The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXX, NO. 5 • OCTOBER 2015

Asymmetric Learning from Financial Information

CAMELIA M. KUHNEN*

ABSTRACT

This study asks whether investors learn differently from gains versus losses. I find experimental evidence that indicates that being in the negative domain leads individuals to form overly pessimistic beliefs about available investment options. This pessimism bias is driven by people reacting more to low outcomes in the negative domain relative to the positive domain. Such asymmetric learning may help explain documented empirical patterns regarding the differential role of poor versus good economic conditions on investment behavior and household economic choices.

DO INVESTORS LEARN THE same way when they face positive outcomes as when they face negative outcomes? Do economic agents form beliefs using the same learning rules during recessions as during booms? Converging findings from finance and neuroscience suggest that this may not be the case.

Recent empirical finance work indicates that learning by market participants may differ depending on whether the economic conditions are good or bad. Economic downturns are characterized by stronger reactions to negative news by equity markets, higher risk premia, and more pessimistic expectations by corporate executives (Andersen et al. (2007), Bollerslev and Todorov (2011), Ben-David, Graham, and Harvey (2013)). Poor stock market outcomes receive disproportionately pessimistic press coverage (Garcia (2012)). Households that witness bad economic times become reluctant to invest in equities and have pessimistic beliefs about future stock returns (Malmendier and Nagel (2011)). After floods or earthquakes, people are more likely to buy insurance against

*Kuhnen is with the Kenan-Flagler Business School at the University of North Carolina at Chapel Hill. I thank Bruno Biais (the Editor), two anonymous referees, Nicholas Barberis, Marianne Bertrand, Cary Frydman, Cam Harvey, David Hofmann, Eric Hughson, Jonathan Parker, and Richard Todd, as well as seminar participants at Northwestern University, University of Washington, University of Utah, Santa Clara University, University of Oregon, Caltech, University of Southern California, Massachusetts Institute of Technology, Stanford University, Arizona State University, UCLA, University of Minnesota, UC Davis, Duke University, University of North Carolina, New York University, and the Consumer Financial Protection Bureau and participants at the 2011 Society for Neuroeconomics annual meeting, the 2012 Western Finance Association meeting, the 2012 Boulder Consumer Finance meeting, the 2012 Miami Behavioral Finance conference, the 2013 UNC Jackson Hole Winter Finance conference, the 2013 Behavioral Economics annual meeting, and the 2013 NBER Household Finance Summer Institute for helpful comments and discussion. Alexandra Baleanu provided excellent research assistance. All remaining errors are mine.

DOI: 10.1111/jofi.12223

such events, even though the probability of their occurrence does not change (Froot (2001)). This empirical evidence suggests that bad times, characterized by a preponderance of negative outcomes, may have a particularly strong influence on people's beliefs about the future.

Moreover, neuroscience evidence indicates that the brain processes deployed when people learn from their environment differ depending on whether they are faced with positive or negative outcomes (Kuhnen and Knutson (2005), Knutson and Bossaerts (2007)). Memory processes are different for details related to positive contexts than for those related to negative contexts (Eppinger, Herbert, and Kray (2010), Mather and Schoeke (2011)), in that negative contexts lead to a more narrow focus than positive ones. People's emotional reactions are stronger in the face of losses, relative to gains, and this is particularly true when the stakes are higher (Sokol-Hessner, Camerer, and Phelps (2013)). This biology-based evidence suggests that people perceive and incorporate negative outcomes differently than positive ones.

Here, I use an experimental setting to examine whether people indeed learn differently from gains or positive news relative to losses or negative news. I find that, when they are in the negative domain, people form overly pessimistic beliefs about the available financial assets, particularly if they are actively investing. This pessimism bias is driven by an overreaction to low outcomes in the negative domain relative to the positive domain. These results are robust to alternative explanations and they replicate out of sample.

The idea that learning may be different in the gain and loss domains is different from and complementary to the well-documented phenomenon of loss-aversion suggested by Kahneman and Tversky (1979), whereby the disutility of losing an amount of money is greater, in absolute terms, than the utility of winning that amount. A large body of work has provided evidence for this difference in preferences in the gain and loss domains. The findings that I document here suggest that gains and losses are different not only in terms of how they shape the value function, but also in terms of how they are incorporated in the formation of beliefs.

To investigate whether learning is different when people face negative outcomes relative to when they face positive ones, adult participants from a U.S. university were invited to a study that required the completion of two financial decision making tasks. In the Active task, subjects made 60 decisions, split into 10 separate blocks of six trials each, to invest in one of two securities: a stock with risky payoffs coming from one of two distributions (good and bad), one that was better than the other in the sense of first-order stochastic dominance, and a bond with a known payoff. In each trial, participants observed the dividend paid by the stock after making their asset choice, and then were asked to provide an estimate of the probability that the stock was paying from the good distribution. In the Passive task, subjects were only asked to provide the probability estimate that the stock was paying from the good distribution, after observing its payoff in each of 60 trials, which were also split into 10 separate learning blocks of six trials each. In either task, two types of conditions—gain or loss—were possible. In the gain condition, the two securities provided positive payoffs only. In the loss condition, the two securities provided negative payoffs only. Subjects were paid based on their investment payoffs and the accuracy of the probability estimates provided.

Importantly, the learning problem faced by subjects was exactly the same in gain condition blocks as in loss condition blocks. The only difference was that the two possible stock payoffs had a minus sign in front of them in the loss condition relative to the gain condition (i.e., −\$10 or −\$2 in the former vs. +\$2 or +\$10 in the latter condition). Hence, people's estimate regarding the probability that the stock was paying from the good dividend distribution, namely that distribution where the high outcome for that condition was more likely to occur than the low outcome, should not depend on whether they are in a block where they learn from negative outcomes, or in one where they learn from positive outcomes.

However, I find that subjects learn differently in the gain condition relative to the loss condition. Subjective probability estimates that the stock is paying from the good dividend distribution are 3% to 5% lower in the loss condition than in the gain condition, controlling for the objective Bayesian posterior probability that the stock is the good one given the dividends observed by participants. That is, subjective beliefs about the risky asset are overly pessimistic in the loss condition. Moreover, the deviation of subjective probability estimates from the objective Bayesian posterior that the stock is the good one is 2% to 4% larger in the loss condition relative to the gain condition. In other words, belief errors are on average larger when people learn from negative payoffs than from positive ones. The pessimism bias and resulting larger deviation in subjective posteriors from Bayesian beliefs in the loss relative to the gain condition are generated by the fact that people update more from a low outcome in the loss condition (i.e., a −\$10 dividend) than in the gain condition (i.e., a +\$2 dividend). There is no difference between the two conditions in terms of updating beliefs from high outcomes (i.e., either −\$2 or a +\$10 dividend, in the loss and gain conditions, respectively).

I then conduct several robustness tests and find that the loss versus gain condition effect on subjective beliefs is robust in-sample, whether I analyze data from early or late learning blocks, or, within each learning block of six trials, from early or late trials, or whether I conduct my analysis with or without subject fixed effects.

Moreover, I show that the effect replicates out of sample, in a population more than twice as large as in the original group of subjects, and in a different country (Romania). There, too, I find that the loss condition induces larger errors in subjective beliefs, and an overreaction to low outcomes, just as was found in the U.S. sample.

Finally, I examine several alternative explanations for the documented learning effects induced by the loss versus gain context. I test whether in the loss condition, relative to the gain condition, people may start with different priors that the stock is good, whether they may have different risk attitudes and whether their beliefs may have a different impact on asset choices across the two contexts. Finally, I test whether the experimental task that I use in fact engages the subjects' learning processes. All of these tests provide evidence that supports the documented loss versus gain condition effect on subjective beliefs.

Learning in financial markets has been the focus of a small but growing experimental literature. Kluger and Wyatt (2004) document the existence of heterogeneity across traders with respect to their ability to learn according to Bayes's rule, and the impact of this heterogeneity on asset prices. Asparouhova et al. (2010) find that investors unable to perform correct probability computations prefer to hold portfolios with unambiguous returns and do not directly influence asset prices. Payzan-LeNestour (2010) shows that Bayesian learning is a reasonably good model for investment decisions in complex settings. Bruguier, Quartz, and Bossaerts (2010) show that the ability to forecast price patterns in financial markets depends on traders' capacity to understand others' intentions, and not on the ability to solve abstract mathematical problems. Kogan (2009) and Carlin, Kogan, and Lowery (2013) show that strategic considerations influence learning and trading in experimental asset markets. In addition, there exists a novel body of theoretical work focused on understanding the role of bounded rationality and nonstandard preferences in the formation of beliefs by economic agents.[1] This body of work assumes that individuals learn according to Bayes's rule, given a possibly incorrect prior belief and possibly sparse new information. The focus of this paper is complementary to this literature, as the evidence presented here sheds light on the process by which people incorporate newly available information into beliefs starting from objective priors, and documents domain-specific departures from Bayesian learning.

The novel contribution of this paper, therefore, is to show that the ability to learn from financial information is different in the gain and the loss domains. In particular, in the loss domain, people form beliefs about available investment options, which are overly pessimistic and further away from Bayesian beliefs, relative to the gain domain. I describe the experimental design in Section I. The main result, as well as the replication study and tests of alternative explanations, are presented in Section II. In Section III, I discuss the implications of the pessimism bias induced by the loss domain for underinvestment in the context of household finance, corporate finance, and development economics, and suggest avenues for further research building on this finding. Finally, in Section IV, I conclude.

I. Experimental Design

Eighty-seven individuals (37 males, 50 females, mean age 20 years, 1.6 years standard deviation) were recruited at Northwestern University (Evanston, IL, USA) and participated in the experiment. Each participant completed two financial decision-making tasks, referred to as the Active task and the Passive task, during which information about two securities, a stock and a bond, was

[1] See Barberis, Shleifer, and Vishny (1998), Bossaerts (2004), Brunnermeier and Parker (2005), Gabaix et al. (2006), Van Nieuwerburgh and Veldkamp (2010), and Gennaioli and Shleifer (2010).

presented. Whether a participant was presented with the Active task first, or the Passive task first, was determined at random.

Each task included two types of conditions: gain or loss. In the gain condition, the two securities provided positive payoffs only. The stock payoffs were +\$10 or +\$2, while the bond payoff was +\$6. In the loss condition, the two securities provided negative payoffs only. The stock payoffs were −\$10 or −\$2, while the bond payoff was −\$6.

In either condition, the stock paid dividends from either a good distribution or from a bad distribution. The good distribution is that where the high outcome occurs with 70% probability in each trial, while the low outcome occurs with 30% probability. The bad distribution is that where these probabilities are reversed: the high outcome occurs with 30% probability, and the low outcome occurs with 70% probability in each trial.

Each participant went through 60 trials in the Active task, and 60 trials in the Passive task. Trials are split into "learning blocks" of six: for these six trials, the learning problem is the same. That is, the computer either pays dividends from the good stock distribution in each of these six trials, or it pays from the bad distribution in each of the six trials. At the beginning of each learning block, the computer randomly selects (with 50%/50% probabilities) whether the dividend distribution to be used in the following six trials will be the good or the bad one.

There are 10 learning blocks in the Active task, and 10 learning blocks in the Passive task. In either task, there are five blocks in the gain condition, and five blocks in the loss condition. The order of the blocks is pseudo-randomized. An example of a sequence of loss or gain learning blocks the subject may face during either the Active task or the Passive task, as well as a summary of the experimental design, are shown in Table I.

In the Active task, participants made 60 decisions (six in each of the 10 learning blocks) to invest in one of the two securities, the stock or the bond, then observed the stock payoff (irrespective of their choice) and provided an estimate of the probability that the stock was paying from the good distribution. Figure 1 shows the time line of a typical trial in the Active task, in either the gain and or the loss conditions (top and bottom panel, respectively).

In the Passive task, participants were only asked to provide the probability estimate that the stock was paying from the good distribution, after observing its payoff in each of 60 trials (split into 10 learning blocks of six trials each, as in the Active task). Figure 2 shows the time line of a typical trial in the Passive task, for both the gain and the loss conditions.

In the Active task participants were paid based on their investment payoffs and the accuracy of the probability estimates provided. Specifically, they received one tenth of accumulated dividends, plus 10 cents for each probability estimate within 5% of the objective Bayesian value. In the Passive task, participants were paid based solely on the accuracy of the probability estimates provided, by receiving 10 cents for each estimate within 5% of the correct value. Information regarding the accuracy of each subject's probability estimates and the corresponding payment was only provided at the end of each of the two

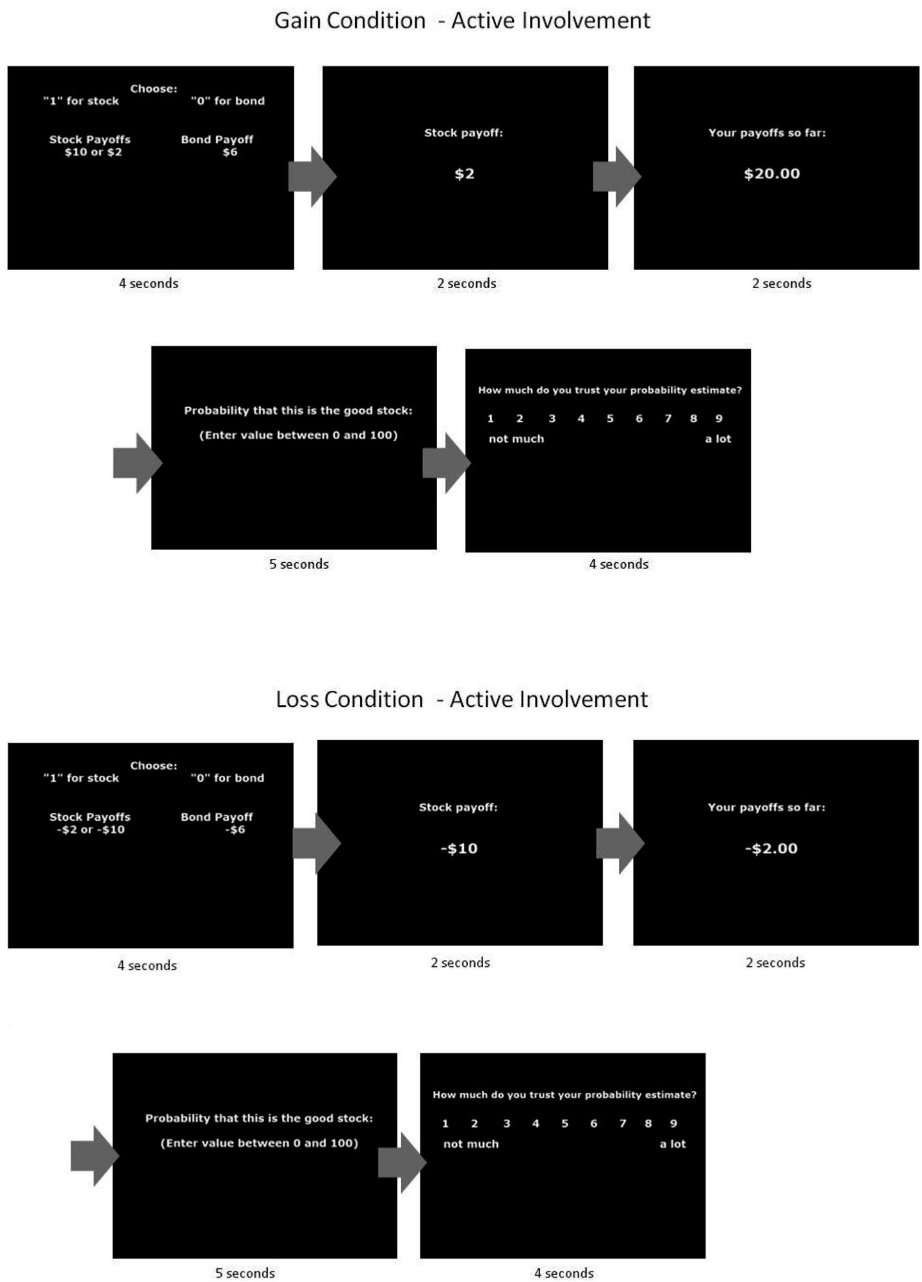

Figure 1. Active task. This figure provides an example of a gain condition trial (top panel) and a loss condition trial (bottom panel). In either type of trial, subjects first choose between the stock and the bond. Then they observe the dividend paid by the stock that trial, no matter which asset they chose, and then are reminded of how much they have earned so far from the payoffs of the assets chosen so far in the Active investment task. Finally, they are asked to provide an estimate for the probability that the stock is paying from the good dividend distribution, and their confidence in this estimate.

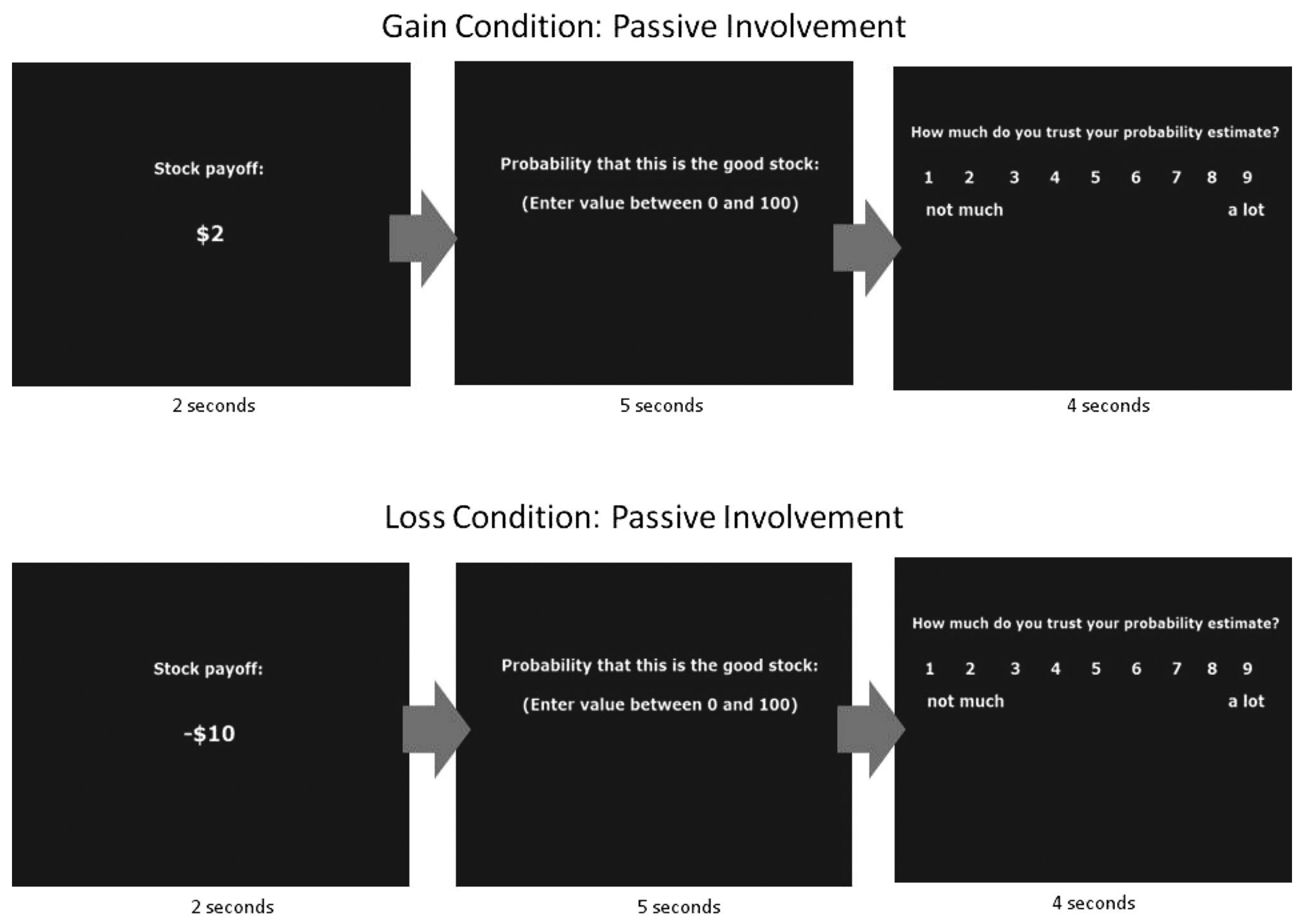

Figure 2. Passive task. This figure provides an example of a gain condition trial (top panel) and a loss condition trial (bottom panel). In either type of trial, subjects observe the dividend paid by the stock that trial. Then they are asked to provide an estimate for the probability that the stock is paying from the good dividend distribution, and their confidence in this estimate.

Table I

Experimental Design

Each participant goes through 60 trials in the Active task and 60 trials in the Passive task. Whether the participant does the Active task first or the Passive task first is determined at random. Trials are split into learning blocks of six trials: for these six trials, the learning problem is the same. That is, the computer either pays dividends from the good stock distribution in each of these six trials, or it pays from the bad distribution in each of the six trials. The good distribution is that where the high dividend occurs with 70% probability in each trial, while the low outcome occurs with 30% probability. The bad distribution is that where these probabilities are reversed: the high outcome occurs with 30% probability, and the low outcome occurs with 70% probability in each trial. At the beginning of each learning block, the computer randomly selects (with 50%/50% probabilities) whether the dividend distribution to be used in the following six trials will be the good or the bad one. There are 10 learning blocks in the Active task, and 10 learning blocks in the Passive task. In either task, there are five blocks in the gain condition and five blocks in the loss condition. The order of the blocks is pseudo-randomized. An example of a sequence of loss or gain blocks that a participant may face is shown below.

Active Task	See Figure 1 for examples of trials	Condition
Block 1	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss
Block 2	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain
Block 3	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain
.	.	.
.	.	.
.	.	.
Block 9	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain
Block 10	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss
Passive Task	See Figure 2 for examples of trials	Condition
Block 1	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain
Block 2	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss
Block 3	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain
.	.	.
.	.	.
.	.	.
Block 9	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss
Block 10	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss

tasks. This was done to avoid feedback effects that could have changed the participants' strategy or answers during the progression of each of the two tasks.

This information was presented to participants at the beginning of the experiment, and is summarized in the participant instructions sheet included in the Appendix. The experiment lasted 1.5 hours and the average payment per person was $30.48.

The value of the objective Bayesian posterior that the stock is paying from the good distribution can be easily calculated. Specifically, after observing t high outcomes in n trials so far, the Bayesian posterior that the stock is the good one is given by: $\frac{1}{1+\frac{1-p}{p}*(\frac{q}{1-q})^{n-2t}}$, where $p = 50\%$ is the prior that the stock is the good one (before any dividends are observed in that learning block) and $q = 70\%$ is the probability that a good stock pays the high (rather than low) dividend in each trial. The Appendix provides the value of the objective Bayesian posterior for all $\{n, t\}$ pairs possible in the experiment. This Bayesian posterior is my benchmark for measuring how close the subjects' expressed probability estimates are to the objectively correct beliefs.

For each participant, I also obtained measures of their financial literacy and risk aversion. I obtained these measures by asking subjects two questions regarding a portfolio allocation problem, after they completed the Active and Passive investment tasks. These questions are described in the Appendix. Briefly, the first question asked how much of a $10,000 portfolio the participant would allocate to the stock market and how much to a savings account. This answer provides a proxy for their risk preference, measured outside of the financial learning experiment. The second question asked the person to calculate the expected value of the portfolio they selected, and through multiple-choice answers could detect whether people lacked an understanding of probabilities, of the difference between net and gross returns, or of the difference between stocks and savings accounts. This yielded a financial knowledge score of zero to three, depending on whether the participant's answer showed an understanding of none, one, two, or all three of these concepts.

II. Empirical Findings

A. Main Result

I find that participants' beliefs regarding the likelihood that the stock pays from the good dividend distribution are different in the loss relative to the gain condition. Specifically, in the loss condition these subjective beliefs are overly pessimistic and further away from Bayesian objective posteriors, relative to the gain condition, particularly when people are actively investing.

These effects can be seen in Figure 3. The x-axis represents each value of the objective Bayesian posterior belief that can be encountered in the experiment, since there is a finite set of outcome history paths that can be observed by participants, all of which are listed in the Appendix. The y-axis represents the average of the subjective probability estimates produced by participants in the experiment, when observing the outcome histories that yield each of the objective Bayesian posteriors on the x-axis.

If participants were perfect Bayesian learners, their subjective posteriors would line up perfectly, along the 45% line, with the objective Bayesian posteriors. However, Figure 3 indicates that this is not the case, whether we examine beliefs expressed in the Active task (left panel) or in the Passive task (right panel). In either the gain or the loss condition, subjective beliefs deviate from the objective value, and, in accordance with the premise of this paper, these deviations are different across the two conditions.

Specifically, both panels of Figure 3 show that, in the loss condition, subjective posteriors that the stock is paying from the good dividend distribution are lower than in the gain condition. In other words, in the loss condition participants are more pessimistic regarding the likelihood that the stock is the good one.

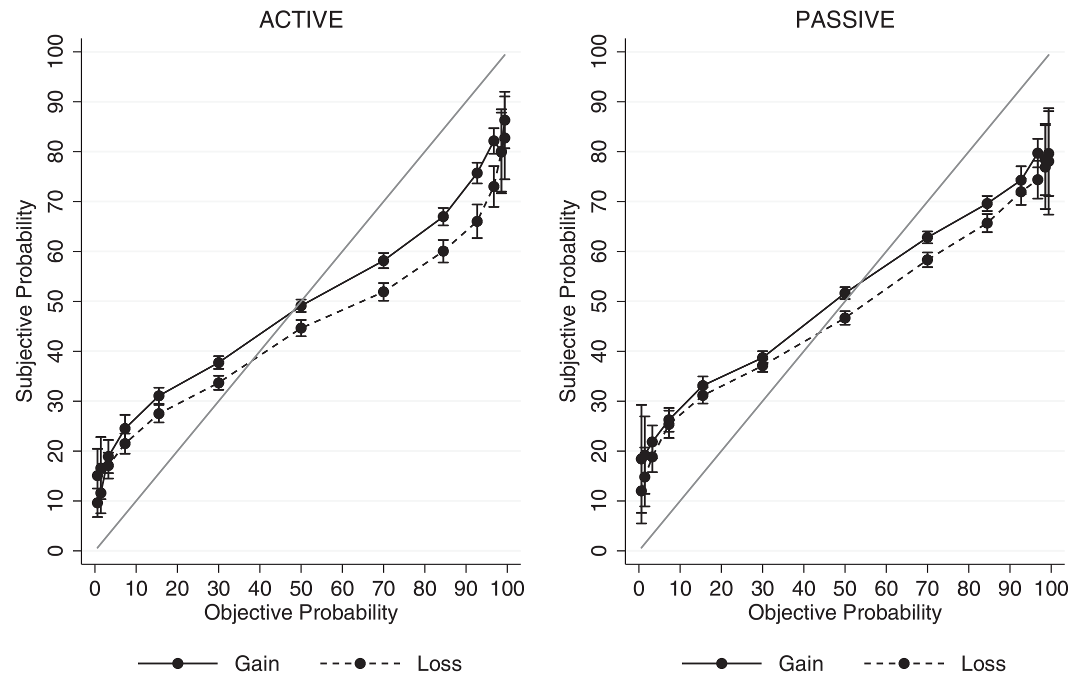

Figure 3. Average subjective estimates for the probability that the stock is paying from the good dividend distribution, as a function of the objective Bayesian probability. The objective Bayesian posteriors that the stock is good that are possible in the experiment are listed in the Appendix, together with the various combinations of high and low outcomes observed during a learning block that lead to such posteriors. If subjective posteriors were Bayesian, they would equal the objective probabilities and thus would line up on the 45° line. Subjective probability estimates provided by participants for each level of the objectively correct Bayesian posterior are shown on dot-dash lines for loss condition trials, and on dot-solid lines for gain condition trials. The left panel presents belief data from the Active task, while the right panel presents belief data from the Passive task.

Table II

Differences in Probability Estimates in the Loss and Gain Conditions

This table shows that subjective posterior beliefs are lower in the loss condition, relative to the gain condition, controlling for objectively correct Bayesian beliefs. The dependent variable in the regression models below, $Probability\ Estimate_{it}$, is the subjective posterior belief that the stock is paying from the good dividend distribution that participant i expressed in trial t. Independent variables include the $Loss\ trial_{it}$ indicator variable, which is equal to one if trial t is in a loss condition block and zero otherwise, as well as $Objective\ Posterior_t$, which is the correct Bayesian probability that the stock is good, given the information seen by the participant up to trial t in the learning block. Subject fixed effects are included in the last three specifications. Standard errors are robust to heteroskedasticity and are clustered by subject. t-statistics are in parentheses. *** indicates significance at the 1% level.

Dependent Variable	$Probability\ Estimate_{it}$ (Subjective Posterior Belief)			
	All Trials	All Trials	Active Trials	Passive Trials
$Loss\ trial_{it}$	–3.94	–3.95	–4.79	–3.13
	(–5.98)***	(–5.96)***	(–5.09)***	(–3.90)***
$Objective\ Posterior_t$	0.56	0.56	0.56	0.56
	(27.37)***	(27.50)***	(27.35)***	(20.92)***
Constant	21.60			
	(18.14)***			
Subject fixed effects	No	Yes	Yes	Yes
R^2	0.547	0.565	0.574	0.575
Observations	10,377	10,377	5,177	5,200

The figure suggests that this pessimism bias, or the wedge between subjective beliefs in the loss and gain domains, is larger in the Active task than in the Passive task, that is, when people are actively involved in making asset choices. These effects are also shown in the regression models in Table II. Controlling for the value of the objective Bayesian probability that the stock is good, I find that beliefs expressed by subjects in trials in the loss condition are on average 3.94% lower (i.e., more pessimistic) than in trials in the gain condition ($p < 0.01$). The condition effect is similar when I estimate the model with or without subject fixed effects. Moreover, the last two columns in Table II show that the difference between subjective beliefs expressed in the loss relative to the gain condition is more pronounced for active trials (-4.79%, $p < 0.01$) than passive trials (-3.13%, $p < 0.01$).

Importantly, the pessimism induced by the loss condition leads to subjects' having larger probability estimation errors, on average, in the loss domain relative to the gain domain. That is, the deviation of subjective beliefs from the objective Bayesian posteriors is higher in the loss condition than in the gain condition. This result is shown in Table III. The errors in the subjective probability estimates, measured relative to the objective Bayesian posteriors that the stock is paying from the good distribution, are 1.86% larger in the loss condition relative to the gain condition ($p < 0.01$). The difference in probability estimation errors between the gain and loss conditions is twice as large in the Active task (2.56%, $p < 0.01$) relative to the Passive task (1.16%, $p < 0.1$).

Table III

Differences in Probability Estimation Errors in the Loss and Gain Conditions

This table shows that probability estimation errors vary across domains. The dependent variable in the regression models in the table, *Absolute Probability Error*$_{it}$, represents the gap between the subjective estimate of the probability that the stock is paying from the good distribution and the objective Bayesian probability. Independent variables include the *Loss trial*$_{it}$ indicator variable, which is equal to one if trial t is in a loss condition block and zero otherwise, as well as *Passive trial*$_{it}$, which is an indicator equal to one if trial t for subject i belongs to the Passive task. Subject fixed effects are included in the bottom panel. Standard errors are robust to heteroskedasticity and are clustered by subject. t-statistics are in parentheses. ***,**,* indicate significance at the 1%, 5% and 10% level, respectively.

Dependent Variable	*Absolute Probability Error*$_{it}$							
	All Trials	Active Trials Only	Passive Trials Only	All Trials	Gain Trials Only	Loss Trials Only	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%
Main Specification								
Loss trial$_{it}$	1.86	2.56	1.16				–1.41	4.31
	(3.88)***	(3.95)***	(1.83)*				(–1.71)*	(6.71)***
Passive trial$_{it}$				–1.06	–0.36	–1.76		
				(–1.72)*	(–0.51)	(–2.23)**		
Constant	13.99	14.17	13.81	15.45	14.17	16.73	15.45	12.93
	(24.56)***	(24.50)***	(18.37)***	(26.24)***	(24.50)***	(22.16)***	(19.57)***	(21.73)***
R^2	0.004	0.007	0.001	0.001	–0.000	0.003	0.002	0.018
Observations	10,377	5,177	5,200	10,377	5,193	5,184	4,439	5,938
Subject fixed effects specification								
Loss trial$_{it}$	1.87	2.59	1.16				–1.46	4.35
	(3.88)***	(4.02)***	(1.81)*				(–2.33)**	(7.07)***
Passive trial$_{it}$				–1.09	–0.38	–1.79		
				(–1.74)*	(–0.53)	(–2.26)**		
R^2	0.129	0.124	0.200	0.127	0.128	0.154	0.185	0.151
Observations	10,377	5,177	5,200	10,377	5,193	5,184	4,439	5,938

In general, also, people make lower estimation errors (by 1.06%, $p < 0.1$) in the Passive task relative to the Active task. This effect is driven by choices in the loss condition, where the errors are lower by 1.76% in Passive versus Active trials ($p < 0.05$). The effect of the loss condition on probability estimation errors is similar if I conduct the estimation without or with subject fixed effects (top and bottom panels of Table III, respectively).

Furthermore, the regression models in Table III confirm a pattern that can also be seen immediately in Figure 3, namely, that the biggest deviation of subjective posteriors from the Bayesian ones happens in the loss condition when the objective probability that the stock pays from the good dividend distribution is high. In those situations, participants' subjective posterior probabilities are the most pessimistic relative to the objective values. Comparing the first and last columns in Table III illustrates this effect. Specifically, while on average the absolute value of probability estimation errors is 1.86% higher in the loss condition trials relative to other trials, this difference increases to 4.31% for loss trials with high values ($\geq 50\%$) of the objective posterior probability that the stock pays dividends from the better distribution. These are trials in the loss condition where subjects have seen more $-\$2$ dividends than $-\$10$ dividends, so the stock is more likely to be paying from the good distribution and hence the situation faced by the subject is likely not the worst possible. It is in these situations that people's beliefs are particularly far, in a pessimistic way, from Bayesian beliefs.

Also, the fact that in the loss condition subjective beliefs are lower than in the gain condition helps bring the subjective probability estimates marginally closer (by 1.41%, $p < 0.1$, see Table III) to the Bayesian ones for low values of the objective probabilities, where, generally speaking, participants' estimates are too high. As seen in Figure 3, across the Active or Passive task and gain or loss trials, subjects update their priors in such a way that the expressed posterior probability that the stock is paying from the good distribution is significantly higher (by 12% on average) than the objective Bayesian posterior for low values of this objective probability, and significantly lower (by 13% on average) than the objective Bayesian posterior for high values of this objective probability, a result that replicates the experimental patterns documented in Kuhnen and Knutson (2011) and is suggestive of conservatism, or regression to the mean, in probability estimation (e.g., Phillips and Edwards (1966)). This relationship between subjective and objective posterior beliefs resembles the relationship between decision weights and objective probabilities postulated by Prospect Theory, but in my experiment reflects errors in updating priors, and is different from the idea that decision makers overweight rare events and underweight frequent ones (Kahneman and Tversky (1979), Prelec (1998)).

So far, thus, the evidence shows that on average subjective beliefs are more pessimistic, and further from Bayesian beliefs, in the loss condition relative to the gain condition, particularly when people actively choose investments. But is this effect of the loss condition robust? In-sample tests indicate that this is the case. Specifically, Table IV shows that the effect of the loss versus gain condition on probability estimation errors is present throughout the experiment, whether I examine trials that come early or late in a learning block, or blocks that come

Table IV

In-Sample Robustness Checks

This table shows that the loss versus gain condition effect on probability estimation errors is similar for early and late trials within each learning block, and for early and late learning blocks in the experiment (Panel A), and that for each of these subsamples it is particularly pronounced in the Active task (Panel B). t-statistics are in parentheses. ***,** indicate significance at the 1% and 5% level, respectively.

Panel A: All Trials

Dependent Variable	*Absolute Probability Error$_{it}$*				
	All Trials	First Three Trials of Each Learning Block	Last Three Trials of Each Learning Block	Trials in First Half of Active and Passive Tasks	Trials in Second Half of Active and Passive Tasks
Loss trial$_{it}$	1.86	1.89	1.83	1.74	1.99
	(3.88)***	(3.45)***	(3.07)***	(3.08)***	(2.93)***
Constant	13.99	13.07	14.91	13.76	14.21
	(24.56)***	(23.64)***	(20.63)***	(23.42)***	(21.22)***
Adj. R^2	0.004	0.005	0.003	0.003	0.004
Observations	10,377	5,192	5,185	5,183	5,194

Panel B: Active Trials Only

Dependent Variable	*Absolute Probability Error$_{it}$*				
	All Trials in Active Task	First Three Trials of Each Learning Block of Active Task	Last Three Trials of Each Learning Block of Active Task	Trials in First Half of Active Task	Trials in Second Half of Active Task
Loss trial$_{it}$	2.56	2.47	2.65	2.69	2.44
	(3.95)***	(3.41)***	(3.31)***	(3.31)***	(2.58)**
Constant	14.17	13.72	14.63	14.00	14.34
	(24.50)***	(22.11)***	(19.78)***	(21.05)***	(19.64)***
R^2	0.007	0.007	0.006	0.008	0.006
Observations	5,177	2,593	2,584	2,583	2,594

either early or late in the experiment (Panel A). Moreover, the loss condition effect across these different subsamples is particularly large for active trials (Panel B). Another question is whether participants may be aware of the effect of the experimental conditions on their beliefs. The answer is no. At the end of each trial, participants were asked to provide a confidence number (from one to nine, with one meaning not confident at all and nine meaning very confident) to indicate how much they trust the subjective probability estimate produced in that trial. I find no significant differences between the average confidence of participants in active versus passive trials (5.31 vs. 5.39, respectively), or during loss versus gain trials (5.27 vs. 5.43, respectively).

An important question is how the participants' probability estimates evolve such that they end up overly pessimistic during loss trials relative to gain trials, and more so during Active investment, as shown in Figure 3 and Table III. To answer this question, for each participant i and trial t, I calculate the change from trial t to trial $t + 1$ in his estimate of the probability that the stock is the good one, that is, the difference between the subjective posterior and prior belief that the stock pays from the good dividend distribution. I then test whether, on average, across all participants and trials, probability updating differs across contexts. The results are presented in Table V and show that participants put significantly more weight on low outcomes during loss trials than during gain trials, particularly during Active investment and when they have expressed high priors that the stock is good. Specifically, the table shows that observing a low dividend reduces participants' probability estimate that the stock is the good one by 0.69% ($p < 0.05$) more in the loss condition relative to the gain condition. The difference between the loss and the gain condition with respect to the effect of a low dividend on probability updating becomes 1.92% ($p < 0.01$) if the participants' priors were above 50%, and 3.92% ($p < 0.01$) during Active investment trials. In other words, the data suggest that the reason people are overly pessimistic during periods of negative outcomes when their own money is at stake is that during such situations, relative to other contexts, their beliefs are more strongly influenced by low dividends.

To summarize, the main result of the paper is that, when people consider available investment options, they form beliefs about these investments that are overly pessimistic, overly sensitive to low outcomes, and further away from the objective Bayesian beliefs in the loss domain relative to the gain domain, that is, when learning from negative outcomes as opposed to learning from positive outcomes.

B. Replication Study

So far, the evidence shows that the main result is robust in-sample. However, it is critical to show that it also replicates out-of-sample. To see whether this is indeed the case, I ran the same experiment at Babes-Bolyai University in Romania, a top institution in that country, in a sample of 203 participants (53 males, 150 females, mean age 21 years, 2.03 years standard deviation). The results show the same context-induced effects on learning as documented in the original sample from Northwestern University.

Table V

Differences in Probability Updating in the Loss and Gain Conditions

This table shows the average change from trial to trial in participants' estimate of the probability that the stock is good (i.e., posterior minus prior belief) for loss and gain trials, for high versus low dividend realizations, for high versus low subjective priors, and for Active versus Passive trials. Differences in belief updating between the loss and gain conditions that are significant at $p < 0.05$ or $p < 0.01$ are indicated by ** and ***, respectively.

	Probability Estimate$_{t+1}$ − *Probability Estimate*$_t$					
	High Dividend in Trial $t+1$	Low Dividend in Trial $t+1$	Low Dividend in Trial $t+1$, *Probability Estimate*$_t$ < 50%	Low Dividend in Trial $t+1$, *Probability Estimate*$_t$ ≥ 50%	Low Dividend in Trial $t+1$, *Probability Estimate*$_t$ ≥ 50% Passive Trials	Low Dividend in Trial $t+1$, *Probability Estimate*$_t$ ≥ 50% Active Trials
Loss condition	11.17%	−12.50%	−7.51%	−16.33%	−14.73%	−17.99%
Gain condition	11.38%	−11.81%	−7.44%	−14.41%	−14.74%	−14.07%
Loss − Gain	−0.21%	−0.69%**	−0.07%	−1.92%***	−0.01%	−3.92%***

Specifically, as can be seen in Table VI, probability estimation errors are 2.25% larger ($p < 0.01$) in the loss versus the gain condition. The loss versus gain condition difference is larger for active trials (2.89%, $p < 0.01$) than for passive trials (1.62%, $p < 0.01$), and is largest (7.81%) for trials where the objective probability that the stock is good is greater than 50%. For trials with lower objective posteriors, being in the loss condition lowers the estimation errors. All these effects are robust to estimation with and without subject fixed effects, and are very similar to those observed in the Northwestern sample (see Table III).

Moreover, as shown in Table VII, the difference in posteriors between the loss and the gain conditions is driven by a difference between the two conditions in terms of how people update their beliefs after seeing low outcomes, as found in the Northwestern University sample (see Table V). Specifically, in the Romanian sample also there is no difference between the gain and loss conditions in how people update their beliefs about the stock after seeing a high outcome. However, after seeing a low outcome, participants lower their estimate that the stock is paying from the good distribution 1.12% more ($p < 0.05$) in the loss condition relative to the gain condition. This increased reactivity to low outcomes in the loss condition is particularly pronounced (3.55%, $p < 0.01$) for trials where the stock is more likely to be good, for either active (3.27%, $p < 0.01$) or passive trials (3.87%, $p < 0.01$).

It is interesting to note that in the Romanian replication sample average learning errors are larger than in the U.S. sample. Comparing the results in the first column of Tables III and VI, aside from any effects of the loss versus gain condition manipulation, Romanian subjects make estimation errors that are 13.13% larger than those of U.S. participants (27.12% vs. 13.99% average error, respectively, for the two samples). This difference may come from differences across the two countries in people's comfort with financial investment decisions or in exposure to financial concepts during young adulthood, and it is a topic worth investigation in future work. That being said, while Romanian participants do not learn as well as their U.S. counterparts in this task, they do exhibit the same domain-induced difference in beliefs (i.e., larger estimation errors and overreaction to low outcomes in the loss domain), which is the effect of interest for this replication exercise.

The results from the Romanian replication sample therefore provide reassurance that the effect of the loss versus gain context manipulation on subjective beliefs documented in the U.S. sample is capturing a real phenomenon that is not confined to the original sample.

C. Alternative Explanations

While the replication study speaks to the out-of-sample validity of the main result, it is also important to turn toward thinking about and testing alternative explanations for the documented effect of the loss versus gain context on the beliefs expressed by participants. I discuss and test these alternatives below.

Table VI

Estimation Errors in the Replication Sample (Romanian Participants)

The analysis in this table mirrors that done for the main subject sample (U.S. participants) in Table III. The table shows the effect of the loss versus gain condition on probability estimation errors across different subsamples of trials, without and with subject fixed effects. t-statistics are in parentheses. *** and ** indicate significance at the 1% and 5% level, respectively.

Dependent Variable	*Absolute Probability Error$_{it}$*							
	All Trials	Active Trials Only	Passive Trials Only	All Trials	Gain Trials Only	Loss Trials Only	Trials with Objective Posteriors $< 50\%$	Trials with Objective Posteriors $\geq 50\%$
				Main Specification				
Loss trial$_{it}$	2.25	2.89	1.62				–5.21	7.81
	(4.33)***	(4.07)***	(2.57)**				(–6.18)***	(10.40)***
Passive trial$_{it}$				–1.33	–0.69	–1.96		
				(–2.20)**	(–0.99)	(–2.54)**		
Constant	27.12	27.47	26.78	28.91	27.47	30.36	31.96	23.48
	(54.24)***	(46.44)***	(42.57)***	(54.98)***	(46.44)***	(44.98)***	(36.22)***	(50.63)***
R^2	0.002	0.004	0.001	0.001	0.000	0.002	0.012	0.032
Observations	23,804	11,804	12,000	23,804	11,926	11,878	10,135	13,669
				Subject Fixed Effects Specification				
Loss trial$_{it}$	2.25	2.92	1.62				–4.70	7.58
	(4.31)***	(4.07)***	(2.54)**				(–5.81)***	(10.47)***
Passive trial$_{it}$				–1.23	–0.55	–1.92		
				(–2.02)**	(–0.77)	(–2.45)**		
R^2	0.070	0.093	0.101	0.068	0.088	0.086	0.155	0.123
Observations	23,804	11,804	12,000	23,804	11,926	11,878	10,135	13,669

Table VII

Updating in the Replication Sample (Romanian Participants)

The analysis in this table mirrors that done for the main subject sample (U.S. participants) in Table V. The table shows how subjects update their probability estimates after seeing high or low stock outcomes in the gain or the loss condition. Differences in belief updating between the loss and gain conditions that are significant at $p < 0.05$ or $p < 0.01$ are indicated by ** and ***, respectively.

	$Probability\ Estimate_{t+1} - Probability\ Estimate_t$					
	High Dividend in Trial $t+1$	Low Dividend in Trial $t+1$	Low Dividend in Trial $t+1$, *Probability* $Estimate_t < 50\%$	Low Dividend in Trial $t+1$, *Probability* $Estimate_t \geq 50\%$	Low Dividend in Trial $t+1$, *Probability* $Estimate_t \geq 50\%$ Passive Trials	Low Dividend in Trial $t+1$, *Probability* $Estimate_t \geq 50\%$ Active Trials
Loss condition	19.28%	−21.96%	−7.02%	−30.14%	−31.30%	−29.00%
Gain condition	19.06%	−20.83%	−5.60%	−26.59%	−27.43%	−25.73%
Loss-Gain	0.22%	−1.12%**	−1.42%***	−3.55%***	−3.87%***	−3.27%***

C.1. Are Priors Different in Gain versus Loss Blocks?

An important concern is that perhaps subjects start with different priors that the stock is paying from the good distribution, in the loss versus the gain domain (e.g., Bossaerts (2004)). However, the experimental instructions clearly told subjects that, in the beginning of each block of six trials, the computer randomly decided whether the stock will pay from the good or the bad dividend distribution in that block, and hence the prior probability that the stock was the good one was 50%. Careful explanations and training were provided to ensure that subjects were aware that this prior probability was 50%, no matter whether each new block was in the gain condition, or in the loss condition.

Nonetheless, perhaps subjects somehow started with a bias in priors. Specifically, when faced with a new block in the loss condition, they may have started with a prior that was below 50% that the stock would pay dividends from the good distribution. For example, in the gain condition their prior may be the correct one, that is, 50%, but in the loss condition it may mistakenly be only 47%.

If the pessimism effect that I document in their posterior beliefs (namely, posteriors in the loss condition are about 3% lower than in the gain condition, for the same information set) simply comes from having these two different priors that the stock is good in the loss condition and in the gain condition, then the wedge between posteriors in the loss and the gain domains after seeing an equal number of high and low dividends should always be equal to that initial difference between their prior in the gain condition and their prior in the loss condition. In other words, if updating was done in a Bayesian fashion in both conditions, but the priors were 50% and 47%, in the gain and loss domains, respectively, then we would have that the posterior belief that the stock is the good one, after seeing an equal number of trials t with high outcomes and with low outcomes, would be 50% in the gain condition, and 47% in the loss condition, for any t.[2] However, I find that the wedge in posteriors is not constant over time, but in fact increases with the number of observed outcomes; namely, such as it is bigger the more low outcomes the subject observes.

Table VIII shows this result. There, posterior beliefs after seeing an equal number of high and low outcomes are not statistically different from 50% (i.e., the correct value) when people face the gain condition. In the loss condition, however, posteriors in such situations become more and more different from 50%, as people observe more low outcomes. Specifically, in the loss condition, after seeing one low outcome and one high outcome in the first two trials in a block, people's average estimate that the probability that the stock is good is 46.81%, which is 2.90% ($p < 0.01$) below the average estimate offered in the gain condition. After seeing two high and two low outcomes in the first four trials in a block, the average estimate of this probability drops to 45.36% in the loss condition, and is 5.22% ($p < 0.01$) lower than the estimate produced after

[2] A quick illustration of this can be found in an Excel spreadsheet available on the author's web page at http://public.kenan-flagler.unc.edu/faculty/kuhnenc/RESEARCH/posteriors_calculation.xls.

Table VIII

Test of Different Priors in the Loss and Gain Conditions

This table shows the expressed posteriors, and the difference between those elicited in the loss and gain domains, in situations where one low and one high outcome were observed in two trials, or two low and two high outcomes were observed in four trials, or three low and three high outcomes were observed in six trials. In all these situations the correct posterior that the stock is good is 50%. If people start with more pessimistic priors in the loss domain, relative to the gain domain, such different priors would imply a constant wedge between the posteriors expressed in the loss and gain domains, in situations where an equal number of high and low outcomes have been observed. A chi-square test rejects the null hypothesis that the loss versus gain condition differences found in the three subsamples are equal ($p < 0.05$).

	Average Subjective Posterior Estimate of the Probability that the Stock Is Good		
	Gain Condition	Loss Condition	Loss Versus Gain Difference
After 2 trials, with 1 high and 1 low outcome	49.71% (not significantly different from 50%)	46.81% (significantly different from 50%, $p < 0.01$)	−2.90%*** (significantly different from 0, $p < 0.01$)
After 4 trials, with 2 high and 2 low outcomes	50.58% (not significantly different from 50%)	45.36% (significantly different from 50%, $p < 0.01$)	−5.22%*** (significantly different from 0, $p < 0.01$)
After 6 trials, with 3 high and 3 low outcomes	51.83% (not significantly different from 50%)	43.83% (significantly different from 50%, $p < 0.01$)	−8.00%*** (significantly different from 0, $p < 0.01$)

similar trial sequences in the gain condition. Finally, after seing three high and three low outcomes during the six-trial block, people's average estimate that the stock is paying from the good distribution is 43.83% in the loss condition, and is 8.00% ($p < 0.01$) lower than the average estimate for an equivalent set of trials in the gain condition. A chi-square test rejects the null hypothesis that the wedge between posteriors expressed in the loss and the gain conditions is the same across these three subsamples of the data ($p < 0.05$).

In other words, posterior beliefs in the loss condition diverge from the gain condition posteriors in a manner inconsistent with simply having a more pessimistic prior in loss blocks, but performing similar Bayesian learning in both conditions. Hence, different priors cannot explain the pessimism bias induced by the loss condition.

C.2. Are Risk Preferences Different in Gain versus Loss Blocks?

Another potential concern is that the expressed beliefs of the participants reflect their risk preferences rather than actual beliefs, and that preferences may depend on the experimental condition. The small stakes in each trial should lead participants to act in a risk-neutral manner in each trial of the experiment (Rabin (2000)), but perhaps this is not the case.

Table IX

Test of Different Risk Preferences in the Loss and Gain Conditions

This table shows results of a linear probability model where the dependent variable is an indicator equal to one for trials where subjects chose the stock rather than the bond. The independent variable is the indicator *Loss trial*$_{it}$, which is equal to one if trial t faced by subject i belongs to a Loss condition block, and is equal to zero if the trial belongs to a Gain condition block. The sample is limited to the first trial in each of the 10 Active investment blocks that each subject faced. These are trials where asset choices (stock vs. bond) are made in the absence of any news regarding the stock's dividends, and thus reflect subjects' risk preferences without learning confounds. Standard errors are robust to heteroskedasticity and are clustered by subject. t-statistics are in parentheses. *** indicates significance at the 1% level.

Dependent Variable	Indicator Equal to 1 if Stock Chosen by Subject i in Trial t	
Loss trial$_{it}$	–0.05	–0.05
	(–1.19)	(–1.13)
Constant	0.62	
	(15.94)***	
Subject fixed effects	No	Yes
R^2	0.001	0.340
Observations	861	861

As a first test regarding this concern, I analyze whether risk preferences correlate with people's subjective posterior beliefs that the stock is good. As a measure of risk preferences that is not confounded by learning issues, I use the participants' answers to the postexperiment question regarding how they would like to divide $10,000 in their portfolio between a stock index fund and a savings account. The amount invested in the risky asset is my proxy for their risk tolerance. I find that the correlation between the participants' risk tolerance and their beliefs that the stock is good is −0.01, and is not significantly different from 0. This is true across all trials, as well as separately for gain, loss, active, or passive trials. Hence, I do not find that more risk-tolerant participants report higher estimates for the probability that the stock is paying from the good dividend distribution. Thus, the subjective beliefs that participants produce are orthogonal to their risk preferences.

While this is reassuring, another related concern needs to be addressed. Namely, the observed increase in pessimism in reported beliefs in the loss condition, relative to the gain condition, may solely be an indication of differential risk aversion in the two domains, rather than an indication that belief formation may be different across the two domains. To see whether this is the case, I examine the choices made by subjects in the first trial of each new learning block of six trials, when they have not yet learned any news about the dividends paid by the stock, to see if they preferentially choose the stock, rather than the bond, in one domain relative to the other. Such a behavior would indicate that risk preferences may be different in gain and loss blocks. The evidence in Table IX shows that there is no significant difference induced by the loss versus gain condition on people's investment choice in the first trial of each block, before any learning occurs. Specifically, in the table, I estimate a linear probability model (a logit or probit model would yield the same result) of the decision to select the stock, instead of the bond, on the first trial of each block of six, as a function of whether that block is in the loss condition or in the gain condition. I estimate this model without and with subject fixed effects. Either way, the estimated effect of the $\textit{Loss trial}_{it}$ indicator variable (i.e., 0.05) is not significantly different from zero. Subjects choose the stock in 62% of the first trials in each block, and this preference for the risky investment does not vary significantly between the gain and the loss conditions. Hence, the fact that subjects have different posterior beliefs in the loss versus the gain domain is unlikely to be caused by them having different risk preferences in these two domains.

C.3. Are Subjective Beliefs Irrelevant Quantities to Study?

Yet a different concern is whether the subjective beliefs elicited in the task are meaningful quantities to study. Perhaps they have no relation to how people in fact choose assets in the experiment. I show that this is not the case. People act based on these subjective beliefs. They are significantly more likely to choose the stock if they believe that the probability of it paying from the good distribution is higher.

Table X presents the results of a linear probability model (a logit or probit model would yield the same results) of the decision to select the stock instead of the bond in Active task trials. Since the goal is to test whether the subjects' expressed beliefs about the likelihood that the stock is good in fact influence their choices, the sample used here includes choices made in trials 2 through 6 of each of the 10 blocks that each person faced during the Active investment task. This is because, as discussed earlier, the first subjective belief elicited in each block is obtained at the end of trial 1 in that block, after that first choice has been made. Hence, trial 1 choices cannot be used in this analysis.

Column (1) of the table shows a strong and significant effect of the subjective belief expressed at the end of trial $t-1$ (captured by the variable $\textit{ProbabilityEstimate}_{i,t-1}$) on the subject's asset choice in trial t. Specifically, a 1% increase in the subjective belief that the stock is good will increase the chance that subject i selects the stock rather than the bond in the subsequent trial by 1% ($p < 0.01$). In the same model, I control for whether the trial is in the loss or the gain condition, as indicated by the dummy variable $\textit{Loss trial}_{it}$, and find no condition effect on the type of asset chosen. The model in column (1) as well as those in the rest of the table include subject fixed effects.

In column (2) I estimate the role of being in the loss versus the gain condition for the asset choice in trials 2 through 6 of each block, without controlling for the subjects' beliefs, and find that in the loss condition participants are 6% less likely to choose the stock ($p < 0.05$). The results in columns (1) and (2) therefore indicate that, while people are more reluctant to choose the risky asset in trials 2 through 6 in the loss condition, relative to the gain condition, this effect is completely driven by the difference in the subjective beliefs these individuals have, that is, by the fact that the variable $\textit{Probability Estimate}_{i,t-1}$ is different for gain and loss trials, as discussed earlier in the paper.

Table X

Test of Whether Subjective Beliefs Drive Investment Choices

This table presents the results of a linear probability model of the decision to select the stock instead of the bond in Active task trials. The sample includes asset choices from trials 2 through 6 of each block (since subjects' beliefs are elicited starting at the end of trial 1) in the Active task. The effects of expressed subjective beliefs on choices are similar across the gain versus loss condition, early or late within a block of six trials, or in earlier or later blocks. *Probability Estimate*$_{i,t-1}$ is expressed in percentage points. Standard errors are robust to heteroskedasticity and are clustered by subject. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5% and 10% level, respectively.

Dependent Variable	Indicator Equal to 1 if Stock Chosen by Subject i in Trial t				
	(1)	(2)	(3)	(4)	(5)
Probability Estimate$_{i,t-1}$	0.01		0.01	0.01	0.01
	(11.80)***		(9.18)***	(12.08)***	(7.95)***
Loss trial$_{it}$	–0.01	–0.06	0.01		
	(–0.41)	(–2.08)**	(0.23)		
Probability Estimate$_{i,t-1}$ × *Loss trial*$_{it}$			–0.00		
			(–0.50)		
First half of block				0.08	
				(2.11)**	
Probability Estimate$_{i,t-1}$ × *First half of block*				–0.00	
				(–1.95)*	
First half of active task					–0.04
					(–1.06)
Probability Estimate$_{i,t-1}$ × *First half of active task*					0.00
					(1.59)
Subject fixed effects	Yes	Yes	Yes	Yes	Yes
R^2	0.171	0.003	0.223	0.224	0.224
Observations	4,279	4,312	4,279	4,279	4,279

In the remaining three columns of Table X, I test whether the relationship between the participants' expressed beliefs and their asset choice may vary in strength, depending on whether they are in loss blocks versus gain blocks (column (3)), whether they face earlier or later trials within a learning block (column (4)), or whether they are in earlier or later learning blocks in the Active task (column (5)). The results indicate that the influence of beliefs on choices, across all of these settings, is as before: each 1% increase in the subjective belief that the stock is good leads to a 1% increase in the probability that the subject chooses the stock on the next trial ($p < 0.01$). Moreover, similar to the pattern documented earlier that, once controlling for beliefs, the loss versus gain condition do not induce differences in asset choice, the results in column (5) show that there is no difference in asset choices in the first half (i.e., the first five learning blocks) in the Active task, versus the second half (i.e., the last five learning blocks). In other words, simply moving on to later learning blocks in the experiment does not induce a change in the preference for the stock versus the bond. In column (4), there is a significant effect of the dummy variable *First half of block*, which indicates that in trials 2 and 3, relative to trials 4, 5, and 6 in a block, participants are about 8% more likely to choose the stock ($p < 0.05$). In other words, at the beginning of each learning block, there is a slightly higher tendency for people to choose the stock, relative to the end of the learning block. This is an effect orthogonal to that of the variable of interest, namely the effect of subjective beliefs that the stock is good on the actual asset choices made by participants.

Overall, the results in Table X show that subjective beliefs indeed drive asset choices and that the influence of beliefs on choices is similar across all active trials, irrespective of whether they are in the loss or gain condition, early or late in a learning block, or in early or late blocks during the Active task. Furthermore, the fact that participants are less likely to select the risky asset in the loss condition relative to the gain condition is entirely driven by the difference in beliefs that they display across these two conditions. Put differently, the reluctance to pick the stock in the loss condition simply comes from people having more pessimistic beliefs in that condition about how likely it is that the stock is paying from the good dividend distribution.

C.4. Is the Task Truly Testing Learning from Financial News?

Another concern is whether the experimental task used here is truly about learning from financial news. To test this, I examine whether subjects' errors in the expressed probability that the stock is good (i.e., their ability to learn from news in the task) are related to measures of the subjects' learning capacity outside of my experiment. I have two such measures, one innate and one acquired personal characteristic.

The first measure of learning capacity is an indicator variable equal to one if the person happens to have a genetic variant that has been previously related to working memory capacity, that is, the ability to store and use items in memory in the short run. The gene in question is referred to as *COMT* (catechol-*O*-methyltransferase). People carrying the *Met/Met* allele combination have been shown to have better working memory than those carrying the other two allele combinations, namely *Val/Val* or *Val/Met* (Frank et al. (2007), Doll, Hutchison, and Frank (2011)). In the sample, 19 of the 87 subjects are *COMT Met/Met* carriers and hence have a genetic predisposition toward better working memory function.[3]

The second measure of learning capacity is an indicator equal to one if the person answered correctly the portfolio expected value question (administered

[3] Genotyping was done by ACGT Inc. (Wheeling, IL, USA), a commercial provider of DNA analysis services, according to standard procedures described elsewhere (Frank et al. (2007)). The resulting distribution of $COMT$ genotypes of the 87 participants comprised 19 Met/Met, 34 Val/Met, and 34 Val/Val participants and was consistent with that expected under Hardy Weinberg equilibrium ($\chi^2 = 3.29$, $df = 1$, $p > 0.05$). The sample size here is similar to those of other studies targeting the $COMT$ gene, and the incidence of the Met/Met genotype (22%) is also in line with prior work (28% of 68 participants in Frank et al. (2007), 18% of 74 participants in Doll, Hutchison, and Frank (2011)). Hence, the participant group used in this study is representative and large enough to identify the effect of the $COMT$ gene on financial decision making.

Table XI

Test of Whether the Task Tests Learning from Financial News

In this table, the dependent variable is the size of the gap between the subjective probability estimate and the objective Bayesian posterior for each trial t faced by subject i. The independent variables are measures of the subjects' learning capacity, either innate or acquired. $COMT$ Met/Met_i is an indicator equal to one for participants who have the *Met/Met* variant of the $COMT$ gene, which has previously been shown to correspond to better working memory capacity. *High financial knowledge*$_i$ is an indicator equal to one for subjects who correctly answered the question concerning the calculation of the expected value of a portfolio, which was administered immediately after the experimental task. Standard errors are robust to heteroskedasticity and are clustered by subject. t-statistics are in parentheses. ***,** and * indicate significance at the 1%, 5% and 10% level, respectively.

Dependent Variable	*Absolute Probability Error*$_{it}$ (1)	(2)	(3)
COMT Met/Met$_i$	–2.50		–2.09
	(–2.29)**		(–1.80)*
High financial knowledge$_i$		–3.00	–2.79
		(–2.61)**	(–2.46)**
Constant	15.46	16.47	16.82
	(22.07)***	(16.89)***	(17.05)***
R^2	0.005	0.010	0.013
Observations	10,377	10,377	10,377

immediately after the experimental task) described briefly in Section I and in more detail in the Appendix. For 45 of the 87 subjects, their answer to that question is correct, and hence they are classified as "High financial knowledge" subjects. It is natural to expect that individuals more familiar with concepts relating to stocks, bonds, returns, or probabilities will be better able to learn from stock outcomes in the experimental task.

Therefore, if the experimental task that I use indeed engages people's learning systems, instead of, for example, recruiting automatic or random verbal and behavioral responses, then it should be the case that the people whose beliefs are closer to Bayesian beliefs (i.e., those who seem to learn better in the task) are more likely to be those individuals with the genetic predisposition toward better working memory, as well as those who have acquired more familiarity with financial concepts.

This is indeed what the results in Table XI indicate. The regression model in column (1) shows that individuals who do not have the *COMT Met/Met* genotype, which is advantageous for working memory make estimation errors averaging 15.46%, whereas their peers who have the advantageous genetic combination make estimation errors that are 2.50% lower ($p < 0.05$). The results in column (2) indicate that individuals who have not so far acquired high financial knowledge make estimation errors averaging 16.47%, but those with high financial knowledge make errors that are 3.00% smaller ($p < 0.05$). This finding is consistent with the result in Lusardi and Mitchell (2007) that more financial literacy leads to better financial outcomes. Finally, in the regression model in column (3), I use as predictors of learning in the task both the *COMT*

Met/Met$_i$ and *High financial knowledge*$_i$ indicators simultaneously and find that their effects on learning errors in the task are similar to those estimated in the univariate models in columns (1) and (2). In other words, better learning in the experimental task is exhibited by people with a stronger innate or acquired propensity to learn from financial news.

Overall, therefore, the evidence shows that the asymmetry in subjective beliefs between the loss and the gain domain is a robust phenomenon both in-sample and across samples, that it occurs at the time of updating from low outcomes, that it is not driven by asymmetric priors or different risk preferences in the two domains, that these subjective beliefs actually drive investment choices, and that the experimental task engages participants' learning systems.

III. Discussion and Implications

The presence of a pessimism bias in times of scarcity, or negative news, may have significant effects on economic decisions outside of the laboratory.

First, this pessimism bias may lead to underinvestment in risky assets during bad economic times either by households or by corporate decision makers, consistent with the finding in Malmendier and Nagel (2011) that individuals who have lived through bad stock market times do not invest as much as others. This belief-induced reluctance to take risks during bad times is complementary to the idea that preferences—namely, risk aversion—may be countercyclical (e.g., Routledge and Zin (2010)).

Second, this pessimism bias may lead to underinvestment in human capital, which may keep people in scarcity or in a poor economic situation. This is consistent with the recent suggestion made by Banerjee and Duflo (2011) and Mullainathan and Shafir (2013) that a mind-set of poverty can lead to a feedback loop, in that it reduces people's interest in investing in education or engaging in productive endeavors. This aspiration gap between the poor and the rich may be in part driven by a pessimism bias that the poor experience due to their prolonged exposure to negative outcomes.

Future research is needed to further examine the hypothesis that a pessimism bias induced by bad economic environments leads to underinvestment at an aggregate level. Suggestive evidence for this hypothesis is provided by Hoberg and Phillips (2010) and Ben-David, Graham, and Harvey (2013).

Moreover, it will be important to conduct empirical studies to find whether the same overreaction to low outcomes in the negative domain that is observed here in an experimental context also applies in real life situations where professionals learn about the economic environment. For example, do stock analysts, macroeconomic forecasters, or mutual fund managers exhibit this pessimism bias in learning from corporate earnings announcements or from macroeconomic news?

Another fruitful direction is to understand how the pessimism bias can be undone. For example, could interventions be designed to increase the salience of high or better-than-expected outcomes in negative times? Or perhaps there are personal characteristics that can be cultivated that may mitigate this context-induced bias in learning.

Last but not least, the idea that learning may be asymmetric during booms and recessions, in the sense that people react more to low outcomes in bad relative to good times, may be included in asset pricing models to yield novel predictions about price discovery and evolution as a function of the economic conditions faced by traders.

IV. Conclusion

This paper documents the existence of asymmetries in learning from financial news. I find that, when people consider available investment options, they form beliefs about these investments that are overly pessimistic, overly sensitive to low outcomes, and further away from the objective Bayesian beliefs in the loss domain relative to the gain domain, that is, when learning from negative outcomes relative to learning from positive outcomes.

The evidence shows that the asymmetry in subjective beliefs between the loss and the gain domains is a robust phenomenon in-sample and across samples in two countries, the United States and Romania. It occurs at the time of updating from low outcomes and is not driven by asymmetric priors or different risk preferences in the two domains. The subjective beliefs actually drive investment choices, and the experimental task engages participants' learning systems.

The pessimism in beliefs induced by the loss domain that I show here can help shed light on differences between poor and good economic times in the investment behavior of economic agents, from households to firms, and provides an explanation for the underinvestment in productive activities or in human capital by those who have experienced chronic poverty or bad economic environments. Thus, this result has broad implications across household finance, corporate finance, and development economics.

Initial submission: October 23, 2012; Final version received: July 25, 2014
Editor: Bruno Biais

Appendix:

A. Participant Instructions

Welcome to our financial decision making study!

In this study you will work on two investment tasks. In one task you will repeatedly invest in one of two securities: a risky security (i.e., a stock with risky payoffs) and a riskless security (i.e., a bond with a known payoff), and will provide estimates as to how good an investment the risky security is. In the other task you are only asked to provide estimates as to how good an investment the risky security is, after observing its payoffs.

In either task, there are two types of conditions you can face: the GAIN and the LOSS conditions. In the GAIN condition, the two securities will only provide POSITIVE payoffs. In the LOSS condition, the two securities will only provide NEGATIVE payoffs.

Details for the Investment Choice and Investment Evaluation Task:
Specific details for the GAIN condition:

In the GAIN condition, on any trial, if you choose to invest in the bond, you get a payoff of \$6 for sure at the end of the trial. If you choose to invest in the stock, you will receive a dividend which can be either \$10 or \$2.

The stock can either be good or bad, and this will determine the likelihood of its dividend being high or low. If the stock is good then the probability of receiving the \$10 dividend is 70% and the probability of receiving the \$2 dividend is 30%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is good, then on each trial the odds of the dividend being \$10 are 70%, and the odds of it being \$2 are 30%. If the stock is bad then the probability of receiving the \$10 dividend is 30% and the probability of receiving the \$2 dividend is 70%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is bad, then on each trial the odds of the dividend being \$10 are 30%, and the odds of it being \$2 are 70%.

Specific details for the LOSS condition:

In the LOSS condition, on any trial, if you choose to invest in the bond, you get a payoff of −\$6 for sure at the end of the trial. If you choose to invest in the stock, you will receive a dividend which can be either −\$10 or −\$2.

The stock can either be good or bad, and this will determine the likelihood of its dividend being high or low. If the stock is good then the probability of receiving the −\$10 dividend is 30% and the probability of receiving the −\$2 dividend is 70%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is good, then on each trial the odds of the dividend being −\$10 are 30%, and the odds of it being −\$2 are 70%. If the stock is bad then the probability of receiving the −\$10 dividend is 70% and the probability of receiving the −\$2 dividend is 30%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is bad, then on each trial the odds of the dividend being −\$10 are 70%, and the odds of it being −\$2 are 30%.

In both GAIN and LOSS conditions:

In each condition, at the beginning of each block of six trials, you do not know which type of stock the computer selected for that block. You may be facing the good stock, or the bad stock, with equal probability.

On each trial in the block you will decide whether you want to invest in the stock for that trial and accumulate the dividend paid by the stock, or invest in the riskless security and add the known payoff to your task earnings.

You will then see the dividend paid by the stock, no matter if you chose the stock or the bond.

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

(1) what you think is the probability that the stock is the good one (the answer must be a number between 0 and 100—do not add the % sign, just type in the value);

(2) how much you trust your ability to come up with the correct probability estimate that the stock is good. In other words, we want to know how confident you are that the probability you estimated is correct (answer is between 1 and 9, with 1 meaning you have the lowest amount of confidence in your estimate, and 9 meaning you have the highest level of confidence in your ability to come up with the right probability estimate).

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

As you observe the dividends paid by the stock you will update your belief whether or not the stock is good. It may be that after a series of good dividends, you think the probability of the stock being good is 75%. However, how much you trust your ability to calculate this probability could vary. Sometimes you may not be too confident in the probability estimate you calculated and sometimes you may be highly confident in this estimate. For instance, at the very beginning of each block, the probability of the stock being good is 50% and you should be highly confident in this number because you are told that the computer just picked at random the type of stock you will see in the block, and nothing else has happened since then.

Every time you provide us with a probability estimate that is within 5% of the correct value (e.g., correct probability is 80% and you say 84%, or 75%) we will add 10 cents to your payment for taking part in this study.

Throughout the task you will be told how much you have accumulated through dividends paid by the stock or bond you chose up to that point.

Details for the Investment Evaluation Task:

This task is exactly as the task described above, except for the fact that you will not be making any investment choices. You will observe the dividends paid by the stock in either the GAIN or the LOSS conditions, and you will be asked to provide us with your probability estimate that the stock is good, and your confidence in this estimate. In this task, therefore, your payment only depends on the accuracy of your probability estimates.

Your final pay for completing the investment tasks will be:

$23 + 1/10 × Investment Payoffs + 1/10 × Number of accurate probability estimates, where Investment Payoffs = Dividends of securities you chose in the experiment, in both the GAIN and the LOSS conditions.

Please note: cell phones must be off. No drinks, food, or chewing gum is allowed during the experiment. Thank you!

B. Objective Bayesian Posterior Beliefs

The table below provides all possible values for the objectively correct Bayesian posterior that the stock is paying from the good dividend distribution, starting with a 50% to 50% prior, and after observing each possible dividend history path in a learning block. Every trial a new dividend (high or low) is revealed. There are six trials in each learning block.

The objective Bayesian posterior that the stock is the good one, after observing t high outcomes in n trials so far is given by: $\frac{1}{1+\frac{1-p}{p}*(\frac{q}{1-q})^{n-2t}}$, where $p = 50\%$ is the prior that the stock is good (before any dividends are observed in that learning block) and $q = 70\%$ is the probability that a good stock pays the high (rather than the low) dividend in each trial.

n Trials So Far	t High Outcomes So Far	Probability {stock is good | t high outcomes in n trials}
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%

C. Measures of Financial Literacy and Risk Preferences

To get measures of financial literacy and risk preferences, each participant was asked the following questions after the completion of the experimental tasks: “Imagine you have saved $10,000. You can now invest this money over the next year using two investment options: a U.S. stock index mutual fund, which tracks the performance of the U.S. stock market, and a savings account. The annual return per dollar invested in the stock index fund will be either +40% or −20%, with equal probability. In other words, it is equally likely that for each dollar you invest in the stock market, at the end of the one year investment period, you will have either gained 40 cents, or lost 20 cents. For the savings account, the known and certain rate of return for a one year investment is 5%. In other words, for each dollar you put in the savings account today, for sure you will gain 5 cents at the end of the one year investment period. We assume that whatever amount you do not invest in stocks will be invested in the savings account and will earn the risk-free rate of return. Given this information, how much of the $10,000 will you invest in the U.S. stock index fund? Choose an answer that you would be comfortable with if this was a real-life investment decision. The answer should be a number between $0 and $10,000."

After each participant wrote their answer to this question, they were asked the following: "Let's say that when you answered the prior question you decided to invest x dollars out of the $10,000 amount in the U.S. stock index fund, and therefore you put $(10,000 - x)$ dollars in the savings account. Recall that over the next year the rate of return on the stock index fund will be +40% or −20%, with equal probability. For the savings account, the rate of return is 5% for sure. What is the amount of money you expect to have at the end of this one year investment period? Please choose one of the answers below. If you choose the correct answer, you will get a $5 bonus added to your pay for this experiment. [A] $0.5\,(0.4\,x - 0.2\,x) + 0.05\,(10{,}000 - x)$; [B] $1.4\,x + 0.8\,x + 1.05\,(10{,}000 - x)$; [C] $0.4\,(10{,}000 - x) - 0.2\,(10{,}000 - x) + 0.05\,x$; [D] $0.5\,[\,0.4\,(10{,}000 - x) - 0.2\,(10{,}000 - x)] + 0.05\,x$; [E] $0.4\,x - 0.2\,x + 0.05\,(10{,}000 - x)$; [F] $0.5\,(1.4\,x + 0.8\,x) + 1.05\,(10{,}000 - x)$; [G] $1.4\,(10{,}000 - x) + 0.8\,(10{,}000 - x) + 1.05\,x$; [H] $0.5\,[\,1.4\,(10{,}000 - x) + 0.8\,(10{,}000 - x)] + 1.05\,x$."

The correct answer to this question is [F]. The actual choices (if other than [F]) made by participants indicate three different types of errors that can occur when calculating the expected value of their portfolio holdings: a lack of understanding of statements regarding probabilities (answers [B], [C], [E], [G]); a lack of understanding of the difference between net and gross returns (answers [A],[C], [D], and [E]); and confusing the stock versus risk-free asset investments (answers [C], [D], [G], and [H]). Therefore, a financial knowledge score varying between zero and three can be constructed, based on the number of different types of errors contained in the answer provided by each participant (i.e., zero errors for answer [F], one error for answers [A], [B], and [H], two errors for answers [D], [E], and [G], and three for answer [C]). Hence a financial knowledge score of three indicates a perfect answer, while a score of zero indicates that the participant's answer included all three possible types of errors. Of the 87 participants, 45 made no errors, 24 made one type of error only, 17 made two types of errors, and one person made all three possible types of errors.

REFERENCES

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2007, Real-time price discovery in global stock, bond and foreign exchange markets, *Journal of International Economics* 73, 251–277.

Asparouhova, Elena, Peter Bossaerts, Jon Eguia, and William Zame, 2010, Cognitive biases, ambiguity aversion and asset pricing in financial markets, Working paper, University of Utah.

Banerjee, Abhijit, and Esther Duflo, 2011, *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty* (Public Affairs, New York, NY).

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.

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

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears and risk premia, *Journal of Finance* 66, 2165–2211.

Bossaerts, Peter, 2004, Filtering returns for unspecified biases in priors when testing asset pricing theory, *Review of Economic Studies* 71, 63–86.

Bruguier, Antoine J., Steven R. Quartz, and Peter Bossaerts, 2010, Exploring the nature of trader intuition, *Journal of Finance* 65, 1703–1723.

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

Carlin, Bruce, Shimon Kogan, and Richard Lowery, 2013, Trading complex assets, *Journal of Finance* 68, 1937–1960.

Doll, Bradley B., Kent E. Hutchison, and Michael J. Frank, 2011, Dopaminergic genes predict individual differences in susceptibility to confirmation bias, *Journal of Neuroscience* 31, 6188–6198.

Eppinger, Ben, Michael Herbert, and Jutta Kray, 2010, We remember the good things: Age differences in learning and memory, *Neurobiology of Learning and Memory* 93, 515–521.

Frank, Michael J., Ahmed A. Moustafa, Heather M. Haughey, Tim Curran, and Kent E. Hutchison, 2007, Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning, *Proceeding of the National Academy of Sciences* 104, 16311–16316.

Froot, Kenneth A., 2001, The market for catastrophe risk: A clinical examination, *Journal of Financial Economics* 60, 529–571.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg, 2006, Costly information acquisition: Experimental analysis of a boundedly rational model, *American Economic Review* 96, 1043–1068.

Garcia, Diego, 2012, The kinks of financial journalism, Working paper, University of North Carolina at Chapel Hill.

Gennaioli, Nicola, and Andrei Shleifer, 2010, What comes to mind, *Quarterly Journal of Economics* 125, 1399–1433.

Hoberg, Gerard, and Gordon Phillips, 2010, Real and financial industry booms and busts, *Journal of Finance* 65, 45–86.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.

Kluger, Brian D., and Steve B. Wyatt, 2004, Are judgment errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem, *Journal of Finance* 59, 969–997.

Knutson, Brian, and Peter Bossaerts, 2007, Neural antecedents of financial decisions, *Journal of Neuroscience* 27, 8174–8177.

Kogan, Shimon, 2009, Distinguishing the effect of overconfidence from rational best-response on information aggregation, *Review of Financial Studies* 22, 1889–1914.

Kuhnen, Camelia M., and Brian Knutson, 2005, The neural basis of financial risk taking, *Neuron* 47, 763–770.

Kuhnen, Camelia M., and Brian Knutson, 2011, The influence of affect on beliefs, preferences, and financial decisions, *Journal of Financial and Quantitative Analysis* 46, 605–626.

Lusardi, Ana-Maria, and Olivia S. Mitchell, 2007, Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth, *Journal of Monetary Economics* 54, 205–224.

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

Mather, Mara, and Andrej Schoeke, 2011, Positive outcomes enhance incidental learning for both younger and older adults, *Frontiers in Neuroscience* 5, 1–10.

Mullainathan, Sendhil, and Eldar Shafir, 2013, *Scarcity: Why Having Too Little Means So Much* (Times Books, New York, NY).

Payzan-LeNestour, Elise, 2010, Bayesian learning in unstable settings: Experimental evidence based on the bandit problem, Swiss Finance Institute Research Paper No 10-28, 1–41.

Phillips, Lawrence D., and Ward Edwards, 1966, Conservatism in a simple probability inference task, *Journal of Experimental Psychology* 72, 346–354.

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

Rabin, Matthew, 2000, Risk aversion and expected-utility theory: A calibration theorem, *Econometrica* 68, 1281–1292.

Routledge, Brian R., and Stanley E. Zin, 2010, Generalized disappointment aversion and asset prices, *Journal of Finance* 65, 1303–1332.

Sokol-Hessner, Peter, Colin F. Camerer, and Elizabeth A. Phelps, 2013, Emotion regulation reduces loss aversion and decreases amygdala responses to losses, *Social Cognitive and Affective Neuroscience* 8, 341–350.

Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2010, Information acquisition and underdiversification, *Review of Economic Studies* 77, 779–805.