias without relying on a specific memory elicitation method. Moreover,

we show that the belief distortion of subjects who invested in the stock is larger if more negative outcomes occurred, which is consistent with our model.

Result 2. Subjects with a larger memory bias form significantly more optimistic beliefs about their investment.

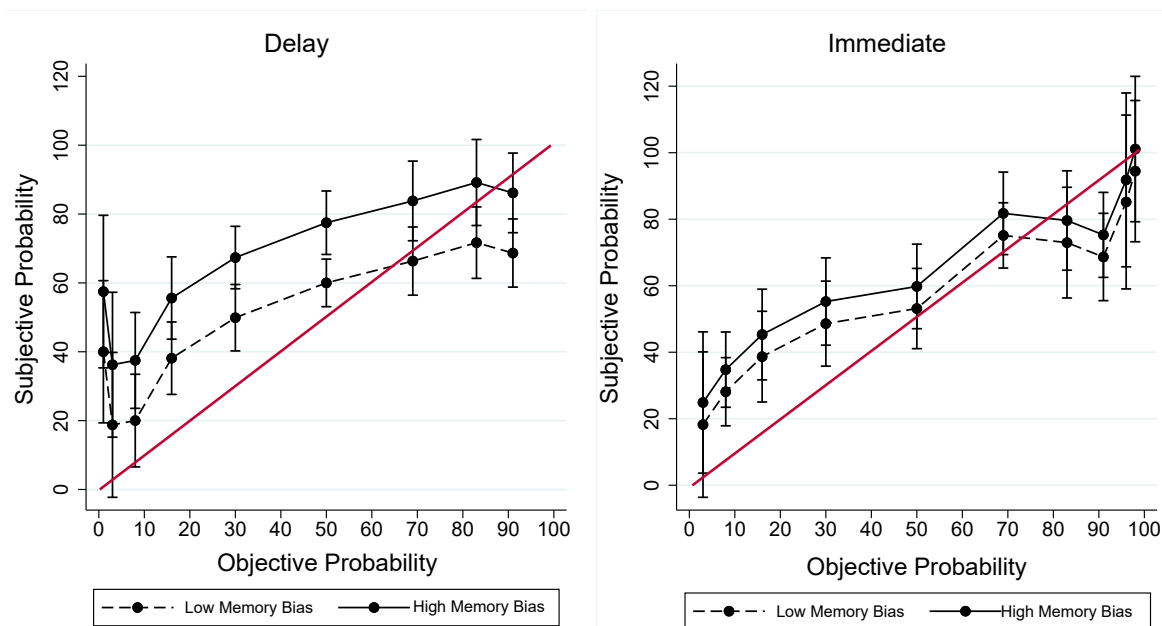
First, Figure 1 gives a graphical presentation of subjects' beliefs in our experiment if they chose to invest. The figure displays subjects' beliefs relative to the objective Bayesian probabilities. The x -axis indicates the value of each possible objective Bayesian posterior belief, i.e., the objective 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 (left panel) subjective beliefs are overly optimistic regarding the likelihood that the stock is good. Moreover, 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 predicts subjective beliefs. The regression models in Table 4 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 for positive and negative outcomes serve as independent variables.²⁰ 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 ($p < 0.001$). That is, subjects who recollect a 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 over-remember ($p < 0.001$). This also means that subjects who remember a lower number

²⁰When 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 regression analyses throughout the paper, we include both memory bias variables to account for subjects' deviations from the total number of 12 outcomes.

Figure 1: 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). 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.



of losses than observed form also more optimistic beliefs about the stock.²¹ 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.

Second, we show that subjects' weighting of observed outcomes when forming beliefs is consistent with our model and supports our finding of a systematic memory bias without relying on a specific memory elicitation method. To begin with, we find that subjects' belief distortion is positively correlated with the number of observed signals that might be inconsistent with their self-image, as indicated in Proposition 2. Figure 2 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

²¹Note 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 4: 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*. *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)	(2)	(3)	(4)
	Subjective Probability	Subjective Probability	Belief Distortion	Belief Distortion
Memory Bias (for Pos. Outcomes)	6.196*** (0.92)		0.424*** (0.05)	
Memory Bias (for Neg. Outcomes)		-7.958*** (1.02)		-0.525*** (0.06)
Objective Probability	0.684*** (0.05)	0.701*** (0.05)		
Constant	16.751*** (5.16)	15.204*** (5.00)	-0.185 (0.27)	-0.223 (0.26)
Session	Yes	Yes	Yes	Yes
N	188	188	182	182
R ²	0.52	0.55	0.31	0.35

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.

Moreover, we are interested in subjects' memory bias parameters from our model, q^- and q^+ . In the model, q^- is the probability of under-remembering a negative outcome of 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%. Regressions in Table 5 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 (5) and (6) and control for session fixed effects. Thus, the regression coefficients represent our estimations for q^- and q^+ from our model based on subjects' relative underweighting of negative

Figure 2: 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^- .



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) are likely to underweight negative relative to positive outcomes when forming beliefs ($p > 0.001$), while subjects who did not invest in the stock (columns 2 and 4), do not display a significant probability to underweight outcomes in a systematic way. 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 3). Here, q^- can be interpreted as the magnitude of subjects' memory bias when updating beliefs. We find that subjects who invested in the stock show a higher belief distortion when their memory bias (q^-) is higher. After observing a negative outcome in the *Delay* condition, subjects who invested in the stock forget this outcome with a probability of 59.9% ($p < 0.001$). Moreover, this probability is larger for subjects in the *Delay* condition than for subjects in the *Immediate* condition (47.3%, $p < 0.001$). Thus, we find that subjects' belief distortion stems from relatively underweighting negative outcomes when forming beliefs about the stock. This finding does not rely on a specific memory elic-

Table 5: 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 (5) and (6). *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. (<i>Del.</i> , inv.)	(2) Belief Dist. (<i>Del.</i> , not inv.)	(3) Belief Dist. (<i>Imm.</i> , inv.)	(4) Belief Dist. (<i>Imm.</i> , not inv.)
$\ln(\frac{\theta}{1-\theta})2n_t^-$	0.599*** (0.07)		0.473*** (0.07)	
$-\ln(\frac{\theta}{1-\theta})2n_t^+$		0.341 (0.27)		0.147 (0.36)
Constant	-3.189*** (0.54)	0.686 (1.81)	-1.880*** (0.50)	1.830 (2.02)
Session	Yes	Yes	Yes	Yes
N	73	18	73	19
R^2	0.55	0.58	0.46	0.40

itation method. This finding and the observation that subjects' elicited memory bias is significantly correlated with their belief distortion (Table 4) is supportive of our Proposition 3.

3.3.3 Consequences for Investment Decisions

The evidence so far shows that subjects relatively under-remember negative outcomes compared to positive outcomes if they invested in the stock. This memory bias is associated with overly optimistic beliefs about the investment. In this section, we show that subjects' memory is related to future investment choices. We find that subjects in the *Delay* condition invest significantly more in the observed stock compared to subjects in the *Immediate* condition, even in cases when this is objectively a mistake. This investment mistake is strongly related to subjects' elicited memory bias.

Result 3. Subjects with a larger memory bias are significantly more likely to re-invest suboptimally.

Table 6 shows the proportion of subjects who chose to invest in the stock after observing the investment outcomes and *T*-tests for differences in means across our treatments. In the *Immediate* condition, subjects were asked to decide immediately after the observation and in the *Delay* condition with a one week delay. Comparing the results in columns 1 and 2 illustrates that subjects decide

differently in the two conditions, although they show the same preferences for investing in the risky stock before the treatment (first choice in the experiment).

Table 6: Memory-Based Investment Decisions

	First choice		Second choice			
	Investment in stock	N	Investment in stock	N	Investment in stock (suboptimal)	N
	(1)		(2)		(3)	
<i>Immediate</i>	80.41	97	40.21	97	29.82	75
<i>Delay</i>	80.43	92	56.52	92	54.39	69
Difference (T-test)	p = 0.997		p = 0.025		p = 0.008	

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

An important question is whether this effect is associated with the quality of subjects' choices. Therefore, column 3 reports results for subjects who made 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 more than half of the subjects in the *Delay* condition (54.39%) invest in the stock despite its lower expected outcome, while in the *Immediate* condition only 29.82% of the subjects invest in such a manner. The difference between the two treatments is significant at the 1% level. Thus, subjects seem to be more likely to avoid suboptimal decisions when immediately deciding whether to invest or not (*Immediate* condition), whereas subjects who have to rely on their memory from last week (*Delay* condition) are likely to invest, 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 7: Subjective Memory Bias and Suboptimal Investment 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). *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) Investment in Stock (Suboptimal)	(2) Investment in Stock (Suboptimal)
Memory Bias (for Pos. Outcomes)	1.563*** (0.19)	
Memory Bias (for Neg. Outcomes)		0.661*** (0.09)
Constant	0.050*** (0.052)	0.0452*** (0.047)
Session	Yes	Yes
N	188	188
Pseudo R^2	0.12	0.10

Furthermore, regression analyses in Table 7 show that subjects' suboptimal investment decisions are correlated with their elicited individual memory bias. The table reports results from Logit regressions with a dummy variable equal to one for subjects who chose to invest in the stock with a lower expected outcome than the bond after observing the outcomes as dependent variable (second choice). We use subjects' memory bias for positive and negative outcomes as independent variables and control for session fixed effects. Column 1 indicates that subjects' probability to invest in the stock with a lower expected outcome than the bond increases by 56.3% with each positive outcome they over-remember ($p < 0.001$). In other words, subjects who recollect a higher number of positive outcomes than actually observed have a significantly higher probability to invest in the stock with a lower expected outcome. We find the opposite direction for negative investment outcomes (column 2). Subjects' probability to invest in the stock with the lower expected outcome is on average 33.9% lower for every loss they over-remember ($p < 0.005$). Note that the average subject remembers a lower number of negative outcomes than actually occurred (Table 3). In this regard, each under-remembered loss increased the probability to invest in the stock with a lower expected outcome by 33.9%. Together, these findings document that subjects' memory bias is related to suboptimal investment choices.

3.4 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 timing of the experimental tasks in the *Immediate* condition. 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. Our detailed robustness checks are provided in Appendix C.

3.4.1 Different Measures of Memory Bias

The main findings in Table 3 of Section 3.3.1 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 14) and a memory bias of the absolute difference between recalled positive and negative outcomes (Table 15). Our results are qualitatively unchanged. See Appendix C.1 for more details.

3.4.2 Subjects' Risk Preferences

Our findings are robust to differences in subjects' self-reported risk preferences (Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner, 2011). First, subjects' memory bias is not correlated with self-reported risk preferences (Table 16 of Appendix C.2). Second, our main findings are robust to controlling for individual risk preferences (Table 17 of Appendix C.2). See Appendix C.2 for more details.

3.4.3 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 18 of Appendix C.3 reports the robustness of our findings in Table 3 of Section 3.3.1, 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 19 of Appendix C.3). Third, Table 17 shows the robustness of our main findings of Section 3.3.2 and 3.3.3 to differences in subjects' financial literacy. The results are qualitatively unchanged.

3.4.4 Timing of Experimental Tasks

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 Appendix C.4).

3.4.5 No 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 21 in Appendix C.5 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 subsample 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.

4 Exploring the Mechanism

The findings of Experiment 1 are consistent with the essential notion of our model, namely that self-image concerns form the basis for how outcomes are remembered. In a second experiment (Experiment 2), we provide additional support for a self-serving memory bias by exploring the underlying mechanism more directly.

4.1 Experimental Design

We maintain the structure and key features of Experiment 1, but implement additional treatment arms to pinpoint the underlying mechanism of the investor memory bias. Experiment 2 also tests the robustness of our findings by replicating the results of Experiment 1 in different laboratories. In this section, we outline how the setting of Experiment 2 deviates from the original design (Experiment 1).

We implement a *Baseline* condition, which replicates our Experiment 1 with a few exceptions. Like in Experiment 1, we randomly assign subjects to an *Immediate* and *Delay* treatment group. However, we implement a few changes.

First, at the beginning of the experiment we let subjects choose between two risky stocks rather than between a risky stock and a riskless bond. A choice between two stocks rules out any potential bias due to subjects' endogenous selection into investing in the stock. Like in Experiment 1, each stock is a good stock or a bad stock with equal probability, which is independently determined. 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 and negative outcomes with a 60% probability each period.

Second, before subjects make the investment decision, we show them the outcomes of the two risky stocks from three previous periods. These three outcomes give subjects noisy information about the two stocks from which they can choose from. 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.

Third, the stocks' positive and negative outcomes are symmetric in size. The positive outcomes

are either 2, 4 or 6 GBP/USD and the negative outcomes are either -6 , -4 , or -2 GBP/USD.²² That is, given that the outcome in a period is positive, it is randomly drawn from $\{2, 4, 6\}$ with equal probability. Given that the outcome in a period is negative, it is randomly drawn from $\{-6, -4, -2\}$ with equal probability. The determination of the outcomes is independent across periods. We use 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.

Fourth, like in Experiment 1 subjects participate in three tasks in random order, after observing the stocks' outcomes: a surprise memory elicitation task, a belief elicitation task, and a second investment task. In the second investment task of Experiment 2, however, subjects are asked to choose between investing 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.

Fifth, we want to rule out anticipatory utility affecting subjects' memory of outcomes from the first investment task. From a theory perspective, subjects who care about expected future utility flows are happier now if they overestimate the probability that their investments pay off well in the future (Brunnermeier and Parker, 2005). 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 $t + 1$.

Importantly, we implement two additional conditions which aim at exploring underlying motivational mechanisms of the investor memory bias, namely the role of motivated memory suppression and the role of active choice. In each condition we randomly assign subjects to an *Immediate* or *Delay* treatment group. Table 8 summarizes our treatment conditions as well as the number of subjects per treatment condition.

4.1.1 Memory Suppression (High Stake)

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

²²We conducted the sessions in two labs, respectively in the UK and the USA, which will be explained in detail in the next subsection.

Table 8: Experimental Conditions of Experiment 2

Conditions	Treatment Groups	N
<i>Baseline</i>	<i>Immediate</i>	91
	<i>Delay</i>	94
<i>HighStake</i>	<i>Immediate</i>	77
	<i>Delay</i>	77
<i>NoChoice</i>	<i>Immediate</i>	75
	<i>Delay</i>	84

Notes: This table provides an overview of the experimental conditions and the number of subjects participated.

which case the true information could never be retrieved. Thus, in situations where people have set their mind on suppressing 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 in our documented investor memory bias, we conduct a high stakes version of the *Baseline* condition. In the *HighStake* condition, we manipulate the monetary benefit of reporting accurate memory. 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. With a comparison between these two conditions, we can test whether the true outcomes are correctly coded in subjects' memory but are simply suppressed in the retrieval process. If so, high incentives would induce subjects to report memory more accurately.

In addition, we aim at exploring 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). 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.

4.1.2 Active Choice vs. No 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, Chen,

McDonel, and Sherman, 1982). Outcomes from active choices not only mean changes in wealth, but also represent successes or failures of the investor. 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.

Second, we compare subjects' recall in the *Immediate* versus *Delay* treatment groups in the *Baseline* and *NoChoice* condition. In the *Baseline* condition we expect a positive memory bias in line with individuals enhancing their own positive self-image. In contrast, in a situation with less or no self-image concerns, 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). Thus, with limited motivational reasons in the *NoChoice* condition, we expect an opposite pattern: no bias in the *Delay* treatment group of the *NoChoice* condition and a negativity bias in subjects' recall in the *Immediate* treatment group.

4.2 Incentives and Procedures

Experiment 2 was conducted online as a virtual lab experiment 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, between May and July, 2021.²³ The number of subjects who participated in the experiment by lab can be found in Appendix D.

Like in Experiment 1, subjects were paid a show-up fee for participating in the study. The show-up fee for each session at CESS was 3 GBP. The show-up fee for the first session was 3 USD, and for the second session was 7 USD at Xlab in order to comply with the lab requirements. All subjects received the payoff from the first investment decision, which was 8 GBP/USD plus 10% of

²³The experiment and its procedure were ethically approved by the Office for Protection of Human Subjects (OPHS) at UC Berkeley (protocol number 2021-06-14375) and by the Ethics Committee of the Centre for Experimental Social Sciences (CESS) at Oxford University (protocol number VE_0009). We obtained subjects' informed consent before they participated in the experiment.

the dividends accumulated over 12 periods. Additionally, subjects were paid based on their payoff in one randomly chosen task: the memory elicitation task paid 8 GBP/USD if their answer was correct (50 GBP/USD in the *HighStake* condition); the belief elicitation task paid 8 GBP/USD if their answer was within 5% of the objective Bayesian value; the second investment task paid 8 GBP/USD plus 10% of the dividends accumulated over 12 periods.

The experiment was followed by a questionnaire. 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, Falk, Huffman, Sunde, Schupp, and Wagner, 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 completed both sessions of the experiment.²⁴ On average, subjects earned 15.50 GBP at CESS / 23 USD at Xlab. For each subject, the two sessions took about 30 minutes each. The experiment is programmed and conducted with oTree (Chen, Schonger, and Wickens, 2016).

4.3 Results

The findings of Experiment 2 provide direct evidence for the motivational mechanism underlying the memory bias.

4.3.1 Replication of the Findings from Experiment 1

Before we explore the mechanism of the memory bias, we replicate our results 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. Further, the memory bias is related to overly optimistic subjective beliefs about the investment and suboptimal re-investing decisions.

Table 9 summarizes our findings. 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

²⁴Note that few subjects took part in the first session, but not the second session. Attrition was at 8.5%. 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. 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.

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 9: 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 re-invested 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)	(2)	(3)
	Memory Bias (for Pos. Out.)	Subjective Probability	Re-Investment (Subopt.)
Treatment	0.433*** (0.13)		
Memory Bias (for Pos. Outcomes)		2.880*** (0.51)	0.162* (0.10)
Objective Probability		0.340*** (0.02)	
Constant	-0.206** (0.09)	40.030*** (1.54)	-0.927*** (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. 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 re-invest. 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 re-invest in the stock suboptimally. In other words, subjects' probability to re-invest 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.

4.3.2 The Role of Memory Suppression

Compared to the *Baseline* condition, we increased the financial incentive for a correct answer in the memory elicitation task from 8 GBP/USD to 50 GBP/USD in the *HighStake* condition. A comparison between these two conditions can isolate subjects' memory suppression as an underlying mechanism of the memory bias. A reduced memory bias in the *HighStake* condition would support the idea that motivated memory suppression is at play.

First, Table 10 illustrates subjects' mean memory bias for positive outcomes in the *Delay* treatment group and the treatment effect (difference in mean memory bias between *Delay* and *Immediate* group) across the three experimental conditions. 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.

In support of motivated memory suppression, we find that subjects show a significant positivity 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). Table 11 provides evidence for such a channel. We use subjects' memory bias for positive outcomes as dependent variable. 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.

We split the sample based on subjects' responses to a validated survey scale, which measures

Table 10: Subjective Memory Bias Across Experimental Conditions

Condition	Mean bias in <i>Delay</i> (vs. true number of outcomes)	Treatment effect (difference between <i>Delay</i> and <i>Immediate</i>)	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 vs. passive choice

Notes: This table displays subjects' memory bias in the *Delay* and *Immediate* group, separately for subjects in the *Baseline*, *HighStakes*, 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*.

the tendency to deceive oneself for 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. 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.

Table 11: 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)*. 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). *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
HighStakes	-0.469** (0.19)	0.115 (0.22)
Constant	0.182 (0.12)	0.058 (0.16)
N	172	167
R ²	0.04	0.00

4.3.3 The Role of Active Choice

Subjects in the *NoChoice* condition do not make an investment choice in the first task, but receive a randomly assigned risky stock from the two available options. By comparing this condition with the *Baseline* condition, in which subjects make an active investment decision, we can investigate the role of an active choice for subjects' memory bias. A reduced memory bias in the *NoChoice* condition would support the notion of a motivated memory bias. That is, we expect a significant positive memory bias with active choices (*Baseline* condition), but no or a negative bias in a passive choice setting (*NoChoice* condition).

Table 12 displays subjects' memory bias in the *Delay* and *Immediate* treatment groups, separately for the *Baseline* and *NoChoice* condition. The table reports mean values and T-tests results against the null hypothesis that the memory bias is zero (column 1 and 2) and T-test results of the difference in means between *Delay* and *Immediate* (column 3). The table provides two key results. First, 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.

Second, the memory 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. This is in line with the notion that people's attention to the bad rather than attention to the good is dominant when motivational reasons are limited (Baumeister, Bratslavsky, Finkenauer, and Vohs, 2001).

Table 12: Memory Bias in Active vs. Passive Choice Settings

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

Notes: 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*.

4.4 The Role of Anticipated Regret

A 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, Ward, Monterosso, Lyubomirsky, White, and Lehman, 2002).

Table 13: 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, Ward, Monterosso, Lyubomirsky, White, and Lehman, 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 (0.08)	-0.067 (0.13)	-0.024 (0.14)	0.074 (0.17)
Constant	0.067 (0.29)	0.359 (0.44)	0.038 (0.49)	-0.308 (0.60)
N	490	178	154	158
R ²	0.00	0.00	0.00	0.00

Table 13 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).

5 Conclusion

This paper investigates how memory shapes individuals' financial decisions. We find a self-serving memory bias for investment gains and losses. A key characteristic of the self-serving memory bias is the distinction between memory of events with different levels of self-relevance. Our results show that subjects relatively under-remember investment losses compared to gains, if they actively invested in the stock. By contrast, subjects who decided not to invest in the stock, or subjects who were allocated a stock randomly, do not display this memory bias. These findings are robust to different

measures of memory bias.

The self-serving memory bias is related to subjective beliefs and subsequent investment decisions. We find that subjects who invested in the stock relatively underweight negative compared to positive outcomes from memory, resulting in overly optimistic beliefs. This belief distortion increases with the number of observed negative outcomes. These findings are consistent with our model in which image-concerns form the basis for how information is remembered. We further show that those subjects do not adjust their behavior to account for the fallibility of their memory when making investment choices, which leads to investment mistakes in our experiment. They are likely to re-invest in the stock even when doing so leads to a lower expected return. A remarkable share of subjects in our *Delay* group (54% in Experiment 1 and 35% in Experiment 2) invest in the asset despite its lower expected outcome.

We explicitly designed an experiment in which subjects have to rely on their recollection of financial information and we directly elicit subjects' memory of this information at different points in time. This allows us to isolate biased memory as a microfoundation of subjective beliefs and choices, which can have important implications for financial decision-making. 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 take financial decisions that are biased by their subjective memory.

The memory bias documented in this paper opens several avenues for further research. First, we study self-serving memory bias. This complements prior work in economics showing experimental evidence for selective memory based on representativeness (Bordalo, Coffman, Gennaioli, Schwerter, and Shleifer, 2021). Theoretical and empirical investigations of the interaction between memory biases based on heuristics and motivation seem promising avenues for future research.

Second, we designed our experiment such 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, 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.

Fourth, self-image concerns 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.

Fifth, we decided to conduct a laboratory experiment that provides 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 memory bias in the field.

References

- 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.
- Andersen, Steffen, Tobin Hanspal, and Kasper Meisner Nielsen, 2019, Once Bitten, Twice Shy: The Power of Personal Experiences in Risk Taking, *Journal of Financial Economics* 132, 97–117.
- 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, 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.
- 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.
- , 2004, Willpower and Personal Rules, *Journal of Political Economy* 112, 848–886.
- , 2016, Mindful Economics: The Production, Consumption, and Value of Beliefs, *Journal of Economic Perspectives* 30, 141–64.
- Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch, 2014, root: Hamburg Registration and Organization Online Tool, *European Economic Review* 71, 117–120.
- Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, Frederik Schwerter, and Andrei Shleifer, 2021, Memory and Representativeness, *Psychological Review* 128, 71–85.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2012, Salience Theory of Choice under Risk, *Quarterly Journal of Economics* 127, 1243–1285.
- , 2013, Salience and Asset Prices, *American Economic Review* 103, 623–628.
- , 2020, Memory, Attention, and Choice, *Quarterly Journal of Economics* 135, 1399–1442.
- 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.
- Charness, Gary, Uri Gneezy, and Brianna Halladay, 2016, Experimental Methods: Pay One or Pay All, *Journal of Economic Behavior & Organization* 131, 141–150.
- 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.
- Cubitt, Robin P., Chris Starmer, and Robert Sugden, 1998, On the Validity of the Random Lottery Incentive System, *Experimental Economics* 1, 115–131.
- 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, Individual 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.
- 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.
- 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.
- Frydman, Cary, and Colin Camerer, 2016, Neural Evidence of Regret and its Implications for Investor Behavior, *Review of Financial Studies* 29, 3108–3139.
- 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.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2018, Time Varying Risk Aversion, *Journal of Financial Economics* 128, 403–421.
- Hamann, Stephan B., 2001, Cognitive and Neural Mechanisms of Emotional Memory, *Trends in Cognitive Sciences* 5, 394–400.
- , 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* pp. 1–42.
- Hartzmark, Samuel M., Samuel D. Hirshman, and Alex Imas, 2021, Ownership, Learning, and Beliefs, *Quarterly Journal of Economics* 136, 1665–1717.
- Hey, John D., and Jinkwon Lee, 2005, Do Subjects Separate or Are They Sophisticated, *Experimental Economics* 8, 233–265.
- Holt, Charles A., and Susan K. Laury, 2002, Risk Aversion and Incentive Effects, *American Economic Review* 92, 1644–1655.
- , 2005, Risk Aversion and Incentive Effects: New Data without Order Effects, *American Economic Review* 95, 902–912.
- Huffman, David, Collin Raymond, and Julia Shvets, 2019, Persistent Overconfidence and Biased Memory: Evidence from Managers, Working Paper.
- Imas, Alex, 2016, The Realization Effect: Risk-Taking after Realized versus Paper Losses, *American Economic Review* 106, 2086–2109.
- Kahan, Dan M., 2013, Ideology, Motivated Reasoning, and Cognitive Reflection: An Experimental Study, *Judgment and Decision Making* 8, 407–424.
- , Ellen Peters, Erica C. Dawson, and Paul Slovic, 2017, Motivated Numeracy and Enlightened Self-Government, *Behavioural Public Policy* 1, 54–86.
- Kahana, Michael, 2012, *Foundation of Human Memory* (Oxford University Press: Oxford, UK).
- 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.
- Knüpfer, Samuli, Elias Rantapuska, and Matti Sarvimäki, 2017, Formative Experiences and Portfolio Choice: Evidence from the Finnish Great Depression, *Journal of Finance* 72, 133–166.
- 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.
- , Sarah Rudolf, 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, and Marcia K. Johnson, 2000, Choice-Supportive Source Monitoring: Do Our Decisions Seem Better to us as We Age?, *Psychology and Aging* 15, 596–606.

- 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, 2019, Asset Pricing with Fading Memory, Working Paper.
- Paulhus, Delroy L., 1991, *Measures of Personality and Social Psychological Attitudes* . chap. Measurement and Control of Response Bias, pp. 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.
- , 2002, Inference by Believers in the Law of Small Numbers, *Quarterly Journal of Economics* 117, 775–816.
- , 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.
- Schacter, Daniel L., 1999, The Seven Sins of Memory: Insights from Psychology and Cognitive Neuroscience, *American Psychologist* 54, 182–203.
- Schwartz, Barry, Andrew Ward, John Monterosso, Sonja Lyubomirsky, Katherine White, and Darin R. Lehman, 2002, Maximizing versus Satisficing: Happiness is a Matter of Choice, *Journal of Personality and Social Psychology* 83, 1178–1197.
- Shiller, Robert J., 2017, Narrative Economics, *American Economic Review* 107, 967–1004.
- Starmer, Chris, and Robert Sugden, 1991, Does the Random-Lottery Incentive System Elicit True Preferences? An Experimental Investigation, *American Economic Review* 81, 971–978.
- Strahilevitz, Michal Ann, 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.
- Suhonen, Niko, and Jani Saastamoinen, 2017, How Do Prior Gains and Losses Affect Subsequent Risk Taking? New Evidence from Individual-Level Horse Race Bets, *Management Science* 64, 2797–2808.
- 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 Eric J. Johnson, 1990, Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice, *Management Science* 36, 643–660.

Tulving, Endel, 1972, *Organization of Memory* . chap. Episodic and Semantic Memory, pp. 381–403 (Academic Press: Oxford, UK).

——— , and David Murray, 1985, Elements of Episodic Memory, *Canadian Psychology* 26, 235–238.

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.

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

Appendices

A Experimental Instructions

A.1 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-

A.2 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 esti-

mates 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-

A.3 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-

A.4 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 3×40 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-

A.5 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.

B.1 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 10×12 EUR i.e., 120 EUR. If you don't answer any question correctly, you will earn nothing.

-Next Page-

B.2 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-

B.3 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?

-Next Page-

C Robustness of Experiment 1

C.1 Different Measures of Memory Bias

Table 14 displays the results for the memory bias measured based on the fraction of remembered positive and negative outcomes. It reports a memory bias of positive and negative outcomes for subjects who invested in the stock.²⁵ They significantly over-remember the fraction of gains and under-remember the fraction of losses (column 1, column 2, column 5, and column 6). Similar to our main results, there is no significant memory bias for subjects who did not invest in the stock (column 3, column 4, column 7, and column 8). In line with our previous results, Table 14 indicates a significant difference in subjects' memory bias, if they invested, between our conditions.

Table 15 documents results for the memory bias based on the absolute difference between recalled positive and negative outcomes. It is estimated at the individual level as the difference between the recalled and actual number of positive minus negative outcomes. The table reports similar results as we have shown previously. If subjects invested in the stock, they significantly over-remember positive outcomes and under-remember negative outcomes (columns 1, column 2, column 5, and column 6). Subjects incorrectly remember 1.62 more positive than negative outcomes compared to the actually observed outcomes ($p < 0.001$) in the *Delay* treatment (column 1 and column 2). However, if they did not invest, they do not show a memory bias (column 3 and column 4). In line with our previous results, the table indicates a significant difference in subjects' memory bias, if they invested, between our *Delay* and *Immediate* condition, suggesting a memory effect (column 9).

²⁵When 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. The table shows that the fraction for positive and negative outcomes on average adds up to zero. That is, on average subjects reported a correct number of total outcomes.

Table 14: Subjective Memory Bias based on the Recalled Fraction

	<i>Delay</i>				<i>Immediate</i>				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, 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 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 15: Subjective Memory Bias based on the Recalled Absolute Difference

	<i>Delay</i>				<i>Immediate</i>				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).

C.2 Subjects' Risk Preferences

The table shows results from linear regressions with subjects' memory bias for positive (column 1) and negative outcomes (column 2) as dependent variable. We use subjects' risk preferences as independent variable and control for session fixed effects. The results indicate that neither the memory bias for positive outcomes (column 1) nor the memory bias for negative outcomes (column 2) is significantly correlated with subjects' risk preferences.

Table 16: 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, Falk, Huffman, Sunde, Schupp, and Wagner, 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)	(2)
	Memory Bias (for Pos. Outcomes)	Memory Bias (for Neg. Outcomes)
Risk Tolerance	-0.003 (0.05)	0.030 (0.04)
Constant	-0.136 (0.43)	-0.082 (0.38)
Session	Yes	Yes
N	187	187
R ²	0.06	0.03

The table reports results from linear regressions with subjects' belief distortion relative to the Bayesian posterior as dependent variable (columns 1 and 2) as well as Logit regressions with a dummy variable equal to one for suboptimal investment in the observed stock as dependent variable (columns 3 and 4). We use subjects' memory bias for positive and negative outcomes as independent variables and control for subjects' risk preferences, subjects' financial literacy as well as session fixed effects.

In line with our previous results, the memory bias for positive outcomes is positively correlated and the memory bias for negative outcomes is negatively correlated with subjects' belief distortion ($p < 0.001$). Thus, our main finding is robust to controlling for risk preferences. 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. Moreover, subjects' memory bias is correlated with their probability to invest suboptimally, i.e. to invest in the stock with a lower expected outcome ($p < 0.005$). Hence, our finding that subjects who recollect a higher

number of positive outcomes (and a lower number of negative outcomes) than actually observed have a significantly higher probability to invest suboptimally, holds when controlling for subjects' risk preferences. Again, controlling for risk preferences did not affect this result. Further, subjects' risk preferences are not associated with their belief distortion (columns 1 and 2) or with suboptimal investment decisions (columns 3 and 4).

Table 17: 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 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), *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, Falk, Huffman, Sunde, Schupp, and Wagner, 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)	(2)	(3)	(4)
	Belief Distortion	Belief Distortion	Investment (Subopt.)	Investment (Subopt.)
Memory Bias (for Pos. Outcomes)	0.427*** (0.05)		1.607*** (0.21)	
Memory Bias (for Neg. Outcomes)		-0.528*** (0.06)		0.635*** (0.09)
Risk Tolerance	-0.019 (0.04)	-0.003 (0.04)	1.103 (0.08)	1.111 (0.08)
Financial Literacy	-0.155 (0.21)	-0.166 (0.21)	1.449 (0.62)	1.500 (0.63)
Constant	-0.046 (0.32)	-0.149 (0.31)	0.026*** (0.03)	0.022*** (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

C.3 Subjects' Financial Literacy

Table 18 reports the robustness of our findings in Table 3 of Section 3.3.1, by separately testing the significance of memory bias for subjects with different levels of financial literacy. Columns 1 to 4 display results for subjects who indicated the correct answer to the financial literacy question and columns 5 to 8 display the results for subjects who indicated an incorrect answer. The sample is limited to subjects who invested in the stock before the observation of outcomes (first choice). We find that subjects show a memory bias, irrespective of their financial literacy. They remember significantly more positive and significantly less negative outcomes than actually observed in the *Delay* treatment. Thus, higher financial knowledge does not alleviate the memory bias. This is in line with the argument that although biases in inference due to bounded rationality or limited attention decrease with cognitive abilities and sophistication, this seems not to be the case in situations when people need to rationalize away contradictory evidence, compartmentalize knowledge, and deceive themselves (Bénabou and Tirole, 2016). In such circumstances, people who are more analytically sophisticated, educated, or numerate can be more prone to making distorted inferences to protect valued beliefs (Bénabou and Tirole, 2016; Kahan, 2013; Kahan, Peters, Dawson, and Slovic, 2017).

Further, Table 19 reports that subjects' memory bias is not correlated with their financial knowledge. The table shows results from linear regressions with subjects' memory bias for positive (column 1) and negative outcomes (column 2) as dependent variable. We use subjects' measure of financial literacy as independent variable and control for session fixed effects. The results indicate that neither the memory bias for positive outcomes (column 1) nor the memory bias for negative outcomes (column 2) is significantly correlated with subjects' financial knowledge.

Table 17 shows the robustness of our main findings of Section 3.3.2 and 3.3.3 to differences in subjects' financial literacy. The table reports results from linear regressions with subjects' belief distortion compared to the Bayesian posterior as dependent variable (columns 1 and 2) as well as Logit regressions with a dummy variable equal to one for suboptimal investment in the observed stock as dependent variable (columns 3 and 4). We use subjects' memory bias for positive and negative outcomes as independent variables and control for subjects' financial literacy, subjects' risk preferences as well as session fixed effects. Subjects' financial literacy is a dummy variable equal to one for subjects who indicated a correct answer to the financial literacy question.

The results are qualitatively unchanged. Subjects' memory bias for positive outcomes is positively correlated and subjects' memory bias for negative outcomes is negatively correlated with their belief distortion ($p < 0.001$). Hence, when controlling for subjects' financial literacy, our main finding

that subjects who remember a higher number of positive outcomes (or a lower number of negative outcomes) than observed, form overly optimistic beliefs, remains. In addition, subjects' memory bias is correlated with their probability to invest in the stock with a lower expected outcome ($p < 0.005$). Thus, our finding that subjects who remember a higher number of positive outcomes (and a lower number of negative outcomes) than actually occurred, have a significantly higher probability to invest suboptimally, holds. In addition, financial literacy is correlated neither with subjects' belief distortion (columns 1 and 2) nor with suboptimal investment decisions (columns 3 and 4).

Table 18: Subjective Memory Bias and Individual Financial Literacy

	<i>Financial Literacy Question Correct</i>				<i>Financial Literacy Question Incorrect</i>			
	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.

Table 19: 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 (0.29)	-0.217 (0.25)
Constant	-0.237 (0.38)	0.126 (0.34)
Session	Yes	Yes
N	187	187
R ²	0.06	0.03

C.4 Timing of Experimental Tasks

Here, we show the robustness of our findings to differences in timing of the experimental tasks for our main findings of Section 3.3.2 and 3.3.3. Table 20 reports results from linear regressions with subjects' belief distortion compared to the Bayesian posterior as dependent variable (columns 1 and 2) as well as Logit regressions with a dummy variable equal to one for suboptimal investment in the observed stock as dependent variable (columns 3 and 4). We use subjects' memory bias for positive and negative outcomes as independent variables and control for timing as well as session fixed effects. 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 they perform the tasks in week $t + 1$.

Our main results hold. Again, subjects' memory bias for positive outcomes is positively correlated and subjects' memory bias for negative outcomes is negatively correlated with their belief distortion ($p < 0.001$). Thus, when controlling for the timing of experimental tasks in the *Immediate* condition, our main finding that subjects who remember a higher number of positive outcomes (or a lower number of negative outcomes) than observed, form overly optimistic beliefs, remains. In addition, subjects' memory bias is correlated with their probability to invest in the stock with a lower expected outcome ($p < 0.005$). Our finding that subjects who remember a higher number of positive outcomes (and a lower number of negative outcomes) than actually occurred have a significantly higher probability to invest suboptimally, remains unchanged.

Table 20: 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 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), *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)	(2)	(3)	(4)
	Belief Distortion	Belief Distortion	Investment (Subopt.)	Investment (Subopt.)
Memory Bias (for Pos. Outcomes)	0.424*** (0.05)		1.563*** (0.19)	
Memory Bias (for Neg. Outcomes)		-0.525*** (0.06)		0.661*** (0.09)
Timing	-0.401 (0.36)	-0.544 (0.35)	0.140* (0.16)	0.128* (0.14)
Constant	0.216 (0.24)	0.322 (0.23)	0.358** (0.15)	0.353** (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

C.5 No Memory Elicitation

Table 21: Comparison of Subjective Beliefs and Choices between *Delay* and *NoRecall*

	<i>Delay</i>		<i>NoRecall</i>		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).

D Sample of Experiment 2

Table 22: Number of subjects across labs and experimental conditions of Experiment 2

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