

Stock Ownership and Learning from Financial Information*

Camelia M. Kuhnen[†] Sarah Rudolf[‡] Bernd Weber[§]

Abstract

We document that people's prior portfolio choices influence the way they learn from new information about available investment options. Specifically, we find that people update more from information which is consistent with their prior choice. This behavioral effect is mirrored by a bias in activation in brain centers important for valuation, as these centers are more responsive to new information about investment options which matches the participants' prior portfolio choice. These findings can help shed light on puzzling patterns in investor behavior, such as the low participation rate of households in equity markets and people's reluctance to sell losing stocks.

*Funding for this work was provided by the German Research Council (DFG; We 4427/3-1), the Frankfurt Institute of Risk Management and Regulation, and the Wells Fargo Center for Corporate Finance at the University of North Carolina at Chapel Hill. We thank participants at the 2014 meeting of the Society for Neuroeconomics, the 2015 Consumer Financial Protection Bureau Research Conference and the 2015 MIT Sloan Interdisciplinary Symposium on Decision Neuroscience for valuable comments and discussion.

[†]University of North Carolina, Kenan-Flagler Business School & NBER, 300 Kenan Center Dr., MC #4407, Chapel Hill, NC 27599, USA, E-mail: camelia_kuhnen@kenan-flagler.unc.edu (Corresponding author)

[‡]University of Bern, Institute of Psychology, Social Psychology and Social Neuroscience, Fabrikstrasse 8, 3012 Bern, Switzerland, E-mail: sarah.rudorf@psy.unibe.ch.

[§]University of Bonn, Center for Economics and Neuroscience, Nachtigallenweg 86 and University Hospital Bonn, Department of Epileptology, Sigmund-Freud Str. 25, 53127 Bonn, Germany, E-mail: bweber@lifeandbrain.com.

I. Introduction

One of the most puzzling patterns in household finance and a topic of current policy debates is that the majority of people in the U.S. and Europe do not invest in the stock market (Campbell (2006), Calvet et al. (2007)), which results in lower wealth accumulation and consumption over the life span (Mankiw and Zeldes (1991)). While in part the low participation rates in equity markets may be driven by insufficient provision of financial services to those willing to invest, it is also possible that some individuals are unable to process financial information and thus choose to avoid the stock market (Grinblatt et al. (2011), Van Rooij et al. (2011)).

Here, we test a specific mechanism that could lead people to have incorrect beliefs about the outcomes of stock investments, which in turn could change their willingness to participate in equity markets. In particular, using behavioral and brain imaging data, we test whether people's ability to learn from new financial information may mistakenly depend on their prior investment choices, in a manner that would make those not currently holding stocks to be more pessimistic about the potential outcomes of these risky assets, and thus less willing to invest.

An implicit assumption in the finance literature is that market participants are able to learn the same way from new information about available investments, irrespective of the composition of their portfolio. While theoretical work has shown that previous portfolio choices may influence investors' *utility function*, based on prospect theory (Barberis et al. (2001)) or realization theory (Barberis and Xiong (2012)), it is possible that these prior choices might also change investors' *beliefs* or the learning rules they use to incorporate financial market news. Experimental evidence from a sample of college students with limited financial expertise (Kuhnen and Knutson (2011)) suggests that people may update beliefs in a way to be consistent with their prior investment choices. However, it is important to determine whether this bias in beliefs induced by previous portfolio decisions generalizes to individuals who are responsible for household financial decisions and have available means

to invest. If that is the case, brain imaging might shed light on the underlying processes driving this effect.

Thus, in this paper, we seek to understand whether indeed previous investment choices influence people’s ability to use financial information correctly, and to uncover the brain mechanisms underlying this effect. In an experimental setting with working-age, highly educated and high-earning adults faced with real financial incentives, we find that prior investment decisions interfere with these individuals’ ability to correctly update their beliefs about the quality of financial assets available to them. In particular, controlling for subjective beliefs prior to making an investment choice, when participants choose to hold a stock with risky payoffs, they subsequently form more optimistic beliefs about the quality of the stock, relative to when they choose to hold a safe bond. Moreover, stock holders update beliefs about the stock quality more after observing a high dividend of the stock, rather than a low dividend, whereas bond holders update their beliefs more after observing a low dividend of the stock, rather than a high one. This shows that people are more likely to learn from new information which ex-post justifies their prior investment choice. Importantly, in the financial investment task used in the study, it is optimal for participants to learn objectively from all new outcomes, as their prior choices do not constrain them from changing their portfolio going forward.

Furthermore, we show that this behavioral effect whereby prior choices interfere with people’s ability to objectively learn from new information is accompanied by a bias induced by prior choices in the reactivity of brain regions important for valuation – namely, the ventromedial prefrontal cortex (vmPFC), the posterior cingulate cortex (PCC) and the ventral striatum – to new information. We find that these areas preferentially encode new information which matches the participants’ prior investment choice. Specifically, when participants choose to have the stock in their portfolio in a given trial, activation in these areas is significantly higher if the new dividend paid by the stock in that trial is high, relative to when they choose to hold the bond. We then show that this muted brain reactivity shown by bond

holders in response to high stock outcomes can help predict the errors in beliefs subsequently expressed by these individuals regarding the quality of the stock's payoff distribution.

Existing work has documented that people form beliefs indicating that they view themselves as better individuals, with a better future, than deemed by objective data. Specifically, Eil and Rao (2011) and Mobius et al. (2014) provide experimental evidence that people update more from good news regarding their own beauty or intelligence, relative to bad news. Sharot et al. (2011) show that people underestimate the likelihood that adverse life events, such as illness, will happen to them, and found that this unfounded optimism about one's odds of encountering bad events was related to a diminished brain response to undesirable information. Mayraz (2013) provides a theoretical model and experimental evidence supporting the idea of wishful thinking, namely, that people believe that the future outcomes that will occur will be those that maximize their payoffs. Somewhat optimistic beliefs have been shown to lead to better overall utility, due to anticipatory effects, even though they may lead to suboptimal actions (Brunnermeier and Parker (2005)). We add to this literature by documenting that people's prior choices induce a bias in how they learn from new information, in that they update more from news which confirms, rather than disconfirms, their prior choices, in a setting where it is optimal to update beliefs correctly. Hence, the evidence in this paper brings forward a novel mechanism through which people can form incorrect beliefs about future outcomes.

Our paper contributes to the finance literature regarding individual investor behavior. Our findings can help shed light on the puzzling fact that a large share of households do not participate in the stock market (Campbell (2006), Calvet et al. (2007)), a behavior which is detrimental to wealth accumulation. Our findings suggest that one potential reason for this outcome may be that people who are not currently stock holders update their beliefs insufficiently if stocks perform well, and hence will be overly pessimistic about future payoffs in the stock market. These pessimistic beliefs will in turn deter these individuals from investing in equities, which will lead to limited stock market participation in the population.

Our results may also help explain another puzzling aspect of investor behavior, namely, the disposition effect (Odean (1998)), which refers to the fact that investors seem to be reluctant to sell stocks that have not performed well. Our results suggest that a potential reason for this pattern is that these investors do not update sufficiently their beliefs after observing low outcomes of stocks they have previously chosen. This belief-related channel for the disposition effect complements the findings in Frydman et al. (2014), who find evidence suggesting that people’s utility function is affected negatively by the realization of losses in their portfolio. Our results also complement the recent experimental literature in finance regarding learning in markets, which shows that investors exhibit deviations from Bayesian learning (Kluger and Wyatt (2004), Payzan-LeNestour and Bossaerts (forthcoming), Asparouhova et al. (forthcoming)), and that strategic considerations regarding others’ behavior influence learning and trading (Kogan (2009), Bruguier et al. (2010), and Carlin et al. (2013)).

II. Experimental design

Participants completed a financial decision making task based on the experimental protocol in (Kuhnen (forthcoming)). Each participant made 96 decisions, split into 16 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 which 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.

The task included gain and loss blocks, as learning may differ across these domains (Kuhnen (forthcoming)). In gain blocks, the two securities provided positive payoffs only. The stock payoffs were + €10 or + €2 (gain & low variance condition), or €0 or + €12 (gain & high variance condition), while the bond payoff was + €6. In loss blocks, the two

securities provided negative payoffs only. The stock payoffs were - €10 or - €2 (loss & low variance condition), or €0 or - €12 (loss & high variance condition), while the bond payoff was - €6 (see Table I).

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

For each block of six trials, the participants' learning problem is the same. That is, participants know that the computer will either pay dividends from the good stock distribution in each of these six trials, or it will pay from the bad distribution in each of the six trials. At the beginning of each learning block, the computer randomly selected (with 50%-50% probabilities) whether the dividend distribution to be used in the following six trials would be the good or the bad one.

Importantly, participants were carefully instructed that the probability that the stock will pay from the good distribution is 50% for each learning block. Participants were given this information and the experiment proceeded when they confirmed that they understood that they would start each block with this 50% prior. Based on this 50% prior belief, as well as on the dividend revealed each trial, participants were asked to provide a posterior probability estimate that the stock is paying from the good distribution in that block. In total, each subject faced 16 learning blocks, split equally in gain vs. loss, and high vs. low variance conditions. The order of the blocks was pseudo-randomized (see Figure 1 for trial examples).

Participants were paid based on their investment payoffs and the accuracy of the probability estimates provided. Specifically, they received the accumulated payoffs of the investments they chose throughout the task, plus ten Euro cents for each probability estimate within 5% of the correct answer, namely, the objective Bayesian posterior probability. Information re-

garding the accuracy of each subject’s probability estimates and the corresponding payment was only provided at the end of the task. This was done to avoid feedback effects that could have changed the participants’ strategy or answers during the progression of the task.

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: $\frac{1}{1 + \frac{1-p}{p} * (\frac{q}{1-q})^{n-2t}}$, 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 the low) dividend in each trial. Table A1 in the Appendix provides the value of the objective Bayesian posterior for all $\{n, t\}$ pairs possible in the experiment. This Bayesian posterior is our benchmark for measuring how close the subjects’ expressed probability estimates are from the objectively correct beliefs.

The 46 participants in the study – all male (to avoid gender effects, see Almenberg and Dreber (2012)), age $40.08 \pm \text{sd. } 6.53$ years, age range 29-49 years – were recruited in Bonn, Germany. Participants gave written informed consent, as required by human subjects protection rules. The experiment lasted approximately 2 hours. Subjects received €25 at the outset, as well as the payoff accumulated during the financial learning task. If the final payoff was less than €20 per hour, the difference was compensated for. On average, participants’ compensation for this study was €41.4.

Importantly, the sample was purposely selected to be representative of educated, high-earning individuals across the age span: 78% of the participants have a college degree, 49% are home owners, and 60% earn more than €4000 per month.

III. Results

A. Behavioral results

We first document that the subjective beliefs we elicit from participants each trial are meaningful quantities, in the sense that they predict these individuals' investment choices in subsequent trials. In Figure 2 we present the mean frequency with which the stock, and not the bond, was chosen, as a function of the participants' estimate in the prior trial regarding the probability that the stock pays from the good dividend distribution. This belief can either be strictly below 50%, exactly equal to 50%, or strictly above 50%. Individuals who are risk-neutral and who maximize expected value would never choose the stock in the first case, would randomize between stock and bond in the second case, and would always choose the stock in the third case. We find that the mean frequencies with which participants choose the stock in these three cases are 26.66%, 48.45% and 73.65%, respectively. Thus, when participants' belief that the stock is good is 50% or strictly greater than 50%, the probability that they would subsequently choose the stock is greater by 22% and 47%, respectively, relative to situations when their belief is strictly below 50%. These differences are significant at $p < 0.01$.

While the relationship between subjective beliefs and subsequent choices is strong and in the expected direction, as more optimistic beliefs predict a higher likelihood that the stock will be chosen in the next trial, it is also the case that there is a substantial degree of stochasticity in participants' choices, as observed in many other economic settings (see Krajbich et al. (2014)). In the probit model in the first column of Table II we show that the probability estimates of the participants explain 12% of variation in the investment choice. The estimated effects are in line with those documented in Figure 2. Specifically, we find that the probability that the stock is chosen by participant i in trial t is higher in situations when the person's belief at the end of trial $t - 1$ was equal to 50% or strictly above 50%, relative to those when this belief was strictly below 50%, by 23% and 47%, respectively.

Adding subject fixed effects as explanatory variables for the decision to choose the stock, not the bond, in a given trial, increases the R-squared of the investment choice model to 19%.

This means that in the data there exist many situations where we observe the same participant expressing the same subjective probability that the stock is good, but the person's choice in the next trial is sometimes the stock and some other times the bond. We can thus exploit this within-subject variation in the investment decision, controlling for prior subjective beliefs, to conduct our main analysis – namely, to study the effect of the investment choice on how people learn from new information after they make this choice.

As predicted, we find that prior investment choices influence participants' posterior beliefs about the stock dividend distribution, and the manner in which they update from new dividend information. Panel (a) of Figure 3 shows that controlling for the beliefs expressed by subjects regarding the quality of the stock before the investment choice is made, when individuals choose the stock they form more positive posterior beliefs about the stock dividend distribution compared to when they choose the bond ($p < 0.01$). Specifically, in trials where, before the choice, participants' prior belief that the stock is the good one was strictly less than 50%, the average posterior belief that the stock is good is 42.16% for stock holders, and it is 31.21% for bond holders. In trials where participants' prior before the choice was equal to 50%, the average posterior belief that the stock is the good one is 51.69% for stock holders, and 48.59% for bond holders. Finally, in trials where participants' prior before the choice was strictly higher than 50%, the average posterior belief that the stock is the good one is 68.71% for stock holders, and 60.43% for bond holders.

We find similar effects in a multivariate regression analysis presented in Table III. The dependent variable is the probability estimate produced by a subject in a particular trial, and the independent variables include subject fixed effects, the subject's probability estimate in the prior trial, an interaction term that indicates whether a person chose the stock in the trial and the stock paid a high dividend that trial, as well as indicators for whether the

trial is in the gain vs. loss condition, and in the low vs. high variance condition. This analysis yields within-subject estimates of the effect of the stock vs. bond choice, as well as of the effect of observing a high vs. a low dividend, conditional on the asset choice, on the probability estimates produced by participants in each trial. Confirming the findings in panel (a) of Figure 3, the regression analysis in Table III shows that, controlling for prior beliefs expressed right before the asset choice, the mere choice of the stock leads participants to express a posterior belief about the probability that the stock is the good one that is 4.02% higher ($p < 0.01$) than the posterior expressed in instances when the same individuals chose the bond.

Moreover, the interaction term that captures situations when a participant chose the stock and the stock that trial paid a high dividend is a significant predictor of the participant's posterior belief about the probability that the stock pays from the good dividend distribution, even after controlling for an indicator variable for whether the dividend was high. Specifically, the coefficient estimate for this interaction term in Table III indicates that upon the revelation of a high stock dividend in a given trial, a participant provides a posterior estimate of the probability that the stock is good which is 4.85% higher ($p < 0.05$) in instances when the participant had chosen to hold the stock that trial, relative to instances when the same person, with the same prior, had chosen to hold the bond that trial. Therefore, the results in Table III show that not only do stock holders end up with more optimistic assessments of the stock quality, confirming the effects shown in Figure 3, but they also react differently to dividend information relative to bond holders, by updating more from high dividends.

The analysis in panel (b) of Figure 3 further illustrates that updating upon the release of new dividend information is different for stock holders relative to bond holders, even though in this task it is optimal for subjects to learn objectively from all new outcomes, as their prior choices do not constrain them from changing their portfolio going forward.¹ To under-

¹Individuals who update incorrectly in this task earn significantly less than those who update and choose assets optimally. To measure the extent of losses from suboptimal learning and behavior, we ran simulations

stand how people update their beliefs when given new information about the stock dividend distribution, we estimate regression models where the dependent variable is the change in a subject’s probability estimate from the prior to the current trial, and the independent variable is their probability estimate from the prior trial. We estimate these regressions for trials where the stock dividend is low (left side of the panel), and for trials where the stock dividend is high (right side of the panel), separately for bond holders and for stock holders.

As expected, the figure shows that the average update regarding the probability that the stock is paying from the good distribution when a high dividend is revealed is positive, in accordance with Bayesian learning. However, this update is significantly greater for stock holders, relative to bond holders, on average by 9% (see Table III), controlling for these individuals’ beliefs about the stock quality prior to them making the asset choice ($p < 0.05$). The average update regarding the probability that the stock is paying from the good distribution when a low dividend is revealed is negative, in accordance to Bayesian learning, but it is more negative for bond holders (by 4% on average, $p < 0.05$, see Table III), relative to stock holders.

Therefore, the results in panel (b) of Figure 3 and in Table III indicate that investors update more from new information which ex-post justifies their prior investment choice. Importantly, this effect is different from the classic confirmation bias (Myanatt et al. (1977)), which refers to people’s tendency to choose information sources that can help confirm a particular hypothesis, instead of seeking sources that can help reject it. Our results suggest that when people obtain explicit falsifying information that does not match their prior choice, they fail to use it to reject incorrect hypotheses. Also, in the regression model in Table III

for different possible decision strategies to estimate the subjects’ compensation. Individuals who always choose the bond earn just their initial endowment (€25) due to the equal number of gain and loss trials. Based on simulations (1000 subjects per strategy), perfect Bayesian investors earn on average €65.4, whereas individuals who always choose the stock earn on average €26.4. Those who choose assets based solely on the most recent outcome earn on average €58. Aside from these earnings that come from accumulating the payoffs of the chosen assets, Bayesian learners would also earn €9.60 for 96 correct probability estimations. Individuals who never update and always report a 50% posterior would be correct on average on 28.8 trials and get €2.88. Therefore if participants were to always choose the bond and never update the 50% prior, they would earn a total of €27.88, whereas Bayesian optimizers would earn €75.

we replicate the result in (Kuhnen (forthcoming)) that probability estimates are higher in the gain domain than in the loss domain, that is, people are overly pessimistic when learning from negative financial outcomes. The same regression shows that whether the stock dividend distribution has high or low variance does not have a significant impact on subjective probability estimates.

We also investigate in which instances people form posterior beliefs that are the closest to the objective, Bayesian posterior probability that the stock is paying from the good distribution. The analysis presented in Figure 4 documents that there are two types of situations when subjective estimates are particularly close to Bayesian values, that is, when learning performance is best. Specifically, we find that instances when subjective posteriors are the closest to the objective Bayesian beliefs are as follows: (1) when objectively the stock is likely to be the good one (i.e., objective probability > 50%), stock holders who see a high dividend are those who form the most accurate posteriors; (2) when objectively the stock is unlikely to be the good one (i.e., objective probability < 50%), bond holders who see a low dividend are those who form the most accurate posteriors. In other words, stock holders are particularly good at evaluating stocks during good times, when good news is revealed about stocks, whereas bond holders are particularly good at evaluation stocks during bad times, when bad news is revealed about these investments.

B. Brain imaging results

In support of our main behavioral result – namely, that prior portfolio choices influence learning – we find that these choices bias the brain response to new information. As shown in the analysis in Figure 5, stock holders are more likely than bond holders to experience an increase in activation in valuation-related brain areas when a high stock dividend is revealed.

To conduct the analysis in Figure 5, we estimated a generalized linear regression that modeled brain activation, as measured by the blood-oxygen-level dependent (BOLD) response, at the time of the stock outcome presentation separately by the prior investment

(stock or bond) and the presented stock outcome (high or low dividend), while controlling for the prior probability that the stock is good (see Appendix). A two-by-two ANOVA of the whole-brain data revealed that a cluster extending in the ventromedial prefrontal cortex and the anterior cingulate cortex, the bilateral ventral striatum, the posterior cingulate cortex and a region in the inferior parietal lobule were sensitive to an interaction of the prior investment and the stock outcome ($p < 0.05$, FWE-corrected, see Table A2 in the Appendix).

To illustrate the nature of this interaction we plotted the mean beta parameters in our regions of interest in Figure 5, and found that if subjects had chosen the stock at the beginning of the trial they showed a stronger BOLD response to high stock dividends than if they had chosen the bond. Given the within-subject design of the experiment, a subject is considered a stock holder or a bond holder for a particular trial only, as he is free to switch his investment in the next trial.

As Figure 5 indicates, we find that activity in the ventromedial prefrontal cortex, ventral striatum and posterior cingulate cortex is greater for high stock dividends when the subject holds a stock than a bond. The left panel shows statistical parametric maps of the two-by-two ANOVA. The right panel shows the mean and standard error of the beta estimates for stock vs. bond holders, and for high and low stock dividends separately, in these three regions of interest. These estimates indicate that these areas preferentially encode new information that matches the participants' prior investment choices, a finding that mirrors the behavioral effect of prior choices on participants' beliefs documented in Figure 3 and Table III.

C. Predicting errors in beliefs

We find that the brain response to new information that contradicts a participant's prior investment choice can help predict the person's ability to form correct beliefs about the quality of the stock. Specifically, when participants were faced with information contradicting their prior choice, higher activation in the vmPFC or the left ventral striatum (but not other

regions) during dividend presentation leads to a significantly lower probability estimation error in that trial (Table IV). These are situations when participants who chose the bond at the beginning of the trial subsequently observed that the stock paid a high dividend that trial. In those types of trials, the average probability estimation error is 15%. The coefficients in the table imply that a standard deviation increase in vmPFC activation, left ventral striatum activation, or in the common component of these two, leads to a reduction in the estimation error by 1%. This suggests that learning from financial information is improved if valuation-related brain centers are able to correctly encode new information, in spite of it being in conflict with the person’s previous investment choice.

These brain imaging results complement the recent literature that identifies the vmPFC, ventral striatum and PCC as the key regions of a subjective value network (Bartra et al. (2013), Clithero and Rangel (2014)). Here we show that they are critically involved in the updating of subjective probability estimates. This is in line with research associating the vmPFC and ventral striatum with the encoding of objective and subjective probability (d’Acremont et al. (2013), Ting et al. (2015)).

Importantly, we document the novel result that the engagement of the vmPFC and ventral striatum at the time when new stock dividend information is presented is biased by irrelevant factors – namely, the subjects’ prior portfolio choices, and this has detrimental consequences for people’s ability to correctly assess their investment opportunities.

IV. Conclusion

We document that prior investment decisions influence people’s ability to correctly form beliefs about the quality of financial assets. Specifically, controlling for prior beliefs, we find that if the participants’ most recent choice is a stock, they will form more optimistic posterior beliefs about the dividend distribution of the stock, and will update their beliefs more after observing a high dividend, rather than a low one. If the participants’ most recent choice is a

bond, they will form less optimistic posterior beliefs about the quality of the stock, and will update their beliefs more after observing a low dividend, rather than a high one. In other words, individuals learn more from financial information which confirms their prior portfolio choice.

This behavioral asymmetry in learning that is induced by participants' prior investment choice is mirrored by an asymmetry in the response of valuation-related brain areas when new dividend information is presented. Specifically, activation in these areas increases significantly more when a high dividend is observed in trials where participants chose to hold the stock, rather than when they chose the bond. Moreover, the degree to which brain activation in these areas increases at the presentation of disconfirming evidence predicts the accuracy of participants' evaluation of the quality of investment options available to them.

Our behavioral result indicating that prior choices impact beliefs, together with the brain-based evidence that we provide supporting this biased learning effect, brings forward a novel mechanism through which people can form incorrect expectations about future outcomes. Here we examine this new mechanism – that is, the impact of prior choices on belief updating – in the context of financial decision making. Future work may investigate whether this mechanism applies in other, non-financial situations that require learning from news which may contradict one's past actions.

REFERENCES

- Almenberg, J. and Dreber, A.: 2012, Gender, financial literacy and stock market participation, *Working paper* .
- Asparouhova, E., Bossaerts, P., Eguia, J. and Zame, W.: forthcoming, Asset pricing and asymmetric reasoning, *Journal of Political Economy* .
- Barberis, N., Huang, M. and Santos, T.: 2001, Prospect theory and asset prices, *The Quarterly Journal of Economics* .
- Barberis, N. and Xiong, W.: 2012, Realization utility, *Journal of Financial Economics* **104**(2), 251–271.
- Bartra, O., McGuire, J. T. and Kable, J. W.: 2013, The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value, *NeuroImage* **76**, 412–427.
- Bruguier, A. J., Quartz, S. R. and Bossaerts, P.: 2010, Exploring the nature of "trader intuition", *Journal of Finance* **65**(5), 1703–1723.
- Brunnermeier, M. and Parker, J.: 2005, Optimal expectations, *American Economic Review* **95**(4), 1092–1118.
- Calvet, L. E., Campbell, J. Y. and Sodini, P.: 2007, Down or out: Assessing the welfare costs of household investment mistakes, *Journal of Political Economy* **115**(5), 707–747.
- Campbell, J.: 2006, Household finance, *Journal of Finance* **61**(4), 1553–1604.
- Carlin, B., Kogan, S. and Lowery, R.: 2013, Trading complex assets, *Journal of Finance* **68**(5), 1937–1960.
- Clithero, J. A. and Rangel, A.: 2014, Information parcellation of the network involved in the computation of subjective value, *SCAN* **9**, 1289–1302.

- d'Acremont, M., Schultz, W. and Bossaerts, P.: 2013, The human brain encodes event frequencies while forming subjective beliefs, *Journal of Neuroscience* **33**(26), 10887–10897.
- Eil, D. and Rao, J.: 2011, The good news-bad news effect: Asymmetric processing of objective information about yourself, *American Economic Journal: Microeconomics* **3**(2), 114–138.
- Frydman, C., Barberis, N., Camerer, C., Bossaerts, P. and Rangel, A.: 2014, Using neural data to test a theory of investor behavior: An application to realization utility, *Journal of Finance* **69**(2), 907–946.
- Grinblatt, M., Keloharju, M. and Linnainmaa, J.: 2011, Iq and stock market participation, *Journal of Finance* **66**(6), 2121–2164.
- Kluger, B. D. and Wyatt, S. B.: 2004, Are judgment errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem, *Journal of Finance* **59**, 969–997.
- Kogan, S.: 2009, Distinguishing the effect of overconfidence from rational best-response on information aggregation, *Review of Financial Studies* **22**(5), 1889–1914.
- Krajbich, I., Oud, B. and Fehr, E.: 2014, Benefits of neuroeconomic modeling: New policy interventions and predictors of preference, *American Economic Review: Papers & Proceedings* **104**(5), 501–506.
- Kuhnen, C. M.: forthcoming, Asymmetric learning from financial information, *Journal of Finance* .
- Kuhnen, C. M. and Knutson, B.: 2011, The influence of affect on beliefs, preferences, and financial decisions, *Journal of Financial and Quantitative Analysis* **46**(3), 605–626.
- Mankiw, G. N. and Zeldes, S. P.: 1991, The consumption of stockholders and nonstockholders, *Journal of Financial Economics* **29**, 97–112.

- Mayraz, G.: 2013, Wishful thinking, *Working paper* .
- Mobius, M. M., Niederle, M., Niehaus, P. and Rosenblat, T. S.: 2014, Managing self-confidence, *Working paper* .
- Myanatt, C. R., Doherty, M. E. and Tweney, R. D.: 1977, Confirmation bias in a simulated research environment: An experimental study of scientific inference, *Quarterly Journal of Experimental Psychology* **29**, 85–95.
- Odean, T.: 1998, Are investors reluctant to realize their losses?, *The Journal of Finance* **53**(5).
- Payzan-LeNestour, E. and Bossaerts, P.: forthcoming, Learning about unstable, publicly unobservable payoffs, *Review of Financial Studies* .
- Sharot, T., Korn, C. W. and Dolan, R. J.: 2011, How unrealistic optimism is maintained in the face of reality, *Nature Neuroscience* **14**(11), 1475–1479.
- Ting, C.-C., Yu, C.-C., Maloney, L. T. and Wu, S.-W.: 2015, Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference, *Journal of Neuroscience* **35**(4), 1792–1805.
- Van Rooij, M., Lusardi, A. and Alessie, R.: 2011, Financial literacy and stock market participation, *Journal of Financial Economics* **101**(2), 449–472.

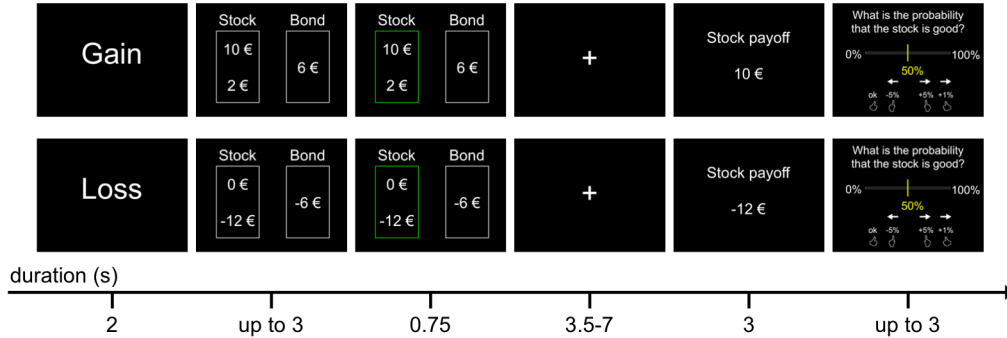


Figure 1: Timeline of trials. At the beginning of each block, an indications for the gain or loss condition is displayed for 2s, followed by a jittered interstimulus interval (ISI) of 1-3s (not shown). Then, the first choice screen is presented and the subject has 3s to make a choice. After the subject has made a choice, a green frame is presented around the chosen option for 0.75s, followed by an anticipation period during which a fixation cross is presented for a jittered period of 3.5-7s (M 6.02 \pm SD 0.41s). Next, irrespective of the subject's choice, the stock dividend is shown for 3s, followed by an ISI. The subject is then asked to estimate the probability that the current stock is a good stock. The estimation is self-paced and followed by an ISI. Finally, the subject's updated balance is presented for 3s (not shown). After a final ISI, the next trial begins. Each learning block consists of 6 trials and can be a gain or loss block, with high or low variance in the dividend distribution, rendering four task conditions. The figure shows one example of a gain block with a low variance stock (upper panel) and one example of a loss block with a high variance stock (lower panel). For each block, the stock was randomly assigned to be good or bad. If the stock was good (bad), it paid the high outcome with 70% (30%) probability and the low outcome with 30% (70%) probability. To make optimal investment choices, subjects need to update their beliefs about the stock's dividend distribution after observing each dividend paid by the stock.

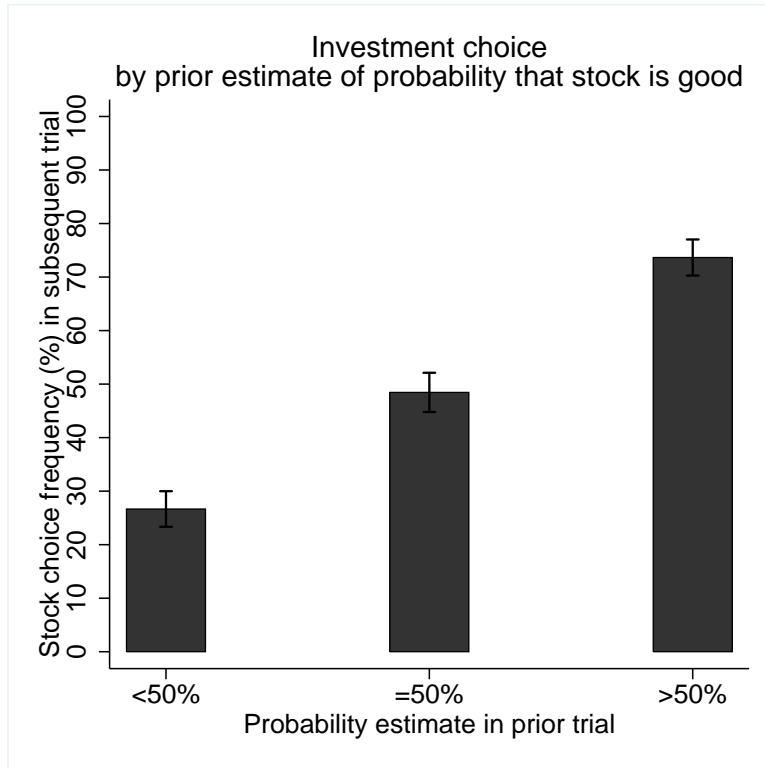
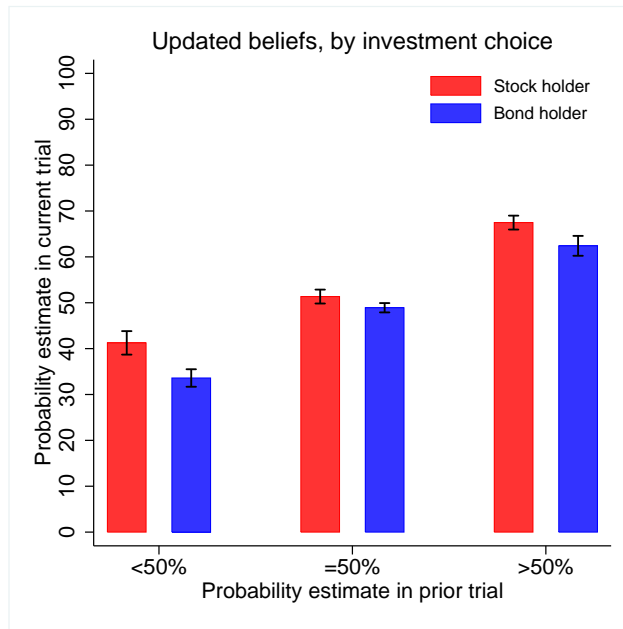
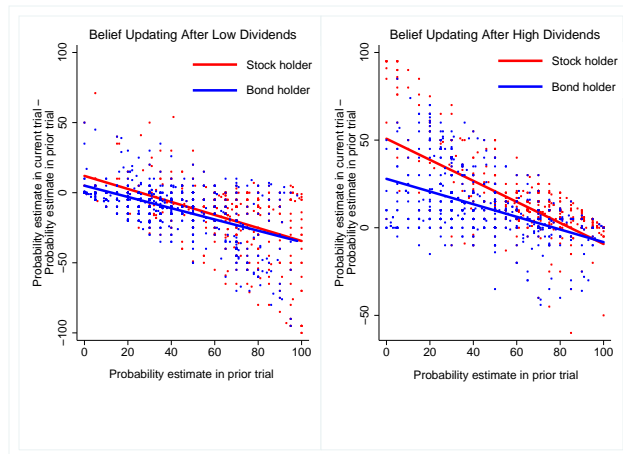


Figure 2: Subjective probability estimates predict subsequent investment choices. We split the data based on the participants' estimate in the prior trial for the probability that the stock is paying from the good dividend distribution. This prior belief can either be strictly below 50%, exactly equal to 50%, or strictly above 50%. Risk neutral expected value maximizing agents would never choose the stock in the first case, would randomize between stock and bond in the second case, and would always choose the stock in the third case. In our sample, the mean frequencies (shown above together with their standard errors) with which participants choose the stock, rather than the bond, in these three cases are: 26.66%, 48.45% and 73.65%, respectively (differences are significant at $p < 0.01$).



(a) Investment choices influence posterior beliefs.



(b) Investment choices influence the updating process.

Figure 3: Panel (a) shows that controlling for the beliefs held before the investment choice is made, when individuals choose the stock they form more positive posterior beliefs about the stock dividend distribution compared to when they choose the bond (bars show means and standard errors). Panel (b) shows that updating upon the release of new dividend information is different for stock holders relative to bond holders. The solid lines represent lines of best fit for regressions where the dependent variable is the change in the subject's probability estimate from the prior to the current trial, and the independent variable is their probability estimate from the prior trial. For each value of prior belief expressed by subjects, these linear predictions indicate the average belief update produced across all subjects, either after observing a low dividend in the current trial (left side of Panel (b)), or after observing a high dividend in the current trial (right side of Panel (b)), separately for those who chose the bond in the current trial (blue line) or for those who chose the stock in the current trial (red line).

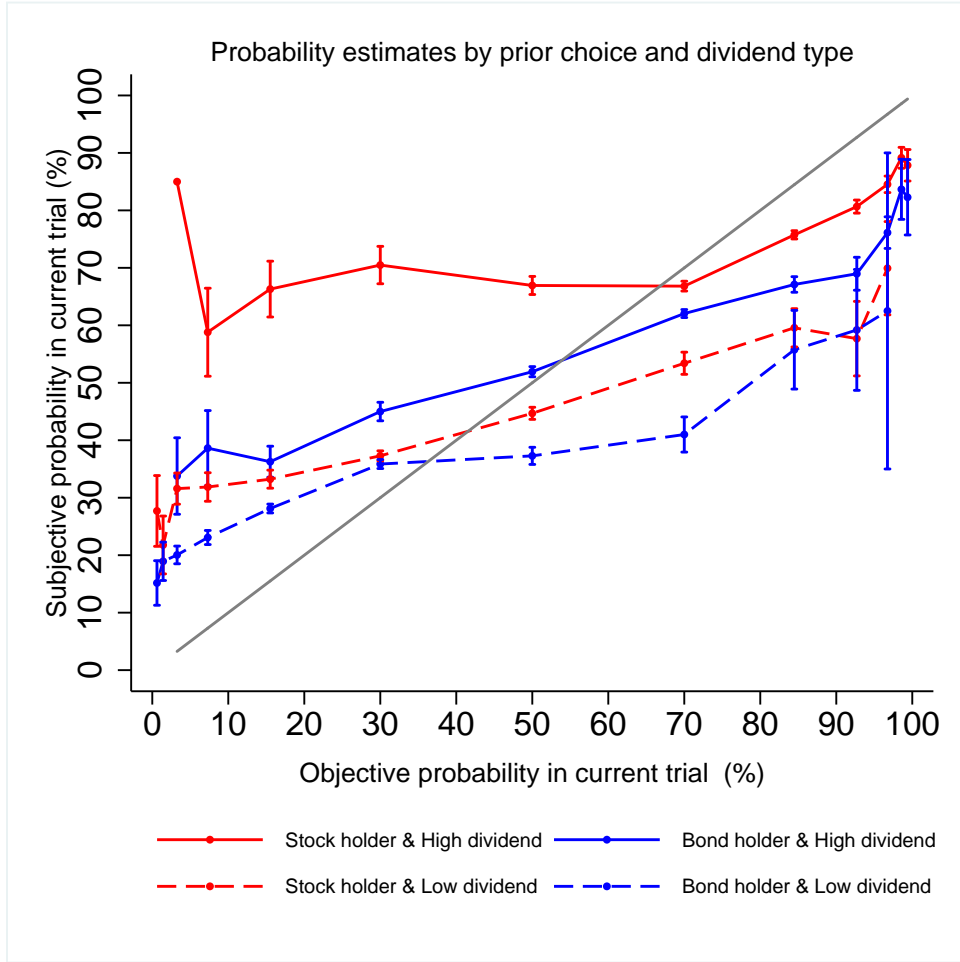


Figure 4: Average subjective estimates for the probability that the stock is paying from the good dividend distribution, for each possible value of the objective Bayesian probability, by prior choice (stock vs. bond) and dividend type (high vs. low). The objective Bayesian posteriors that the stock is good which 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 grey 45° line. Subjective probability estimates provided by participants for each level of the objectively correct Bayesian posterior, along with their standard errors, are shown in red for situations when participants chose the stock and blue for situations when participants chose the bond. Solid lines refer to trials when the stock pays the high dividend, and dashed lines refer to trials when the stock pays the low dividend. The instances when the subjective posterior beliefs are the closest to the objective Bayesian beliefs are as follows: (1) when objectively the stock is likely to be the good one (Objective probability > 50%), stock holders who see a high dividend are those who form the most accurate posteriors; (2) when objectively the stock is unlikely to be the good one (Objective probability < 50%), bond holders who see a low dividend are those who form the most accurate posteriors.

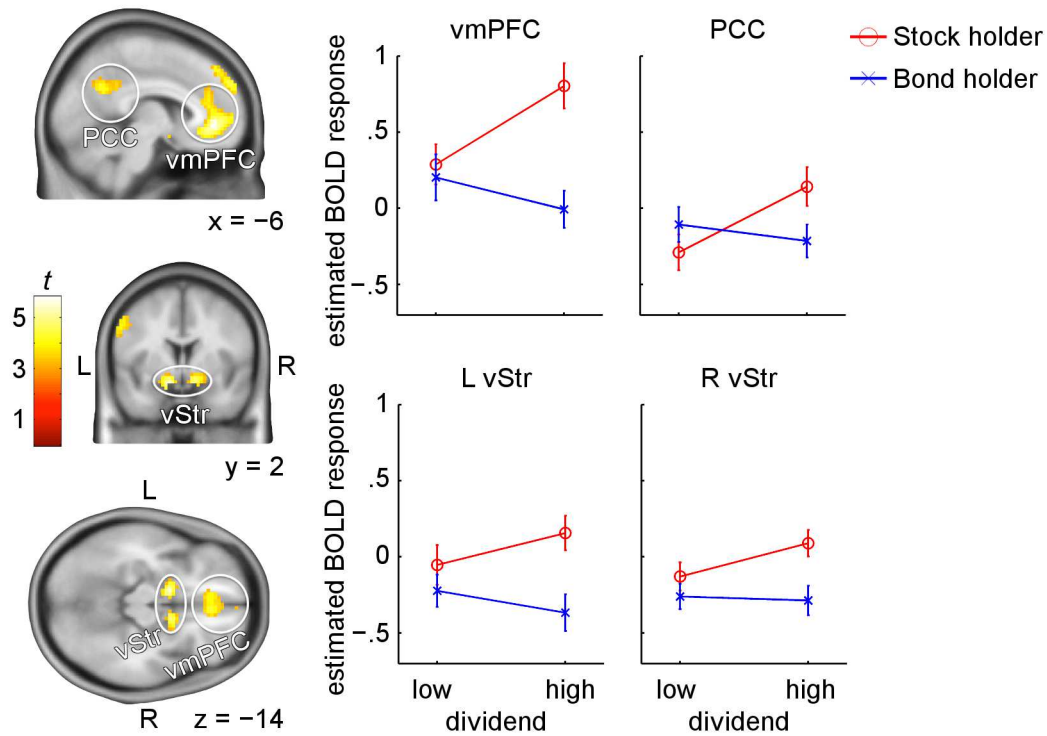


Figure 5: Stock ownership effect on the neural encoding of new stock information. Activity in the vmPFC/ACC, the bilateral vStr and the PCC is greater for high stock dividends when subjects hold the stock vs. when they hold the bond. The figures on the left show statistical parametric maps of the two-by-two ANOVA displayed at $p < 0.001$, uncorrected, projected on a template brain in MNI space, and color coded for the t -values as indicated by the color bar on the left. The graphs on the right show the mean (M) and standard error (SE) of the beta estimates for prior investments and high and low stock dividends separately, in the indicated anatomically defined regions of interest. (vmPFC=ventromedial prefrontal cortex; ACC=anterior cingulate cortex; PCC=posterior cingulate cortex; vStr=ventral striatum.)

Table I: Experimental design. Subjects made 96 decisions to invest in one of two securities: a stock with risky payoffs coming from one of two distributions, one better than the other, and a bond with a known payoff. After each choice subjects provided an estimate of the probability that the stock was paying from the better distribution. Subjects were paid based on their investment payoffs and the accuracy of the probability estimates provided. The 96 trials are split into 16 blocks of 6 trials each: 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 dividend occurs with 30% probability. The bad distribution is that where these probabilities are reversed: the high dividend occurs with 30% probability, and the low dividend 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. See Figure 1 for examples of trials.

Condition		Stock Payoffs	Bond Payoff	Number of blocks	Trials per block
Gain	Low variance	+ €10 or + €2	+ €6	4	6
Gain	High variance	+ €12 or + €0	+ €6	4	6
Loss	Low variance	- €10 or - €2	- €6	4	6
Loss	High variance	- €12 or - €0	- €6	4	6

Table II: Investment choices and subjective prior beliefs. The dependent variable in the probit models below is equal to 1 if in trial t participant i chose the stock, and not the bond. The independent variables in the specification in the first column are indicator variables for whether in the prior trial the participant's estimated for probability that the stock is paying from the good dividend distribution is equal to 50% ($ProbabilityEstimate_{i,t-1} = 50\%$) or strictly greater than 50% ($ProbabilityEstimate_{i,t-1} > 50\%$). The reference (omitted) category refers to trials prior to which the subjective estimate was strictly below 50% ($ProbabilityEstimate_{i,t-1} < 50\%$). In the specification in the second column, subject fixed effects are added as explanatory variables. Marginal effects are reported. Standard errors are clustered by subject. t -statistics are reported in parentheses. *, ** and *** denote statistical significance at $p < 0.1$, $p < 0.05$, and $p < 0.01$, respectively.

Dependent variable	<i>StockChoice_{it}</i>	
<i>ProbabilityEstimate_{i,t-1} = 50%</i>	0.23 (5.81)***	0.25 (6.23)***
<i>ProbabilityEstimate_{i,t-1} > 50%</i>	0.47 (8.99)***	0.49 (8.61)***
Subject Fixed Effects	No	Yes
<i>PseudoR²</i>	0.12	0.19
Observations	4395	4395

Table III: Stock holders react more strongly than bond holders to high dividends, and are more positive, when estimating the probability that the stock is paying dividends from the good distribution. The dependent variable in the linear regression in the table, $ProbabilityEstimate_{it}$, is the probability estimate produced by subject i in trial t . The independent variables of interest are $HighDividend_{it}$, which is equal to 1 if the stock paid a high dividend that trial and 0 otherwise, $StockHolder_{it}$, which is equal to 1 if the participant chose to hold the stock at the beginning of the trial and 0 otherwise, and their interaction ($HighDividend_{it} \times StockHolder_{it}$). The regression includes subjects-fixed effects and controls for the experimental condition faced by each participant i in each trial t (i.e., gain vs. loss, low vs. high variance conditions), as well as for the subject's probability estimate in the prior trial. Standard errors are clustered by subject. t -statistics are reported in parentheses. *, ** and *** denote statistical significance at $p < 0.1$, $p < 0.05$, and $p < 0.01$, respectively.

Dependent variable	$ProbabilityEstimate_{it}$
$HighDividend_{it} \times StockHolder_{it}$	4.85 (2.31)**
$HighDividend_{it}$	25.09 (11.77)***
$StockHolder_{it}$	4.02 (3.06)***
$ProbabilityEstimate_{it-1}$	0.53 (7.62)***
$GainCondition_{it}$	2.03 (2.21)**
$LowVarianceCondition_{it}$	0.45 (0.95)
Subject Fixed Effects	Yes
R^2	0.67
Observations	4395

Table IV: The degree to which participants make errors in assessing the stock when confronted with dividend information that contradicts their prior choice is predicted by the engagement of the vmPFC and the left ventral striatum during the presentation of that information (no other regions have significant effects, results omitted here for brevity). The data included in this analysis refers to instances when participants chose the bond at the beginning of a trial, and then observed that the stock paid a high dividend that trial. The three regression specifications in the table indicate that higher activation in the vmPFC, left ventral striatum, and their first principal component, all measured at the time of the dividend presentation in the trial, leads to a smaller probability estimation error at the end of the trial. The dependent variable, *Probability Estimation Error_{it}*, is defined as the absolute value of the difference between the estimate provided by subject *i* in trial *t* regarding the probability that the stock is paying dividends from the good distribution, and the objective Bayesian value of that probability. Standard errors are robust to heteroskedasticity and are clustered at the subject level. The regressions in the three specifications in the table include fixed-effects for each subject, as well as for each level of objective probability, and controls for the experimental condition faced by participant *i* in trial *t* (i.e., gain vs. loss, low vs. high variance conditions). *t*-statistics are reported in parentheses. * and ** denote statistical significance at $p < 0.1$ and $p < 0.05$, respectively.

Dependent variable	<i>Probability Estimation Error_{it}</i>		
<i>vmPFC_{it}</i> at dividend presentation	-1.03 (-1.96)*		
<i>vSTR_{it}</i> at dividend presentation	-1.28 (-1.97)*		
1 st principal component of <i>vmPFC_{it}</i> and <i>vSTR_{it}</i> at dividend presentation	-0.84 (-2.39)**		
Condition Fixed Effects	Yes	Yes	Yes
Objective Probability Fixed Effects	Yes	Yes	Yes
Subject Fixed Effects	Yes	Yes	Yes
R^2	0.38	0.38	0.38
Observations	1014	1014	1014

Appendix

Brain imaging data acquisition and analysis

fMRI session and financial decision making task To measure brain activation during financial learning we used a task that requires subjects to update their beliefs about a stock’s dividend distribution on a trial-by-trial basis in order to make optimal choices, as in Kuhnen (forthcoming). Instructions for the task were presented to the subjects as a standardized slide presentation at a desktop computer. The subjects then proceeded with a training session for the task. The design and timing of the training session were identical to the task used during scanning, but the training was shortened to one gain and one loss block (order randomized across subjects). Next, the subject was placed in the MR scanner and accustomed with the choice button device (four buttons, one for each thumb and index finger). The subject viewed the experimental screen over video goggles that were adjusted to the subject’s sight (subjects with an ametropia of more than ± 5 dpt were excluded).

MRI data acquisition. All MRI sessions were run on a Siemens Trio 3.0 T scanner with a standard eight-channel head coil. Scan sessions started off with a localizer scan followed by a structural scan that included T1-weighted images (TR, 1570 ms; TE, 3.42 ms; flip angle, 15; 1 mm slices). While subjects played the financial learning task, T2*-weighted echoplanar images (EPIs) were collected (TR, 2500 ms; TE, 30 ms; flip angle, 90; 37 3 mm slices in ascending order; field of view, 192 mm; voxel size, 3 x 3 x 3.3 mm; approx. 840 volumes). The task was implemented in Presentation (Neurobehavioral Systems; www.neurobs.com).

fMRI data preprocessing Preprocessing of the functional images was implemented in the MATLAB (MathWorks) based software Statistical Parametric Mapping 8 (SPM8, version r5236). It included realignment, normalization on MNI standard (Evans et al. (1993)) using SPM8’s optimized segmentation of the T1 image and the mean realigned EPI into gray and white matter, cerebrospinal fluid, bone matter, soft tissue and air tissue classifications and the application of these deformations on the remaining EPI images, as well as spatial smoothing with an 8 mm full width half maximum Gauss kernel.

fMRI analysis The statistical fMRI analysis was also implemented in SPM8 (version r5236). For the first-level analysis we used a general linear model (GLM) which was estimated with SPM8’s canonical hemodynamic response function and included a high-pass filter of 128 Hz as well as correction for autocorrelations. SPM8’s internal masking threshold for the estimation of beta parameters was set to 0.4.

The first goal of the fMRI analysis was to identify brain regions that are sensitive to prior investment choices during the processing of new stock information. For this we computed a GLM that included the following events for every trial of the financial decision making task: onset of the choice screen, onset of the stock outcome presentation, onset of the estimation screen, and onset of the accumulated payoff presentation. Each event was modeled by four onset regressors as stick functions. The onset of the choice screen was modeled by regressors for the four conditions: (1) gain & low variance, (2) gain & high variance, (3) loss & low variance, (4) loss & high variance. The onset of the stock outcome presentation was modeled by indicator variables for the subject’s prior investment choice and the presented stock outcome, creating four onset regressors: (1) stock choice & high dividend, (2) bond choice & high dividend, (3) stock choice & low dividend, (4) bond choice & low dividend. Because we were primarily interested in the effect of prior investment decisions on the BOLD signal when new stock information was presented, these four onset regressors were defined as regressors of interest. To ensure that the BOLD signal was not simply driven by changes in the objective Bayesian probability that the stock is good we added this variable as a parametric modulator of the regressors of interest. The onsets of the estimation screen and the accumulated payoff presentation were modeled by the same indicator variables, but without parametric modulations. Together with six motion parameters, they were added to the GLM as nuisance regressors.

For each subject, we computed contrast images for the four regressors of interest and tested them in a two-by-two ANOVA at the group level with a within-subject factor for prior investment (stock and bond) and a within-subject factor for stock outcome (high and low dividend). The ANOVA was designed as a full factorial model with dependent measurements. We tested the positive interaction of the two factors and applied whole-brain correction for multiple comparisons based on familywise error (FWE) control. We report results that survive a FWE-corrected threshold of $p < 0.05$ at the peak or cluster level in Table A2. We expected activation in brain areas associated with outcome and reward encoding, such as the vmPFC, the ventral striatum, and the PCC (e.g. Bartra (2013)). To illustrate the results we extracted the mean beta parameters for each subject and each regressor of interest from anatomically defined ROI masks for the

vmPFC, the bilateral striatum, and the PCC derived from the Automated Anatomical Labeling (AAL) atlas (Maldjian et al. (2003), Tzourio-Mazoyer et al. (2002)) using the MarsBaR toolbox (Brett et al. (2002)). We then plotted the mean and standard error of the beta parameters across all subjects against the four regressors of interest in Figure 5.

The second goal of the fMRI analysis was to investigate the effect of prior investments on the processing of new stock information on a trial-by-trial basis. To analyze the fMRI time courses, we first defined Volumes of Interest (VOIs) in the vmPFC and the bilateral ventral striatum based on a whole-brain meta-analysis of subjective value effects at the decision stage as reported by Bartra et al. (2013) (their figure 6A).

For each subject and each VOI, we extracted the first BOLD eigenvariate from the preprocessed image files. As part of the extraction each BOLD eigenvariate was whitened, high-pass filtered, and corrected for confounds like scanner drifts. From the BOLD eigenvariates we selected, trial by trial, the values that corresponded to the expected peak BOLD responses of the same four events as estimated in the GLM: onset of the choice screen, onset of the estimation screen, onset of the accumulated payoff presentation, and onset of the stock outcome presentation. We used the latter in the prediction analysis in Table IV.

References

- Bartra, O., McGuire, J. T. & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, 76 (0), 412427.
- Brett, M., Anton, J.-L., Valabregue, R., & Poline, J.-B. (2002). Region of interest analysis using an SPM toolbox. Presented at the 8th International Conference on Functional Mapping of the Human Brain in Sendai, Japan. Available on CD-ROM in *NeuroImage*, 16(2).
- Evans AC, Collins DL, Mills SR, Brown ED, Kelly RL, Peters TM (1993) 3D statistical neuroanatomical models from 305 MRI volumes. In: IEEE Conference Record Nuclear Science Symposium and Medical Imaging Conference, pp 18131817. IEEE.
- Kuhnen, C. M., forthcoming. Asymmetric Learning from Financial Information. *The Journal of Finance*.
- Maldjian, J.A., Laurienti, P.J., Kraft, R.A., Burdette, J.H., 2003. An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fmri data sets. *NeuroImage* 19, 1233 1239.
- Tzourio-Mazoyer, N. et al. (2002) Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage* 15, 273289

Table A1. Objective Bayesian Posterior Beliefs The table provides all possible values for the objectively correct Bayesian posterior that the stock is paying from the good dividend distribution, starting with a 50%-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 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%

Table A2. Brain regions positively correlated with the interaction effect of prior investments (stock or bond) and stock dividend (high or low dividend) at the time of stock outcome presentation. Results from the two-by-two ANOVA are shown. Height threshold, $t > 3.14$; extent threshold, $kE > 10$. Asterisks denote activations that survive whole-brain correction for multiple comparisons at $p < 0.05$ based on FWE control at the peak level (*) or at the cluster level (**). ACC, anterior cingulate cortex; vmPFC, ventromedial prefrontal cortex; dlPFC, dorsolateral prefrontal cortex; PCC, posterior cingulate cortex; IPL, inferior parietal lobule.

Region	Side	MNI coordinates			Cluster size k_E	Max stat t
		x	y	z		
ACC, vmPFC, dlPFC	L	-6	41	-4	1228	5.82*
ventral striatum	L	-12	2	-14	52	5.55*
ventral striatum	R	12	2	-14	49	4.80*
PCC	L	-9	-58	29	201	4.28**
IPL, postcentral gyrus	L	-54	-28	52	160	3.90**