Universität Basel

Fakultät für Psychologie

On Beliefs and Learning in Investment Decisions

Inauguraldissertation zur Erlangung der Würde eines Doktors der Philosophie vorgelegt der Fakultät für Psychologie der Universität Basel von

Kevin Trutmann

aus Küssnacht a.R.

Basel, 2023

Universität Basel

Fakultät für Psychologie

Genehmigt von der Fakultät für Psychologie auf Antrag von

Prof. Dr. Jörg Rieskamp

Prof. Dr. Sebastian Gluth

Datum des Doktoratsexamen:

Dekan:in der Fakultät für Psychologie

Universität Basel

Fakultät für Psychologie

Erklärung zur wissenschaftlichen Lauterkeit

Ich erkläre hiermit, dass die vorliegende Arbeit ohne die Hilfe Dritter und ohne Benutzung anderer als der angegebenen Hilfsmittel selbstständig verfasst habe. Zu Hilfe genommene Quellen sind als solche gekennzeichnet. Die veröffentlichten oder zur Veröffentlichung in Zeitschriften eingereichten Manuskripte wurden in Zusammenarbeit mit den Koautoren erstellt und von keinem der Beteiligten an anderer Stelle publiziert, zur Publikation eingereicht, oder einer anderen Prüfungsbehörde als Qualifikationsarbeit vorgelegt. Es handelt sich dabei um folgende Manuskripte:

1. Belief Updating and Investment Decisions: The Impact of Good or Bad News Varies With Prior Returns

2. Take your Time: How Delayed Information and Restricted Decision Opportunities Improve Belief Formation in Investment Decisions

3. Degree of Involvement in Decisions and the Likelihood to Stop Investing among Professionals

Spezifizierung des eigenen Forschungsbeitrags zu den Manuskripten:

1. Eigener Beitrag nach CRediT[1]:

☒ Conceptualization ☒ Data curation ☒ Formal Analysis
☐ Funding acquisition ☒ Investigation ☒ Methodology
☒ Project administration ☒ Resources ☒ Software
☐ Supervision ☒ Validation ☒ Visualization
☒ Writing – original draft
☒ Writing – review & editing

Das Manuskript wurde bisher für keine anderen Qualifikationsarbeiten eingereicht.

2. Eigener Beitrag nach CRediT[1]:

☒ Conceptualization ☒ Data curation ☒ Formal Analysis
☐ Funding acquisition ☒ Investigation ☒ Methodology
☒ Project administration ☒ Resources ☒ Software
☐ Supervision ☒ Validation ☒ Visualization
☒ Writing – original draft
☒ Writing – review & editing

Das Manuskript wurde bisher für keine anderen Qualifikationsarbeiten eingereicht.

3. Eigener Beitrag nach CRediT[1]:

☐ Conceptualization ☒ Data curation ☒ Formal Analysis
☐ Funding acquisition ☒ Investigation ☐ Methodology
☐ Project administration ☐ Resources ☐ Software
☐ Supervision ☒ Validation ☒ Visualization
☒ Writing – original draft
☒ Writing – review & editing

Das Manuskript wurde bisher für keine anderen Qualifikationsarbeiten eingereicht.

[1] https://casrai.org/credit/

Universität Basel

Fakultät für Psychologie

Open-Science Aspekte der Manuskripte:

1. Preregistration: ☒ ja ☐ nein
   Open-Access-Publikation: ☐ ja ☐ nein
   Open-Access-Data/Analyse: ☒ ja ☐ nein
   Ort/URL der Daten und Analysen: https://osf.io/j6fus

2. Preregistration: ☒ ja ☐ nein
   Open-Access-Publikation: ☒ ja ☐ nein
   Open-Access-Data/Analyse: ☒ ja ☐ nein
   Ort/URL der Daten und Analysen: https://osf.io/kxqda

3. Preregistration: ☐ ja ☒ nein
   Open-Access-Publikation: ☐ ja ☐ nein
   Open-Access-Data/Analyse: ☒ ja ☐ nein
   Ort/URL der Daten und Analysen: https://osf.io/5hm24

Ort, Datum Basel, 26.04.2023

Signatur ………………………………………………………

Vorname Nachname Kevin Trutmann

Acknowledgements

Getting a PhD is a long journey. Not only that, but it is one where rough patches of drought-ridden straits, thorny undergrowth, and slippery mountaintops are to be expected. It is all the more important to make sure not to embark on this journey alone. I am thankful for all the people who kept me focused, entertained, nudged me back on the path when I swerved or even carried me for a bit when I felt I couldn't do it.

First and foremost of course I want to thank my wife Angie. I remember you once looking over my shoulder, seeing all the code and command lines opened on my screen and asking "do *you* even know what you're doing here?", making me pause and think for a moment "*do* I?". Still, you always supported me wherever possible and, of course, when love is shared it multiplies, so you gave me the greatest gift of all: Our sons Nelio and Nino. Spending time with you all was and still is the best way to replenish my energy and I am immensely grateful for having you.

Next, I want to thank Jörg Rieskamp for giving me the opportunity to do this PhD and for his helpful feedback along the way. Also Sebastian Gluth, who served as my secondary PdD advisor and who played a big part in making this whole journey possible. Steve, thank you for being a great colleague, coworker and friend. You were the best guide I could hope for when, as a psychologist, straying into the sometimes murky and unknown waters of economics. The book "Microeconomic Theory" you gave me served me greatly by raising my computer screen at home by a few centimeters so I could sit with healthy posture. Joking aside, it was also very insightful when reading it!

Speaking (or writing) of reading, I also want to thank Dimitris, with whom I read more of "Microeconomic Theory", as well as "Statistical Rethinking", than was strictly necessary for our research. We called these reading sessions our "mathematical self-defense reading group", and I think our stats and theory Jiu-jitsu greatly profited from it. The discussions with you were always extremely interesting and enlightening and I'm looking forward to many more years of friendship.

Thank you also to all the other people who helped me along the way, either directly or indirectly by creating a stimulating environment. Laura, Dominik, Rebecca, Nathaniel, Jana and so many more who have come and gone over the course of these years, thank you so much! Lastly, I also want to thank my family, especially my parents for always being supportive and, well, creating me. This whole "existing" thing is certainly an interesting concept, so thank you! I could not have done it without you. Literally.

Abstract

Whether we enjoy the fruits of a good decision or experience the fallout of a bad one, we are always haunted by our past to some extent. Whatever the outcome may be, it is crucial to learn from this past, especially if a new, maybe similar, decision awaits. The question arises which aspects can be used to improve the expected outcome of the new decision and which aspects are better disregarded. This thesis consists of three manuscripts in which I investigate how this question is answered in the domain of investment decisions. In doing so I also explore three potential remedies to averse effects of clinging to prior decisions: Improved information, increased mental distance and decreased personal involvement.

The first manuscript looks specifically at how previous gains or losses influence the way in which a decision maker learns from a new signal. I show that this investment position (i.e., whether the investment has created a gain or a loss so far) interacts with the favorability of the new information (i.e., whether it is good or bad news). This interaction can lead to beliefs that may underlie well known and profit harming investment behaviors such as the Disposition Effect. I develop an extended Reinforcement Learning model that captures this interaction in the learning rate and show that it represents the data of an investment experiment better than competing models. In a second part of the experiment I also show that this interaction effect can only be overcome if participants are provided with very clear information such as revealing the true probability of the next price move.

The second manuscript builds upon the findings of the first. It presents a different investment experiment in which I aim to improve participants' belief updating by increasing the perceived mental distance from their initial investment decision as well as the development of their investment. In the first of two treatment conditions participants were blocked from changing their investment for a number of rounds. In the second treatment condition participants could neither change their investment nor track its development for the same number of periods. I show that in the treatment conditions participants' beliefs indeed approach those of a Bayesian updater calculating the objective probabilities. However, only the condition in which both the investment and information was blocked was sufficient to render this improvement significant.

Finally, the third manuscript continues the theme of involvement. It uses a sample of professional decision makers to replicate a study aimed at investigating the sunk cost fallacy. In the experimental task, participants could choose one of six lotteries. These lotteries were either framed as an investment into a startup or as the hiring of a consultant, that would then invest on the participants' behalf. Participants were generally less likely to change their chosen investment option after a bad outcome if their previous decision was framed as their own investment (i.e. they were directly involved) rather than as hiring a consultant (and thereby having "someone to blame"). Additionally I find that this effect was only significant for younger participants, which could indicate that these negative effects of involvement may fade with increasing experience.

Contents

Chapter 1

Introduction

"If seven maids with seven mops
Swept it for half a year,
Do you suppose," the Walrus
said, "That they could get it
clear?"
"I doubt it," said the Carpenter,
And shed a bitter tear.

–Lewis Carroll,
The Walrus and the Carpenter

The process of taking in information and using it to arrive at a decision is far from being well understood, despite it being ubiquitous in everyday life. One might have to decide within a split second what to do to avoid a collision while driving; One might ponder endlessly about what food to order from the menu. From the essential to the mundane, we must assume that there are common mechanisms at play, influencing how we ultimately arrive at a decision. Firstly, one needs an internal model of the world, ideally one which corresponds well with the physical outside world and is informed by all the relevant information gathered so far. Next, one needs to make a prediction about possible outcomes. Will I be able to brake in time or should I swerve to the side? Will I like the vegetarian burger, or will I regret not getting the pasta? This internal model and the beliefs about future outcomes also need to be regularly updated with any new relevant information that occurs (Rangel et al., 2008). A friend might mentions that they did not like the pasta that much the last time they were at this restaurant. This information can be used to adjust ones beliefs. It is now unlikely that you will regret not getting the pasta. The focus of this cumulative dissertation will lay on the last aspect described in this example: What happens when new information is gathered, how does it influence our beliefs and thereby our decisions, and how is this influence changed by the decisions that came before?

Multiple approaches can be taken when studying these mechanisms. One is to rely on situations in which repeated decisions are naturally made and data can be readily collected. For example, one can look at the data of people investing in the stock market and try to understand their decisions from the recorded trades. A clear drawback of this approach is that the environment in which the decisions are made and recorded is inherently noisy, and the researcher has no control over it. If these aspects are of primary importance, a better idea might be to run a laboratory study with a simplified task, such as decisions between risky lotteries. There, each aspect of the experiment is under the full control of the researcher. These risky decision tasks have become the workhorse of economic psychology and have lead to many insights into peoples risk calculation and decision making processes. Nonetheless, it is a strong abstraction from situations one might encounter in real life. Between these two extremes — from extreme realism to extreme control — simulated investment tasks offer a great compromise. While they can closely simulate decisions that are encountered by many people in daily life, they also allow the researcher some control by setting the attributes of the simulated asset market. For this reason the first two manuscripts make use of such a simulated investment task framework, while the last manuscript also strikes a similar balance between realism and control in its experimental design.

The first manuscript in this dissertation (Chapter 2) focuses on establishing the influence of the context in which new information is received.[1] In detail, it shows how, in the domain of investment decisions, a good or bad situation (having made a gain or a loss with an investment) interacts with good or bad news (a value in- or decrease of the investment) to influence how the new information impacts investors' beliefs. The second manuscript (Chapter 3) builds upon this finding. It postulates that the influence of these situational factors ("context") may be due to the investor's strong involvement with their investment and their decision to invest. In line with this hypothesis it shows that an intervention in which

[1] Note, that here the word "context" is not used to describe alternative options in the choice environment such as in the "context effects" (see for example Busemeyer et al., 2007). Instead, the "context" is here specifically comprised of two variables described further in the first manuscript: Whether the information is received while in a good or bad situation and whether the type of investment renders the information favorable or not.

participants have to wait a number of rounds before receiving information about how their investment fared does significantly improve belief formation. Lastly, the third manuscript (Chapter 4) investigates the role of involvement more directly. In an experiment mirroring that of Chang et al. (2016) a sample of professional decision makers are less willing to let go of a losing investment if it is framed as "their decision to invest into a startup" rather than "them hiring a consultant to invest in their stead". The remainder of this chapter will place the manuscripts in this thesis in their broader context in the literature as well as provide further background information.

1.1 Investment Decisions: Preferences and Beliefs

When pondering about whether to buy or sell an investment, one must generally consider two main aspects. The first is ones preferences: How great would it be to make a high profit and how much would it hurt if the investment was a failure? Much effort has been spent on researching peoples preferences in risky decisions. This effort has lead to much insight about why people behave in ways that prevent them from maximizing their payoff, based on "suboptimal preferences" (Tversky & Kahneman, 1992; Barberis & Xiong, 2012; Odean, 1998). The second aspect is the probabilities that one ascribes to the possible outcomes, or in other words the "beliefs" about the outcome probabilities. As much pleasure as a high return might bring, if it is very unlikely to happen the investment may not be worth the risk.

Whether one watches each trade closely or checks their investments only once in a while, the assets value will most likely have in- or decreased by some amount. Each of these price movements can be viewed as an entry in a series of binary outcomes (Oskarsson et al., 2009) the next of which is to be predicted. An ever growing body of literature now investigates the role that these outcome expectations and beliefs play in investment decisions (e.g., Grosshans et al., 2020; Greenwood & Shleifer, 2014; Kuhnen, 2015; Kuhnen et al., 2017; Fischer & Maier, 2019). The research in this thesis points toward the possibility that part of the mismatch between the observed behavior and that which would maximize profits could be due to the way expectations are formed.

1.2 Motivated Beliefs

While risk preferences are generally considered to be relatively constant (especially within a certain domain, Frey et al., 2017), beliefs and expectations are highly dependent on the current situation as well as the surrounding information and are therefore more malleable. One explanation for how such beliefs may divert from the ones held by a rational actor is presented by Fischer & Maier (2019), who describe a mechanism of "motivated belief distortion". In this mechanism, a decision makers' beliefs are linked to their valuation, that is, their preferences. In detail, motivated belief distortion would lead someone to believe what seems most favorable to them in terms of utility (Fischer & Maier, 2019; Bénabou, 2015), especially in ego-relevant tasks (Drobner, 2022, but see also Ertac, 2011 who generally find pessimism in ego-relevant feedback). Some evidence even suggests that people are not a "victim" of these mechanisms, but are actively willing to distort their beliefs, given the right circumstances (Saccardo & Serra-Garcia, 2023).

Bénabou & Tirole (2016) distinguish three strategies of self-deception leading to motivated beliefs: strategic ignorance, reality denial and self-signaling. The manuscripts in this thesis are mainly concerned with the reality denial strategy, as these strategies work not by completely ignoring or avoiding information, but under- or overweighting its importance. Whether such reality denial takes place depends on a number of factors, one of which I propose to be the involvement of a decision maker with their decision, or the contrary, the psychological distance to the previous decision. Psychological distance is a well established concept stemming from construal level theory (Trope & Liberman, 2003) and can take on multiple forms. Most important for this thesis are the temporal and social distance as well as the informational and affective distance concept proposed by Fiedler (2007). Especially for the concepts of affective and social distance, respectively, aspects such as ego-relevance of the outcome (Drobner, 2022) and personal responsibility (Chang et al., 2016; Martens & Orzen, 2021) can play a critical role.

One mechanism which may then decrease psychological distance and thereby drive beliefs to get distorted is that of *cognitive dissonance* (Festinger, 1962). As an example, assume an investor finds themselves having made a "paper loss", that is, a loss which they have not yet "realized" by selling the asset. This loss is at odds with a potential, ego relevant, inner narrative of being a good investor who makes good investment decisions. To resolve this dissonance, the investor can either correct the narrative in accordance with the new evidence (and admit to having made a mistake), or they can adjust their belief about the investment itself to fit the narrative (believing that the price will surely bounce back, and that the investment was a good idea after all). It is clearly in the best interest of ones self worth to adjust the beliefs and convince oneself of a brighter future than may be warranted.

Such a mechanism of belief updating could for example lead to overconfidence which impacts real investment decisions and can lead to behaviors such as the disposition effect. The disposition effect is the tendency of investors to sell gains very quickly while holding on to losses for longer (Shefrin & Statman, 1985). As stock prices often continue their momentum (at least historically for three to 12 months; Jegadeesh & Titman, 1993) this behavior can hurt an investor's overall return. This connection between overconfidence and the disposition effect was also drawn by Hanspal (2017), who finds a group of Danish investors with a high disposition effect measure to have significantly higher expectations about future returns than a matched group with lower disposition effect values.

A similar but distinct belief-based mechanism is mentioned by Odean (1998) who proposed a belief in mean reversion of the price to potentially underlie the disposition effect. An investor who believes in mean reversion will not only assume that a loss will be recovered, as was the case for cognitive dissonance, but will assume that any gain is also only due to temporary fluctuations of the price. Overall, such an investor believes that the price will always tend to revert back to the initial buying price (i.e., a belief in "buying price reversion"). This assumption is closely linked to the concept of mental accounting, where making a new investment can be viewed as "opening up a new mental account". An investor will then track gains and losses not on the level of the portfolio but relative to the initial balance of that account (Shefrin & Statman, 1985; Thaler, 1999).

Under the assumption of such a belief in buying price reversion it would be reasonable to hold on to losses and sell gains as quickly as possible, hence leading to the disposition effect. Jiao (2017) finds evidence that such mean reversion beliefs significantly contributed towards the disposition effect in an experimental investment task. He however reports notable between subject variability, with 12.5% of participants reporting extrapolative beliefs within their investments. Similarly mixed evidence can be found in a seminal experiment by Weber & Camerer (1998), in which participants invested in price paths with upward or downward trends, trying to discern which might be the most profitable. Here the authors find that people do indeed buy more stocks after losses than after gains, hinting at a belief in mean reversion. However, in a second condition in which stocks where sold automatically after each period and could be bought back at no additional cost participants more willingly let go of their losses, indicating that the disposition effect stemmed from a reluctance to sell rather than from misguided beliefs. One could however argue that participants' knowledge about the "one shot nature" of the automatically sold investments would remove the need for mental accounting and could thereby have influenced how they formed their expectations, mitigating a belief in mean reversion that could otherwise have developed. While the authors also show that participants where proficient in recognizing which assets had an upward trend, they did not directly elicit participants' expectations about the price movements during the study.

1.3 Reinforcement Learning

When investigating how people learn and adapt to new information it is essential to have a model of this updating process. The most prevalent of these models found in the literature are Bayesian updating and reinforcement learning. Bayesian updating is based upon the principles of probability theory and describes how one would optimally incorporate new information into a prior held belief. Reinforcement learning simplifies this task by proposing that an updated belief always consists of a weighted average between the previously held belief, and the newly acquired evidence. While it is often used to update the value a decision maker ascribes to an option, the same technique can be used to reflect the formation of expectations and learning about the probability of an outcome (Sutton & Barto, 2018). While the mechanism behind reinforcement learning is computationally rather simple, it is still capable of explaining behavioral phenomena such as recency and the underestimation of rare events in decisions from experience (see e.g. Hertwig & Pleskac, 2010). Even extending beyond the investing paradigm, Erev & Roth (1998) find that people are best described as reinforcement learners in a multitude of economic games. Lastly, the prediction errors that are necessary for reinforcement learning strongly correlate with measurable brain activity, further lending it some biological foundation (Schultz et al., 1997; Rangel et al., 2008; Schultz, 2015). For these reasons the general approach and main analysis in the first manuscript (Chapter 2) are based on the assumption that the study participants update their beliefs using a reinforcement learning mechanism.

Two points are worth noting here: While this thesis focuses on the updating of expectations and beliefs throughout an investment an investor would not only learn within, i.e., during investments, but also between them. As an example, Malmendier & Nagel (2016) show that experiencing extreme macro-economic conditions, such as a financial crisis with high inflation, can have a long lasting impact on how investors form their expectations about inflation in the future, especially when these experiences were made at a younger age. Secondly, while investors can and will often consult outside sources of information the information in the studies in this thesis is located entirely within the value changes of the investment itself. This allows for a simpler design of the study while retaining the most important aspects of the investment decisions.

Context Dependent Reinforcement Learning

When applying reinforcement learning models, one constant learning rate parameter is generally used to capture the weight given to new evidence. However, the reality denial strategy of self-deception as defined by Bénabou & Tirole (2016) can also be expressed as a form of distorted learning; A piece of information about reality is learned, but not fully integrated into ones model of the environment. The weight that is given to the new information is not necessarily static, but can change with the context of the information. In this way, motivated beliefs may be based on motivated learning. Note also that in contrast to the aforementioned belief in buying price reversion being held *a priori*, motivated learning could induce such a belief only once the investor finds themselves in the relevant situation. In other words, being invested oneself could be crucial for the formation of this belief pattern.

In line with this assumption of a dynamic learning rate, Kuhnen (2015) finds that investors learn differently from gains compared to losses. Specifically, participants in her study updated their beliefs more strongly when seeing a low outcome in the loss domain. Studies by Knutson & Bossaerts (2007) and Seymour et al. (2007) further corroborate the idea of dynamic learning rates as they find evidence that different brain areas are involved in processing gains as opposed to losses. Kaustia & Knüpfer (2008) on the other hand show that involvement plays an important role in this effect of dynamic learning as well. They show that investors who personally experienced higher returns from participating in an initial public offering (IPO) of a stock also decided to hold more shares in these companies later on. The same pattern is found in retirement funds by Choi et al. (2009). Here, people invested more into their retirement fund when they personally experienced positive returns from prior investments. Lastly, Kuhnen et al. (2017) show that future expectations are especially influenced by ones own previous decisions. In her experiment participants disregarded information that was not in line with their previous investment decisions, leading to "sticky portfolios". While these differences in belief formation may happen by giving more weight to some lived experience than to others during the updating process, there exists also the possibility of "motivated memory", i.e., beliefs being influenced by remembering certain experiences better than others (Amelio & Zimmermann, 2023).

Studies such as the one by Kuhnen (2015) show that the domain, that is, having made a previous loss or gain plays an important role in belief updating. Sharot & Garrett (2016) argue that the valence, that is, whether information is favorable or unfavorable, is an additional important factor. There is however no unequivocal evidence on whether there is a general "processing advantage" (and thereby a stronger influence on beliefs) of negative or positive information (Unkelbach et al., 2020). We solve this conundrum in the first manuscript (Chapter 2) by combining the two aspects mentioned above and proposing a further differentiation of favorable and unfavorable information into that received after having made a gain and that received after having made a loss.

1.4 First Manuscript: Belief Updating and Investment Decisions

The focus of the first manuscript (Chapter 2) lays on the way in which investors form their beliefs and the influence of two particular factors: First, the current position of the investment, i.e. whether so far a gain or a loss has been made. Second, whether new information is favorable or unfavorable towards the investment. Using these two factors, belief updating can happen in a way that ultimately leads to beliefs that resemble a belief in buying price reversion. As described above, the disposition effect can be shown to be a natural consequence of a belief in buying price reversion, which in turn can be the consequence of motivated beliefs.

I adapt a task introduced by Frydman et al. (2014) in which participants can invest in a simulated asset. This asset has two states, expressed through an upward or downward drift in its value. The current state of the asset is governed by a Markov-Chain process, which is set up such that in each round the state remains the same but has some small probability of switching. Participants were informed about this system. They therefore knew that they could gather whether the asset was more likely to currently have an up- or downward drift from observing the price changes, yet they could never be sure it would remain that way. The main outcome measures in this task were twofold: First, the investments participants made. This could be either owning one share, no shares or short selling one share. Short selling is an investment technique in which an asset is loaned, sold and then bought back later, hopefully at a lower price, to return it to the loaner. This gave our participants the opportunity to also profit from downward drifts. The second main outcome was participants' belief reports. Here I asked participants to indicate how likely they thought it to be that the price of the asset would increase in the next round.

After establishing that participants do follow their beliefs with their investments, I investigate how belief updates — that is, the difference between belief reports from one round to the next — are influenced by the factors introduced above. Looking at the updating magnitudes I do find a strong interaction between the investment position (i.e., whether the investment so far has made a gain or a loss) and price movement favorability (i.e., whether the price move in- or decreased the portfolio's value). However, this interaction interestingly did not quite follow the hypothesized pattern. While in a gain position beliefs are updated more strongly from unfavorable information (in line with the hypothesis), this is also true for a loss position (contrary to the hypothesis). There is however an (expected) strong effect of price movement favorability, in which unfavorable information is updated more strongly in general. Removing this effect would lead the pattern of belief updating to match the hypothesis.

In addition to the direct analysis of the price updates, I also translated the concept of the differential learning rates into an extended reinforcement learning model. The context sensitive reinforcement learning (CSRL) model incorporates different learning rates for each of the possible combinations of the two factors described above, as well as a separate rate for the state of not being invested. Using Bayesian model estimation and a hierarchical implementation of the model I confirm a credible interaction between investment position and information favorability. Further, using the modelling approach, the learning rates follow the expected pattern: Beliefs are updated less strongly when facing favorable information after having made a gain and unfavorable information after having made a loss.

In a second phase of the experiment I aimed to mitigate the impact of these belief updating effects by providing participants with more precise information, thereby hopefully supporting the process of belief updating. In a first treatment condition participants were shown the pre-calculated Bayesian probability of a price increase in the next round. This condition technically does not add any new information that a savvy participant could not have calculated. The second treatment condition however revealed the current state of the asset completely, therefore adding information that before could have only been inferred with uncertainty from price moves. Taking the disposition effect as a target measure, I find the additional information to improve participants' disposition effect measures by bringing them closer to the benchmark of a risk-neutral Bayesian investor. However, only the intervention in which the state of the asset is completely revealed lead to a statistically significant improvement. This result could be indicative of participants knowingly distorting their expectations following a "reality denial" strategy and thereby changing their investment decisions despite "knowing better".

In summary I show that the impact of new information can be influenced by outside factors such as the current position of the investment as well as the favorability of the information. In detail, these factors seem to interact in a way that leads to beliefs in buying price reversion which can have detrimental consequences for investment decisions. Very strong information (i.e., revealing the true state of the asset) was necessary to overcome this interaction. While the effect seems to be robust overall, it is worth mentioning that there seems to be a high heterogeneity between participants in all aspects, but also in their estimated learning rates. This between person heterogeneity appears as a fruitful point of interest for future investigations. It is also in line with the results by Jiao (2017), who, despite a prevalent belief in mean reversion, finds a substantial subgroup to hold extrapolative beliefs. This heterogeneity also allows for subgroups of investors who may pursue even simpler heuristics such as a following a win-stay-lose-shift strategy. In addition, jointly estimating the investors' preference functions was outside the scope of this project, but may also prove promising to gain further understanding of investors' the decision processes.

1.5 Second Manuscript: Take your Time

The second manuscript (Chapter 3) builds upon the findings reported in Chapter 2. Again, participants were given the opportunity to invest into an asset with either an up- or downward drift, the nature of which they had to deduce from the evidence given by the observed price movements. In this experiment however the price paths were kept shorter, as the drift direction of the asset was stable. This setup is similar to that used in an earlier study by Weber & Camerer (1998). After an initial period of three rounds participants in the baseline condition were free to invest or short sell the asset as they pleased from round to round. In a second condition, called the *blocked trades* condition participants were forced to select, after the initial three rounds, an investment which would then be held for the rest of the price-path (five rounds) except for a final decision. The third condition, the *delayed information* condition had the same restriction as the blocked trades condition. However, additionally participants could not watch the price move from round to round, but only received them in list form after all rounds had already played out.

The goal of these interventions was to increase the psychological distance between the investor's past and current decision. The blocked trades condition aimed to do so by letting investors focus on updating their beliefs using the new information, rather than making a new consequential decision in each round. In the delayed information condition this distance could be increased even more, as participants were not able to track each price move over time (and thereby "drive the roller coaster"), but only ponder the results of their initial investment after some investment periods had passed. Overall, such increased distance would imply a lower involvement with the price development, which in turn could lead to less overreaction in belief updating and investment decisions (Kaustia & Knüpfer, 2008). For these reasons I expected the interventions to improve participants' belief updating.

To measure the effect of these interventions, I focus on the round before the final investment decision. There, I calculate the distance between participants' reported beliefs in a price increase and that of a Bayesian updater. The second target measure was the success of the investments made in that final round. This success rate was measured by comparing whether the investment (holding or short selling a share) corresponded to the drift of the given asset (up or down). Interestingly, I do find that participants' beliefs approach that of a Bayesian updater in both treatment conditions, yet this difference only reaches statistical significance in the delayed information condition. As the effect was also rather small, this impact on the beliefs does not translate into more successful investment decisions in the final investment period.

Interestingly, when analyzing the data further, I do find the drift, that is, whether the asset had an up- or downward drift, as a more reliable predictor for investment success. Overall, participants made better investment decisions in the final round of an upward drift and the interventions used in the treatment conditions lead to greater improvements during downward drifts. Note, that downward drifts require short selling to make a profitable investment. Short selling adds a layer of complexity by flipping the meaning of price movements around, that is, price increases are unfavorable while price decreases are favorable for such an investment. The finding that the treatment conditions improved beliefs significantly during a downward drift could indicate that this additional complication may be in part responsible for participants' beliefs differing strongly from those of a Bayesian updater (i.e., participants making the best of the limited resources they have left; Simon, 1990). An interesting avenue for future studies could therefore be to evaluate the interaction between the complexity of an investment endeavor and the accuracy of belief formation.

1.6 Third Manuscript: Degree of Involvement

In the third manuscript (Chapter 4) the focus is laid on the aspect of involvement. I evaluate whether experience could mitigate the negative effects of feeling responsible for a decision that lead to a unfavorable outcome. In other words, the main focus of this study was on the influence of involvement on the effects of sunk costs (Martens & Orzen, 2021).

The basic structure of the experiment follows that described in Chang et al. (2016): Participants are asked to choose one of five investment options. After receiving feedback about the investment's performance, they are asked to either stay with the investment or switch to another investment option. Cognitive dissonance would dictate the prevalence of the sunk cost fallacy, possibly through belief based mechanisms such as the ones shown in the previous manuscripts. An investor who has received feedback that the investment so far was a failure is inclined to overweight any positive aspect of the investment and convince themselves that it was a good idea to invest in the first place. Following Chang et al. (2016) the study contains a second condition in which participants are given a "way out" of this dissonance. In this condition, participants are not asked to make an investment themselves, but pick a "consultant" to make the investment for them. The return options and their probabilities remain identical, and this is not hidden from participants. Hence, the only difference between conditions is that participants do not view the investment as their own decision, but as the decision of the consultant they picked (and who they could easily "fire" in the next round). Participants in this study were recruited from the controlling team of a large Swiss infrastructure company. The study therefore constitutes a "lab in the field" replication of Chang et al. (2016) with the main differences being the participant sample as well as the framing of the decisions. In detail, while the Chang et al. (2016) paper frames the decisions to their student participants as investments in the stock market, the cover story in the present study is about startup-businesses.

Ultimately, the results replicate the findings of Chang et al. (2016), showing that having "someone to blame", i.e., a consultant in this case, increases the chance of ending an unprofitable investment. However, I also find that this effect vanishes with age. This indicates that, with more experience investors may be more willing to let go of an "error", even when they have no one to blame but themselves.

1.7 Discussion

The overarching theme of the manuscripts in this dissertation stems from past decisions and circumstances influencing our beliefs and decision making. The first manuscript points out the influence of circumstances on belief formation and how these beliefs in turn influence the decisions made based upon them. It also implements an intervention based on providing more and clearer information. Very clear information, that is, displaying the current drift of the asset in the study, was necessary to overcome the negative effects of context sensitive belief updating. The second manuscript builds upon these findings and implements a new intervention based on the involvement (or its opposite, creating psychological distance). Here again, only the strongest intervention, that is, blocking investments as well as delaying the information, lead to significant improvements in belief formation. While this study provides high external validity due to the investment task used, it still may prove fruitful for future endeavors to abstract the intervention further. In detail, it would be interesting to directly compare learning between tasks with a sampling (low involvement) and a repeated choice (high involvement) paradigm (Gonzalez & Dutt, 2011). Lastly, the final manuscript is concerned with another phenomenon based in involvement and past decisions: the sunk cost fallacy. This study replicates a previous finding by Chang et al. (2016) in a sample of professional decision makers and expands the analysis to include age as a proxy measure for experience. I show that the effects of sunk costs seem to decrease with age and conclude that experience with a certain decision environment can indeed improve decisions by inhibiting the influence of irrelevant circumstantial factors such as who made the original decision.

While the CSRL model certainly is a step in the right direction in incorporating the influences on belief updating, it certainly does not capture everything perfectly. The experiments used here are done in a controlled laboratory environment and participants are well prepared by the instructions to find out the optimal solution to the tasks given. Investment decisions in the real world are messier, adding many more outcome conditions and information sources. Still, current models of belief formation and decision making struggle to capture the full extent of even the simplified decisions in the laboratory. While there are many possible reasons for these shortcomings, the following paragraphs will discuss three of them as well as their implications for future investigations.

First, the process of belief formation, belief reporting as well as decision making is inherently noisy. The noise in the decision stage is often captured and explicitly modeled by a stochastic choice rule in which options are chosen with a probability that is proportional to its value, but never certain (McElreath, 2020; Farrell & Lewandowsky, 2018). Noise in the elicitation of beliefs can be mitigated to a certain extent by choosing the "right truth serum for the right occasion" (Trautmann & van de Kuilen, 2015), which I believe to have done in the reported experiments. Lastly, the noise in the updating of beliefs and expectations can also be modeled to a certain extent. The context sensitivity aspect of the CSRL model is one factor that may reduce what would otherwise be interpreted as noise in the updating process. However, there may be more factors that lay outside of the scope of a model such as momentary attention or interactions with the aforementioned motivated memory (Amelio & Zimmermann, 2023). Distinguishing the different sources of noise — updating, reporting and deciding — is a task that requires thoughtful experimental design and must be kept in mind for future investigations.

The second factor that current models may struggle to capture lies in the heterogeneity between people. When investigating individual participants' learning rate posterior distributions (omitted in the manuscript), many seem to adhere very clearly to the proposed CSRL model, while others distributions are highly dispersed. This indicates that some participants may use other strategies than the proposed belief updating mechanism. As an example, simplifications such as counting heuristics or win-stay-lose-shift strategies can often yield acceptable results in simplified laboratory tasks while being cognitively far less taxing. Future endeavors may find it promising to classify these inter-personal differences. This could help in understanding the strategy-selection problem which has to be solved before participants can start solving the task itself. Such a classification could also bring the opportunity to focus on a specific subgroup of investors at a time and thereby obtain clearer results and a better understanding of the underlying mechanisms.

The last problem, and most likely the hardest to address, is the question of whether belief formation and decision making truly works in the clear sequential way that is often implicitly assumed. While sequential updating and reinforcement learning seem to be sensible strategies for the tasks in the studies presented here, decisions in the field are often spread out over a longer period. It is therefore unclear whether intermediate values such as beliefs are truly "kept in mind" at all times or whether beliefs and values are only constructed and reconstructed once they are needed. A similar argument has already been made for utility calculations and preferences, shaking at one of the fundamental axioms of economic decision theory, the axiom of completeness (Murawski & Bossaerts, 2016). This point may be addressed by spreading learning experiments over a longer period in which working memory plays less of a role and reconstruction of previously remembered expectations becomes more important. Preliminary evidence for a "spacing effect" in value learning (Wimmer et al., 2018) indicates that it may indeed prove fruitful to combine the model presented in the first manuscript with models of reconstructive memory (e.g. Albrecht et al., 2021).

1.8 Conclusion

The manuscripts in this thesis add to the literature by exploring the effect past decisions and circumstances have on how new information is used and how decisions are made. They provide three possible remedies for averse effects of past decisions and contextual influences based in providing information, increasing psychological distance and reducing personal involvement. Additionally the first manuscript provides and evaluates a new formal model for capturing the situational effects in belief updating and explaining their consequences. A better understanding of the effects investigated here can not only help in finding the situations in which they are detrimental but also in developing interventions that ultimately can help decision makers in aligning their beliefs with their given information and their decisions with their best interest.

Chapter 2

Belief Updating and Investment Decisions: The Impact of Good or Bad News Varies With Prior Returns

Kevin Trutmann, Steve Heinke, Jörg Rieskamp

The manuscript has been submitted for publication.

Abstract: Investors' belief updating differs for investments in a gain position versus those in a loss position and by the favorability of the news, leading to a anomalies in investment decisions. We propose a context-sensitive reinforcement learning model unifying these empirical findings. In a preregistered experiment we show that the model captures investors' belief dynamics best. We observe stronger changes in beliefs from negative information in gain compared to loss positions, leading to profit-harming decisions. Providing participants with additional information mitigates these effects. Our findings have implications for theories incorporating belief dynamics and applications improving financial and economic decisions.

2.1 Introduction

Investors' expectations about the value of an asset should be moved only by new information that is truly predictive for the asset's future profits. In contrast, we demonstrate a context sensitivity in belief updating independent of the information's value for future earnings. Combining prior literature we propose a unifying belief formation mechanism incorporating differences in the belief updating dependent on the investment position (i.e., gain or loss), information favorability (i.e., favorable or unfavorable) and their interaction. The belief updating dynamics resemble a belief in "buying-price reversion" (also called "mean reversion"; Odean, 1998; Asparouhova et al., 2009; Jiao, 2017), i.e. investors think that the price will always revert back to the one at which the investment was initially made. We formalize this idea of belief updating dynamics by extending the reinforcement learning framework using context sensitive learning parameters. Estimating this model confirms the expected pattern in belief dynamics and reveals that the model complexity is justified by the data. These patterns in belief updating can serve as part of a unifying mechanism for well-known profit-harming trading patterns in professional and retail investors alike.

Past research indicates that investors have difficulties forming expectations accurately and they do not strictly follow normative principles of Bayesian learning (Knutson et al., 2011; Adam et al., 2020). Such "differential updating" (Bénabou & Tirole, 2016) comes in a number of forms: Not only do past experiences in financial risk taking influence the formation of expectations (Seru et al., 2010; Malmendier & Nagel, 2011; Knüpfer et al., 2017; Andersen et al., 2019; Malmendier et al., 2020), but there also appear to be effects of the context, which can lead to over- or underreaction to news (Enke et al., 2020) or joint evaluation with prior outcomes (Heinke et al., 2020). The literature mostly focuses on the investment position—whether the investment has so far made a gain or a loss—and on the favorability of the information—whether the information is positive or negative for the investors' portfolio value—as two context factors. First, although the direction of the effect is not unequivocal, the literature reports that, once a person has invested, belief updating differs for gain and loss positions (i.e., whether the investment's return so far is positive or negative; Kuhnen, 2015; Grosshans et al., 2020), even on a neural level (Knutson et al., 2011). Second, negative or "unfavorable" information generally seems to change expectations more strongly than favorable information (Kuhnen, 2015; Kuhnen et al., 2017; Kieren et al., 2022). Last, investors seem to incorporate information more strongly when they are invested in the asset of interest (Andersen et al., 2019; Gödker et al., 2019; Hartzmark et al., 2021). Overall, information processing seems to be guided by "motivated belief distortions" (i.e., the "utility" of holding a certain belief will play a role in deciding whether the belief is adopted; Brunnermeier & Parker, 2005; Fischer & Maier, 2019; Zimmermann, 2020). Allowing the influence of information to depend on investment position and information favorability is also sufficient to allow for beliefs to develop in ways that diminish the profitability of investments based on those beliefs.

The anomalies in belief formation described above can be unified by our proposed conceptual framework of context-sensitive belief updating. A natural way to formalize this framework is by extending a standard reinforcement learning model into a context-sensitive reinforcement learning (CSRL) model. Upon receiving new information, a reinforcement learner first calculates the difference between the expected and the received signal, also called the *prediction error*. This error is then weighted by a *learning rate* and added to the value of the next prediction, thereby nudging it in the direction of the new information (Sutton & Barto, 2018). The learning rate is usually considered to be a constant value between 0 and 1, and it adjusts how much influence new information has on the next prediction. However, to account for the observation that belief updating in investment decisions is context dependent, we allow for different learning rates. We assume that the learning rate varies depending on context, that is, whether information is favorable or unfavorable in regard to an investment and whether an investment has so far produced a gain or a loss. We first hypothesize that unfavorable information is generally updated more strongly (Kuhnen & Knutson, 2011; Kuhnen, 2015, ; Although there exists also some evidence to the contrary. See e.g. Bénabou & Tirole, 2016 for an overview). Next, we expect that people generally update their beliefs more strongly when they are invested (Hartzmark et al., 2021). Last, and most importantly, we predict an interaction between the investment position (i.e., gain or loss) and information favorability (i.e., favorable or unfavorable) that coincides with a belief in buying-price reversion (Asparouhova et al., 2009; Jiao, 2017). When an investor has made a gain with an investment, we expect unfavorable information to have a stronger impact on beliefs compared to favorable information. In contrast, when an investor has made a loss with an investment, we expect unfavorable information to have a weaker impact on beliefs compared to favorable information. In summary, our conceptual framework therefore differentiates between five contexts: Not being invested at all, and all combinations of investment position (gain or loss) and information favorability (favorable or unfavorable).

Portfolio investment decisions generally follow investors' beliefs, in trading decisions in the field (Giglio et al., 2021) as well as in laboratory experiments (Grosshans et al., 2020). We therefore expect that the hypothesized belief-updating pattern has direct consequences for investment decisions. Two opposing effects arise after having made a gain or a loss with an investment: stronger updating from unfavorable information after having made a gain increases the likelihood of liquidating an investment at a gain; stronger updating from favorable information after having made a loss decreases the likelihood of liquidating an investment at a loss. This investment pattern—selling gains too early and holding on to losses for too long—is well documented in the finance literature as the disposition effect (Shefrin & Statman, 1985). It is formally defined as the difference between the proportion of the realized (i.e., liquidated) gain positions (relative to all experienced gain positions) and the proportion of the realized loss positions (relative to all experienced loss positions). As positive returns can in fact signal further favorable price development, the disposition effect does lead investors to "leave money on the table" (Odean, 1998). The disposition effect has been found not only in laboratory experiments (e.g., Weber & Camerer, 1998; Fischer & Maier, 2019) but also across multiple asset classes, such as individual stock trading (e.g., Shefrin & Statman, 1985; Odean, 1998), mutual funds (Frazzini, 2006), and real estate markets (Genesove & Mayer, 2001).

A second consequence of updating expectations according to our conceptual framework might be that decisions to sell an asset are less predictive of future price developments than decisions to buy. The reason for this is that the interaction between investment position and information favorability can take place only when a person is invested in an asset. When not invested, there is neither a gain or loss position nor favorable or unfavorable news, and belief updating can therefore not be influenced by these contexts. In line with this prediction, Grosshans et al. (2020) observed in an investment experiment that buying (i.e., investing) decisions are more predictive for future price changes than selling decisions and thereby also more profitable. Note that the profitability of a selling decision can be measured by calculating the losses that were avoided and the profits that were forgone by liquidating the investment. A similar finding is reported by Akepanidtaworn et al. (2021), who analyzed the investment decisions of institutional investors. They found that while the buying decisions of these investors clearly outperform a random buying strategy, their liquidation decisions underperform, even when compared to selling assets at random. In sum, as with the disposition effect, investors who initially made a good purchase lose out on its potential by bad timing of their decision to close the investment.

We test the predictions of our unifying framework with a preregistered experiment building on a standard laboratory investment task (Frydman et al., 2014) which is often used to investigate the patterns described above. In this task participants invest in a stock whose value is subject to either a noisy upward or a noisy downward drift. In each round this drift can, however, also switch direction with a low probability. Note that it is common for this kind of task to restrict participants' options to either holding or not holding the asset (Frydman et al., 2014; Frydman & Rangel, 2014; Grosshans et al., 2020). As investors also experience negative price trends, however, we give our participants the option to short sell the asset. Short selling not only allows participants to react to the detection of a downward trend in the same way as for an upward trend but also allows us to disentangle the two reasons for not holding a positive number of shares: A participant could either not be certain enough that the price will increase or simply believe the price is falling. After each investment decision we also elicit participants' expectations about a price increase in the subsequent round. This elicitation allows us to determine the difference in expectations between rounds, thereby directly assessing the updating of participants' beliefs. The experimental setup further allows us to contrast the observed behavior with the actions and beliefs of a risk-neutral Bayesian investor. In a second phase of the study, we leverage the potential of our experimental setup and aim to demonstrate the causal effect of these belief-based mechanisms: By providing additional information about the likelihood of a price increase, we support participants in updating their beliefs, thereby demonstrating and mitigating the adverse effects of context-sensitive updating on investment decisions.

To gain a broader perspective and ensure the robustness of the results, we analyze the observed belief updating and its effect on investment behavior with two conceptually different approaches: First, we calculate the difference between the belief updates reported by participants and those of a Bayesian learner. This approach allows us to directly investigate the magnitudes of the belief updates and compare them to results of the normative Bayesian updating model. Next, we estimate the parameter values of the proposed context sensitive reinforcement learning model using participants' reported beliefs. This estimation allows us to translate the magnitudes of the belief updates into implied learning rates. We confirm the viability of this method using a parameter recovery exercise. Letting the learning rate vary by context, we aim to uncover the differences in updating strengths predicted by the conceptual framework. Thus, we predict that the estimated learning rates will differ depending on the investment position and the favorability of new information. This structural model approach also allows us to test the CSRL model against more parsimonious competitors. We find that the CSRL explains the data best and the model's complexity is justified by its explanatory power. Independent of the method, the common robust finding is an interaction effect between the investment position and the favorability of the new information on belief updating. This interaction effect on the updating strength is also reflected in participants' investment decisions. Compared to the rather low hypothetical disposition effect value of a risk-neutral Bayesian investor, the disposition effect value shown by participants is significantly higher. Further, we observe that when an investor closes an investment, the subsequent price movements would have been more profitable (i.e., more profit was forgone by closing the investment) when closing from a loss position compared to a gain. In the second part of the experiment we mitigate the effects of the context on belief updating by providing additional information to participants. This leads to a decrease in the disposition effect, approaching that of a risk-neutral Bayesian investor, and thereby leads to a larger profit from the investment decisions. We do find, however, that the information provided needs to be rather precise to have a significant impact.

To the best of our knowledge we are the first to investigate the combination of the effects of investment position, information favorability and their interaction on belief updating. The resulting framework provides a unifying explanation for frequently observed anomalies in belief formation and investment decisions. Nevertheless, there is a large body of research examining how past experience influences changes in expectations. People's expectations in turn have a large impact on their stock market participation and portfolio composition, which influence their overall wealth levels (Malmendier & Nagel, 2011; Knüpfer et al., 2017). There is strong evidence from the field (Choi et al., 2009; Bucher-Koenen & Ziegelmeyer, 2014; Hoffmann & Post, 2017; Guiso et al., 2018) as well as from the lab (Kuhnen & Knutson, 2011; Kuhnen et al., 2017; Hartzmark et al., 2021) showing that past experiences have a stronger effect on people's expectations when their own investments are involved and they therefore have "skin in the game." The seminal work by Kuhnen (2015) shows an interaction between the domain of the returns (positive or negative, akin to the investment position) and information favorability in a lottery investment task. In her experiment, participants experienced either always negative or always positive returns and had to learn whether returns stemmed from a good (high positive or low negative returns) or a bad (low positive or high negative returns) distribution. The main finding in Kuhnen (2015) is an overly strong pessimism from favorable information in the negative domain. We extend this strand of literature to financial decisions, were gain and loss positions arise endogenously.

We further add to the belief distortion literature, investigating why people deviate from the normative or Bayesian way of belief updating? The possible explanations can be subsumed by three approaches: First, people could be following a specific motivation when updating their beliefs, a mechanism called *motivated belief distortion* (e.g., Brunnermeier & Parker, 2005; Fischer & Maier, 2019; Zimmermann, 2020). Accordingly, people have preferences among possible beliefs, so that they are motivated to update their beliefs in such a way that favorable beliefs result (e.g., "It would be horrible to believe that there will be a price drop"). In other words, each possible state of belief has some utility of its own and the investor decides what to believe, taking that utility into account. Second, belief updating could be influenced by cognitive dissonance (Festinger, 1962). Here, an investor might feel a displeasurable dissonance between their self-image ("I'm a good investor making good investment decisions!") and the current expectation about the future value of the investment ("The investment's value is likely going to decrease"). Both of these elements can be changed to alleviate the dissonance, but it is easier to change the expectation rather than take a hit to one's positive self-image. A further loss is a bigger threat for ones self-image compared to losing part of an already made gain. For this reason we expect optimistic updating to mainly occur when in a loss position. In keeping with the concept of self-image, Chang et al. (2016) finds that investors are more likely to liquidate a losing investment if the initial investment decision was made by a fund manager rather than by themselves. In this case, the investor can protect their self-image by blaming the fund manager while still admitting that the investment should be stopped. Cognitive dissonance can also arise after having made a gain, through a desire to avoid regret (Loomes & Sugden, 1982; Bell, 1982). Here, the dissonance lies between the belief that the value will further increase and the desire to liquidate to avoid regret in case of a value decrease. As the "cost of being wrong" (i.e. experiencing a favorable price movement) is rather low, these beliefs are more readily distorted (Bénabou & Tirole, 2016). Last, Jiao (2017) provides evidence that investors foster a belief in buying-price reversion (i.e., mean reversion). According to this belief, the price of an asset will return to the price at which the investment was initially made and all deviations from that price are just temporary. This fixation on the initial buying price also serves as an explanation for why interventions such as hiding the buying price can decrease the disposition effect (Frydman & Rangel, 2014). It is implicitly assumed that this belief in buying-price reversion is held *ex ante*; that is, investors believe it to be an inherent trait of an asset. Note, however, that motivated belief distortion and cognitive dissonance may influence belief updating in such a way that a belief in buying-price reversion emerges *ex post*, and therefore only after an investment was made. In other words, an investor may start to believe that the price will fall (rise) again only after having made a gain (loss) with a held investment.

All of these three approaches can be captured by letting the learning rates of a reinforcement learning model depend on the investment position and the favorability of the new information. The learning model specifies how the probability of an event is learned, described as "probability learning" (Eldar et al., 2018; Sutton & Barto, 2018; Fontanesi et al., 2019). In this domain of probability learning, these models are also known as weighting and updating (Albrecht et al., 2021) or anchoring-and-adjusting models (e.g. Hogarth & Einhorn, 1992, allthough here the model also includes an additional subjective valuation of the new evidence). Regarding the impact of new information, Gershman (2015) reports stronger updating for negative prediction errors underlining the advantage of using multiple learning rates for different situations. Knutson et al. (2011) argue that different neural systems may be responsible for learning from gains and losses and show that the interindividual differences in learning rates for gains versus losses even correlate with real-life consequences such as asset holding and debt. Learning rates further do not have to be stable over different situations. Lee et al. (2020) show that learning rates can be adapted to best account for the "observation stochasticity" (i.e., noise) in a given environment. In general, reinforcement learning models align well with many effects found in the finance literature such as strong recency effects on beliefs and risk taking (Malmendier & Nagel, 2011, 2016; Dessaint & Matray, 2017). As it does not require any assumptions on the knowledge about the underlying distribution of outcomes and is very flexible (see, e.g., Olschewski et al., 2021), reinforcement learning can be appiedl to a plethora of environments with a small computational burden on the learner (Erev & Roth, 1998; Sutton & Barto, 2018). Reinforcement learning is also biologically well founded, as a prediction error signal can be found as measurable brain activity in humans as well as in other species (Rangel et al., 2008; Schultz, 2015). These signals are commonly associated with the reward-processing areas of the brain, which in turn have been shown to react to price changes in an investment task (Frydman & Camerer, 2016). We add to this literature by proposing and testing a novel reinforcement learning model for investment decisions in which the learning rates depend on the current context of the investment.

The remaining manuscript is structured as follows: Section 2.2 introduces our experimental design in further detail. Section 2.3 presents the results of the experiment. Section 2.4 follows by reporting the results of our structural model estimation. Last, all results are then discussed in Section 2.5, which will also provide our conclusion.

2.2 Experimental Design

To test the proposed conceptual framework, our task needs to fulfill two major requirements: First, the task should provide the opportunity to learn about the future returns of an asset. The amount of information that can be learned by participants should be clearly quantifiable, allowing for a clean analysis. Second, participants should be able to hold investments over multiple rounds, allowing them to create gain and loss positions by their own choices and reliably exhibit the expected anomalies in investment behavior. We implemented a version of the well-studied experimental paradigm used by Frydman et al. (2014) to meet these criteria. Additionally, we also elicited participants' beliefs in each round. In a second phase of the experiment we aimed to manipulate the strength of the context-sensitive belief updating. To do so, we provided participants with three different levels of information about the next price movement in the second phase. The more information is presented, the less participants have to rely on their ability to learn. This within-subject design of splitting the experiment into two phases enabled us to test the effect of our experimental manipulations on belief formation while holding individual abilities, preferences, and other interindividual differences constant.

Investment task: In detail, we gave participants the option to invest in an asset whose value was governed by a hidden Markov chain. If the chain was in the "good" state, there was a .65 probability of the price increasing; if the chain was in the "bad" state, the probability of a value increase was .35.[1] The magnitude of each price movement was drawn uniformly from the set $\{5, 10, 15\}$, rendering it uninformative. In each round, the state of the Markov chain remained the same with a probability of .8 or switched with a probability of .2. Participants were informed about this mechanism and that this implied that the price would experience a "drift." The direction of this drift could be learned from the direction of the observed price movements. Whereas investment decisions in the field are often based on information other than the price movements, this compact design allowed us to test the proposed mechanism without having to consider complicating outside sources of information. In each round, participants could change their investment without transaction costs and decided either to hold one share of the asset, not to hold any shares, or to short sell one share of the asset (i.e., a "short" investment). Thus, participants' decision whether to invest at all and the gains or losses they experienced where all endogenous. This feature of the experimental design is crucial for mechanisms based on motivated beliefs and cognitive dissonance. While a simplification of real world investment decisions, the design also comprises the necessary features to pursue our research questions and better reflects real investment situations than a standard lottery task.

Belief elicitation: After participants made their investment decisions, we elicited their beliefs about the probability of a price increase. To ensure accuarcy we implemented a incentivized belief elicitation. However, some of these tasks can be skewed by participants' risk preferences or require an explanation of the Bayesian probabilities. For this reason we used a lottery matching task which ensures incentive compatibility even under strong risk aversion (Trautmann & van de Kuilen, 2015). Participants were presented with a lottery with possible payouts of 0 or 10 experimental points and a winning probability of p_w. In each round, p_w was randomly drawn from a uniform distribution, but the result of this draw remained hidden. Participants were asked to indicate the minimum value of the winning probability p_w for which they preferred playing the lottery over a bet on the next price move (receiving 10 points for a price increase or 0 points otherwise). This probability was reported using a slider ranging from 0 to 100%. If $p_w > p_m$, the lottery was played; if $p_w < p_m$, the bet on a price increase was taken. This system was communicated to participants in the instructions. It ensured that participants could maximize their expected payoff only by reporting their true beliefs. This

[1] Note that these probabilities are symmetrical around .5, meaning that the probability of a price increase $P(up)$ in the other state was always $1 - P(up)$.

was also explained to participants in the instructions. Points won in this procedure were added to the final payoffs only at the end of the experiment and were weighted by a factor of .1.

Sequence of events in one round: Participants first saw their portfolio at current values and chose whether and how to invest. Next, they reported their beliefs about an upcoming price increase. After 6 and 5 s of no investment decision or belief report, respectively, a message reminded participants to please continue with the task. Finally, the price update was displayed, indicating the movement and the new price.[2] Note that no new information was added between participants' investment decisions and their belief reporting, rendering the order of these actions irrelevant.

Second phase: The goal of the second phase was an experimental manipulation to reduce the impact of belief updating effects on investment behavior. To rule out that observed changes in the behavior are driven by individual differences (e.g. preferences, cognitive abilities etc) we implemented this as a within participant comparison. The experimental procedure was identical to the first phase, with the exception of the further information that was provided to some participants, supporting them in the rational updating of their beliefs. In detail, we implemented three between-subjects conditions: Participants in the *full-information* condition received the true probability of a price increase in each round. This essentially revealed whether the asset was currently in the good or bad state, which is more precise information than what even a Bayesian learner could infer from the price signal. Short of revealing the actual next price movement, this therefore represented the "best possible" informational situation for the participants. In the *partial-information* condition, the participants received the probability of a price increase according to the Bayesian updating solution. This information would be redundant to a Bayesian learner and therefore should not change the belief of someone following the "normative" Bayesian updating process. Participants in this condition were informed that these probabilities were determined by incorporating all information about the price structure they had seen in the instructions as well as the past price movements. Note that both of these conditions left participants with the possibility of simply transcribing the probabilities shown in the trading interface to the belief report page. As the belief reports were incentivized to be as accurate

[2]Screenshots of the experimental setup and instructions (in German) can be found in Section A.5 of the internet appendix.

a prediction as possible, such "transcription" would constitute reward-maximizing behavior. The *no-information* condition served as a control group; here the task used in the first phase, where the participants received no further information, was repeated. The trading interface was kept identical in all three conditions and in both phases of the experiment.[3]

Portfolio and payoff: Each participant played two blocks of Phase 1 and two of Phase 2, with each block lasting 75 rounds. Each block for each participant was initialized completely independently: A new price path was generated, with the asset's price always starting at 1,000 points. To avoid inertia and systematic anchoring effects, the initial portfolio—holding one share, no shares, or having short sold one share—at the start of each block was also randomly determined with equal probability. The value of the starting portfolios was set to 2,500 points, and points that were not invested were displayed as cash. Note that participants could adjust their initial portfolio without transaction costs before the first price movement. We calculated the payoff for each block as the difference between the initial value of the portfolio (including cash) and its value after the final round of the block. This difference was then summed over all four blocks and 5% of this sum was added to a baseline payment of CHF 15. This procedure allowed participants to lose money over all four blocks while still ensuring that they received a minimum participation fee of CHF 10 in all cases.

Bayesian Benchmark Investor: Some of the numbers presented here are strongly dependent on the experimental task and its implementation, making them hard to interpret without reference. For this reason we will sometimes compare participants' performance to that of a risk-neutral Bayesian investor which serves as a benchmark. A Bayesian investor would infer the current state of the asset by optimally utilizing the information provided in the instructions as well as the observed price.[4] Risk neutrality implies that this investor would hold or short sell the asset whenever the probability of a good state (and thereby the probability of a price increase) is estimated to be above or below .5, respectively. Although we do not assume our participants to be risk neutral, this benchmark allows for an easily understandable comparison. Risk aversion would simply lead to fewer investments, as there would be a "range" of uncertainty around the .5 probability mark at which

[3]The probability indicator field was replaced with a dash in Phase 1 as well as in the no-information condition of Phase 2.

[4]We describe the way in which such an investor would draw inferences about the state of the assets in Appendix A.1.

a risk-averse investor would refrain from investing. The experimental setup allows us to calculate and compare the benchmark values for each individual price path in which participants could invest. Note that a Bayesian investor would show the same behavior in the no-information and partial-information conditions in the second phase. In contrast, in the full-information condition with the additional information, even a Bayesian investor would not need to do any belief updating but could just rely on the given information.

Control variables: To further investigate the individual differences between participants, we also asked them to indicate their age, gender, program of study, stock ownership, and trading experience. We further used a test with four progressive matrix tasks (incentivized by paying CHF 0.60 per correct answer; Civelli & Deck, 2018) as an approximation of intelligence and measured risk preferences with the 7-item risk questionnaire of the German Socio-Economic Panel (11-point Likert scale; Richter et al., 2017). Last, we checked participants' level of engagement in the study with a 7-point Likert scale and similarly asked them to self-report their level of ambiguity and loss aversion.

Procedures and demographics: We implemented the experiment using the oTree framework (Chen et al., 2016) and each session lasted approximately 1.5 hr. After signing the informed consent all participants read through the on-screen instructions. To ensure sufficient knowledge about the investment task and belief elicitation, we conducted a comprehension quiz and had participants go through 30 training rounds, which were all payoff irrelevant. After this, participants started the investment task. The elicitation of the other control variables followed, and the session concluded with the exit questionnaire. We recruited 192 participants (mean age = 22.92 years, $SD = 3.24$; 112 women) who all completed the study. A general overview of the sample demographics is reported in the appendix in Table A.2.1. Participants on average earned CHF 49.8 in our experiment, including the CHF 10 participation fee that was paid independent of task performance.

2.3 Empirical Findings

The objective of this section is to analyze the belief updates in the five different contexts discussed in section 2.1 and their effects on the investment decisions, before reporting the results of the structural model. We begin by analyzing the data of Phase 1, in which all participants received no additional information about the price developments.[5] The first analysis focuse on the investment decisions and works also as a sanity check that participants understood the task. We then continue by relating the magnitudes of the belief updates to the previously discussed investment contexts and decisions. To complete this investigation we then turn to the effect of the experimental intervention of phase 2. Note that for brevity we frequently use the term "belief" to refer to the participants' reported belief about the probability of a price increase in the next round.

Investment Positions

Participants' investment positions broadly follow those of a risk-neutral Bayesian investor. There are, however, some notable differences: Consider that the Markov chain governing the price development has an overall probability of .5 of being in a good state. This would lead the risk-neutral Bayesian investor to hold the asset on average in half of the rounds while short selling it in the other half. In contrast, participants held the asset on average in 49.5% of the rounds in Phase 1 while shorting it in 32.5%. Such a reluctance to short sell might be the result of the more complex nature of a trading strategy that includes short selling.

Next, a risk-neutral Bayesian investor (i.e. our benchmark investor) would have made an average of 57.5 trades over the course of the 150 trading rounds in Phase 1 of the experiment, therefore switching between portfolios relatively often.[6] Here participants behaved very similarly, trading on average in 58.9 of the 150 rounds, albeit with a much higher dispersion (a standard deviation of 7.6 for the benchmark investor and 26.3 for participants). Note, that this measure includes trades where the portfolio "jumps" over holding no shares from holding a share to short selling and vice versa. As a risk-neutral Bayesian investor would always hold or short, all of their trades would constitute such jumps. Participants, in contrast, jumped on average in only 24.8 rounds, with the rest of their trades either starting a new investment (starting from holding no shares) or liquidating an investment (moving to holding no shares). These trading frequencies mean that participants on average held the asset for 4.58 rounds while shorting it for 3.61 rounds. For comparison,

[5] The results reported here include data from all 192 participants who completed the study. Twenty-eight participants reported a level of engagement in the study of 3 or smaller (on a 7-point Likert scale). Excluding this group from our analyses in accordance with the preregistration does not change the qualitative conclusions. Robustness checks of the main analyses excluding those participants are, however, reported in Table A.5.6 and Table A.5.7 in Appendix A.5

[6] Note, that risk aversion would lower this number, as a risk averse agent would refrain from investing during insecure times.

the risk-neutral Bayesian investor would switch from holding to shorting or vice versa after an average of 2.64 and 2.6 rounds, respectively.

Last, we test whether the first observed price change and later the underlying state of the asset did in fact influence participants' investment decisions. This seems to be the case, as participant adjust their initial portfolio in line with the first signal, i.e. after observing the price going up (down) in the first period 75% (25%) participants held the asset, 10% (54%) were shorted and 5% (20%) remained not invested ($\chi^2 = 219.59, p < .001$ for the differences in the investment positions after observing up or downs). A similar pattern emerges when looking at the state of the asset: the ratio of holding over shorting the asset is significantly higher in the good (52.3% held, 31.6% shorted) compared to the bad (48.1% held, 34.2% shorted) state, $\chi^2 = 42.05, p < .001$. A second way to show the effect of the states can be seen in Figure A.2.1 in Appendix A.2, which reports an event analysis regarding participants' investments before and after the asset switches states from good to bad or vice versa. The event analysis reveals that participants did react quickly to state switches by adjusting their portfolios. In summary, participants therefore did make sensible decisions that were adapted to the circumstances of the environment, that is, whether the asset was currently in the good or bad state. These results indicate that participants did understand the task and generally acted in their best interest.

Belief Updating

We now turn to analyzing participants' belief updating and its dependence on investment position and information favorability. In general, the reasonable investment decisions reported above are paralleled by sensible beliefs: Participants on average reported higher beliefs whenever the Markov chain was in the good compared to the bad state (.513 and .479, respectively, one-sided t test, $p < .001$). Nonetheless, throughout Phase 1, belief reports vary widely between participants. A Bayesian updater's beliefs would fluctuate around the central value of .5 with a standard deviation of .031. In contrast, for participants' belief reports these standard deviations range from .029 to .49. The dispersion of the reported belief was independent of participants' engagement with the study (Spearman rank correlation test, $\rho = .066, p = .360$). The average fluctuation in the beliefs of participants (again measured as the standard deviation of the reported beliefs) being higher than that of the Bayesian benchmark, with an average value of .22, is also an indication for participants' general tendency to overupdate.[7] This tendency can also be seen when comparing the magnitude of individual updates to those suggested by a Bayesian learner: When participants updated in the same direction as a Bayesian agent, 86.2% of their updates were larger, and only 13.8% were smaller.

To investigate the formation of beliefs, we calculate belief updates as the difference between a belief report at time t and the report in the previous round $t-1$. To compare the magnitudes of these updates between contexts, regardless of whether expectations increased or decreased, one would be tempted to calculate the absolute values of these updates. However, the sign of the absolute value will be the same, no matter whether expectations had *increased* or *decreased*. This problem becomes even more severe considering belief updates that go against the signal ("inverse updates"), for example, *lowering* the expectation of a price increase by after seeing a price *increase*.[8] Such inverse belief updates would clearly deviate from the rational benchmark, yet taking the absolute value would treat them the same way as though they had been in line with the benchmark. To preserve the information of the sign of these inverse updates, we use a different procedure: We flip the sign of any update that was made after seeing a price decrease. This way, belief increases (decreases) after seeing a price decrease (increase, i.e., inverse updates) will always end up as negative values, preserving the information that would be lost when taking the absolute value. Next, to ensure that we are measuring participants' updating "beyond" that of a Bayesian learner, we analyze the difference between participants' reported belief updates, Δ_{Report}, and that of a hypothetical Bayesian updater, Δ_{Bayes}, by calculating $\Delta_{\text{Report}} - \Delta_{\text{Bayes}}$. This is important, as a similar, but weaker, pattern in belief updating to the hypothesized one can also be produced by a Bayesian updater. More precisely, after an investor has made a gain (loss) from a certain price movement, a second movement in the same direction will also be favorable (unfavorable) but have a weaker impact on beliefs than the first movement. Note that taking the absolute value at this stage without flipping the signs of updates from price decreases would lead to even more information being lost, as over- and underupdates relative to the Bayesian updates would both be counted as positive values.

[7] See Figure A.5.6 in the internet appendix for a histogram of the standard deviations of participants' belief reports.

[8] For a closer investigation on these inverse updates in our participants see A.5 in the internet appendix. Note however that the median inverse update had a small magnitude of four percentage points and therefore most likely constitutes reporting errors. in- or exclusion of these updates does not change our qualitative results, as shown in A.5, also in the internet appendix.

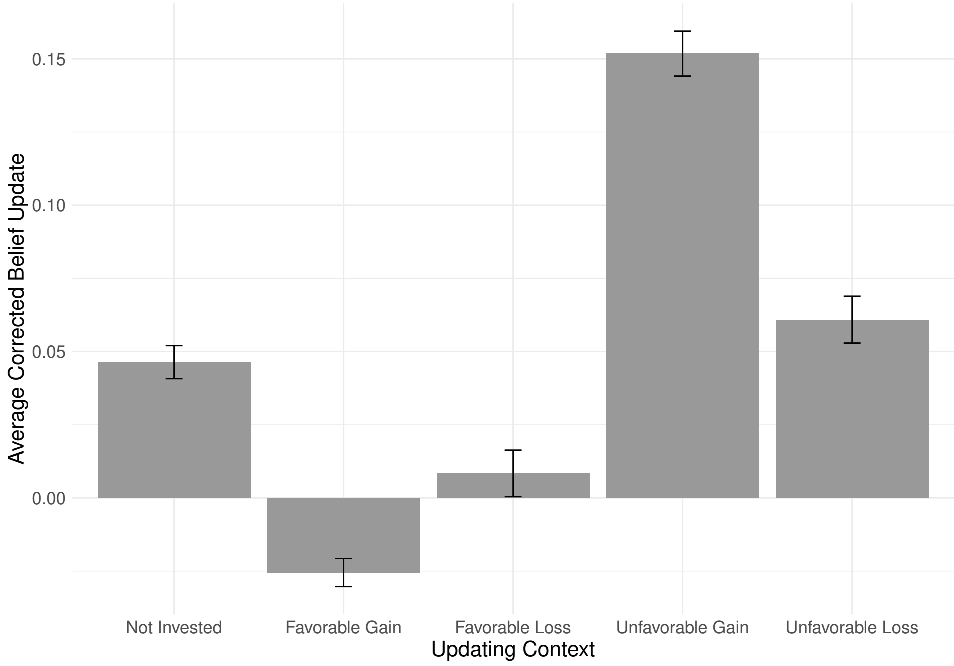

Figure 2.1: Belief updating in Phase 1 dependent on the current position (gain, loss, or not invested) and the last price move (favorable or unfavorable). Error bars show group-level 90% confidence intervals. For the same figure when excluding rounds with inverse updates, see Figure A.5.10 in Appendix A.5. The x axis shows the context in which the update was made (i.e., from not being invested, from a favorable information while in a gain position etc.). The y axis shows the average difference between the belief updates reported by participants and those of a Bayesian updater ($\Delta_{\text{Report}} - \Delta_{\text{Bayes}}$). Further, the sign of updates from price decreases was flipped, such that the numbers represent the "magnitude" of the updates.

Figure 2.1 displays these corrected belief reports for each context defined in our hypothesis. Generally, participants seem to overupdate by a value of .05 compared to a Bayesian updater (SD = .26, two-sided t test, $p < .001$). Contrary to our expectations, we find no difference in the average updating magnitude between rounds in which participants were invested and rounds in which they were not (two-sided t test, $p = .755$). When focusing on rounds in which participants were invested, favorable information seems to be on average under- or even inversely updated ($-.01$, SD = .22). This stands in stark contrast to unfavorable information, which is updated with an average magnitude of .11 (SD = .28). Taken together, unfavorable information leads to significantly stronger updates compared to favorable information (two-sided t test, $p < .001$). This result aligns with Kuhnen (2015), who also reports stronger updating for unfavorable information. We also find update magnitudes on average to be greater in gain compared to loss positions (.05, SD = .26, and .03, SD = .26, respectively, two-sided t test, $p = .026$), driven by the strong overupdating of unfavorable information in a gain position.

To test for the significance of these findings, we use a regression analysis with cluster robust standard errors (clustered on the participant level), including all belief updates where participants were invested (Table 2.1). In line with the pattern visible in Figure 2.1, we find a significant interaction between the investment position and information favorability. Belief updates from favorable information are generally stronger when participants find themselves in a loss compared to a gain position (.008, SD = .26, and −.026, SD = .2, respectively, two-sided t test, p = .007). The opposite is true for unfavorable information, where updates are stronger in a gain position, thereby driving the interaction (.15, SD = .29 in gain and .06, SD = .27 in loss positions, two-sided t test, $p < .001$).[9] Table 2.1 also confirms the two main effects of investment position and information favorability, although the magnitudes of the estimated parameters are clearly overshadowed by the interaction term. In sum, the hypothesized contextual sensitivity of belief updates does in fact emerge in the reported beliefs, leading to a significant interaction effect between position and information favorability on updating magnitude.

Investment Behavior

To influence an investment's profitability, context-dependent beliefs first have to impact investment decisions. Hence, we first check whether participants do trade in accordance with their reported beliefs. Table 2.2 reports the results of an ordered logistic regression on investments after the trading decision in round t (shorted, not invested, held, coded as −1, 0, and 1 respectively). Model 1 includes as an independent variables only the reported beliefs in a price increase.[10] This model shows that the reported beliefs are highly predictive for investments with an effect size of .05 ($p < .001$).[11] Overall, the significant parameter estimate for the reported beliefs in Model 1 indicates that participants did overall base their investments on their beliefs.

[9]Table A.5.6 and Figure A.5.10 in Appendix A.5 in the internet appendix contain the same analysis when excluding rounds with inverse updates. Table A.2.2 in Appendix A.2 confirms that the general pattern still holds when analyzing only held or shorted investments.

[10]Recall the sequence of events in a round: After the investment decision, participants were asked for their belief about the probability of a price increase at the transition from round t to $t+1$. Only at the end of round t did participants observe the new price that they could trade upon in round $t+1$. Recall also that there are no transaction costs in our task. Including transaction costs could have lead participants to stick to a portfolio they no longer prefered.

[11]For a brief example of how to interpret these effect sizes, see Appendix A.5 in the internet appendix. Appendix A.5 in the internet appendix further shows a visualization of the model's predicted investment probabilities at different levels of reported belief.

Table 2.1: Bayes Corrected Belief Updates

Favorability	$-.05\,(.01, p < .01)^{***}$
Position	$.09\,(.014, p < .01)^{***}$
Favorability × Position	$-.124\,(.018, p < .01)^{***}$
Constant	$.025\,(.051, p = .63)$
Control	Yes
Obs. (Participants)	14,558 (191)
Adjusted R^2	.072

Note. $^{***}p < .01$; Ordinary least squares regression clustered on participant level; standard errors and p-values are reported in parentheses. The table also reports the number of observations (Obs.) and participants as well as the adjusted R^2. Only investments lasting longer than one round are considered, which leads to the exclusion of one participant from this analysis. *Dependent Variable:* Belief updates while participants were invested, normalized by the change in Bayesian probability. The beliefs were reported on a scale of 0 to 100 and the updating values are calculated as $(q_t - q_{t-1})/100$, where q_t is the belief q in round t. Further, to obtain the magnitude of the updates rather than their direction, the sign of updates after a price decrease was flipped. *Independent Variables:* All variables are dummy coded. "Favorability" indicates whether the change in price of the asset was favorable (coded as 1) or unfavorable (coded as 0) to the participant's investment. A price decrease would, for example, be favorable if the participant had short sold the asset. "Position" indicates whether the current price of the held (shorted) asset is above (below) the initial buying price. A gain position was coded as 1, losses as 0. *Control Variables:* Age and gender are included as control variables but are omitted from the table.

As noted before, a similar pattern to our hypothesis would arise from the updates of a Bayesian learner. We therefore test how much additional variance is explained by the reported beliefs, beyond what a Bayesian learner could already explain. To this end, Model 2 in Table 2.2 separates the reported belief into a Bayesian and a residual component. We first regress the reported beliefs on the Bayesian probability using an ordinary least squares (OLS) regression. The residual values of this regression can be interpreted as the overupdating orthogonal (i.e., independent) to the Bayesian probability. These two signals, the residual values of the OLS regression and the Bayesian probability, are then used as predictors in the ordered logistic regression. The results of Model 2 in Table 2.2 at first confirm that the Bayesian probability has a strong and significant positive effect (.24) on the investment decisions ($p < .001$). While the effect of the residuals on investments is .04 and thus smaller than that of the Bayesian probability, it does also have a significant influence ($p < .001$). This indicates that the belief reports are an informative signal for the investment decision in addition to the Bayesian probability. In other words: Although assuming that investment decisions are based on Bayesian beliefs provides a reasonable predictor for participants' portfolios, their

Table 2.2: Portfolio Allocation After Trades in Round t

	(1)	(2)
Reported belief	$0.05\,(< .01, p < .01)^{***}$	-
Residual beliefs	-	$.04\,(< .001, p < .01)^{***}$
Bayesian probability	-	$.22\,(.004, p < .01)^{***}$
Intercept not invested	$2.28\,(.39, p < .01)^{***}$	$10.92\,(.44, p < .01)^{***}$
Intercept short	$1.16\,(.39, p < .01)^{***}$	$9.76\,(.44, p < .01)^{***}$
Control	Yes	Yes
Obs. (Participants)	28,799 (192)	28,799 (192)
AIC	46,339.93	45,673.05

Note. $^{***}p < .01$; Ordered logistic regressions with random effects for participants; standard errors are reported in parentheses. The table further displays the number of observations (Obs.), participants, and the Akaike Information Criterion (AIC) of each model. Model 2 aims to check whether the reported beliefs in round t predict the next investment above the predictive value of the Bayesian probability in round t. For a similar analysis per trade (i.e., short selling, liquidating, etc.) see Table A.2.3 in Appendix A.2. *Dependent Variable:* Participants' portfolio allocation (i.e., short selling, not investing, or holding a share) after trades in round t. Note that this is identical to the portfolio allocation at the start of round $t+1$. *Independent Variables:* "Reported belief" is the beliefs about an upcoming price increase reported by participants (on a scale of 0 to 100) in round t. "Residual beliefs" are the residual values of an ordinary least squares linear model regressing the reported beliefs on the Bayesian probability. "Bayesian probability" is the probability of a price increase as calculated by a Bayesian learner. *Control Variables:* Age and gender as well as the round number and an inertia dummy (i.e., the last portfolio allocation) are included as control variables but are omitted from the table. Round number had a significantly negative, and inertia a significant positive effect (both $p < .01$).

true beliefs differ enough from the Bayesian such that knowing their beliefs significantly improves predictions about their portfolio choices. This result also hints at a non-Bayesian mechanism, which we aim to capture with our structural model of context-dependent belief formation.

Our hypothesized context-dependent belief formation can take place only when a decision maker is invested. It therefore should mainly affect decisions to liquidate an investment. To get a more detailed picture of this effect, we calculate the average belief whenever participants decided to change their portfolio.[12] Interestingly, short sales were on average liquidated while the reported beliefs were still significantly below .5 (.46, $SD = .2$, two-sided t test, $p < .001$). At this probability of a price increase, a short sale would still be profitable on average. Participants liquidating

[12] A complete table of these average beliefs grouped by decision (opening or closing an investment) and position (gain or loss) as well as the hypothetical averages of a Bayesian updater can be found in Table A.2.4 in Appendix A.2.

short sales at this belief therefore indicates that preferences in the option valuation process, for example, risk aversion, might also affect the investment decisions.[13] In contrast, when deciding to sell a share that they hold, participants reported an average belief of .47 (SD = .21). These decisions were therefore in line with their beliefs, which were also significantly below .5 (two-sided t test, $p < .001$). A more detailed picture arises when differentiating liquidations further into those from a gain and those from a loss position. Further, we can compare the average reported beliefs of participants to the average Bayesian probability. This allows us to see where (in terms of investment position and information favorability) deviations from the Bayesian benchmark are especially influential for investment decisions.

We first look at liquidations from losses. Here, the pattern reported before emerges: When closing any investment (holding or short selling) from a loss position, participants on average thought a price decrease to be more likely than a price increase. A risk-neutral Bayesian investor would however "agree" with their decision: The average belief of a Bayesian learner when closing a short sale in the loss position would have been .51 (compared to .49 reported by participants); the average belief of a Bayesian learner when selling a held share in the loss position would have been .48 (compared to .46 reported by participants). Turning to liquidations from the gain position, participants still believed a price decrease to be more likely in general. Here, however, the risk-neutral Bayesian investor would not agree with participants' decisions: The average Bayesian probability when closing a short sale in the gain position was .48, indicating that the short sale could still be profitable (compared to .45 reported by participants); the average Bayesian probability when selling a held share in the gain position was .52, indicating that holding on to the share could still be profitable (compared to .48 reported by participants). In general, this analysis therefore confirms our previous findings, in that deviations from the benchmark belief are strongest when in the gain position.

Such "early exits"—that is, selling before the Bayesian probability would suggest doing so—from gain positions directly affect the overall profitability of participants' investment decisions. This can be illustrated by calculating the "success rate" of participants' investment decisions in the following way: Define a decision to make an investment as "successful" whenever the following price movement is favorable, and a decision to liquidate an investment whenever the following price movement would have been unfavorable. Take as an example an investor who has

[13] We explore this option by analyzing the individual differences between participants in Appendix A.2 and find none of our additional variables to be sufficient to explain our findings.

just liquidated a short sale: If the price now increases, then the liquidation was a successful action, as the unfavorable (when short sold) price movement was avoided. Note that this analysis also quantifies the information contained in the observed action (i.e., how well the observed action predicts the subsequent price movement). Decisions by participants to open an investment had a rather low success rate (i.e., correctly predicting the next price movement) of 50.3%. Note that this number excludes jumps (from short selling to holding and vice versa), as participants were already invested when deciding to jump. According to our conceptual framework, the contextual factors—investment position and information favorability—will influence belief updating, and thereby decisions, only while invested (i.e., for jumps or decisions to liquidate). This would imply that these types of decisions would be less profitable than those made when not invested. We do not find this to be the case, as the average success rate of decisions to jump or liquidate was 50.6%. Taking a closer look at these decisions to jump or liquidate, however, reveals a familiar pattern: Decisions that were made from a loss position were in general far more successful, with an average success rate of 52.1% as compared to 48.5% for decisions from a gain position ($\chi^2 = 13.77$, $p < .01$). This underperformance of decisions from gain positions can be explained by an overreaction to unfavorable information, as depicted in Figure 2.1. Such an overreaction in the gain position will lead to liquidation of the investment when it would have been better to keep it.

This eagerness to liquidate investments in the gain position, as well as a reluctance to liquidate from a loss position, is also what we conjectured to drive the disposition effect. The disposition effect value is calculated by subtracting the fraction of liquidated losses from the fraction of liquidated gains. Formally, this means we calculate $DE = \frac{\#RG}{\#PG} - \frac{\#RL}{\#PL}$. Here, G and L stand for gains and losses and R and P indicate whether the position has been realized (i.e., liquidated) or is a paper gain/paper loss (i.e., unrealized). This measure can therefore hypothetically vary between -1 and 1. Note also that for this analysis, a jump is also counted as a liquidation. Participants liquidated on average 48.97% ($SD = 30.9$) of their loss positions and 32.27% ($SD = 21.9$) of their gain positions. For comparison, a risk-neutral Bayesian investor would liquidate on average 99% ($SD = 0.1$) of their loss positions and 16.3% ($SD = 2.74$) of their gain positions.[14] Relative to

[14] A risk-neutral Bayesian investor will invest as soon as the probability of a favorable price movement is above .5. If the price movement is, however, unfavorable, this almost always pushes the probability of a subsequent favorable price movement below .5, leading the investor to liquidate the investment in 99% of the cases. If the investor has made a gain, however, this is a further

this benchmark, participants were therefore reluctant to liquidate losses and were realizing gains too early.[15] This consequently leads to a significantly higher disposition effect value of participants ($-.17$, $SD = .39$) relative to that of the Bayesian benchmark investor ($-.83$, $SD = .02$, two-sided t test, $p < .001$), meaning that participants did in fact hold losses for too long and sell gains too early.

Last, all these effects are also reflected in participants' payoffs being lower than that of a risk-neutral Bayesian investor in Phase 1. When we look only at the payoff from investment decisions (i.e., ignoring the incentivized belief reports and fixed payoff sums), participants on average gained only CHF 0.71 ($SD = 3.63$) per block of the investment task. This stands in contrast to a risk-neutral Bayesian investor who would have gained CHF 2.23 ($SD = 3.94$) per investment block. One possible explanation for this difference in payoff would be risk aversion in our participants. Risk aversion would, however, lower the payoff by preventing participants from investing in the first place. While our participants were not always invested, they were so in a majority of the rounds. The low success rate of participants' investment decisions demonstrated above thus appears as a more plausible explanation compared to pure risk aversion.

Summarizing, we do find evidence for context-sensitive belief updating and its impact on investment decisions, as hypothesized by our conceptual framework. Interestingly, we find the hypothesized diminishing effect on the profitability of decisions to liquidate an investment only when these decisions were made from the gain position. Although such a lower profitability of decisions in a gain position is only half of the components of the disposition effect, we do find a significantly stronger disposition effect in our participants compared to the benchmark investor.

Phase 2

The results reported so far underline the context sensitivity of belief updating and its impact on investment decisions. In Phase 2 of the experiment, we aim for a more causal test of our proposed mechanism by manipulating the necessity of belief updating, while holding individual differences among participants constant with our within subject design. After the two blocks of Phase 1, participants

signal that the investment was a good one. Thus multiple unfavorable price movements may be necessary to push the probability of a favorable price movement below .5, leading the investor to liquidate the investment. This was the case in 16.3% of rounds for the risk-neutral Bayesian investor.

[15] Figure A.5.8 and Figure A.5.7 in the internet appendix visualize this result.

continued the experiment with two additional blocks in either the no-information (i.e., baseline or control condition), the partial-information, or the full-information condition.

Information Provision and Belief Updating

Participants in the two treatment conditions were provided with additional information, namely, the Bayesian probability (*partial information*) or the true probability (*full information*) of a price increase in the next round. We first test whether participants did incorporate this additional information into their beliefs in the partial-information condition. To do so we calculate the difference between the reported beliefs and the Bayesian probability. Had participants reported the same probability estimates as the displayed Bayesian probability in the partial-information condition, these difference values should be zero. Further, if participants even partially incorporated the displayed information they may still vary around zero, but with a smaller standard deviation than in the no-information condition. Indeed, we do find a significantly smaller variance of these "belief errors" (Levene's test, $p < .001$) in the partial-information condition, indicating that participants reported beliefs closer to the Bayesian probabilities. In the full-information condition we find that most participants used the shown probabilities as orientations for their reported beliefs. [16] However, despite our belief elicitation being incentivized, and simply reporting the shown probabilities (i.e., .65 in the good and .35 in the bad state) clearly yielding the highest payoff, not all participants did so. This shows that although participants did in fact incorporate the information into their expectations once it was provided, they did not adopt it completely.

We next check how this added information impacted the interaction effect between investment position and information favorability on belief updating. We performed an OLS regression analogous to that shown in Table 2.1 with the reported beliefs from Phase 2. Model 1 in Table 2.3 shows that the interaction between position and favorability also reproduces in a smaller group ($p = .009$). It further confirms that the effect does not disappear through experience after the 150 rounds of Phase 1. Model 2 calculates the same regression for the partial-information condition. Whereas the favorability of the information maintains its effect on belief

[16] Figure A.3.1 in Appendix A.3 shows a visual representation of this result. Further, Figure A.3.2 in the same appendix shows the average belief updating split by information favorability and investment position.

formation, investment position and the interaction term do lose their significance. Note that the estimated effect of the interaction nearly halves from $-.081$ to $-.047$ between Models 1 and 2. A more direct comparison is provided in Model 3: It pools the data used in Models 1 and 2 and adds a dummy variable for the condition. Whereas the interaction term has lost its significance between Models 1 and 2, in Model 3 this difference does not show up as a three-way interaction between investment position, information favorability, and condition.[17] Further, neither the condition dummy itself nor any interactions with it are significant. This indicates that the effects we report from Phase 1 persist on a lower level even when providing the Bayesian probability.

Last, Model 4 includes only participants in the full-information condition. Note that following the displayed probabilities, we expect only updating magnitudes of either 0 or .3 (in case of the state remaining the same or switching, respectively). This is a bimodal distribution, which violates the assumption of normally distributed errors of the OLS regression. To mitigate this problem we add a variable indicating whether the state of the chain has switched. As in all other models, favorability of the information seems to remain a significant factor in the way participants update their beliefs. However, the interaction between investment position and information favorability also does not reach significance in this model.

Overall, the provision of additional information in the full-information condition seems to have mitigated the interaction effect between information favorability and investment position. Similarly, this interaction term is no longer significant in the partial-information condition, although this mitigation is not strong enough to show up as a three-way interaction with a condition dummy. The only effect remaining throughout all conditions is that of information favorability. Here, participants consistently update their beliefs less strongly when witnessing favorable as compared to unfavorable information, although the strength of the effect declines with the amount of information provided.

[17] A visual representation of these results can be found in Figure A.3.3 in Appendix A.3.

Table 2.3: Bayes Corrected Belief Updates

	(1) *No*	(2) *Partial*	(3) *No & Partial*	(4) *Full*
Favorability	$-.073\,(.02, p < .01)^{***}$	$-.049\,(.014, p < .01)^{***}$	$-.072\,(.02, p < .01)^{***}$	$-.038\,(0.016, p = .02)^{**}$
Position	$.047\,(.023, p = .04)^{**}$	$.016\,(.025, p = .53)$	$.05\,(.02, p = .04)^{**}$	$.024\,(.017, p = .15)$
Condition			$-.001\,(.021, p = .99)$	
State switch				$.037\,(.012, p < .01)^{***}$
Position × Favorability	$-.081\,(.031, p < .01)^{***}$	$-.047\,(.027, p = .08)^{*}$	$-.081\,(.031, p < .01)^{***}$	$-.02\,(.018, p = .32)$
Position × Condition			$-.036\,(.035, p = .3)$	
Favorability × Condition			$.024\,(.024, p = .32)$	
Position × Favorability × Condition			$.033\,(.041, p = .42)$	
Constant	$.025\,(.109, p = .81)$	$.085\,(.072, p = .24)$	$.06\,(.06, p = .32)$	$.056\,(.041, p = .17)$
Control	Yes	Yes	Yes	Yes
Obs. (Participants)	4,761 (64)	5,067 (64)	9,838 (128)	5,683 (64)
Adjusted R^2	.03	.06	.05	.02

Note. $^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$; ordinary least squares regression clustered on participant level; standard errors are reported in parentheses. Model 1 includes only participants from the no-information condition, Model 2 includes only participants from the partial-information condition, Model 3 pools participants from the no-information and the partial-information conditions, and Model 4 includes only participants from the full-information condition. The table also reports the number of observations (Obs.) and participants included as well as the adjusted R^2 for each model. *Dependent Variable:* Belief update normalized by the change in Bayesian probability. The beliefs were reported on a scale of 0 to 100 and the updating values are calculated as $(q_t - q_{t-1})/100$, where q_t is the belief q in round t. Further, to obtain the magnitude of the updates rather than their direction, the sign of updates after a price decrease was flipped. *Independent Variables:* All variables are dummy coded. "Favorability" indicates whether the change in price of the asset was favorable (coded as 1) or unfavorable (coded as 0) to the participant's investment. A price decrease would, for example, be favorable if the participant had short sold the asset. "Position" indicates whether the current price of the held (shorted) asset is above (below) the initial buying price. A gain position was coded as 1, losses as 0. "Condition" indicates whether the participant was in the partial-information condition in Model 3. "State Switch" is coded as either -1 or 1 for a state switch and 0 for no state switch being indicated in the full-information condition. *Control Variables:* Age and gender are included as control variables but are omitted from the table.

Information Provision and Investment Behavior

We continue the analysis by inspecting whether the improved belief updating between conditions do also translate to better investment decisions, and thereby increases participants' profits. Recall that the success rate of participants' investment decisions in Phase 1 only differed substantially between decisions made from gain and loss positions. We find the same pattern for success rates in the no-information condition. Although jumps exhibited a success rate of 49.1% from gain positions and 50.5% from loss positions, a larger difference emerges for decisions to close an investment: Here the success rates are 48.9% and 54.2% for gain and loss positions, respectively. We find similar results for participants in the partial-information condition (see Table A.5.4 in Appendix A.5 of the internet appendix for a tabulated overview of these results). Comparing the full-information condition to the no-information condition, we see a substantial improvement especially after jumps: Jumps from a gain position were successful in 56.7% of cases, compared to 49.1% in the no-information condition; jumps from a loss position were successful in 55.8% of cases, compared to 50.5% in the no-information condition. Thus, only when provided with full information on the state of the price-generating process in the full-information condition could participants improve the profitability of their decisions.[18]

The next investment pattern implied by context-sensitive belief updating was an increased rate of selling at a gain position compared to a loss position, that is, the disposition effect. Given the increased availability of information in the two treatment conditions, we expect an improvement in the disposition effect values toward their respective benchmarks. Note that in the full-information condition, the expected disposition effect value of a risk-neutral investor is 0. This is because the only reason to sell in this condition is a state switch. As state switches are independent of the investor's losses or gains, this leads to the same propensity to sell gains as losses, and thereby to a disposition effect of 0.

To measure the effect of the provided information on the disposition effect within each participant, we calculate the difference between the disposition effect values in Phases 1 and 2.[19] The values used here are the difference between the

[18]See Table A.5.5 in the internet appendix for a tabulated overview of participants' reported beliefs when making these decisions, as reported in Table A.2.4 for Phase 1 of the experiment.

[19]This can be done because the disposition effect values did not differ significantly between groups during Phase 1. See Table A.2.5 in Appendix A.2.

Table 2.4: Change in DE Between Experimental Phases

Constant (no information)	$-0.13\,(0.16, p = .43)$
Partial information	$-0.04\,(0.05, p = .51)$
Full information	$-0.67\,(0.05, p < .001)^{***}$
Control	Yes
Observations	192
Adjusted R^2	.51

Note. $^{***}p < .01$; Linear regression on the difference in the disposition effect (DE) between Phases 1 and 2. All participants received the same treatment (no information) during Phase 1, whereas the different groups received increasing amounts of information about future price moves in Phase 2. *Dependent Variable:* The disposition effect measure calculated as the difference between the propensity of selling a gain and that of selling a loss. This model uses the difference between the disposition effect values of the participants and that of a risk-neutral Bayesian investor on the same price path (i.e., the "corrected disposition effect values"). The model is calculated on the difference between these corrected disposition effect values in the first and the second phase of the experiment. The same model using the uncorrected disposition effect values can be found in the internet appendix in Table A.5.7. *Independent Variables:* Dummy coded variables indicating the condition in Phase 2. *Control Variables:* Age and gender are included as control variables but are omitted from the table.

participants' disposition effect and that of a risk-neutral Bayesian investor. As expected, participants in the no-information condition did not improve significantly between phases (one-sided t test, $p = .178$), meaning they did not come closer to the benchmark disposition effect. This confirms that experience alone was not enough to alleviate the disposition effect and that improvements in the other conditions are due to the provided information. Indeed, participants in the partial-information condition did improve significantly (one-sided t test, $p = .038$). However, as Table 2.4 reports, the improvement does not significantly surpass that of the slight improvement observed in the no-information condition due to experience. Last, a different picture emerges for the disposition effect values of the full-information condition. Here, values did improve significantly toward the rational benchmark, both in a separate model ($p < .001$) and in the full regression as described in Table 2.4. Overall, these findings regarding the disposition effect confirm the general pattern found in Phase 2: The observed context-sensitive belief formation and its effects on investment decisions seem to withstand the provision of the Bayesian probability. Only very clear information, such as the true probability of a price increase provided in the full-information condition, allows participants to overcome the influence of these effects.

All the effects described above culminate in participants' payoffs. While those in the no-information condition earned a nonsignificant average of CHF 0.2 less in the second compared to the first phase, participants in the partial-information condition did in fact improve their earnings. Yet, although participants in the partial-information condition earned on average CHF 1.78 more in the second compared to the first phase, this difference was not statistically significant (one-sided t test, $p = .121$). Participants in the full-information condition, however, improved their earnings between the two phases by an average of CHF 13.32. This improvement is both significantly greater compared to zero and to the improvement in the no-information condition (both one-sided t tests, $p < .001$).

In summary, we observe that providing additional information improves not only belief updating but also investment decisions and thereby overall profits. Participants who did not receive any further information did not change their belief updating or their behavior between experimental phases. Interestingly, providing the Bayesian probability of a further price increase in the partial-information condition led to weaker results than expected. Although it did seem to impact the way participants updated their beliefs, its effects in regard to investment behavior did not surpass that of mere experience effects found in the no-information condition. Only when we provided very clear information about the future price movements, such as in the full-information condition, did participants strongly change their expectations as well as their investment behavior.

2.4 Context-Sensitive Reinforcement Learning Model

We suggest capturing the context sensitivity of belief updating with an augmented reinforcement learning model. The previous analysis focused on the difference in updating magnitudes, and thereby the result of the learning process. However the implied non-linear dynamics in the observed belief formation can be better captured by a structural modelling approach. Beyond that, one can test such a structural model more rigorously against simpler versions, and thus investigate whether the model complexity is justified by the data. First, we translate the proposed conceptual framework into a reinforcement learning model with context-sensitive learning rates. A test on model parameter recovery suggests that such a model can be identified with our type of data and analysis approach. In such a parameter recovery exercise one simulates data (in our case reported beliefs) for the experimental task using known model parameters and then recovers them by fitting the model on the simulated data (for details of this step see Appendix A.5). Note that successful recovery of the parameters of the full model in our case implies that the model itself can be distinguished. As our competitor models are nested models of the full model, they can not mimic data from a more complex one. Next, we estimate the learning rates of the model using participants' belief reports. Lastly, we investigate the explanatory power of this CSRL model compared to reduced versions including a standard reinforcement learning model with a single fixed learning rate.

We model a reinforcement learner interested in the probability, $q_t \in [0,1]$, of observing a price increase at the end of each round, t. Price increases are coded as $x_t = 1$, and price decreases as $x_t = 0$.[20] Equation 2.1 describes such a learner, who at the start of a round holds a prior belief, q_{t-1}, and incorporates new information, x_{t-1}. This is done by adding to the prior belief a weighted difference between the new information and the prior belief (*prediction error*; Sutton & Barto, 2018). Notice that the updated belief in round t becomes the prior belief used to update the belief in the next round, $t+1$. Whereas classic reinforcement learning models assume a constant weight by which the prediction error is integrated into the new value, we allow the learning rate $\eta_c \in [0,1]$ to vary for different contexts $c \in \{i, g, l, f, u\}$, that is, not invested, gain, or loss position and favorable or unfavorable information.

$$q_t = q_{t-1} + \eta_c(x_{t-1} - q_{t-1}) \tag{2.1}$$

We estimate the learning rates of the model described in Equation 2.1 using a hierarchical Bayesian model estimation technique. The intuition behind such a model is that although the individual parameters may vary between participants, they are not completely independent but form a group-level distribution themselves. In practical estimation terms, this means that on the group level, we estimate for each of the five contexts c—that is, not being invested and the combinations of investment position and information favorability—the mean and standard deviation, μ_c^η and σ_c^η, of the learning rates.[21] Each participant is then assumed to possess their own set of learning rates, η_c, which are drawn from the distributions defined by the

[20] Note, that adding a subjective transformation of the evidence as in Hogarth & Einhorn (1992) could, depending on the transformation, yield similar results as varying learning rates. As this transformation would have to depend on the context as well, while still maintaining a free learning rate parameter in the model, varying learning rates directly is a more parsimonious approach.

[21] Note that the superscript η here only serves the purpose of making clear that these parameters relate to the individual learning rates η of the participants.

group-level parameters, μ_c^η and σ_c^η. In sum, we estimate for each context c and with all N participants the set of parameters $\theta_c = \{\mu_c^\eta, \sigma_c^\eta, \eta_{c,n=1}, \eta_{c,n=2}, \ldots, \eta_{c,n=N}\}$.

The Bayesian approach in turn means that we first make our prior assumptions explicit by defining prior probability distributions and likelihood functions for each parameter. Bringing the priors and the model to the data, we arrive at a posterior probability distribution for the set of parameters. All priors for the learning rates were set to $\Phi(N(-.5, .5))$. Here Φ is the cumulative density function of the standard normal and is used to compress the learning rate values to the interval $[0, 1]$. This transformation results in the effective prior being centered around a value of .33, which is a sensible value given previous findings (Gershman, 2015; Fontanesi et al., 2019).

The Bayesian posterior distribution of a parameter indicates the probability of each value given the specified priors, likelihood functions, and the observed data. As the calculation of such a distribution has no closed-form solution for most priors and likelihood functions, we use a Markov Chain Monte Carlo algorithm that numerically approximates the posterior distribution (McElreath, 2020). Detailed information about the model estimation method and our specifications are documented in Appendix A.4. Overall, the hierarchical Bayesian approach holds multiple advantages over other analysis methods: It requires the explicit specification of prior distributions and a mechanism to calculate the likelihood. Further, the results are not point estimates but probability distributions, which are more informative and allow for a more intuitive interpretation of parameters. Last, the hierarchical approach informs each individual case with the group-level data. This can lead to more precise estimates, especially if only a few data points were present for certain individuals (McElreath, 2020, a phenomenon called "shrinkage").

Figure 2.2 displays for each context the posterior distributions of the group-level parameters for the mean learning rate, μ_c^η. The dot and whiskers beneath the distributions indicates their mean value and the 90% credible interval. The latter is the interval in which 90% of all posterior samples lie. It can thus be directly interpreted as the interval containing the true parameter value with a probability of .9. Note that all estimated learning rates are clearly below the priors, which were centered around a value of .33. This confirms that the data were informative enough to overrule the priors. One reason for these low estimates may lie in the rounds in which participants updated their beliefs in the direction opposite the price movement (inverse updates). An inverse update would be best captured by a negative learning rate. However, we restrict the learning rates to the interval

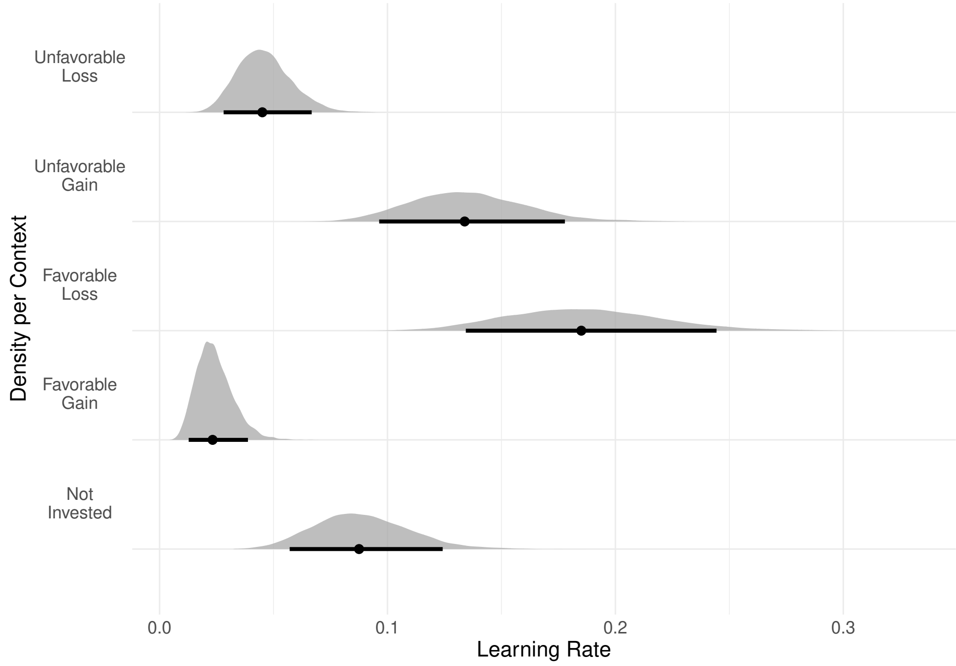

Figure 2.2: Posterior probability distributions for the group-level mean parameters of the learning rates, μ_c^η. The dots and whiskers indicate the mean values and the 90% credible intervals, that is, the inner area in which the value of the parameter is estimated to lie with a probability of .9. These values are also tabulated in Table 2.5.

$[0, 1]$. Thus, the closest estimate within the given range is 0, pulling down the overall estimates of the learning rate. Nonetheless, all 90% credible intervals are reasonably far from zero, indicating that true learning rates below zero are rather improbable.

The average estimated learning rate when not invested was .082 according to its posterior distribution. Contrary to our predictions, this estimate not generally lower than those for rounds where the participants were invested. The pattern observed in the learning rates when invested are however in line with our conceptual framework: The average learning rate for favorable information is estimated to be .019 when in a gain position and .176 in a loss position; the average learning rate for unfavorable information is estimated to be .127 when in a gain position and .033 in a loss position; We directly test whether this constitutes a credible interaction between investment position and information favorability in the following way: First, we calculate the difference between gain, g, and loss, l, samples separately for unfavorable, u, and favorable, f, information, $D_f = S_{gf} - S_{lf}$ and $D_u = S_{gu} - S_{lu}$. The resulting distributions of samples, D_f and D_u, can be directly interpreted as the posterior distributions of a difference in learning rates between gain and loss positions. Next, we can calculate the difference between the samples of these new distributions to evaluate the interaction $D_I = D_f - D_u$. Overall, this leads to a difference-in-difference calculation, which tests whether the different effects of investment position differ between information favorabilities. The 90% credible interval of this new posterior distribution does not include zero (mean -0.251, 90% credible interval $[-.339, -.17]$), indicating a credible interaction between the contexts. This analysis confirms that both factors—investment position and information favorability—are necessary to accurately describe participants' belief updating.[22] Note also that the standard deviation of the group-level parameter distribution is estimated to be rather high (See Table A.5.1 in the internet appendix), indicating some heterogeneity between participants.

[22]See Appendix A.5 in the internet appendix for an illustrative example of how the estimated parameters would affect investment behavior.

Table 2.5: Parameter Estimates per Model

Par.	CSRL	Invested	Favorability	Gain vs. Loss	Standard
μ_n^η	.082 [.049, .124]	.141 [.087, .205]	.078 [.041, .125]	.075 [.036, .127]	
μ_{fg}^η	.019 [.008, .036]				
μ_{fl}^η	.176 [.115, .247]				
μ_{ug}^η	.127 [.082, .179]				
μ_{ul}^η	.033 [.016, .056]				
μ_i^η		.101 [.058, .155]			
μ_f^η			.140 [.085, .207]		
μ_u^η			.095 [.576, .141]		
μ_g^η				.003 [.001, .007]	
μ_l^η				.108 [.066, .159]	
μ^η					.108 [.059, .166]
BF	-	> 100	> 100	> 100	> 100

Note. Results of the hierarchical Bayesian model estimation of the reinforcement learning models. The models were all fit on the belief reports of Phase 1. The table displays the sample mean and (in brackets) 90% credible intervals of the posterior for the mean parameter of the group-level learning rate distribution μ^η. The "Par." column indicates for which parameter the row reports the calculated values. The subscripts in the parameters are coded as n = not invested, i = invested, u = unfavorable, f = favorable, g = gain, and l = Loss. The row labeled BF displays the Bayes factors of a comparison between the given model and the full context-sensitive reinforcement learning (CSRL) model. A $BF > 1$ indicates evidence in favor of the CSRL model.

Through the increased number of parameters, the CSRL model has more degrees of freedom, that is, it has a higher complexity compared to a standard reinforcement learning model with a single learning parameter. We thus now investigate whether this complexity is justified by the data by comparing the explanatory power of the CSRL model to reduced versions. We first apply the same estimation procedure—including the same prior distributions and data—to reduced forms of the model: a standard reinforcement learning model with a single learning rate and models differentiating either only between portfolio states (i.e., invested or not), investment positions (gain vs. loss), or information favorability.[23] Table 2.5 shows the resulting average estimated group-level mean learning rates (μ^{η}) as well as the comparisons to the full CSRL model. To compare the models we calculate Bayes factors by estimating the marginal likelihood of the non-hierarchical implementation of each model using bridge-sampling (Meng & Wong, 1996). This approach reveals under which model the data is more likely to occur (see e.g., Andraszewicz et al., 2015) and also punishes for model complexity (it contains an "automatic occam's razor"; Kass & Raftery, 1995). It therefore constitutes a rigours test to determine whether the complexity of the CSRL model is justified. A Bayes Factor of above one favors the model at test in explaining the data best. We find that comparisons between each model and the full CSRL model consistently yield $BF > 100$, indicating strong evidence in favor of the CSRL model.

The parameter estimates of the models with a reduced number of parameters can also be used as a way to test the main effect of each factor. Using this method we find no evidence for a difference in learning rate between being invested and not being invested (mean difference in parameter estimates .04, 90% credible interval $[-.034, .118]$). This is also expressed in a Bayes factor of $BF = 22.07$ in favor of the model with a single learning rate when comparing to the one differentiating between being invested or not invested. In contrast, we do find separate effects for investment position and information favorability. The learning rate in a loss position is estimated to be higher by .071 compared to the gain position (90% credible interval $[.032, .123]$), indicating that participants reacted stronger to new information when in a loss position. Regarding price movement favorability, the learning rate is estimated .105 higher when witnessing an unfavorable price movement compared to a favorable one (90% credible interval $[.062, .156]$). Lastly, comparing only

[23]Note that the models estimating the difference between investment positions as well as favorability do use a third learning rate for the not-invested context. This is because the latter cannot be assigned to one of the two contexts of interest (e.g., gain or loss).

the reduced models with multiple learning rates among each other results in a clear favorite in the model differentiating between information favorability ($BF > 100$).

In summary, this analysis confirms that investment position and information favorability generally interact in influencing belief updating. Even when estimated separately, investment positions and information favorability seem to influence the strength of the belief updating independently, as well as interacting with each other. While the pattern of learning rates estimated here does differ somewhat from the one found when analyzing the belief updates directly, they do agree on one aspect: A meaningful interaction between investment position and information favorability influencing the strength of the belief updates.

2.5 Discussion

We combine effects from the literature (e.g. Knutson et al., 2011; Kuhnen, 2015; Jiao, 2017) and propose a novel conceptual framework of how people update beliefs when facing new information depending on an interaction between their investment position (gain or loss) and information favorability. The conceptual framework of the belief-updating process predicts profit-harming investment behaviors regularly observed in the laboratory and field, such as a difference in the profitability of opening or closing investments (Akepanidtaworn et al., 2021) and the disposition effect (Odean, 1998). We confirm the feasibility of this framework and its consequences in an preregistered experiment utilizing a standard investment task (Frydman et al., 2014) where we additionally elicited participants' beliefs about the likelihood of a price increase in the next round. We analyze the reported beliefs using regression models as well as by estimating the parameters of the CSRL cognitive model. Differences in the strength of belief updating in the contexts of investment position and favorability of the information, as well as the importance of the interaction emerges as a robust finding in both approaches. Moreover, a model comparison reveals that our proposed model finds more support by the data than simpler models, even when punishing for complexity. In Phase 2 of the experiment we aimed for a more causal test of the effect of belief formation on investments by providing additional information about the probability of observing a price increase. We find that with our interventions the context-sensitivity of belief formations can only be mitigated and improve investment decisions once very precise information about the next price movement is provided. The within subject nature of our experimental design holds individual differences such as risk preferences, motivation, intelligence, and comprehension of the task constant. This allows a more causal interpretation of the observed impact of information provision on belief updating.

We first tested our hypotheses on belief updating in terms of direct updating values and compared them to those of a Bayesian learner. Generally the participants in our experiment updated their beliefs more strongly than a Bayesian learner, even more so for unfavorable information, which is in line with the findings by Kuhnen (2015). We also find a significant interaction between investment position (gain vs. loss) and favorability of the information. This interaction arose despite our incentivization of belief reports, which can itself lead to more rational belief formation (Zimmermann, 2020), and persisted even when participants were shown the Bayesian probability of a price increase. Moreover, controlling for individual abilities and preferences by including them in our models does not change the significance of the interaction for belief formation. One pattern we do not observe in our data, using either of our analysis approaches, is that of participants updating their beliefs more strongly for assets they owned (Hartzmark et al., 2021). A possible reason may lay in the mechanism proposed by Hartzmark et al. (2021), which is based on attention. Because our participants only had to focus on a single asset, a lack of attentional resources is unlikely, which may be why ownership of the asset did not influence updating strength. Attentional effects, and with it ownership, might however gain importance in a more complex scenario than the one used in our experiment.

As belief updates lie at the core of our hypotheses, we fit a cognitive learning model, the CSRL model, to our data as a second method to investigate and ensure the robustness of our results. We find that the CSRL model provides a good fit to the data and its complexity is well justified by the data compared to simpler versions. The literature using models including differential updating is rather novel but a promising direction of research (see Gershman, 2015). To the best of our knowledge our CSRL model is the first to implement these effects in an investment context and using an explicit dependence on these circumstantial factors.

Our results are further in line with explanations based on a belief in buying-price reversion (Jiao, 2017). Note, however, that we do not propose that people hold a belief in buying-price reversion *a priori*. Rather, the proposed context-sensitive belief updating leads to expectations that resemble a belief in mean reversion, although they emerge only once an investor is already invested. It may prove fruitful for future research to explore other contexts that may influence learning in settings beyond financial decision making.

Turning to the behavioral consequences of our findings regarding belief updating, we do find a strong difference in profitability (i.e., success rates) between liquidations or jumps from gain and loss positions. In contrast, the general finding of opening investments being more profitable than closing them (Akepanidtaworn et al., 2021) does not replicate in our experiment. Note that Akepanidtaworn et al. (2021) do not find this profitability difference for investors following a momentum strategy. In our experimental setting, the actions of a risk-neutral Bayesian investor would strongly resemble a momentum strategy. It might therefore not be surprising that we find only a negligible effect in this environment. Only decisions to jump between investments were improved by the information we provided, and only substantially so in the full-information condition. This seems sensible, given that belief updates were also only strongly influenced when participants were presented with the objective probability of a price increase. In our experiment we can therefore only partially confirm a fundamental difference between the profitability of buying and selling decisions, namely, when also differentiating between the original investment position from which the decision was made.

Participants in our experiment were more willing to liquidate losses than gains, leading to overall negative disposition effect values. These values were, however, clearly above that of a risk-neutral Bayesian investor, indicating that participants could have increased their earnings by holding gains longer and selling losses earlier. Here again, providing information in the partial-information condition did not move disposition effect values much, whereas the full-information condition did have a very strong impact. Our results complement other experimental findings pointing toward the importance of belief updating. For instance, Kuhnen et al. (2017) proposes motivated beliefs to be the reason why people stick to their portfolios in the losses, which conversely can lead to the disposition effect. Our participants had higher disposition effect values compared to a risk-neutral Bayesian investor even when provided with the true probabilities of a price increase. Thus, while the data of the experiment underline the importance of belief updating in investment behavior, there remains room for preference-based explanations of the disposition effect(e.g. Tversky & Kahneman, 1992; Barberis & Xiong, 2012), which do, however, have their own problems and limitations (see, e.g. Hens & Vlcek, 2011).

In summary, we do find a strong effect of the context on belief updating and demonstrate how it can lead to adverse investment behavior. The effect of the context on belief updates can, however, be generalized toward other sequential decision-making scenarios: Take as an example a project manager who receives mixed news about an otherwise well-running project. Overweighting the more unfavorable aspects of the news may drive the manager to be overly skeptical about the project. The sunk cost fallacy Arkes & Blumer (1985) is another well-established effect that may stem from belief updating: Our participants gave less weight to negative information after incurring a loss; someone influenced by the sunk cost fallacy may similarly disregard negative information (and overweight positive information) after having incurred a loss in terms of time, money, or effort (i.e., a sunk cost). Findings such as those by Asparouhova et al. (2009) also indicate that such belief effects are rather universal as they can also appear when a person thinks about abstract probabilities. It could thus be fruitful for future research to expand and apply our conceptual framework across different domains. In particular, the CSRL model provides a useful quantitative framework to derive clear and testable predictions.

Preferences are largely assumed to be stable traits; belief updating, in contrast, may be improved by education or providing better information. It is thus valuable to know its mechanisms, what influences it, and how decisions are affected by it. Our results stress the importance of making relevant information clear, precise, and easily accessible, especially when financial decisions are to be made. There are also policy implications, for example, for trading interfaces. Here interventions such as hiding the original buying price of an asset—thus making the gain and losses of an investment position less present to the investor—have already shown to be fruitful (Frydman & Rangel, 2014), and the promotion of limit orders (Fischbacher et al., 2017) may help by making selling decisions before expectations can be affected by contextual factors. Overall our experimental data reveal the importance of the context sensitivity of belief updating and investment behavior. They show that providing clearer information can indeed play an important role in aiding investors in achieving their goals.

Chapter 3

Take your Time: How Delayed Information and Restricted Decision Opportunities Improve Belief Formation in Investment Decisions

Kevin Trutmann, Steve Heinke, Jörg Rieskamp

The manuscript has been published in *Finance Research Letters*.

Abstract: Investors' belief-updating is often influenced by factors such as the current investment position and whether information is subjectively favorable. Such motivated beliefs can lead to profit harming decisions. We argue that the degree of involvement with the development of an investment is a driver of such motivated beliefs. In a pre-registered experiment we aim to lower involvement by delaying information and committing participants to a portfolio. We show that this brings participants' beliefs significantly closer to a Bayesian benchmark. Separating information processing and belief-updating from decisions thus appears as a promising and easy to implement intervention to improve financial decisions.

3.1 Introduction

Processing information and forming beliefs about future asset returns is an essential aspect of successful financial decisions. Investor's information processing is however often influenced by contexts irrelevant for predicting future returns. For instance, the way investors update their beliefs given new information depends on whether their current investment position reflects a gain or a loss and whether the information they receive is favorable with respect to their investment (Kuhnen & Knutson, 2011; Kuhnen, 2015; Trutmann et al., 2022). Such context-dependent belief updating processes can have a detrimental impact on investor's decisions and profits (Grosshans et al., 2020). At first glance, it appears natural to assume that people's belief updating will profit from receiving information instantly. Likewise it should be beneficial to respond immediately to new information. However, piece-wise and instantaneous information processing might lead to less deliberated decisions (Imas et al., 2022). The objective of the present work is to investigate whether people's belief updating might benefit from delayed information provision and restricted decision making opportunities. By separating information processing from decision making we expect improvements in the beliefs which can in turn lead to better investment decisions.

The hypothesized underlying mechanism motivating our intervention is that cognitive distance to a prior decision helps mitigate the effects of the context in belief formation such as motivated beliefs, regret aversion and cognitive dissonance. When people immediately receive information affecting their profits, they may update their beliefs very subjectively (e.g. expecting the value of an investment to return to the initial buying value, Trutmann et al., 2022). In contrast, delayed information provision and restricted decision-making opportunities separate the information processing and the decision. Delaying information therefore provides time to process new information thoroughly and may have similar positive effects as a dedicated waiting period (Imas et al., 2022). Consequently, these changes to the decision environment might reduce the involvement of an investor with their prior decisions. This in turn can improve belief updating and thereby also subsequent decisions. Based on these assumptions we predict that, first, delayed information provision and restricted decision-making opportunities brings people's beliefs closer to a Bayesian belief which serves as our rational benchmark. Second, such improved beliefs might translate into to more profitable subsequent investment decisions.

Past research has illustrated that it can be beneficial for subsequent decision making to distance oneself from previous decisions (e.g. Chang et al., 2016; Heinke et al., 2021; Rotaru et al., 2021). These studies focus mainly on preference based explanations, disregarding the potential role of beliefs. We argue that cognitive distance also affects and improves decision makers' beliefs by bringing it closer to a rational Bayesian belief. Another strand of literature shows that manipulating frequency of decisions and information can improve decision performance (e.g. Benartzi & Thaler, 1995; Gneezy & Potters, 1997; Thaler et al., 1997). Our pre-registered experiment and analysis differs from these studies in three substantial ways: First, while previous investigations have mainly focused on the anticipatory effect of the manipulations we focus on their influence on beliefs and decisions *after* the intervention has been experienced. In other words, rather than investigating how decision makers react to the anticipation of a changed environment, we study the effects of having experienced the environment. As most investment decisions have to be repeated or reconsidered, this is an important open research question in the literature. Second, our within-subject study design controls for preferences and other inter-personal differences, allowing for a more causal interpretation of the effects on beliefs. Third, the elicitation of participants' beliefs about future price movements allows us to differentiate mechanisms based in risk preferences from improvements in beliefs. Not only is this important because investment decisions are largely based on the investor's expectations (Trutmann et al., 2022; Grosshans et al., 2020), but also because knowledge about the cognitive mechanisms behind a behavioral change is valuable when planning new interventions to change and improve behavior.

3.2 Experimental Design

We want to test the hypothesis that separating investment decisions from information processing improves investors' beliefs. Therefore, we need an investment task where participants learn something about the future value of an asset and which also allows to determine a rational benchmark belief. This section describes briefly the pre-registered study. Detailed descriptions, instructions, screenshots of the task as well as the pre-registration and further tables can be found on OSF (osf.io/kxqda/).

Investment Task We adapt a standard investment task with a single asset with an up- or downward trending price (Weber & Camerer, 1998). Participants could either buy or short sell up to four shares of the asset over a "block" of eight rounds with no transaction costs. Throughout a block the assets' price followed either an upward or downward drift (chosen with equal probability), implemented by a .65 probability of a price in- or decrease respectively. Participants were instructed about all details of the task. The posterior probability of an upward or downward trend in this task can be determined by Bayesian probability updating, which in turn allows to predict the probability of a price increase for the upcoming round.

A standard round consisted of the following sequence of events: First, participants decide whether to buy or sell any shares. Next, using a slider they indicated their estimated probability for a price increase from the current to the subsequent round. Notice, that we ask participants for this probability as it is immediately relevant for their decision. However, as this probability is only influenced by the probability of an upward drift it perfectly correlates with the latter. Lastly, participants observe the price change.

For every participant eight different price paths were generated, each with its own 50% chance of an up- or downward drift. Over the course of the experiment, participants experienced each of the eight individual price paths in each of the three conditions, resulting in 24 path-condition combinations (blocks). The order of blocks was randomized for each participant.

Conditions After an initial three rounds in which participants could not change their randomly allocated portfolio the conditions varied in the following ways: In the *Baseline* condition participants played five additional standard rounds as described above. In the *Blocked Trades* condition participants had to decide upon one portfolio which was then held for the upcoming five rounds. They observed the price change and reported their beliefs in a price increase in each of the five rounds. Lastly, the *Delayed Information* condition also required participants to commit to their portfolio over five rounds. Here the price updates were however only displayed in a list format at the end of the five rounds and participants only report their belief about a price increase in the final round. In all conditions, participants could then make one last investment for a final round where their beliefs about the final price update were elicited. Note that at the final round participants in all conditions had received the same information. The only difference between the blocked trades and delayed information condition was therefore the time at which beliefs could be updated, as both could only act upon their beliefs in the final round.

Procedures & Sample The online experiment was implemented using oTree (Chen et al., 2016) and conducted through the prolific.co platform in March of 2022. The main task was preceded by a detailed and interactive instruction, a comprehension quiz as well as multiple training rounds. This was to make sure that participants had fully understood the task as well as the option of short selling the asset (the median participant answered 6 out of 8 quiz questions correctly). The study took around one hour to complete. Participants were incentivized by paying out the returns of a randomly selected block and earned an average of \$13.55. Guided by a power analysis (see the pre-registration) we aimed at a sample of around 300 participants, giving us a power $> .95$ to detect a small effect of Cohen's $d = .2$. The final sample consists of 315 participants (156 female, mean age 33.36, SD 10.86) with 68.57% of the sample declaring having at least some previous experience with financial investments.

3.3 Results

We start the pre-registered analysis by investigating how increasing the cognitive distance from prior decisions reduces the deviation from the rational Bayesian benchmark. Next we take a closer look at participants' trading decisions and study whether the treatments improved profits. Finally, as up- and downward drifting asset prices might differ in their demand for cognitive resources, we study whether the improvement by the treatments was stronger in either case.

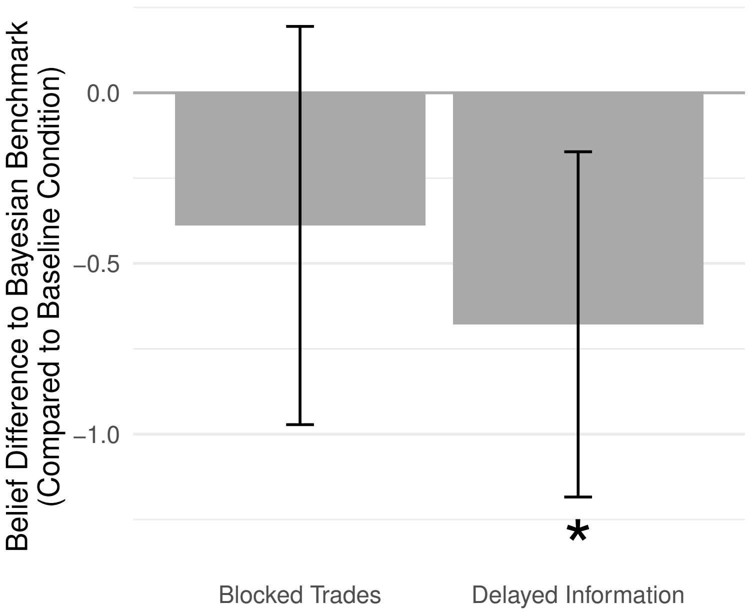

Figure 3.1: Treatment effect on the average absolute difference between the Bayesian Benchmark and participants' reported beliefs for the final investment decision. The bars show the average difference to the baseline condition, with lower values implying stronger improvement. The average differences are the respective coefficients of treatment dummies from an estimated OLS regression, error bars indicate 90% confidence intervals with standard errors clustered by participant and price path, and the asterisk denotes p-Values < .05 for a one-sided t-test.

Beliefs We measure the accuracy of beliefs by calculating the absolute distance of the reported beliefs at the end of a block to the Bayesian benchmark. The average of this measure in the baseline condition was 18.47 percentage points. A decrease in this value means beliefs closer the Bayesian updater, which we interpret as an improvement in the belief formation. Figure 3.1 shows the effect of the treatments on this distance compared to the baseline condition, estimated by a OLS regression with treatment dummies and standard errors clustered on participant and price path. Blocking trades brought participants' beliefs (as compared to the baseline condition) on average 0.39 percentage points closer to the Bayesian benchmark ($p = 0.096$, one sided test). Delaying the information in turn reduced the deviation from the Bayesian benchmark on average by 0.68 percentage points ($p = 0.004$, one sided test). Thus, Figure 3.1 reveals that both treatments improved the beliefs slightly, whereas only delaying information is statistically significant.

Investment Decisions Over the five trading rounds and the extra round at the end of a price path participants on average invested (i.e. held or sold short)

2.76 shares in the baseline condition. In general, we find that participants' beliefs ($p < .001$) as well as the price drift direction ($p < .001$) significantly influences investment decisions in an ordered logistic regression, underlining that participants generally understood the task.[1] At the start of the investment period they invested slightly more when facing blocked trades (0.07 shares, $p = 0.017$) or delayed information (0.11 shares, $p = 0.009$). This is in line with other results investigating the anticipatory effect of varying the length of evaluation periods and risk taking (e.g. Gneezy & Potters, 1997; Thaler et al., 1997; Larson et al., 2016).

In contrast to this previous work we focus on the end of the investment phase of a block, once participant experienced the price changes and the specific treatments. To investigate the success rate of the decisions after our interventions we calculate each participants' "drift hit rate": the number of shares held in the final round that are in line with the current asset's price drift. For example, shorting two shares during a downward drift would be counted as a correct investment and coded as a value of plus two. In contrast, shorting two shares during an upward drift counts as an unsuccessful decision and is coded with a value of minus two. The binary version of this measure is only concerned with whether the general portfolio (holding or shorting shares) is in line with the current price drift (coded as zero or one, hence binary). Note here, that random investments would generate an average drift hit rate of zero for both measures. Table 3.1 compares these measures between conditions using an OLS regression and a logistic regression for the two measures respectively. In the baseline condition participants surpassed the random benchmark significantly in both measures. However, we find no significant improvement for the two treatment conditions.

[1] The table can be found on the project's OSF repository (osf.io/kxqda/).

Table 3.1: Treatment Effect on Drift Hit Rate

	(1)	(2)
Baseline	0.527***	0.377***
	(0.071)	(0.051)
Blocked Trades	-0.058	-0.058
	(0.043)	(0.033)
Delayed Information	-0.088	-0.065
	(0.055)	(0.034)
R^2	0.0001	-
AIC	-	8373.69
Obs. (Clusters)	7560 (315, 8)	6161 (315, 8)

Note. * = $p < .05$, ** = $p < .01$ *** = $p < .001$. One sided p-values with standard errors (in parentheses) clustered on participant and path level; Treatments are dummy coded; Model (1): OLS regression with Drift Hit Rate as outcome variable; Model (2): Logistic regression with Binary Drift Hit Rate as outcome variable.

Note that our task involves a constant asset price trend and therefore trend chasing is a profitable strategy. As a consequence of this specific setting, having stronger beliefs than a Bayesian updater can be advantageous in this setting. If an investor is for example strongly convinced of an upward trend where such a trend is indeed present, it may help reduce the investor's perceived uncertainty and thereby lead them to invest more in a profitable opportunity. As a result, participant's beliefs may deviate less from a Bayesian benchmark in our treatment conditions while they simultaneously make more cautious decisions and thereby participate less in profitable investments.

Regarding participants' overall earnings, using an OLS regression we find the average points earned over all eight rounds in the baseline condition to be significantly positive with a value of 9.06 ($p < .001$). The treatments did lead to slight, though statistically insignificant, improvements in participants' earnings by 2.21 (blocked trades, $p = 0.067$) and 1.6 (delayed information, $p = 0.253$) points. Note that earnings are calculated over all eight rounds of a block whereas the hit rates reported above are only concerned with the final round.

Effects of Price Drift Profiting from a downward trending asset requires short selling. This strategy may however be more cognitively demanding as the meaning of price moves are "flipped" (i.e. falling prices now being positive). We therefore further investigate the underpinnings of our results by adding the price drift direction as a dummy coded predictor to the model underlying Figure 3.1, investigating the belief difference to a Bayesian benchmark.[2] Doing so reveals a significant improvement from both treatments during a downward drift (by 0.63 and 0.57 percentage points respectively). The interaction terms reveal that the difference in condition effects seen in Figure 3.1 are due to blocked information being substantially (0.48 percentage points), albeit not significantly ($p = 0.096$), less effective during upward drifts.

Regarding participants' investment decisions, a strong effect of the price drift alone can be observed when adding it as a predictor to the models presented in Table 3.1. Here, an upward price drift lead to significant improvements in both measures (1.59 and 1.39 for the binary version) compared to a downward price drift. In addition, there is a significant negative interaction between an upward price drift and the blocked trades treatment in both measures. This interaction indicates that the intervention of blocking trading decisions was more effective during a downward compared to an upward price drift. Thus blocking trades was more sensitive to the price drift compared to delaying information both for beliefs and investment decisions.

3.4 Discussion

In this study we investigated whether increasing the cognitive distance to prior decisions improves belief formation and thereby the success rate of subsequent investment decisions. We find both of our treatments to bring beliefs closer to those of a Bayesian benchmark. This improvement was statistically significant when information provision was delayed. Such a stronger effect of delayed information provision is in line with motivated reasoning, as delayed information allows investors even more distance to their previous decisions than only blocking investments. Further supporting this interpretation, the "weaker" intervention of blocking trades only improved beliefs during downward price drifts, while delaying information improved beliefs independent of price drift. While the uncovered effects are small, they could potentially influence longer investments by compounding over time. Small effect sizes might also indicate that further mechanisms which do not change with involvement also play a role in the discrepancy between participants' beliefs and the Bayesian benchmark.

[2] These tables can also be found on the projects OSF repository (osf.io/kxqda/).

In our setting, which we adapted from a simple standard investment task, the improvement of the beliefs did not lead to higher profitability of all investments decision. A partial explanation may be that overly strong beliefs in the correct direction (e.g. being very sure about an upward trend that truly exists) was beneficial in our task. In environments where trend chasing can be harmful to profits, such as the impending burst of a bubble (Hefti et al., 2018), having rational beliefs about future price changes may be even more important. While the rather small effects on beliefs might be the result of the very simple and stylized environment they also leave room for mechanisms based in investor preferences. Moreover, beyond the effects on beliefs, our results are the first to shed light on the effects of restricted trading and information provision *after* investor experienced them, rather than anticipating them. Lastly, we generally find a significantly worse investment performance during downward price drifts in our baseline condition, and blocking trades being significantly more helpful there than during an upward drift. Falling prices therefore overall seem more demanding for participants, and future investigations may focus specifically on how to aid investors to make the best out of falling markets.

In general our results fall in line with previous literature which, for a variety of reasons, recommends lay investors to check their portfolios less frequently (Benartzi & Thaler, 1995; Odean, 1999; Mosenhauer, 2020), albeit for a different reason. Ideally, expanding on the small effects we found, investigations into different interventions using delays and bundling of information can be a further step toward better investment decisions.

Chapter 4

Degree of Involvement in Decisions and the Likelihood to Stop Investing among Professionals

Kevin Trutmann, Steve Heinke, Céline Rudin

The manuscript has been submitted for publication.

Abstract: The likelihood of stopping an investment project differs after an experienced gain or loss. We investigate how the degree of involvement in prior decisions affects the subsequent decision to change an investment. To this end we conduct a lab-in-the-field experiment with professional participants from the finance and controlling department of a large infrastructure company. In line with the hypothesis and prior findings from student samples we find that lower involvement in the decision is associated with a higher likelihood of changing the investment project after a loss. However, this difference disappears with age, which we interpret as experience in the professional career.

4.1 Introduction

Many start-up founders speak openly about their failures and what they learned from them. For instance, Andrew Wilkinson described on Twitter how he burned 10 million dollars of his own money on an ultimately failed workflow management platform.[1] At the core of his and others' sequence of risk-taking decisions lies a stopping dilemma: Should they continue a project with uncertain large rewards in the future while losing money along the way or stop the losses by not proceeding with the venture and give up the chance for a large reward? Experiencing the outcome of a prior decision that one has been involved in might lead to cognitive or emotional costs when evaluating whether to continue the endeavor, which can in turn lead to suboptimal stopping (Chang et al., 2016; Cohen & Erev, 2021; Martens & Orzen, 2021).

In this study we investigate how the degree of involvement in decisions which lead to experiences of gains or losses affects the subsequent decision to stop the investment. Our analysis focuses on a specific sample of professionals in the finance and controlling department of one of the largest Swiss infrastructure provider, who regularly deal with such high stakes investment decisions. Moreover, we investigate whether age lowers the impact of personal involvement on decisions of this type, as this correlates highly with seniority and thus professional experience in their career.

A common observation in these stopping dilemmas is that human (e.g., Thaler, 1980) and other animal (e.g., pigeons; Navarro & Fantino, 2005) decision makers tend to over-commit resources—money, effort, or time—to failing enterprises. They are more likely to continue a project once large initial investments are already made, an effect known as the sunk cost fallacy (Thaler, 1980; Arkes & Blumer, 1985), escalation of commitment (Staw, 1976, 1981), or the Concorde fallacy (Dawkins & Carlisle, 1976; Arkes & Ayton, 1999; Navarro & Fantino, 2005). While money that was burned due to stopping too late is rather salient, the opposite—giving up too early—gets less attention, as that counterfactual is never realized once a project is stopped. Early stopping is, however, a regularly reported behavior in experimental studies. For example, Zikmund-Fisher (2004) observes in a sequential risk-taking task where winning is rare that individuals stopped after six rounds, even though the optimal strategy suggests stopping after 10. This phenomenon is also known as learned helplessness (Seligman, 1972; Teodorescu & Erev, 2014),

[1]https://twitter.com/awilkinson/status/1376985854229504007

de-escalation of commitment (Heath, 1995; McCain, 1986), the reverse sunk cost effect (Zeelenberg & van Dijk, 1997), or early termination in optional stopping tasks (such as the secretary problem; Seale & Rapoport, 1997; Brockner et al., 1979; Hoelzl & Loewenstein, 2005; Lejuez et al., 2002; Wallsten et al., 2005) (see Cohen & Erev, 2021 for a framework reconciling these seemingly opposing effects).

Motivated beliefs (Fischer & Maier, 2019; Jiao, 2017) or emotions (Corgnet et al., 2020) might work as a mechanism to amplify the behavioral bias in risk taking after an initial decision. As such, decision makers may, for example, avoid cognitive or emotional costs such as feeling regret (Strahilevitz et al., 2011) or admitting a mistake (e.g., cognitive dissonance; Festinger, 1962). In line with such a cognitive dissonance mechanism, Chang et al. (2016) finds in both experimental data with students as well as field data of retail investors that investors are more likely to liquidate a losing investment if the initial investment decision was made by a fund manager rather than by themselves. Thus when investors buy a stock themselves, they are less likely to sell losing positions than winning ones—a phenomenon called the disposition effect—which can be caused by motivated belief distortions (Fischer & Maier, 2019; Jiao, 2017; Trutmann et al., 2022, 2023). Two studies especially corroborate this finding: In the field Lehenkari (2011) analyzes data from the Finish stock market and finds that investors show a significantly higher disposition effect when they were personally responsible for the investment. In a laboratory setting Martens & Orzen (2021) modified an experiment on escalation of commitment by Staw (1976) with real experienced outcomes and consequential decisions and also confirm that participants react differently to financial losses when they have been responsible for the prior decision.

Many of the aforementioned experimental investigations rely on student samples or retail customers. However, a crucial part of investment decisions and in particular those with larger budgets are made by professionally trained and educated executives in large-scale enterprises. Such individuals go through specialized training, and as they advance in seniority they learn from their experiences and feedback in their career. On one hand, it is possible for feedback to increase a bias toward profit-harming behavior if the experienced outcome of the profit-harming decision is on average positive (Cohen et al., 2020; Yakobi et al., 2020). On the other hand however, if the experience effect of feedback offsets these biases, professionals might show less profit-harming behavior (Erev & Roth, 2014). In this latter case the assumption of unbiased institutional investors might hold. The empirical results from the field are similarly contradictory regarding whether certain behavioral biases do in fact disappear for professionals in financial decision making: On the one hand, professionals are better at judging the quality of public signals in information cascades (Alevy et al., 2006) and are less prone to trading overpriced assets in a laboratory asset market where bubbles usually occur (Weitzel et al., 2020). On the other hand, a persuasive amount of literature documents biases among experts in corporate finance settings (Malmendier, 2018), such as CEOs in charge of mergers (Malmendier et al., 2011) or restructuring (C. F. Camerer & Malmendier, 2007), who are influenced in their decisions by major life experiences and their own personal traits, such as overconfidence. Mutual fund investors further fall victim to profit harming behaviors such as following short-lived price patterns (Griffin et al., 2003) being myopically loss averse (Haigh & List, 2005) and other biases similarly held by student samples (Abdellaoui et al., 2013; Schwaiger et al., 2020).

We investigate how the degree of involvement in prior decisions affects the subsequent decision to change an investment and how age, as a proxy for professional experience, may mitigate this bias. For this purpose we conducted a lab-in-the-field online experiment[2] with 167 employees in the finance and controlling departments of one of the largest infrastructure companies in Switzerland. The sample mostly consists of higher ranked employees working as project manager, controller, accountant or product owner in the the headquarter of the company. Thus, participants in our study frequently face and experience the outcomes of many high stake financial decisions.

In our experiment, participants faced a sequential risk-taking task. In each of 15 rounds they decided whether to invest in one of four risky options. They always had a choice: Keep the previous investment, switch to a different option, or stay out of investing completely. After each round concluded, participants received direct performance feedback. All options yielded on average the same positive return, had two of six positive outcomes (all with equal probabilities), and differed only in their variance between outcomes. Conceptually our work resembles that of Chang et al. (2016), with a between-group manipulation of the perceived degree of involvement in the decision by the cover story and the labeling of the risky options: While one group had to think of themselves as managers who invest in a start-up (*manager*

[2]Following the definition of Gneezy and Imas (2017) a lab-in-the-field study is conducted in a naturalistic environment targeting the theoretically relevant population but using a standardized, validated lab paradigm.

condition, high decision involvement), the other group had to hire a consultant who would invest for them in a start-up (*consultant condition*, low decision involvement). For a normative profit-maximizing decision maker the likelihood of switching to another option should not be affected by either the experience of a loss or gain nor by being in the manager or *consultant condition*. Nonetheless, we observe that participants in the low-involvement treatment (*consultant condition*) switch investment options after having experienced a loss more often compared to those in the high-involvement treatment (*manager condition*). This tendency is particularly strong for those below the median age of 34, while there is hardly any effect of the condition for those above the median age. Higher education further seems to mitigate the effect, while participants with a business background in their education or training do not behave differently from those without such a background.

We thus contribute to the literature twofold: first, we investigate the effect of involvement into a decision for a sample of professionals. This is important, as most investment decisions bearing consequences for the general public are made by such professional decision makers overseeing a large budget. Second, to the best of our knowledge we are among the first to investigate the role of age, and with it professional experiences in the effects of sunk costs. More professional experiences may bring with it the ability of detaching oneself from the decision being made, thereby lowering the effect of involvement on the decision.

In the following we describe the experimental design and our hypotheses in detail in section 4.2. We present the results in section 4.3 and conclude with a brief discussion of their immediate implications in section 4.4.

4.2 Experimental Design

We conducted a lab-in-the-field online experiment with an investment task, following a between-subjects design in which we changed the framing of the investment opportunities. This section describes the details of our experimental study.

Investment Task

Participants were initially endowed with 10,000,000 experimental currency units (ECU). In each of the 15 rounds they first decided whether they want to invest or keep the current cash balance. If they wanted to invest, they chose one of the four options and decided how much of their current cash will be allocated to this investment option. In each round participants could invest in at most one option. At the end of the round one of the returns was drawn from the respective investment option with the same likelihood for each of the six returns, by rolling a virtual fair die. Subsequently, the invested amount was multiplied by the return, and added to the participants' overall budget. Participants only saw the outcome of their chosen investment option.[3]

The five options remained the same over the whole 15 rounds and always consisted of four investment options and one option not to invest. All investment options consisted of six equally likely potential percentage returns of the invested amount. Table 4.1 displays the return outcomes of each investment option.

Table 4.1: Return Outcomes of the Investment Options

Investment Option	Potential return outcomes (%)	Mean return (%)	Variance
I	$-10, -5, -3, -2, +2, +20$	.33	108.27
II	$-10, -7, -2, -1, +7, +15$	.33	69.47
III	$-10, -6, -1, -0, +9, +10$	.33	63.47
IV	$-6, -4, -3, -1, +7, +9$	.33	38.27

All four investment options offered the same positive average return of .33%. Despite the positive overall return, each option contained only two positive outcomes meaning that the distribution of the outcomes was positively skewed. In other words, while participants would on average make a positive return, they would experience a negative return more often than a positive one. The options differed in the spread between the lowest and highest possible return as well as in the variances of these returns, ranging from 38.27 to 108.27. As the risk-free option of not investing always yielded a return of zero, it was therefore on average slightly more attractive to invest, as there was a potential to earn large significant returns.

Participants were allocated into two between-subject conditions which differed only by the label of the investment options (similar to the approach used in Chang et al., 2016): In the *manager condition* participants learned that they were investors and could invest in one of the four start-ups (high degree of involvement in the decision). In the *consultant condition* participants were informed that they were investors and could choose an investment consultant who would in turn invest in start-ups on their behalf (low degree of involvement in the decision). In

[3] See Appendix C.3 for screenshots of the task.

both conditions participants were informed about the potential outcomes for each investment option. The investment options were same for both conditions, only the label changed as a consequence of the cover story. Thus, participants knew the potential six returns if they choose one particular investment option and that one of the six returns was drawn randomly with equal likelihood. Participants were informed that their budget at the end of round 15 would determine their payoff.

Procedure

The online study was conducted in German and programmed using JavaScript and HTML. The invitation E-mail to participate in the study was distributed in July and August 2019 across the email distribution list of the finance and controlling department of one of the largest infrastructure companies in Switzerland. Most recipients on this distribution list work as project manager, controller, accountant or product owner in the headquarter of the infrastructure company. Thus, our study sample consisted of higher ranked employees in this major Swiss infrastructure company, who regularly make investment decisions for small and large projects up to multiple millions USD and monitor their outcomes. In total 167 employees completed the study. After participants gave their informed consent they were randomly assigned to one of the two treatment conditions. This led to our sample consisting of 93 participants in the manager and 74 in the *consultant condition.* This imbalance in sample size is within the range of what is expected when participants are randomly assigned to treatments. This sample allocation resulted in a power of at least .8 to detect medium sized effects of $d = .39$ or greater.

First, participants were provided with instructions for the experiment, which were followed by a short comprehension quiz and three practice rounds to ensure their understanding of the task. At the end of the 15 rounds of the investment task, participants were informed about their total payoff. Afterward, they could state their investment strategy in an open-form text field. The study ended with an exit questionnaire eliciting age, education, and self-reported risk-taking behavior (Richter et al., 2017) as additional control variables.

The top five participants with the highest final payoff additionally received a book coupon worth CHF 20 (around USD 22 at the time of the experiment) and this was communicated to participants before they started the study. As a true financial incentive for professional employees in a high wage country like Switzerland would surpass the budgetary constraints of most studies, this mechanism was rather an incentive to motivate participation. As the incentives affect all participants in the same way, we conclude that our between group comparisons are not affected by a weak incentivation. In Section 4.4 we elaborate and discuss whether and how this may influence our results and how we come to this conclusion.

Table 4.2: Summary Statistics of Participants

Variable	Total (N=167)	Manager (N=93)	Consultant (N=74)	Difference	p value
Age	36.892	35.430	38.730	3.300	0.048**
Gender	0.395	0.366	0.432	0.067	0.383
Higher education	0.611	0.591	0.635	0.044	0.567
Business background	0.389	0.398	0.378	-0.019	0.799
Risk taking	5.808	5.839	5.770	-0.068	0.840
Nonnative speaker	0.054	0.032	0.081	0.049	0.167

Note. Means of demographic indicators for total, manager, and *consultant condition*. The difference column reports the differences between the treatment groups, and the last column reports the p values of a two-sided t test, $^{**}p < 0.05$. *Age*, measured in years; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise; *higher education*, dummy variable taking on the value of 1 if the participant's highest education is a tertiary degree, 0 otherwise; *business background*, dummy variable taking on the value of 1 if the participant's highest education included some business aspects, 0 otherwise; *risk taking*, self-reported willingness to take risks in general on a scale from 0 (no willingness to take risks) to 10 (high willingness to take risks); *nonnative speaker*, dummy variable taking on the value of 1 if the participant's mother tongue is any other language than German, 0 otherwise.

Table 4.2 shows summary statistics for the total participant pool and for each treatment group separately. Overall, the two treatment groups are very similar, apart from participants in the *consultant condition* being a bit older.

Hypotheses

Involvement: Assuming a similar mechanism as the one proposed in Chang et al. (2016) our first hypothesis is the following: a reduced perceived involvement in a decision lowers the emotional costs of changing ones mind (i.e. the cognitive dissonance). Thus, we expect the likelihood of changing a project after losses to be higher in the *consultant condition*, due to a "denial of responsibility" (McGrath, 2017).

Professional Experience: Age correlates with professional experience. This type of seniority in a field of high stakes decision making can improve the decisions made (see e.g. von den Driesch et al., 2015), possibly through an improved ability for cognitive closure after a negative experience (Beike et al., 2009). Therefore, we measure the professional experience through the age of the participants in our sample. We expect the effect of the treatment to be smaller for participants above the median age compared to those below. This is because we expect older participants to be less affected by negative prior outcomes.

4.3 Results

In this section we investigate the investment decisions made in our online experiment. The presentation of the results is organized along the hypotheses and focuses on the effect of the degree of involvement, which was manipulated by asking participants to think of themselves as a manager or as someone who hires an investment consultant. We start with comparing the general investment behavior and testing for treatment differences. Next, we look at the likelihood of changing the investment option due to prior losses or gains. This is a direct test for the involvement hypothesis, as we expect the likelihood of changing after prior losses to be higher in the *consultant condition*. Finally, we check for robustness of the treatment differences to the effects of age. This serves as proxy for professional seniority and expertise and thus testing the professional experience hypothesis.

Investment Decisions

Overall, in most rounds participants chose to invest in one of the risky options. Only in 2.4% (SE: 0.3%) of all rounds did they not invest at all. There are, however, strong treatment differences: Participants in the *manager condition* decided not to invest at all in 4.2% (SE: 1.4%) of the rounds, which is significantly more than the 0.1% (SE: 1.4%) in the consultant condition (two-sided t test, p = .004). On average, one investment lasted 3.089 (SE: 0.109) rounds, i.e. the numbers of consecutive rounds participants were invested in one specific investment option, and this investment length did not differ significantly between the manager (3.13 rounds, SE: 0.161) and consultant (3.049 rounds, SE: 0.148) conditions.

Participants chose the least risky option (option IV) only in 15% (SE: 0.7) of the rounds, while the second least risky option (option III) was the most favored one (30% of the rounds, SE: 0.9) and options I and II were in between (24%, SE: 0.9 and 27%, SE: 0.9, of the rounds, respectively). Moreover, the amount invested in the least risky option (option IV) was significantly lower than the amount invested in the most risky (option I) and most favored (option III) options (see Appendix Table C.1.2). Furthermore, the average amount invested increased significantly throughout the rounds, on average by 113,000 ECU per round (see Appendix Table C.1.2). We do not find any differences among the treatment groups regarding the likelihood of investing in one specific option nor in the invested amount (see Appendix Tables C.1.1 and C.1.2). In other words, participants' preferences among the set of investment options did not differ between treatment conditions. The average profit per round was not significantly different from 0, with an average of $-4,936$ ECU (SE: 10,637); nor was there any treatment effect on the average profit. This amounted to an average profit of $-73,299$ ECU over all 15 investment rounds, again with no measurable treatment differences. Note, however, that there exists a notable gender difference, with women on average earning 50,610 ECU per round (753,899 ECU over all 15 rounds) more than men (see Appendix Table C.1.3).

Result 1: *There are no significant differences in the manager and consultant conditions with respect to chosen investment options, average investment length, or general profitability. Participants more often decided not to invest at all in the manager condition compared to the consultant condition, although at a very low level.*

Involvement Hypothesis

To investigate treatment group differences in the likelihood of switching between investment options, we constructed a dummy variable *change*, which takes on a value of 1 if the participant starts to invest in a new option and 0 otherwise. Figure 4.1 shows the average of this change dummy for each treatment group, separately for whether the previous round resulted in a gain or a loss and for all rounds. On average, participants changed the investment option in 20.5% (SE: 1.1) of the rounds in the *manager condition*, which is less often than in the *consultant condition*, where the option was changed in 22.9% (SE: 1.3) of the rounds. This differences is however not significant (two sided t-test $p = .430$, robust SE clustered on participant level). Conditional on observing a gain in the previous round, participants in the *manager condition* changed their investment in 19.2% (SE: 1.8) of rounds. In comparison, after a gain in the *consultant condition* a change was observed in only 15.5% (SE: 1.8) of rounds. We observe the opposite pattern following a prior loss. Here, participants in the *manager condition* behaved similarly to after a prior gain,

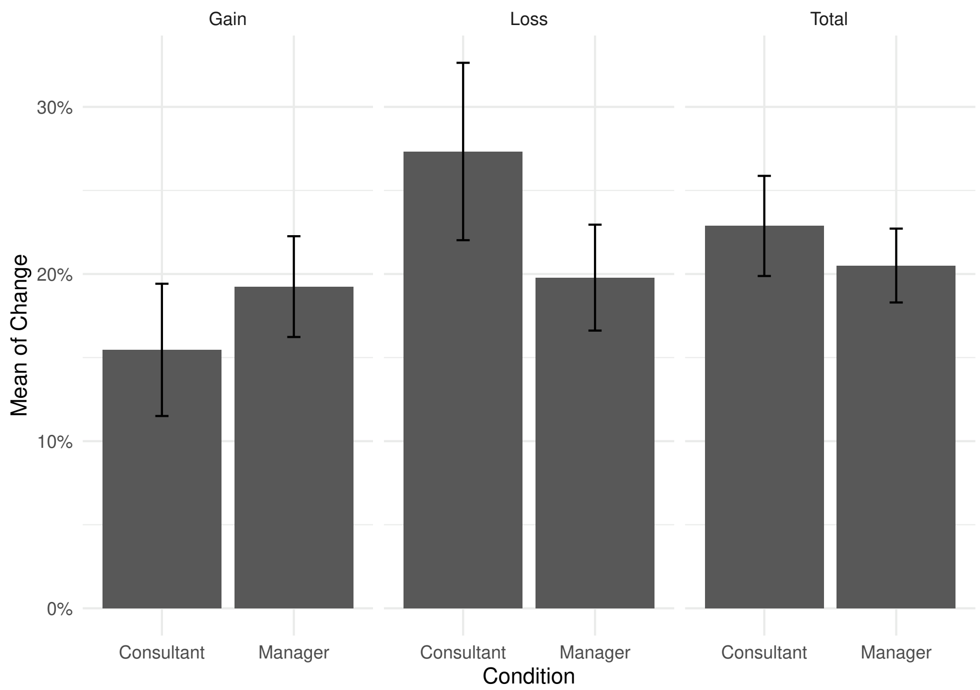

Figure 4.1: This figure displays separately for the consultant and *manager conditions* the likelihood of changing an investment option due to a gain, a loss, and overall. As the variable change takes on a value of 1 if the participant starts to invest in a new option and 0 otherwise its average represents the propensity to change in each group. Error bars depict the robust standard errors clustered on the participant level.

changing investments in 19.8% (SE: 1.5) of the rounds. In contrast, in the *consultant condition* this happened in 27.3% (SE: 1.8) of the rounds. Thus participants in the *manager condition* changed their investments in $19 - 20\%$ of all rounds, regardless of whether they experienced a prior loss or gain, while the likelihood to change investment options in the *consultant condition* depended heavily on the previous outcome, with a higher likelihood to switch after a prior loss.

Table 4.3 reports the corresponding probit regression analysis to Figure 4.1 and the tests for significance of the differences in the likelihood of changing the investment option mentioned above. We also control for age and gender effects as these variables are known to have a small but robust influence on decision making under risk (Filippin & Crosetto, 2016; Frey et al., 2021). In addition, we take into account the structure of the data where we observe multiple responses from the same individual. Thus observations can not be considered as completely independent, affecting the covariance matrix in the regressions. Ignoring this may

Table 4.3: Prior Losses and the Probability of Changing the Investment Option

Variable	(1) Manager	(2) Consultant	(3) Pooled
Loss $(t-1)$	0.00704	0.407***	0.0102
	(0.111)	(0.150)	(0.111)
Consultant			-0.107
			(0.151)
Consultant*Loss $(t-1)$			0.393**
			(0.186)
Controls	Yes	Yes	Yes
Observations	1,247	1,034	2,281
Pseudo R^2	0.00923	0.0256	0.0172
Log-likelihood	-610.8	-540.0	-1,152

Note. Probit regression with cluster robust standard errors (in parentheses) clustered on the participant level. Estimates for the constant parameter and control variables are omitted but were included in the regression. Two-sided t statistics are shown, $^{**}p < .05$; $^{***}p < .01$. *Outcome variable: Change*, dummy variable taking on the value of 1 if the chosen investment option in the current round is different from in the previous round, 0 otherwise. *Independent variables: Loss(t − 1)*, dummy variable taking on the value of 1 if the result of the previous round was a loss, 0 otherwise; *consultant*, dummy variable taking on the value of 1 if the participant was in the consultant (low-involvement) condition, 0 otherwise. *Control variables: Age* and *gender.*

lead to an overestimation of the precision of the estimated effect, which we combat by using a clustered sandwich estimator for the standard errors (see Cameron & Miller, 2015).

While in the *manager condition* the decision to change the investment is not significantly affected by a prior gain or loss (model 1), a change in investment is significantly more likely after an experienced loss in the *consultant condition* (model 2). Comparing the two conditions (model 3), the likelihood of a change after a prior loss is significantly higher among participants in the *consultant condition*, indicated by a significant interaction term in the model.

The observed differences in how losses affect the decision to continue an investment in both conditions are not subject to how prior losses are measured or to preferences. Similar results are obtained if the prior loss dummy is replaced by the proportion of rounds in which a loss was observed (see Table C.1.4) or if one controls for the self-reported willingness to take risks in investments (see Table C.1.5).

The treatment differences and in particular the strong interaction between prior loss and the low-involvement treatment support the involvement hypothesis: A lower degree of involvement in the decision leads to a higher likelihood of switching the investment option. This result can also be seen as a conceptual replication of Chang et al. (2016) in the context of business investments. They report that private investors become more likely to sell a losing investment position if it is a fund rather than a stock they picked themselves. Moreover, we interpret the results as evidence against an escalation of commitment explanation, as we do not observe a significant decrease in the likelihood of changing due to prior losses in either condition. The increase in the likelihood of changing after a loss in the low-involvement treatment (*consultant condition*) speaks more in favor of under-commitment to a project. Broadly speaking, this result supports a scapegoat explanation, as participants can blame a hypothetical consultant from the cover story for a loss.

Result 2: *We observe significantly more changes in the investment options under low involvement (consultant condition), in particular after a prior loss.*

Professional Experience Hypothesis

The age of of managers reflects their seniority and thus the cumulative professional experience and training received throughout their career path. A higher stock of personal experience may mitigate the need for a hypothetical "scapegoat", thereby reducing the observed differences between the manager and the *consultant condition* in the likelihood of switching investments. We investigate this idea by looking at the effects of age, higher education, and whether the education contained a business component.

Figure 4.2 depicts the average number of changes per treatment separately for participants younger and older than the median age of 34 years. Overall, older participants were committed significantly longer to their investment projects and switched in only 15.5% (SE: 1.1) of the rounds. In comparison, younger participants changed investment options in 27.0% (SE: 1.3) of the rounds. Those participants older than 34 chose to change the option with almost equal likelihood in the two conditions, that is, in 15.8% (SE: 1.5) of the rounds in the *manager condition* and 15.2% (SE: 1.5) of the rounds in the *consultant condition.* Participants age 34 or younger tended to change the investment option significantly less often in the

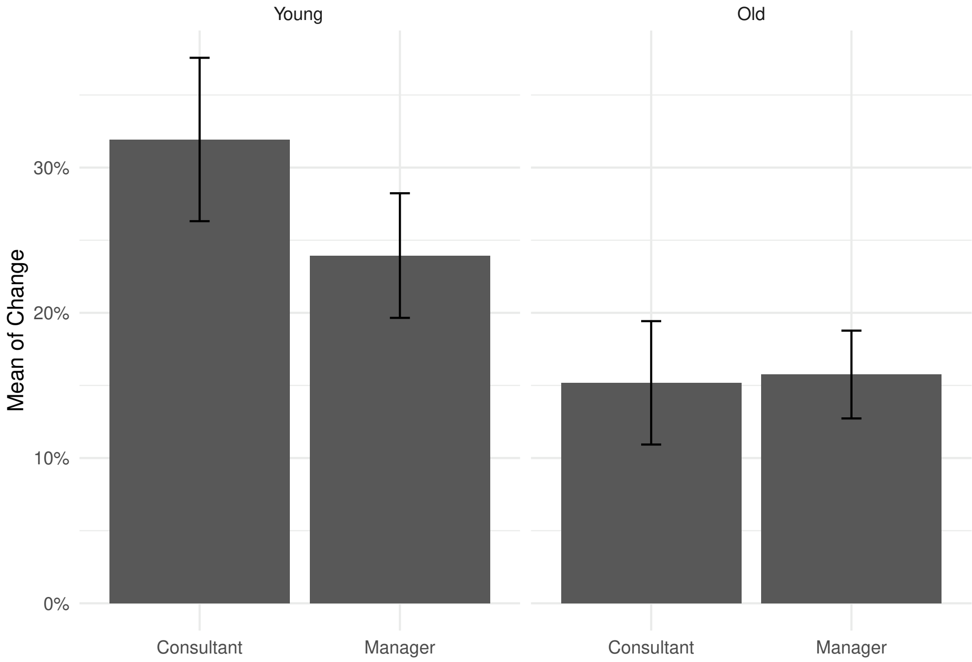

Figure 4.2: This figure displays separately for the *consultant* and *manager condition*s the likelihood of changing an investment option for those who are below the median age of 34 years (young) or above (old). The variable change takes on a value of 1 if the participant starts to invest in a new option and 0 otherwise, meaning that its average represents the propensity to change investments. Error bars depict the robust standard errors clustered on the participant level.

manager condition (23.9% of the rounds, SE: 1.6) than in the *consultant condition* (31.9% of the rounds, SE: 2.1).

To confirm the results visualized in Figure 4.2 and to test for treatment differences as well as control for gender, we conducted probit regressions explaining the likelihood of switching investment options by age. Table 4.4 reports the results of this investigation: Model 1 looks at those with an age of 34 or younger, where we find a significantly higher likelihood of changing the investment option after a loss in the *consultant condition*. Model 2, in contrast, looks at those with an age above 34, where we do not find any notable differences between the two treatment groups. In model 3 we pool the data and add a three-way interaction of prior loss, *consultant condition* dummy, and age in years. Here, the higher likelihood of a change due to prior losses in the low-involvement treatment (*consultant condition*) remains but decreases strongly with age. As Model 1 in Table C.1.7 shows, similar

Table 4.4: The Effect of Age on the Treatment Differences

Variable	(1) Young	(2) Old	(3) Pooled	(4) Pooled (Age)
Loss $(t-1)$	0.0650	-0.0463	0.0276	0.00960
	(0.133)	(0.183)	(0.109)	(0.111)
Consultant	-0.111	-0.0526	-0.149	-0.553
	(0.192)	(0.229)	(0.182)	(0.505)
Consultant*Loss $(t-1)$	0.627***	0.120	0.669***	1.520***
	(0.230)	(0.291)	(0.217)	(0.578)
Old			-0.367**	
			(0.152)	
Consultant*Old			0.166	
			(0.252)	
Consultant*Loss $(t-1)$*Old			-0.630**	
			(0.294)	
Age				-0.00392
				(0.00892)
Consultant*Age				0.0112
				(0.0135)
Consultant*Loss $(t-1)$*Age				-0.0292**
				(0.0148)
Controls	Yes	Yes	Yes	Yes
Observations	1,209	1,072	2,281	2,281
Pseudo R^2	0.0506	0.00124	0.0466	0.0228
Log-likelihood	-661.3	-452.9	-1,118	-1,146

Note. Probit regression with robust standard errors in parentheses clustered on the participant level. Estimates for the constant parameter and control variables are omitted but were included in the regression. One-sided t statistics, **$p < .05$; ***$p < .01$. *Outcome variable: Change*, dummy variable taking on the value of 1 if the investment option in the current round is different from in the previous round, 0 otherwise. *Independent variables: Loss $(t-1)$*, dummy variable taking on the value of 1 if the result of the previous round was a loss, 0 otherwise; *consultant*, dummy variable taking on the value of 1 if the participant is in the *consultant condition*, 0 otherwise; *old*, dummy variable taking on the value of 1 if the age of the participant is above 34 years (the median age of the sample), 0 otherwise; *age*, measured in years. *Control variable: Gender.*

results are obtained when losses are operationalized as the proportion of losses from all the experiences with a particular investment option.

Result 3: *We observe significantly fewer changes in the investment options for older participants.*

Gender & Education are factors that are regularly associated with individual differences in risk-taking (Frey et al., 2021). To rule out that they drive the effect, we conducted an exploratory analysis for age and educational effects in Appendix C.1. We observe significantly fewer changes in the investment options for female participants. We also find female participants to change investments significantly more often due to prior losses in the consultant compared to the *manager condition*, an effect which we do not find for male participants. However, these gender differences with respect to treatment effects are not statistically significant. We also do not observe differences in the likelihood to change the investment options for participants who obtained a tertiary degree or have a business background in their education.

Summarizing, one can say that we robustly observe a higher likelihood of a change due to a loss in the consultant compared to the *manager condition*. Moreover, age seems to be the only individual attribute that consistently and significantly reduces the treatment effect. As our participant pool consists of finance professionals, age can be interpreted as a proxy for experiences in the job, which brings with it the knowledge and strategies for making better decisions.

4.4 Discussion

We investigate differences in the likelihood of changing an investment project due to prior gains or losses as a function of the perceived degree of involvement of the decision maker. Moreover, we inquire whether age as a measure for professional experiences mitigates observed differences between the treatment groups. We conducted a lab-in-the-field online experiment with 167 employees of the finance and controlling department in one of the largest infrastructure companies in Switzerland. One main objective of this manuscript is to conceptually replicate Chang et al. (2016) with a professional sample in a lab-in-the-field study. We therefore followed the initial study design by Chang et al. (2016) and altered the labels of the investment options in two between-participant conditions, one condition inducing high personal involvement, *manager condition*, and one with more perceived distance to the decision, *consultant condition*. Only under low involvement (*consultant condition*) do we find a significantly higher likelihood of switching to another investment after prior losses compared to prior gains. However, these effects of prior outcomes are limited to participants at or below the median age of 34 years, as we do not find such effects in the choices of those above the age of 34. Further, higher education mitigates the effect of prior outcomes, but a business background in education or training does not seem to make a significant difference.

Limitations. One limitation to the cover story approach might be that the framing has different meanings for different participants and is prone to individual interpretation. While we can not fully rule out that such differences in interpretations may have played a role in the observed treatment effect, the instructions introduced both treatment groups to the true return generating mechanism behind the returns. The instruction text about this mechanism also did not differ between conditions and did not personify the mechanism by including the cover story in its description. Nevertheless, a better understanding of the driving factors behind these effects is a relevant and promising endeavor for future research, as the effects can be found in the laboratory, as well as in the field for lay investors and professionals. To us the personal involvement mechanism appears to be the most plausible explanation. It employs very few assumptions and is in line with other research that studies personal involvement Heinke et al. (2020); Rotaru et al. (2021) and we replicate the findings of Chang et al. (2016) who similarly interpret their results in terms of a blaming and involvement mechanism.

One aspect that did differ between the treatment groups was their size which was due to the truly random allocation of participants. However, we do not believe this imbalance to be detrimental to our findings. In detail, while participants in the *manager condition* were significantly younger than those in the *consultant condition* we do observe more change of investments after a loss in the latter. We further find that younger participants are more likely to change their investments after a loss. We therefore find an increased likelihood to change in the *consultant condition, despite* this group being older, and thus more likely *not* to change investments after a loss.

Further aspect that might impact the outcome are the statistical power of the experiment as well as the incentivization. While the sample size and thereby the statistical power was restricted by the professional sample, our sample does allow for adequate power of .8 for medium effects of size $d = .39$. Further, a low power to detect small effects does not impact the interpretation of the significant differences we find throughout the experiment, as power calculation are concerned with the probability of missing an existent effect. Regarding the incentives, first, the tournament character of only the top five participants winning a book coupon, might explain why we observe a small share of investments in the least risky option. However, this incentive structure affects both treatment conditions in the same way, and thus disappears when comparing both treatments. Second, the incentives were small, in particular for higher ranked professionals working in Switzerland. The literature is, however, rather inconclusive whether low-stake/hypothetical decisions only add noise or systematically bias decisions (C. F. Camerer & Hogarth, 1999; Holt & Laury, 2002; Kühberger et al., 2002; Harrison et al., 2005; Gneezy et al., 2015; C. Camerer & Mobbs, 2017; Hackethal et al., 2022). While noisier decisions would make it harder to detect significant effects, a decision bias would also affect all participants in a similar direction and should therefore cancel out when comparing decisions between condition groups.

Lastly, it is interesting that participants mostly chose the option which included a zero return as an outcome, as this return is strictly speaking not negative. Nonetheless, this aspect also does not impact our conclusions as the outcomes of each investment option also did not differ between conditions.

Relation to the Literature. We find a lower commitment to an investment project in the low-involvement treatment (*consultant condition*). This can be interpreted as younger participants taking advantage of the possibility to use the cover-story of hired consultant as a hypothetical scapegoat who can be blamed for the negative outcome from a random process. These results contribute to the literature in two aspects:

First, we conceptually replicate Chang et al. (2016) by experimentally manipulating the perceived degree of involvement in a decision. We observe a higher likelihood of stopping a losing investment when the investment option is labeled as hiring a consultant (low involvement). In the high-involvement treatment (*manager condition*), we do not observe a difference in the likelihood of switching investment options after either a gain or a loss. Presumably, we do not observe undercommitment in the high-involvement treatment because of the lack of "someone (else) to blame" after losses. Moreover, our findings on the role of prior losses and personal involvement are in line with those by Martens & Orzen (2021), even though the studies differ in the subject pool (see next point), the frequency of decisions made, and how the degree of involvement is manipulated. This speaks in favor of personal involvement in prior decisions having a robust impact on decision making.

Second, we document that these behavioral differences exist also among professionals but disappear with age and level of education. As with increasing age one often climbs the career ladder and gains more experience, we interpret the observed age effect as a proxy for professional seniority. However, it might also be a pure and general effect of aging and gaining experiences through life, that drives these results. While we cannot separate this with our data, our interpretation is in line with the broader literature, such as Malmendier (2018), who points out that a plethora of documented biases in professional decision making could be the result of early career errors. If the outcome of profit-harming decisions is on average negative (i.e. more often negative than positive), such biases in decisions might disappear with feedback and experience, as predicted by Erev & Roth (2014).

Also regarding the frequency of outcomes, Cohen & Erev (2021) observe that if gains appear more (less) often than losses decision makers prolong their projects for too long (short) in order to break even. For this reason we expect frequent changes of options in our experiment, as only two out of six possible outcomes are positive, despite the lotteries having a positive expected value. It might prove interesting to investigate further whether and how the involvement in a decision interacts with the frequency of observed gains or losses and subsequently over- or under-commitment of resources to a project.

Lastly, while we focus our analysis on the propensity to switch the investment, one may also ask why investors switch between options in our settings at all. All information is presented to participants in the instructions, yet their experience throughout the experiment still influences their decision. While this behavior is at odds with what one would expect form a rational investor, it is in line with results reported by Heinke et al. (2022) who find that experiences, may even overrule provided information. In line with our finding of the treatment effect vanishing with age, Dessaint & Matray (2017) show that in field data experienced managers "overreact" less to close-by hurricanes, even though here as well the statistical probabilities of a hurricane incident occurring are known. Also note that participants not having a "rational" reason to switch investments does not diminish our results as this was the case for both treatment conditions equally.

In conclusion, we observe that the likelihood of changing an investment project due to prior gains or losses is a function of the perceived degree of involvement of the decision maker. The effect of involvement as well as treating gains and losses differently, disappears with age, which can be seen as a measure for professional experience. One straightforward implication of our findings is the importance of feedback to enhance learning from experience and reduce biases in decisions. However, feedback in itself might also increase biases if profit-harming decisions are on average rewarding (Cohen et al., 2020). Thus, ideally the profit-maximizing option is the one with the highest frequency of positive rewards (Erev & Roth, 2014). While such considerations in designing a decision environment may be important, our results also suggest that general experience of the decision maker also plays a central role, even more so than specific economic and business training. Being able to separate the personal involvement into a decision from evaluating the outcome is crucial and seems to be a skill acquired over the life- and career span. Thus, to improve decision-making quality along a career path, enhancing reflection on experiences and thereby learning presents itself as a promising approach.

References

Abdellaoui, M., Bleichrodt, H., & Kammoun, H. (2013). Do financial professionals behave according to prospect theory? An experimental study. *Theory and Decision*, *74*(3), 411–429. https://doi.org/10.1007/s11238-011-9282-3

Adam, K., Matveev, D., & Nagel, S. (2020). Do survey expectations of stock returns reflect risk adjustments? *Journal of Monetary Economics*. https://doi.org/10.1016/j.jmoneco.2020.04.010

Akepanidtaworn, K., Di Mascio, R., Imas, A., & Schmidt, L. (2021). *Selling fast and buying slow.* (NBER Working Paper) https://doi.org/10.3386/w29076

Albrecht, R., Jenny, M. A., Nilsson, H., & Rieskamp, J. (2021). The similarity-updating model of probability judgment and belief revision. *Psychological Review*, *128*, 1088–1111. https://doi.org/10.1037/rev0000299

Alevy, J. E., Haigh, M. S., & List, J. A. (2006). Information cascades: Evidence from a field experiment with financial market professionals. *Journal of Finance*, *62*(1), 151-180. https://doi.org/10.1111/j.1540-6261.2007.01204.x

Amelio, A., & Zimmermann, F. (2023). Motivated memory in economics—a review. *Games*, *14*(1), 15. https://doi.org/10.3390/g14010015

Andersen, S., Hanspal, T., & Nielsen, K. M. (2019). Once bitten, twice shy: The power of personal experiences in risk taking. *Journal of Financial Economics*, *132*(3), 97–117. https://doi.org/10.1016/j.jfineco.2018.10.018

Andraszewicz, S., Scheibehenne, B., Rieskamp, J., Grasman, R., Verhagen, J., & Wagenmakers, E.-J. (2015). An introduction to Bayesian hypothesis testing for management research. *Journal of Management*, *41*(2), 521–543. https://doi.org/10.1177/0149206314560412

Arkes, H. R., & Ayton, P. (1999). The sunk cost and Concorde effects: Are humans less rational than lower animals? *Psychological Bulletin*, *125*(5), 591-600. https://doi.org/10.1037/0033-2909.125.5.591

Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, *35*(1), 124–140. https://doi.org/10.1016/0749-5978(85)90049-4

Asparouhova, E., Hertzel, M., & Lemmon, M. (2009). Inference from streaks in random outcomes: Experimental evidence on beliefs in regime shifting and the law of small numbers. *Management Science*, *55*(11), 1766–1782. https://doi.org/10.1287/mnsc.1090.1059

Barberis, N., & Xiong, W. (2012). Realization utility. *Journal of Financial Economics*, *104*, 251–271. https://doi.org/10.1016/j.jfineco.2011.10.005

Beike, D. R., Markman, K. D., & Karadogan, F. (2009). What we regret most are lost opportunities: A theory of regret intensity. *Personality and Social Psychology Bulletin*, *35*(3), 385-397. https://doi.org/10.1177/0146167208328329

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, *30*(5), 961–981. https://doi.org/10.1287/opre.30.5.961

Bénabou, R. (2015). The economics of motivated beliefs. *Revue d'économie politique*(5), 665–685. https://doi.org/10.3917/redp.255.0665

Bénabou, R., & Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, *30*(3), 141–164. https://doi.org/10.1257/jep.30.3.141

Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *The quarterly journal of Economics*, *110*(1), 73–92. https://doi.org/10.2307/2118511

Brockner, J., Shaw, M. C., & Rubin, J. Z. (1979). Factors affecting withdrawal from an escalating conflict: Quitting before it's too late. *Journal of Experimental Social Psychology*, *15*(5), 492–503. https://doi.org/10.1016/0022-1031(79)90011-8

Brunnermeier, M. K., & Parker, J. A. (2005). Optimal expectations. *American Economic Review*, *95*(4), 1092–1118. https://doi.org/10.1257/0002828054825493

Bucher-Koenen, T., & Ziegelmeyer, M. (2014). Once burned, twice shy? Financial literacy and wealth losses during the financial crisis. *Review of Finance*, *18*(6), 2215–2246. https://doi.org/10.1093/rof/rft052

Busemeyer, J. R., Barkan, R., Mehta, S., & Chaturvedi, A. (2007). Context effects and models of preferential choice: implications for consumer behavior. *Marketing Theory*, *7*(1), 39–58. https://doi.org/10.1177/1470593107073844

Camerer, C., & Mobbs, D. (2017). Differences in behavior and brain activity during hypothetical and real choices. *Trends in Cognitive Sciences*, *21*(1), 46–56. https://doi.org/10.1016/j.tics.2016.11.001

Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, *19*(1), 7–42. https://doi.org/10.1023/A:1007850605129

Camerer, C. F., & Malmendier, U. (2007). Behavioral economics of organisations. In P. Diamond & H. Vartiainen (Eds.), *Behavioral Economics and its Applications* (p. 235-290). Princeton University Press.

Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, *50*(2), 317–372. https://doi.org/10.3368/jhr.50.2.317

Chang, T. Y., Solomon, D. H., & Westerfield, M. M. (2016). Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect. *The Journal of Finance*, *71*(1), 267–302. https://doi.org/10.1111/jofi.12311

Chen, D. L., Schonger, M., & Wickens, C. (2016). otree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97. https://doi.org/10.1016/j.jbef.2015.12.001

Choi, J. J., Laibson, D., Madrian, B. C., & Metrick, A. (2009). Reinforcement learning and savings behavior. *The Journal of Finance*, *64*(6), 2515–2534. https://doi.org/10.1111/j.1540-6261.2009.01509.x

Civelli, A., & Deck, C. (2018). A flexible and customizable method for assessing cognitive abilities. *Review of Behavioral Economics*, *5*(2), 123-147. http://dx.doi.org/10.1561/105.00000081

Cohen, D., & Erev, I. (2021). Over and under commitment to a course of action in decisions from experience. *Journal of Experimental Psychology: General*, *150*(12), 2455. https://doi.org/10.1037/xge0001066

Cohen, D., Plonsky, O., & Erev, I. (2020). On the impact of experience on probability weighting in decisions under risk. *Decision*, *7*(2), 153-162. https://doi.org/10.1037/dec0000118

Corgnet, B., Cornand, C., & Hanaki, N. (2020). *Tail events, emotions and risk taking.* (SSRN) http://dx.doi.org/10.2139/ssrn.3606079

Dawkins, R., & Carlisle, T. R. (1976). Parental investment, mate desertion and a fallacy. *Nature*, *262*(5564), 131–133. https://doi.org/10.1038/262131a0

Dessaint, O., & Matray, A. (2017). Do managers overreact to salient risks? evidence from hurricane strikes. *Journal of Financial Economics*, *126*(1), 97–121. https://doi.org/10.1016/j.jfineco.2017.07.002

Drobner, C. (2022). Motivated beliefs and anticipation of uncertainty resolution. *American Economic Review: Insights*, *4*(1), 89–105. https://doi.org/10.1257/aeri.20200829

Eldar, E., Roth, C., Dayan, P., & Dolan, R. J. (2018). Decodability of reward learning signals predicts mood fluctuations. *Current Biology*, *28*(9), 1433-1439.e7. https://doi.org/10.1016/j.cub.2018.03.038

Enke, B., Schwerter, F., & Zimmermann, F. (2020). *Associative memory and belief formation.* (NBER Working Paper) https://doi.org/10.3386/w26664

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, *88*(4), 848–881. http://www.jstor.org/stable/117009

Erev, I., & Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(Suppl 3), 10818–10825. https://doi.org/10.1073/pnas.1402846111

Ertac, S. (2011). Does self-relevance affect information processing? experimental evidence on the response to performance and non-performance feedback. *Journal of Economic Behavior & Organization*, *80*(3), 532–545. https://doi.org/10.1016/j.jebo.2011.05.012

Farrell, S., & Lewandowsky, S. (2018). *Computational modeling of cognition and behavior*. Cambridge University Press.

Fellner, G., & Maciejovsky, B. (2007). Risk attitude and market behavior: Evidence from experimental asset markets. *Journal of Economic Psychology*, *28*(3), 338–350. https://doi.org/10.1016/j.joep.2007.01.006

Festinger, L. (1962). *A theory of cognitive dissonance*. Stanford University Press.

Fiedler, K. (2007). Construal level theory as an integrative framework for behavioral decision-making research and consumer psychology. *Journal of Consumer Psychology*, *17*(2), 101–106. https://doi.org/10.1016/S1057-7408(07)70015-3

Filippin, A., & Crosetto, P. (2016). A reconsideration of gender differences in risk attitudes. *Management Science*, *62*(11), 3138–3160. https://doi.org/10.1287/mnsc.2015.2294

Fischbacher, U., Hoffmann, G., & Schudy, S. (2017). The causal effect of stop-loss and take-gain orders on the disposition effect. *The Review of Financial Studies*, *30*(6), 2110–2129. https://doi.org/10.1093/rfs/hhx016

Fischer, D. S., & Maier, J. K. (2019). *Decomposing the disposition effect*. (Unpublished)

Fontanesi, L., Palminteri, S., & Lebreton, M. (2019). Decomposing the effects of context valence and feedback information on speed and accuracy during reinforcement learning: a meta-analytical approach using diffusion decision modeling. *Cognitive, Affective & Behavioral Neuroscience*, *19*(3), 490–502. https://doi.org/10.3758/s13415-019-00723-1

Frazzini, A. (2006). The disposition effect and underreaction to news. *The Journal of Finance*, *61*(4), 2017–2046. https://doi.org/10.1111/j.1540-6261.2006.00896.x

Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. (2017). Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, *3*(10), e1701381. https://doi.org/10.1126/sciadv.1701381

Frey, R., Richter, D., Schupp, J., Hertwig, R., & Mata, R. (2021). Identifying robust correlates of risk preference: A systematic approach using specification curve analysis. *Journal of Personality and Social Psychology*, *120*(2), 538–557. https://doi.org/10.1037/pspp0000287

Frydman, C., Barberis, N., Camerer, C., Bossaerts, P., & Rangel, A. (2014). Using neural data to test a theory of investor behavior: An application to realization utility. *The Journal of Finance*, *69*(2), 907–946. https://doi.org/10.1111/jofi.12126

Frydman, C., & Camerer, C. (2016). Neural evidence of regret and its implications for investor behavior. *The Review of Financial Studies*, *29*(11), 3108–3139. https://doi.org/10.1093/rfs/hhw010

Frydman, C., & Rangel, A. (2014). Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price. *Journal of Economic Behavior & Organization*, *107*, 541–552. https://doi.org/10.1016/j.jebo.2014.01.017

Genesove, D., & Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market. *The Quarterly Journal of Economics*, *116*(4), 1233–1260. https://doi.org/10.1162/003355301753265561

Gershman, S. J. (2015). Do learning rates adapt to the distribution of rewards? *Psychonomic Bulletin & Review*, *22*(5), 1320–1327. https://doi.org/10.3758/s13423-014-0790-3

Giglio, S., Maggiori, M., Stroebel, J., & Utkus, S. (2021). Five facts about beliefs and portfolios. *American Economic Review*, *111*(5), 1481–1522. https://doi.org/10.1257/aer.20200243

Gneezy, U., & Imas, A. (2017). Lab in the field: Measuring preferences in the wild. In A. V. Baneriee & E. Duflo (Eds.), *Handbook of economic field experiments* (Vol. 1, pp. 439–464). Elsevier. https://doi.org/10.1016/bs.hefe.2016.08.003

Gneezy, U., Imas, A., & List, J. (2015). *Estimating individual ambiguity aversion: A simple approach.* (NBER Working Paper)

Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, *112*(2), 631–645. https://doi.org/10.1162/003355397555217

Gödker, K., Jiao, P., & Smeets, P. (2019). *Investor memory.* (SSRN) http://dx.doi.org/10.2139/ssrn.3348315

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: integrating sampling and repeated decisions from experience. *Psychological Review*, *118*(4), 523. https://doi.org/10.1037/a0024558

Greenwood, R., & Shleifer, A. (2014). Expectations of returns and expected returns. *Review of Financial Studies*, *27*(3), 714–746. https://doi.org/10.1093/rfs/hht082

Griffin, J. M., Harris, J. H., & Topaloglu, S. (2003). The dynamics of institutional and individual trading. *The Journal of Finance*, *58*(6), 2285–2320. https://doi.org/10.1046/j.1540-6261.2003.00606.x

Grinblatt, M., Keloharju, M., & Linnainmaa, J. T. (2012). IQ, trading behavior, and performance. *Journal of Financial Economics*, *104*(2), 339–362. https://doi.org/10.1016/j.jfineco.2011.05.016

Grosshans, D., Langnickel, F., & Zeisberger, S. (2020). *Is buying more forward-looking than selling? The role of beliefs in investment decisions.* (SSRN) http://dx.doi.org/10.2139/ssrn.2972112

Guiso, L., Sapienza, P., & Zingales, L. (2018). Time varying risk aversion. *Journal of Financial Economics*, *128*(3), 403–421. https://doi.org/10.1016/j.jfineco.2018.02.007

Hackethal, A., Kirchler, M., Laudenbach, C., Razen, M., & Weber, A. (2022). On the role of monetary incentives in risk preference elicitation experiments. *Journal of Risk and Uncertainty*, *66*, 1–25. https://doi.org/10.1007/s11166-022-09377-w

Haigh, M. S., & List, J. A. (2005). Do professional traders exhibit myopic loss aversion? An experimental analysis. *The Journal of Finance*, *60*(1), 523–534. https://doi.org/10.1111/j.1540-6261.2005.00737.x

Hanspal, T. (2017). *Essays in household finance.* (Publication No. 07.2017) [Doctoral Dissertation, Copenhagen Business School]. https://hdl.handle.net/10398/9457

Harrison, G. W., Johnson, E., McInnes, M. M., & Rutström, E. E. (2005). Risk aversion and incentive effects: Comment. *American Economic Review*, *95*(3), 897–901. https://doi.org/10.1257/0002828054201378

Hartzmark, S. M., Hirshman, S. D., & Imas, A. (2021). Ownership, learning, and beliefs. *The Quarterly Journal of Economics*, *136*(3), 1665–1717. https://doi.org/10.1093/qje/qjab010

Heath, C. (1995). Escalation and de-escalation of commitment in response to sunk costs: The role of budgeting in mental accounting. *Organizational Behavior and Human Decision Processes*, *62*(1), 38–54. https://doi.org/10.1006/obhd.1995.1029

Hefti, A., Heinke, S., & Schneider, F. (2018). *Mental capabilities, heterogeneous trading patterns and performance in an experimental asset market.* (SSRN) http://dx.doi.org/10.2139/ssrn.2832767

Heinke, S., Leuenberger, A., & Rieskamp, J. (2020). *This time is different: On similarity and risk taking after experienced gains and losses.* (SSRN) http://dx.doi.org/10.2139/ssrn.3691829

Heinke, S., Olschewski, S., & Rieskamp, J. (2022). *Experiences and price dynamics.* (SSRN) https://dx.doi.org/10.2139/ssrn.4279001

Heinke, S., Trutmann, K., & Rudin, C. (2021). *Degree of involvement in decisions and the likelihood to stop investing among professionals.* (SSRN) https://dx.doi.org/10.2139/ssrn.4266346

Hens, T., & Vlcek, M. (2011). Does prospect theory explain the disposition effect? *Journal of Behavioral Finance*, *12*(3), 141–157. https://doi.org/10.1080/15427560.2011.601976

Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: why small samples? *Cognition*, *115*(2), 225–237. https://doi.org/10.1016/j.cognition.2009.12.009

Hoelzl, E., & Loewenstein, G. (2005). Wearing out your shoes to prevent someone else from stepping into them: Anticipated regret and social takeover in sequential decisions. *Organizational Behavior and Human Decision Processes*, *98*(1), 15–27. https://doi.org/10.1016/j.obhdp.2005.04.004

Hoffmann, A. O., & Post, T. (2017). How return and risk experiences shape investor beliefs and preferences. *Accounting & Finance*, *57*(3), 759–788. https://doi.org/10.1111/acfi.12169

Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, *24*(1), 1–55. https://doi.org/10.1016/0010-0285(92)90002-J

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644–1655. https://doi.org/10.1257/000282802762024700

Imas, A., Kuhn, M. A., & Mironova, V. (2022). Waiting to choose: The role of deliberation in intertemporal choice. *American Economic Journal: Microeconomics*, *14*(3), 414–40. https://doi.org/10.1257/mic.20180233

Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, *48*(1), 65–91. https://doi.org/10.1111/j.1540-6261.1993.tb04702.x

Jiao, P. (2017). Belief in mean reversion and the disposition effect: An experimental test. *Journal of Behavioral Finance*, *18*(1), 29–44. https://doi.org/10.1080/15427560.2017.1274754

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773-795. https://doi.org/10.1080/01621459.1995.10476572

Kaustia, M., & Knüpfer, S. (2008). Do investors overweight personal experience? evidence from ipo subscriptions. *The Journal of Finance*, *63*(6), 2679–2702. https://doi.org/10.1111/j.1540-6261.2008.01411.x

Kieren, P., Müller-Detthard, J., & Weber, M. (2022). *Can agents add and subtract when forming beliefs.* (SSRN) http://dx.doi.org/10.2139/ssrn.3644226

Knüpfer, S., Rantapuska, E., & Sarvimäki, M. (2017). Formative experiences and portfolio choice: Evidence from the finnish great depression. *The Journal of Finance*, *72*(1), 133–166. https://doi.org/10.1111/jofi.12469

Knutson, B., & Bossaerts, P. (2007). Neural antecedents of financial decisions. *The Journal of Neuroscience*, *27*(31), 8174–8177. https://doi.org/10.1523/JNEUROSCI.1564-07.2007

Knutson, B., Samanez-Larkin, G. R., & Kuhnen, C. M. (2011). Gain and loss learning differentially contribute to life financial outcomes. *PLoS One*, *6*(9), Article e24390. https://doi.org/10.1371/journal.pone.0024390

Kühberger, A., Schulte-Mecklenbeck, M., & Perner, J. (2002). Framing decisions: Hypothetical and real. *Organizational Behavior and Human Decision Processes*, *89*(2), 1162–1175. https://doi.org/10.1016/S0749-5978(02)00021-3

Kuhnen, C. M. (2015). Asymmetric learning from financial information. *The Journal of Finance*, *70*(5), 2029–2062. https://doi.org/10.1111/jofi.12223

Kuhnen, C. M., & Knutson, B. (2011). The influence of affect on beliefs, preferences, and financial decisions. *Journal of Finance and Quantitative Analysis*, *46*(3), 605–626. https://doi.org/10.1017/S0022109011000123

Kuhnen, C. M., Rudorf, S., & Weber, B. (2017). *The effect of prior choices on expectations and subsequent portfolio decisions.* (NBER Working Paper) https://doi.org/10.3386/w23438

Larson, F., List, J. A., & Metcalfe, R. D. (2016). *Can myopic loss aversion explain the equity premium puzzle? evidence from a natural field experiment with professional traders.* (NBER Working Paper) https://doi.org/10.3386/w22605

Lee, S., Gold, J. I., & Kable, J. W. (2020). The human as delta-rule learner. *Decision*, *7*(1), 55-66. https://doi.org/10.1037/dec0000112

Lehenkari, M. (2011). In search of the underlying mechanism of the disposition effect. *Journal of Behavioral Decision Making*, *25*(2), 196–209. https://doi.org/10.1002/bdm.727

Lejuez, C. W., Read, J. P., Kahler, C. W., Richards, J. B., Ramsey, S. E., Stuart, G. L., ... Brown, R. A. (2002). Evaluation of a behavioral measure of risk taking: the balloon analogue risk task (BART). *Journal of Experimental Psychology: Applied*, *8*(2), 75-84. https://doi.org/10.1037/1076-898X.8.2.75

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, *92*(368), 805–824. https://doi.org/10.2307/2232669

Malmendier, U. (2018). Behavioral corporate finance. In B. D. Bernheim, S. DellaVigna, & D. Laibson (Eds.), *Handbook of behavioral economics: Applications and foundations 1* (Vol. 1, pp. 277–379). Elsevier. https://doi.org/10.1016/bs.hesbe.2018.08.001

Malmendier, U., & Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, *126*(1), 373–416. https://doi.org/doi/10.1093/qje/qjq004

Malmendier, U., & Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, *131*(1), 53–87. https://doi.org/10.1093/qje/qjv037

Malmendier, U., Pouzo, D., & Vanasco, V. (2020). Investor experiences and financial market dynamics. *Journal of Financial Economics*, *136*(3), 597–622. https://doi.org/10.1016/j.jfineco.2019.11.002

Malmendier, U., Tate, G., & Yan, J. (2011). Overconfidence and early-life experiences: The effect of managerial traits on corporate financial policies. *The Journal of Finance*, *66*(5), 1687–1733. https://doi.org/10.1111/j.1540-6261.2011.01685.x

Martens, N., & Orzen, H. (2021). Escalating commitment to a failing course of action—A re-examination. *European Economic Review*, *137*, Article 103811. https://doi.org/10.1016/j.euroecorev.2021.103811

McCain, B. E. (1986). Continuing investment under conditions of failure: A laboratory study of the limits to escalation. *Journal of Applied Psychology*, *71*(2), 280-284. https://doi.org/10.1037/0021-9010.71.2.280

McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in r and stan*. CRC Press.

McGrath, A. (2017). Dealing with dissonance: A review of cognitive dissonance reduction. *Social and Personality Psychology Compass*, *11*(12), e12362. https://doi.org/10.1111/spc3.12362

Meng, X.-L., & Wong, W. H. (1996). Simulating ratios of normalizing constants via a simple identity: a theoretical exploration. *Statistica Sinica*, *6*(4), 831–860. https://www.jstor.org/stable/24306045

Mosenhauer, M. (2020). *Information management against excessive stock trading: More or less? or both?* (EconStor Working Paper)

Murawski, C., & Bossaerts, P. (2016). How humans solve complex problems: The case of the knapsack problem. *Scientific reports*, *6*(1), Article 34851. https://doi.org/10.1038/srep34851

Navarro, A. D., & Fantino, E. (2005). The sunk cost effect in pigeons and humans. *Journal of the Experimental Analysis of Behavior*, *83*(1), 1–13. https://doi.org/10.1901/jeab.2005.21-04

Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance*, *53*(5), 1775–1798. https://doi.org/10.1111/0022-1082.00072

Odean, T. (1999). Do investors trade too much? *American Economic Review*, *89*(5), 1279–1298. https://doi.org/10.1257/aer.89.5.1279

Olschewski, S., Diao, L., & Rieskamp, J. (2021). Reinforcement learning about asset variability and correlation in repeated portfolio decisions. *Journal of Behavioral and Experimental Finance*, *32*, Article 100559. https://doi.org/10.1016/j.jbef.2021.100559

Oskarsson, A. T., Van Boven, L., McClelland, G. H., & Hastie, R. (2009). What's next? judging sequences of binary events. *Psychological Bulletin*, *135*(2), 262. https://psycnet.apa.org/doi/10.1037/a0014821

R Core Team. (2021). R: A language and environment for statistical computing [Computer software manual]. https://www.R-project.org/

Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, *9*(7), 545–556. https://doi.org/10.1038/nrn2357

Richter, D., Rohrer, J., Metzing, M., Nestler, W., Weinhardt, M., & Schupp, J. (2017). *SOEP scales manual: (updated for SOEP-Core v32.1)*. DIV/SOEP. https://www.diw.de/documents/publikationen/73/diw_01.c.571151.de/diw_ssp0423.pdf

Rotaru, K., Kalev, P. S., Yadav, N., & Bossaerts, P. (2021). Transferring cognitive talent across domains to reduce the disposition effect in investment. *Scientific Reports*, *11*(1), 23068. https://doi.org/10.1038/s41598-021-02596-2

Saccardo, S., & Serra-Garcia, M. (2023). Enabling or limiting cognitive flexibility? evidence of demand for moral commitment. *American Economic Review*, *113*(2), 396–429. https://doi.org/10.1257/aer.20201333

Schultz, W. (2015). Neuronal reward and decision signals: From theories to data. *Physiological Reviews*, *95*(3), 853–951. https://doi.org/10.1152/physrev.00023.2014

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599. https://doi.org/10.1126/science.275.5306.1593

Schwaiger, R., Kirchler, M., Lindner, F., & Weitzel, U. (2020). Determinants of investor expectations and satisfaction. A study with financial professionals. *Journal of Economic Dynamics and Control*, *110*, Article 103675. https://doi.org/10.1016/j.jedc.2019.03.002

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the "secretary problem". *Organizational Behavior and Human Decision Processes*, *69*(3), 221–236. https://doi.org/10.1006/obhd.1997.2683

Seligman, M. E. (1972). Learned helplessness. *Annual Review of Medicine*, *23*(1), 407–412. https://doi.org/10.1146/annurev.me.23.020172.002203

Seru, A., Shumway, T., & Stoffman, N. (2010). Learning by trading. *The Review of Financial Studies*, *23*(2), 705–739. https://doi.org/10.1093/rfs/hhp060

Seymour, B., Daw, N., Dayan, P., Singer, T., & Dolan, R. (2007). Differential encoding of losses and gains in the human striatum. *Journal of Neuroscience*, *27*(18), 4826–4831. https://doi.org/10.1523/JNEUROSCI.0400-07.2007

Sharot, T., & Garrett, N. (2016). Forming beliefs: Why valence matters. *Trends in Cognitive Sciences*, *20*(1), 25–33. https://doi.org/10.1016/j.tics.2015.11.002

Shefrin, H., & Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance*, *40*(3), 777–790. https://doi.org/10.1111/j.1540-6261.1985.tb05002.x

Simon, H. A. (1990). Bounded rationality. In J. Eatwell, M. Milgate, & P. Newman (Eds.), *Utility and probability* (pp. 15–18). London: Palgrave Macmillan UK. 10.1007/978-1-349-20568-4_5

Spektor, M. S., & Kellen, D. (2018). The relative merit of empirical priors in non-identifiable and sloppy models: Applications to models of learning and decision-making. *Psychonomic Bulletin & Review*, *25*(6), 2047–2068. https://doi.org/10.3758/s13423-018-1446-5

Stan Development Team. (2020). *RStan: the R interface to Stan.* (R package version 2.21.2). http://mc-stan.org/

Staw, B. M. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, *16*(1), 27–44. https://doi.org/10.1016/0030-5073(76)90005-2

Staw, B. M. (1981). The escalation of commitment to a course of action. *Academy of Management Review*, *6*(4), 577–587. https://doi.org/10.2307/257636

Strahilevitz, M. A., Odean, T., & Barber, B. M. (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. https://doi.org/10.1509/jmkr.48.SPL.S102

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction.* MIT Press.

Teodorescu, K., & Erev, I. (2014). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science*, *25*(10), 1861–1869. https://doi.org/10.1177/0956797614543022

Thaler, R. H. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior & Organization*, *1*(1), 39–60. https://doi.org/10.1016/0167-2681(80)90051-7

Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, *12*(3), 183–206. https://doi.org/10.1002/(SICI)1099-0771(199909)12:3<183::AID-BDM318>3.0.CO;2-F

Thaler, R. H., Tversky, A., Kahneman, D., & Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, *112*(2), 647–661. https://doi.org/10.1162/003355397555226

Trautmann, S. T., & van de Kuilen, G. (2015). Belief elicitation: A horse race among truth serums. *The Economic Journal*, *125*(589), 2116–2135. https://doi.org/10.1111/ecoj.12160

Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review*, *110*(3), 403. https://doi.org/10.1037/0033-295X.110.3.403

Trutmann, K., Heinke, S., & Rieskamp, J. (2022). *Belief updating and investment decisions: The impact of good or bad news varies with prior returns.* (SSRN) https://dx.doi.org/10.2139/ssrn.3935798

Trutmann, K., Heinke, S., & Rieskamp, J. (2023). Take your time: How delayed information and restricted decision opportunities improve belief formation in investment decisions. *Finance Research Letters*, *51*, 103442. https://doi.org/10.1016/j.frl.2022.103442

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *4*(5), 297–323. https://doi.org/10.1007/BF00122574

Unkelbach, C., Alves, H., & Koch, A. (2020). Negativity bias, positivity bias, and valence asymmetries: Explaining the differential processing of positive and negative information. In B. Gawronski (Ed.), *Advances in experimental social psychology* (p. 115-187). Academic Press. https://doi.org/10.1016/bs.aesp.2020.04.005

von den Driesch, T., Da Costa, M. E. S., Flatten, T. C., & Brettel, M. (2015). How ceo experience, personality, and network affect firms' dynamic capabilities. *European Management Journal*, *33*(4), 245–256. https://doi.org/10.1016/j.emj.2015.01.003

Wallsten, T. S., Pleskac, T. J., & Lejuez, C. W. (2005). Modeling behavior in a clinically diagnostic sequential risk-taking task. *Psychological Review*, *112*(4), 862-880. https://doi.org/10.1037/0033-295X.112.4.862

Weber, M., & Camerer, C. F. (1998). The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization*, *33*(2), 167–184. https://doi.org/10.1016/S0167-2681(97)00089-9

Weitzel, U., Huber, C., Huber, J., Kirchler, M., Lindner, F., & Rose, J. (2020). Bubbles and financial professionals. *The Review of Financial Studies*, *33*(6), 2659–2696. https://doi.org/10.1093/rfs/hhz093

Wimmer, G. E., Li, J. K., Gorgolewski, K. J., & Poldrack, R. A. (2018). Reward learning over weeks versus minutes increases the neural representation of value in the human brain. *Journal of Neuroscience*, *38*(35), 7649–7666. https://doi.org/10.1523/JNEUROSCI.0075-18.2018

Yakobi, O., Cohen, D., Naveh, E., & Erev, I. (2020). Reliance on small samples and the value of taxing reckless behaviors. *Judgment & Decision Making*, *15*(2), 266-281. https://doi.org/10.1017/S1930297500007403

Zeelenberg, M., & van Dijk, E. (1997). A reverse sunk cost effect in risky decision making: Sometimes we have too much invested to gamble. *Journal of Economic Psychology*, *18*(6), 677–691. https://doi.org/10.1016/S0167-4870(97)00029-9

Zikmund-Fisher, B. J. (2004). De-escalation after repeated negative feedback: Emergent expectations of failure. *Journal of Behavioral Decision Making*, *17*(5), 365–379. https://doi.org/10.1002/bdm.478

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, *110*(2), 337–363. https://doi.org/10.1257/aer.20180728

Appendix A

Appendix Manuscript I

A.1 Benchmark Investor

This section describes the process of calculating the Bayesian probability of a price increase using the information known about the price-generating process as well as the price history. The state of the asset is denoted as $s \in \{good, bad\}$, and price increases and decreases as $up = 1$ and $up = 0$, respectively. If the asset is in a good state at time t, the probability of a price increase is $P(up_t = 1|s = good) = .65$, resulting in an upward drift over time. The opposite is true for the bad state, meaning $P(up_t = 1|s = bad) = .35$. Given the knowledge of the investor about this structure one can apply Bayes's rule to calculate the Bayesian belief, B_{t+1}, that the price will increase in the next round $t + 1$. To do so, first one needs to determine the probability of the asset being in the good state, Q_t, given the observed price movement, up_t, and the prior belief that the asset was in a good state in the previous round, Q_{t-1}.

$$
\begin{aligned}
Q_t \quad &= P_t(s = good|up_t) \\
&= \frac{Q_{t-1} \times P(U_t|s_{good})}{P(up_t)} \\
&= \frac{Q_{t-1} \times (.35 + .3up_t)}{Q_{t-1} \times (.35 + .3up_t) + (1 - Q_{t-1}) \times (.65 - .3up_t)} \qquad \text{(A.1.1)}
\end{aligned}
$$

Note that the expression in equation (A.1.1) holds for price increases and decreases. In the former case, $up_t = 1$ and thus $P(up_t = 1|s = good) = .65$, whereas in the latter case, $up_t = 0$, the probability of observing a price decrease in the good state is $P(up_t = 0|s = good) = .35$. The same logic applies vice versa to the case of the asset being in the bad state.

Next, one also has to consider the possibility of a state switch. This can be incorporated in the following way:

$$Q'_t = .8Q_t + .2(1 - Q_t) = .2 + .6Q_t \tag{A.1.2}$$

This binds the prior being in the good state Q'_t for the next round to the interval $[.32, .68]$. Given these updated beliefs, the probability of observing a price increase, B_{t+1}, in round $t + 1$ can be calculated by considering the likelihood that the asset will increase both if it is in a good state and if it is in a bad state:

$$\begin{aligned} B_{t+1} &= .65 \times Q_t + .35 \times (1 - Q_t) \\ &= .35 + .3 \times Q_t \end{aligned} \tag{A.1.3}$$

This confines the Bayesian probability of observing a price increase further to the interval of $[.446, .554]$. Thus, despite price movements being informative, the environment in our experiment is uncertain and information is noisy.

A.2 Phase 1

Additional Figures and Tables

Table A.2.1: Descriptive Statistics of the Sample

Condition	Age	Gender	Exp.	Student	Err. Quiz	IQ Score	Risk Av.
No information	22.95 (2.76)	0.61	0.48	0.95	1.64 (1.72)	2.56 (0.97)	4.91 (2.02)
Partial information	22.92 (3.27)	0.58	0.53	0.95	1.83 (1.89)	2.77 (0.96)	5.44 (2.02)
Full information	22.89 (3.67)	0.56	0.48	0.92	2.08 (1.90)	2.70 (0.99)	5.30 (1.99)
Total	22.92 (3.24)	0.58	0.5	0.94	1.85 (1.84)	2.68 (0.97)	5.21 (2.01)

Note. Descriptive overview of the sample. Standard errors are reported in parentheses where applicable. "Age" is reported in years. "Gender" refers to the fraction of female participants. "Exp." is the fraction who reported having prior experience making investment decisions. "Student" is the fraction who reported to be enrolled as student. "Err. Quiz" describes the number of errors in the comprehension quiz. "IQ Score" is the score on the progressive matrices task. "Risk Av." (risk aversion) describes the average answers to the German Socio-Economic Panel general risk assessment question.

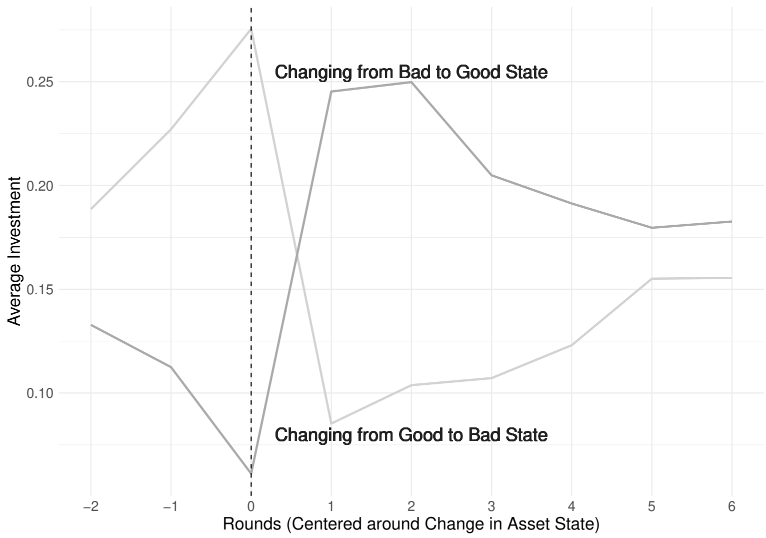

Figure A.2.1: Event analysis of the average portfolio allocation around a state switch. The x axis shows rounds since the first price movement from a new state was observed. Round 0 is therefore the first round in which a participant can react to a state switch and adjust their portfolio allocation accordingly The y axis shows average investments at the given time point. Investments are coded as -1 for a short investment 0 for no investment and 1 for holding the asset. Further, the line is split between whether a switch from the bad to the good state or vice versa happened.

Table A.2.2: Belief Updates for Long and Short Investments

	Long	Short
Favorable price move	$-.102\,(.011, p < .001)^{***}$	$-.137\,(.014, p < .001)^{***}$
Gain position	$.017\,(.012, p = .17)$	$.049\,(.013, p < .001)^{***}$
Interaction	$-.126\,(.018, p < .001)^{***}$	$-.12\,(.024, p < .001)^{***}$
Constant	$.006\,(.05, p = .9)$	$.021\,(.064, p = .74)$
Obs. (Participants)	9,222 (191)	5,336 (191)
Adjusted R^2	.056	.11

Note. $^{***}p < .01$; linear regression predicting the average intensity of belief updates by investment domain and price movement (effect coded, standard errors and p-values in parentheses). The table also reports the adjusted R^2 and number of observations (Obs.) as well as participants. The beliefs were reported on a scale of 0 to 100 and the updating values are calculated as $q_{t-1} - q_t$ where q_t is the belief q in round t. Further, updating values for downward price movements were flipped such that all updates should have a positive sign. The analysis uses cluster robust standard errors clustered by participant. Age and gender are added as control variables but omitted in the table.

Table A.2.3: Influence on Trades

	Buy Long	Sell Short	Liq. Long	Liq. Short
Beliefs	0.048 $(0.002, p < .001)^{***}$	-0.053 $(0.002, p < .001)^{***}$	-0.049 $(0.002, p < .001)^{***}$	0.054 $(0.002, p < .001)^{***}$
Bayesian probability	12.644 $(1.205, p < .001)^{***}$	-17.848 $(1.397, p < .001)^{***}$	-14.872 $(1.133, p < .001)^{***}$	21.403 $(1.511, p < .001)^{***}$
Loss position	–	–	-0.760 $(0.121, p < .001)^{***}$	-0.456 $(0.166)^{**}$
Interaction position belief	–	–	0.012 $(0.002, p < .001)^{***}$	-0.001 $(0.003, p = .73)$
Constant	9.109 $(0.610, p < .001)^{***}$	-9.960 $(0.695, p < .001)^{***}$	-9.166 $(0.592, p < .001)^{***}$	13.140 $(0.723)^{***}$
Obs. (Participants)	5,182 (190)	5,182 (190)	13,325 (192)	8,755 (188)

Note. $^{**}p < .05$; $^{***}p < .01$; mixed logistic regressions predicting trades with the Bayesian probability, the reported beliefs of the participants, and an interaction with investment position in case of liquidation (liq.). Standard errors and p-values are reported in parentheses. Further a random effect term per participant was added. The table also reports the number of observations (Obs.) and participants included in the calculation.

Table A.2.4: Beliefs When Making Trading Decision

Investment	Trading Decision	Position	Mean Reported	*SD*	Mean Bayesian
Hold (+1 share)	Start Invest.		.599	.194	.509
	Liquidate	Gain	.480	.214	.522
		Loss	.459	.196	.479
Short (−1 share)	Start Invest.		.347	.216	.488
	Liquidate	Gain	.448	.207	.475
		Loss	.495	.183	.518

Note. The mean and standard deviation of the beliefs reported by participants in the first phase of the experiment. They are grouped by the trading decisions they made in that round. The last column provides the mean probability calculated by a Bayesian updater in the specific situations for comparison.

Table A.2.5: DE to Benchmark in Phase 1

Constant (no information)	$0.787\,(.208, p < .001)^{***}$
Partial information	$-.056\,(.069, p = .42)$
Full information	$-.078\,(.069, p = .26)$
Observations	192
Adjusted R^2	$-.003$

Note. $^{***}p < .01$; linear regression on the disposition effect (DE) measures for Phase 1 using the participants' condition as dummy variable (standard errors and p-values in parentheses). Age and gender are added as control variables but omitted in the table. Although all measures are significantly above the rational benchmark value, they do not differ significantly, as all participants received the same treatment during Phase 1.

Individual Differences

The size of our sample (192 participants) also allows us to investigate if the effects reported above may be moderated by individual differences. In that regard, previous literature has shown effects of intelligence and risk preferences (e.g., Fellner & Maciejovsky, 2007; Grinblatt et al., 2012) on investment behavior. To test for individual differences, we first added additional variables to the updating regression in Table 2.1 as a three-way interaction with the already present interaction term (the resulting tables can be found in the internet appendix). However, neither age, gender or intelligence nor the reported ambiguity or loss aversion significantly affected the main results. Only two of the additional variables entered a significant three-way interaction with investment position and price movement favorability. These were the number of wrong answers in the initial quiz after the tutorial ($p = .040$)

and the reported engagement with the study ($p < .001$). This may indicate that the effect was stronger for participants who initially struggled to grasp the instructions, possibly due to a lack of motivation. Including this interaction does not, however, take away from the original two-way interaction, which remains significant even in this case ($p < .001$).

All of the mentioned variables as well as the reported investment experience further had no significant impact on disposition effect values. In one questionnaire participants indicated on a 5-point Likert scale how much nine strategies applied to their behavior. None of the suggested strategies had a significant direct impact on participants' payoff. However, those strategies that are expected to correlate negatively with the disposition effect do show a significant effect in the hypothesized direction (see Table A.5.2 in the internet appendix for a detailed analysis).

A.3 Phase 2

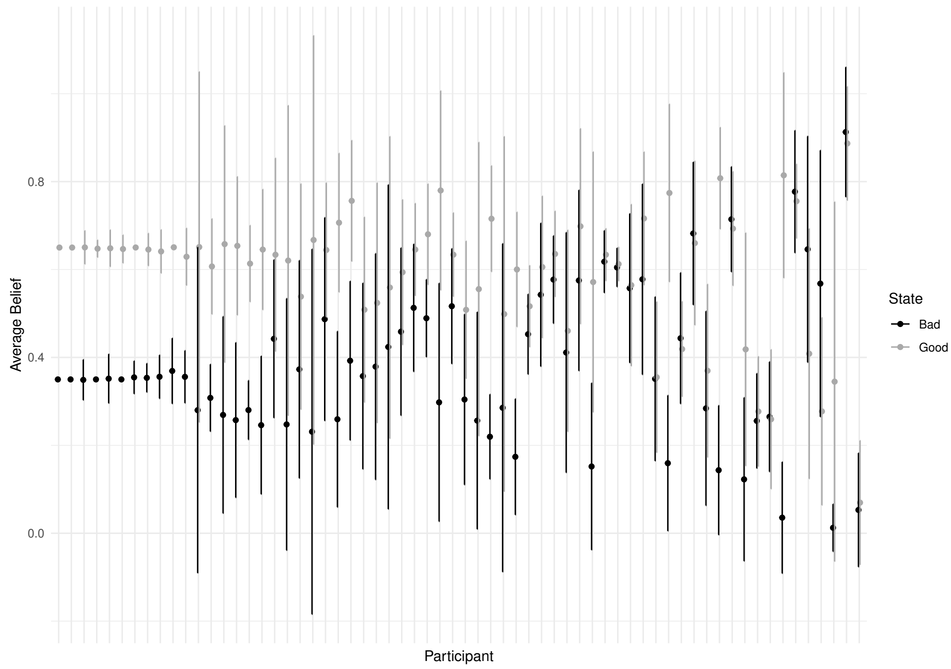

Figure A.3.1: Figure showing the average belief of each participant along the x-axis split between the good and the bad state in the full-information condition. Participants were sorted by their distance to the expected responses of .35 in the bad and .65 in the good state. Error bars indicate one standard deviation.

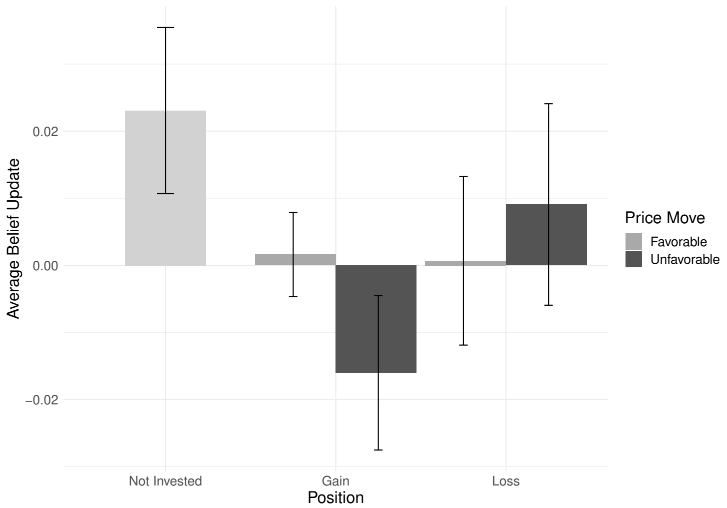

Figure A.3.2: Belief updating of the full-information group in the second phase of the experiment split by investment position and price movement. Depicted are the averages of the raw, nonflipped updating values. Note here that all updates should be 0 if the state remains the same, or either .3 or −.3 in case of a state switch. As state switches are independent of gains, losses, or the favorability of the price move and since both switching probabilities are identical, we would expect these updates to cancel out in all situations. Therefore the benchmark values in this figure are all average updates of 0.

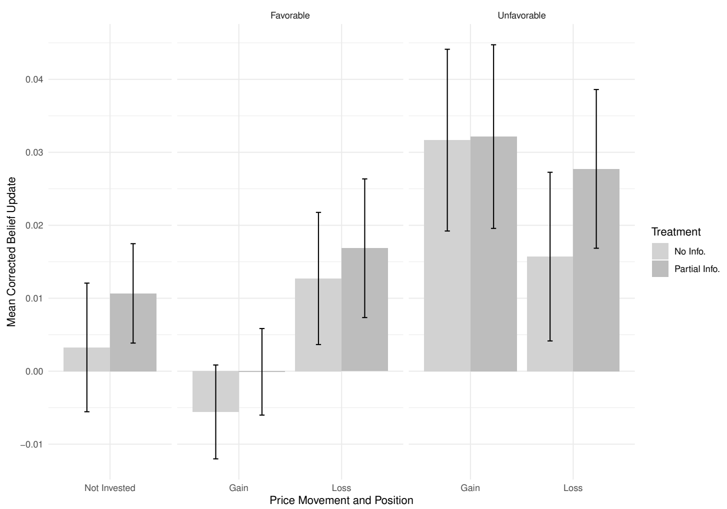

Figure A.3.3: Belief updating in Phase 2 according to the current position (gain, loss, or not invested), the last price move (favorable or unfavorable), and condition (Partial or no nformation). For the same figure when excluding inverse updates see Figure A.5.9 in the internet appendix. The x axis shows the combination of the portfolio allocation from which the update was made (i.e., not invested, gain, or loss) and the favorability of the price move. Further, each of these cells is split into the no-information and partial-information condition. The y axis shows the average of the difference between a Bayesian investor's belief updates and those reported by participants ($\Delta_{\text{Report}} - \Delta_{\text{Bayes}}$). Further, the sign of updates from negative price moves was flipped, as described in the main text. Error bars show group-level 90% confidence intervals.

A.4 Model Estimation Methods

All models were implemented as a hierarchical Bayesian model using the *rstan* (Stan Development Team, 2020) package for the R programming language (R Core Team, 2021). As mentioned in the main text, this approach assumes that each participant's parameters are drawn from a group-level distribution. As an example, if the participant has a certain learning rate η_c in context c, we assume that η_c is drawn from a group-level distribution that is defined by its own mean and standard deviation parameters. This means that not only is the group-level average informed by the individual estimates, but in turn the parameter estimates for individual participants are informed by the estimates of the whole group. This partial pooling of information can lead to a favorable effect called *shrinkage* (McElreath, 2020).

Starting with the learning rate itself, this hierarchical structure was implemented for both the mean and standard deviation of the individual learning rates. The prior distribution for the mean learning rates was defined by a normal distribution with a mean of $-.5$ and standard deviation of .5. For use in the model, the samples from this learning rate distribution were then transformed using the cumulative density function of the standard normal distribution, thus ensuring all values to be within the interval $[0, 1]$. This led the values to be centered around .33 (which is reasonable given previous findings, e.g., Gershman, 2015; Fontanesi et al., 2019), while also putting enough probability weight on other values to represent uncertainty. The group-level standard deviation parameters were drawn in turn from a gamma distribution with values of 1.2 and 3, allowing for a wide range of values to be estimated. The individual learning rates of participant i in context c were therefore defined by:

$$
\begin{aligned}
\mu_c^\eta &\sim \Phi(N(-.5, -5)) \\
\sigma_c^\eta &\sim \gamma(1.2, 3) \\
\eta_{i,c} &\sim N(\mu_c^\eta, \sigma_c^\eta)
\end{aligned}
$$

The belief report at time t, r_t, was then assumed to be a noisy representation of the belief q_t modeled as a normal distribution around q_t with a standard deviation σ_i. This standard deviation parameter was also modeled hierarchically and thus was assumed to be a draw from a population-level gamma distribution. The priors of the shape α_σ and rate β_σ parameters of this group-level distribution were themselves sampled from gamma distributions. To calculate the likelihood of a belief report, the normal distribution around q_t was truncated to the interval $[0, 1]$. Thus, the likelihood of a belief report was defined as

$$\begin{aligned} \alpha_\sigma &\sim \gamma(10, .3) \\ \sigma_\sigma &\sim \gamma(15, .2) \\ \sigma_i &\sim \gamma(\alpha_\sigma, \beta_\sigma) \\ r_t &\sim N(q_t, \sigma_i), r_t \in [0, 1] \end{aligned}$$

where q_t is calculated by the CSRL model as described in the main text. Note that in the full model, situations in which participants were invested but had not made a return yet were counted as "not invested." In the model differentiating only the invested and not invested case, these situations were, however, classified as invested. As we do not find any difference between the learning rates in the latter model, this detail of the implementation seems not to have had a strong influence.

To estimate the posterior probability distributions of the parameters, the *rstan* package uses a Markov Chain Monte Carlo algorithm. In our case we used four chains with $21,000$ iterations each, of which the first $1,000$ were discarded as warm-up. This high number of samples was necessary to obtain a stable estimate of Bayes factors.

A.5 Internet Appendix

Parameter Recovery

To make founded claims about the models tested in this paper two prerequisites must be fulfilled (Spektor & Kellen, 2018): First, the model must be identifiable, meaning that different parameters can not trade of with each other, leading to uncertain results. Second, the data delivered by the task must contain enough information to accurately estimate the parameters of the model. To test these two prerequisites, we use a parameter recovery exercise for our task and the CSRL model.

In detail, we simulate 150 participants with known model parameters which are then recovered through fitting the model to this simulated data. More participants will always add information and thus improve parameter estimates. Therefore we use this number as a "lower bound" to check whether successful parameter recovery is generally possible. To simulate participants' behavior their belief updating and choices were simulated using the following functions: In each price path, beliefs in a price increase started at .5 (reported as 50%). After this, they were updated using a CSRL model. This model used a learning rate of .2 for favorable or unfavorable information in a loss and gain position respectively and .15 for favorable and unfavorable information in gain and loss positions respectively. However, to simulate the expected hierarchical nature of the data, an additional value drawn from a normal distribution with mean zero and a standard deviation of .02 was added to these learning rates for each participant separately. Lastly, the reports of the calculated beliefs were also simulated to be noisy around the true value with a normal distribution centered at zero and a standard deviation of 5.

Based on the updated belief our simulated participants would calculate a prospective value of each possible investment option. For this we calculated the utility of an outcome x as $u(x) = x^{\alpha}$ with $\alpha = .88$. By weighting these values with the outcome probabilities for each investment option one arrives at the utility value for each investment option i, $u(i)$. Lastly, our simulated participants made an investment decision with decision probabilities calculated using the softmax function. This function calculates the choice probability for each option using its utility value $u(i)$ as $p(u(i)) = \frac{e^{u(i)}}{\sum e^{u(i)}}$. This choice function thereby adds non-deterministic choices to the simulated sample.

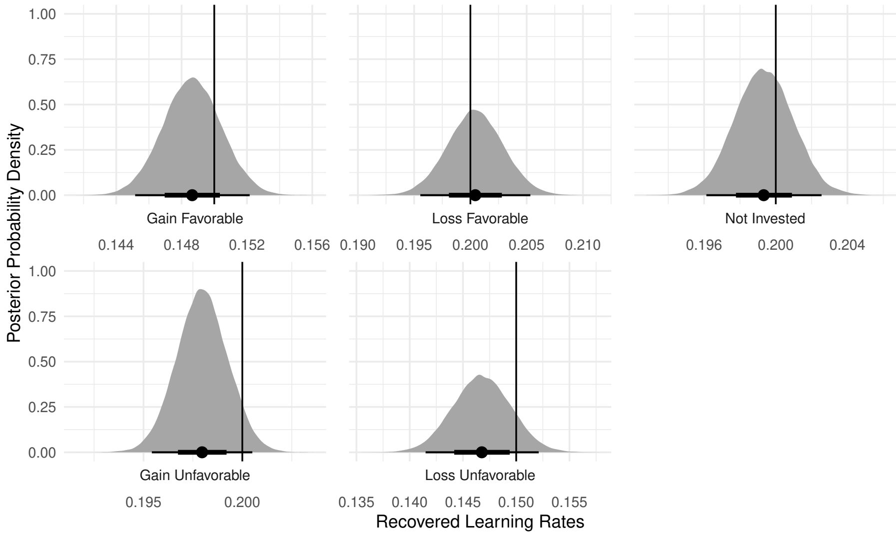

Figure A.5.1: Posterior probability distributions for the recovered group-level mean parameters of the learning rates, μ_c^η. True values were .2 and .15. The dots and whiskers indicate the mean values and the 90% credible intervals, that is, the inner area in which the value of the parameter is estimated to lie with a probability of .9.

For the recovery of the parameters the same procedure as described in Appendix A.4 was used. Figure A.5.1 shows the group level estimates of the recovered parameters. All of the estimated values for the group level average learning rate μ_c^η contain the true value within their 90% credibility intervals. Despite the noise added on several levels (i.e. different learning rates per participant, stochastic investment decisions and noise added to the belief reports), these estimates are very precise, indicated by the narrow credibility intervals. We conclude that the task in therefore well suited to gain insight into contextual differences between learning rates.

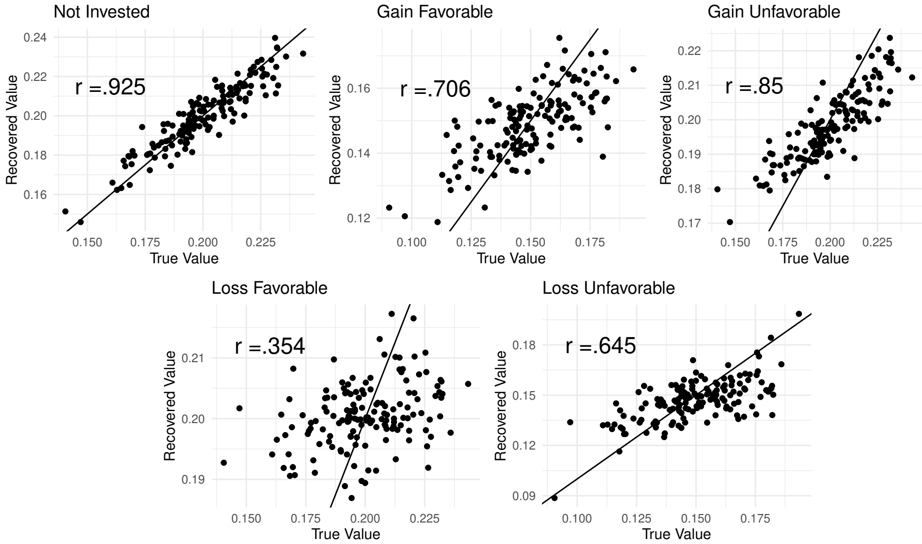

Figure A.5.2: Scatterplots showing the correlation between each simulated participant's true learning rates and the recovered value (median of the posterior distribution). The black line has a slope of 1 indicating perfect recovery.

Figure A.5.2 shows the parameter recovery on the individual level. While all recovered parameters correlate highly significantly ($p < .001$) with their original value, they do so less strongly in those contexts that occur less often (i.e. encountering a favorable update after having made a loss). This indicates that estimates of parameters on the individual level should be interpreted carefully. The focus of our investigation lies on the group level parameters however, which we were able to recover with high precision.

Model Estimation Results

Table A.5.1: Standard Deviation Estimates for CSRL Model Learning Rates

Parameter	Mean	5th Percentile	95th Percentile
Not Inv.	4.32	3.80	4.93
Fav. Gain	4.53	3.79	5.44
Fav. Loss	4.22	3.66	4.91
Unfav. Gain	3.87	3.39	4.46
Unfav. Loss	4.01	3.42	4.74

Note. This table shows the mean, 5th, and 95th percentiles for the parameter distribution of the group-level standard deviation of the context-sensitive reinforcement learning (CSRL) model learning rates, σ^{η}. "Not Inv." is not invested. "Fav." and "Unfav." are favorable and unfavorable information, respectively. Note that these parameters were estimated on an "unbound" scale and are the standard deviations for distributions that were transformed by the cumulative probability density function of the standard normal distribution before learning rates were drawn from them.

Illustrative Example

To further understand the effects of the interaction in belief updating on investment behavior, consider Figure A.5.3, which shows the updating in the first five rounds of an illustrative price path. We assume three rounds of increasing prices (indicated by the + signs in the figure), followed by two rounds with falling prices (−). All beliefs start at .5 and the reinforcement learner uses the mean estimates of the learning rates displayed in Figure 2.2. Note that the belief of a Bayesian learner will also increase steeply at first but will then taper when approaching a value of .65. A CSRL model learner who is not invested and therefore uses only one constant learning rate makes adjustments in the beliefs that are of similar magnitude to that of a Bayesian learner. However, given enough price increases, a reinforcement learner would approach a belief of 1 and therefore would be certain that another price increase has to follow. After witnessing the subsequent price decreases, both the Bayesian learner and the reinforcement learner who is not invested end up with similar beliefs about the next price move being an increase.

Turning to the invested case, we see a clear difference between updating in a gain and a loss position. The investor in the loss position will over-update the favorable information, leading to an overly optimistic expectaion for a price increase. The following unfavorable information is then under-updated, which means that the investor would remain over-optimistic. The investor in the gain position in turn will under-update favorable information as compared to the Bayesian benchmark and then over-update the unfavorable information. Toghether, this will result in a belief

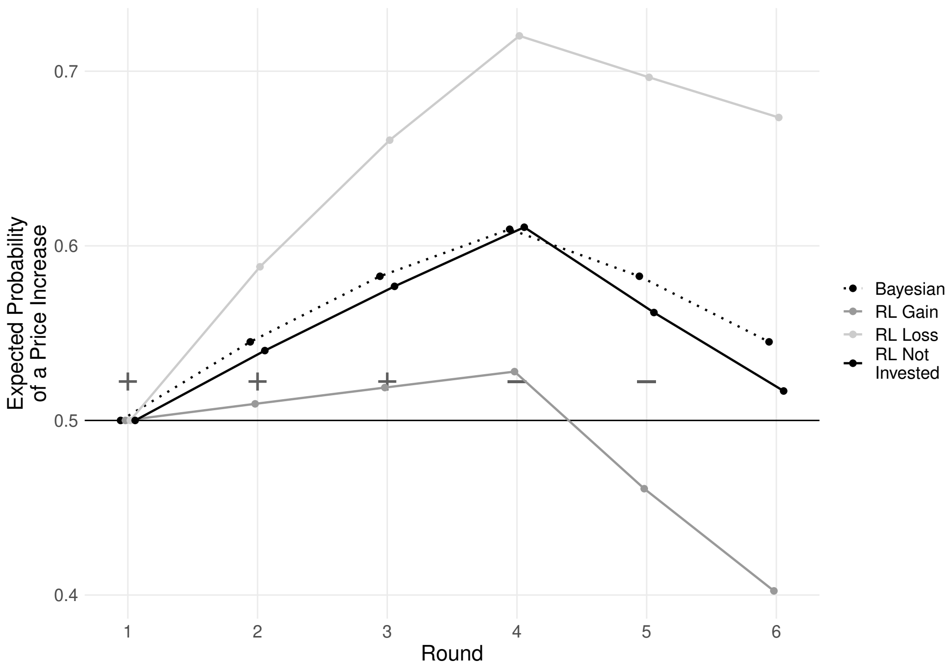

Figure A.5.3: Figure comparing the beliefs of a Bayesian updater (dashed line) against those created by our proposed reinforcement learning (RL) model (solid lines). The learning rates used for the reinforcement learner are the median values of the posterior distribution of each group-level parameter. The beliefs of a reinforcement learner are shown for a gain and a loss position as well as for not invested at all. Given that the investor does not short sell the asset, the first three price movements in this example are favorable (indicated by the + sign in the graph) followed by two unfavorable movements (indicated by a − sign).

in a price increase below .5, most likely leading to a liquidation decision by a risk-neutral investor. Although not all situations will be as clear-cut as this example, the average result of these updating patterns will be an increased likelihood of liquidating investments in a gain compared to a loss position.

Elaboration of Ordered Logistic Regression Results

The effect size of an ordered logistic regression reflects the log odds of observing a category or any of the "lower" categories. In the case of our investment portfolios, the categories were ordered, from lowest to highest, as "short sold," "not invested," and "holding the asset." The log odds here therefore concern whether one's portfolio Y is in category y or "lower," $\ln\left(\frac{P(Y \leq y)}{1-P(Y \leq y)}\right)$. The estimated parameter can now be interpreted as the difference in log odds per unit increase of the respective predictor. In the case of Model 1 in Table 2.2, this predictor would be the reported belief. The intercept of a given category then indicates the log odds of observing this category or a lower one, given a predictor (belief) of 0, $\ln\left(\frac{P(Y\leq y|q=0)}{1-P(Y\leq y|q=0)}\right)$.

To provide an example, we use the values of Model 1 in Table 2.2. The intercept of not invested is 2.28. Calculating the corresponding probability to this log odds value indicates that given a reported belief of 0, the probability of either being not invested or shorting is 91%. Further, given that the log odds space is linear, a given predicted probability can be calculated by using the intercept, estimated parameter, and predictor, similar to a linear model. As an example, we can calculate the probability of short selling (intercept 1.16) when holding a belief (estimated parameter of .05) of .5 (i.e., 50 on the participants' rating scale). We arrive at a log odds value of $1.16 - 50 \times .05 = -1.34$, which translates to a probability of short selling of 26.1%. As a second example, we can easily calculate the belief at which the probability of being shorted is .5. This is because a probability of .5 implies log odds of $\ln\left(\frac{.5}{.5}\right) = 0$. Such a log odds value of 0 is reached when the reported belief is $1.16/.05 = 23.2$ percentage points.

Additional Tables and Figures

Table A.5.2: Effect of Strategies

Strategy	(1)	(2)
I bought and sold the stock randomly.	$-.12^{*}$	$.15^{**}$
I trusted my gut feeling about the price development.	.018	$.15^{**}$
I invested whenever I was convinced that a price increase would happen with over 50% probability.	.035	$-.24^{***}$
I shorted whenever I was convinced that a price decrease would happen with over 50% probability.	.081	$-.25^{***}$
I only invested when I was very sure that the price would increase.	.053	.08
I only shorted when I was very sure that the price would decrease.	.020	$.21^{***}$
I tried to hold/not buy back the stock for as long as possible.	$-.07$	$.18^{**}$
I tried to sell/buy back the stock only after having made a gain.	.021	$.31^{***}$
I tried to sell/buy back the stock as fast as possible after having made a loss.	.026	$-.06$

Note. $^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$; participants were asked to indicate on a Likert scale how much the nine presented strategies applied to them (translated from German). The table depicts the ρ value of a Spearman correlation test. Column (1) shows the correlations between the answers and the final payoff of Phase 1 and column (2) does the same with the disposition effect values.

Table A.5.3: Updating With Extra Variables

Variable	Age	Gender	Intelligence	Engagement	Amb. Av.	Loss Av.
Favorability	$-.052\,(.011)^{***}$	$-.052\,(.011)^{***}$	$-.052\,(.011)^{***}$	$-.053\,(.011)^{***}$	$-0.052\,(.011)^{***}$	$-.052\,(.011)^{***}$
Position	$.090\,(.015)^{***}$	$.090\,(.015)^{***}$	$.090\,(.015)^{***}$	$.090\,(.015)^{***}$	$0.090\,(.015)^{***}$	$.090\,(.015)^{***}$
Favorability × Position	$-.101\,(.101)$	$-.124\,(.022)^{***}$	$-.141\,(.025)^{***}$	$-.05\,(.027)^{*}$	$-0.121\,(.022)^{***}$	$-.099\,(.023)^{***}$
3WI	$-.001\,(.004)$	$-.002\,(.023)$	$.006\,(.006)$	$-.015\,(.004)^{***}$	$-0.001\,(.003)$	$-.004\,(.003)^{*}$
Constant	$.016\,(.082)$	$.024\,(.052)$	$.024\,(.051)$	$.028\,(.050)$	$0.025\,(.051)$	$.023\,(.051)$
Adjusted R^2	.073	.073	.073	.074	.073	.073

Note. $^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$; estimates and standard errors of a linear ordinary least squares regression on participants' reported belief updates as shown in the main text in Table 2.1. These models, however, add a three-way interaction (3WI) between the price movement favorability, the investment position, and one additional variable, which is denoted in the table heading row. Ambiguity and loss aversion are abbreviated "amb. av." and "loss av.," respectively. All models include age and gender as control variables.

Table A.5.4: Percentage of Successful Trades

Trade	Position	No Info.	Partial Info.	Full Info.
Open		52.2%	50.3%	53.4%
Jump	Gain	49.1%	49.5%	56.7%
	Loss	50.5%	51.4%	55.8%
Close	Gain	48.9%	48.4%	46.7%
	Loss	54.2%	53.4%	53.7%

Note. The percentages here refer to the fractions of trades that resulted in either an increase in the portfolio's value or the avoidance of a decrease. For example, a trader can decide to sell a previously held asset. If the assets price decreases in the next round, that is counted as a successful trade. This table reports only the values in Phase 2, where the conditions did differ.

Table A.5.5: Beliefs at Time of Trades

Investment	Trading Decision	Position	No Info.	Partial Info.	Full Info.
Hold (+1 share)	Start Investment		.618(.180)	.583(.176)	.626 (.223)
	Liquidate	Gain	.448 (.221)	.478 (.162)	.475 (.246)
		Loss	.409 (.204)	.481 (.173)	.410 (.222)
Short (−1 share)	Start Investment		.317 (.204)	.370 (.214)	.307 (.234)
	Liquidate	Gain	.407 (.193)	.456 (.156)	.436 (.254)
		Loss	.443 (.218)	.482 (.138)	.471 (.248)

Note. Means and standard deviation (in parentheses) of participants' reported beliefs about a price increase in the rounds of a trade, separately for the three conditions. This table reports only the values in Phase 2, where the conditions did differ.

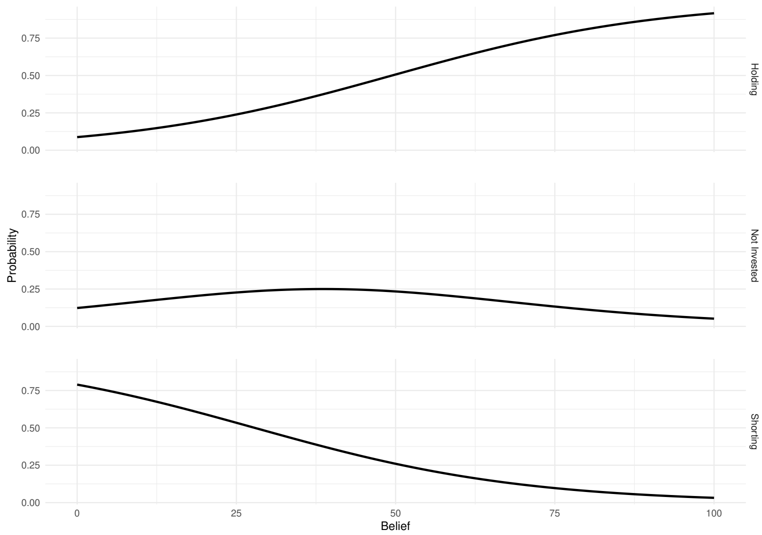

A visual representation of the model-estimated probabilities of portfolio allocation for different beliefs. The x axis shows the reported belief in a price increase on a scale of 0 to 100%. The y axis shows the probability of each portfolio (i.e., holding, shorting, or not being invested) according to the ordered logistic regression model at a specific belief.

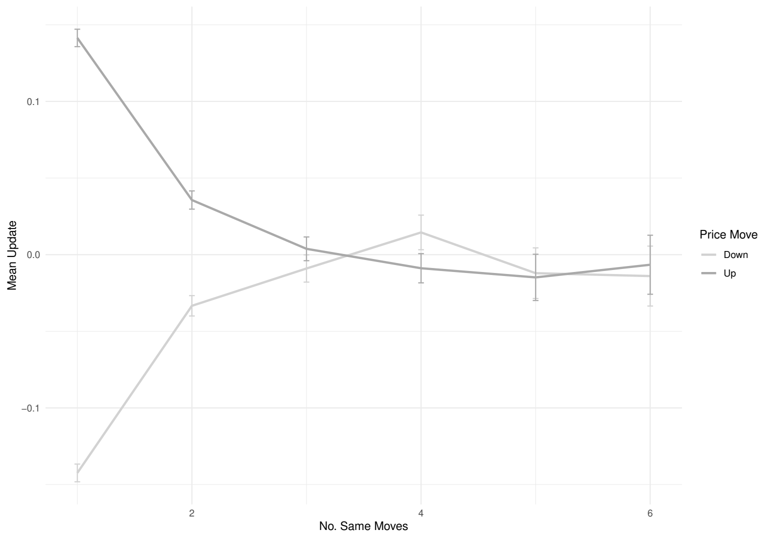

Figure A.5.4: Average belief updates after repeated price moves. The x axis shows the number of repeated price movements. The y axis shows the average belief update after the price move.

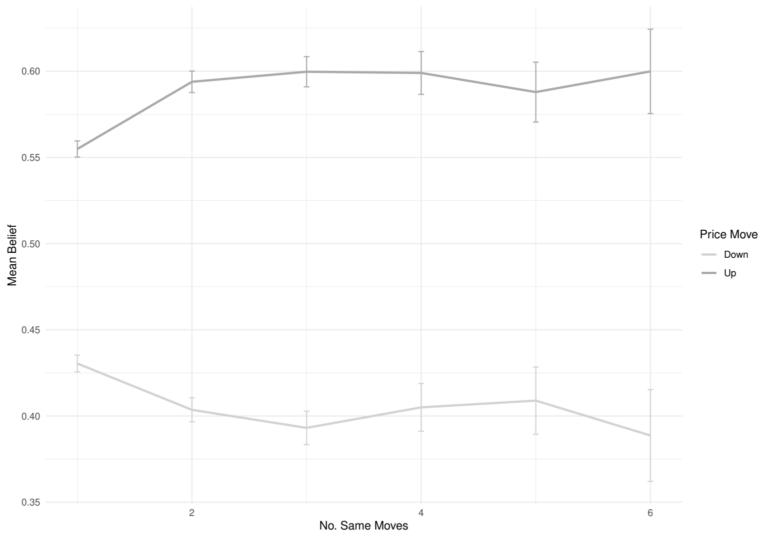

Figure A.5.5: Average beliefs after repeated price moves. The x axis shows the number of repeated price movements. The y axis shows the average belief after the price move.

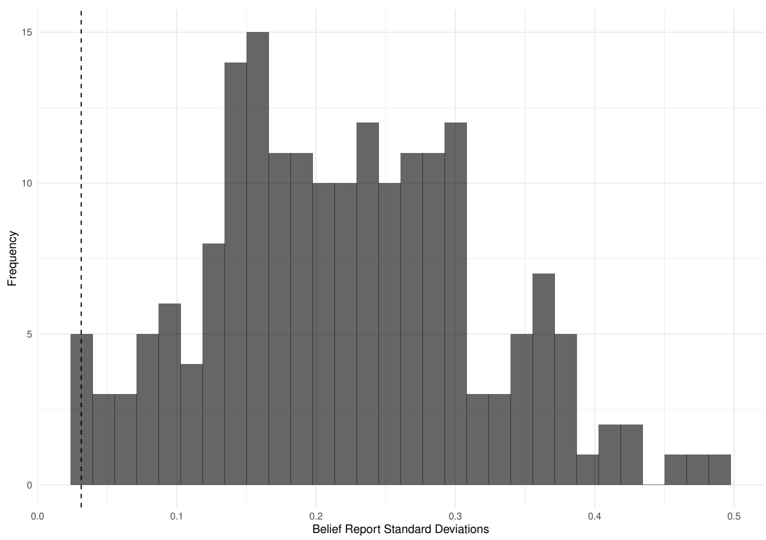

Figure A.5.6: Histogram of the standard deviations of belief reports between subjects. The dashed line indicates the value of a Bayesian updater for comparison, which is .0314.

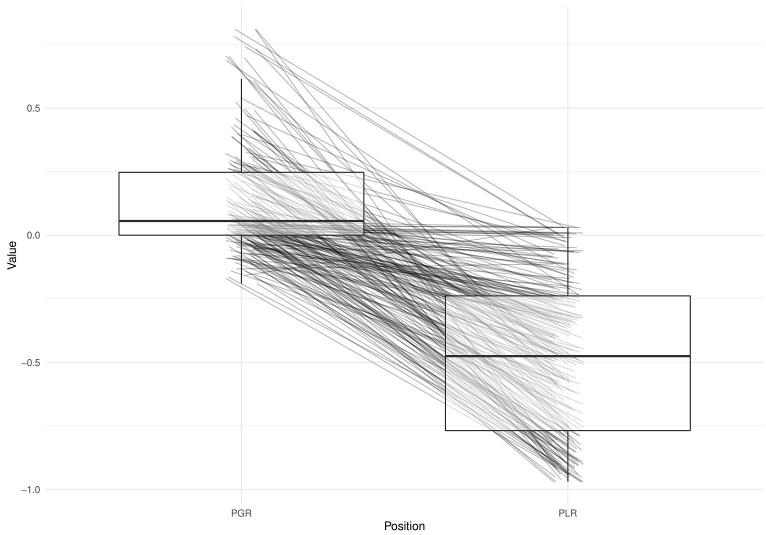

Figure A.5.7: Boxplot showing the proportion of realized gains (PGR) and losses (PLR) for each participant in Phase 1 relative to their respective benchmarks. The boxplot displays the median (line), 25th and 75th percentile (box), and 1.5 the interquartile range (whiskers). The diagonal lines connect the values of individual participants. Their steepness can be interpreted as the difference between the rational benchmarks of the PGR and PLR and therefore the disposition effect value (i.e., the difference from its rational benchmark value).

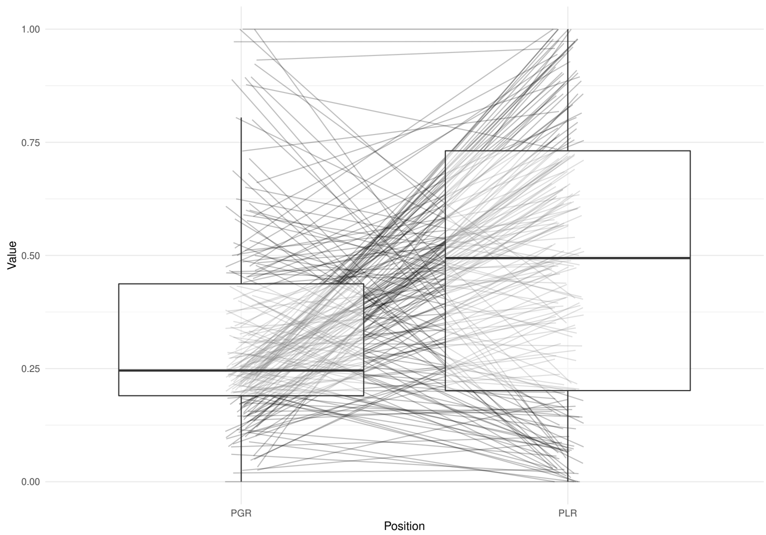

Figure A.5.8: Boxplots showing the raw values for the proportion of realized gains (PGR) and losses (PLR) for each participant in Phase 1. The boxplot displays the median (line), 25th and 75th percentile (box), and 1.5 the interquartile range (whiskers). The diagonal lines connect the values of individual participants. Their steepness can be interpreted as the difference between PGR and PLR and therefore the raw disposition effect value.

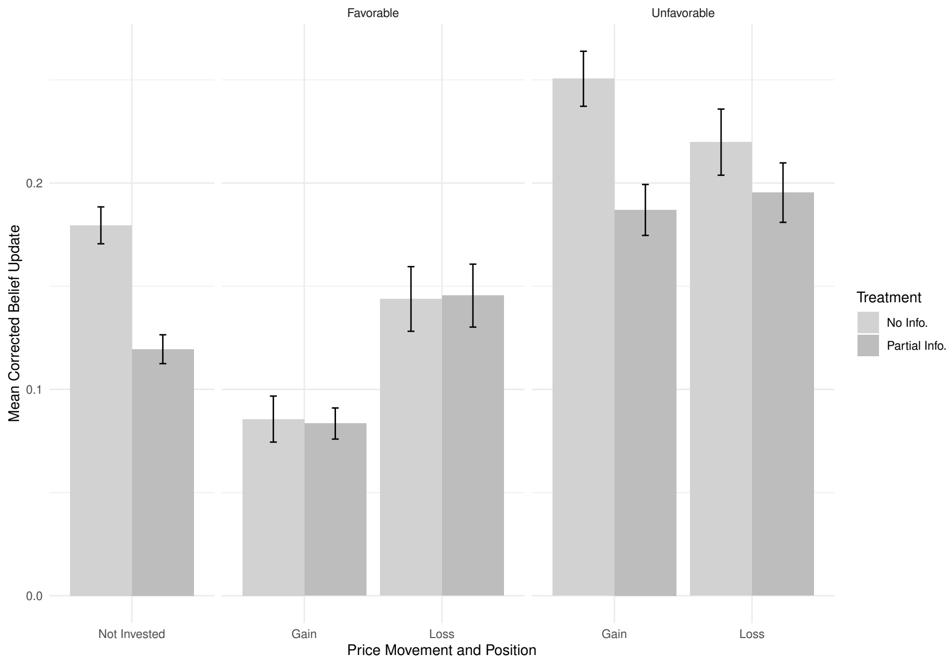

Figure A.5.9: Belief updating in Phase 2 by the current position (gain, loss, or not invested), the last price move (favorable or unfavorable), and the condition (partial or no information). The figure excludes inverse updates. The x axis shows the combination of the portfolio allocation from which the update was made (i.e., not invested, gain, or loss) and the favorability of the price move. Further, each of these cells is split into the no-information and the partial-information conditions. The y axis shows the average of the difference between a Bayesian investor's belief updates and those reported by participants ($\Delta_{\text{Report}} - \Delta_{\text{Bayes}}$). Further, the sign of updates from negative price moves was flipped, as described in the main text. Error bars show group-level 90% confidence intervals.

Robustness Checks

Table A.5.6: Robustness Checks for Belief Updates

Favorable price move	$-.057\,(.014)^{***}$
Gain position	$.063\,(.017)^{***}$
Interaction	$-.126\,(.02)^{***}$
Constant	$.154\,(.097)^{***}$
Obs. (Participants)	6,698 (163)
Adjusted R^2	.093

Note. $^{***}p < .01$; linear regression predicting the average intensity of belief updates by investment domain and price movement (dummy coded). The table also reports the adjusted R^2 and number of observations (Obs.) and participants. The beliefs were reported on a scale of 0 to 100 and the updating values are calculated as $q_{t-1} - q_t$, where q_t is the belief q in round t. Further, updating values for downward price movements were flipped such that all updates should have a positive sign. The analysis uses cluster robust standard errors clustered by participant. Age and gender are added as control variables but omitted in the table. The regression excludes both inverse updates as well as participants who reported an engagement with the study lower than 3 (reported on a 7-point Likert scale).

Table A.5.7: Difference in DE Values in Phases 1 and 2

	(1)	(2)
Constant (no information)	$-.18\,(.16)$	$-0.03\,(0.04)$
Partial information	$-.03\,(.05)$	$-.04\,(.06)$
Full information	$.16\,(.05)^{***}$	$-.69\,(.06)^{***}$
Control	Yes	Yes
Observations	162	192
Adjusted R^2	.53	.51

Note. $^{***}p < .01$; linear regression on the difference in the disposition effect (DE) values between Phases 1 and 2 (standard errors in parentheses). All participants received the same treatment (no information) during Phase 1; the different groups received increasing amounts of information about future price moves in Phase 2. This table reports the analysis when using the raw disposition effect values (1), as opposed to subtracting the value of a risk-neutral Bayesian investor, and when excluding those participants who reported an engagement with the study below 3 on a 7-point Likert scale (2).

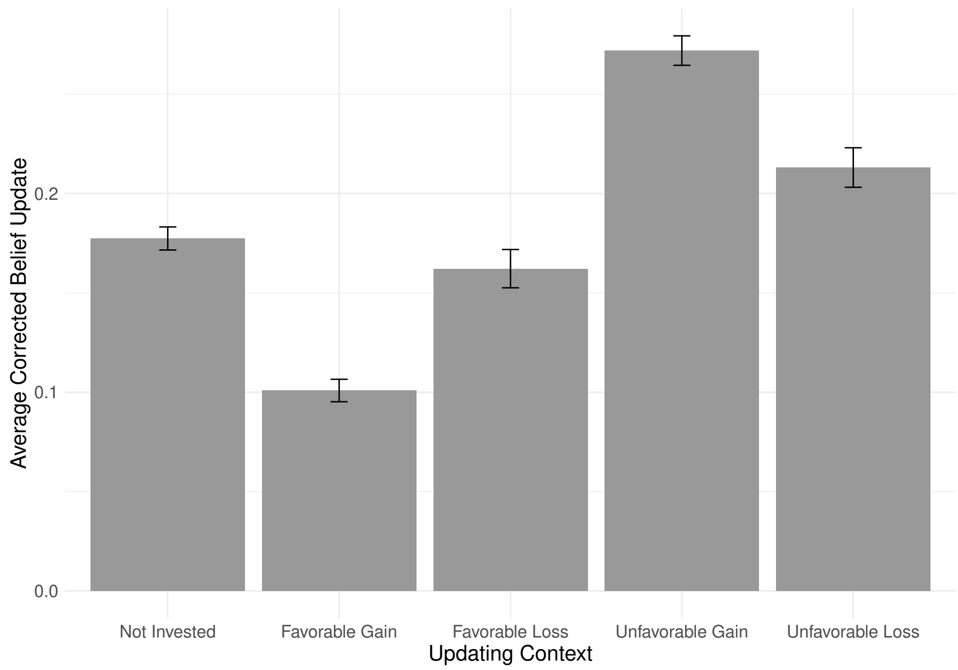

Figure A.5.10: Belief updating during the first two blocks, excluding updates counter to the direction of the last price move. The updating magnitude is corrected for that of a Bayesian-updating investor. Error bars show group-level 90% confidence intervals. The pattern remains when excluding those whose frequency of inverse updates is above the median (i.e. focusing on the half of participants with the lowest amount of inverse updates).

Inverse Updates

A substantial number of belief updates (44.4%) were counter to the direction of the price move. This means that participants increased their expectations about a price increase despite having just witnessed a price decrease and vice versa. In this section we further analyze their characteristics and whether there are individual differences in the number of "inverse updates" during Phase 1. It should, however, be noted that the main findings of the paper are robust to inclusion or exclusion of said rounds. Further it is important to mention that the median inverse update had a magnitude of four percentage points. This means that most of the inverse updates may constitute a "trembling hand" error when using the slider to indicate the current belief.

We nonetheless further investigate the nature of these inverse updates. Figure A.5.4 shows the average updates after streaks of one to six price movements in the same direction. On average, participants updated their beliefs counter to the price movement after a streak of four price movements in the same direction. This could be interpreted as evidence for a possible ceiling effect, but this is an unlikely explanation, as the average belief at that point is still far from either 1 or 0 (see Figure A.5.5).

Although there was no significant gender effect ($p = .338$), the number of inverse updates per participant did in fact correlate with participants' age ($r = .14, p = .049$). This correlation however vanishes if one participant is excluded, who constituted an outlier in terms of age (44 years, $r = .078, p = .282$). Of the remaining control variables, that is, engagement with the study, result of the matrix intelligence test, and reported general risk attitude, none were significantly correlated with the number of inverse updates ($p > .05$). One value that did correlate significantly with the number of inverse updates was that of the disposition effect measure itself ($p < .001$), which, however, does not explain why these participants updated their beliefs in the way they did.

Task Screenshots and Translation

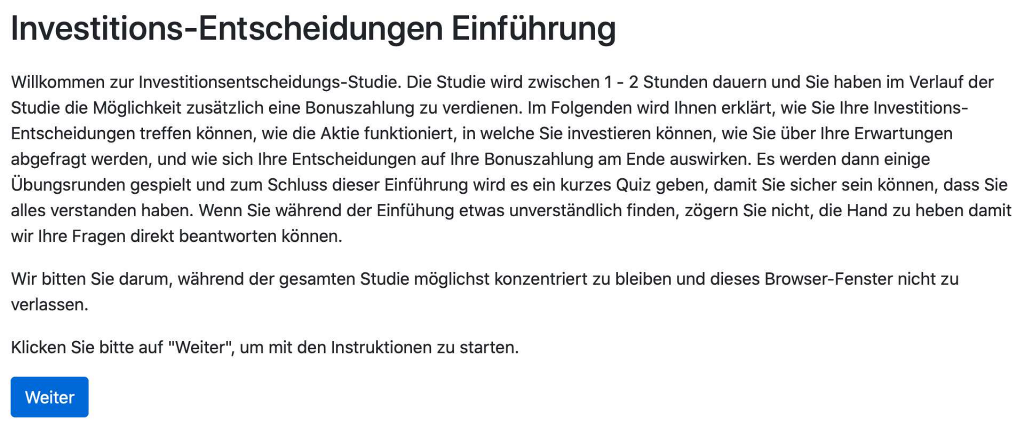
Investitions-Entscheidungen Einführung

Willkommen zur Investitionsentscheidungs-Studie. Die Studie wird zwischen 1 - 2 Stunden dauern und Sie haben im Verlauf der Studie die Möglichkeit zusätzlich eine Bonuszahlung zu verdienen. Im Folgenden wird Ihnen erklärt, wie Sie Ihre Investitions-Entscheidungen treffen können, wie die Aktie funktioniert, in welche Sie investieren können, wie Sie über Ihre Erwartungen abgefragt werden, und wie sich Ihre Entscheidungen auf Ihre Bonuszahlung am Ende auswirken. Es werden dann einige Übungsrunden gespielt und zum Schluss dieser Einführung wird es ein kurzes Quiz geben, damit Sie sicher sein können, dass Sie alles verstanden haben. Wenn Sie während der Einfühung etwas unverständlich finden, zögern Sie nicht, die Hand zu heben damit wir Ihre Fragen direkt beantworten können.

Wir bitten Sie darum, während der gesamten Studie möglichst konzentriert zu bleiben und dieses Browser-Fenster nicht zu verlassen.

Klicken Sie bitte auf "Weiter", um mit den Instruktionen zu starten.

Weiter

Investment-Decisions Instructions Welcome to this investment-decision study. The study will take between 1 - 2 hours and you will have the opportunity to earn a bonus throughout the study. In the following it will be explained how you can make your investment decisions, how the asset works in which you can invest, how you will be asked about your expectations and how these decisions influence your bonus payment in the end. Then there will be some training rounds and in the end of the instructions there will be a brief quiz for you to make sure that you understood everything. If you find something confusing during the instructions do not hesitate to raise your hand so we can answer your questions.

We kindly ask you to keep as focused as possible throughout the whole study and to remain within this browser window.

Press "Continue" to start the instructions.

Investitions-Entscheidungen Einführung

Handelsentscheidungen treffen

In dieser Studie werden Sie in 4 "Blöcken" über jeweils 75 Runden Investitionsentscheidungen treffen. Es wird dabei eine Aktie geben, in die Sie investieren können. Zu Beginn wird der Wert dieser Aktie 1000 Punkte sein. Sie können in jeder Periode entscheiden wie Ihr Portfolio aussehen soll. Dabei haben Sie drei Möglichkeiten: Sie können eine Aktie halten, Sie können keine Aktie halten, oder Sie können einen Blankoverkauf tätigen. Letzteres wird später genauer erläutert.

Unten sehen Sie den Bildschirm mit welchem Sie während der Aufgabe interagieren werden. Es werden hier die einzelnen Spalten kurz beschrieben. Schauen Sie sich diese am besten jeweils kurz in der Tabelle an, während dem Sie sich die Liste durchlesen:

- **Spalte I:** Zeigt die gehaltenen Anteile der Aktie. '+1' bedeutet also, dass Sie in die Aktie investiert haben indem Sie Sie gekauft haben und nun besitzen. Eine '0' bedeutet, dass Sie die Aktie momentan nicht halten. Eine '-1' bedeutet, dass Sie einen Blankoverkauf getätigt haben. Dies wird Ihnen später erklärt. Ob Sie die Aktie halten, nicht halten, oder einen Blankoverkauf tätigen können Sie jeweils entscheiden.
- **Spalte II:** Zeigt den Einstandskurs an. Dies ist der Preis zu welchem Sie die Aktie gekauft haben, falls Sie diese halten.
- **Spalte III:** Diese Spalte zeigt den aktuellen Kurs der Aktie an. Dies ist der Preis, zu welchem Sie in dieser Runde die Aktie kaufen oder verkaufen können. Wie sich dieser Preis entwickelt wird später genauer erklärt.
- **Spalte IV:** Diese Spalte kann zu einem späteren Zeitpunkt weitere Informationen über die zukünftige Preisentwicklung enthalten. Der Inhalt der Spalte wird Ihnen dann genauer erläutert.
- **Spalte V:** Diese Spalte zeigt Ihnen an, wie viel Gewinn oder Verlust Sie mit der aktuell gehaltenen Aktie bisher gemacht haben, seit Sie diese gekauft haben. Die Zahl in der Klammer drückt das selbe Konzept in Prozenten aus. 100% Gewinn würde bedeuten, dass sich der Preis seit dem Kauf verdoppelt hat. -50% würde wiederum bedeuten, das sich der Preis halbiert hat, und Sie somit 50% Verlust gemacht haben.
- **Spalte VI:** Zeigt den aktuellen Kurswert der einzelnen Positionen an. Wenn Sie also die Aktie mit Wert 100 halten, würde hier 100 angezeigt. Ebenso sehen Sie hier Ihr aktuelles Bargeld, sowie den gesamten Wert Ihres Bargeldes plus dem Ihres Portfolios in der "Gesamt" Zeile.

Unterhalb der Tabelle finden Sie die Schaltflächen, mit welchen Sie Ihre Handelsentscheidungen treffen können. Um sich mit dem Prozess vertraut zu machen, kaufen Sie bitte die Aktie, indem Sie auf "+1 Aktie Halten" klicken.
Der Instruktionstext wird dann hier weitergeführt und Sie können nach unten scrollen.

	I	II	III	IV	V	VI
	Anteil	**Einstandskurs**	**Aktueller Kurs**	**Aktueller Zustand**	**Gewinne/Verluste Absolut (in %)**	**Kurswert**
Bargeld						2500
Aktie	0	0	1000	-	0 (0.0)	0
Gesamt						**2500**

-1 Aktie Halten | Keine Aktie Halten | +1 Aktie Halten

Making investment decision In this study you will make investment decisions in 4 "blocks" consisting of 75 rounds each. There will be an asset in which you invest. At first the value of that asset will always be 1000 points. You can decide in each period what your portfolio should look like. There are three possibilities: You can hold a share, you can hold no share or you can short sell. The latter will be explained shortly.

Below you can see the screen with which you will interact throughout the task. The columns are described below. Please also look at them in the table wen you read their descriptions below.

- **Column I** Shows the number of shares of the asset. '+1' means that you have invested into the asset by buying a share which you now hold. '0' means that you don't hold any shares at the moment. '-1' means that you have short sold the asset. This will be explained later. Whether you want to hold a share, no shares or short sell a share is up to you to decide in each round.
- **Column II** Shows the base value. This is the value at which you bought or short sold the asset in case you invested.
- **Column III** This column shows the current value of the asset. This is the price at which you could buy or sell a share in the current round. How this price moves is also explained later.
- **Column IV** This column can contain further information about future values at a later point. The content of this column will be explained at the relevant time.
- **Column V** This column displays whether you have made any gains or losses so far with the current investment. The number in parentheses displays the same concept in relative numbers. 100% gain means that the price has doubled since your initial investment. -50% would mean that the price has halved since and you therefore have made a loss of 50%.
- **Column VI** Shows the current value of the individual positions. Therefore, if you hold a share with a value of 100, this would display 100. The column also displays your current cash, as well as the value of your cash plus your portfolio in the "total" row.

Below the table you can find the buttons with which you can make your investment decisions. To make yourself comfortable with the process, please buy one share by clicking the "hold +1 shares" button. The instruction text will then be continued here and you will be able to scroll down.

Sehr gut. Sie haben nun in die Aktie investiert. Der Wert der gekauften Aktie wurde Ihrem "Bargeld" Feld abgezogen.

Um Ihnen während der Aufgabe später eine Orientierung über Ihre Zeit zu geben, wird nach 6 Sekunden eine blinkende Warnmeldung erscheinen, wie Sie sie unten sehen können. Wenn Sie Ihre Entscheidungen jeweils treffen, bevor die Meldung erscheint, werden Sie in der vorgegebenen Zeit fertig. Lassen Sie sich jedoch nicht unter Druck setzen. 6 Sekunden reichen, um eine überlegte Entscheidung zu treffen. Wenn Sie möchten, können Sie die Zeit in den folgenden Übungsrunden einmal verstreichen lassen, um ein Gefühl dafür zu erhalten.

So sieht diese Warnmeldung aus: "Bitte entscheiden Sie sich jetzt!"

Nachdem Sie sich für eine Aktion entschieden haben wird Ihnen in der Aufgabe für 3 Sekunden ein Update über den Preis der Aktie gegeben. Wie sich der Preis entwickelt, wird auf den nächsten Seiten genauer erklärt. Dieses Update sieht zum Beispiel so aus:

> Der Preis der Aktie **steigt** um **5** und beträgt nun **1005**

Nehmen Sie sich kurz Zeit um sich anzuschauen, wie sich die Tabelle nach Ihrem Kauf und dem Preis-Update verändert hat.

Wenn Sie soweit alles verstanden haben, gehen Sie bitte davon aus, dass Sie die Aktie nun verkaufen möchten. Klicken Sie bitte die Entsprechende Schaltfläche.

	I	II	III	IV	V	VI
	Anteil	**Einstandskurs**	**Aktueller Kurs**	**Aktueller Zustand**	**Gewinne/Verluste Absolut (in %)**	**Kurswert**
Bargeld						1500
Aktie	+1	1000	1005	-	5 (0.5)	1005
Gesamt						**2505**

-1 Aktie Halten | Keine Aktie Halten | +1 Aktie Halten

Very good. You have now invested into the asset. The value of the bought share was subtracted from the "cash" field.

To give you some guidelines regarding the time during the task later on a blinking indicator text like the one you can see below will appear after six seconds. If you always make your decisions before the appearance of this indicator text you will finish the study within the planned time. However, do not feel put under pressure by the indicator. Six seconds are enough to make a well considered decision. If you like, you can let the time pass during some rounds in the training rounds to get a better feeling for it.

This is what the warning will look like: "Please decide now!"

After having decided for an action you will be given an update on the asset's price for three seconds. How the price develops will be explained in the upcoming pages. The update will look like this:

The price of the asset **rises** by **5** and is now **1005**

Please take a moment to look at what the table will look like after you bought the asset.

If you have understood everything so far, please assume that you would like to sell the asset now. Please click on the corresponding button below.

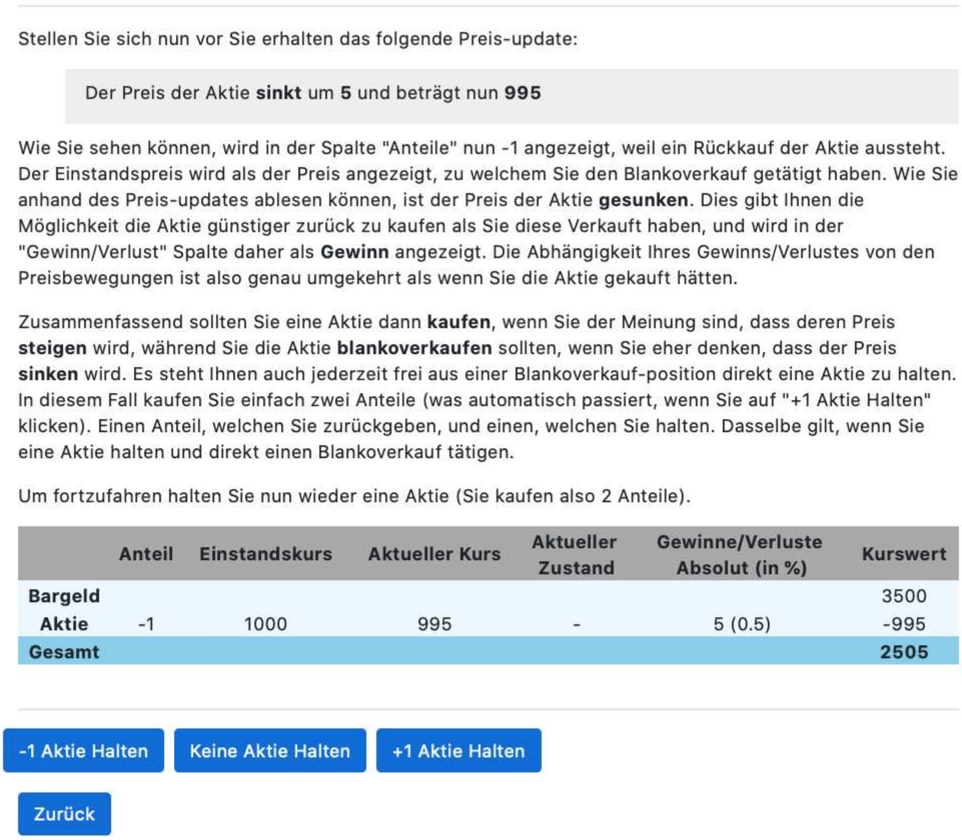

Stellen Sie sich nun vor Sie erhalten das folgende Preis-update:

Der Preis der Aktie **sinkt** um **5** und beträgt nun **995**

Wie Sie sehen können, wird in der Spalte "Anteile" nun -1 angezeigt, weil ein Rückkauf der Aktie aussteht. Der Einstandspreis wird als der Preis angezeigt, zu welchem Sie den Blankoverkauf getätigt haben. Wie Sie anhand des Preis-updates ablesen können, ist der Preis der Aktie **gesunken**. Dies gibt Ihnen die Möglichkeit die Aktie günstiger zurück zu kaufen als Sie diese Verkauft haben, und wird in der "Gewinn/Verlust" Spalte daher als **Gewinn** angezeigt. Die Abhängigkeit Ihres Gewinns/Verlustes von den Preisbewegungen ist also genau umgekehrt als wenn Sie die Aktie gekauft hätten.

Zusammenfassend sollten Sie eine Aktie dann **kaufen**, wenn Sie der Meinung sind, dass deren Preis **steigen** wird, während Sie die Aktie **blankoverkaufen** sollten, wenn Sie eher denken, dass der Preis **sinken** wird. Es steht Ihnen auch jederzeit frei aus einer Blankoverkauf-position direkt eine Aktie zu halten. In diesem Fall kaufen Sie einfach zwei Anteile (was automatisch passiert, wenn Sie auf "+1 Aktie Halten" klicken). Einen Anteil, welchen Sie zurückgeben, und einen, welchen Sie halten. Dasselbe gilt, wenn Sie eine Aktie halten und direkt einen Blankoverkauf tätigen.

Um fortzufahren halten Sie nun wieder eine Aktie (Sie kaufen also 2 Anteile).

	Anteil	Einstandskurs	Aktueller Kurs	Aktueller Zustand	Gewinne/Verluste Absolut (in %)	Kurswert
Bargeld						3500
Aktie	-1	1000	995	-	5 (0.5)	-995
Gesamt						**2505**

-1 Aktie Halten | Keine Aktie Halten | +1 Aktie Halten

Zurück

Imagine you have received the following price-update:

The price of the asset **decreases** by **5** and is now **995**

As you can see, the "shares" column shows a value of -1, because you have short sold one share. The buying price now displays the price at which the asset was short sold. As you can read in the price update, the price of the asset **decreased**. This gives you the chance to re-buy the asset for cheaper than what you sold it for, and it is therefore marked as a **gain** in the "gains/losses" column. As you can see, the dependence of gains and losses on the price movement now is the other way around compared to if you had bought a share.

In summary you should **buy** a share when you think the price will **increase** and you should **short sell** a share when you think the price will **decrease**. You are also free to jump directly from short selling to holding. In that case you just buy two shares (which happens automatically when you click on the "hold +1 shares" button): One which you give back from short selling and one which you now hold. The same is true when you hold a share and want to short sell directly.

To continue please now hold one share (i.e. buy two shares).

Investitions-Entscheidungen Einführung

Preis der Aktie

Im folgenden wird erklärt, wie der Preis der Aktie sich im Verlauf der Aufgabe verändert. Die Aktie ist ein Anteil an einer Firma, welche Gewinne, aber auch Verluste erzielen kann. Die Firma ist immer entweder in einem guten oder schlechten Zustand. Wenn es der Firma gut geht, macht diese mit einer höheren Wahrscheinlichkeit (65%) Gewinne, weswegen der Aktienpreis eher steigen wird. Wenn es der Firma hingegen schlecht geht, wird der Aktienpreis eher sinken (65%). Zu Beginn der Aufgabe wird zufällig bestimmt in welchem Zustand die Firma startet (mit gleichen Wahrscheinlichkeiten).

Die Art wie sich der Zustand der Firma im Verlauf der Aufgabe entwickelt sieht folgendermassen aus: Mit einer Wahrscheinlichkeit von 80% wird der Zustand gleich bleiben wie er in der letzten Runde war. Das heisst, dass ein guter Zustand dann in einer weiteren Runde guten Zustands fortgeführt würde, und ein schlechter Zustand mit weiterhin schlechtem Zustand. Mit einer Wahrscheinlichkeit von 20% verändert sich jedoch der Zustand vom guten zum schlechten oder umgekehrt. Die folgende Grafik und Tabelle zeigen auf zwei verschiedene Arten wie sich die Zustände entwickeln.

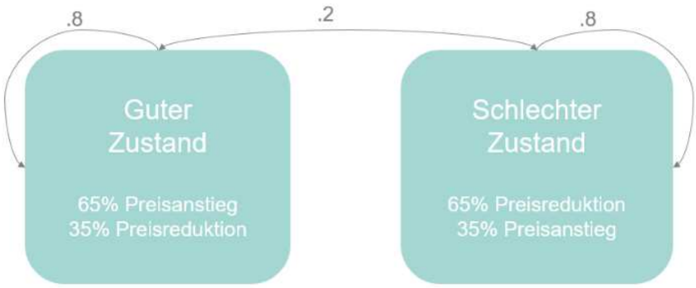

		Nachher:	
		Guter Zustand	**Schlechter Zustand**
Vorher:	**Guter Zustand**	80%	20%
	Schlechter Zustand	20%	80%

Zurück Weiter

Price of the asset In the following text we will explain how the price will develop throughout the task. The asset is a share for a company which can make gains or losses. The company is always either in a good or a bad state. When the company is doing great, there is a higher probability of the price increasing (56%). If the company is doing poorly however, there is a higher probability (65%) of the price decreasing. At the start of the task it will be decided randomly in which state the company starts (with equal probability).

The state of the company will develop in the following way: With a probability of 80% the state will remain the same as it was in the previous round. This means that a good state in one round would be continued in the next round and a bad state would be continued with a bad state. With a probability of 20% the state will flip from good to bad or vice versa. The following graph and table display this principle in two different ways.

Der "gute" Zustand bedeutet dabei nichts Anderes, als dass der Preis der Aktie in dieser Runde mit einer Wahrscheinlichkeit von 65% steigen und mit 35% fallen wird. Dies bedeutet, dass wenn die Firma 100 Perioden im guten Zustand wäre, der Preis der Aktie durchschnittlich 65 mal steigen und 35 mal fallen wird. Befindet sich die Firma im schlechten Zustand, dann steigt der Preis mit 35% Wahrscheinlichkeit an und fällt mit 65% Wahrscheinlichkeit. Er würde also von 100 Perioden im schlechten Zustand durchschnittlich 35 mal ansteigen und 65 mal fallen.

Die Grösse der Preisbewegungen können entweder 5, 10 oder 15 sein. Dies gilt für ansteigende wie auch fallende Preise. Welchen der drei möglichen Schritte der Preis beim nächsten Update macht wird zufällig und mit gleicher Wahrscheinlichkeit gewählt. Lediglich ob der Preis ansteigt oder sinkt wird durch den oben beschriebenen Mechanismus beeinflusst. Als Beispiel: Angenommen der aktuelle Preis der Aktie ist 1000 und der Preis-mechanismus führt zu einem Anstieg des Preises. Es ist nun gleich wahrscheinlich, dass der nächste Preis der Aktie 1005, 1010 oder 1015 sein wird.

Zurück Weiter

The "good" state only means that the price in this round will rise with a probability of 65% and fall with a probability of 35%. This means that if the company were in the good state for 100 periods, the price would on average increase in 65 of these rounds and decrease in 35 of them. If the company is in the "bad" state, the price rises with a probability of 35% and falls with a probability of 65%. This means that out of 100 periods, the price would on average increase in 35 of these rounds and decrease in 65 of them.

The magnitude of the price movements can be either 5, 10 or 15 points. This is true for increases as well as decreases. Which of the three possible magnitudes is realized in the next price update is decided randomly and with equal probability. Only whether the price will increase or decrease is decided by the mechanism described above. As an example: Imagine the current price of the asset is 1000 and the mechanism leads to a price increase. It is now equally as likely for the next price to be 1005, 1010 or 1015.

Investitions-Entscheidungen Einführung

Wette und Lotterie

Zwischen den Handelsentscheidungen wird Ihnen jeweils eine weitere Frage gestellt, welche nun hier genauer erklärt wird.

In jeder Runde gibt es eine **Lotterie**, in welcher Sie entweder 10 Bonuspunkte oder gar nichts erhalten können. Die Wahrscheinlichkeit, mit welcher Sie die Bonuspunkte erhalten (Gewinnwahrscheinlichkeit) wird Ihnen aber nicht gezeigt. Sie wird zufällig festgelegt, wobei alle Gewinnwahrscheinlichkeiten zwischen 0 und 100% gleich häufig auftreten.

Als Alternative zu dieser Lotterie können Sie auf den Preisanstieg in der nächsten Runde **Wetten**. Diese Wette würde dann so aussehen:

> Wenn der Preis beim nächsten Preis-Update ansteigen sollte (egal um wie viel), erhalten Sie 10 Bonuspunkte gutgeschrieben, welche am Ende in Ihre Auszahlung einfliessen.

In der Lotterie zu gewinnen oder in der Wette zu gewinnen erbringt Ihnen also gleich viele Punkte. Folgendes ist nun die Frage, welche Sie beantworten sollen:

> Welche Gewinnwahrscheinlichkeit muss die Lotterie **mindestens** haben, damit Sie nicht lieber auf den Preisanstieg wetten.

Sollte Ihnen die tatsächliche Gewinnwahrscheinlichkeit der Lotterie nicht ausreichen (die Gewinnwahrscheinlichkeit der Lotterie also *kleiner* sein als die von Ihnen angegebene mindest-Gewinnwahrscheinlichkeit), dann wetten Sie automatisch auf den Preisanstieg und erhalten Bonuspunkte falls dieser eintritt. Sollte die Gewinnwahrscheinlichkeit der Lotterie aber *gleich gross* oder sogar *grösser* sein als die von Ihnen angegebene mindest-Gewinnwahrscheinlichkeit, dann wird die Lotterie ausgespielt, und Sie erhalten mit der tatsächlichen Wahrscheinlichkeit der Lotterie die Bonuspunkte.

Im Folgenden sehen Sie den Slider mit welchem Sie Ihre Antwort abgeben können. Bevor Sie darauf klicken ist nur die Linie sichtbar. Setzen Sie den Slider nun bitte auf **60%** und klicken Sie auf "Weiter".

> *0%* ———— *100%*
>
> Ich würde lieber die Lotterie spielen, wenn diese eine Gewinnwahrscheinlichkeit von **XX%** oder höher hat.

Zurück   Weiter

Bet and Lottery Between investment decisions you will be asked to answer another question which will be explained further now.

In each round there is a **lottery**, in which you can earn either 10 bonus points or nothing. The probability with which you win the bonus points (winning probability) will not be displayed to you. It will be chosen randomly from the range of 0 to 100%, all with equal probability.

As an alternative to this lottery, you can **bet** on a price increase in the next round. This bet would look as follows:

If the price increases in the next round (no matter by how much) you will receive 10 bonus points, which are paid out at the end of the study.

Therefore, winning the bet or winning the lottery would earn you the same amount of points. The question you have to answer is the following:

What winning probability does the lottery have to have **at least** for you to prefer it to betting on the price increase.

In case the true winning probability is not high enough for you (it is smaller than the minimum probability you entered) you will automatically bet on a price increase and receive the bonus points in case of the price increasing. If the winning probability is equal or greater than the minimum you set, then you will play the lottery and receive the bonus points with the true winning probability of the lottery.

Here you can see the slider with which you can set your minimum probability. Before you click on it, only a line is visible. Please set the slider to 60% now and click "continue".

Zwei Punkte sind speziell wichtig:

Punkt 1: Ihre Antwort hängt nur von der Wahrscheinlichkeit des nächsten Preisanstieges ab. Stellen Sie sich vor, dass Sie mit Sicherheit wüssten, dass der Preis ansteigen wird: Die Lotterie müsste ebenfalls eine garantierte Auszahlung bieten (also 100% Gewinnchance), damit Sie nicht eindeutig die Wette bevorzugen. Falls Sie mit Sicherheit wüssten, dass der Preis sinkt, dann wäre Ihnen jede Lotterie lieber, die Ihnen wenigstens irgend eine Gewinnwahrscheinlichkeit bietet (oder gleichgültig gegenüber einer Lotterie, die keine Gewinnchance bietet). Sie würden in dieser Situation also 0% angeben.

Während der Aufgabe werden Sie vermutlich nicht 100% von einem Preisanstieg oder einer Preisreduktion überzeugt sein. Mit dem Slider können Sie daher den Grad Ihrer Überzeugung oder Ihrer Vermutung ausdrücken. Je weiter Sie den Slider nach links oder rechts bewegen, desto sicherer sind Sie sich, dass der Preis ansteigen oder sinken wird.

Punkt 2: Es ist für Sie immer vorteilhaft, wahrheitsgetreu anzugeben, was Sie von der Preisbewegung erwarten. Falls Sie eine *zu tiefe* Wahrscheinlichkeit angeben, kann es sein, dass Sie eine Lotterie spielen, die eine schlechtere Gewinnchance hatte, als auf den Preisanstieg zu wetten. Falls Sie einen *zu hohen* Wert angeben, kann es sein, dass Sie auf den Preisanstieg wetten, obwohl die Lotterie eine bessere Gewinnchance gehabt hätte.

Am Ende des Experiments wird Ihnen neben den Punkten, welche Sie in der Investitionsaufgabe erhalten haben auch der Ertrag aus den Wetten und Lotterien mitgeteilt. 10,0% dieses Ertrages werden zu Ihrer Gesamtpunktezahl gezählt, aus welcher sich Ihre Bonuszahlung am Ende der Studie berechnet.

Zurück Weiter

Two points are especially important: **Point 1:** Your answer only depends on the probability of a price increase. Imagine that you knew for sure that the price would increase: The lottery would have to match that guaranteed payment (100% winning probability) for you to not prefer the bet. If you knew for sure that the price will fall, you would prefer any winning probability to betting on a price increase (or be indifferent to a winning probability of 0%). Therefore you would indicate a minimum probability of 0% in this case.

During the task you will probably not be 100% convinced that there will be a price increase or decrease. With the slider you can indicate the decree of your belief. The further left or right you put the slider, the more sure you are that the price will decrease or increase.

Point 2: It is always preferable to you to truthfully report what price development you expect. If you indicate a probability that is *too low*, it can happen that you play a lottery with worse winning probability than you would expect from betting on a price increase. If you report a minimum probability that is *too high* it could happen that you bet on the price increase despite the lottery having a higher winning probability.

At the end of the experiment you will receive your earnings from the bets and lottery plays alongside the earnings from the investment task. 10% of these points will be added to the total points which your payment for the study is based on.

Investitions-Entscheidungen Einführung

Ablauf und Auszahlung

Die Aufgabe wird Ihnen in 4 "Blöcken" präsentiert. Jeder Block besteht aus 75 Runden und ist unabhängig. Das heisst, dass Ihre Entscheidungen und die Entwicklung der Preise im vorherigen Block keinerlei Einfluss auf den aktuellen Block haben. Zwischen den Blöcken können sich Einzelheiten der Aufgabe ändern, worauf Sie vor dem Start jedes Blocks jedoch hingewiesen und instruiert werden.

Nach der letzten Runde jedes Blocks wird es noch ein letztes Preis-Update geben. Falls Sie die Aktie in der letzten Runde noch hielten oder noch schuldeten, wird diese nach dem Preis-Update automatisch verkauft oder zurückgekauft, so dass Sie am Ende des Blocks nur noch Bargeld halten. Die Auszahlung geschieht dann nach folgendem Schema:

Wichtig für die Auszahlung ist die Differenz zwischen dem Wert Ihres Portfolios (Bargeld plus Aktie) zu Beginn des Blocks und dem Wert zum Schluss (also Ihr Gewinn oder Verlust). Hinzu kommen 10,0% der Punkte welche Sie für die Einschätzung der Wahrscheinlichkeiten erhalten. Von diese Summe an Punkten wird dann mit einem Wechselkurs von 5,0% zu einer Basis-Auszahlung von CHF 15 hinzugezählt, respektive abgezogen.

Ein Beispiel: Wenn die Wertsteigerung Ihres Portfolios plus der Bonus aus den Wahrscheinlichkeitsschätzungen im Verlauf der Aufgabe 100 Punkte beträgt, dann wird Ihre Auszahlung CHF 20,00 sein (CHF 15 Basis-Auszahlung plus 5,0% von 100).

Wenn Sie die Instruktionen soweit verstanden haben, können Sie auf "Weiter" klicken um 30 Übungsrunden zu starten. Diese dienen nur zu Übungszwecken, und haben keinerlei Einfluss auf Ihre Auszahlung am Ende der Studie.

Zurück Weiter

Process and Payment The task will be presented in four "blocks". Each block consists of 75 rounds and is independent from the others. This means that your decisions and the development of the price in previous blocks will have no influence on the current block. Between blocks details of the task can change, which will be indicated to you before the start of each block.

After the last round of each block there will be a final price-update. If you still hold or short sold the asset in that final round, it will be automatically sold or bought so that you will only hold cash in the end. The payment will work in the following way:

Important for the payment is only the difference between the value of your portfolio (cash plus asset) at the start of the block and the value at the end (so your gains and losses). In addition 10% of the points you earned by indicating the probabilities will be added. This sum will then be added or subtracted from the base-payment of CHF 15 with an exchange rate of 5%.

As an example: If the value gain of your portfolio plus the bonus from the probability tasks was 100 points, then your payment will be CHF 20 (CHF 15 base-payment plus % of 100 points).

If you have understood the instructions you may click "continue" to start 30 training rounds. These rounds only serve as preparation and do not influence your payment at the end of the study.


Investitions-Entscheidungen

Bitte entscheiden Sie sich jetzt!

	Anteile	Einstandskurs	Aktueller Kurs	Wahrscheinlichkeit für Anstieg	Gewinne/Verluste Absolut (in %)	Kurswert
Bargeld						2500
Aktie	0	0	1000	-	0 (0)	0
Gesamt						**2500**

-1 Aktien Halten | Keine Aktie Halten | +1 Aktie Halten

Lotterie/Wette

Sie werden entweder auf einen Preisanstieg wetten, oder eine Lotterie spielen. Wie hoch müsste die Gewinnwahrscheinlichkeit der Lotterie sein, damit Sie diese der Wette vorziehen?

Bitte entscheiden Sie sich jetzt!

0% *100%*

Ich würde lieber die Lotterie spielen, wenn diese eine Gewinnwahrscheinlichkeit von **78%** oder höher hat.

Weiter

Update

Der Preis der Aktie **sinkt** um **5** und beträgt nun **995**

Weiter

Appendix B

Appendix Manuscript II

B.1 Extended Procedure and Methods

Participants could invest into an asset with a constant upward or downward drift. The asset's price would change in each round and increase either with a probability of .65 (upward drift) or .35 (downward drift). The magnitude of the price changes were 5, 10 or 15 points with equal probability. This means that only the direction of a price update was informative, but not the magnitude of the change. A price path always started at 10'000 points. Participants' portfolios could always consist of holding between four and "negative four" (i.e. short sell four) shares of the asset and an initial portfolio was allocated randomly with equal probabilities. Note first, that a risk neutral profit maximizing Bayesian investor would always use the full potential of the allowed portfolio, that is, they would always either hold or short sell four shares. As participants were fully informed about the price generating mechanism it was always possible by Bayesian probability updating to calculate the probability of an upward drift and thereby an upcoming price increase. However, note also that this calculation would not be necessary for a risk neutral investor, as a simple counting heuristic ("were there more up- or downward moves?") would suffice to know whether an up- or downward trend was more likely.

The sequence of events in a (standard) round of the experiment would look as follows: First, participants were asked to decide whether to buy or sell any shares. Next, they used a slider (ranging from 0 to 100) to indicate how probable they thought a price increase from this to the next round to be. Lastly, the price update would be displayed in a sentence (e.g. "The assets price decreased by 5 and is now 995").

Eight price paths were generated per participant, each of which having an upward drift with a probability of .5 and a downward drift otherwise. Over the course of the experiment, participants experience each of the eight price paths in each of the three conditions (described below), resulting in 24 path-condition combinations ("blocks"). This threefold presentation of the price paths was done because the noise of the randomly generated price paths could otherwise drown out any effects of the treatments. To minimize recognition effects, the order of the blocks were randomized for each participant. At the start of each block participants were shown the block number (e.g. "Block 9 of 24") and were informed about the current condition.

Each "block" was comprised of two phases: First participants received a randomly allocated portfolio and started with three rounds in which they could not trade. This means that, when making a first investment decision on a price path, participants would already have observed three rounds and thus gathered some information about the current asset's drift. This also meant that at the end of the first phase participants would always have made an initial gain or loss. The subsequent five rounds changed depending on the experimental condition: In the *Baseline* condition participants could change their portfolio at the start of each round. In the *Blocked Trades* condition participants had to decide upon one portfolio which was then held for the upcoming five rounds. Lastly, the *Delayed Information* condition also required participants to commit to their portfolio. Here however, the price updates were only displayed in list format at the end of the final five rounds and participants reported their beliefs about an upcoming price increase only at once afterwards. Finally, independent of the condition, participants could make one last investment for a final round. Their belief in a price increase was also elicited for this final decision.

The online experiment was implemented using oTree (Chen et al., 2016) and conducted through the prolific.co platform. The main task was preceded by a detailed and interactive instruction, a comprehension quiz as well as multiple training rounds. This was to make sure that participants had fully understood the task as well as the option of short selling the asset. The study took around one hour to complete. Participants were incentivized by paying out the returns of a randomly selected block and earned an average of $13.55. The final sample consists of 315 participants (156 female, mean age 33.36, SD 10.86) with 68.57% of the sample declaring having at least some previous experience with financial investments.

B.2 Instructions and Screenshots

Investment Decision Experiment

Welcome to this Experiment. It will take around one hour to complete and you will have the option to earn an additional bonus payment depending on your performance. This is the tutorial part of the experiment, where you will be shown how to make your investments, how the asset you're investing in works, how you will be asked about your expectations, and how this all affects the bonus payment at the end. You will then be able to play some training rounds to familiarize yourself with the task before starting the actual experiment.

Please try to stay focused during the whole experiment and do not leave this browser window. If possible, complete the study in the fullscreen mode of your browser. To enter fullscreen mode on windows, press the ***F11*** key on your keyboard. If you are using a Mac you can press ***Cmd + Ctrl + F***.

Lastly before we start, there are rare cases in which browsers can load scripts in the wrong order or have other "hiccups" which would prevent you from continuing. If you ever become stuck or think that something is wrong, you can always reload the page (in most browsers this can be done using the ***F5*** key).

In the unlikely event that you get stuck and can't resolve the issue by reloading the page, always feel free to send an email to k.trutmann@unibas.ch.

Please click the continue button to start the tutorial.

Continue

Investment Decision Experiment

Making investment decisions

In this experiment you will make investment decisions in 24 blocks with 9 rounds each. There will be an asset you can invest in. At the start of a block, the asset will always have a value of 1000 Points. In each round you can decide how many shares of this asset you want to hold. You can hold a maximum of 4 shares and a minimum of -4. How holding negative shares works will be explained in the following pages. Your cash amount will always be enough to buy the maximum of 4 shares.

Below you can see the screen with which you will interact during the task. The different columns are described here. Please make sure you have looked at each row and understood its meaning before proceeding.

- **Column I:** Shows how many shares of the asset you are holding. '+3' thus means you are holding three shares and each price movement will affect the value of your portfolio by three times the movement. If this collumn shows '0' this means you are not holding any shares and the value of your portfolio (which in this case would be purely cash) is not affected by the next price movement.
- **Column II:** This column shows the average buying price of your shares. If the current price is above this value, your investment will have made a profit. It will be reset whenever you sell all your shares, or switch from short selling (holding negative shares) to holding shares or vice versa.
- **Column III:** This column shows the current price of the asset. If you decide to buy or sell in this round, this is the price at which the asset will be bought or sold. How this price develops over time will be explained later.
- **Column IV:** This column shows your returns for this investment (how much value your investment has gained or lost). The value in parenthesis shows the same idea in percentages, where 100% would mean you have doubled your investment and -50% means it has lost half its value.
- **Column V:** This column shows the current values in your portfolio. It shows your cash, the value of the assets you are holding (i.e. its price times the amount of assets) and the total amount.

Beneath the table you can find the interface with which you can make your trades. To make yourself familiar with the process, please **buy two shares now** by entering '2' in the field and clicking the 'Buy' button. In the experiment you will then have to confirm that order. Here it is enough to just click 'Buy'.
Once you have done that the text will continue here and you will be able to scroll down.

	I	II	III	IV	V
	Shares	**Average Buying Price**	**Current Price**	**Returns (in %)**	**Total Value**
Cash					5000
Asset	+2	1000	1000	0 (0.0)	2000
Total					**7000**

Your Order:

Amount

Buy

Sell

Continue

confirm that order. Here it is enough to just click 'Buy'.
Once you have done that the text will continue here and you will be able to scroll down.

Very good. You now doubled your investment from holding two to holding four shares of the asset. The value of the purchased shares was subtracted from the cash field.

To keep you on track with the time there will be a blinking warning after 7 seconds in the task to tell you to make a decision. If you make your decisions before this warning appears, you will be able to finish the experiment in the recommended time. However, don't let this put any pressure on you. After some rounds 7 seconds will be enough time to make a well considered decision. If you want you can also use some of the training rounds to get a feeling for the timing by waiting until the warning appears.

This is what the warning will look like: "Please decide now!"

After deciding on your trade you will see an update about the price for 5 seconds. How these updates are decided will be described on the following page. This is what such an update could look like:

> The price of the asset **increases** by **5** and is now **1005**

Please take some time to look at how the table has changed, now that you bought additional shares and the price has moved.

A little hint on the side: If you are happy with your portfolio as it is (i.e. you do not want to sell or buy anything) you can either just click the 'Buy' or 'Sell' button without typing anything into the field or you can use the 'Continue' button below.

If you have understood everything so far, sell all your shares by entering the amount of shares you are holding and clicking on the 'Sell' button.

	I	II	III	IV	V
	Shares	**Average Buying Price**	**Current Price**	**Returns (in %)**	**Total Value**
Cash					3000
Asset	+4	1000	1005	20 (0.5)	4020
Total					**7020**

Your Order:

4

Buy

Sell

Continue

Investment Decision Experiment

Short selling

When investing a profit is made by buying an asset as cheaply as possible and selling it for as much as possible. However, the order in which those actions are taken do not play a role. This means you can sell shares that you do not own yet, usually by borrowing them from a third person (in this case this would be the experiments bank, lending you the shares you need). You then "owe" these shares, which is why it is displayed as negative shares in the interface. This is what is called "short selling".

Imagine having sold one share (that you borrowed from the bank) for 100 points. You will have received these 100 points in cash, but you also still have to buy one share to give back to the bank. When short selling you thus make a profit when the price of the asset **decreases**, because then you can buy the asset back for cheaper than what you have sold it for.

Imagine now that in our example the price drops by 20 points. You are now able to buy one share for 80 points and return it to the bank. Having received 100 points earlier and spent 80 points now, you have made a profit of 20 points. On the other hand this means that a price increase is **not** favorable when you have short sold the asset, since you would then have to spend more money to buy it back than what you sold it for.

If you think that you have understood this principle, please look at the table below (which has been reset to holding 2 shares) and *sell four shares*, thus selling two more shares than you currently have. The tutorial will then continue here.

	Shares	Average Buying Price	Current Price	Returns (in %)	Total Value
Cash					5000
Asset	+2	1000	1000	0 (0.0)	2000
Total					**5000**

Your Order:

Amount

Buy

Sell

Continue

Back

If you think that you have understood this principle, please look at the table below (which has been reset to holding 2 shares) and *sell four shares*, thus selling two more shares than you currently have. The tutorial will then continue here.

You receive the following price update:

> The price of the asset **decreases** by **5** and is now **995**

As you can see, the "Shares" column now shows negative two shares. The average buying price is now displayed as the average price at which you made your short sales. As you can see from the price update, the price of the asset has **decreased**, which is favorable for your short sale investment. This is also reflected in the "Returns" column, which shows a positive return. When short selling, the dependence of your returns on the price movements is therefore reversed from that of a standard investment.

In summary, you should **buy** shares when you expect the price to **increase** and **short sell** the asset when you expect the price to **decrease**. You are also free to "jump" from holding shares to short selling shares and vice versa at any time.

To continue please make an order such that you will hold 4 shares.

	Shares	Average Buying Price	Current Price	Returns (in %)	Total Value
Cash					6000
Asset	-2	1000	995	10 (1)	1010
Total					**5005**

Your Order:

Amount

Buy

Sell

Continue

Back

Investment Decision Experiment

Price development

We now turn to the way in which the assets price can develop. To illustrate, imagine that the asset you are buying is a stock that is based on a company which can make a profit but can also lose money. You will either be able to invest in the stock of a company with a **65%** probability of a price increase, resulting in a general upwards trend, or one with a **35%** probability of a price increase, resulting in a general downward trend. Remember that even if you are given the stock from a company where the price is more likely to decrease, you can still profit from this by short selling the asset.

In each round, the price can move by either **5, 10 or 15** points. This is true for both increases as well as decreases. Which of these magnitudes the next price move will take on is determined by chance and with equal probability. Thus, whether the price in- or decreases is determined by the kind of asset while the magnitude of the change is determined randomly.

As an example: Assuming the current price of the asset is 1000 points and the price will increase. It is then equally likely that the next price will be 1005, 1010 or 1015.

You will be investing under different conditions throughout different "blocks" in the experiment. The page that indicates when a new price path is started will also tell you what the conditions for the upcoming block are.

- The first 3 rounds will always be "pre-decided", meaning you are handed a fixed portfolio and you can only observe the price movement.
- After that, in some cases you will be able to make a decision for each individual round of the remaining rounds.
- In some cases you will have to make an investment that will then be "fixed" for the next 5 rounds (i.e. if you decide to hold 2 shares, you will hold 2 shares throughout the next 5 price movements). There will be an indicator below the table whenever you are deciding for more than one round.
- In some of the cases in which you can not freely trade you will be informed about the price movements in each round, in others you will only be informed about the price movements after they all already happened.

It therefore makes sense to **pay good attention** to which decision you are currently making. You will be able to experience all possible conditions throughout the 3 training blocks.

Back Continue

Investment Decision Experiment

Expectations about Price Movements

Between your trading decisions you will also be asked about your expectations about the further development of the price. You will be asked to answer the following question:

> How likely do you think is it that the price will be higher in the next round than it is now?

If you answer with *50%* you indicate that you have no idea whether the price could increase or decrease (as would for example be the case at the start of a block). After some rounds you will however likely have some idea about what kind of asset this is and how its price develops.
Thus, shifting the slider with which you provide your answer *towards the right (towards 100%)* indicates that you belief the price will *increase* (whereby setting it to 100% would mean you are absolutely convinced that the price will increase in the next round).
Shifting the slider *towards the left (towards 0%)* on the other hand indicates that you think the price will *decrease in the next round* (whereby 0% would indicate that you are absolutely convinced that the price will *not* increase, and therefore decrease).

Below you can see the slider with which you can provide your answer. Before clicking on it, you will only see a line. Please place the slider on **60%** and click the "Continue" button.

> *0%* ——— *100%*
>
> "I think the probability of a price increase in the next round is **XX%**."

Back Continue

Investment Decision Experiment

Expectations about Price Movements

Between your trading decisions you will also be asked about your expectations about the further development of the price. You will be asked to answer the following question:

> How likely do you think is it that the price will be higher in the next round than it is now?

If you answer with *50%* you indicate that you have no idea whether the price could increase or decrease (as would for example be the case at the start of a block). After some rounds you will however likely have some idea about what kind of asset this is and how its price develops.
Thus, shifting the slider with which you provide your answer *towards the right (towards 100%)* indicates that you belief the price will *increase* (whereby setting it to 100% would mean you are absolutely convinced that the price will increase in the next round).
Shifting the slider *towards the left (towards 0%)* on the other hand indicates that you think the price will *decrease in the next round* (whereby 0% would indicate that you are absolutely convinced that the price will *not* increase, and therefore decrease).

Below you can see the slider with which you can provide your answer. Before clicking on it, you will only see a line. Please place the slider on **60%** and click the "Continue" button.

0%  *100%*

"I think the probability of a price increase in the next round is **60%**."

During the experiment you should not worry about the exact number you are providing but rather use the slider and the number to provide your best estimate of what you expect to happen.

Back Continue

Investment Decision Experiment

Blocks and Payoff

The task will consist of **24 price paths** to invest in, each of which consisting 9 rounds. Each price path is independent from that before and only determined by the probability of a price increase of the given asset. You will be informed whenever a new price path starts.

After the last round of each block, there will be one last price update. If you are holding any shares after this last price update, they will automatically be sold, such that you will only hold cash in the end. Similarly any short positions will automatically be bought back. The payoff will be calculated in the following way:

The payoff will be based on **the difference** between the value of your portfolio (cash plus assets) *at the start* of each block and the value of your portfolio *at the end* of the block. This therefore reflects the winnings or losses of your investments. In the end, you will be presented with a list with all your winnings in each block. One of the blocks' winnings (or losses) will be randomly chosen to be payed out.
This chosen amount will then be added or subtracted from the base payoff of $2.5 with an exchange rate of 1.5%. Additionally you will receive the $8.13 for completing the study, no matter your performance.

As an example:
If in the block that was randomly chosen to be payed out you have gained 100 points, your payoff will be $12.13. It will consist of $8.13 for completing the study, $2.5 base payoff plus 1.5% of the 100 points you earned.

If you have understood these instructions, please click the "Continue" button to start a short comprehension quiz which will be followed by the 3 training blocks.

Back Continue

Comprehension Quiz

Here you find some questions about the experiment. Please try to answer them to the best of your understanding.

A decrease in the assets price is...

- ...more likely during an upward trend.
- ...more likely during a downward trend.
- ...equally likely during an upward as during a downward trend.

Next

Comprehension Quiz

Here you find some questions about the experiment. Please try to answer them to the best of your understanding.

A price increase of 15 points is...

- ...always equally as likely as an increase of 10.
- ...more likely than a price increase of 10 if the asset has an upward trend.
- ...always equally as likely as a decrease of 15 points.

Next

Comprehension Quiz

Here you find some questions about the experiment. Please try to answer them to the best of your understanding.

When the asset has a downward trend...

- ...there is a 50% chance that the price will in- or decrease.
- ...there is a 35% chance that the price will increase and a 65% chance that it will decrease.
- ...there is a 35% chance that the price will decrease and a 65% chance that it will increase.

Next

Comprehension Quiz

Here you find some questions about the experiment. Please try to answer them to the best of your understanding.

Imagine that you have short sold the asset by two shares (i.e. you hold -2 shares). Which statement is correct?

- If the asset has an upward trend, the value of your portfolio is more likely to increase now.
- If the price decreases by 10 points, the value of your portfolio will decrease by 20 points.
- If the price decreases by 15 points, the value of your portfolio will increase by 30 points.

Next

Training Block 1 of 3

A new block of the training rounds starts now.
These rounds will not influence your final payoff.

Everything will be reset to its starting value, including the price of the asset. It is again randomly determined whether the asset will have **a 35% or a 65% probability** for a price increase.

- Your portfolio will be fixed for the first 3 rounds.
- After that you will be asked to **make a decision for the subsequent 5 rounds**.
- You will be informed about the price changes **at the end of these 5 rounds**.
- The price updates will be presented to you **in list-form**.
- Then you can make one last decision in the last round.

Below you can find a graphical representation of the decisions you will make. Each arch indicates the length for which you will make the decision in that round. Only the first investment decision is pre-made, the other decisions are up to you.

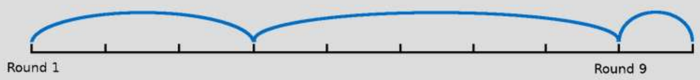

Click the "Next" button to start this block.

Next

Investment-decisions

Please continue now!

	Shares	Average Buying Price	Current Price	Returns (in %)	Total Value
Cash					7068
Asset	-2	1000.0	1000	0 (0.0)	-2000
Total					**5068**

Next

Expectation

Please use the slider below to indicate how likely you think it is that the price will be higher in the next round than it is now.

- **50%** means you think a price increase and a price decrease to be equally likely.
- **Moving the slider towards 0%** means you believe it to be more probable that the price will decrease.
- **Moving the slider towards 100%** means you believe it to be more probable that the price will increase.

0% *100%*

"I think the probability of a price increase in the next round is **XX%**."

Continue

Update

The assets price **decreased** by **15** and is now **985**

Next

Investment-decisions

This is a training round!

	Shares	Average Buying Price	Current Price	Returns (in %)	Total Value
Cash					7068
Asset	-2	1000.0	980	40 (4.0)	-1960
Total					**5108**

After this decision, your portfolio will stay the same for the next **5 rounds**!

Your Order:

Amount

Buy

Sell

Continue

Update

5 rounds have now passed.

The price has developped in the following way:

The assets price **decreased** by **10** and is now **970**

The assets price **decreased** by **15** and is now **955**

The assets price **decreased** by **5** and is now **950**

The assets price **decreased** by **5** and is now **945**

The assets price **decreased** by **15** and is now **930**

Next

Investment-decisions

This is a training round!

	Shares	Average Buying Price	Current Price	Returns (in %)	Total Value
Cash					7068
Asset	-2	1000.0	930	140 (14.0)	-1860
Total					**5208**

Your Order:

Amount

Buy

Sell

Continue

Training Block 2 of 3

A new block of the training rounds starts now.
These rounds will not influence your final payoff.

Everything will be reset to its starting value, including the price of the asset. It is again randomly determined whether the asset will have **a 35% or a 65% probability** for a price increase.

- Your portfolio will be fixed for the first 3 rounds.
- After that you will be asked to **make a decision for the subsequent 5 rounds**.
- You will be informed about the price changes **in each round**.
- Then you can make one last decision in the last round.

Below you can find a graphical representation of the decisions you will make. Each arch indicates the length for which you will make the decision in that round. Only the first investment decision is pre-made, the other decisions are up to you.

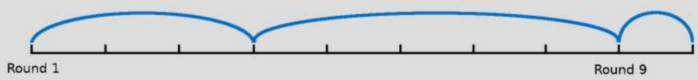

Click the "Next" button to start this block.

Next

Investment-decisions

Please continue now!

	Shares	Average Buying Price	Current Price	Returns (in %)	Total Value
Cash					1086
Asset	+4	1000.0	1045	180 (18.0)	4180
Total					**5266**

Next

Training Block 3 of 3

A new block of the training rounds starts now.
These rounds will not influence your final payoff.

Everything will be reset to its starting value, including the price of the asset. It is again randomly determined whether the asset will have **a 35% or a 65% probability** for a price increase.

- Your portfolio will be fixed for the first 3 rounds.
- After that you will have **full control** of your portfolio until the end of the block.

Below you can find a graphical representation of the decisions you will make. Each arch indicates the length for which you will make the decision in that round. Only the first investment decision is pre-made, the other decisions are up to you.

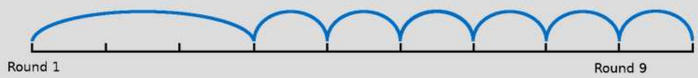

Click the "Next" button to start this block.

Next

Investment-decisions

This is a training round! **Please decide now!**

	Shares	Average Buying Price	Current Price	Returns (in %)	Total Value
Cash					8946
Asset	-4	1000.0	1000	0 (0.0)	-4000
Total					**4946**

Buy 8 shares

Change Order Confirm Order

End of the tutorial

This was the last training Block. The experiment will now begin.
The following rounds are relevant for your final payoff.

Please click the "Next" button whenever you are ready.

Next

B.3 Additional Tables

Table B.3.1: Treatment Effects on Belief Measure

Baseline	18.475***
	(0.543)
Blocked Trades	-0.389
	(0.298)
Delayed Information	-0.678**
	(0.258)
R^2	< .001
Obs. (Clusters)	7560 (315, 8)

Note. * = $p < .05$, ** = $p < .01$ *** = $p < .001$. This is the table underlying Figure 1 in the manuscript. OLS regression using one sided p-values with standard errors (in parentheses) clustered on participant and path level; Treatments are dummy coded;

Table B.3.2: Effect of Belief on Investments

Belief	.053***
	(.002)
Drift	1.782***
	(.023)
Interaction	-.017***
	(.005)
AIC	60623.16
Obs.	17640

Note. * = $p < .05$, ** = $p < .01$ *** = $p < .001$. This table reports the results of an ordered logistic regression, testing the effect of beliefs, drift and their interaction on participants' investment decision. Omitted are the intercept estimates (all significant). The model contains random effects per participants.

Table B.3.3: Treatment Effects on Investments in Round Three

Baseline	2.737***
	(0.046)
Blocked Trades	0.073*
	(0.034)
Delayed Information	0.11**
	(0.046)
R^2	< .001
Obs. (Clusters)	7560 (315, 8)

Note. * = $p < .05$, ** = $p < .01$ *** = $p < .001$. This model compares the (absolute) number of shares held at the end of the third round, before the conditions differed. As participants invest more when facing longer investment and evaluation times, these results replicate the known findings on myopic loss aversion. OLS regression using one sided p-values with standard errors (in parentheses) clustered on participant and path level; Treatments are dummy coded;

Table B.3.4: Treatment Effects including Drift

	(1)	(2)	(3)
Baseline	18.829***	-0.264**	-0.247***
	(0.582)	(0.101)	(0.073)
Blocked Trades	-0.625*	0.051	0.016
	(0.291)	(0.052)	(0.041)
Delayed Information	-0.572**	-0.069	-0.049
	(0.205)	(0.079)	(0.053)
Upward Drift	-0.715	1.594***	1.387***
	(0.671)	(0.11)	(0.097)
Drift × Blocked Trades	0.477	-0.221***	-0.186**
	(0.365)	(0.071)	(0.075)
Drift × Delayed Info	-0.215	-0.038	-0.062
	(0.679)	(0.135)	(0.089)
R^2	0.0009	0.062	-
AIC	-	-	7779.78
Obs. (Clusters)	7560 (315, 8)	7560 (315, 8)	6161 (315, 8)

Note. * = $p < .05$, ** = $p < .01$ *** = $p < .001$. One sided p-values with standard errors (in parentheses) clustered on participant and path level; Treatments are dummy coded; Model (1): OLS regression with Distance to the Bayesian Belief as outcome variable; Model (2): OLS regression with Drift Hit Rate as outcome variable; Model (3): Logistic regression with Binary Drift Hit Rate as outcome variable.

Appendix C

Appendix Manuscript III

C.1 Additional Results

Table C.1.1: Likelihood of Investing in an Investment Option

Variable	(1) Set I	(2) Set II	(3) Set III	(4) Set IV
Consultant	0.213*	−0.0950	−0.116	0.202
	(0.126)	(0.109)	(0.139)	(0.133)
Age	−0.00271	0.00590	−0.00610	−0.000517
	(0.00464)	(0.00546)	(0.00545)	(0.00615)
Gender	−0.448***	0.101	0.369***	−0.0160
	(0.111)	(0.110)	(0.143)	(0.133)
Constant	−0.524***	−0.816***	−0.389	−1.107***
	(0.170)	(0.228)	(0.240)	(0.244)
Observations	2,505	2,505	2,505	2,505
Pseudo R^2	0.0240	0.00280	0.0210	0.00500
Log-likelihood	−1,363	−1,468	−1,510	−1,052

Note. Probit regression with robust standard errors in parentheses clustered on the participant level. Two-sided t statistics, $^{*}p < .1$; $^{***}p < .01$. *Outcome variables: Set I/II/III/IV*, dummy variable taking on the value of 1 if the investment option is from set I, II, III, or IV, 0 otherwise. *Independent variables: Consultant*, dummy variable taking on the value of 1 if the participant is in the *consultant condition*, 0 otherwise; *age*, measured in years; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise.

Table C.1.2: Ivested Amount

Variable	Invested amount (100 k)
Round	1.130***
	(0.198)
Consultant	5.960
	(4.185)
Set I	10.94**
	(4.679)
Set II	1.227
	(4.165)
Set III	19.34***
	(4.941)
Age	−0.274
	(0.204)
Gender	−7.093
	(4.382)
Constant	42.15***
	(9.682)
Observations	2,444
R^2	0.096

Note. Ordinary least squares regression with robust standard errors in parentheses clustered on the participant level. Two-sided t statistics, $^{**}p < .05$; $^{***}p < .01$. *Outcome variable: Invested amount*, measured in 100,000 experimental currency units. *Independent variables: Round*, number of rounds played, from 1 to 15; *consultant*, dummy variable taking on the value of 1 if the participant is in the *consultant condition*, 0 otherwise; *set I/II/III*, dummy variable taking on the value of 1 if the investment option is from set I, II, or III, 0 otherwise; *age*, measured in years; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise.

Table C.1.3: Profits

Variable	(1) Profit per round	(2)	(3) Sum of profits	(4)
Consultant		−943.8		−15,524
		(20,163)		(302,357)
Age	863.4	871.4	12,785	12,915
	(1,020)	(1,010)	(15,172)	(15,078)
Gender	50,610**	50,716**	752,142**	753,899**
	(23,861)	(23,637)	(356,518)	(354,249)
Constant	−56,892	−56,808	−842,222	−840,822
	(50,787)	(51,060)	(755,554)	(761,762)
Observations	2,480	2,480	167	167
R^2	0.003	0.003	0.031	0.031

Note. Ordinary least squares regression with robust standard errors in parentheses clustered on the participant level. Two-sided t statistics, $^{**}p < .05$. *Outcome variables: Profit per round*, profit made in round measured in experimental currency units (ECU); *sum of profits*, profits made over all 15 rounds measured in ECU. *Independent variables: Consultant*, dummy variable taking on the value of 1 if the participant is in the consultant condition, 0 otherwise; *age*, measured in years; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise.

Table C.1.4: Cumulative Prior Losses and Likelihood of Changing

Variable	(1) Manager	(2) Consultant	(3) Pooled
Proportion rounds loss $(t-1)$	-0.149	0.288*	-0.133
	(0.162)	(0.152)	(0.158)
Consultant			-0.0747
			(0.161)
Consultant × Proportion Rounds Loss			0.398*
			(0.218)
Age	0.000296	-0.00730	-0.00339
	(0.00881)	(0.00680)	(0.00584)
Gender	-0.182	-0.0309	-0.108
	(0.139)	(0.125)	(0.0927)
Constant	-0.795**	-0.674**	-0.702***
	(0.332)	(0.285)	(0.231)
Observations	1,025	843	1,868
Pseudo R^2	0.00458	0.00739	0.00705
Log-likelihood	-471.0	-436.5	-909.1

Note. Probit regression with robust standard errors in parentheses clustered on the participant level. Two-sided t statistics, *$p < .1$; **$p < .05$; ***$p < .01$. *Outcome variable: Change*, dummy variable taking on the value of 1 if the investment option in the current round is different from in the previous round, 0 otherwise. *Independent variables: Proportion rounds loss*, number of rounds with loss in this investment option over number of rounds invested in this investment option; *consultant*, dummy variable taking on the value of 1 if the participant is in the consultant condition, 0 otherwise; *age*, measured in years; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise.

Table C.1.5: Main Results—Robustness Check, Risk Preferences

Variable	(1) Manager	(2) Consultant	(3) Pooled
Loss $(t-1)$	−0.00541	0.430***	−0.0121
	(0.107)	(0.153)	(0.108)
Consultant			−0.117
			(0.153)
Consultant × Loss $(t-1)$			0.427**
			(0.185)
Age	−0.00367	−0.00255	−0.00458
	(0.00911)	(0.00697)	(0.00642)
Gender	−0.238	0.0309	−0.130
	(0.177)	(0.140)	(0.121)
Risk preferences, investment	0.0289	0.0930**	0.0501
	(0.0409)	(0.0375)	(0.0315)
Constant	−0.771	−1.311***	−0.868**
	(0.473)	(0.424)	(0.377)
Observations	1,247	1,034	2,281
Pseudo R^2	0.00923	0.0256	0.0172
Log-likelihood	−610.8	−540.0	−1,152

Note. Probit regression with robust standard errors in parentheses clustered on the participant level. Two-sided t statistics, $^{**}p < .05$; $^{***}p < .01$. *Outcome variable: Change*, dummy variable taking on the value of 1 if the investment option in the current round is different from in the previous round, 0 otherwise. *Independent variables: Loss (t − 1)*, dummy variable taking on the value of 1 if the result of the previous round was a loss, 0 otherwise; *consultant*, dummy variable taking on the value of 1 if the participant is in the consultant condition, 0 otherwise. *age*, measured in years; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise. *Risk preferences, investment*, self-reported risk preferences for investments on a scale of 0 (unlikely to take risk) to 10 (very likely to take risk).

Exploratory Analisys

Gender. As women earned significantly more in this investment task, it is of interest to check whether gender also affects the behavioral differences in the two conditions. Female participants changed their investment option in 16.0% of the rounds in the *manager condition* and in 20.5% of the rounds under the consultant frame. Men did so in 23.1% of the cases in the *manager condition* and in 24.7% of the rounds when they had to hire a consultant. Thus there exists a gender difference in the general likelihood of changing, where men changed the chosen investment option significantly more often (23.8%) than women (18.2%; two sided t-test, $p = .043$, robust standard errors clustered by individual).

Table C.1.6 reports for male (model 1) and female (model 2) participants separately the results of a probit regression checking for treatment and context (gain or loss) differences in the likelihood of changing. In Model (3) we directly test for the interaction of the treatment and gender differences. The overall level of changing the investment option among women is significantly lower in both treatments compared to that of men. Note that no interaction with treatment and gender is significant, which indicates that the treatment effect is similar between male and female participants. Testing whether the treatment differences are affected by gender, we include a three-way interaction term between treatment condition dummy, prior experienced loss, and gender. The coefficient of this three-way interaction is not statistically significant. In sum, we conclude that we do not find evidence that the higher likelihood of changing after prior losses in the *consultant condition* is affected by gender differences, albeit on a lower base level of change for female participants.

Table C.1.6: The Role of Gender

Variable	(1) Male	(2) Female	(3) Pooled
Loss $(t-1)$	-0.0206	-0.199	-0.0174
	(0.200)	(0.202)	(0.200)
Consultant	-0.0850	0.0472	-0.0860
	(0.158)	(0.106)	(0.157)
Consultant $\times$ Loss $(t-1)$	0.180	0.678***	0.181
	(0.252)	(0.257)	(0.251)
Gender			-0.456**
			(0.211)
Consultant $\times$ Gender			-0.195
			(0.291)
Loss $(t-1)$ $\times$ Gender			0.247
			(0.217)
Consultant $\times$ Loss $(t-1)$ $\times$ Gender			0.499
			(0.359)
Age	-0.00666	-0.0111	-0.00755
	(0.00722)	(0.00685)	(0.00597)
Constant	-0.469	-0.784**	-0.437*
	(0.298)	(0.312)	(0.256)
Controls	Yes	Yes	Yes
Observations	1,362	919	2,281
Pseudo R^2	0.00434	0.0501	0.0252
Log-likelihood	-734.6	-408.0	$-1,143$

Note. Probit regression with robust standard errors in parentheses clustered on the participant level. Two-sided t statistics, *$p < .1$; **$p < .05$; ***$p < .01$. *Outcome variable: Change*, dummy variable taking on the value of 1 if the investment option in the current round differs from the previous round, 0 otherwise. *Independent variables: Loss $(t-1)$*, dummy variable taking on the value of 1 if the result of the previous round was a loss, 0 otherwise; *consultant*, dummy variable taking on the value of 1 if the participant is in the *consultant condition*, 0 otherwise; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise; *age*, measured in years.

Education. Having a higher education or even a business degree might affect the way one makes decisions. Having at least a tertiary degree (model 2 in Table C.1.7) nor a business background (model 3 in Table C.1.7) reduces the treatment differences in the likelihood of switching the investment option due to prior losses, an education with.

Table C.1.7: Age, Educational Background, and Treatment Differences

Variable	(1) Proportion rounds loss	(2) Higher education	(3) Business background
Proportion rounds loss	−0.137		
	(0.159)		
Consultant × Proportion Rounds loss	1.655***		
	(0.501)		
Consultant × Age	0.0113		
	(0.0137)		
Consultant × Proportion Rounds Loss × Age	−0.0334**		
	(0.0136)		
Loss($t-1$)		0.0248	0.0210
		(0.112)	(0.115)
Consultant × Loss ($t-1$)		0.638***	0.297
		(0.204)	(0.190)
Higher education		0.203	
		(0.129)	
Consultant × Loss ($t-1$) × Higher Education		−0.393	
		(0.209)	
Business background			0.159
			(0.142)
Consultant × Loss ($t-1$) × Business Background			0.256
			(0.241)
Consultant	−0.507	−0.118	−0.116
	(0.471)	(0.151)	(0.155)

Table continued on next page

Variable	(1) Proportion rounds loss	(2) Higher education	(3) Business background
Age	0.00105	−0.00707	−0.00700
	(0.00875)	(0.00618)	(0.00598)
Gender	−0.112	−0.235**	−0.150
	(0.0940)	(0.0988)	(0.107)
Constant	−0.855***	−0.663**	−0.635**
	(0.317)	(0.281)	(0.272)
Controls	Yes	Yes	Yes
Observations	1,868	2,281	2,281
Pseudo R^2	0.0117	0.0215	0.0245
Log-likelihood	−904.9	−1,147	−1,144

Note. Probit regression with robust standard errors in parentheses clustered on the participant level. Two-sided t statistics, $^{**}p < .05$; $^{***}p < .01$. *Outcome variable: Change*, dummy variable taking on the value of 1 if the investment option in the current round is different from in the previous round, 0 otherwise. *Independent variables: Consultant*, dummy variable taking on the value of 1 if the participant is in the *consultant condition*, 0 otherwise; *proportion rounds loss*, the number of rounds in which the participant experienced a loss from this investment option over the number of rounds in which they invested in this option; *Loss $(t-1)$*, dummy variable taking on the value of 1 if the result of the previous round was a loss, 0 otherwise; *higher education*, dummy variable taking on the value of 1 if the participant's highest education is a tertiary degree, 0 otherwise; *business background*, dummy variable taking on the value of 1 if the participant's education included some business aspects, 0 otherwise; *age*, measured in years; *gender*, dummy variable taking on the value of 1 if the participant is female, 0 otherwise.

C.2 Invitation Email

Dear sir or madam, you always wanted to test your investment skills? In that case I would like to draw your attention to my thesis in the field of economic psychology. There, I investigate the investment behavior of decision makers. For this thesis I am conducting a study. The study takes around 30 minutes and you can participate online. In this way I wanted to ask you kindly whether you would be willing to participate in this experiment. You can find the study under the following link [EXPERIMENT LINK]. Among the five best participants I will also raffle off 5 x CHF 20 vouchers. Whoever wants to participate at this raffle can write me an email to [EMAIL], where I will keep your email address separate from your data. The data will be saved using a random key which can only be recognized by the participant. In the end I will send the list of rankings to the participants including only those random keys. I am looking forward to your participation. Thank you and best regards.

C.3 Task Screenshots

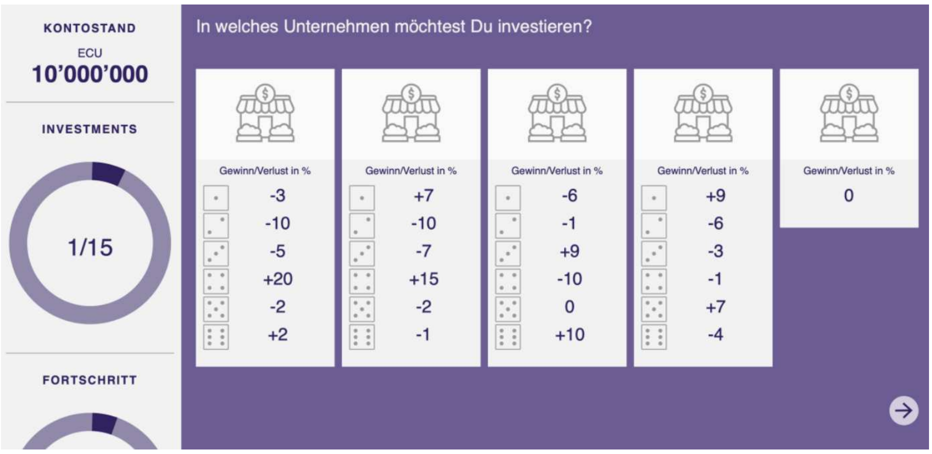

Figure C.3.1: Screenshot of the decisions screen in the *manager condition* where participants saw an overview of their investment options. The sidebar shows the current balance ("Kontostand") as well as the investments and progress in the study so far. The title asks "Which company would you like to invest in?" and each option shows the "Gains/Losses in %".

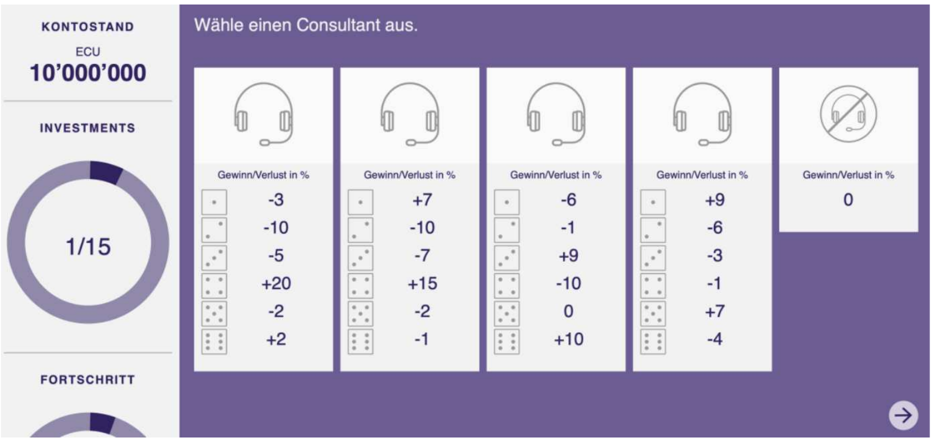

Figure C.3.2: Screenshot of the decisions screen in the *consultant condition* where participants saw an overview of their investment options. The sidebar shows the current balance ("Kontostand") as well as the investments and progress in the study so far. The title reads "Please choose a consultant" and each option shows the "Gains/Losses in %".

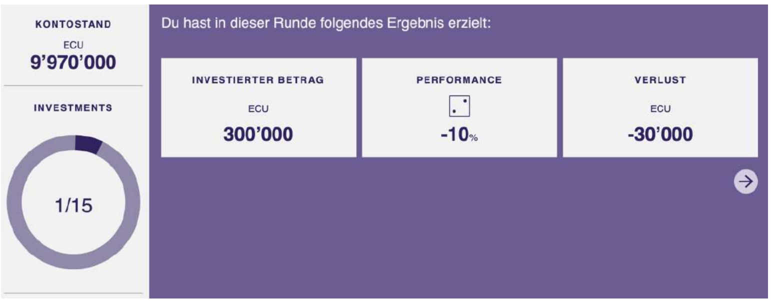

Figure C.3.3: Screenshot of the outcome screen after a round. The title reads "In this round you have achieved the following result:" and the three tiles show "Amount invested", "Performance" and "Loss" (/"Gain") respectively.

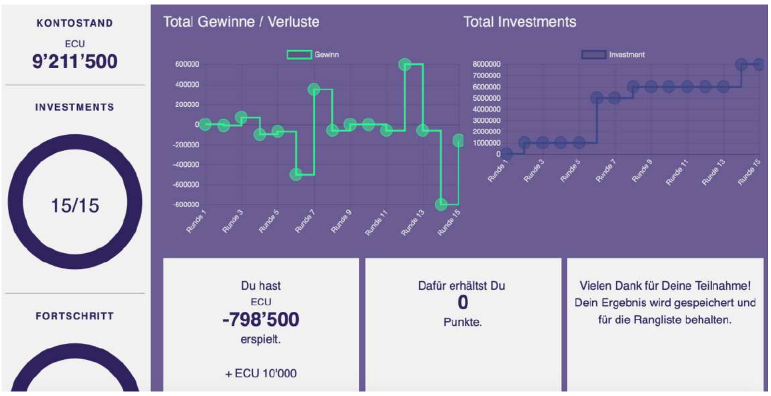

Figure C.3.4: Screenshot of the overview Screen after all 15 rounds had been completed. The left graph shows the "Total gains/losses" over all 15 rounds. The right graph shows the "Total investments" over all 15 rounds. The three tiles below read "You have made -798’500 EUC.", "For this you receive 0 points." and "Thank you for participating! Your result will be saved and kept for the ranking list."

Curriculum Vitae

Kevin Trutmann

Education

2018 – 2023	**PhD in Economic Psychology**, University of Basel
2016 – 2018	**MSc in Psychology**, University of Basel *Thesis: Using Cognitive Models of Social Decision Making in Search for Neural Signals of Altruism* *Supervisor: Prof. Sebastian Gluth*
2013 – 2016	**BSc in Psychology**, University of Basel *Thesis: Mechanisms of Caffeine and its Effects on Cognitive Functions* *Supervisors: Dr. Klara Spalek & Prof. Dominic de Quervain*

Publications & Preprints

Trutmann, K., Heinke, S., & Rudin, C. (in press). Degree of Personal Responsibility in Decisions and the Likelihood to Abandon an Investment among Professionals: Evidence from a Lab-in-the-Field Experiment. *Journal of Behavioral Finance.* https://doi.org/10.1080/15427560.2023.2228549

Trutmann, K., Heinke, S., & Rieskamp, J. (2022). *Belief Updating and Investment Decisions: The Impact of Good or Bad News Varies With Prior Returns.* SSRN Electronic Journal. https://dx.doi.org/10.2139/ssrn.3935798

Trutmann, K., Heinke, S., & Rieskamp, J. (2023). Take your time: How delayed information and restricted decision opportunities improve belief formation in investment decisions. *Finance Research Letters*, *51*, 103442. https://doi.org/10.1016/j.frl.2022.103442

Ongoing Projects

"The Foundations of Successful Financial Decision Making", SNF Sinergia Grant No. 177277, in collaboration with Steve Heinke, Silvia Maier, Alexander Ziegler, Alexandra Bagaïni, Isabella Kooij, Stephan Nebe, Nick Sidorenko, Thorsten Hens, Rui Mata, Philippe N. Tobler and Jörg Rieskamp

Conference Contributions

22–24. August 2021	SPUDM (Subjective Probability, Utility, and Decision Making) Conference, Poster Presentation "Previous Gains and Losses Influence Belief Formation in Investment Decisions"
26–29. July 2021	CogSci (Cognitive Science Society) Conference, Poster Presentation "The Effect of Investment Position on Belief Formation and Trading Behavior"
16–18. June 2021	Experimental Finance (EF) Conference, Talk "The Effect of Investment Position on Belief Formation and Trading Behavior"

28. May 2021	Bernoulli Workshop, Basel, "Lightning" Talk, "Learning Short- and Long-term Trends"
7–9. October 2020	Society for Neuroeconomics (SNE) Conference, Poster Presentation "The Effect of Investment Position on Belief Formation and Trading Behavior"
10–12. September 2020	Economic Science Association (ESA) Conference, Talk "The Effect of Investment Position on Belief Formation and Trading Behavior"

Research Internships

June – September 2015	Internship, Department of Biopsychology, University of Zürich
June – July 2017	Internship, Department of Methodology and Mathematical Psychology, University of Tübingen, Germany

Teaching & Supervision

2019 – 2022	Supervision of 10 BSc Theses, University of Basel
2014 – 2016	R-statistics Tutor, University of Basel

Extracurricular Activities

June – July 2019	Sloan-NOMIS Summer School on the Cognitive Foundations of Economic Behavior, Vitznau, Switzerland
June 2019	Experimental Finance Summer School, Copenhagen, Denmark