

NBER WORKING PAPER SERIES

THE EFFECT OF PRIOR CHOICES ON EXPECTATIONS AND SUBSEQUENT  
PORTFOLIO DECISIONS

Camelia M. Kuhnen  
Sarah Rudolf  
Bernd Weber

Working Paper 23438  
<http://www.nber.org/papers/w23438>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
May 2017

Funding for this work was provided by the German Research Council (DFG; We 4427/3-2), 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Camelia M. Kuhnen, Sarah Rudolf, and Bernd Weber. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Effect of Prior Choices on Expectations and Subsequent Portfolio Decisions  
Camelia M. Kuhnen, Sarah Rudolf, and Bernd Weber  
NBER Working Paper No. 23438  
May 2017  
JEL No. G02,G11

**ABSTRACT**

We document that prior portfolio choices influence investors' expectations about asset values, and their future choices. We find that people update more from information consistent with their prior choices, leading to sticky portfolios over time. These effects are related to how the brain's valuation centers encode new information about assets and about the trader's own success. These findings provide microfoundations for theoretical models where agents learn jointly about their skill and about asset values, leading to disagreement, and offer a common explanation for several puzzling investor behaviors, specifically, households' low stock market participation rate, and the disposition and repurchase effects.

Camelia M. Kuhnen  
University of North Carolina  
Kenan-Flagler Business School  
300 Kenan Center Drive, MC #4407  
Chapel Hill, NC 27599  
and NBER  
camelia\_kuhnen@kenan-flagler.unc.edu

Bernd Weber  
Center for Economics and Neuroscience  
Nachtigallenweg 86 53127 Bonn  
Germany  
bweber@lifeandbrain.com

Sarah Rudolf  
University of Bern, Dept. Psychology  
Fabrikstrasse 8  
CH-3012 Bern, Switzerland  
sarah.rudolf@psy.unibe.ch

# 1 Introduction

An implicit assumption in the finance literature is that market participants learn the same way from new information about available investments, irrespective of the composition of their portfolio. However, there are several puzzling aspects of investor behavior that raise questions about this assumption. People exhibit the disposition (e.g., Odean (1998)) and repurchase effects (e.g., Strahilevitz, Barber, and Odean (2011)), as they are reluctant to sell stocks that have not performed well, and to repurchase stocks that have increased in price after selling them. Moreover, the literature has documented surprisingly low levels of stock market participation, as the majority of people in the U.S. and Europe do not invest in equities (Campbell (2006), Calvet, Campbell, and Sodini (2007)).

Here, we propose and provide evidence for a common mechanism that can help explain these patterns in household financial decisions. Namely, we show that people’s prior investment choices shape how they *learn* from new information about investment options, and make subsequent portfolio decisions. We find that people’s expectations about assets are biased by whether their prior choices seem right ex-post, in that disconfirming evidence tends to be ignored, and because of this bias in updating beliefs about the available assets, portfolio choices are sticky over time. We document these behavioral effects using a controlled experimental setting, and we provide further evidence supporting this biased updating mechanism by examining participants’ brain activation patterns during their portfolio choices.

Theoretical work has shown that previous portfolio choices can influence investors’ *utility function*. These models include regret theory (Shefrin and Statman (1985), Barberis, Huang, and Santos (2001)), prospect theory (Barberis, Huang, and Thaler (2006)), as well as realization utility theory (Barberis and Xiong (2012)). Utility-based mechanisms can help explain the disposition effect, the repurchase effect and the low stock market participation puzzle, but to do so they need to be combined with narrow framing (Barberis and Huang (2008)). Experimental work has documented evidence in support of utility-based mechanisms. For example, the results in Frydman and Camerer (2016) show that regret theory

can help explain people’s propensity to exhibit the repurchase effect, and those in Frydman, Barberis, Camerer, Bossaerts, and Rangel (2014) indicate that realization utility can help explain the disposition effect. However, it is possible that prior portfolio choices might also change investors’ *beliefs* or the learning rules they use to incorporate financial market news.

The general idea that investors can be different in how they form expectations is at the heart of the large finance literature on investor disagreement and its effects on asset prices and trading volume. In a comprehensive summary of this literature, Hong and Stein (2007) point out that investors’ beliefs may differ not only because of access to different information sets or different priors, but also because of differential interpretation of the same information. In more recent work, Banerjee and Kremer (2010) show theoretically that disagreement in the interpretation of common information can impact trading volume and price volatility. Cookson and Niessner (2016) provide empirical evidence consistent with the idea that disagreement in investor expectations is driven by both different information sets as well as by differential interpretation of the same information, by studying opinions expressed by investors on a social media investing platform. As Cookson and Niessner (2016) note, an important missing piece in the literature on investor disagreement is why people differ in evaluating information. Our paper provides evidence for one mechanism that leads to investors being different when forming expectations about financial assets.

Specifically, we document that, at the behavioral and brain level, different people react to the same information differently because of their prior choices, as the outcomes of these provide signals about the person’s trading ability, and not just about the quality of the available investments. Our results provide microfoundations for models such as Daniel, Hirshleifer, and Subrahmanyam (1998) and Gervais and Odean (2001), where self-attribution bias causes investors to make different inferences about stocks, based on their beliefs regarding their ability as traders (i.e., beliefs about the precision of the signals they receive regarding the available assets). In these models, people are more inclined to attribute past successes, rather than past failures, to their ability. As a result, investors will become over-

confident about their talent, and will put too much weight on past successes when deciding their future choices. This prediction is in line with the recent empirical evidence in Campbell, Ramadorai, and Ranish (2014), who use detailed data on household investment choices and provide evidence that people repeat investment behaviors that happen to have performed well for them in the recent past. The authors propose that this might be due to people “feeling better” about themselves after receiving positive feedback about prior investment decisions, but the data do not allow them to investigate the exact mechanism driving households to repeat past winning actions. Hence, while the literature posits that previous choices may impact how people assess their opportunities going forward, no clear evidence for this assertion has yet been provided. Our goal in this paper is to investigate whether and how expectations and subsequent choices are influenced by prior actions.

We study the effect of prior choices on expectations and future portfolio decisions using an experimental setting with working-age, highly educated and high-earning adults faced with real financial incentives. Participants in the experiment made decisions regarding which one of two assets they wanted to hold in their portfolio at a given point in time. The two assets were a risk-free bond with a known payoff, and a stock that paid dividends from either a good or a bad distribution. The two dividend distributions had the same support, but the probability of high outcomes was higher if the stock was paying dividends according to the good distribution, rather than the bad one. Participants learned the stock dividend in each time period, or trial, no matter what portfolio choice they had made in that trial. Everybody started from a common prior that there was a 50% chance that the stock paid from the good dividend distribution, and then as trials went by, new dividend information was revealed and thus people’s expectations about the quality of the stock evolved over time. After 6 time periods, which constitute a learning block, the information about assets was refreshed, and participants had to learn again starting from the 50% prior that the stock was paying from the good dividend distribution. In total, each person completed 16 learning blocks of 6 trials each, for a total of 96 trials. Importantly, in this experiment it was optimal for participants

to learn objectively from all new outcomes, as their prior choices did not constrain them from changing their portfolio going forward.

We find that controlling for people's prior expectations about the quality of the stock, in trials when participants held the stock they updated more positively from observing a high dividend than in trials when they held the bond. However, if a low dividend was observed, when people held the bond they updated more negatively their views about the stock relative to situations when they held the stock. Importantly, these effects are identified within person and thus are not confounded by unobserved risk preferences or other fixed characteristics of the individual. As a result, if they held the stock, people ended up with more optimistic beliefs about the quality of the stock than if they held the bond, irrespective of the available objective information about the stock's dividends. This in turn induced stickiness over time in portfolio choices. After choosing the stock in one trial, people had a significantly higher chance to choose the stock again, rather than switch to the bond. For those who chose the bond, the opposite pattern was true: they had a significantly higher chance to hold the bond subsequently, and not switch to holding the stock. The random choice that people made in the first trial of each learning block, based solely on the 50% common prior for the probability that the stock paid dividends from the good distribution that block, was a significant predictor of people's subsequent choices, particularly in the next trial. This stickiness in choices over time was driven by the fact that when making portfolio decisions, people relied less on the new information about the stock if this information contradicted the prior choice, and relied more on their prior belief.

In support of our behavioral results we also provide brain imaging evidence that learning about asset payoffs depends on how these payoffs match with people's prior choices. That is, at the brain level we observe that there exists learning not just about the stock dividend but also about one's success in selecting investments. Brain areas previously found critical for learning and valuation encode more strongly news about the stock that matches people's prior choices, that is, they are more sensitive to stock outcome information in cases when this

information also provides a favorable signal about the person's ability to select investments. This differential activation at the moment when the stock dividend news is revealed can help predict, before people make their next choice, whether they will rely more on the newly observed dividend information when updating their expectations about the stock and when making their next portfolio choice.

These results add 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, Campbell, and Sodini (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 indicate 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 purchased, since these low outcomes do not line up ex-post with people's choice to invest in those stocks. Moreover, our evidence suggests that the repurchase effect (Strahilevitz, Barber, and Odean (2011)) can occur because investors do not update beliefs upwards sufficiently after observing high outcomes of stocks they no longer own, since these high outcomes contradict the earlier decision to sell those stocks. Our beliefs-related mechanism for the disposition and repurchase effects is complementary to the findings in Frydman, Barberis, Camerer, Bossaerts, and Rangel (2014) and Frydman and Camerer (2016), who find evidence that people's utility function is affected negatively by realizing losses and by missing possible gains.

Our results also provide microfoundations for models of investor behavior such as Daniel,

Hirshleifer, and Subrahmanyam (1998) and Gervais and Odean (2001), where people learn not just about asset payoffs but also about their own ability as traders. Investors in these models form incorrect beliefs about asset values because of their bias to update more from the success, rather than the failure, of their prior portfolio choice. These individual-level errors in learning can have aggregate outcomes. This general point was made by Daniel and Hirshleifer (2015), who note that differences among investors in terms of their beliefs about their ability can lead to high trading volume and predictability in asset prices. A specific implication of the biased learning mechanism we document here is that in a world with short-selling constraints, in the near term markets will over-react to good news and under-react to bad news, since market participants made the choice to own stocks and thus will be biased to update more from good news, rather than bad news, about these assets.

Our paper complements 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 (2015), Asparouhova, Bossaerts, Eguia, and Zame (2015)), and that strategic considerations regarding others' behavior influence learning and trading (Kogan (2009), Bruguier, Quartz, and Bossaerts (2010), and Carlin, Kogan, and Lowery (2013)). 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, and rely more on these news when making subsequent 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.<sup>1</sup>

---

<sup>1</sup>Outside the finance literature, existing work shows 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, Niederle, Niehaus, and Rosenblat (2014) provide experimental evidence that people update more from good news regarding their own beauty or intelligence, relative to bad news. Sharot, Korn, and Dolan (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.



The closest paper to our work is Kuhnen and Knutson (2011). While the primary focus of that paper was on the role of affect manipulations on financial decision making, the authors found that errors in updating priors about stocks were larger in situations where stockholders experienced a financial loss, or bondholders missed on a financial gain. While this evidence suggests that investment outcomes are parsed differently depending on how they line up with prior portfolio choices, it is confounded by the fact that realized gains and losses may impact learning differently, aside from any effect of prior choices and whether they are confirmed by new evidence. The fact that gains and losses shape updating rules differently was shown in Kuhnen (2015), where learning from negative payoffs was found to induce overly pessimistic expectations, relative to learning from positive payoffs – an effect completely orthogonal to that of prior decisions, which is the focus of our paper. Here we isolate the effect of prior choices on expectations formation from any effect related to the sign of outcomes that investors learn from, by having each subject learn separately from gains, and from losses, in different segments of the experiment. Moreover, to investigate the robustness of the effect of prior choices on learning across different types of environments, we change the variance of the distribution of dividends that subjects learn about, keeping the sign of the outcomes the same (i.e., keeping the outcomes in the gain domain, or in the loss domain). We also analyze the effect of prior investments on future choices, not just expectations, and document that people’s portfolios are sticky, and this is driven by the fact that new information is less relied upon when it disagrees with one’s prior investment. Furthermore, here we also document the existence of a biological mechanism, related to the functioning of valuation areas in the brain, that underlies the effect of prior choices on decision making, and show that this biological mechanism can offer microfoundations for theoretical models of individual trading that involve learning about assets at the same time as learning about one’s ability as an investor.

## 2 Experimental design

Participants completed a financial decision making task where they made 96 portfolio choices, split into 16 separate blocks of six trials each. In each of the 96 trials, each subject had to select one of two securities to hold in their portfolio for the duration of the trial: 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, no matter which asset they chose, participants observed the dividend paid by the stock, after making their asset pick, 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. This effect was shown in Kuhnen (2015) but it is independent from the effect of prior choices on learning and subsequent actions, which is the focus of this paper. The way we measure subjective expectations and portfolio choices in our experiment is similar to the procedure in Kuhnen (2015), but we study a completely different determinant of expectations formation relative to that studied in this prior work.

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 1).

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, par-

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 learning 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. The 16 learning blocks that each subject faced were split equally in gain vs. loss, and high vs. low variance conditions. The order of the blocks was randomized across subjects. Figure 1 presents two trial examples – one belonging to a learning block in the gain and low variance condition, and one belonging to a learning block in the loss and high variance condition.

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 regarding the accuracy of each subject’s probability estimates and the corresponding payment was only provided at the end of the task, after all learning blocks were completed. 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. Using this Bayesian posterior we can measure how close the subjects' expressed probability estimates are to the objectively correct beliefs, and how well the subjective updates in beliefs each trial track the objective updates based on the new dividend information.

The 46 participants in the study were all male, to avoid gender effects, with ages ranging from 29 to 49 years (mean 40.08 years, sd. 6.53 years) and 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.

## 3 Results

### 3.1 Expectations

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. This can be seen in Figure 2, where we plot the average propensity to choose the stock in a trial as a function of the subjective belief that the stock is paying from the good dividend distribution provided by participants at the end of the prior trial. This belief can either be below, equal to, or above 50%. Individuals who are risk-neutral and who maximize expected value would choose the bond in the first case, would randomize between

stock and bond in the second case, and would choose the stock in the third case.

We group the subjective probability estimates gathered from participants at the end of the prior trial in several categories, to examine whether the likelihood of the stock being chosen in the current trial indeed exhibits a significant jump when these estimates cross 50%. To simplify the exposition, we use five categories for the subjective estimates that the stock is good expressed at the end of the prior trial, namely  $[0\%,30\%)$ , 30%, 50%, 70%, and  $(70\%, 100]$ , but using a finer categorization scheme will not change the results. As shown in the figure, in situations when people had a 50% prior that the stock was good, in 47% of these cases these individuals decided to hold the stock in the next trial, and not the bond. This propensity to choose the stock is significantly higher, at  $p < 0.01$ , than that observed in situations where people's prior was either below 30%, or between 30% and 50% (in these two scenarios, the chance that people would choose the stock in the next trial was 22%, and 29%, respectively) and significantly lower, also at  $p < 0.01$ , than that observed in situations where people's prior was either between 50% and 70%, or above 70% (i.e., 72% and 75%, respectively). Hence, Figure 2 shows that subjective beliefs are significant drivers of subsequent portfolio choices.

However, it is important to establish that these expectations are meaningful predictors of asset choices across all experimental conditions, and after accounting for differences across subjects that can simultaneously influence a person's the degree of optimism about stocks, as well as one's portfolio choices. The probit models in Table 2 shows that, even after controlling for the type of experimental condition of each trial faced by participants, as well as for subject fixed effects, the chance that the stock, and not the bond, was chosen in a given trial is increasing in participants' estimate in the prior trial regarding the probability that the stock pays from the good dividend distribution.

Confirming the results in Figure 2, the estimates in Table 2 show that there exists a jump in the probability that the stock is chosen in a given trial, once the subjective belief that the stock is good crosses 50%: at this belief level, the chance that the participant picks

the stock, rather than the bond, is 27% to 31% ( $p < 0.01$ ) higher than if the belief was in the lowest category, namely, less than 30%. For subjective beliefs strictly above 50%, the increase in the chance that the stock is selected, relative to situations when subjective beliefs are in the lowest category, is between 48% and 51% ( $p < 0.01$ ). Gradually across the three specifications in the table, we add as control variables indicators for whether the particular block that the trial belongs to is in the gain, loss, high variance or low variance conditions, as well as subject fixed effects. The inclusion of these variables does not change the result that participants' propensity to choose the stock in a given trial increases significantly with these individuals' estimate for the probability that the stock is paying from the good distribution of dividends.

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, Oud, and Fehr (2014)). In the probit model in the first column of Table 2 we show that the probability estimates of the participants explain 12% of variation in the investment choice. Adding condition fixed effects (second column) does not increase the model fit, but adding subject fixed effects (third column) as explanatory variables for the decision to choose the stock and not the bond in a given trial, increases the  $R^2$  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, and how they make subsequent choices.

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 outcomes, even though in this task it is optimal for subjects to use objectively all information, as their prior choices do not constrain them from changing their portfolio going forward.<sup>2</sup>

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, given the same new information about the stock dividend. We assign the subjective beliefs about the stock being good, measured at the end of trial  $t - 1$ , into the same five categories listed above. We then split the observations based on whether in trial  $t$  the person decided to hold the stock or the bond (indicated in the figure by colors red and blue, respectively), and on whether the stock dividend revealed during trial  $t$ , after the choice was made, was high or low (indicated by solid lines, and dashed lines, respectively). In other words, for each level of subjective beliefs about the stock at the end of trial  $t - 1$ , we calculate the subjective beliefs expressed by participants at the end of trial  $t$ , in four scenarios: they chose the stock and the dividend in trial  $t$  was high, they chose the stock and the dividend in trial  $t$  was low, they chose the bond and the dividend in trial  $t$  was high, and they chose the bond and the dividend in trial  $t$  was low. Figure 3 shows that no matter what the subjective prior (in trial  $t - 1$ ) was, after a high dividend is revealed in trial  $t$ , participants who chose to hold the stock in that trial form significantly more optimistic estimates, by

---

<sup>2</sup>Individuals 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 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.

9.27% on average ( $p < 0.01$ ), for the probability that the stock is good, relative to those who chose the stock. When the dividend revealed in trial  $t$  is low, then participants who chose the bond form significantly more pessimistic estimates, by 3.27% on average ( $p < 0.01$ ), for the probability that the stock is paying from the good dividend distribution, relative to those who chose the stock that trial.

These results suggest that updating subjective priors about the quality of the stock depends on people's choice, and on the new information revealed. We examine this more closely by conducting the analysis shown in Figure 4. The procedure is similar to that used for the analysis in Figure 3, with one exception: the variable shown on the y-axis is not the subjective probability estimate expressed in trial  $t$ , but rather, the subjective update in people's own estimate between trials  $t - 1$  and  $t$ . Controlling, as before, for subjective priors in trial  $t$ , we find as expected that upon seeing a high dividend, people update their prior upward (i.e., the subjective update is a positive quantity), and upon seeing a low dividend, people update their prior downward (i.e., the subjective update is a negative quantity). The interesting finding here is that the size of the update depends on people's choice in that trial. When a high dividend is presented, those who chose the stock that trial have significantly more positive updates, by 9.07% ( $p < 0.01$ ) on average, in their estimates for the probability that the stock is good, relative to those who chose the bond. When a low dividend is presented, those who chose the bond that trial have significantly more negative updates, by 2.98% ( $p < 0.05$ ) on average, in their estimates for the probability that the stock is good, relative to those who chose the stock that trial.

The results in Figures 3 and 4 illustrate in a clear manner that expectations formation depends on how new information matches with people's prior choices. However, in this analysis we have not accounted for possible differences across participants in their propensity to choose a certain asset, or in their estimates about the quality of the stock. It is important to show that the same person, with the same prior about the stock dividend distribution, and seeing the same dividend news in a trial, will update differently from this news depending



on whether they held the stock or the bond that trial.

Moreover, if indeed we observe that prior choices change expectations within subjects, it is also important to understand whether this effect is confounded by the type of learning environment faced by the person: whether they are updating about a stock that has all positive outcomes (i.e., they are learning in the gain domain), or all negative outcomes (i.e., they are learning in the loss domain), or has high or low variance in the dividend distribution. Therefore, to assess the effect of the degree to which new information matches one's prior choice on the person's subsequent beliefs, we need to control for subjective priors, for subject fixed effects, and for the exact experimental condition faced by a person in that trial.

We do so in the linear regression models in Table 3. In the top panel, the dependent variable is the probability estimate that the stock is good provided by a subject in trial  $t$  of a learning block, and in the bottom panel, the dependent variable is the subjective update in this belief between trial  $t - 1$  and  $t$ . The independent variables of interest are interaction terms between indicators for whether the person chose the stock or the bond in trial  $t$ , and whether the dividend observed in that trial was high or low.

As can be seen in the first regression model in the top panel, when the dividend is low, we find that, all else equal, if a person chose the bond in the beginning of that trial, their estimate for the probability that the stock is good is 4.25% ( $p < 0.01$ ) lower than if they chose the stock. When the dividend shown in trial  $t$  is high, all else equal, if a person chose the stock in the beginning of that trial, their estimate for the probability that the stock is good is 9.09% ( $p < 0.01$ ) higher than if they chose the bond. This effect is obtained by calculating the difference between the coefficients on the indicators *Stock holder x High dividend* and *Bond holder x High dividend*, namely 34.39% and 25.30%. These two coefficients separately show how much higher subjective probabilities are in the two situations the indicators capture, relative to the omitted indicator in the regression, namely *Bond holder x Low dividend*.

In columns 2-5 of Table 3 we estimate the effects of the interaction terms between the prior choice and the type of new dividend revealed on subjective beliefs separately for each

of the four experimental conditions that participants faced. Whether subjects have to learn from or experience solely positive stock outcomes, or solely negative stock outcomes, or face dividend distributions characterized by low or high variance, we find that the extent to which new information matches prior choices has a significant impact on how people form expectations about the stock – similar in magnitude and statistical significance to the results found in the main specification in the first column in the table. In other words, these results indicate that the effect of prior choices on how people use new information about assets to form their expectations about these assets going forward is robust to the type of learning environment faced by investors.

The results in the first regression model in the bottom panel of Table 3 show that upon observing a low dividend, the update in a person’s belief that the stock is paying from the good dividend distribution is 3.89% lower if the person chose the bond earlier in that trial, than if they chose the stock. Upon observing a high dividend, the update in subjective beliefs is 8.93% higher (i.e., 33.78%-24.85%) for stock holders relative to bond holders. These differences are significant at  $p < 0.01$ , and are obtained holding fixed the subjective prior belief of the individual. In other words, the results in the bottom panel of Table 3 show a differential reaction, starting from the same prior, by the same person, to news about the stock dividend, depending on whether the person in that trial held the stock or the bond. The effects documented in the two panels of the table are similar to those shown in the simpler analyses in Figures 3 and 4, and demonstrate that people update their expectations about the quality of assets available to them differently depending on whether news about these assets seems to ex-post confirm these people’s prior investment decisions.

Importantly, this effect is different from the classic confirmation bias (Myanatt, Doherty, and Tweney (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.

## 3.2 Subsequent portfolio choices

We now investigate whether prior choices influence participants' portfolio decisions going forward. So far, we have documented the people's expressed beliefs about the quality of the stock depend on whether the new information learned about the stock matches their prior choice. However, it is important to show that the expressed beliefs are truthful, in the sense that they drive people's choices – that is, we need to document that participants' asset choices going forward are biased by whether past choices were confirmed by new dividend information.

If choices were to be biased in this manner, we should observe that people persevere with the same choice from trial to trial, since they do not incorporate disconfirming evidence as readily as confirming outcomes. In other words, we should see that, controlling for people's prior beliefs before they made a choice in trial  $t$ , and for the new information revealed in trial  $t$ , the choice made in trial  $t + 1$  will be positively correlated with the choice made in trial  $t$ . The analysis in the probit models in Table 4 confirms this prediction. The dependent variable is an indicator equal to one if the person chose the stock in trial  $t + 1$ , and zero if they chose the bond. The independent variables of interest are interaction terms between indicators for whether the individual chose the stock or the bond in trial  $t$  and indicators for whether the new stock dividend revealed in trial  $t$  was high or low. The controls include fixed effects for subjective priors at the end of trial  $t - 1$ , for the experimental condition encountered by the subject in the learning block to which the trial belongs (i.e., gain, loss, high or low variance in the dividend distribution), and subject fixed effects. We examine the relationship between the prior and the subsequent choice, for each level of dividend revealed in trial  $t$  (i.e., high or low), for the entire set of observations (in column 1), and separately for each experimental condition (columns 2-5). The effects we obtain are very similar across all these subsets of the data, therefore we here report the estimates obtained using all trials, which are those in column 1 of Table 4.

We find that if the dividend in trial  $t$  is low, a person who chose the stock in trial  $t$

has a 32% ( $p < 0.01$ ) higher chance to choose the stock again in trial  $t + 1$ , relative to a person who chose the bond in trial  $t$ . A similar effect is observed after a high dividend is revealed: those who held the stock in trial  $t$  have a 21% ( $p < 0.01$ ) higher chance to hold the stock in the next trial, compared to people who in trial  $t$  held the bond. This estimate is obtained as the difference between the coefficients on the interaction terms *Stock holder & High dividend* and *Bond holder & High dividend*, i.e.,  $0.55 - 0.34$ , or 21%. In other words, stock holders tend to continue holding the stock, and bond holders tend to continue holding the bond, given the same prior beliefs about the stock, and the same new information about the stock's dividends.

One concern that needs to be addressed is that the stickiness of choices from trial  $t$  to trial  $t + 1$  is not driven by people reacting differently to dividend news in trial  $t$  based on whether this news matches their prior choice, but rather, it is driven by measurement error in the subjective probability estimates that people form before they make their choice in trial  $t$ . To alleviate this concern, we will examine whether the asset choice made by a person in the first trial of each of the 16 learning blocks they go through will influence which asset they choose for the rest of that block, controlling for the actual information observed during block. Subjects know that in the beginning of each learning block of 6 trials the learning problem is refreshed, and there is a 50% probability that they will face the good or the bad stock for the incoming 6 trials. All participants thus start with the same 50% prior, before any stock dividends are presented later on. Given this information, the expected payoff from either choosing the stock or the bond in that first trial of the learning block is exactly the same. A risk-neutral agent would therefore randomize which asset to hold during the first trial, and this is what we in fact see our subjects do. Hence, we can see whether that randomly made first choice in a learning block is a predictor of what asset the person will choose in the remaining 5 trials in the block, controlling for the dividend information that the person is shown later during the block. Since we conduct this analysis by using within-subject variation in their first choice in each learning block, the results are not driven

by unobserved subject fixed effects such as risk attitudes.

As expected, we find that subsequent choices in a learning block depend on the random asset choice the person made in trial 1 of the 6 trials in the block. This result is shown in the probit models in Table 5. The dependent variable is an indicator equal to 1 if the subject chose the stock in trial  $t \in 2, 3, 4, 5, 6$  of a certain learning block, and 0 if they chose the bond. The independent variable of interest is an indicator equal to 1 if in trial 1 of that same learning block the subject chose the stock, and 0 if they chose the bond. We include experimental condition fixed effects, subject fixed effects, and fixed effects for the objective probability that the stock is paying from the good dividend distribution, based on the dividends revealed up to and including trial  $t - 1$  in that learning block. These fixed effects capture the information that the person has seen about the stock and should drive, for a Bayesian learner, their choice in trial  $t$ . However, as the results in the table show, even controlling for what objectively was revealed to the subject about the stock up to trial  $t$ , the asset choice made at random by the individual in the first trial of the block has a significant and positive effect on the asset choice made in trial  $t$ . If the person chose the stock in the first trial, then during trials 2 through 6 of that learning block they have a 10% ( $p < 0.01$ ) higher chance to chose the stock. We study how long-lived the effect of the first choice is, by examining separately in the last three columns in Table 5 the choices made in trial 2, trials 3 and 4, and trials 5 and 6, respectively. We find that the effect of the choice made in the first trial on subsequent choices is particularly strong – namely, 24% ( $p < 0.01$ ) – for choices in trial 2, drops to 10% ( $p < 0.05$ ) for trials 3 and 4, and then is not significantly different from zero in trials 5 and 6. In other words, first choices induce stickiness in subsequent choices, but the effect diminishes over time.

While the stickiness in choices is consistent with the learning mechanism we have in mind, it is necessary that we document in a more direct manner that upon making an asset choice, people discount new evidence that does not match their prior investment decision. We shows this graphically in a population-based analysis in Figure 5, and in probit models

using within-subject variation in Table 6.

Our goal for the analysis in Figure 5 was to examine whether people's propensity to choose the stock, rather than the bond, is not influenced as much if they observe a stock outcome that does not match their choice in the prior trial, than if they observe an outcome that matches their prior choice. Consider situations where people said at the end of trial  $t - 1$  that they think that there is a 50% probability that the stock is paying from the good dividend distribution. From the analysis in Figure 2 we know that in these situations, 47% of the time people choose the stock. We find that among people who chose the stock in trial  $t$  and had the 50% prior right before making this choice, 84% of them chose the stock again in trial  $t + 1$  if a high dividend was presented in trial  $t$ , and 46% of them chose the stock in trial  $t + 1$  if a low dividend was presented in trial  $t$ . Therefore, relative to the baseline propensity of choosing the stock of 47% prior to seeing the new information, the propensity to choose the stock does not change much (i.e., drops by 1%) if those who chose the stock see a low dividend, but it increases dramatically (i.e., by 84%-47%, or 37%) if they see a high dividend. However, among people who chose the bond in trial  $t$  and had the 50% prior right before making this choice, 59% of them chose the stock in trial  $t + 1$  if a high dividend was presented in trial  $t$ , and 14% of them chose the stock in trial  $t + 1$  if a low dividend was presented in trial  $t$ . Therefore, relative to its 47% baseline value, for bond holders the propensity to choose the stock drops dramatically (i.e., by 47%-14%, or 33%) if they see a low dividend, but does not change as much (i.e., it increases by 59%-47%, or 12%) if they see a high dividend. In the figure we show these changes in people's propensity to choose the stock, based on whether the new information matches their prior choice, for each category of subjective estimates of the probability that the stock is good, not just for the case of the 50% prior, and find similar patterns across all categories of subjective priors. Namely, the people whose propensity to choose the stock is increased most by seeing high dividends are those who previously chose the stock, and the people whose propensity to choose the stock is decreased most by seeing a low dividend are those who previously chose the bond.

Another strategy for investigating whether subsequent choices depend differentially on the new dividend information based on whether it matches the prior choice is to estimate a within-subject model predicting a person's choice in trial  $t$ , where we examine the dependence of this choice on the information revealed in trial  $t$  and on the subjective belief about the quality of the stock that the person expressed at the end of trial  $t - 1$ , right before making their choice for trial  $t$ . If people tend to ignore information that disconfirms their prior choice, then when making the asset choice in trial  $t + 1$ , they will rely more on the new dividend information in trial  $t$  if that confirms their choice in that trial, than when it disconfirms it. In case of disconfirming evidence, as opposed to confirming outcomes, people's subsequent choice will rely more on the subjective beliefs about the stock that they had at the end of trial  $t - 1$ , before the new evidence was presented in trial  $t$ . In other words, when the evidence does not match their choice, people will rely more on their old beliefs, as if there was no information revealed in trial  $t$ , when making their portfolio choice in trial  $t + 1$ . When, instead, the dividend news in trial  $t$  matches their choice in trial  $t$ , people will rely more on this news when making their next portfolio choice.

The analysis in Table 6 confirms these predictions. The dependent variable in the probit models in the table is an indicator equal to 1 if the person chose the stock in trial  $t + 1$  of a learning block, and 0 if they chose the bond. The models include condition and subject fixed effects, as in prior analyses. In the first two columns we split the data based on whether in trial  $t$  the dividend presented matched the choice made earlier in that trial by the person. We define a match as occurring if the person chose the stock in trial  $t$  and then was informed that the stock outcome that trial was high, or if the person chose the bond in trial  $t$  and then was informed that the stock outcome was low. In the case of matching trials (first column), we find that there is a positive and significant correlation between the dividend shown in trial  $t$  and the person's likelihood of choosing the stock in trial  $t + 1$ . Specifically, if the dividend was high, all else equal the person would have a 59% higher chance of choosing the stock on the next trial. However, looking at trials where the dividend did not match

the person’s choice (second column), we observe no significant correlation between the newly revealed dividend and the person’s subsequent choice. The only significant predictor of the asset choice in trial  $t + 1$  in these cases is the person’s subjective estimate for the probability that the stock is good, as expressed before the trial  $t$  dividend is revealed. An increase in this probability of 1% leads to an increase of 0.6% in the chance that the person chooses the stock in trial  $t + 1$ . The impact of this old subjective prior on the choice in trial  $t + 1$  is double in the case when the new information does not line up with the person’s choice in trial  $t$ , relative to the case when the information matches the choice – namely, 0.6% vs. 0.3%, respectively. A chi-squared test shows that these effects are significantly different at  $p < 0.05$ .

To further examine whether people rely less on new information if it disagrees with their prior investment when making their next portfolio choice, we assess the degree to which the choice in trial  $t + 1$  is correlated with what the person’s subjective assessment about the quality of the stock ought to be, given their subjective estimate in trial  $t - 1$  and the new information in trial  $t$ , assuming that they update from this new information in an objective, Bayesian manner. It is easy to show that if a person’s prior that the stock is paying from the good distribution is some value  $p$ , the person’s posterior probability that the stock is good should be  $7p/(4p + 3)$  if the new dividend is high, and  $3p/(7 - 4p)$  if the new dividend is low. The less people update their priors in a Bayesian manner, the less correlated their choice in the next trial will be with this Bayesian subjective posterior probability. As shown in the last two columns in Table 6, while this subjective posterior — obtained by correctly updating a person’s own prior after seeing the new dividend in trial  $t$  — is correlated with the person’s choice in trial  $t + 1$ , this relationship is much weaker in cases when the dividend in trial  $t$  disconfirmed the person’s choice that trial. Specifically, we find that an increase in this Bayesian subjective posterior of 1% leads to an increase of 1.1% ( $p < 0.01$ ) in the probability that the person will choose the stock in trial  $t + 1$  if the news received in trial  $t$  matched their choice that trial, but it only corresponds to an increase of 0.4% in this probability if



the news received in trial  $t$  did not match their choice in trial  $t$ . A chi-squared test shows that these effects are significantly different at  $p < 0.01$ . In other words, people rely less on what should be their correctly updated beliefs about the quality of the stock when making their next portfolio decision, if they received information that does not line up with their previous decision.

### **3.3 Evidence on learning mechanism from brain imaging data**

The evidence shown so far indicates that the extent to which new information is in line with investors' prior choices will influence how these individuals update their beliefs about available assets, and which assets they will choose in the future. Specifically, we have found that subjective expectations and portfolio choices are sticky over time, as investors update more from information that matches, rather than disconfirms, their previous choices.

We interpret this evidence as providing support for models of investor behavior such as Daniel, Hirshleifer, and Subrahmanyam (1998) and Gervais and Odean (2001), where the degree to which people update from new information about financial assets depends on their beliefs about their own ability to get precise signals about asset values. In those models, the assumption is that people prefer to think that they have high ability as traders, and thus will preferentially update their beliefs about their skill after observing information consistent with their prior choices. In other words, this prior theoretical work assumes that learning about one's skill can interfere with learning about asset payoffs.

The behavioral evidence we have presented so far lines up well with this assumption. Ideally, however, it would be helpful to in fact identify whether at the neural level two distinct learning processes— one about a person's ability, and one about the quality of the available financial assets — take place simultaneously and interfere with each other. To accomplish this goal, we had the participants in our experiment perform the investment task while undergoing functional magnetic resonance imaging. The brain imaging question we need to answer is whether the areas of the brain that are involved in learning and valuation

have different activation patterns when the information revealed to the trader matches the person’s prior choice, relative to when it disconfirms it.

A large body of work in neuroscience and neuroeconomics has identified the brain areas that are critical for updating from new information and calculating the value of available actions. Specifically, the ventromedial prefrontal cortex (vmPFC) and the ventral striatum (vSTR) have been found to be the key regions of a subjective value network (e.g., Bartra, McGuire, and Kable (2013), Clithero and Rangel (2014)), responsible for the encoding of objective and subjective probabilities in choice valuation tasks (d’Acremont, Schultz, and Bossaerts (2013), Ting, Yu, Maloney, and Wu (2015)).

Therefore, our goal was to investigate whether at the time when new dividend information is released in a trial, activation in these two areas will depend on whether the new dividend is in line with the person’s choice in that trial. That is, does learning about asset payoffs, at the neural level, depend on whether these payoffs confirm or contradict the investor’s prior choice? If the answer is yes, then this would provide biological microfoundations for our behavioral effects documented earlier, and for the prior theoretical models that assume that updating about one’s skill as an investor can interfere with updating about asset payoffs.

We find that indeed, activation in the two key learning areas, the vmPFC and the vSTR, depends on both the type of dividend observed, as well as on whether that dividend confirms the person’s choice that trial. This result is shown in Figure 6, which pools data on brain activation and behavior from all subjects. As it is standard in these types of analyses, brain activation was measured as the blood-oxygen-level dependent (BOLD) signal recorded by the brain scanner continuously during the task, using standard imaging procedures (see the Appendix for more details).

On the left side of the figure we show – colored in yellow – the areas in the brain that we find to have significantly higher activation ( $p < 0.001$ ) at the time of the presentation of new stock dividend information if that information confirms the person’s choice made earlier in the trial. These areas are the vmPFC and the vSTR (left and right), as expected in light

of the prior literature on valuation. Activation here is therefore significantly greater when stock holders see a high stock dividend, rather than a low one, and when bond holders see a low stock dividend, rather than a high one.

On the right side of Figure 6 we investigate the distribution of these effects across subjects. For each participant, we calculate the average activation response in the areas identified in the group analysis above, in four types of situations, based on whether they were a stock or a bond holder that trial, and whether the stock dividend presented at the end of the trial was high or low. Estimates in red refer to observations from trials where the subjects were stock holders, and estimates in blue refer to observations from trials where the subjects were bond holders. We calculate the mean and the standard error of these brain responses across all participants. The results of this individual level brain activation analysis match those obtained using pooled data across all participants, namely, we observe that subjects have significantly higher activation in the learning and valuation areas of the brain when they are stock holders and see a high stock dividend, rather than a low one, and when they are bond holders and observe a low stock dividend, rather than a high one. In other words, learning about the payoff of the asset is influenced by whether or not the new stock outcome conveys good or bad information about the person's ability to make good investment choices.

Importantly, these brain-level effects showing the importance of prior choices on how people learn from new information are obtained even after we control for the probability (subjectively estimated, or the objective Bayesian value) that the stock is paying from the good distribution, given all the information observed about the stock prior to that trial. That is, controlling for prior beliefs, we show that updating from new information about the stock at the brain level is dependent on whether the choice recently made is confirmed or not by the new dividend. In the analysis, we also control for the payoff of the participant that trial, to make sure the effect we isolate is specifically about the new dividend matching or not the prior choice, rather than reflecting an asymmetry in responding to realized gains or losses, for example.

While analyzing the brain response to high versus low dividends provides a simple way to see whether prior choices influence the brain response to new information, we can use a finer analysis to shed light on this question. Specifically, we can calculate the objective Bayesian value of the update in the probability that the stock is paying from a good distribution, in any given trial. For example, if the prior probability that the stock is good was 50%, and now a high dividend is observed, this probability should jump to 70% (see Table A1 in the appendix for all possible values of this probability, and how they depend on how many high dividends were observed so far during a learning block). In this case, the objective update in the probability that the stock is good is an increase of 20%. If, however, the prior probability was 70% – which would be the case, for example, if the person observed a high dividend in the first trial in the block, and now they are in trial two – and a high dividend is seen in the current trial, then the probability should jump to 84.48%. Therefore, in this scenario, the objective update in the probability that the stock is the good one is an increase of 14.48%. For the person to learn correctly about the quality of the stock, their brain valuation areas need to do more than simply keep track of whether the dividend is high or low: namely, they would need to keep track of how informative this dividend is, which is captured by the value of the probability update.

Therefore, we investigated whether prior choices change the strength with which the brain encodes the value of the probability update that should be correctly inferred from the new dividend information in a given trial. The results of this analysis are presented in Figure 7. The top panel presents the results of a group analysis, where all the subjects' data are pooled together, and the goal is to identify where in the brain there exist a significant difference in activation in response to objective probability updates, based on whether the person was a stock holder or a bond holder. The areas shown in yellow are those where this difference is statistically significant ( $p < 0.001$ ). Unsurprisingly, the vmPFC and the vSTR (left and right) show this differential response to objective probability updates, based on the person's recent portfolio choice, in that they have increased activation when the update

matches their choice.

This effect is illustrated in more detail in the bottom panel of Figure 7. For each subject, we calculate the average response in each of the two areas identified in the group analysis above, separately for different values of the objective update, and for situations where the person was a stock holder that trial, or a bond holder. We calculate the mean response in the vmPFC and the vSTR for each of these situations, as well as the standard error of this response, across participants. We group the value of objective updates in four categories: -20%, -15% to -0.08%, 0.8% to 15%, and 20%. The number of observations in each category is roughly similar.<sup>3</sup> Estimates referring to observations when participants were stock holders are shown in red, and those referring to observations when participants were bond holders are shown in blue.

The patterns of activation in the vmPFC and the left and right vSTR are similar: when people hold the stock, the activation in these valuation areas is positive when the objective probability update regarding the quality of the stock is positive, and negative otherwise. When instead people hold the bond, the relationship between the activation in these areas and the objective update changes sign, such that they brain response is more negative when the objective update about the quality of the stock is more positive. These effects are identified controlling for the prior probability that the stock is good, and for the payoff each subject received in each trial, to isolate the relationship between activation in valuation areas and objective changes in expectations about asset values, without concerns for confounding effects of priors, or realized gains or losses.

Overall, the results in Figures 6 and 7 indicate that valuation areas in the brain exhibit a particularly strong reaction when new information about the stock is presented that matches the person's prior choice, that is, when the person is shown ex-post to have made the right investment that trial. However, to validate that these brain results indeed speak to the

---

<sup>3</sup>Updates of +20% or -20% are frequent because they occur in trials following an equal number of high and low dividends (e.g., one high and one low, or two high and two low dividends) when the objective probability is 50%. Upon seeing a new dividend, this probability will become either 70% or 30%, depending on whether the dividend is high or low, and thus the update will be either +20% or -20%.

learning channel we have in mind, we need to show that this bias in activation due to prior choices actually influences how people change their expectations, and their subsequent choices. That is, we need to show that the stronger is the activation in brain areas that keep track of one's ex-post success, the tighter is the connection between a person's subjective probability estimate and the new information, and between their next choice and the new information about the stock.

The analysis in Table 7 confirms this prediction. In the first column we estimate a linear regression where the dependent variable is the probability estimate produced by a subject at the end of a trial. We control for the person's prior belief (i.e., that expressed at the end of the previous trial), for condition effects, and include subject fixed effects also. As expected, all else equal, if the person observes a high dividend in the current trial, their estimate for the probability that the stock is good is on average 26.55% ( $p < 0.01$ ) higher than if they observe a low dividend. That is, unsurprisingly, there is a positive correlation between the person's expectation about the stock quality, and the size of the newly revealed dividend. The novel result in this regression is given by the interaction term between the indicator for whether the dividend is high, and an indicator for whether at the time when the dividend was presented, the subject exhibited high (i.e., above the median) activation in the vmPFC area identified in the analysis in Figure 6. The coefficient estimate on this interaction term is positive and significant (2.28,  $p < 0.01$ ). This means that subjective probability estimates about the stock quality are more tightly related to the new dividend information revealed in that trial (by about 10% of the baseline effect) when the person has a stronger brain response encoding whether the new dividend matches their prior choice.

In the second column the dependent variable is the update in people's subjective beliefs about the quality of the stock, from trial  $t - 1$  to  $t$ . The independent variables of interest are the value of the objective, Bayesian value of the update in the probability that the stock is good given the available dividend information, and an interaction of this objective update with the indicator for whether the subject exhibited high (i.e., above the median) activation

in the vmPFC area. As expected if subjects pay attention to the information revealed during the experiment, the subjective update in beliefs from trial to trial is positively and significantly correlated with the objective update – an increase in the latter of 1% leads to an increase in the former by 0.78%. However, the effect is larger, by 0.08% ( $p < 0.01$ ), or about a tenth of the baseline value, if the subject exhibited high (i.e., above the median) activation in the vmPFC area when the dividend news was presented during the trial.

In the last two columns of Table 7 we examine whether subsequent choices are more tightly dependent on the dividend information in trial  $t$  if there is more activation in the vmPFC area related to encoding whether one’s choice was right. We find evidence consistent with this prediction. Namely, people’s choices in trial  $t + 1$  are significantly more dependent, by about a quarter of the baseline effect ( $p < 0.05$ ), on the dividend learned in trial  $t$ , or on the objective update in trial  $t$  in the probability that the stock is good, if the brain signals more strongly that the person’s choice that trial was confirmed by the new evidence.

In other words, the results in Table 7 indicate that new information about the stock dividend distribution is relied upon more strongly when people update expectations about the quality of the stock, and when they decide their next investment choice, the more sensitive their valuation brain regions are to whether the person’s choice is ex-post shown to be right.

## 4 Conclusion

We document that prior investment decisions influence people’s ability to correctly form beliefs about the quality of financial assets, and their subsequent portfolio choices. Specifically, controlling for prior beliefs, we find that if participants’ most recent choice is a stock, rather than a safe bond, they will form more optimistic posterior beliefs about the dividend distribution of the stock. This occurs because stock holders update their beliefs about the stock significantly more after seeing a high dividend, compared to seeing a low one, and bond holders do the opposite, in that they update more from observing low dividends. Future asset

choices reflect these biased beliefs, as people rely less on recently observed information when making subsequent portfolio choices, if this information disconfirms their prior investment decisions. As a result, portfolio choices are sticky over time.

This behavioral asymmetry in learning and portfolio selection decisions that is induced by participants' prior investment choices is mirrored by an asymmetry in the response of brain areas important for learning and valuation, at the time when new dividend information is presented. Specifically, these areas encode more strongly information that confirms, rather than contradicts, people's prior choices. The extent to which the brain is sensitive to whether the news matches an individual's prior choice helps predict whether this new information will be relied on, rather than ignored, when the person updates their expectations about the quality of the stock, and makes their next portfolio choice.

Our evidence that people's ability to learn from news about assets depends on whether the news confirms these individuals' prior investment decisions provides a common explanation for several puzzling financial behaviors, namely, the disposition effect, the repurchase effect, and households' low stock market participation rate. Also, by documenting that updating from news differs based on people's prior portfolio choices, our results provide microfoundations for disagreement models where traders receive information that not only speaks to the quality of the assets in the market, but also, to their ability to value these assets.

An important question to be addressed in future research is under which conditions people will be able to learn and make choices free of any bias induced by how well new information ex-post justifies their prior decisions. These effects of prior choices may get weaker with enough experience, assuming that the environment is stable enough that people can assess their trading ability over time, and that this ability does not change. That being said, being proven right may fundamentally be something that people covet, and could be ingrained in our brains' functioning. In this case, our prior actions will have a bearing on our future view of the world and on our choices, financial or not, even after we accumulate significant objective information about our options, and our skills.



## References

- Asparouhova, E., P. Bossaerts, J. Eguia, and W. Zame (2015). Asset Pricing and Asymmetric Reasoning. *Journal of Political Economy* 123(1), 66–122.
- Banerjee, S. and I. Kremer (2010). Disagreement and Learning: Dynamic Patterns of Trade. *Journal of Finance* 65(4), 1269–1302.
- Barberis, N. and M. Huang (2008). *Handbook of the equity risk premium*, Chapter The Loss Aversion/Narrow Framing Approach to the Equity Premium Puzzle, pp. 199–236. Elsevier.
- Barberis, N., M. Huang, and T. Santos (2001). Prospect theory and asset prices. *The Quarterly Journal of Economics*.
- Barberis, N., M. Huang, and R. Thaler (2006). Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review*, 1069–1090.
- Barberis, N. and W. Xiong (2012). Realization utility. *Journal of Financial Economics* 104(2), 251–271.
- Bartra, O., J. T. McGuire, and J. W. Kable (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., S. R. Quartz, and P. Bossaerts (2010, October). Exploring the nature of "trader intuition". *Journal of Finance* 65(5), 1703–1723.
- Calvet, L. E., J. Y. Campbell, and P. Sodini (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.
- Campbell, J. Y., T. Ramadorai, and B. Ranish (2014). Getting better or feeling better? How equity investors respond to investment experience. *NBER Working Paper 20000*.

- Carlin, B., S. Kogan, and R. Lowery (2013, October). Trading complex assets. *Journal of Finance* 68(5), 1937–1960.
- Clithero, J. A. and A. Rangel (2014). Information parcellation of the network involved in the computation of subjective value. *SCAN* 9, 1289–1302.
- Cookson, J. A. and M. Niessner (2016). Why Dont We Agree? Evidence from a Social Network of Investors. *Working paper*.
- d’Acremont, M., W. Schultz, and P. Bossaerts (2013). The human brain encodes event frequencies while forming subjective beliefs. *Journal of Neuroscience* 33(26), 10887–10897.
- Daniel, K. and D. Hirshleifer (2015). Overconfident Investors, Predictable Returns, and Excessive Trading. *Journal of Economic Perspectives* 29(4), 61–88.
- Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998). Investor Psychology and Security Market Under- and Overreactions. *Journal of Finance* 53(6), 1839–1885.
- Eil, D. and J. Rao (2011). The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself. *American Economic Journal: Microeconomics* 3(2), 114–138.
- Frydman, C., N. Barberis, C. Camerer, P. Bossaerts, and A. Rangel (2014). Using Neural Data to Test a Theory of Investor Behavior: An Application to Realization Utility. *Journal of Finance* 69(2), 907–946.
- Frydman, C. and C. Camerer (2016). Neural Evidence of Regret and Its Implications for Investor Behavior. *Review of Financial Studies* 29(11), 3108–3139.
- Gervais, S. and T. Odean (2001). Learning to be overconfident. *Review of Financial Studies* 14(1), 1–27.
- Hong, H. and J. C. Stein (2007). Disagreement and the Stock Market. *Journal of Economic Perspectives* 21(2), 109–128.

- Kluger, B. D. and S. B. Wyatt (2004, June). 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., B. Oud, and E. Fehr (2014). Benefits of Neuroeconomic Modeling: New Policy Interventions and Predictors of Preference. *American Economic Review: Papers & Proceedings* 104(5), 501–506.
- Kuhnen, C. M. (2015). Asymmetric Learning from Financial Information. *Journal of Finance* 70(5), 2029–2062.
- Kuhnen, C. M. and B. Knutson (2011). The influence of affect on beliefs, preferences, and financial decisions. *Journal of Financial and Quantitative Analysis* 46(3), 605–626.
- Mayraz, G. (2013). Wishful thinking. *Working paper*.
- Mobius, M. M., M. Niederle, P. Niehaus, and T. S. Rosenblat (2014). Managing self-confidence. *Working paper*.
- Myanatt, C. R., M. E. Doherty, and R. D. Tweney (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, October). Are investors reluctant to realize their losses? *The Journal of Finance* 53(5).
- Payzan-LeNestour, E. and P. Bossaerts (2015). Learning About Unstable, Publicly Unobservable Payoffs. *Review of Financial Studies* 28(7), 1874–1913.
- Sharot, T., C. W. Korn, and R. J. Dolan (2011). How unrealistic optimism is maintained in the face of reality. *Nature Neuroscience* 14(11), 1475–1479.

- Shefrin, H. and M. Statman (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance* 40(3), 777–90.
- Strahilevitz, M. A., B. M. Barber, and T. Odean (2011). Once burned, twice shy: How naive learning, counterfactuals, and regret affect the repurchase of stocks previously sold. *Journal of Marketing Research* 48(SPL), 102–120.
- Ting, C.-C., C.-C. Yu, L. T. Maloney, and S.-W. Wu (2015). Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference. *Journal of Neuroscience* 35(4), 1792–1805.




Figure 1: Timeline of trials. At the beginning of each learning block, an indication for the gain or loss condition is displayed for 2s, followed by a brief pause of 1-3s. Then, the choice screen for the first trial in the block 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 asset for 0.75s, followed by an anticipation period during which a fixation cross is presented for 3.5-7s. Next, irrespective of the subject's choice, the stock dividend is shown for 3s. The subject is then asked to estimate the probability that the stock is paying from the good dividend distribution. The estimation is self-paced. Finally, the subject's accumulated balance so far is presented for 3s (not shown). After a brief pause of 1-3s, 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 trial in a block in the gain and low variance condition (upper panel) and one example of a trial in a block in the loss and a high variance condition (lower panel). For each block, the stock is randomly assigned to pay from a good or a bad dividend distribution. If the stock is paying from the good (bad) distribution, it pays the high outcome with 70% (30%) probability and the low outcome with 30% (70%) probability. To make investment choices yielding the highest expected payoff at the end of each trial, it is in subjects' best interest to correctly update their beliefs about the stock's dividend distribution after observing each dividend paid by the stock.

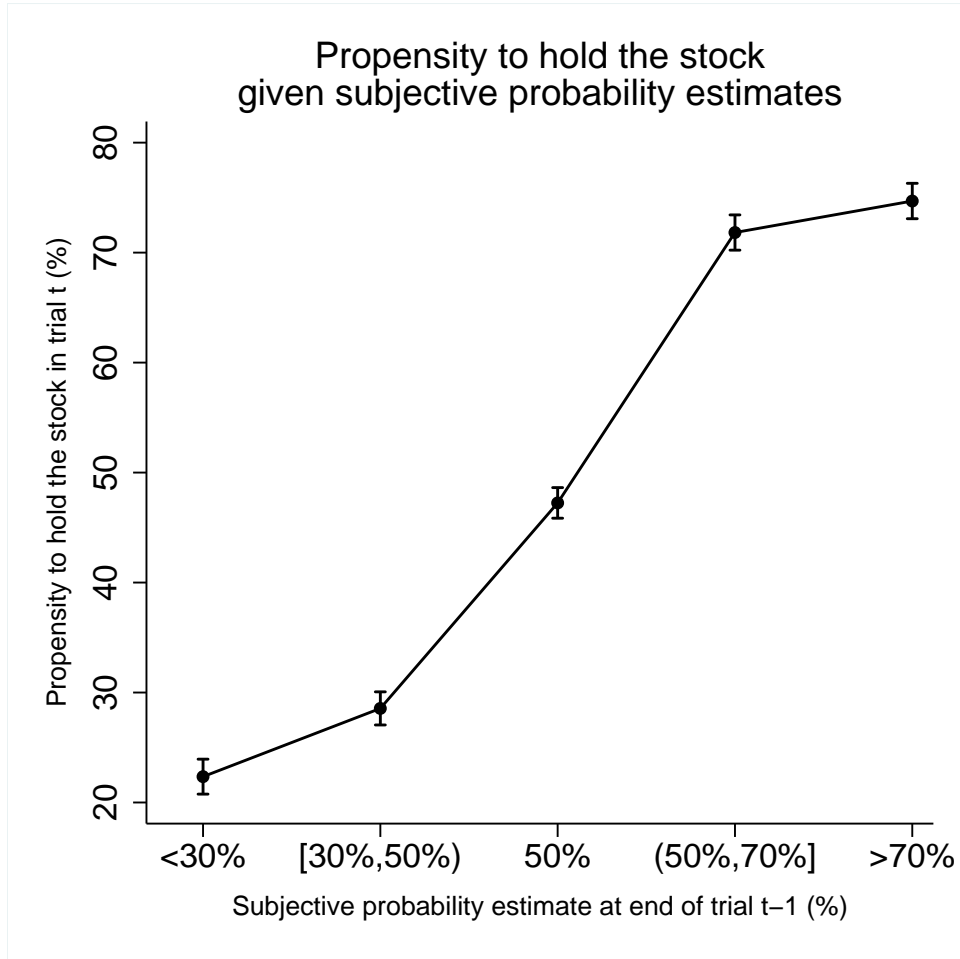


Figure 2: Portfolio choices depend on subjective beliefs. Data on whether the stock or the bond is chosen by each participant in a given trial is pooled across all subjects and split based on the value of the subjective estimates of the probability that the stock is good, expressed at the end of the prior trial. We group these values in the five categories listed on the x-axis. The frequency with which the stock is chosen for each of the values of subjective prior beliefs is shown on the y-axis, along with standard error bars.

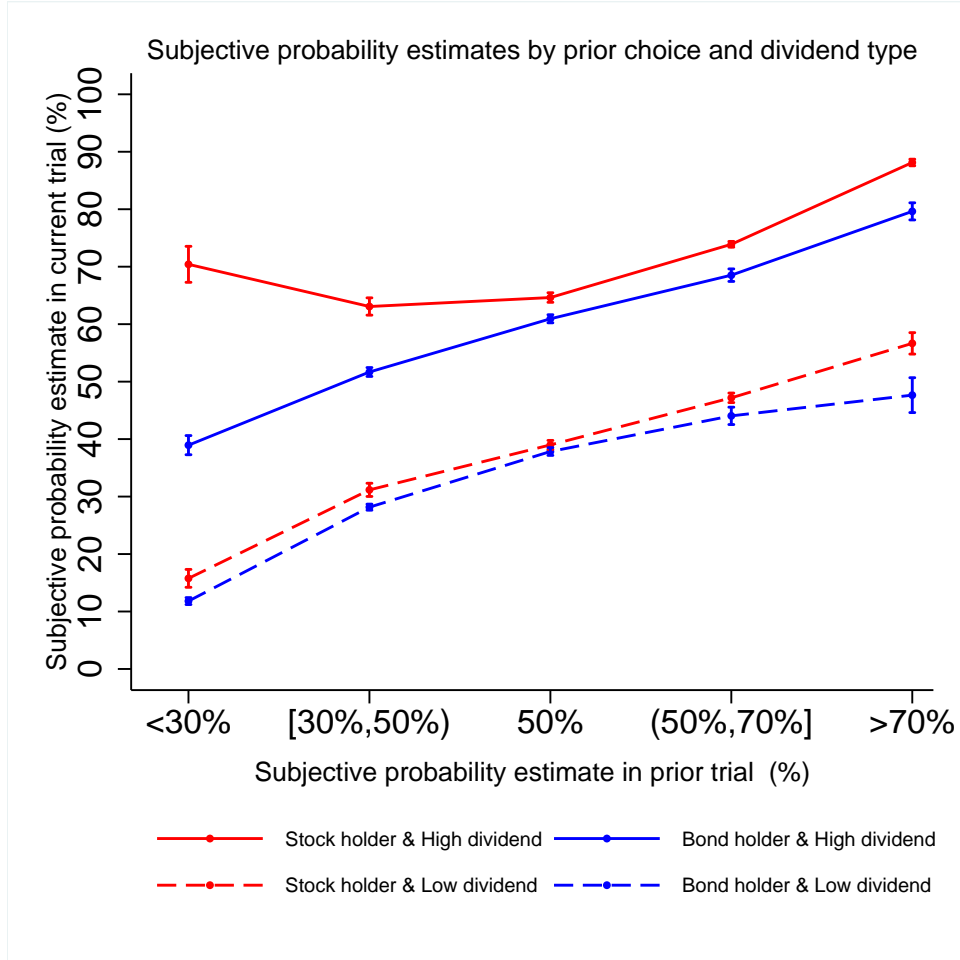


Figure 3: Subjective posteriors depend on prior choices and the new dividend. Data on the subjective probability estimate provided by each participant in a given trial is pooled across all subjects and split based on the value of the subjective estimates of the probability that the stock is good, expressed at the end of the prior trial, as well as on the choice made (stock or bond) and on the type of dividend observed (high or low) in the current trial. We group the values of subjective beliefs in the five categories listed on the x-axis. The average subjective estimates for the probability that the stock is good, expressed at the end of the current trial, is shown on the y-axis for each of the values of subjective prior beliefs, along with standard error bars.

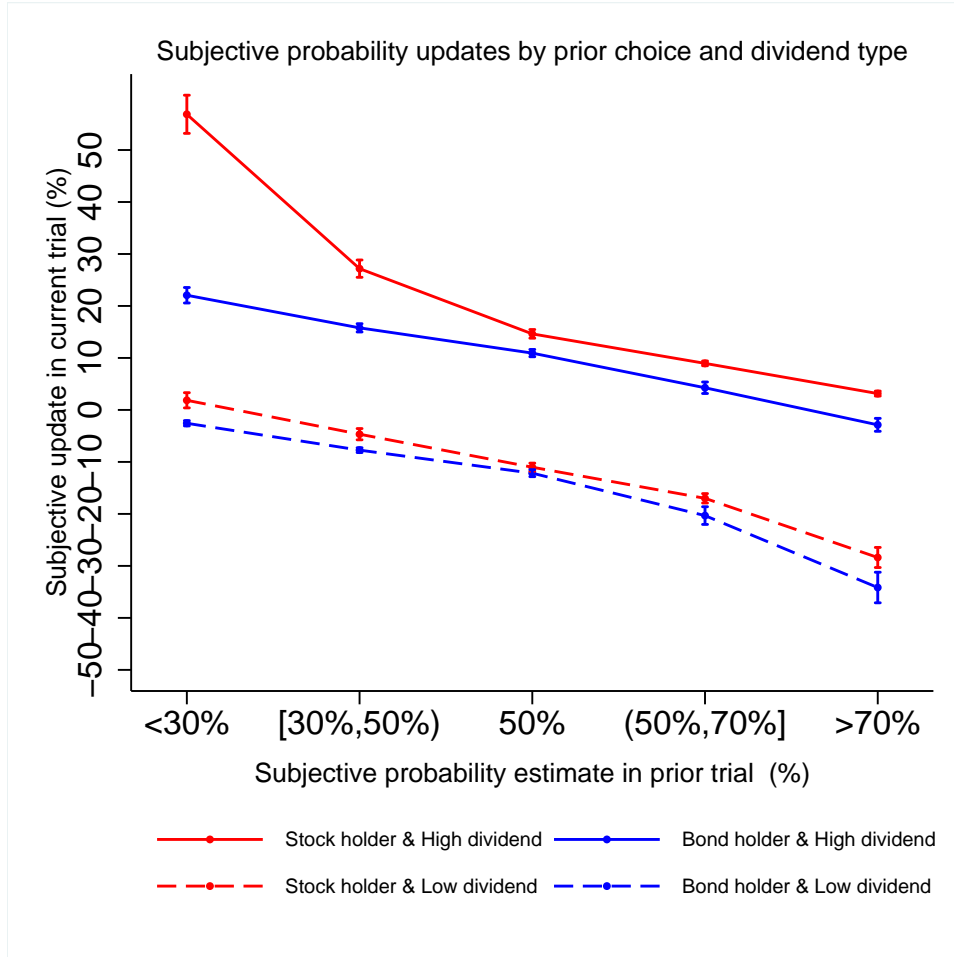


Figure 4: Subjective updates depend on prior choices and the new dividend. Data on the subjective update of the probability estimate provided by each participant from one trial to the next is pooled across all subjects and split based on the value of the subjective estimates of the probability that the stock is good, expressed at the end of the prior trial, as well as on the choice made (stock or bond) and on the type of dividend observed (high or low) in the current trial. We group the values of subjective beliefs in the five categories listed on the x-axis. The average subjective update in the probability that the stock is good, expressed at the end of the current trial, is shown on the y-axis for each of the values of subjective prior beliefs, along with standard error bars.



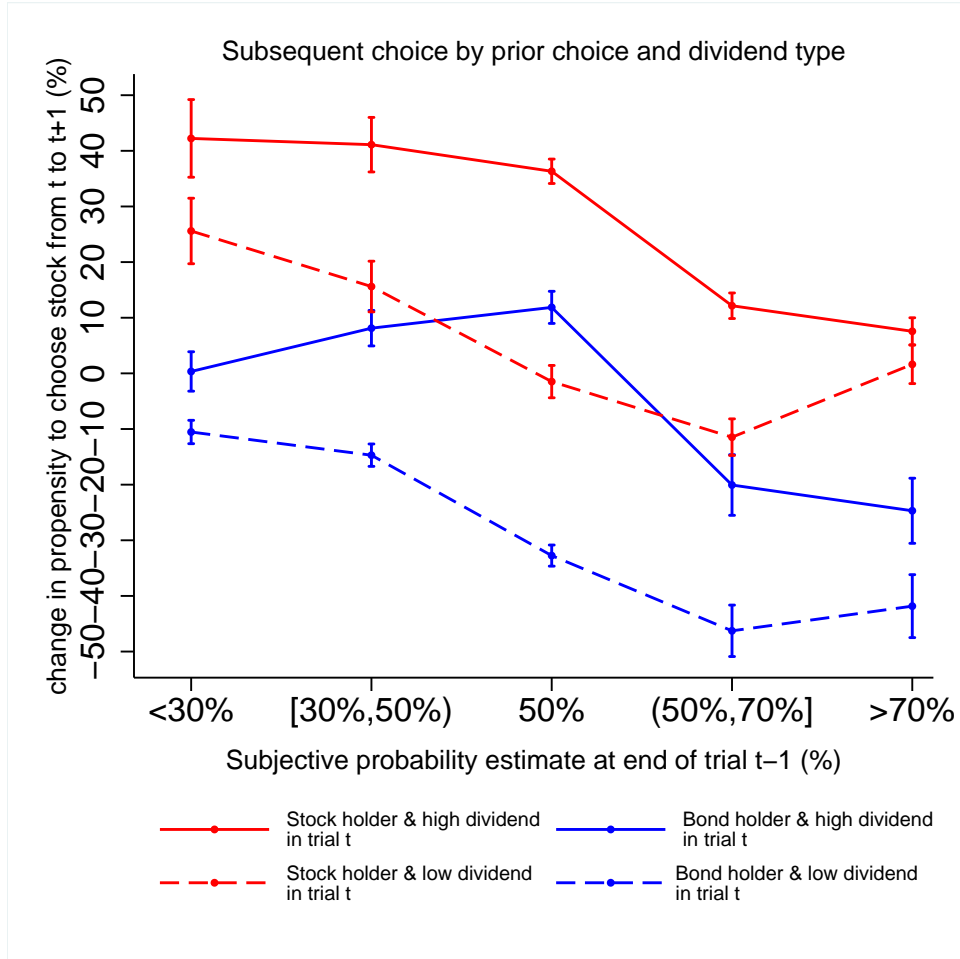


Figure 5: Portfolio choices are sticky. Data on whether the stock or the bond is chosen by each participant in a given trial is pooled across all subjects and split based on the value of the subjective estimates of the probability that the stock is good, expressed at the end of the prior trial, as well as on the choice made (stock or bond) and on the type of dividend observed (high or low) in the current trial. We group the values of subjective beliefs in the five categories listed on the x-axis. The change from the prior to the current trial in the propensity to choose the stock is shown on the y-axis for each of the values of subjective prior beliefs, along with standard error bars. For each subjective prior category, the baseline propensity to hold the stock is that shown in Figure 2.

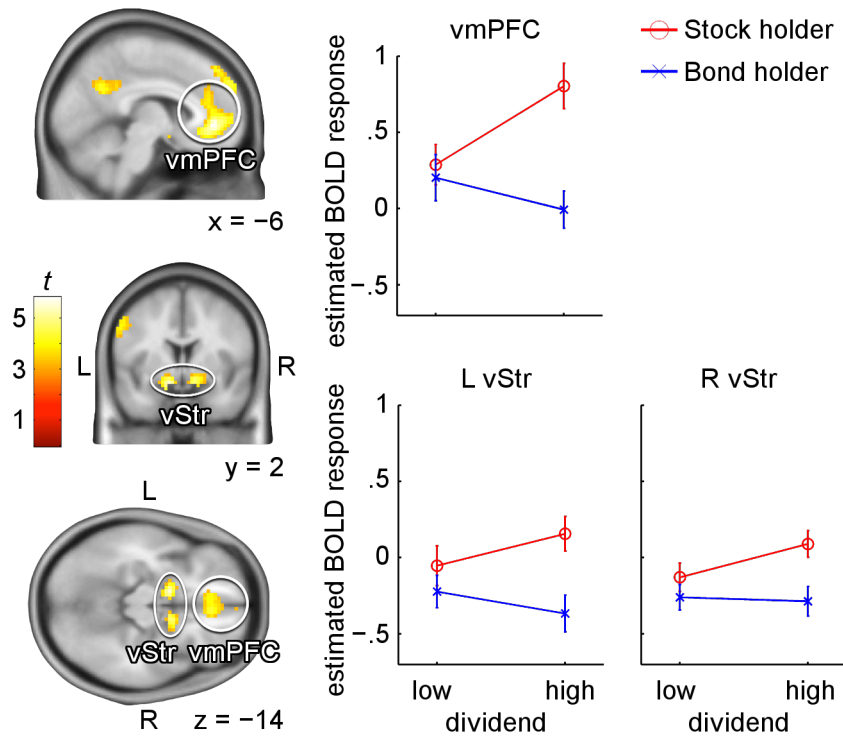


Figure 6: The brain activity at the time of stock outcome presentation reflects the interaction effect of prior investments (stock or bond holder) and stock dividend (high or low dividend). Activity in the ventromedial prefrontal cortex (vmPFC) and the left (L) and right (R) ventral Striatum (vSTR) is greater if the newly revealed dividend matches the prior asset choice. The individual subject regression models underlying these results control for the probability that the stock is good before the new stock outcome was presented and the received payoff in the current trial. More details regarding the brain imaging data analysis are available in the Appendix.

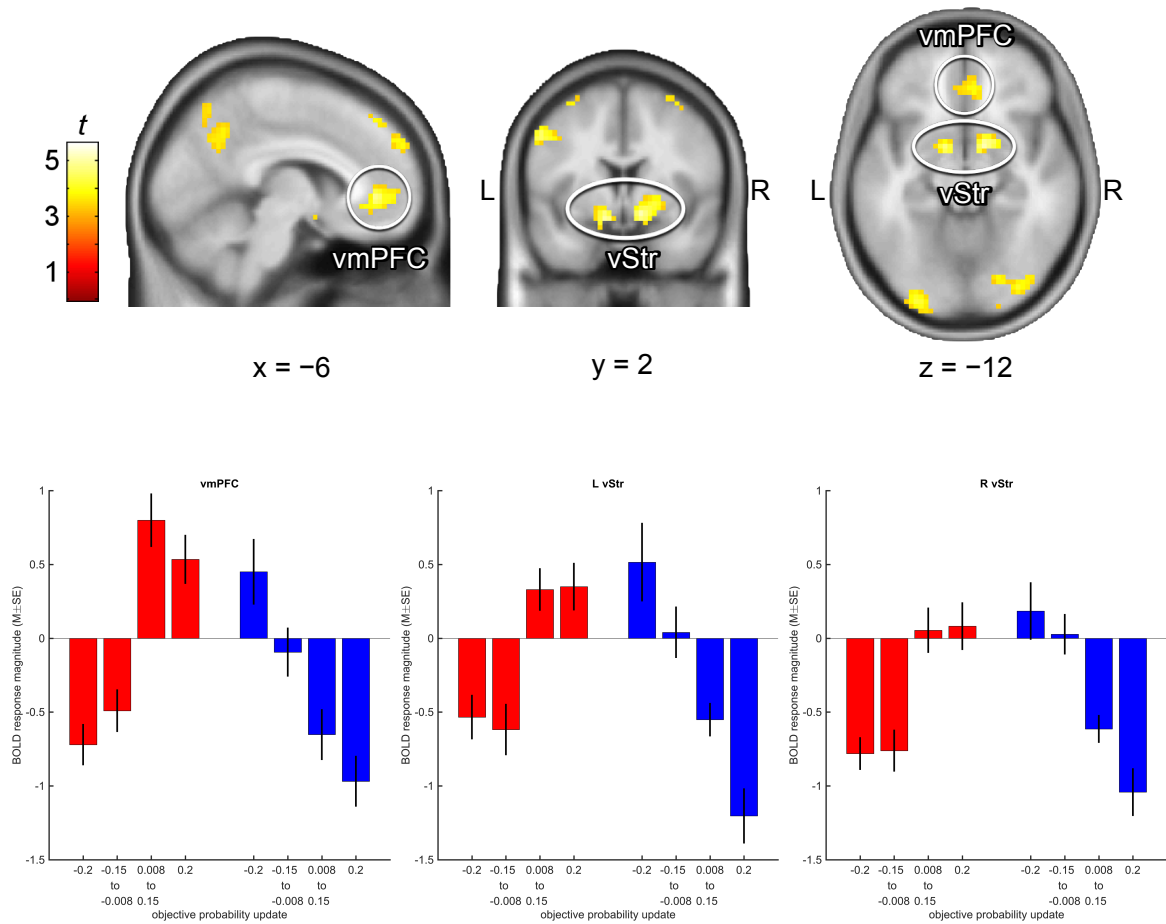


Figure 7: The brain encoding of objective probability updates differs by the prior investment choice. Activity in the ventromedial prefrontal cortex (vmPFC) and the left (L) and right (R) ventral Striatum (vSTR) is positively correlated with the objective update regarding the probability that the stock is paying from the good dividend distribution, given the newly revealed stock dividend in the trial, if the person holds the stock that trial. If the person holds the bond, activity in these areas is negatively correlated with the value of the objective update in the probability that the stock is good. The individual subject regression models underlying these results control for the probability that the stock is good before the new stock outcome was presented and the received payoff in the current trial. More details regarding the brain imaging data analysis are available in the Appendix.

Table 1: 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. These parameters are common knowledge for participants in the experiment. 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 2: Investment choices depend on subjective beliefs. The dependent variable in the probit models below,  $Stock\ holder_{i,b,t}$ , is equal to 1 if in trial  $t$  of learning block  $b$  participant  $i$  chose the stock, and not the bond. In the first column, we include as independent variables fixed effects for the levels of subjective estimates, expressed at the end of trial  $t - 1$ , of the probability that the stock is paying from the good dividend distribution. In columns 2 and 3 we add controls for the experimental condition faced by the subject in that learning block (i.e. gain or loss condition, low or high variance condition), and subject fixed effects. 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	$Stock\ holder_{i,b,t}$		
$ProbabilityEstimate_{i,b,t-1} \in [30\%, 50\%)$	0.08 (1.09)	0.07 (1.05)	0.12 (2.01)**
$ProbabilityEstimate_{i,b,t-1} = 50\%$	0.27 (4.26)***	0.27 (4.20)***	0.31 (5.75)***
$ProbabilityEstimate_{i,b,t-1} \in (50\%, 70\%]$	0.47 (7.87)***	0.46 (7.80)***	0.49 (10.02)***
$ProbabilityEstimate_{i,b,t-1} > 70\%$	0.48 (9.03)***	0.48 (8.77)***	0.51 (9.93)***
$Gain\ condition_{i,b,t}$		0.07 (2.39)**	0.08 (2.42)**
$Low\ variance\ condition_{i,b,t}$		0.01 (0.53)	0.01 (0.45)
Subject FEs	No	No	Yes
$R^2$	0.12	0.12	0.19
Observations	4395	4395	4395

Table 3: Subjective probability estimates, and subjective updates after new dividends depend on prior choices. The dependent variable in the linear regressions in the top panel in the table,  $ProbabilityEstimate_{i,b,t}$ , is the probability estimate produced by subject  $i$  at the end of trial  $t$  in learning block  $b$ . In the bottom panel of the table, the dependent variable is the update between trial  $t - 1$  and  $t$  in the person's subjective estimate for the probability that the stock is paying from the good distribution (i.e., it is  $ProbabilityEstimate_{i,b,t} - ProbabilityEstimate_{i,b,t-1}$ ). The independent variables of interest in both panels are the interaction terms between the person's choice that trial (indicator variables  $StockHolder_{i,b,t}$  or  $BondHolder_{i,b,t}$ ) and the level of the stock dividend revealed at the end of the trial (indicator variables  $HighDividend_{i,b,t}$  or  $LowDividend_{i,b,t}$ ). The reference (omitted) category of events refers to situations when the person was a bond holder and the stock paid a low dividend. The regressions include subject fixed effects and control for the experimental condition faced by each participant  $i$  in learning block  $b$  (i.e., gain vs. loss, low vs. high variance), as well as for the subject's probability estimate at the end of trial  $t - 1$  in that learning block. 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	$ProbabilityEstimate_{i,b,t}$				
	— Experimental condition —				
	All	Gain	Loss	Low variance	High variance
$Stock\ holder_{i,b,t} \ \& \ High\ dividend_{i,b,t}$	34.39 (11.10)***	36.64 (11.96)***	32.68 (9.94)***	32.30 (11.52)***	36.39 (10.37)***
$Stock\ holder_{i,b,t} \ \& \ Low\ dividend_{i,b,t}$	4.25 (2.95)***	4.88 (2.91)***	4.45 (2.29)**	3.97 (2.87)***	4.73 (2.82)***
$Bond\ holder_{i,b,t} \ \& \ High\ dividend_{i,b,t}$	25.30 (11.73)***	26.84 (12.36)***	23.79 (9.41)***	23.19 (10.70)***	27.40 (11.69)***
$ProbabilityEstimate_{i,b,t-1} \in [30\%, 50\%)$	12.52 (5.62)***	12.83 (4.91)***	9.79 (4.20)***	13.98 (6.53)***	10.64 (4.19)***
$ProbabilityEstimate_{i,b,t-1} = 50\%$	19.86 (7.21)***	19.61 (6.25)***	17.39 (6.28)***	21.81 (7.92)***	17.47 (5.76)***
$ProbabilityEstimate_{i,b,t-1} \in (50\%, 70\%]$	26.88 (7.85)***	25.57 (7.27)***	23.99 (6.66)***	28.81 (8.61)***	24.12 (6.27)***
$ProbabilityEstimate_{i,b,t-1} > 70\%$	37.00 (7.48)***	34.25 (6.60)***	35.22 (6.87)***	39.60 (8.50)***	33.30 (6.13)***
$Gain\ condition_{i,b,t}$	2.14 (2.14)**			1.82 (1.55)	2.49 (2.28)**
$Low\ variance\ condition_{i,b,t}$	0.16 (0.32)	0.01 (0.02)	0.44 (0.53)		
Subject FEs	Yes	Yes	Yes	Yes	Yes
$R^2$	0.649	0.688	0.632	0.674	0.630
Observations	4395	2205	2190	2199	2196

Dependent variable	$ProbabilityEstimate_{i,b,t} - ProbabilityEstimate_{i,b,t-1}$				
	— Experimental condition —				
	All	Gain	Loss	Low variance	High variance
$Stock\ holder_{i,b,t} \ \& \ High\ dividend_{i,b,t}$	33.78 (10.63)***	36.06 (11.47)***	32.21 (9.53)***	31.50 (11.22)***	36.05 (9.75)***
$Stock\ holder_{i,b,t} \ \& \ Low\ dividend_{i,b,t}$	3.89 (2.79)***	3.95 (2.72)***	4.82 (2.47)**	3.78 (2.89)***	4.15 (2.47)**
$Bond\ holder_{i,b,t} \ \& \ High\ dividend_{i,b,t}$	24.85 (11.85)***	26.48 (12.66)***	23.39 (9.51)***	23.28 (10.93)***	26.47 (11.72)***
FES for $ProbabilityEstimate_{i,b,t-1}$	Yes	Yes	Yes	Yes	Yes
Condition FEs	Yes	Yes	Yes	Yes	Yes
Subject FEs	Yes	Yes	Yes	Yes	Yes
$R^2$	0.515	0.569	0.487	0.504	0.531
Observations	4395	2205	2190	2199	2196

Table 4: Subsequent investments depend on prior choices and new information. The dependent variable in the probit models in the table, *Stock holder*<sub>*i,b,t+1*</sub>, is an indicator equal to 1 if subject *i* in learning block *b* chose to hold the stock in trial *t* + 1, and 0 if they chose to hold the bond. The independent variables of interest are interaction terms between the person’s choice trial *t* (indicator variables *StockHolder*<sub>*i,b,t*</sub> or *BondHolder*<sub>*i,b,t*</sub>) and the level of the stock dividend revealed at the end of trial *t* (indicator variables *HighDividend*<sub>*i,b,t*</sub> or *LowDividend*<sub>*i,b,t*</sub>). The reference (omitted) category of events refers to situations when in trial *t* the person was a bond holder and the stock paid a low dividend. The models include subject fixed effects and controls for the experimental condition faced by each participant *i* in learning block *b* (i.e., gain vs. loss, low vs. high variance), as well as for the subject’s probability estimate at the end of trial *t* – 1 in that learning block. 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>Stock holder</i> <sub><i>i,b,t+1</i></sub>				
	— Experimental condition —				
	All	Gain	Loss	Low variance	High variance
<i>Stock holder</i> <sub><i>i,b,t</i></sub> & <i>High dividend</i> <sub><i>i,b,t</i></sub>	0.55 (13.03)***	0.51 (10.26)***	0.56 (11.60)***	0.50 (9.72)***	0.57 (14.19)***
<i>Stock holder</i> <sub><i>i,b,t</i></sub> & <i>Low dividend</i> <sub><i>i,b,t</i></sub>	0.32 (10.26)***	0.32 (7.91)***	0.26 (5.49)***	0.28 (7.41)***	0.32 (7.58)***
<i>Bond holder</i> <sub><i>i,b,t</i></sub> & <i>High dividend</i> <sub><i>i,b,t</i></sub>	0.34 (8.22)***	0.35 (7.95)***	0.31 (6.36)***	0.32 (6.32)***	0.35 (7.34)***
FES for <i>ProbabilityEstimate</i> <sub><i>i,b,t-1</i></sub>	Yes	Yes	Yes	Yes	Yes
Condition FEs	Yes	Yes	Yes	Yes	Yes
Subject FEs	Yes	Yes	Yes	Yes	Yes
<i>R</i> <sup>2</sup>	0.233	0.247	0.246	0.250	0.250
Observations	3647	1834	1813	1826	1821

Table 5: Subsequent investments depend on the random first choice in a learning block. The dependent variable in the probit models in the table,  $Stock\ holder_{i,b,t}$ , is an indicator equal to 1 if subject  $i$  in learning block  $b$  chose to hold the stock in trial  $t$ , and 0 if they chose to hold the bond. The independent variable of interest is the indicator  $Stock\ holder\ in\ trial_{i,b,1}$ , which is equal to 1 if in the first trial of learning block  $b$  subject  $i$  chose the stock. At the moment of the first choice in the block, the subject knows that the probability that the stock is good is 50%, and the expected value of choosing the stock is equal to that of choosing the bond. In these situations a risk neutral decision maker will randomly pick one of the assets. The regressions also include fixed effects of the value of the objective Bayesian probability that the stock is paying from the good distribution, based on the information revealed prior to trial  $t$ , fixed effects for various experimental conditions (i.e., gain vs. loss, low vs. high variance), and subject fixed effects. The specifications in the table use observations from various sets of trials in the learning block: all trials after the first one (i.e., trials 2-6), only trial 2, trials 3-4, and trials 5-6, respectively, to assess how long lived the effect of the random first choice is during the rest of the learning block. 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	$Stock\ holder_{i,b,t}$			
	Trials 2-6	Trial 2	Trials 3-4	Trials 5-6
$Stock\ holder\ in\ trial_{i,b,1}$	0.10 (2.90)***	0.24 (3.77)***	0.10 (2.50)**	0.03 (0.66)
FEs for <i>Objective probability that stock is good</i> $_{i,b,t-1}$	Yes	Yes	Yes	Yes
Condition FEs	Yes	Yes	Yes	Yes
Subject Fixed Effects	Yes	Yes	Yes	Yes
$R^2$	0.217	0.252	0.224	0.245
Observations	3644	714	1456	1458



Table 6: Subsequent investment choices rely differently on new information depending on whether it matches prior choices. The dependent variable in the probit models in the table is an indicator equal to 1 if person  $i$  chose the stock in trial  $t + 1$  of learning block  $b$ , and 0 if they chose the bond. The models include experimental condition fixed effects and subject fixed effects. We split the data based on whether in trial  $t$  the dividend presented matched the choice made earlier in that trial by the person. We define a match as occurring if the person chose the stock in trial  $t$  and then was informed that the stock outcome that trial was high, or if the person chose the bond in trial  $t$  and then was informed that the stock outcome was low. In the first two columns the independent variables of interest are the indicator for whether the dividend revealed in trial  $t$  was high or low, and the subjective estimate of the probability that the stock was paying from the good distribution, expressed by participant  $i$  at the end of trial  $t - 1$  in that learning block. In the last two columns the independent variable of interest is the Bayesian value of the probability that the stock is good, given the subjects' own prior belief (expressed at the end of trial  $t - 1$ ) and the dividend information revealed in trial  $t$ . Assuming that people update their own prior  $p$  using this new information in an objective manner, the Bayesian posterior probability that the stock is good at the end of trial  $t$  is  $7p/(4p + 3)$  if the new dividend is high, and  $3p/(7 - 4p)$  if the new dividend is low. 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>Stock holder</i> <sub><math>i,b,t+1</math></sub>			
	Trial $t$ dividend matches choice in trial $t$ ?			
	YES	NO	YES	NO
<i>High dividend</i> <sub><math>i,b,t</math></sub>	0.598 (11.59)***	0.021 (0.38)		
<i>Probability estimate</i> <sub><math>i,b,t-1</math></sub>	0.003 (3.12)***	0.006 (4.38)***		
<i>Bayesian estimate using own prior in trial <math>t - 1</math> and the dividend in trial <math>t</math></i>			0.011 (7.74)***	0.004 (3.36)***
Condition FEs	Yes	Yes	Yes	Yes
Subject Fixed Effects	Yes	Yes	Yes	Yes
$R^2$	0.374	0.089	0.328	0.078
Observations	1965	1682	1965	1682

Table 7: The value of the brain signal for learning about one’s success influences the incorporation of new information about asset payoffs in the updating of expectations, as well as in selecting subsequent investments. The dependent variables across the models in the table include the subjective probability that the stock is paying from the good dividend distribution, expressed by person  $i$  at the end of trial  $t$  in block  $b$  (first column), the subjective update in this beliefs relative to the prior trial (second column), and an indicator equal to 1 if the person chose the stock in trial  $t + 1$ , and 0 if they chose the bond (third and fourth columns). The independent variables of interest in the first and third columns include an indicator for whether the dividend observed in trial  $t$  was high or low, and its interaction with an indicator equal to 1 if at the time when the dividend was presented the participant had high (i.e., above median) activation in the ventromedial prefrontal (vmPFC) area encoding whether news matches one’s prior choice. The independent variables of interest in the second and fourth columns include the value of the objective update relative to the prior trial in the probability that the stock is paying from the good dividend distribution, given the newly observed dividend in trial  $t$ , and its interaction with the same indicator for high activation in the vmPFC area sensitive to encoding whether information confirms prior actions. The regressions also include fixed effects of the value of the subjective belief about the stock in the prior trial, fixed effects for various experimental conditions (i.e., gain vs. loss, low vs. high variance), and subject fixed effects. 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>Probability Estimate</i> <sub><math>i,b,t</math></sub>	<i>Subjective Update</i> <sub><math>i,b,t</math></sub>	<i>Stock holder</i> <sub><math>i,b,t+1</math></sub>	<i>Stock holder</i> <sub><math>i,b,t+1</math></sub>
<i>High dividend</i> <sub><math>i,b,t</math></sub>	26.55 (11.99)***		0.27 (6.14)***	
<i>High dividend</i> <sub><math>i,b,t</math></sub> × <i>HighMatchSignal</i> <sub><math>i,b,t</math></sub> <sup>vmPFC</sup>	2.28 (3.41)***		0.08 (2.28)**	
<i>Objective update</i> <sub><math>i,b,t</math></sub>		0.78 (12.16)***		0.008 (6.07)***
<i>Objective update</i> <sub><math>i,b,t</math></sub> × <i>HighMatchSignal</i> <sub><math>i,b,t</math></sub> <sup>vmPFC</sup>		0.08 (3.12)***		0.002 (2.35)**
<i>HighMatchSignal</i> <sub><math>i,b,t</math></sub> <sup>vmPFC</sup>	-0.15 (-0.23)	0.73 (1.38)	-0.05 (-1.43)	-0.007 (-0.36)
FEs for <i>Probability estimate</i> <sub><math>i,b,t-1</math></sub>	Yes	Yes	Yes	Yes
Condition FEs	Yes	Yes	Yes	Yes
Subject FEs	Yes	Yes	Yes	Yes
$R^2$	0.633	0.449	0.189	0.188
Observations	4416	4416	3663	3663

# Appendix

## 1. Calculating objective Bayesian posterior beliefs

Table A1. 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%-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%

## 2. Brain imaging data acquisition and analysis

**Functional magnetic resonance imaging (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. 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 two learning blocks. 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  $\pm 5$  dpt were excluded).

In the first trial of each block, the stock could be either good or bad with the same probability (50%). With every high (low) stock outcome, it is more (less) likely that the stock is good. After the stock payoff was shown in each trial, the subject was asked to estimate the probability that the stock is good. Estimations within 5% of the correct answer, namely the objective Bayesian value, were rewarded with 0.10 Euro that were added to the final payoff, at the end of the experiment. The last screen of each trial showed the subject's updated balance, that is, the accumulated payoffs from their asset choices so far. This balance did not reflect any additional money earned by correctly estimating the probability that the stock is paying from the good dividend distribution. This additional amount was only revealed to the subject at the end of the experiment, to avoid feedback effects on how people estimated probabilities during the task.

**fMRI data acquisition and analysis** 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).

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.

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.

### Details regarding Figure 6

The main goal of the fMRI analysis was to identify brain regions that are sensitive to prior investment choices during the processing of new stock information. In the analysis in Figure 6 we investigate whether the size of the new stock dividend (high vs. low) revealed at the end of a trial is encoded differently depending on the person's choice earlier in the trial. For this we computed a general linear model 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. The onsets of the estimation screen and the accumulated payoff presentation were modeled by the same indicator variables. As control variables we included the objective Bayesian probability that the stock is good given the information available as of the end of the prior trial, as well as the subjects' received payoff during the trial. Six motion parameters 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

expected activation in brain areas associated with outcome and reward encoding, such as the vmPFC and the ventral striatum (e.g. Bartra (2013)). The left panel of Figure 6 shows that the brain activity at the time of stock outcome presentation reflects the interaction effect of prior investments (stock or bond holder) and stock dividend (high or low dividend). Activity in the ventromedial prefrontal cortex (vmPFC) and the bilateral ventral striatum (vStr) is greater when the new dividend information matches the person choice that trial, namely, when stock holders see a high dividend, and bond holders see a low dividend. The figures on the left show statistical group activation maps of the two-by-two ANOVA. For illustration purposes, maps are thresholded at  $t > 3.14$ , cluster size  $kE > 10$ , projected on a template brain in MNI space, and color coded for the t-values as indicated by the color bar on the left.

In the right panel of Figure 6 we illustrate these results using subject-level estimates of brain activation in the regions of interest, vmPFC and vSTR, as a function of the choice made and the newly revealed dividend. For each subject we extracted the activation, or BOLD, response to each of the four events of interest, captured by the interaction terms between indicators of stock or bond choice, and indicators of whether the stock paid a high or low dividend that trial. We measured this activation response using anatomically defined region of interest masks for the vmPFC and the vSTR using the MarsBaR toolbox and the the AAL atlas (Brett et al. (2002), Tzourio-Mazoyer et al. (2002)). We then plotted the mean (M) and standard error (SE) of these parameters across all subjects against the four regressors of interest in the right panel in Figure 6. The individual subject regression models underlying these results control for the objective probability that the stock is good before the new stock outcome was presented and for the received payoff in the current trial.

### Details regarding Figure 7

The goal of the analysis in Figure 7 is to investigate whether the objective update in the probability that the stock is good is encoded differently depending on the person’s choice earlier in the trial. This probability update is a direct measure of the ”prediction error” realized at the presentation of a new stock outcome, in that it captures how much better or how much worse the stock is relative to what its assessed value was in the prior trial, in light of the new dividend observed in the current trial.

The top panel in 7 shows that activity in the vmPFC and the bilateral vStr at the presentation of a new stock outcome correlates positively with the objective probability update when the subjects hold the stock, but not when they hold the bond. The panel shows statistical group activation maps of the paired t-test of the parametric modulation of the BOLD response by the objective probability update at the stock outcome presentation for trials in which the subject held the stock contrasted with trials in which the subject held the bond. For illustration purposes, maps are thresholded at  $t > 3.28$ , cluster size  $kE > 20$  voxels, projected on a template brain in MNI space and color coded for the t-values as indicated by the color bar on the left.

In the bottom panel of Figure 7 we illustrate these results using subject-level estimates of brain activation in the regions of interest, vmPFC and vSTR, as a function of the choice made at the beginning of a trial and the value of the objective probability update that occurred at the end of the trial. Specifically, the figure shows that activity in the vmPFC and bilateral vStr at the time when the stock outcome is presented increases with the update in the objective probability that the stock is good when subjects hold the stock, but decreases with this update about the quality of the stock when subjects hold the bond. The figure presents the mean and standard error of activation, or BOLD response, across all subjects, in the vmPFC and vSTR areas, separately for situations when the hold the stock, and hold the bond, respectively, and across four categories of the objective probability update:  $-0.2$ ,  $[-0.15, -0.008]$ ,  $[0.008, 0.15]$ , and  $0.2$ . These categories, or update bins, encompass all the possible changes in the probability that the stock is good, upon observing one additional dividend, and each contains roughly a quarter of all trials. The levels of this probability are shown in Table A1 in the Appendix.

For each subject, BOLD response values corresponding to each choice (stock versus bond) and each probability update bin were extracted from a 5 mm sphere around the single-subject voxel with the peak value for the correlation shown in the top panel, within an anatomical mask for the respective region of interest, derived from the AAL atlas (Tzourio-Mazoyer et al. (2002)). This was done to correct for variations in single-subject peaks. Note that this is an illustration of the whole-brain analysis shown in the top panel in Figure 7 and is not subjected to inference statistics. Whereas the whole-brain analysis in the top panel in Figure 7 tests how well the BOLD response fits with the linear function of the objective probability update (i.e. the slope), these graphs visualize the estimated magnitude of the BOLD response separately for bins of

increasing value of the objective probability update. This illustration of the whole-brain analysis was done with the *rfxplot* toolbox (Glscher, 2009). The individual subject regression models underlying these results control for the objective probability that the stock is good before the new stock outcome was presented and for the received payoff in the current trial.

### Details regarding Table 7

The second goal of the fMRI analysis was to investigate at the subject-trial level the effect of the differential encoding in these areas of new dividends depending on prior investments on probability updating and future choices. To analyze the fMRI time courses, we first defined volumes of interest (VOIs) in the vmPFC and the bilateral vSTR based on 8 mm spheres around the voxel with peak correlation value for the match-mismatch contrast. 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 response for the onset of the stock outcome presentation, and used those values in the analysis in Table 7.

### 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.
- Glscher, J. (2009). Visualization of group inference data in functional neuroimaging. *Neuroinformatics*, 7(1), 7382. <http://doi.org/10.1007/s12021-008-9042-x>
- 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