DISCUSSION PAPER SERIES

DP14647
(v. 2)

Why so Negative? Belief Formation and Risk Taking in Boom and Bust Markets

Martin Weber, Pascal Kieren and Jan Mueller-Dethard

FINANCIAL ECONOMICS

ISSN 0265-8003

Why so Negative? Belief Formation and Risk Taking in Boom and Bust Markets

Martin Weber, Pascal Kieren and Jan Mueller-Dethard

Discussion Paper DP14647
First Published 22 April 2020
This Revision 06 May 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Why so Negative? Belief Formation and Risk Taking in Boom and Bust Markets

Abstract

An increasing number of studies depart from the rational expectations assumption to reconcile survey expectations with asset prices. While surveys are helpful to establish a link between subjective beliefs and investment decisions, they do not allow inference about how investors depart from rational expectations. In this paper, we provide direct experimental evidence of how systematic distortions in investors' expectations affect their risk-taking across market cycles. As mechanism, we identify an asymmetry in how individuals update their expectations across boom and bust markets. The documented mechanism is consistent with survey data and provides important implications for recently proposed asset pricing models.

JEL Classification: D83, D84, E32, E44, G01, G11, G41

Keywords: risk-taking, Belief formation, Market Cycles, Return Expectations

Martin Weber - weber@bank.bwl.uni-mannheim.de
University of Mannheim and CEPR

Pascal Kieren - pascal.kieren@gess.uni-mannheim.de
University of Mannheim

Jan Mueller-Dethard - mueller-dethard@bank.bwl.uni-mannheim.de
University of Mannheim

Why so Negative?
Belief Formation and Risk-Taking in Boom and Bust Markets*

Pascal Kieren†, Jan Müller-Dethard‡, and Martin Weber§

May, 2021

Abstract

An increasing number of studies depart from the rational expectations assumption to reconcile survey expectations with asset prices. While surveys are helpful to establish a link between subjective beliefs and investment decisions, they do not allow inference about how investors depart from rational expectations. In this paper, we provide direct experimental evidence of how systematic distortions in investors' expectations affect their risk-taking across market cycles. As mechanism, we identify an asymmetry in how individuals update their expectations across boom and bust markets. The documented mechanism is consistent with survey data and provides important implications for recently proposed asset pricing models.

Keywords: *Risk-Taking, Belief Formation, Market Cycles, Return Expectations*

JEL Classification: *D83, D84, E32, E44, G01, G11, G41*

---

* For valuable comments, we thank Nicholas Barberis, Alex Imas, Camelia Kuhnen, Oliver Spalt, Konrad Stahl, Stefan Trautmann, Georg von Weizsäcker, as well as participants of the Maastricht Behavioral Experimental Economics Symposium 2019, the Experimental Finance 2019, the ESA 2019, the Annual Meeting of the SJDM 2019, the Annual Meeting of the AEA 2020, the Yale Graduate Student Conference 2020, the Annual Meeting of the EEA 2020, and seminar participants at the University of Mannheim. We gratefully acknowledge financial support by the German Research Foundation (DFG grant WE 993/15-1).

† Pascal Kieren, University of Mannheim, Department of Finance, E-mail: pascal.kieren@gess.uni-mannheim.de

‡ Jan Müller-Dethard, University of Mannheim, Department of Finance, E-mail: mueller-dethard@bank.bwl.uni-mannheim.de

§ Martin Weber, University of Mannheim, Department of Finance and the Centre for Economic Policy Research, E-mail: weber@bank.bwl.uni-mannheim.de

Introduction

How do individuals form expectations about future stock returns? The answer to this question is crucial to understand differences in risk-taking over time and in particular across market cycles. A key assumption in models that generate time-variation in risk-taking is that investors have rational expectations, which are immediately updated according to Bayes' rule when new information arrives (Barberis, Huang, & Santos, 2001; Campbell & Cochrane, 1999; Grossman & Shiller, 1981). This assumption implies that agents know the objective probability distribution in equilibrium and are as such fully aware of the countercyclical nature of the equity risk premium (Nagel & Xu, 2019). Yet, a number of recent surveys of investors' expectations show that this is not the case, and that investors – if anything – have rather pro-cyclical expectations: they are more optimistic in boom markets and less optimistic in recessions (Amromin & Sharpe, 2013; Giglio et al., 2020; Greenwood & Shleifer 2014).

In the light of this inconsistency, it is imperative to obtain a deeper understanding of *how* investors form expectations across boom and bust markets, and whether this could ultimately explain observed differences in risk-taking. To reconcile survey expectations with asset prices, a growing body of studies depart from the rational expectation's assumption. Examples include extrapolative expectations (Barberis et al., 2015; Bordalo et al., 2019), learning about underlying trends in price growth (Adam et al., 2016, 2017), learning with fading memory (Nagel and Xu, 2019), or learning from life-time experience (Collin-Dufresne et al., 2016; Malmendier and Nagel, 2011, 2015). However, while recent survey data on expectations are helpful to establish a link between subjective beliefs and investment decisions, they do not allow inference about how investors depart from rational expectations without imposing strong assumptions.

In this paper, we provide direct experimental evidence for the role of expectations for investors' risk-taking behavior across macroeconomic cycles. In an experiment, we can establish a setting in which we have direct control over objective (rational) expectations and can compare them to participants' subjective beliefs. This allows us to document systematic errors in the belief formation process, which we can then relate to the subjects' investment choice. In a series of experimental studies, we show that investors employ different learning rules when forming expectations in bust markets relative to boom markets. Conditional on observing the same information, individuals who learn in bust market environments form significantly more pessimistic beliefs than those who learn in boom market environments. The difference in beliefs subsequently translates to a lower willingness to take risk. We then show that the asymmetry in learning is driven by investors in bust markets putting more weight on unfavorable outcomes as well as taking significantly more time to recover pessimistic expectations from temporary shocks.

In our experiments, we combine a sequential belief updating task (similar to Grether, 1980) with an unrelated incentive-compatible investment task. In the learning task, subjects have to incorporate a sequence of information signals into their beliefs to estimate the likelihood that an asset pays dividends drawn from one of two distributions. In a between-subject design, we vary the frame of the learning environment to be representative of a boom market (boom treatment) or a bust market (bust treatment). The underlying probability distribution, however, from which the information is drawn, is completely identical in both learning environments. In other words, a Bayesian agent should make identical forecasts, irrespective of whether the information is observed in a boom or a bust market. To not rely on a single characterization for boom and bust markets, we examine two different market settings. In the first experiment, individuals exclusively learn either in the positive (boom) or in the negative (bust) domain. In the second experiment, we apply a more realistic market setting in which subjects learn from mixed-outcome distributions with either positive expected value (boom) or negative expected value (bust). After subjects completed the belief updating task, they make an unrelated investment decision in either a risky or an ambiguous lottery, which serves as a between-subject measure of belief- and preference-based risk-taking. In the ambiguous lottery, we purposefully give participants room to form subjective beliefs about the underlying true probability distribution. In the risky lottery, we have perfect control over subjects' return and risk expectations since both probabilities and outcomes are known. As such, investments in the ambiguous lottery are affected by both subjects' risk preferences and their beliefs about the underlying probability distribution, while investments in the risky lottery serve as a measurement tool for risk aversion. The between-subject comparison allows us to cleanly identify the influence of beliefs on individuals' risk-taking behavior. Overall, our design allows us to test 1) whether boom or bust learning environments affect individuals' risk-taking, 2) whether differences in risk-taking are driven by systematic differences in subjective return expectations, and 3) to explore potential mechanisms why individuals' beliefs deviate from Bayesian benchmarks.

Our findings can be summarized as follows. First, we find that subjects who learn to form expectations in bust market environments take significantly less risk in an unrelated ambiguous investment task than those subjects who learn to form beliefs in boom market environments. Once there is room to form subjective beliefs, subjects in the bust treatment invest on average 20% less in the ambiguous lottery compared to subjects in the boom treatment. To validate that the difference in individuals' risk-taking is driven by expectations, we show that subjects in the bust treatment are also substantially more pessimistic about the success probability of the ambiguous lottery (by about 19 percentage points). In the risky lottery, when probabilities are perfectly known, we can directly test whether bust learning environments also affect the subjects' risk aversion. However, we do not find any significant difference between treatments on subjects' investment in an unrelated risky investment option. This indicates that subjects' risk preferences remained stable and were unaffected by the initial boom or bust learning environment.

Next, we provide evidence for a specific mechanism that determines why risk-taking in the ambiguity task is a function of the prior learning environment by investigating whether learning rules are systematically different across treatments. We find that – for every Bayesian posterior in our sample – individuals' subjective beliefs in the bust treatment are on average significantly more pessimistic than the beliefs of individuals in the boom treatment. To understand why this is the case, we examine how individuals' probability estimates evolve such that they end up overly pessimistic in the bust treatment. The data suggests that participants in the bust treatment put significantly more weight on low outcomes when updating their expectations compared to those in the boom treatment. Additionally, we show that the wedge in updating across treatments is strongest when: 1) individuals with optimistic priors observe unfavorable information, which causes individuals in the bust treatment to overweight the signal; or when 2) individuals with pessimistic priors observe favorable information, which causes individuals in the bust treatment to underweight the signal. This asymmetry in learning across treatments indicates that people in bust market phases are not only more responsive to unfavorable outcomes, but also that they take significantly more time to recover pessimistic expectations from temporary shocks. The mechanism we document is consistent with recent survey evidence presented by Giglio et al. (2021), who find that investors' short-run expectations are very sensitive to market crashes. Most importantly, however, the authors show that those investors who were most optimistic before the crash experienced the steepest decline in expectations.

Finally, we provide evidence that the pessimism induced by bust learning environments within our experimental setup even affects subjects' return expectations in the real economy.

When asked to provide a return forecast of the Dow Jones Industrial Average, subjects in the bust treatment are significantly more pessimistic about the future performance of the index than their peers in the boom treatment. In addition to the more pessimistic expectations, we find that subjects who learn in bust markets provide negative return estimates, while those learning in boom markets provide positive return estimates. Given that we are able to systematically manipulate return expectations for real-world market indices even in a short-living learning environment as in our experiment, we believe that the effect reported here is even more generalizable in the real economy, where stakes and involvement are presumably higher.

Our findings contribute to several strands of literature. First, we contribute to an ongoing research effort to understand why risk premiums of many asset classes vary strongly and systematically over time (Campbell and Shiller, 1988; Cochrane, 2011; Shiller, 1981). To rationalize the high volatility of asset prices and the countercyclical equity risk premium, rational expectations models have evolved that introduce modifications into the representative agent's utility function, which effectively generates countercyclical risk aversion (Campbell and Cochrane, 1999; Barberis et al., 2001). Cohn et al. (2015) present experimental evidence supporting this notion, while Guiso et al. (2018) present survey evidence in line with this argument.[1] Recently, a number of studies depart from the rational expectations assumption to reconcile survey expectations with asset prices (e.g. Adam et al., 2017; Barberis et al., 2015; Bordalo et al., 2019; Nagel and Xu, 2019). In these models, asset prices are volatile because subjective expectations (instead of risk-aversion) are time-varying. Our results provide a direct and causal link of how systematic distortions in investors' expectations can affect their willingness to take financial risks. In an experiment, we can circumvent measurement and identification problems inherent in field data by establishing a setting in which we have direct control over objective (rational) expectations and can compare them to participants' subjective beliefs. In line with our results, Amromin and Sharpe (2014) find that households' lower willingness to take risks during recessions is driven by their more pessimistic subjective expectations rather than by countercyclical risk aversion. Similarly, Weber et al. (2013) show that changes in risk-taking of UK online-broker customers over the financial crisis of 2008 were mainly explained by changes in return expectations and to a lesser degree by changes in risk attitudes. It remains to stress, however, that we do not claim that changes in risk-taking are entirely driven by changes in beliefs. While there is a host of factors that differ across market

[1] There are also recent papers who challenge the notion of countercyclical risk aversion as tested in Cohn et al. (2015), such as Alempaki et al. (2019) and König-Kersting & Trautman (2018).

phases (i.e. return volatility, changes in wealth, employment status, etc.), we explicitly examine and isolate the effect of one particular aspect, which is investors' belief formation.[2]

Our study also relates to the findings reported in recent surveys of investor return expectations (Vissing-Jorgensen, 2003; Amromin & Sharpe, 2014; Greenwood & Shleifer, 2014; Giglio et al., 2020, 2021). A common finding is that survey expectations of stock returns are pro-cyclical (i.e. investors are more optimistic during boom markets and more pessimistic during recessions), and as such inconsistent with rational expectation models. A first attempt to reconcile this puzzling finding was made by Adam et al. (2020), who test whether alternative expectation hypotheses proposed in the asset pricing literature are in line with the survey evidence. However, they reject all of them. The findings of our study are consistent with investors' having pro-cyclical return expectations, as participants in our sample are more optimistic when learning in boom markets than when learning in bust markets. Additionally, the mechanism that drives our results is consistent with the observed behavior in Giglio et al. (2021), which suggests that it may provide an interesting starting point for future theories of belief updating featuring pro-cyclical expectations.

Finally, our findings also contribute to the broad literature on behavioral biases in belief formation. Prior studies have shown that people neglect base-rates (Kahneman and Tversky, 1973), display overconfidence (Moore and Healy, 2008), over-extrapolate from recent signals (Bordalo et al., 2018, 2019), or interpret evidence in a manner biased towards current beliefs (Charness and Dave, 2017, Benoît and Dubra, 2018) when forming their expectations. Recent research also studies the role of the outcome domain in belief formation. For example, Kuhnen (2015) shows that agents learn differently from outcomes in the negative domain than from the same outcome history in the positive domain. Our results suggest that the domain of the outcome distribution is not a necessary condition for asymmetric updating. Instead, we show that individuals already asymmetrically update their beliefs when they have different expectations about the outcome distribution, despite receiving the same information.

The remainder of the paper is organized as follows. Section 1 outlines the experimental design, and briefly discusses the most important design aspects. Section 2 describes summary statistics of our sample and randomization checks. Section 3 presents our results, including evidence on the underlying mechanism and external validity. Section 4 concludes.

---

[2] For example, Cohn et al. (2015) use a broader characterization of stock market booms and busts to prime their participants, which are financial professionals. While their approach allows to take more factors into account, it also makes the identification of underlying drivers more challenging.

1. Experimental Design

Seven-hundred fifty-four individuals (458 males, 296 females, mean age 34 years, 10.3 years standard deviation) were recruited from Amazon Mechanical Turk (MTurk) to participate in two online experiments.[3] MTurk advanced to a widely used and accepted recruiting platform for economic experiments. Not only does it offer a larger and more diverse subject pool as compared to lab studies (which frequently rely on students), but it also provides a response quality similar to that of other subject pools (Buhrmester et al., 2011; Goodman et al., 2013).

A. Description of the Experiment

Both experiments consist of two independent parts, a sequential belief updating task and an investment task. The experiments differ with respect to the updating task, but are identical with respect to the investment task. In the sequential belief updating task, we create learning environments which resemble key characteristics of boom and bust markets (see Figure 1).

Figure 1: Learning Environments and Treatments

This figure documents the learnings environments of the first part, the sequential belief updating task, of our two experiments. In both experiments, subjects are randomly assigned to either a Boom or a Bust Treatment. In Experiment 1, subjects learn from sequentially drawn positive (Boom) or negative (Bust) returns about the underlying state of a lottery (good or bad state). In Experiment 2, subjects learn from sequentially drawn positive and negative returns, but either from a lottery with positive (Boom) or negative (Bust) expected value about the underlying state of the lottery (good or bad state).

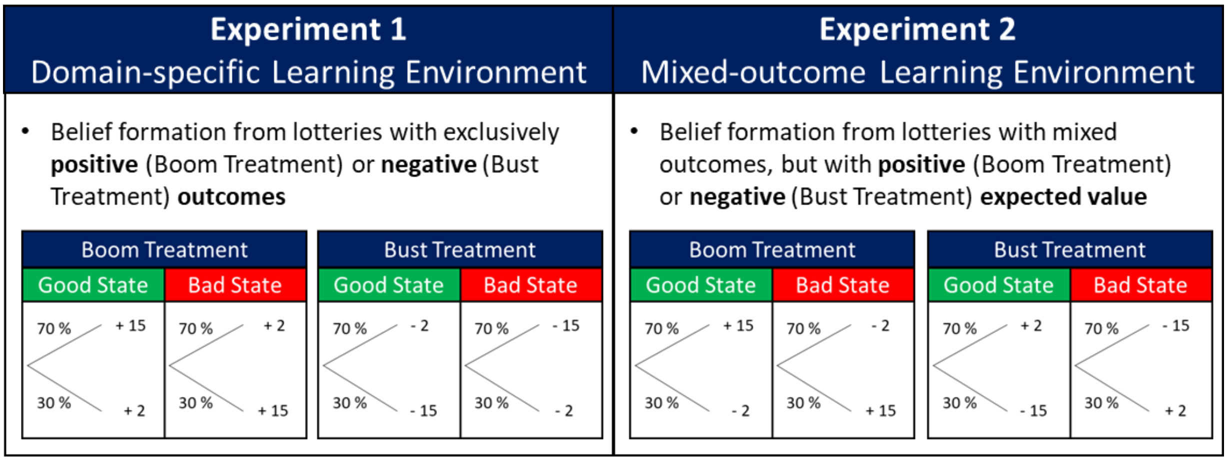

In Experiment 1, we focus on the domain (positive vs. negative returns) in which subjects primarily learn across different market phases. As such, we let subjects learn from

[3] Both experiments were preregistered at www.aspredicted.org. The preregistration of the first experiment can be found under http://aspredicted.org/blind.php?x=2rm3jw, whereas the preregistration of the second experiment can be found under http://aspredicted.org/blind.php?x=v7qa7q.

either exclusively positive outcome lotteries (boom-treatment) or negative outcome lotteries (bust-treatment). However, even in recessions agents occasionally observe positive returns, but the magnitude is on average smaller than the magnitude of observed negative returns. For example, during the last two financial crises, the frequency of observing a negative monthly return of the MSCI AC World index was 66.67 % for the DotCom Crisis and 68.42 % for the 2008 Financial Crisis, while the average realized monthly return was -1.17 % and -2.11 %, respectively, as displayed in Figure 2.[4]

Figure 2: Characteristics of Boom and Bust Market Phases

The figure documents both the relative frequency of observing a negative monthly return of the MSCI All Country World Index as well as the average monthly return for the last two financial recessions. Recessions are defined according to the NBER US Business Cycle Contraction classification. The left y-axis refers to the relative frequency of negative returns. The right y-axis (reversed scale) refers to the average monthly realized returns.

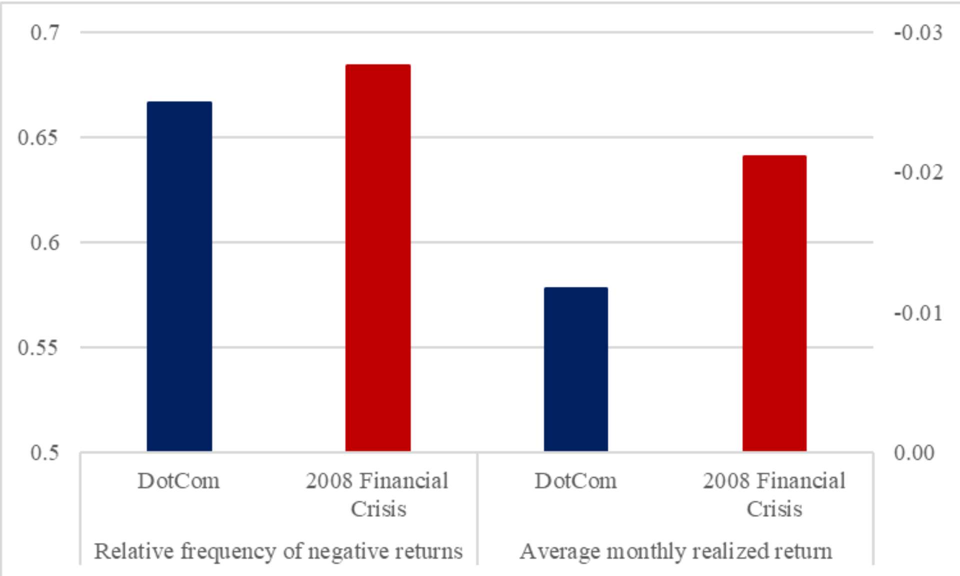

To account for this fact, we conduct another experiment with a presumably more realistic learning environment. In Experiment 2 subjects learn from mixed-outcome lotteries, which either have a positive expected value (boom-treatment) or a negative expected value (bust-treatment).

In the sequential belief updating task of both experiments, subjects receive information about a risky asset, whose dividends are either drawn from a "good distribution" or from a "bad distribution". As depicted in Figure 1, both distributions are binary with identical high and low outcomes. In the good distribution, the higher dividend occurs with a 70 % probability while

[4] Business cycles are defined using the NBER Business Cycle Expansion and Contractions Classification.

the lower dividend occurs with a 30 % probability. In the bad distribution, the probabilities are reversed, i.e. the lower dividend occurs with a 70 % probability while the higher dividend occurs with a 30 % probability. The actual dividends depend on both the experiment and the treatment to which subjects are assigned. In both experiments, subjects are randomly assigned to either a "*boom*" treatment or a "*bust*" treatment. In the first experiment, the dividends of the risky asset are either exclusively positive or negative. The dividends in the boom treatment are either +15, or +2, whereas they are -2, or -15 in the bust treatment. In the second experiment, the dividends of the risky asset are drawn from mixed-outcome lotteries, with either a positive or a negative expected value. The dividends in the boom treatment are either +15, or -2, whereas they are +2, or -15 in the bust treatment. While the dividends across treatments are mirrored, the underlying two distributions of the risky asset from which outcomes are drawn are identical.

In both experiments, subjects make forecasting decisions in two consecutive blocks each consisting of eight rounds. At the beginning of each block, the computer randomly determines the distribution of the risky asset (which can be good or bad). In each of the eight rounds, subjects observe a dividend of the risky asset. Afterwards, we ask them to provide a probability estimate that the risky asset draws from the good distribution and how confident they are about their estimate. As such, subjects will make a total of 16 probability estimates (8 estimates per block). To keep the focus on the updating task and to not test the performance of their memory, we display the prior outcomes in a price-line-chart next to the questions. To ensure that subjects have a sufficient understanding of the sequential belief updating task, they had to correctly answer three comprehension questions before they could continue (see Appendix B and C).

In the second part of each experiment, the investment task, we introduce a between-subject measure of belief- and preference-based risk-taking, presented in Figure 3. Subjects were randomly assigned to invest in either an *ambiguous* or a *risky* lottery with an endowment of 100 Cents (Gneezy and Potters, 1997). In both lotteries, the underlying success probability is 50 %. However, to introduce uncertainty and to provide subjects the freedom to form beliefs, the success probability remains unknown to them in the ambiguous lottery. In both lotteries, subjects can earn 2.5 times the invested amount if the lottery succeeds, whereas they lose the invested amount if the lottery fails. Subjects can keep the amount not invested in the lottery without earning any interest. In addition to the lottery investment, subjects in the ambiguous treatment are asked to provide an estimate of the success probability of the ambiguous lottery. Subjects in the risky treatment are not asked about a probability estimate as the objective success probability is known and clearly communicated.

Figure 3: Between-subject Measure of Belief- and Preference-based Risk-Taking

This figure presents the between-subject measure of belief- and preference-based risk-taking used in the second part, the investment task, of our experiments. The ambiguous lottery is characterized by unknown probabilities, such that subjects have the room to form beliefs about the underlying probability distribution of the ambiguous lottery. The risky lottery is characterized by known probabilities and is used as a measurement tool for between-treatment differences in risk aversion.

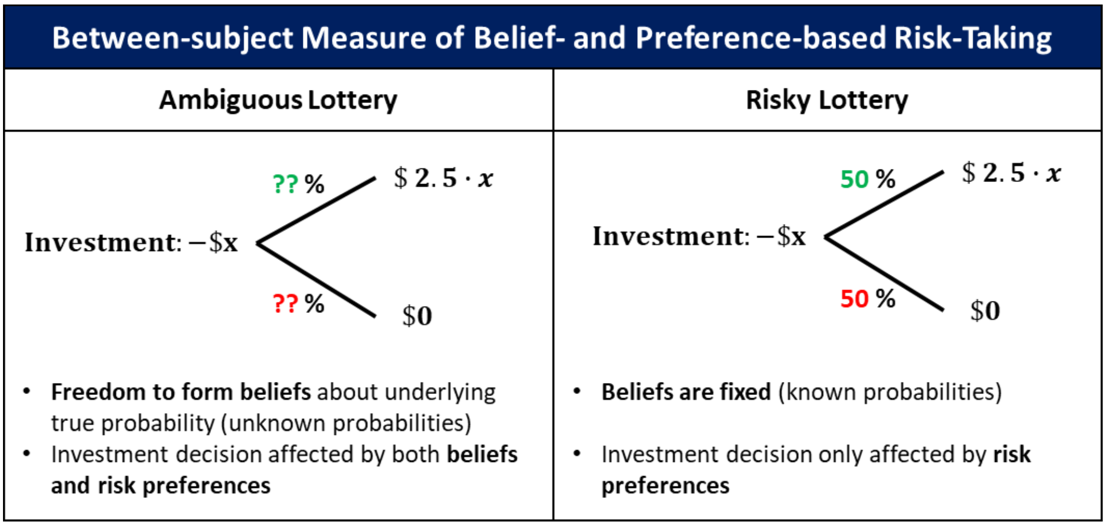

The experiments concluded with a brief survey about subjects' socio-economic background, a 10-item inventory of the standard Life Orientation Test (Scheier, Carver, & Bridges, 1994), self-assessed statistical skills, stock trading experience and whether a participant was invested during the last financial crisis. In addition, subjects were asked to provide a 6-month return forecast of the Dow Jones Industrial Average index on a twelve-point balanced Likert scale. In summary, Figure 4 provides a time line of the experiments, including all described stages.

Both parts of the experiment were incentivized. In the first part, participants were paid based on the accuracy of the probability estimate provided. Specifically, they received 10 cents for each probability estimate within 10 % (+/- 5%) of the objective Bayesian value. In the second part of the experiment, subjects received the amount not invested in the lottery plus the net earnings from their lottery investment. Both studies took approximately 9 minutes to complete and participants earned $1.93 on average.

Figure 4: Structure and Flow of the Experiments

This figure shows a time line of our experiments. Subjects do a sequential belief updating task followed by an independent investment task. The sequential belief updating task consists of two blocks. In each block, subjects have to give eight probability estimates and eight estimates about how confident they are about their forecasts. Both blocks of forecasting are either in a boom market or in a bust market environment. The random assignment of the boom or bust market environment is done at the beginning of the experiment. After the sequential belief updating task, subjects invest either in an ambiguous lottery or in a risky lottery. For the ambiguous lottery, they are in addition asked about an estimate of the underlying success probability. The experiments end with a short survey which consists of a six-month forecast of the Dow Jones Industrial Average, a 10-item Life Orientation Test, and socio-demographic questions.

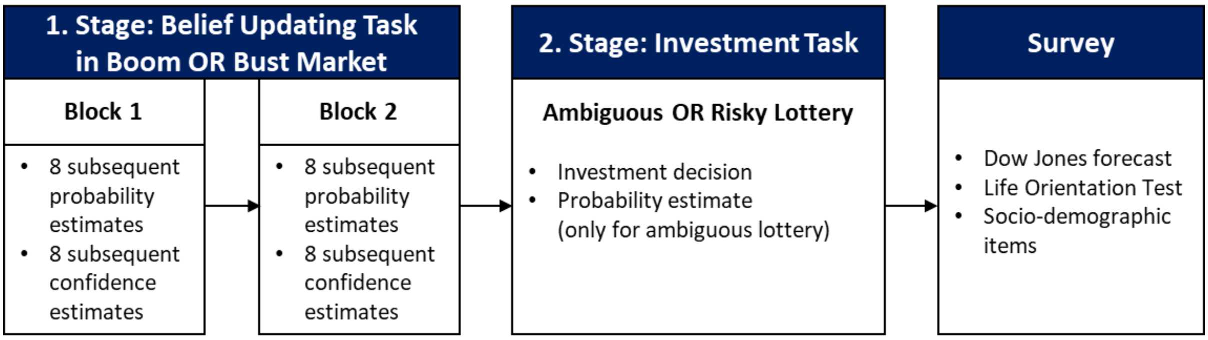

B. Discussion of Important Aspects

Overall, our design allows us to test whether different learning environments across boom and bust markets can account for time variation in risk-taking. In particular, we aim to examine whether differences in the way beliefs are formed across boom and bust markets induce optimism or pessimism, respectively, which ultimately transmissions to differences in risk-taking. As it is imperative for our design to ensure that treatment effects can be pinpointed to changes in risk preferences or changes in expectations, a few aspects warrant a brief discussion. First, feedback regarding the accuracy of subjects' probability estimates was only provided at the very end of the experiment. This was done to not only avoid wealth effects, but also to ensure that subjects do not hedge the lottery investment against their earnings from the updating task, which would inevitably affect their risk-taking. Second, we abstract from using predisposed words like "boom", "bust", or similar financial jargon. This circumvents evoking negative or positive emotions (such as fear), experience effects, and other confounding factors, which would distort a clear identification of belief-induced risk-taking. Third, by introducing a between-subject measure of belief- and preference-based risk-taking, we can directly test whether differences in risk-taking are attributable to changes in risk preferences, beliefs, or both. In the risky treatment, we have perfect control over subjects' return and risk expectations since both probabilities and outcomes are known and clearly communicated. In the ambiguous treatment, we intentionally give participants room to form subjective beliefs as there is uncertainty about the true underlying probability (Klibanoff, Marinacci, & Mukerji, 2005). As such, investments in the ambiguous lottery are affected by both subjects' risk preferences and their beliefs about the probability distribution, while investments in the risky lottery serve as a measurement tool for risk aversion. The between-subject comparison allows us to cleanly identify the influence of our treatments in the learning task on individuals' risk-taking behavior.

2. Summary Statistics and Randomization Checks

Table 1 presents summary statistics, Panel A for Experiment 1 and Panel B for Experiment 2. Overall 754 subjects participated in our studies, with an average age of 35.15 years in Experiment 1 (33.53 years in Experiment 2). Forty-five percent (thirty-four percent) were female. Subjects reported average statistical skills of 4.19 out of 7 (4.47) and are medium experienced in stock trading, with a self-reported average score of 3.64 out of 7 (3.94). Roughly thirty-nine percent (forty-four) were invested during the 2008 Financial Crisis.

Table 1: Summary Statistics on Subjects

This table shows summary statistics for our experimental data. Reported are the mean and the standard deviation (in parentheses) for the whole sample (Column 1) and split across treatments (Column 2 and 3). Column 4 presents randomization checks. Differences in mean were tested using rank-sum tests, or χ^2-tests for binary variables. The p-value is reported in Column 5. *Female* is an indicator variable that equals 1 if a participant is female. *Statistical skills* denotes participants' self-assed statistical skills on a 7-point Likert scale. *Experience in stock trading* is the self-reported experience participants have in stock trading, assessed by a 7-point Likert scale. *Invested financial crisis* is an indicator that equals 1 if participants were invested in the stock market during the last financial crisis.

Panel A **Variable**	**Full sample (N = 350)**	**Boom (N = 174)**	**Bust (N = 176)**	**Difference**	**p-value**
Age	35.15	34.76	35.54	0.78	0.76
	(11.52)	(11.18)	(11.86)		
Female	0.45	0.47	0.43	0.04	0.44
	(0.50)	(0.50)	(0.50)		
Statistical Skills (1-7)	4.19	4.22	4.16	0.06	0.91
	(1.62)	(1.51)	(1.72)		
Experience in Stock Trading (1 - 7)	3.64	3.73	3.56	0.17	0.42
	(1.88)	(1.84)	(1.92)		
Invested Financial Crisis (1 = Yes)	0.39	0.39	0.39	0	1
	(0.49)	(0.49)	(0.49)		

Panel B **Variable**	**Full sample** **(N = 403)**	**Boom** **(N = 207)**	**Bust** **(N = 196)**	**Difference**	**p-value**
Age	33.53	32.73	34.37	1.63	0.07
	(9.03)	(8.46)	(9.55)		
Female	0.34	0.33	0.35	0.02	0.69
	(0.48)	(0.47)	(0.48)		
Statistical Skills (1-7)	4.47	4.40	4.55	0.15	0.42
	(1.67)	(1.69)	(1.65)		
Experience in Stock Trading (1 - 7)	3.94	3.89	3.98	0.09	0.52
	(1.99)	(1.95)	(2.03)		
Invested Financial Crisis (1 = Yes)	0.44	0.41	0.47	0.06	0.24
	(0.50)	(0.49)	(0.50)		

Additionally, we tested whether our randomization successfully resulted in a balanced sample. Table 1 also reports the mean and standard deviation of each variable split by treatment. Differences were tested using rank-sum tests, or χ^2-tests for binary variables. As we find no significant difference between our treatments for any variable, our randomization was successful. As such, we cannot reject the null hypothesis that the socio-economic background of the subjects is balanced between our boom and bust treatment.

3. Results

3.1 Main Result

A. Risk-Taking Behavior

We begin by comparing the average amount invested in the risky and ambiguous lottery conditional on whether subjects initially learned in a boom or bust market environment. Figure 5 displays the average investment share in both lotteries by treatment.

The results reported in Figure 5 provide a simple first test for our research question. Whereas subjects in the bust treatment invest on average 36 out of 100 Cents into the ambiguous lottery, subjects in the boom treatment invest roughly 45 Cents into the ambiguous lottery, a relative difference of about 20 percent ($p < 0.01$, two-sided t-test). As such, we find a significant treatment effect of learning to form beliefs in boom and bust markets on subjects' willingness to take risks. However, we find no such effect for investments in the risky lottery. While subjects in the boom treatment invest on average 39 Cents in the risky lottery, subjects in the bust treatment invest roughly 43 Cents, with no significant difference between the two (p = 0.32, two-sided t-test). Effectively, the second set of findings indicates that our treatment differences are not driven by a change in risk preferences.

Figure 5: Risk-Taking in Boom and Bust Learning Environments I

This figure displays the average investments (0 – 100) of participants in the ambiguous lottery and the risky lottery split by treatment. Reported are 95% confidence intervals.

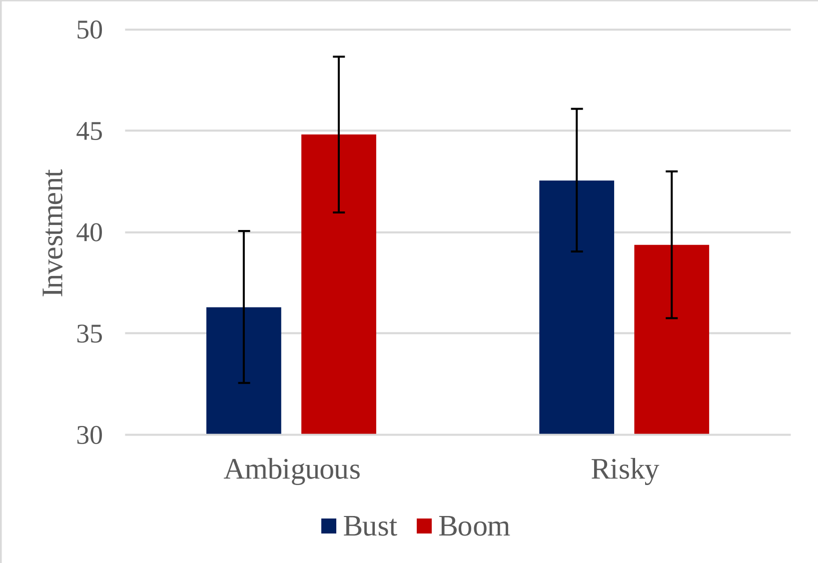

In order to underpin the treatment differences statistically and to control for individual differences in socio-economic and financial background, we specify the following regression model:

$$Investment_i = \beta_0 + \beta_1 Bust_i + \beta_2 Ambiguous_i + \beta_3 Bust_i \text{ x } Ambiguous_i + \sum_{j=1}^{n} \beta_j \boldsymbol{X_{ij}} + \epsilon_i,$$

where the dependent variable $Investment_i$ is the amount individual i invested in the risky/ambiguous asset. $Bust_i$ is a dummy for treatment Bust, while $Ambiguous_i$ is a dummy that denotes that the investment decision was made under ambiguity (i.e. unknown probabilities in the investment task). The interaction $Bust_i$ x $Ambiguous_i$ allows us to examine our main research question, i.e. that subjects who learned to form beliefs in bust environments invest significantly less in the ambiguous lottery where they have room to form subjective expectations (e.g. Klibanoff et al., 2005). Finally, $\boldsymbol{X_{ij}}$ is a set of control variables including gender, age, statistical skills, stock trading experience, a life orientation test, the order of good and bad distributions in the sequential belief updating task, and an indicator whether subjects were invested in the last financial crisis. We estimate our regression model using OLS with robust standard errors. However, results remain the same if we use a Tobit model instead.

Table 2: Risk-Taking in Boom and Bust Learning Environments II

This table examines subjects' risk-taking across treatments. We report the results of OLS regressions for the whole sample with and without controls. The dependent variable is *Investment*, which denotes participants' invested amount (0 – 100) in the lottery they were assigned to. *Bust* is an indicator variable that equals 1 if participants were in the bust treatment. *Ambiguous* is an indicator variable that equals 1 if participants were asked to invest in the ambiguous lottery, and 0 if they invested in the risky lottery. To test for differences across experiments and as such the different learning environments (Experiment 1: domain-specific; Experiment 2: mixed), we run the same regression including a *Mixed* indicator variable that equals 1 if participants learned in the mixed environment and an interaction term. Controls include age, gender, statistical skills, self-reported experience in stock trading, whether subjects were invested in the stock market during the last financial crisis, and the order of outcomes they experienced in the sequential belief updating task. Reported are coefficients and t-statistics (in parentheses) using robust standard errors. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	*Investment*		
	(1)	(2)	(3)
Bust	3.184	2.271	5.505
	(0.99)	(0.72)	(1.22)
Ambiguous	5.438*	5.149*	5.661
	(1.77)	(1.71)	(1.31)
Bust x Ambiguous	-11.69***	-11.23**	-14.18**
	(-2.60)	(-2.54)	(-2.32)
Mixed			1.408
			(0.32)
Bust x Ambiguous x Mixed			5.500
			(0.61)
Constant	39.38***	15.82*	14.99
	(17.83)	(1.70)	(1.51)
Observations	754	753	753
Controls		X	X
R^2	0.011	0.060	0.061

The estimation results reported in Table 2 support our main finding. The negative sign of the interaction term in Column 1 indicates that individuals in the bust treatment invest significantly less in the ambiguous lottery compared to those in the boom treatment ($p < 0.01$). To rule out that the reported differences in risk-taking behavior are driven by changes in risk aversion, we investigate subjects' investment in the risky lottery (baseline). We designed the risk task in such a way that expectations should not matter. Since subjects knew all the parameters of the risky asset, any kind of uncertainty was eliminated, which allows us to directly test the effect of our sequential belief updating task on subjects' risk aversion. However, we do not find any significant difference between treatments on subjects' investment in the risky lottery ($p = 0.323$). This means that we cannot reject the null hypothesis that risk aversion for subjects who formed beliefs in bust market environments is similar compared to subjects who formed beliefs in boom market environments. Both findings also hold if we control for socio-economic and financial background (Column 2). Finally, in Column 3 we investigate whether our reported treatment effect depends on differences in the underlying learning environment between Experiment 1 and 2.[5] To do so, we add an experiment dummy as well as an interaction term with our main variable of interest. The interaction term shows that subjects across both experiments responded similarly to our manipulation, which indicates that our main finding does not critically depend on individuals exclusively observing positive/negative outcomes.

B. Expectations

Next, we seek to validate that the reported differences in individuals' risk-taking behavior are indeed driven by expectations. We designed the ambiguous treatment in such a way that we can assess participants' subjective beliefs about the success probability of the lottery and directly relate them to their investment decision. If expectations are the main driver of differences in risk-taking, we should observe that subjects who learned to form beliefs in our bust treatment are more pessimistic about the success probability of the ambiguous lottery. In addition, we would expect a positive correlation between the subjective probability estimate of the success chance of the ambiguous lottery and the amount invested in the ambiguous lottery.

We first compare the average beliefs about the success probability of the ambiguous lottery between subjects who initially learned in the boom and subjects who initially learned in the bust market environment. Figure 6 displays the results. We find a strong and highly significant effect of our treatment indicator on the subjective success probability of the ambiguous lottery. In particular, those subjects who learned to form expectations in the bust

[5] This is done to ensure that the characterization of a boom or bust learning environment does not solely depend on individuals who exclusively observe positive or negative outcomes, respectively.

treatment are about 19 percentage points ($p < 0.001$) more pessimistic about the success probability than subjects who learned to form beliefs in the boom treatment (average success probability estimate for boom treatment: 68 %; for bust treatment: 49 %). Moreover, the induced pessimism from the sequential belief updating task that spills over to subjects' expectations about the success probability of the ambiguous lottery can be observed irrespective of whether subjects went through the domain-specific (Experiment 1) or the mixed (Experiment 2) learning environment.

Figure 6: Estimated Success Probability of the Ambiguous Lottery

This figure shows subjects' average beliefs about the success probability of the ambiguous lottery split by treatment for the whole sample and for each learning environment (Experiment 1: domain-specific; Experiment 2: mixed) individually. Displayed are 95% confidence intervals.

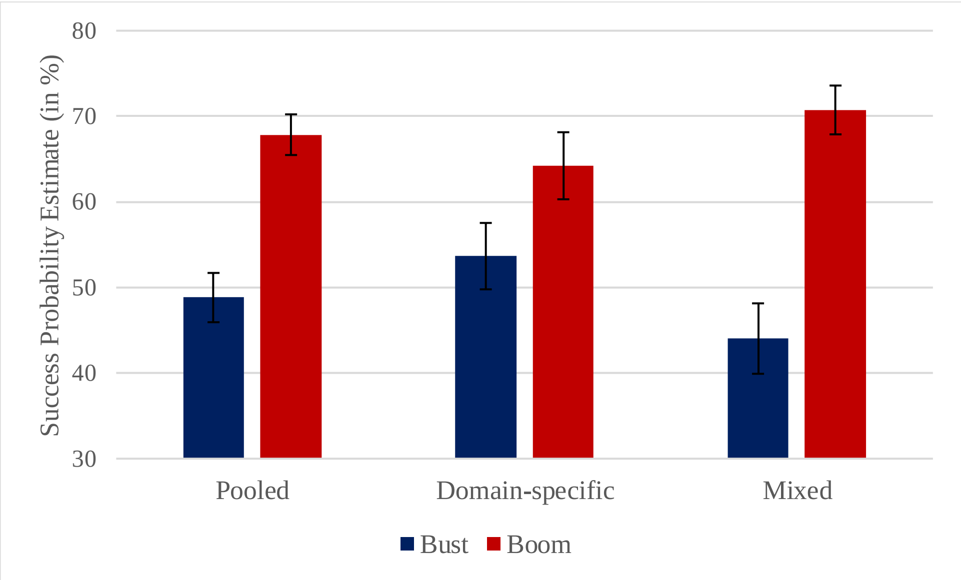

In order to directly test the treatment-belief as well as the belief-investment relation further, we estimate the following two OLS regression models:

$$(1)\ \ Probability_i = \beta_0 + \beta_1 Bust_i + \sum_{j=1}^{n} \beta_j \boldsymbol{X_{ij}} + \epsilon_i$$

$$(2)\ Investment\_Ambiguous_i = \beta_0 + \beta_1 Probability_i + \sum_{j=1}^{n} \beta_j \boldsymbol{X_{ij}} + \epsilon_i,$$

where $Probability_i$ is the subjective success probability of the ambiguous lottery of subject i, and $Investment\_Ambiguous_i$ is the investment of subject i in the ambiguous lottery. Findings for the first model are reported in Table 3 and for the second model in Table 4.

Table 3: Relation Between Treatment Variable and Probability Estimates

This table examines how our treatments affect subjects' beliefs about the success probability of the ambiguous lottery. We report the results of OLS regressions for the whole sample with and without controls. The dependent variable is *Success Probability*, which denotes participants' beliefs about the success probability of the ambiguous lottery. *Bust* is an indicator variable that equals 1 if participants were in the bust treatment. To test for differences across experiments and as such the different learning environments (Experiment 1: domain-specific; Experiment 2: mixed), we run the same regression including a *Mixed* indicator variable that equals 1 if participants learned in the mixed environment and an interaction term. Controls include age, gender, statistical skills, self-reported experience in stock trading and whether subjects were invested in the stock market during the last financial crisis. Reported are coefficients and t-statistics (in parentheses) using robust standard errors. *, **, *** indicates statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	*Success Probability Estimate of Ambiguous Lottery*		
	(1)	(2)	(3)
Bust	-19.03***	-18.86***	-11.71***
	(-8.36)	(-8.59)	(-3.70)
Mixed			4.576
			(1.64)
Bust x Mixed			-13.76***
			(-3.20)
Constant	67.85***	55.83***	54.48***
	(46.81)	(6.15)	(5.84)
Observations	377	377	377
Controls		X	X
R^2	0.158	0.241	0.263

We confirm that the strong and highly significant effect of our treatment indicator on the subjective success probability of the ambiguous lottery holds after controlling for socio-economic and financial background. The finding remains stable and highly significant across both experiments, albeit they are even more pronounced in the second experiment.

Table 4: Relation Between Beliefs About Success Probability and Investment

This table examines whether subjects in our experiment act upon their beliefs about the success probability of the ambiguous lottery. We report the results of OLS regressions for the whole sample with controls. The dependent variable is *Investment Ambiguous*, which captures subjects' invested amount in the ambiguous lottery. *Probability* denotes participants' beliefs about the success probability of the ambiguous lottery. *Bust* is an indicator variable that equals 1 if participants were in the bust treatment. To test for differences across experiments and as such the different learning environments (Experiment 1: domain-specific; Experiment 2: mixed), we run the same regression including a *Mixed* indicator variable that equals 1 if participants learned in the mixed environment and interaction terms. Controls include age, gender, statistical skills, self-reported experience in stock trading and whether subjects were invested in the stock market during the last financial crisis. Reported are coefficients and t-statistics (in parentheses) using robust standard errors. *, **, *** indicates statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	*Investment in Ambiguous Lottery*				
	(1)	(2)	(3)	(4)	(5)
Probability	0.412***	0.409***	0.352***	0.421***	0.331***
	(6.45)	(5.70)	(3.85)	(5.72)	(3.41)
Bust		-0.372		-2.968	-3.937
		(-0.11)		(-0.73)	(-0.95)
Mixed			-4.682	-1.205	-13.31
			(-0.66)	(-0.28)	(-1.30)
Probability x Mixed			0.105		0.179
			(0.88)		(1.29)
Bust x Mixed				5.507	8.864
				(0.90)	(1.28)
Constant	-3.304	-2.985	-0.0281	-3.616	3.638
	(-0.26)	(-0.23)	(-0.00)	(-0.28)	(0.28)
Observations	377	377	377	377	377
Controls	X	X	X	X	X
R^2	0.146	0.146	0.148	0.148	0.152

In Table 4, we test whether differences in subjective expectations regarding the success probability of the ambiguous lottery also translate to changes in risk-taking. In essence, we test whether subjects adhere to a basic economic principle: keeping everything else constant, do subjects increase their investment in an ambiguous lottery when their beliefs about the outcome distribution are more optimistic? Our results across all specifications confirm that subjects act upon their beliefs. In other words, the more optimistic they are about the success probability of the ambiguous lottery, the more they invest ($p < 0.01$). In addition, in Columns (2), (4), and (5), we include the Bust indicator as an additional control variable to exclude the possibility that our manipulation affects factors unrelated to expectations (also interacted with an experiment dummy in Columns (4) and (5)). Even after including the Bust indicator, the effect of subjective probability estimates on investments remains of similar magnitude and statistical significance. Moreover, we find no additional effect of our manipulation on the investment decision. Effectively, this means while our manipulation does induce pessimism, it does not affect factors unrelated to expectations.

3.2 Mechanism

The previous section demonstrated that differences in risk-taking in the ambiguous lottery are a function of whether the prior learning environment resembled characteristics of either a boom or a bust market. Further analyses have shown that the channel through which risk-taking in the ambiguous lottery is affected are beliefs. In particular, subjects who previously went through the bust learning environment state more pessimistic expectations about the success probability of the ambiguous lottery. A central question remains open so far: how does the bust learning environment induce pessimism? This is particularly interesting given that by design of the experiment, there should be no differences in overall optimism/pessimism across treatments since the underlying probability distributions from which subjects learn are identical. In this section, we aim to provide evidence for a specific mechanism behind the effect by investigating whether learning rules are systematically different across treatments.

First, we examine how beliefs in the updating task evolve relative to normative Bayesian benchmarks when learning about the quality of a risky asset. Figure 7 graphs participants' posterior beliefs regarding the likelihood that the risky asset pays dividends from the good distribution for each value of the objective Bayesian posterior belief, split by treatments.[6] If participants were perfect Bayesian learners, their subjective posteriors would line up perfectly with the objective Bayesian posteriors. However, Figure 7 indicates that this is not the case. The blue line, representing learning in bust market environments, is consistently below the red line, representing learning in boom market environments. In other words, for *every* Bayesian posterior in our sample, subjects in the bust treatment form on average more pessimistic

[6] Results do not depend on the experiment. The figures for each experiment separately are available in the Appendix (Figures A1 and A2).

expectations regarding the likelihood that the risky asset is paying dividends from the good distribution than subjects in the boom treatment.

Figure 7: Asymmetric Learning in Boom and Bust Market Environments

This figure displays average subjective probability estimates that the risky asset is paying dividends from the good distribution as a function of the objective Bayesian probability, split by treatment (Boom; Bust). Subjective probability estimates provided by participants for each objective Bayesian probability are shown as a blue line for the Bust treatment, and as a red line for the Boom treatment. If subjective probability estimates were Bayesian, they would equal the objective probabilities and thus would line up on a line starting in the origin with slope one. A table with all possible objective Bayesian probabilities resulting from the combinations of high and low outcomes observed is shown in the Appendix.

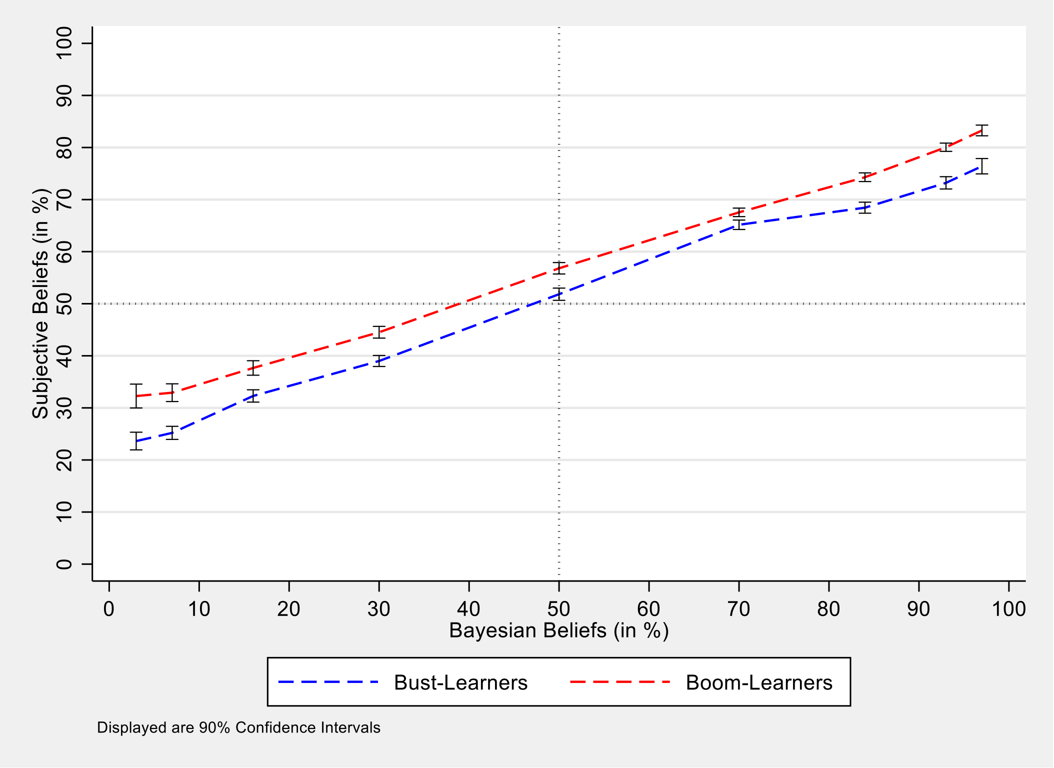

Table 5 examines this pattern in greater detail. We regress subjects' probability estimates on a bust-indicator and the objective Bayesian probability that the risky asset is in the good state. Across both experiments, we find that beliefs expressed by subjects in the bust treatment are on average 5.33% lower (i.e. more pessimistic) than in the boom treatment ($p <$ 0.001). This means that – holding the objective posterior constant – subjects update their priors differently when learning in bust market environments compared to boom market environments. The magnitude of this pattern does not significantly differ across experiments. In other words, the induced pessimism from the bust treatment does not critically depend on whether subjects observe exclusively negative outcomes or mixed outcomes drawn from a distribution with negative expected value. In essence, our results imply that learning differs across boom and bust markets.

Table 5: Asymmetric Learning in Boom and Bust Market Environments

This table reports the results of three OLS regressions on how subjective posterior beliefs about the distribution of the lottery depend on the treatment. We report the results of OLS regressions for the whole sample, and for each experiment individually (Experiment 1: domain-specific; Experiment 2: mixed). The dependent variable in the regression model, *Probability Estimate*, is the subjective posterior belief that the asset is paying from the good distribution. Independent variables include the *Bust* dummy, an indicator variable that equals 1 if participants were in the bust treatment and zero otherwise, as well as *Objective Posterior*, which is the correct Bayesian probability that the stock is good, given the information seen by the participant up to round *t* in the learning block. Controls include age, gender, statistical skills, self-reported experience in stock trading, whether subjects were invested in the stock market during the last financial crisis, and the order of outcomes they experienced in the sequential belief updating task. Reported are coefficients and t-statistics (in parentheses) using robust standard errors clustered at the individual level. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	*Probability Estimate (Subjective Posterior)*		
	Pooled Data	Domain-specific	Mixed
Bust	-5.331***	-5.228***	-6.077***
	(-5.68)	(-3.42)	(-5.34)
Objective Posterior	0.536***	0.488***	0.573***
	(31.15)	(19.88)	(24.42)
Constant	29.86***	26.53***	34.07***
	(7.76)	(4.69)	(6.62)
Observations	7584	3520	4064
R^2	0.520	0.493	0.551

The observed wedge between learning in boom and bust markets is not consistent with Bayesian learning, which predicts no differences conditional on receiving the same information about the underlying distribution. To understand why this is the case, it is important to examine how participants' probability estimates evolve such that they end up overly pessimistic in the bust treatment relative to the boom treatment. To shed light on this question, we calculate, for each participant *i* in round *t+1*, the difference between his subjective posterior and prior belief that the risky asset pays dividends from the good distribution. That is, we calculate the change from round *t* to round *t+1* in subjects' probability estimate. We then test whether, on average, probability updating differs across market phases. Results are presented in Table 6.

Table 6 reveals that participants put significantly more weight on low outcomes in the bust treatment relative to the boom treatment. In particular, the table shows that observing a low dividend reduces participants' probability estimate that the risky asset is the good one by 1.73% ($p < 0.01$) more in the bust treatment than in the boom treatment. Additionally, the table reveals a fundamental asymmetry in how updating depends on prior beliefs across market phases. Whenever prior expectations are inconsistent with the information of the most recent signal, probability updating in the bust treatment is significantly more pessimistic than in the boom treatment. In particular, after observing a low dividend payment while having optimistic prior beliefs (prior > 0.5), individuals in the bust treatment reduce their probability estimate by 5.86% ($p < 0.01$) more than individuals in the boom treatment. Similarly, after observing a high dividend payment conditional on having pessimistic prior beliefs (prior < 0.5), individuals in the bust treatment form beliefs that are on average 5.44% ($p < 0.01$) more pessimistic than individuals who learn in a boom market. This difference in updating from high and low dividend payments between boom and bust markets gets more pronounced the more extreme individuals' priors are. It increases to 6.48% for the case of a low dividend that is observed when having more optimistic prior beliefs (prior > 0.7) and to 7.33% for the case of a high dividend that is observed when having more pessimistic prior beliefs (prior < 0.3). This asymmetry in learning across market phases indicates that people in bust market phases are not only more responsive to unfavorable outcomes, but also that they take significantly more time to recover pessimistic expectations from temporary shocks. Additionally, the asymmetry in the learning rate across treatments manifests quickly, suggesting that investors' belief formation process is responsive to changes in the underlying market regime. Both patterns are consistent with the survey evidence presented by Giglio et al. (2021), who find that investors' short-run expectations are very sensitive to market crashes. Most importantly, however, the authors find that those investors who were most optimistic before the crash experienced the steepest decline in expectations, consistent with the mechanism we discuss. Overall, the beliefs we document are consistent with survey evidence on the pro-cyclicality of investor beliefs (Vissing-Jorgensen, 2003; Amromin and Sharpe, 2013; Greenwood and Shleifer, 2014; Giglio et al., 2020) and provide a micro foundation for why this is the case. The asymmetric learning, we observe, is not consistent with standard Bayesian learning models, which predict no difference in the learning rate. It is also not consistent with behavioral models of motivated beliefs (Kunda, 1990; Brunnermeier and Parker, 2005), which predict asymmetric updating but overall optimism. Moreover, models of associative memory (Enke et al. 2019) cannot rationalize our findings because prior outcomes were accessible and clearly displayed for both treatments.

Table 6: Differences in Belief Updating After High and Low Outcomes Conditional on Boom and Bust Learning Environments

This table shows the average change in subjective probability estimates that the risky asset is paying dividends from the good distribution from round to round for the Boom and Bust treatment split by high and low dividends and high and low subjective priors. Significant differences in belief updating between the Boom and Bust treatment at the 10%, 5%, and 1% level are indicated by *, **, and ***, respectively.

	Probability Estimate in t+1 - Probability Estimate in t							
	High Dividend in Round t+1	Low Dividend in Round t+1	High Dividend in Round t+1, Probability Estimate in t < 50%	High Dividend in Round t+1, Probability Estimate in t > 50%	Low Dividend in Round t+1, Probability Estimate in t < 50%	Low Dividend in Round t+1, Probability Estimate in t > 50%	High Dividend in Round t+1, Probability Estimate in t < 30%	Low Dividend in Round t+1, Probability Estimate in t > 70%
Bust	8.77%	-9.68%	18.14%	4.26%	-1.34%	-17.12%	21.04%	-24.67%
Boom	9.16%	-7.95%	23.58%	5.42%	-1.21%	-11.26%	28.37%	-18.19%
Bust - Boom	-0.39%	-1.73%***	-5.44%***	-1.16%**	-0.13%	-5.86%***	-7.33%***	-6.48%***

3.3 External Validity

In this section, we explore the extent to which the documented asymmetry in learning affects investors' expectations about real-world market indices. A common concern for experimental studies is that the size of parameters estimated cannot be extrapolated to field settings without caveats or relying on strong assumptions (e.g. Harrison, List, and Towe, 2007; or Charness, Gneezy, and Imas, 2013). However, because we are interested in the comparative static effects of boom markets versus bust markets rather than the absolute level of differential learning, we feel confident that our results are generalizable to field settings. To further reduce this concern, we analyze subjects' responses to two additional sets of questions, which deal with expectations outside the experimental setting. The first question tests to which extent the induced pessimism translates to expectations in the real economy. We gave subjects the at the time current level of the Dow Jones Industrial Average, and asked them to provide a 6-month return forecast on a balanced 12-point Likert scale (see Appendix C). Figure 8 graphs subjects' return forecasts split by treatment.

Figure 8: Dow Jones Estimates I

The figure displays subjects' self-reported return expectations of the Dow Jones Industrial Average. Dow Jones return expectations were assessed on a 12-point Likert scale. Results are displayed separately for subjects by treatment (boom/bust) and by experiment (domain-specific/mixed). Displayed are 95 % confidence intervals.

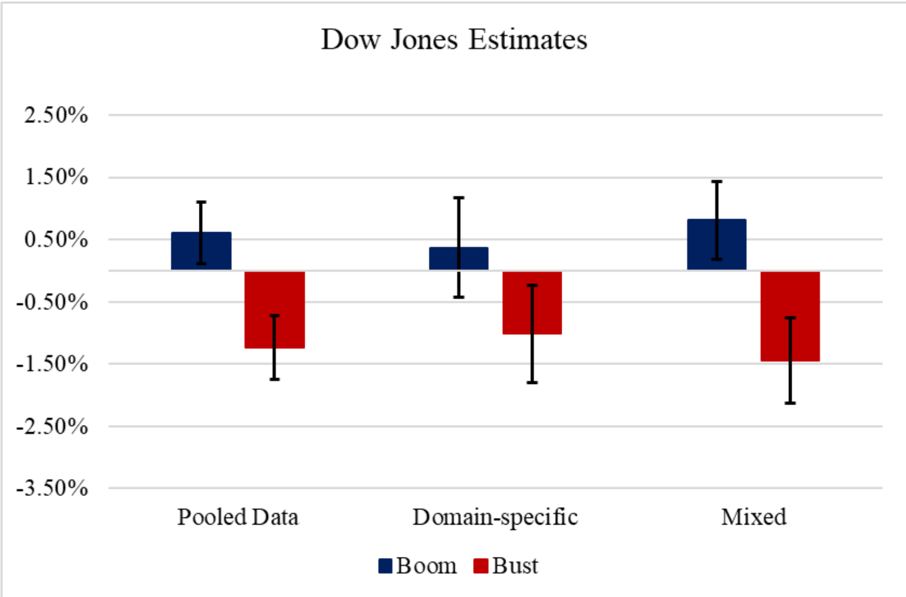

Across both learning environments, we consistently find that subjects in the bust treatment are significantly more pessimistic in their return expectations than subjects in the boom treatment. More strikingly, subjects in the bust treatment provide not only lower return estimates but also negative return estimates, while those in the boom treatment provide positive return estimates on average. Table 7 examines this pattern in greater detail. Return estimates are regressed on a *Bust* indicator as well as our full set of control variables. The negative coefficient on *Bust* across all columns confirms that beliefs about the future performance of the Dow Jones are significantly more pessimistic in the bust treatment. Moreover, we observe that the effect seems to be stronger in absolute magnitude for the negative return estimates, consistent with our documented mechanism that individuals overweight bad outcomes. It remains to stress, that even in such a simple and short-lived learning environment as in our experiment, we are able to systematically manipulate return expectations for real-world market indices.

Table 7: Dow Jones Estimates II

This table reports the results of three OLS regressions on how self-reported return expectations of the Dow Jones Industrial Average are affected by the treatment. We report the results of OLS regressions for the whole sample, and for each experiment individually (Experiment 1: domain-specific; Experiment 2: mixed). The dependent variable in the regression model, *Dow Jones Expectations,* is the self-reported 6-months return expectation for the Dow Jones. The independent variable includes the *Bust* dummy, an indicator variable that equals 1 if participants were in the bust treatment and zero otherwise. Controls include age, gender, statistical skills, self-reported experience in stock trading, whether subjects were invested in the stock market during the last financial crisis, and the order of outcomes they experienced in the sequential belief updating task. Reported are coefficients and t-statistics (in parentheses) using robust standard errors clustered at the individual level. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	*Dow Jones Expectations*		
	Pooled Date	Domain-specific	Mixed
Bust	-0.0191***	-0.0152***	-0.0223***
	(-5.25)	(-2.70)	(-4.68)
Constant	-0.0358**	-0.0476**	-0.0267
	(-2.35)	(-2.03)	(-1.30)
Observations	753	350	403
Controls	X	X	X
R^2	0.052	0.043	0.080

The second set of questions tests to which degree the induced pessimism from the underlying learning environment permeates to different contexts. As a measure of dispositional optimism/pessimism across different life situations, we included a 10-item general Life Orientation Test borrowed from Scheier, Carver, and Bridges (1994), which is frequently used in psychological research (see Appendix C). Figure 9 displays participants' average score in the Life Orientation Test split by treatment. Across all experiments we do not find any significant difference in dispositional optimism/pessimism depending on whether subjects were in the Boom or Bust treatment. Given that asymmetric learning should mostly affect expectations, it is reassuring that we find no treatment differences regarding subjects' inherent psychological traits such as neuroticism, anxiety, self-mastery, or self-esteem.

Figure 9: Life Orientation Test

The figure displays subjects' answers to a general life orientation test. The life orientation test (Scheier, Carver, & Bridges, 1994) is a 10-item inventory where subjects rate statements on a 7-point Likert scale. Displayed is the cumulated score separated by treatment (boom/bust) and by experiment (domain-specific/mixed). Displayed are 95 % confidence intervals.

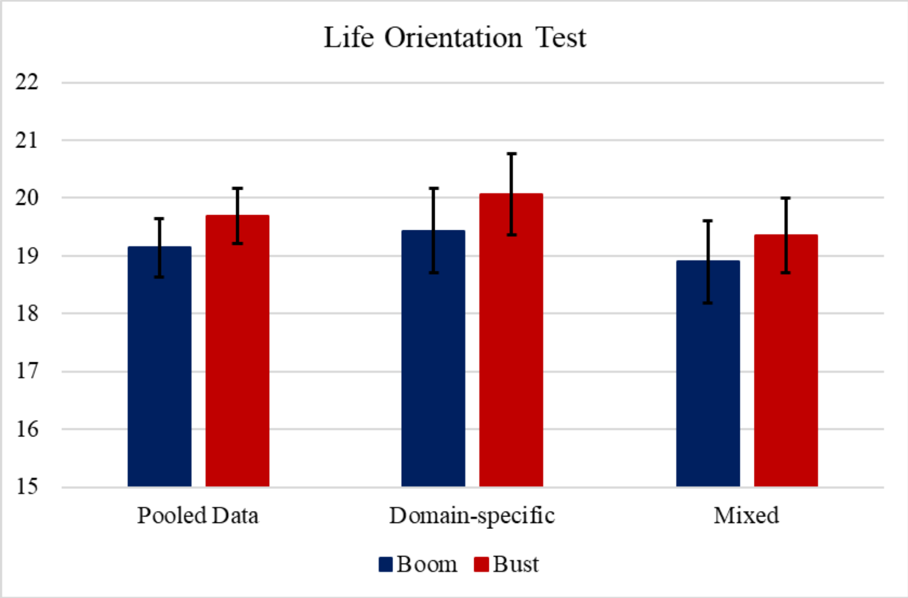

3.4 Further Analyses and Robustness Checks

A. Salience of the Observed Outcome History

An alternative channel through which the risk-taking of our participants might be affected is the salience of the observed outcomes in the first part. More precisely, one may argue that the more frequent exposure to negative (bust treatment) or positive returns (boom treatment) in the sequential belief updating task makes losses and gains more salient, which may subsequently reduce or increase participants' success probability estimate and the investment amount in the ambiguous investment task. To rule out that the salience of the realized returns instead of biased learning drives our results, we investigate the mixed outcome distributions from our second experiment more closely. Irrespective of the boom or bust treatment, participants of the second experiment are exposed to both positive and negative returns with varying frequency. If the salience of the observed outcomes would be the main driver of our results, one would expect that the investment decision in the ambiguous lottery mainly depends on the domain of the most frequently witnessed outcomes instead of the actual treatment. In the second experiment, we can directly test whether this is the case by controlling for the distribution. Whereas participants assigned to the bust treatment are exposed to negative expected value lotteries, they can nonetheless observe a higher share of positive returns if the underlying distribution is the good one. Similarly, participants assigned to the boom treatment are exposed to positive expected value lotteries, but can nonetheless observe a higher share of negative returns if the underlying distribution is the bad one. In other words, whereas the treatment (boom or bust) in the second experiment governs the magnitude of the observed returns, the underlying distributions (good or bad) governs the frequency of positive and negative returns. If salience is the main driver of our results, one would observe a lower (higher) invested amount in the ambiguous lottery when the underlying distribution is the bad (good) one, irrespective of the actual treatment.

Table 8: Underlying Distributions and Subsequent Risk-Taking

This table depicts the average invested amount in the ambiguous lottery and the success probability estimates from the second experiment. Results are displayed separately depending on whether the underlying distributions across the two learning blocks were either both good or both bad and depending on the treatment (boom or bust). Reported are averages and standard errors (in parentheses). *** indicate statistical significance at the 1% level.

	Experiment 2 (Mixed-outcome Lotteries)					
	Distribution (Good / Good)			Distribution (Bad / Bad)		
Variable	**Boom**	**Bust**	**Differ-ence**	**Boom**	**Bust**	**Differ-ence**
Investment Ambiguous Lottery	46.8	36.6	10.2***	44.8	27.5	17.3***
	(32.03)	(32.34)		(30.81)	(30.69)	
Success Probability Estimate (%)	76.3	45.6	30.7***	68.2	36.9	31.3***
	(14.29)	(22.82)		(19.66)	(21.64)	

In Table 8, we report both the average invested amount and the success probability estimate of participants in the boom and bust treatments, split by whether the underlying distributions in the first part of the experiment were either both good or both bad. We find that irrespective of the distribution, participants in the bust treatment are not only more pessimistic about the success probability but also invest less in the ambiguous asset. Even participants assigned to the bust treatment but who witnessed mainly positive returns (i.e. the underlying distributions of both learning blocks were good) are significantly more pessimistic and invest less than those assigned to the boom treatment who experienced mainly negative returns (i.e.

the underlying distributions of both learning blocks were bad). Taken together, we can conclude that the salience of the realized returns is not the main driver of the observed effects.

B. Forecasting Ability and Subsequent Risk-Taking

A frequently expressed concern regarding the impact of investor behavior on market outcomes is that investors' market experience diminishes the importance of psychological forces in financial decisions (e.g. List and Haigh, 2005; or Cipriani and Guarino, 2009). To examine whether the influence of asymmetric learning on investors' risk-taking is diminished for those who are more successful in providing correct forecasts, we define the squared deviation of subjects' probability estimates in each round from the objective posterior probability as a measure of forecasting quality. Next, we conduct median splits with respect to this measure to distinguish above-median forecasters from below-median forecasters. To assess the validity of our measure, we compare the number of correct forecasts (defined in the payment scheme by being in the range of 10 % of the objective forecast) between below- and above-median forecasters. Across both experiments, those subjects who are classified as "above-median" have on average three more correct forecasts than those classified as "below-median" ($p < 0.001$, t-test). Moreover, both measures are highly correlated (Pearson correlation of 0.57, $p < 0.001$).

To better understand to what extent the resulting pessimism through learning from adverse market outcomes is a necessary condition for belief-induced changes in risk-taking, we repeat our main analyses and split by the forecasting ability of our participants in Table 9. Results in Panel A show that our previously drawn conclusion remains stable for the group of above median forecasters. If anything, we find that our reported effects become stronger both in absolute terms and in terms of statistical significance for participants with above-median forecasting ability. In other words, the risk-taking of those agents who achieve more correct forecasts is stronger affected by the learning environment than the risk-taking of agents who achieve less correct forecasts. Besides their investment in the ambiguous lottery, we also consistently find that subjects in the bust treatment are significantly more pessimistic in their assessment of the success probability of the ambiguous lottery. Similar to before, results are stronger for the group of individuals with above-median forecasting performance.

In Panel B of Table 9, we validate that the documented asymmetry in subjects' belief formation is also present in the sub-sample of above-median forecasters. As expected, the bias is less pronounced for subjects with above-median forecasting ability (partly mechanically, due to how splits are conducted). However, and more importantly, the larger pessimism in the bust learning environment still persists and is statistically highly significant.

Table 9: Risk-Taking and Asymmetric Learning Split by Forecasting Ability

Panel A of this table examines subjects' risk-taking and beliefs about the ambiguous asset across treatments split by above- and below-median forecasting ability as defined in the text. We report the results of OLS regressions for the whole sample with controls. For the first two regressions, the dependent variable is *Investment*, which denotes participants' invested amount (0 – 100) in the lottery they were assigned. For the second two regressions, the dependent variable is *Probability Estimate*, which denotes participants' beliefs about the success probability of the ambiguous asset. *Bust* is an indicator variable that equals 1 if participants were in the bust treatment. *Ambiguous* is an indicator variable that equals 1 if participants were asked to invest in the ambiguous lottery, and 0 if they invested in the risky lottery. Panel B of this table investigates subjects' belief updating in the Boom and Bust learning environment split by above- and below-median forecasting ability. The dependent variable in the regression model, *Probability Estimate*, is the subjective posterior belief that the asset is paying from the good distribution. Independent variables include the *Bust* dummy, as defined above, as well as *Objective Posterior*, which is the correct Bayesian probability that the stock is good, given the information seen by the participant up to round *t* in the learning block. Controls include age, gender, statistical skills, self-reported experience in stock trading, whether subjects were invested in the stock market during the last financial crisis, and the order of outcomes they experienced in the forecasting task. Reported are coefficients and t-statistics (in parentheses) using robust standard errors. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Main Finding

	Investment		*Probability Estimate of Ambiguous Lottery*	
	Above Median	Below Median	Above Median	Below Median
Bust	6.126	-1.109	-25.58***	-13.38***
	(1.38)	(-0.25)	(-8.20)	(-4.50)
Ambiguous	10.94***	-1.448		
	(2.65)	(-0.33)		
Bust x Ambiguous	-21.49***	-1.454		
	(-3.54)	(-0.23)		
Constant	1.238	22.65	57.97***	53.54***
	(0.10)	(1.58)	(4.19)	(4.33)
Observations	377	376	187	190
Controls	X	X	X	X
R^2	0.095	0.072	0.333	0.194

Panel B: Belief Updating				
	Probability Estimate (Subjective Posterior)			
	Domain-Specific		Mixed	
	Above Median	Below Median	Above Median	Below Median
Bust	-4.509***	-5.347	-4.925***	-7.787***
	(-3.69)	(-1.61)	(-5.00)	(-3.02)
Objective Posterior	0.645***	0.179***	0.690***	0.188***
	(31.76)	(5.25)	(33.80)	(4.82)
Constant	13.92***	48.56***	28.61***	53.05***
	(2.90)	(4.11)	(6.16)	(4.30)
Observations	2512	1008	3184	880
Controls	X	X	X	X
R^2	0.706	0.125	0.699	0.146

Finally, it remains to address why the risk-taking of the seemingly better performing agents (i.e. the better forecasters) is more affected by the learning environment. One possible explanation could be that our proxy might capture participants' involvement in the experimental task. Effectively, this would suggest that the documented effect is more generalizable outside the experimental environment but limited by the difficulty of the Bayesian updating task. To test whether subjects with above-median forecasting ability are more involved in the experiment, we investigate the time it took to finish the experiment. Interestingly, we find that above-median forecasters spent on average 112 seconds to read the instructions of the forecasting task, while below-median forecasters only spent roughly 86 seconds ($p < 0.05$). Additionally, the overall time to finish the experiment is roughly 580 seconds for above-median forecasters, and about 553 seconds for below-median forecasters ($p < 0.10$). The difference is largely driven by the additional time above-median forecasters spent to read the instructions more carefully, which hints at a significantly higher involvement in the experimental task. Given the strength of the belief-updating bias even in the group of more sophisticated forecasters paired with the higher involvement of the aforementioned group in our experiment, we believe that the effect of different learning environments on risk-taking might be even more pronounced in the real economy.

4. Conclusion

This paper presents direct experimental evidence for the role of expectations for investors' risk-taking behavior across macroeconomic cycles. In a series of experimental studies, we place subjects in learning environments which resemble key characteristics of boom and bust markets and then measure their risk-taking in investment decisions under risk (i.e. known probabilities) or under ambiguity (i.e. unknown probabilities). In our experiments, we have direct control over objective (rational) expectations and can compare them to participants' subjective beliefs. This allows us to document systematic errors in the belief formation process, which we can then relate to the subjects' investment choice.

We show that – conditional on observing the same information – individuals who learn in bust market environments form significantly more pessimistic beliefs than those who learn in boom market environments. The difference in beliefs subsequently translates to a lower willingness to take risk in investments under ambiguity. However, no such effect can be observed for investments under risk. Because we have perfect control over risk preferences in the risky investment decision (where expectations are fixed) and a good measure of expectations in the ambiguous investment decision, we can attribute the observed differences in risk-taking behavior to changes in expectations. Finally, we provide evidence that asymmetric belief updating rules present a plausible mechanism for the pro-cyclicality of beliefs. We show that individuals in bust markets put significantly more weight on low outcomes when updating their expectations compared to those in boom markets. This difference in learning is especially pronounced in situations in which prior beliefs point in the opposite direction of new information.

Overall, our results are in line with recent survey evidence on investors' return expectations being pro-cyclical. The asymmetry in investors' belief formation rules may generate self-reinforcing feedback loops which amplify the intensity and the length of market cycles. For example, subsequent declines in stock prices that render investors to form overly pessimistic expectations, could lead to an increased number of sales which pushes prices further down. Additionally, the mechanism that drives our results in a controlled experimental environment is consistent with the observed behavior in Giglio et al. (2021), which suggests that it may provide interesting implications for recently proposed asset pricing models that incorporate survey evidence (e.g., Barberis et al., 2015; Adam, Marcet and Beutel, 2017).

References

Adam, K., Marcet, A., & Nicolini, J. P. (2016). Stock market volatility and learning. *The Journal of Finance*, 71(1), 33-82.

Adam, K., Marcet, A., & Beutel, J. (2017). Stock price booms and expected capital gains. *American Economic Review*, 107(8), 2352-2408.

Adam, K., Matveev, D., & Nagel, S. (2020). Do survey expectations of stock returns reflect risk-adjustments?. *Journal of Monetary Economics*, forthcoming.

Alempaki, D., Starmer, C., & Tufano, F. (2019). On the priming of risk preferences: The role of fear and general affect. *Journal of Economic Psychology*, forthcoming.

Amromin, G., & Sharpe, S. A. (2013). From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science*, *60*(4), 845-866.

Barberis, N., Huang, M., & Santos, T. (2001). Prospect theory and asset prices. *The Quarterly Journal of Economics*, *116*(1), 1-53.

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

Benoît, J. P., & Dubra, J. (2018). When do populations polarize? An explanation. Working Paper.

Bordalo, P., Gennaioli, N., & Shleifer, A. (2018). Diagnostic expectations and credit cycles. *The Journal of Finance*, 73(1), 199-227.

Bordalo, P., Gennaioli, N., Porta, R. L., & Shleifer, A. (2019). Diagnostic expectations and stock returns. *The Journal of Finance*, 74(6), 2839-2874.

Brunnermeier, M. K., & Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4), 1092-1118.

Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data?. *Perspectives on Psychological Science*, *6*(1), 3-5.

Campbell, J. Y., & Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, *107*(2), 205-251.

Cipriani, M., & Guarino, A. (2009). Herd behavior in financial markets: an experiment with financial market professionals. *Journal of the European Economic Association*, 7(1), 206-233.

Charness, G., & Dave, C. (2017). Confirmation bias with motivated beliefs. *Games and Economic Behavior*, 104, 1-23.

Charness, G., Gneezy, U., & Imas, A. (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior & Organization*, 87, 43-51.

Cohn, A., Engelmann, J., Fehr, E., & Maréchal, M. A. (2015). Evidence for countercyclical risk aversion: An experiment with financial professionals. *American Economic Review*, *105*(2), 860-85.

Collin-Dufresne, P., Johannes, M., & Lochstoer, L. A. (2016). Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review*, 106(3), 664-98.

Enke, B., Schwerter, F., & Zimmermann, F. (2020). *Associative memory and belief formation* (No. w26664). National Bureau of Economic Research.

Giglio, S., Maggiori, M., Stroebel, J., & Utkus, S. (2020). *Five facts about beliefs and portfolios*. American Economic Review, forthcoming.

Giglio, S., Maggiori, M., Stroebel, J., & Utkus, S. (2021). The joint dynamics of investor beliefs and trading during the COVID-19 crash. *Proceedings of the National Academy of Sciences*, 118(4).

Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics*, *112*(2), 631-645.

Goodman, J. K., Cryder, C. E., & Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making*, *26*(3), 213-224.

Greenwood, R., & Shleifer, A. (2014). Expectations of returns and expected returns. *The Review of Financial Studies*, *27*(3), 714-746.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, 95(3), 537-557.

Guiso, L., Sapienza, P., & Zingales, L. (2018). Time varying risk aversion. *Journal of Financial Economics*, *128*(3), 403-421.

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.

Harrison, G. W., List, J. A., & Towe, C. (2007). Naturally occurring preferences and exogenous laboratory experiments: A case study of risk aversion. *Econometrica*, 75(2), 433-458.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3), 430-454.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, *73*(6), 1849-1892.

König-Kersting, C., & Trautmann, S. T. (2018). Countercyclical risk aversion: Beyond financial professionals. *Journal of Behavioral and Experimental Finance*, *18*, 94-101.

Kuhnen, C. M. (2015). Asymmetric learning from financial information. *The Journal of Finance*, *70*(5), 2029-2062.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480.

Malmendier, U., & Nagel, S. (2011). Depression babies: do macroeconomic experiences affect risk taking?. *The Quarterly Journal of Economics*, *126*(1), 373-416.

Malmendier, U., & Nagel, S. (2015). Learning from inflation experiences. *The Quarterly Journal of Economics*, *131*(1), 53-87.

Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502.

Nagel, S., & Xu, Z. (2019). Asset pricing with fading memory. *Review of Financial Studies*, forthcoming.

Scheier, M. F., Carver, C. S., & Bridges, M. W. (1994). Distinguishing optimism from neuroticism (and trait anxiety, self-mastery, and self-esteem): a reevaluation of the Life Orientation Test. *Journal of Personality and Social Psychology*, *67*(6), 1063.

Shiller, R. J., & Grossman, S. J. (1981). *The determinants of the variability of stock market prices*. National Bureau of Economic Research.

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

Weber, M., Weber, E. U., & Nosić, A. (2013). Who takes risks when and why: Determinants of changes in investor risk taking. *Review of Finance, 17*(3), 847-883.

Appendix

A. Further Analyses

Figure A1: Asymmetric Learning in Boom and Bust Markets (Experiment 1)

This figure displays average subjective probability estimates that the risky asset is paying dividends from the good distribution as a function of the objective Bayesian probability. Subjective probability estimates provided by participants for each objective Bayesian probability are shown as a blue line for the Bust treatment, and a red line for the Boom treatment. If subjective probability estimates were Bayesian, they would equal the objective probabilities and thus would line up on a line starting in the origin with slope one. Displayed are 90% confidence intervals.

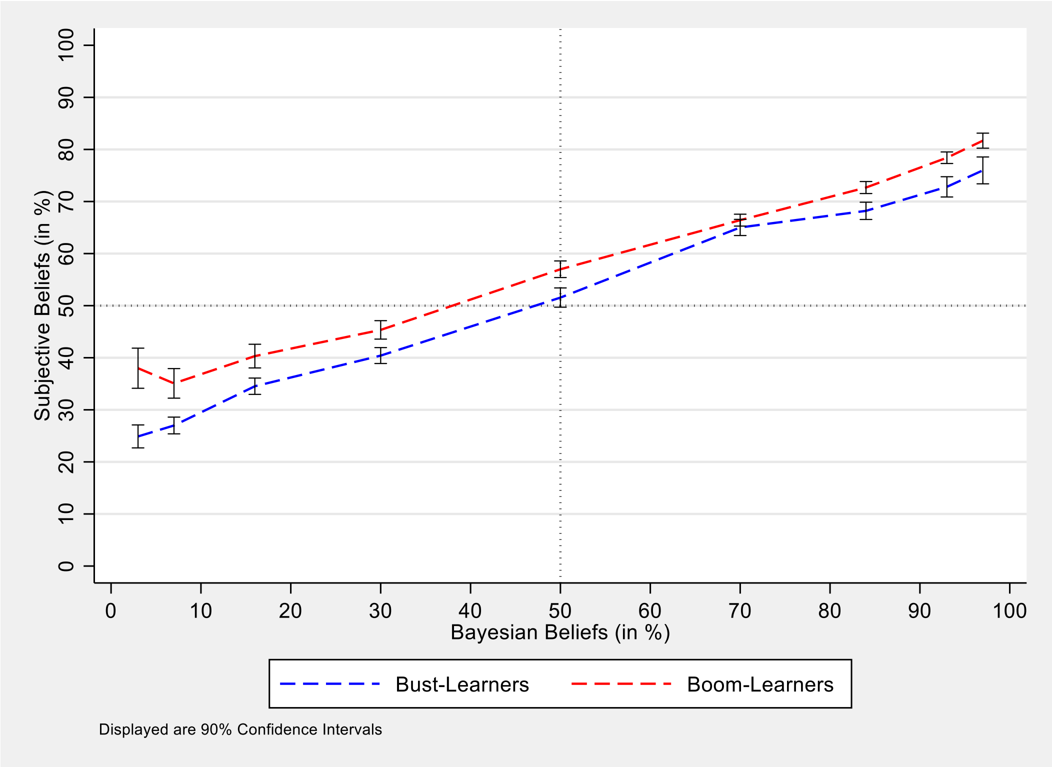

Figure A2: Asymmetric Learning in Boom and Bust Markets (Experiment 2)

This figure displays average subjective probability estimates that the risky asset is paying dividends from the good distribution as a function of the objective Bayesian probability. Subjective probability estimates provided by participants for each objective Bayesian probability are shown as a blue line for the Bust treatment, and a red line for the Boom treatment. If subjective probability estimates were Bayesian, they would equal the objective probabilities and thus would line up on a line starting in the origin with slope one. Displayed are 90% confidence intervals.

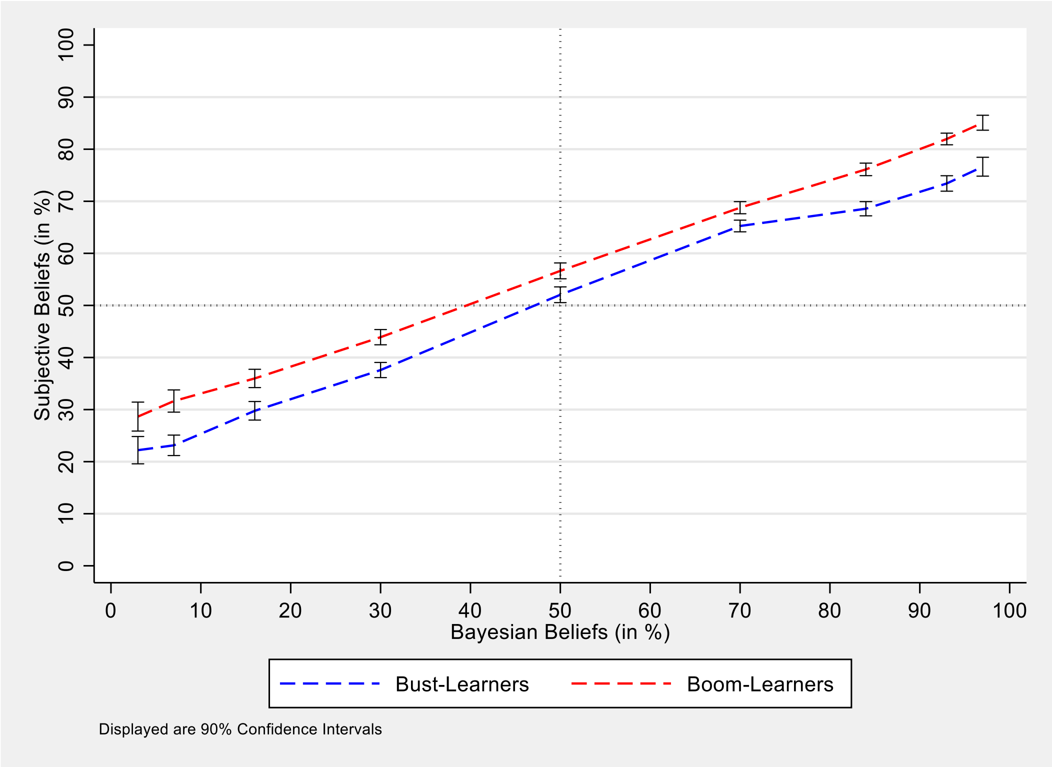

Table A1: Relation Between Beliefs About Success Probability and Investment

This table examines whether subjects in our experiment act upon their beliefs about the success probability of the ambiguous lottery, split by above- and below-median forecasting ability as defined in the text. Dependent variable is *Investment Ambiguous*, which captures subjects' invested amount in the ambiguous lottery. *Success Probability* denotes participants' beliefs about the success probability of the ambiguous lottery. *Bust* is an indicator variable that equals 1 if participants were in the bust treatment. Controls include age, gender, statistical skills, self-reported experience in stock trading and whether subjects were invested in the stock market during the last financial crisis. Reported are coefficients and t-statistics (in parentheses) using robust standard errors. *, **, *** indicates statistical significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable	*Investment in Ambiguous Lottery*					
	Pooled Data		**Domain-specific**		**Mixed**	
	Above Median	Below Median	Above Median	Below Median	Above Median	Below Median
Success Probability	0.387***	0.364***	0.356***	0.369***	0.515***	0.437**
	(4.07)	(3.39)	(2.73)	(2.67)	(3.31)	(2.57)
Bust	-6.384	3.036	-10.71*	1.334	-0.186	5.905
	(-1.42)	(0.63)	(-1.98)	(0.21)	(-0.02)	(0.76)
Constant	-20.72	13.96	-43.07**	27.90	11.33	-3.123
	(-1.21)	(0.75)	(-2.22)	(1.33)	(0.39)	(-0.09)
Observations	187	190	85	92	102	98
R^2	0.22	0.11	0.27	0.18	0.24	0.11

B. Experimental Instructions and Screenshots

Instructions Bayesian Updating (Exemplary for Boom Treatment of Experiment 1)

In this part, we would like to test your forecasting abilities. You will make forecasting decisions in two consecutive blocks each consisting of 8 rounds.

Suppose you find yourself in an environment, in which the value of a risky asset can either increase by 2 or by 15. The probability of either outcome (2 or 15) depends on the state in which the asset is (**good** state or **bad** state). If the risky asset is in the **good** state, then the probability that the risky asset increases in value by 15 is 70% and the probability that it increases in value by 2 is 30%. If the risky asset is in the **bad** state, then the probability that the risky asset increases in value by 15 is 30% and the probability that it increases in value by 2 is 70%.

The computer determines the state at the beginning of each block (consisting of 8 rounds). Within a block, the state does not change and remains fixed. At the beginning of each block, you do not know which state the risky asset is in. The risky asset may be in the good state or in the bad state with equal probability.

At the beginning of each round, you will observe the payoff of the risky asset (2 or 15). After that, we will ask you to provide a probability estimate that the risky asset is in the good state and ask you how sure you are about your probability estimate. While answering these questions, you can observe the price development in a chart next to the question.

There is always an objective correct probability that the risky asset is in the good state. This probability depends on the history of payoffs of the risky asset already. As you observe the payoffs of the risky asset, you will update your beliefs whether or not the risky asset is in the good state.

Every time you provide us with a probability estimate that is within 5% of the correct value (e.g., correct probability is 70% and your answer is between 65% and 75%) we will add 10 Cents to your payment.

Objective Bayesian Posterior Probabilities

This table provides all possible values for the objectively correct probability that the asset is in the good state for every possible combination of rounds and outcomes. The initial prior for good and bad distribution is set to 50%. The objective Bayesian posterior probability that the asset is in the good state, after observing t high outcomes in n rounds so far is given by: $\frac{1}{1+\frac{1-p}{p}\cdot(\frac{q}{1-q})^{n-2t}}$, where p is the initial prior before any outcome is observed that the risky asset is in the good state (50% here), and q is the probability that the value increase of the asset is the higher one (70% here).

n (number of rounds so far)	t (number of high outcomes so far)	Probability [stock is good t high outcomes in n rounds]
0	0	50.00%
1	0	30.00%
1	1	70.00%
2	0	15.52%
2	1	50.00%
2	2	84.48%
3	0	7.30%
3	1	30.00%
3	2	70.00%
3	3	92.70%
4	0	3.26%
4	1	15.52%
4	2	50.00%
4	3	84.48%
4	4	96.74%
5	0	1.43%
5	1	7.30%
5	2	30.00%
5	3	70.00%
5	4	92.70%
5	5	98.57%
6	0	0.62%
6	1	3.26%
6	2	15.52%
6	3	50.00%
6	4	84.48%
6	5	96.74%
6	6	99.38%
7	0	0.26%
7	1	1.43%
7	2	7.30%
7	3	30.00%
7	4	70.00%
7	5	92.70%
7	6	98.57%
7	7	99.74%
8	0	0.11%
8	1	0.62%
8	2	3.26%
8	3	15.52%
8	4	50.00%
8	5	84.48%
8	6	96.74%
8	7	99.38%
8	8	99.89%

Screenshots of Experiment 1

Figures B1 to B3 present the screens of the sequential belief updating task as seen by subjects in the experiment (example block 1, round 5). One round consists of three sequential screens. First, subjects saw the payoff of the risky asset in the respective round. Second, the cumulated payoffs of the risky asset are shown in a price-line-chart and subjects are asked to provide a probability estimate that the risky asset pays from the good distribution. Finally, subjects are asked on a 9-point Likert scale how confident they are in their probability estimate.

Figure B1: Payoff Screen

Figure B2: Probability Estimate Screen

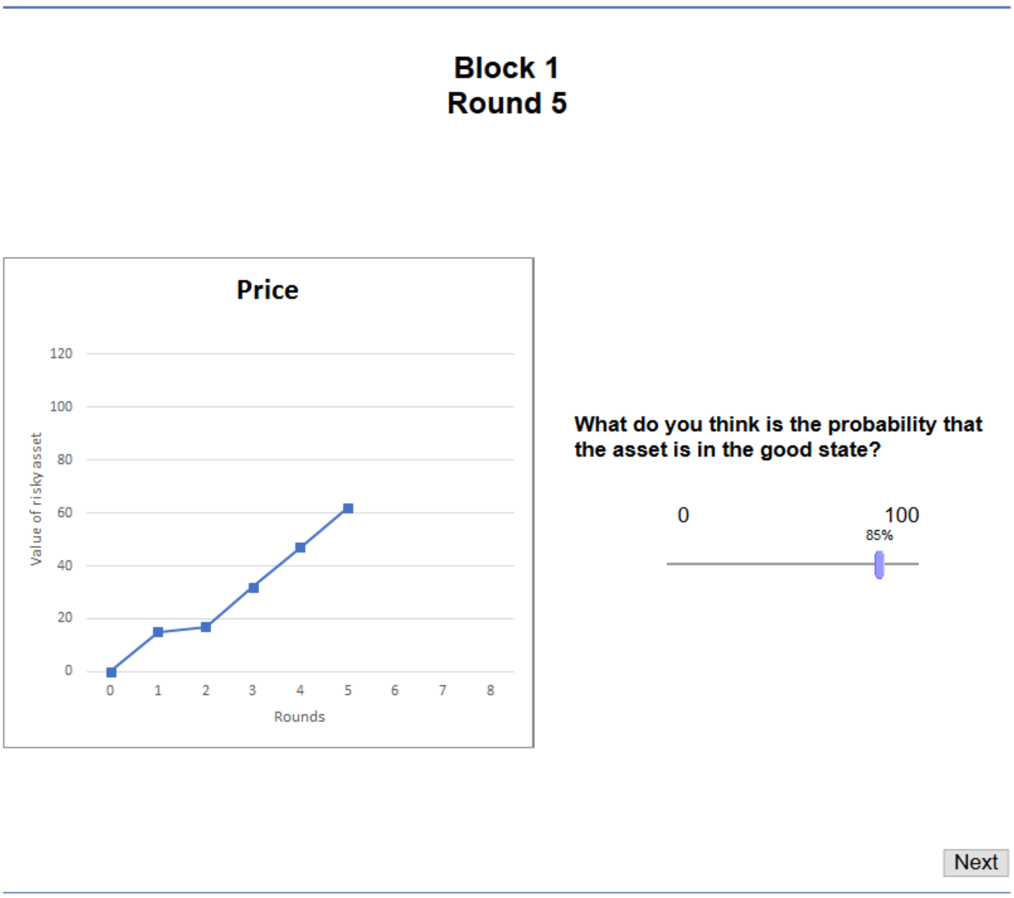

Figure B3: Confidence Level Screen

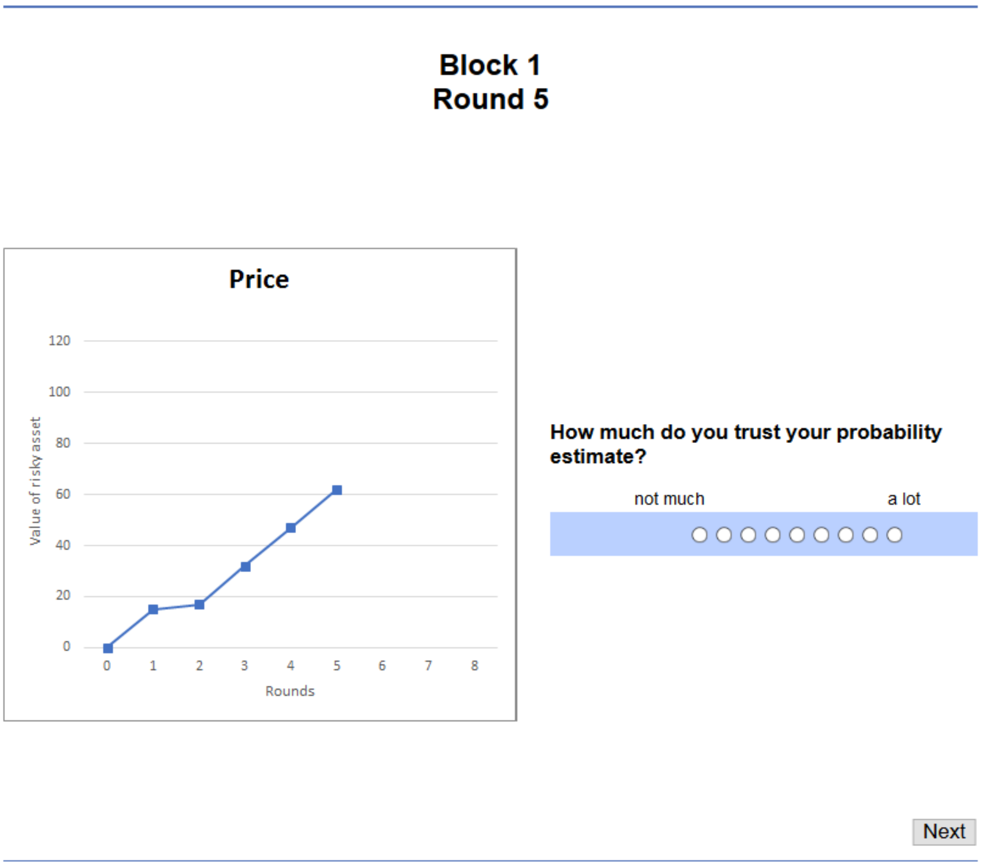

C. Experimental Measures in Experiment 1 and 2

Risky Lottery

Imagine in the stock market there is a risky asset, in which you can invest 100 Cent now. The asset pays you either 2.5 times the amount you invest or it becomes valueless, i.e. your invested amount is lost. **The probability of either outcome is exactly 50%**.

You can keep whatever amount you decide not to invest in the risky asset.

How much of your endowment do you want to invest in the risky asset?

[Dropdown Menu of all possible combinations in 5 Cent steps]

Ambiguous Lottery

Imagine in the stock market there is a risky asset, in which you can invest 100 Cent now. The asset pays you either 2.5 times the amount you invest or it becomes valueless, i.e. your invested amount is lost. **However, the probability of either outcome is unknown**.

You can keep whatever amount you decide not to invest in the risky asset.

How much of your endowment do you want to invest in the risky asset?

[Dropdown Menu of all possible combinations in 5 Cent steps]

Life Orientation Test

Below we report the questions used in the revised version of the Life Orientation Test developed by Scheier, Carver, and Bridges (1994). All questions were answered on a 5-point Likert scale from “do not agree at all” to “fully agree”. Reverse-coded items are indicated by [R]. Filler-items are indicated by [F]. The non-filler items were added to a final score.

1. In uncertain times, I usually expect the best.
2. It’s easy for me to relax. [F]
3. If something can go wrong, it will. [R]
4. I’m always optimistic about my future.
5. I enjoy my friends a lot. [F]
6. It’s important for me to keep busy. [F]
7. I hardly ever expect things to go my way. [R]
8. I don’t get upset too easily. [F]
9. I rarely count on good things happening to me. [R]
10. Overall, I expect more good things to happen to me than bad.

Comprehension Question for Bayesian Updating Task

Below we report the comprehension questions that participants had to answer correctly after reading the instructions to proceed to the Bayesian Updating task. Correct responses are displayed in italic.

1. If you see a series of +15 [-2 for Bust treatment], what is more likely?
    a. *The risky asset is in the good state.*
    b. The risky asset is in the bad state.

2. The correct probability estimate is let's say 0.70. Which probability estimate(s) would be in the range such that you earn 10 cents? [Note: You can check multiple boxes.]
    a. 0.55
    b. *0.67*
    c. *0.75*
    d. 0.85
    e. 0.87

3. At the beginning of each block, the probability that the risky asset is in the good state is 50%.
    a. *True*
    b. False

Dow Jones Return Expectations Question in Experiment 1

The Dow Jones Industrial Average (Stock Market Index of the 30 largest US companies) is currently trading at around 25,343.

In which price range would you expect this index to trade in 6 months from now?

[Dropdown]

- < 23,000
- 23,000 – 23,500
- 23,501 – 24,000
- 24,001 – 24,500
- 24,501 – 25,000
- 25,001 – 25,500
- 25,501 – 26,000
- 26,001 – 26,500
- 26,501 – 27,000
- 27,001 – 27,500
- 27,501 – 28,000
- > 28,000

Dow Jones Return Expectations Question in Experiment 2

The Dow Jones Industrial Average (Stock Market Index of the 30 largest US companies) is currently trading at around 26,770.

In which price range would you expect this index to trade in 6 months from now?

[Dropdown]

- < 24,500
- 24,500 – 25,000
- 25,001 – 25,500
- 25,501 – 26,000
- 26,001 – 26,500
- 26,501 – 27,000
- 27,001 – 27,500
- 27,501 – 28,000
- 28,001 – 28,500
- 28,501 – 29,000
- 29,001 – 29,500
- > 29,500