Asymmetric learning from financial information

Kuhnen, Camelia M.

Northwestern University – Kellogg School of Management

12 June 2012

Asymmetric Learning from Financial Information*

Camelia M. Kuhnen†

June 2012

Abstract

The goal of this study is to ask whether investors learn differently from gains (positive news) versus losses (negative news), whether learning performance is better or worse when people are actively investing in a security or passively observing the security's payoffs, and whether there are personal characteristics that correlate with learning performance. The experimental evidence documented here indicates that the ability to learn from financial information is on average worse in the loss domain, in particular if the investor has personally experienced the prior outcomes of the financial asset considered. Within-individual, learning from gains versus losses, or during active versus passive involvement, are not perfectly correlated, indicating that there exists heterogeneity across people with respect to the type of financial information or context to which they are the most sensitive. Learning performance is determined by acquired financial expertise as well as by genetic factors related to memory and cognitive control.

†Department of Finance, Kellogg School of Management, Northwestern University, 2001 Sheridan Rd., Evanston, IL 60208-2001, c-kuhnen@kellogg.northwestern.edu.

1 Introduction

This paper examines whether investors' ability to learn in financial markets depends on the context and type of information received. Specifically, do people learn differently from gains or positive news, relative to losses or negative news? Does learning performance depend on whether the individual is actively investing in a security, or passively observing the security's payoffs? Finally, are there personal characteristics, either acquired or innate, that predict learning performance?

The idea that learning may be different in the gain and loss domains is different from and complementary to the well-documented phenomenon of loss-aversion suggested by Kahneman and Tversky (1979), whereby the disutility of losing an amount of money is greater, in absolute terms, than the utility of winning that amount. A large body of work has provided ample evidence for this difference in preferences in the gain and loss domain.[1] Nonetheless, recent findings in the neuroeconomics literature suggest that gains and losses are different not only in terms of how they shape the value function, but also, in terms of how they are incorporated in the formation of beliefs.

Specifically, information regarding gains (or positive surprises) and information regarding losses (or negative surprises) is processed by different brain centers important for the generation of emotional states and reactions (Kuhnen and Knutson (2005), Knutson and Bossaerts (2007)). Within a person, the sensitivity of these centers to information may be different, which could lead some individuals to pay more attention to and learn from positive news, and others to focus more on negative news (Samanez-Larkin et al. (2011)). Furthermore, high levels of physiological arousal, which characterize powerful emotional states, have a disruptive effect on memory and cognitive function (Ashby et al. (2002), Mather et al. (2006)), and arousal levels are higher when losses are possible, relative to gains (Sokol-Hessner et al. (2009)). This suggests that memory and learning may be better for details related to positive

---

[1]In the finance literature, see for example Odean (1998), Barberis and Huang (2001) and Barberis et al. (2001).

news than for those related to negative news, in line with experimental findings in psychology (Eppinger et al. (2010), Mather and Schoeke (2011)). Moreover, reducing the intensity of emotional states, for example by taking a broader perspective on financial choices or by lowering the financial stakes of decisions, decreases physiological arousal and the sensitivity to negative outcomes (Sokol-Hessner et al. (2012)), and leads to better choices in experimental gambles (Charness and Levin (2005)) and to more successful trading in financial markets (Lo et al. (2005)). Hence, the natural hypothesis suggested by this prior work relating emotion, brain function and decision making is that investors learn differently in the gain versus the loss domain, and depending on whether they actively engage in trading or passively observe asset payoffs.

In this paper I test and confirm this hypothesis in an experimental setting. I find that the ability to learn from financial information is on average worse in the loss domain, in particular if the investor has personally experienced the prior outcomes of the financial asset considered. In such situations, subjective beliefs about asset payoffs are overtly pessimistic. Within-individual, learning from gains versus losses, or during active versus passive involvement, are not perfectly correlated, indicating that there exists heterogeneity across people with respect to the type of financial information or context to which they are the most sensitive. Moreover, heterogeneity in beliefs is the highest during periods characterized by negative outcomes. To relate these findings back to the neuroeconomics work that informed the hypothesis of the paper, I collect genetic data from the participants in the experiment and show that a specific gene ($COMT$) previously linked to memory function and cognition is a significant driver of the differences in learning performance documented here.

The experimental findings of this paper can help explain several intriguing patterns observed in financial markets. First, empirical work has shown that risk premia rise strongly after large negative moves in prices. This is consistent with an increase in risk aversion after negative shocks (Todorov (2010), Bollerslev and Todorov (2011), Guiso et al. (2011)), but in addition can also be driven by traders being overtly pessimistic during bad times, as

2

suggested by the experimental results shown here. Second, empirically it has been shown that during downturns there is more heterogeneity in the actions of financial market participants, as indicated by the higher volatility in asset prices (e.g., Schwert (1989), Campbell et al. (2001)), higher dispersion in analysts' forecasts (Barinov (2009)), or by the increased variation in strategies deployed by mutual fund managers (Kacperczyk et al. (2011)). The experimental findings presented here suggest that this increase in the heterogeneity of actions of market participants during bad times (i.e., those characterized by a preponderance of negative news) may be driven by the increase in the heterogeneity of beliefs of these individuals during poor market conditions. Third, the result that passive observers learn better than those actively involved in investing, particularly during periods characterized by many negative outcomes, suggests that financial planners and advisors add value to investors by responding more objectively to market outcomes, and thus can help justify the significant demand for these professionals (Elmerick et al. (2002)).

To investigate whether indeed learning is different depending on the type of information received, and on the degree of involvement of agents in financial markets, eighty-seven adults were invited to participate in a study that required the completion of two financial decision making tasks. In the Active task subjects made sixty decisions to invest in one of two securities: a risky security (stock) with risky payoffs coming from one of two distributions, one better than the other in the sense of first-order stochastic dominance, and a riskless security (bond) with a known payoff. After each choice participants provided an estimate of the probability that the risky security was paying from the better distribution. In the Passive task subjects were only asked to provide the probability estimate that the risky security was paying from the better distribution, after observing its payoff in each of sixty trials. In either task, two types of conditions - Gain or Loss - were possible. In the Gain condition, the two securities provided positive payoffs only. In the Loss condition, the two securities provided negative payoffs only. Subjects were paid based on their investment payoffs and the accuracy of the probability estimates provided. Participants were genotyped with respect

3

to the $COMT$ gene, which is known to influence the activity of the prefrontal cortex, a brain area critical for emotion control during cognitive tasks as well as for the functioning of working memory (Dickinson and Elvevg (2009)). A particular variant of this gene, referred to as the $COMT$ Met/Met genotype, has previously been associated with better memory and cognitive performance relative to the other two genotypes, namely Val/Met and Val/Val (Frank et al. (2007), Doll et al. (2011)).[2]

I find that subjects learn significantly better from information about the risky security's payoffs in the Gain condition relative to the Loss condition. The errors in the subjective probability estimates (measured relative to the objective Bayesian posteriors that the stock is paying from the good distribution) are on average 1.86% lower in the Gain condition relative to the Loss condition ($p < 0.001$). The difference in probability estimation errors between the Gain and Loss conditions is twice as large in the Active task (2.56%) relative to the Passive task (1.16%), and is largest for high values of the Bayesian posterior. Probability estimate errors are also significantly lower in the Passive task relative to the Active task, and this effect is driven by the Loss condition, where errors are 1.76% lower in Passive trials. In other words, people learn worse from loss information, especially if they actively invest, and are overtly pessimistic in these circumstances. Also, individuals who err more when constructing their subjective probability estimates make fewer optimal choices (i.e., matching those of a risk-neutral Bayesian-updating agent).

Furthermore, within individual, Gain and Loss learning performance are only partially correlated (ρ=0.7, $p < 0.001$), and Active and Passive learning performance are also only partially correlated (ρ=0.6, $p < 0.001$), indicating that certain individuals are more sensitive to information in specific domains (Gain vs. Loss) or investment conditions (Active vs.

---

[2]Briefly, the catechol-O-methyltransferase ($COMT$) gene on chromosome 22q11 controls the levels of neurotransmitter dopamine in the prefrontal cortex. In humans it contains a highly functional and common variation in its coding sequence in exon 4, namely a substitution of valine (Val) by methionine (Met). This is an instance of a mutation typically referred to as a single nucleotide polymorphism (SNP). The more stable Val allele is associated with greater dopamine degradation and less synaptic dopamine in the prefrontal cortex than the less stable Met allele. As result, the COMT genotype impacts cognition mediated by this brain area, specifically executive control and working memory, with the Met allele being generally associated with better performance. See Dickinson and Elvevg (2009) for a comprehensive review.

Passive). The heterogeneity across participants with respect to their ability to learn from financial information, measured as the standard deviation of subjective beliefs, is significantly higher in the Loss condition (16%) relative to the Gain condition (14%), and during Active (15%) relative to Passive (14%) trials.

I also identify personal characteristics related to the ability to learn across these contexts. Financial knowledge – specifically familiarity with basic concepts such as expected returns and probabilities – helps investors learn better. Gender, age and ethnicity are not significantly related to learning performance. However, participants' genotype with respect to the $COMT$ gene is a significant predictor of the ability to learn from financial information across all contexts, with effects similar in magnitude to those of financial knowledge. Specifically, participants with the $COMT\ Met/Met$ genotype provide probability estimates that are the closest to the correct Bayesian values, in line with prior neuroscience findings that show the positive effect of this particular genetic variant on memory and cognition.[3] Importantly, I also find that financial knowledge and being endowed with the Met/Met genotype lead not only to more accurate subjective beliefs, but also to choosing the optimal asset more often.

The broad idea that learning may be context-dependent is supported by extant empirical findings in economics. For example, learning displays salience and recency effects. In the context of portfolio allocation decisions, Malmendier and Nagel (2011) find that individuals who have more recently experienced low stock-market returns are reluctant to invest in equities and have pessimistic beliefs about future stock returns. Also, people are significantly more likely to buy insurance against natural catastrophes such as floods or earthquakes shortly after such events happen, even though the objective probability of their occurrence does not change (Kunreuther et al. (1978), Palm (1995), Froot (2001)).[4] Other learning biases, such as conservatism and representativeness, can drive anomalous patterns in asset

---

[3]This finding provides a microfoundation for the result documented by Cronqvist and Siegel (2012) that there exists more similarity in the propensity to display investment biases among identical twins, relative to fraternal twins, suggesting that suboptimal behavior may have a genetic component.

[4]See Gennaioli and Shleifer (2010) for a model of "local thinking" that is consistent with the observed importance of salience for belief formation.

prices such as the under- and over-reaction to news (Barberis et al. (1998)).

Learning in financial markets has been the focus of a small but growing experimental literature. Kluger and Wyatt (2004) document the existence of heterogeneity across traders with respect to their ability to learn according to Bayes' rule, and the impact of this heterogeneity on asset prices. Asparouhova et al. (2010) find that investors unable to perform correct probability computations prefer to hold portfolios with unambiguous returns and do not directly influence asset prices. Payzan-LeNestour (2010) shows that Bayesian learning is a reasonably good model for investment decisions in complex settings. Bruguier et al. (2010) show that the ability to forecast price patterns in financial markets depends on traders' capacity to understand others' intentions (i.e., "theory of mind"), and not on the ability to solve abstract mathematical problems. Kogan (2009) and Carlin et al. (forthcoming) show that strategic considerations influence learning and trading in experimental asset markets.

In addition, there exists a novel body of theoretical work focused on understanding the role of bounded rationality and non-standard preferences in the formation of beliefs by economic agents. The existence of costs in the acquisition of information leads to rational inattention, which can help explain the portfolio holdings of individual investors (Van Nieuwerburgh and Veldkamp (2010)), investment strategies of mutual fund managers (Kacperczyk et al. (2011)), or myopic choices in dynamic search problems (Gabaix et al. (2006)). Sparsity in the information set of boundedly-rational decision makers can generate observed patterns in consumption and portfolio choices (Gabaix (2011)). The existence of an anticipatory component in agents' utility function can cause individuals to choose optimistic, but incorrect priors and act on these when making investment or consumption decisions (Brunnermeier and Parker (2005), Brunnermeier et al. (2007)). This body of work assumes that individuals learn according to Bayes' rule, given a (possibly incorrect) prior belief and (possibly sparse) new information. The focus of this paper is complementary to this literature, as the evidence presented here sheds light on the process by which people incorporate newly available information into beliefs starting from objective priors, and documents domain-specific departures

6

from Bayesian learning.

The novel contribution of this paper, therefore, is to show that the ability to learn from financial information is different in the gain and the loss domains, depends on whether the investor has personally experienced the prior outcomes of the financial asset considered, and is driven by acquired familiarity with financial concepts, as well as by a genetic factor related to memory and emotion regulation during cognitive tasks.

2   Experimental design

2.1   Setup

Eighty-seven individuals (43% males, mean age 20 years) were recruited at Northwestern University and participated in the experiment. Table 1 presents the sample summary statistics. Each participant completed two financial decision making tasks, referred to below as the Active and the Passive Task, during which information about two securities, a stock and a bond, was presented. Each task included two types of conditions: Gain or Loss. In the Gain condition, the two securities provided positive payoffs only. The stock payoffs were +$10 or +$2, while the bond payoff was +$6. In the Loss condition, the two securities provided negative payoffs only. The stock payoffs were -$10 or -$2, while the bond payoff was -$6.

In either condition, the stock was either good or bad. If the stock was good, it paid the high dividend with 70% probability on each trial. If it was bad, it paid the high dividend with 30% probability in each trial. In the beginning of each block of 6 trials, it was equally likely that the stock will be good or bad during those trials.

In the Active task participants made sixty decisions to invest in one of the two securities, the stock or the bond, observed the stock payoff (irrespective of their choice) and then provided an estimate of the probability that the stock was paying from the better distribution. Figures 1 and 2 show the time line of a typical trial in the Active task, in the Gain and

Loss conditions, respectively. In the Passive task participants were only asked to provide the probability estimate that the stock was paying from the better distribution, after observing its payoff in each of sixty trials. Figure 3 shows the time line of typical trials in the Passive task, in either the Gain or the Loss conditions.

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

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

For each participant we also obtained a measure of their financial literacy, since the prior literature has documented a positive correlation between this personal characteristic and successful financial outcomes (e.g., Lusardi and Mitchell (2007)). To measure general finance knowledge, each participant was asked the following questions after the completion of the experimental tasks: "Imagine you have saved $10,000. You can now invest this money over the next year using two investment options: a U.S. stock index mutual fund which tracks the performance of the U.S. stock market, and a savings account. The annual return per dollar invested in the stock index fund will be either +40% or -20%, with equal probability. In other words, it is equally likely that for each dollar you invest in the stock market, at the end of the one year investment period, you will have either gained 40 cents, or lost 20 cents. For the savings account, the known and certain rate of return for a one year investment is 5%. In other words, for each dollar you put in the savings account today, for sure you will gain 5 cents at the end of the one year investment period. We assume that whatever amount

you do not invest in stocks will be invested in the savings account and will earn the risk free rate of return. Given this information, how much of the $10,000 will you invest in the U.S. stock index fund? Choose an answer that you would be comfortable with if this was a real-life investment decision. The answer should be a number between $0 and $10,000."

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

The correct answer to this question is [F]. The actual choices (if other than [F]) made by participants indicate three different types of errors that can occur when calculating the expected value of their portfolio holdings: the lack of understanding of statements regarding probabilities (answers [B], [C], [E], [G]); the lack of understanding of the difference between net and gross returns (answers [A],[C], [D] and [E]); and confusing the stock versus risk-free asset investments (answers [C], [D], [G] and [H]). Therefore, a financial knowledge score varying between zero and three can be constructed, based on the number of different types of errors contained in the answer provided by each participant (i.e., zero errors for answer [F], one error for answers [A], [B] and [H], two errors for answers [D], [E] and [G], and three for answer [C]). Hence a financial knowledge score of 3 indicates a perfect answer, while a score of 0 indicates that the participant's answer included all three possible types of errors.

Of the 87 participants, 45 made no errors, 24 made one type of error only, 17 made two types of errors, and 1 person made all three possible types of errors.

2.2 Genotyping

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

3 Results

3.1 Probability estimation errors across domains

The data show that there exist asymmetries in learning across domains, as indicated by the results in Figure 4 and in the regression models in Table 2. The figure plots subjective probability estimates (averaged across all 87 participants) as a function of the objective Bayesian posterior probabilities that the stock was paying from the better dividend distribution. If people learn from dividends exactly in accordance with Bayes' rule, the subjective estimates and objective probabilities would line up perfectly. However, this is not the case in the data.

I find that individuals learn significantly better from information about the risky security's payoffs in the Gain condition relative to the Loss condition. This can be seen by

comparing the left two panels in Figure 4, which refer to trials in the Active and Gain, and Passive and Gain conditions, with the right two panels, which refer to trials in the Active and Loss, and Passive and Loss conditions. On average, as documented in Table 2, the errors in the subjective probability estimates (measured relative to the objective Bayesian posteriors that the stock is paying from the good distribution) are 1.86% lower in the Gain condition relative to the Loss condition ($p < 0.001$). The difference in probability estimation errors between the Gain and Loss conditions is twice as large in the Active task (2.56%) relative to the Passive task (1.16%). In general, probability estimate errors are lower by 1.06% in the Passive task relative to the Active task ($p < 0.1$). This effect is driven by choices in the Loss condition only, where there errors are lower by 1.76% in Passive versus Active trials ($p < 0.05$).

Furthermore, Figure 4 shows that the worst learning overall happens in the Loss condition in particular when the objective probability that the stock pays from the better (but still with negative support) distribution is high. In those situations, participants' subjective posterior probabilities are the most pessimistic relative to the objective values. Comparing the first and last columns in Table 2 illustrates this effect. Specifically, while on average the absolute value of probability estimation errors is 1.86% higher in the Loss condition trials relative to other trials, this difference increases to 4.31% for Loss trials with high values ($\geq 50\%$) of the objective posterior probability that the stock pays dividends from the better distribution. In other words, the data show that when faced with a sequence of mildly negative news (i.e., when the stock is bad, but likely not the worst possible), people are overtly pessimistic.

In general, as seen in Figure 4, across Active/Passive and Gain/Loss trials, subjects update their priors in such a way that the expressed posterior probabilities that the stock is paying from the good distribution is significantly higher (by 12% on average) than the objective Bayesian posterior for low values of this objective probability, and significantly lower (by 13% on average) than the objective Bayesian posterior for high values of this objective probability, a result which replicates the experimental patterns documented in Kuhnen and

Knutson (2011). This relationship between subjective and objective posterior beliefs resembles the relationship between decision weights and objective probabilities postulated by Prospect Theory (e.g., Kahneman and Tversky (1979), Prelec (1998)), but it refers to errors in updating priors, and not to people's tendency to overweight rare events and underweight frequent ones. Note that either type of mistake – updating errors or the use of decision weights – can help explain people's focus on small probability events (as indicated by their willingness to buy insurance or lottery tickets, for example).

While these results show that subjective beliefs are incorrect in some situations, for these findings to matter it has to be the case that these subjective beliefs actually drive choices. Ultimately, therefore, it is necessary to show that more incorrect subjective probability estimates lead to more suboptimal asset choices. Figure 5 shows that this is true in the data. Specifically, there is a strong and significant ($p < 0.001$) negative relationship between the size of probability errors committed by a participant and the number of times they chose the optimal asset (i.e., the asset that a risk-neutral Bayesian agent would choose given the available information set). Increasing the average of the probability error by 1% leads to 0.8 fewer optimal choices.[5]

Moreover, besides acting on their subjective beliefs, individuals do not seem to be aware that these probability estimates are incorrect, or that they tend to err more in certain types of trials. At the end of each trial, participants were asked to provide a confidence number (from 1 to 9, with 1 meaning not confident at all and 9 meaning very confident) to indicate how much they trust the subjective probability estimate produced in that trial. I find no significant differences between the average confidence of participants in active versus passive trials (5.31 vs 5.39, respectively), or during loss versus gain trials (5.27 vs 5.43, respectively).

---

[5]Note that each participant is asked to make 60 asset choices. However, for trials where the objective probability that the stock is paying from the better distribution is exactly 50%, a risk-neutral agent should be indifferent between choosing the stock or the bond. Hence on such trials it can not be determined whether the participant behaved optimally, since either choice would seem correct even though it might have been made for the wrong reason. Hence, I only define optimal choices to be those recorded in trials where the objective prior is either strictly below or strictly above 50%.

3.2 Heterogeneity in learning across participants

Aside from these average effects that indicate learning is most problematic in the loss domain and during active investing, the data also show that there exists a significant degree of heterogeneity across participants in the accuracy of their probability estimates, and also, with respect to the context and type of information from which they learn best. Within individual, Gain and Loss learning performance are only partially correlated ($\rho = 0.73$, $p < 0.001$), and Active and Passive learning performance are also only partially correlated within person ($\rho = 0.57$, $p < 0.001$), as can be seen in Figure 6. These results indicate that certain individuals are more sensitive to information in specific domains (Gain vs. Loss) or investment conditions (Active vs. Passive), as also illustrated in Figure 7.

That being said, most people err more in learning in the Loss condition, and during Active investment, as can be seen in Table 3. Specifically, 57 out of 87 participants display larger learning errors during the Loss condition, relative to the Gain condition. Also, 54 of the 87 participants display larger learning errors during Active investment trials, relative to Passive trials. The most common combination of trial types during which participants make the biggest probability estimation errors is comprised of Loss & Active trials, with 36 out of 87 subjects (i.e., 41%) being worse learners during this combination of trial types.

Furthermore, I find that the heterogeneity in learning abilities across market participants is significantly higher in the Loss domain. Specifically, the standard deviation of probability estimation errors is 13.98% in Gain trials, and 15.86% in Loss trials. A Bartlett's test rejects the null hypothesis that the variances of the absolute probability errors in these two samples are equal ($p < 0.001$). There also exists more heterogeneity in participants' learning abilities during Active trials relative to Passive trials. The standard deviations of absolute values of probability errors in these two conditions are 15.45% and 14.39%, respectively, and a Bartlett's test rejects the null hypothesis that the subsample variances are equal ($p < 0.05$). These differences are summarized in Table 4.

3.3 Predictors of better learning

The results documented above lead to the question of what exactly determines a person's ability to learn from financial information across different domains. A likely driver of this ability is the person's degree of familiarity with finance concepts, which I measure using the financial knowledge score given by their answer to the post-experiment portfolio allocation question described in the Experimental Design section. Indeed, as can be seen in the OLS regression in the first column of Table 5, financial knowledge is a significant and positive predictor of accuracy in probability estimates. An increase of one unit in the financial knowledge score, meaning making one fewer mistakes in the post-experiment portfolio return calculation, leads to a decrease of 1.83% in the size of the probability estimation error in a typical trial. The same regression model shows that the participants' gender, age or ethnicity do not have significant effects on learning ability.

While it is reassuring, and not surprising, that investors' acquired knowledge about basic financial concepts such as expected returns helps them learn better from news in the marketplace, the relatively low R^2 of this OLS model implies that a significant fraction of the variation in learning across people is driven by other factors, some of which may be innate. Indeed, as predicted based on prior neuroeconomics work, I document that variation across people in the ability to learn from and correctly use financial information can be in part explain by their $COMT$ genotype. Specifically, I find that $COMT$ Met/Met individuals express probability estimates closest to the correct Bayesian values. For these individuals, the absolute value of their probability estimation errors is 2.34% lower than for the rest of the participants, as can be seen in the regression model in the second column of Table 5. The results shown in columns three to six indicate that the positive effect of the Met/Met variant on learning is similar in magnitude across active, passive, gain or loss trials. However, as illustrated by the results in the last column on Table 5, the benefit of being a Met/Met type is particularly high in conditions where the learning errors displayed by participants in general are the highest - namely, in loss condition situations where the average participant

is overtly pessimistic (i.e., when the probability that the stock is paying from the better distribution is above 50%). This finding is in line with the hypothesis that the $COMT$ Met/Met genotype improves the functioning of the prefrontal cortex, a brain area which is critical for the regulation of emotional states, as well as for working memory and cognitive control, and this is particularly useful in situations where people are most prone to making errors.

Importantly, financial knowledge and the $COMT$ genotype of participants drive these individuals' subjective beliefs as well as their propensity to select the optimal asset given the available information (i.e., to choose like a risk-neutral Bayesian agent). Figures 8 and 9 illustrate these effects. Here I use the most conservative analysis approach by averaging the data at the participant level and working with 87 observations instead of several thousand as done in the regression models in Table 5. In line with the regression results, Figure 8 shows that participants with high finance knowledge (i.e., those who made no mistake in answering the portfolio return question asked after the experimental task) or those endowed with the $COMT$ Met/Met genotype have lower errors in their probability estimates relative to the rest of the group. Specifically, average estimation errors are 13.5% in the case of high finance knowledge participants, and 16.5% for the rest. Estimation errors are on average 13% for $COMT$ Met/Met individuals and 15.5% for the other types (i.e., Val/Val and Val/Met). These differences are significant at conventional levels ($p < 0.05$), and the magnitudes are similar to those documented in the trial-by-trial regressions in Table 5. Furthermore, producing more accurate beliefs helps participants with high finance knowledge or with the $COMT$ Met/Met genotype to do better in terms of the number of times they select the optimal asset. Figure 9 shows that the performance with respect to making correct choices is significantly better ($p < 0.05$) for high versus low finance knowledge individuals (35.1 vs. 31.2 optimal choices, respectively), and for Met/Met participants versus the others (36.1 vs. 32.5 optimal choices, respectively).

4 Implications for financial markets

The premise of the experiment presented here, built on recent findings in neuroeconomics, is that there are differences in how the brain processes information depending on its nature (i.e., good vs. bad news) or depending on the intensity of the emotional state felt at the time when new financial information is presented. As a result, learning in financial markets may have significant context dependence, and therefore it may be beneficial for investors as well as regulators to be aware of these learning asymmetries. This implication mirrors the recommendations made by Bossaerts (2009) and Lo (2011), who stress the importance of affect for financial risk preferences, and propose that transparency and regulation (e.g., trading restrictions) can prevent the exacerbation of negative outcomes such as those observed during the financial crises of the past decade. An immediate suggestion based on the findings documented here is that investors may benefit from delegating their asset choices to unbiased financial advisors, especially during poor stock market conditions.

Furthermore, besides providing a potential explanation for observed empirical patterns such as the increase in risk premia, volatility and heterogeneity in actions of market participants during bad economic times, the results of this experiment may also be relevant for future theoretical work. For example, models building on these findings could illustrate the importance of the presence of imperfect Bayesian learners for the volatility of asset prices, the informativeness of prices, and the duration of periods where incorrect beliefs about the future payoffs of assets may cause price deviations from fundamental values.

5 Conclusion

This paper documents the existence of asymmetries in learning in financial markets, depending on the type of information received by investors, and the context faced by these individuals. Specifically, I find that learning differs in the gain and the loss domain, and that, on average, it is worse in the loss domain. Learning performance also depends on whether

the investor has personally experienced the prior outcomes of the financial asset. On average, learning is worse during active investing in assets relative to settings where investors passively observe dividend news. Within-individual, learning in the gain vs. loss domain, or in the active vs. passive conditions are not perfectly correlated, indicating that there exists heterogeneity across individuals with respect to the context or type of information to which they are the most sensitive. Acquired familiarity with basic financial concepts, as well as innate variation with respect to a particular gene ($COMT$) previously linked to memory and cognitive performance are significant predictors of the investors' ability to learn from financial information and make optimal asset choices.

References

Ashby, F. G., Valentin, V. V. and Turken, A. U.: 2002, *Emotional cognition: From brain to behaviour*, John Benjamins Publishing Company, Amsterdam, chapter The effects of positive affect and arousal on working memory and executive attention: Neurobiology and computational models, pp. 245–287.

Asparouhova, E., Bossaerts, P., Eguia, J. and Zame, W.: 2010, Cognitive biases, ambiguity aversion and asset pricing in financial markets, *Working paper* .

Barberis, N. and Huang, M.: 2001, Mental accounting, loss aversion and individual stock returns, *Journal of Finance* **56**(4).

Barberis, N., Huang, M. and Santos, T.: 2001, Prospect theory and asset prices, *The Quarterly Journal of Economics* .

Barberis, N., Shleifer, A. and Vishny, R.: 1998, A model of investor sentiment, *Journal of Financial Economics* **49**, 307–343.

Barinov, A.: 2009, Analyst disagreement and aggregate volatility risk, *Working paper* .

Bollerslev, T. and Todorov, V.: 2011, Tails, fears and risk premia, *Journal of Finance* **66**, 2165–2211.

Bossaerts, P.: 2009, What decision neuroscience teaches us about financial decision making, *Annual Review of Financial Economics* **1**, 383–404.

Bruguier, A. J., Quartz, S. R. and Bossaerts, P.: 2010, Exploring the nature of "trader intuition", *Journal of Finance* **65**(5), 17031723.

Brunnermeier, M. K., Christian, G. and Parker, J. A.: 2007, Optimal beliefs, asset prices, and the preference for skewed returns, *American Economic Review* **97**(2), 159–165.

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

Campbell, J. Y., Lettau, M., Malkiel, B. G. and Xu, Y.: 2001, Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk, *Journal of Finance* **56**(1), 1–43.

Carlin, B., Kogan, S. and Lowery, R.: forthcoming, Trading complex assets, *Journal of Finance* .

Charness, G. and Levin, D.: 2005, When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect, *American Economic Review* **95**(4), 1300–1309.

Cronqvist, H. and Siegel, S.: 2012, Why do individuals exhibit investment biases?, *Working paper* .

Dickinson, D. and Elvevg, B.: 2009, Genes, cognition and brain through a COMT lens, *Neuroscience* **164**(1), 72–87.

Doll, B. B., Hutchison, K. E. and Frank, M. J.: 2011, Dopaminergic genes predict individual differences in susceptibility to confirmation bias, *Journal of Neuroscience* **31**(16), 61886198.

Elmerick, S. A., P, M. C. and Fox, J. J.: 2002, Use of financial planners by U.S. households, *Financial Services Review* **11**, 217–231.

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

Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T. and Hutchison, K. E.: 2007, Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning, *Proceeding of the National Academy of Sciences* **104**(41), 16311–16316.

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

Gabaix, X.: 2011, A sparsity-based model of bounded rationality, *Working paper* .

Gabaix, X., Laibson, D., Moloche, G. and Weinberg, S.: 2006, Costly information acquisition: Experimental analysis of a boundedly rational model, *American Economic Review* **96**(4), 1043–1068.

Gennaioli, N. and Shleifer, A.: 2010, What comes to mind, *Quarterly Journal of Economics* **125**(4), 1399–1433.

Guiso, L., Sapienza, P. and Zingales, L.: 2011, Time varying risk aversion, *Working paper* .

Kacperczyk, M., Van Nieuwerburgh, S. and Veldkamp, L.: 2011, Rational attention allocation over the business cycle.

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

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

Knutson, B. and Bossaerts, P.: 2007, Neural antecedents of financial decisions, *Journal of Neuroscience* **27**(31), 8174–8177.

Kogan, S.: 2009, Distinguishing the effect of overconfidence from rational best-response on information aggregation, *Review of Financial Studies* **22**(5), 1889–1914.

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

Kuhnen, C. M. and Knutson, B.: 2011, The influence of affect on beliefs, preferences, and financial decisions, *Journal of Financial and Quantitative Analysis* **46**(3), 605–626.

Kunreuther, H., Ginsberg, R., Miller, L., Sagi, P., Slovic, P., Borkan, B. and Katz, N.: 1978, *Disaster Insurance Protection: Public Policy Lessons*, Wiley Interscience.

Lo, A. W.: 2011, *Handbook on Systemic Risk*, Cambridge University Press, chapter Fear, Greed, and Financial Crises: A Cognitive Neurosciences Perspective.

Lo, A. W., Repin, D. V. and Steenbarger, B. N.: 2005, Fear and greed in financial markets: A clinical study of day-traders, *American Economic Review* **95**(2), 352359.

Lusardi, A. and Mitchell, O. S.: 2007, Baby boomer retirement security: The roles of planning, financial literacy, and housing wealth, *Journal of Monetary Economics* **54**(205-224).

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

Mather, M., Mitchell, K. J., Raye, C. L., Novak, D. L., Greene, E. J. and Johnson, M. K.: 2006, Emotional arousal can impair feature binding in working memory, *Journal of Cognitive Neuroscience* **18**(4), 614–625.

Mather, M. and Schoeke, A.: 2011, Positive outcomes enhance incidental learning for both younger and older adults, *Frontiers in Neuroscience* **5**(129), 1–10.

Odean, T.: 1998, Are investors reluctant to realize their losses?, *The Journal of Finance* **53**(5).

Palm, R.: 1995, *Earthquake Insurance: A Longitudinal Studies of California Homeowners*, Boulder: Westview Press.

Payzan-LeNestour, E.: 2010, Bayesian learning in unstable settings: Experimental evidence based on the bandit problem, *Working paper* .

Prelec, D.: 1998, The probability weighting function, *Econometrica* **66**(3), 497–527.

Samanez-Larkin, G., Kuhnen, C. M. and Knutson, B.: 2011, Gain and loss learning differentially contribute to life financial outcomes, *PLoS ONE* **6**(9).

Schwert, G. W.: 1989, Why does stock market volatility change over time?, *Journal of Finance* **44**(5), 1115–1153.

Sokol-Hessner, P., Camerer, C. F. and Phelps, E. A.: 2012, Emotion regulation reduces loss aversion and decreases amygdala responses to losses, *Social Cognitive and Affective Neuroscience* **forthcoming**.

Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F. and Phelps, E. A.: 2009, Thinking like a trader selectively reduces individuals loss aversion, *Proceedings of the National Academy of Sciences* **106**(13), 5035–5040.

Todorov, V.: 2010, Variance risk premium dynamics: The role of jumps, *Review of Financial Studies* **23**, 345–383.

Van Nieuwerburgh, S. and Veldkamp, L.: 2010, Information acquisition and underdiversification, *Review of Economic Studies* **77**(2), 779–805.

Appendix: Participant Instructions

Welcome to our financial decision making study!

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

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

Details for the Investment Choice and Investment Evaluation Task:

Specific details for the GAIN condition:

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

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

Specific details for the LOSS condition:

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

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

In both GAIN and LOSS conditions:

In each condition, at the beginning of each block of 6 trials, you do not know which type of stock the computer selected for that block. You may be facing the good stock, or the bad stock, with equal probability.

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

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

After that we will ask you to tell us two things: (1) what you think is the probability that the stock is the good one (the answer must be a number between 0 and 100 - do not add the % sign, just type in the value) (2) how much you trust your ability to come up with the correct probability estimate that the stock is good. In other words, we want to know how confident you are that the probability you estimated is correct. (answer is between 1 and 9, with 1 meaning you have the lowest amount of confidence in your estimate, and 9 meaning you have the highest level of confidence in your ability to come up with the right probability estimate)

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

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

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

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

Details for the Investment Evaluation Task:

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

You final pay for completing the investment tasks will be:

$23 + 1/10 * Investment Payoffs + 1/10 * Number of accurate probability estimates,

where Investment Payoffs = Dividends of securities you chose in the experiment, in both the GAIN and the LOSS conditions.

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

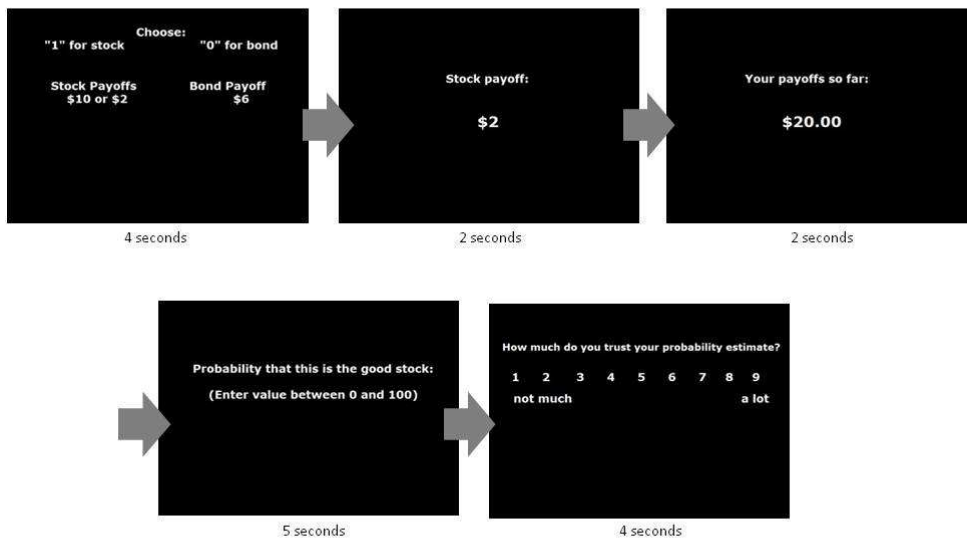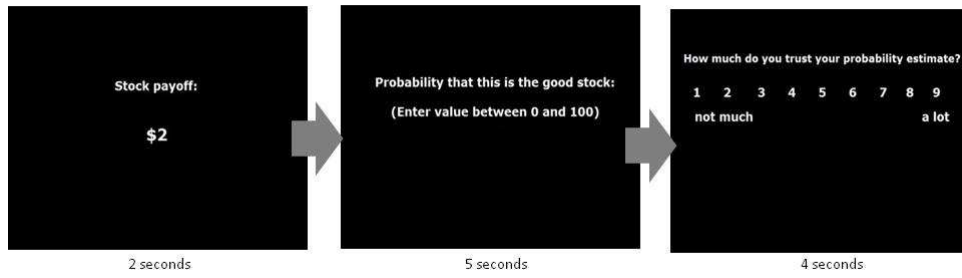Gain Condition - Active Involvement

Choose:
"1" for stock   "0" for bond

Stock Payoffs   Bond Payoff
$10 or $2 $6

4 seconds

Stock payoff:

$2

2 seconds

Your payoffs so far:

$20.00

2 seconds

Probability that this is the good stock:

(Enter value between 0 and 100)

5 seconds

How much do you trust your probability estimate?

1   2   3   4   5   6   7   8   9
not much   a lot

4 seconds

Figure 1: Active task: Gain condition

Loss Condition - Active Involvement

Choose:
"1" for stock   "0" for bond

Stock Payoffs   Bond Payoff
-$2 or -$10   -$6

4 seconds

Stock payoff:

-$10

2 seconds

Your payoffs so far:

-$2.00

2 seconds

Probability that this is the good stock:

(Enter value between 0 and 100)

5 seconds

How much do you trust your probability estimate?

1   2   3   4   5   6   7   8   9
not much   a lot

4 seconds

Figure 2: Active task: Loss condition

27

Figure 3: Passive task: Gain condition and Loss condition

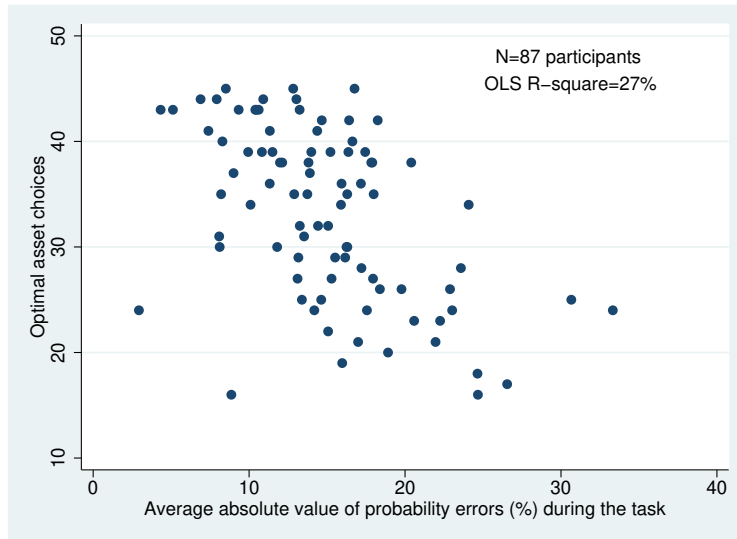Figure 4: Learning differs across conditions.

Figure 5: Correlation between the size of probability estimation errors and the number of optimal asset choices made by the 87 participants. Optimal choices refer to those of a risk-neutral Bayesian agent.


Figure 6: Left: Gain & Loss learning, within-subject correlation. Right: Active & Passive learning, within-subject correlation.
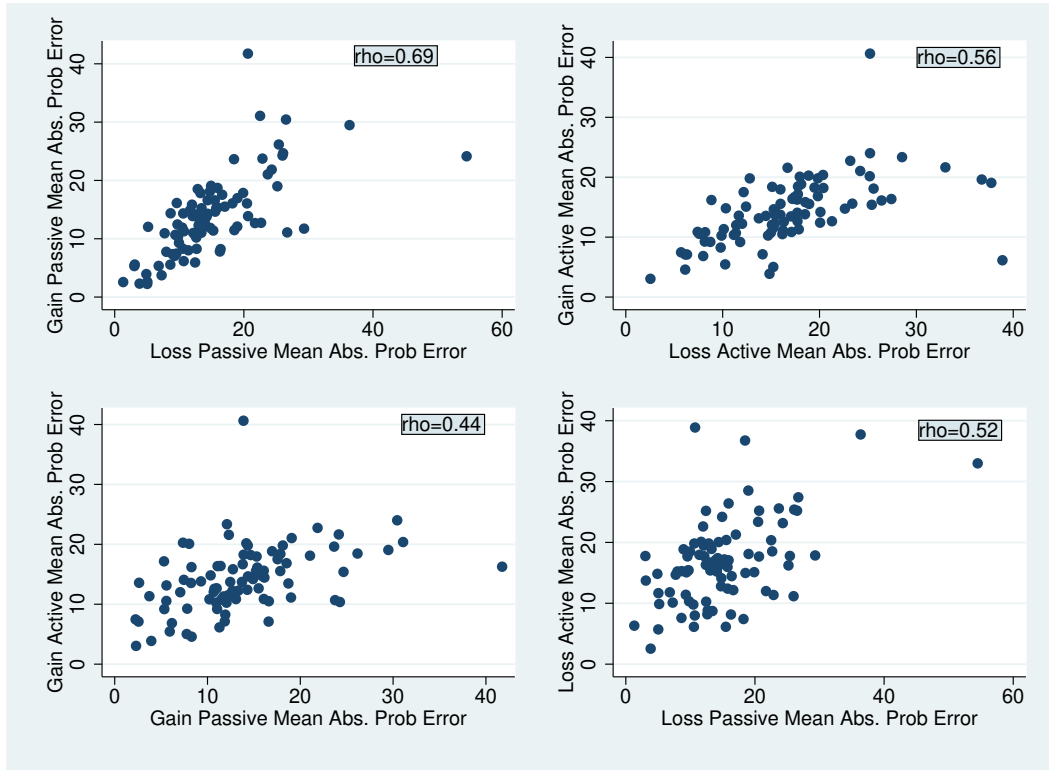
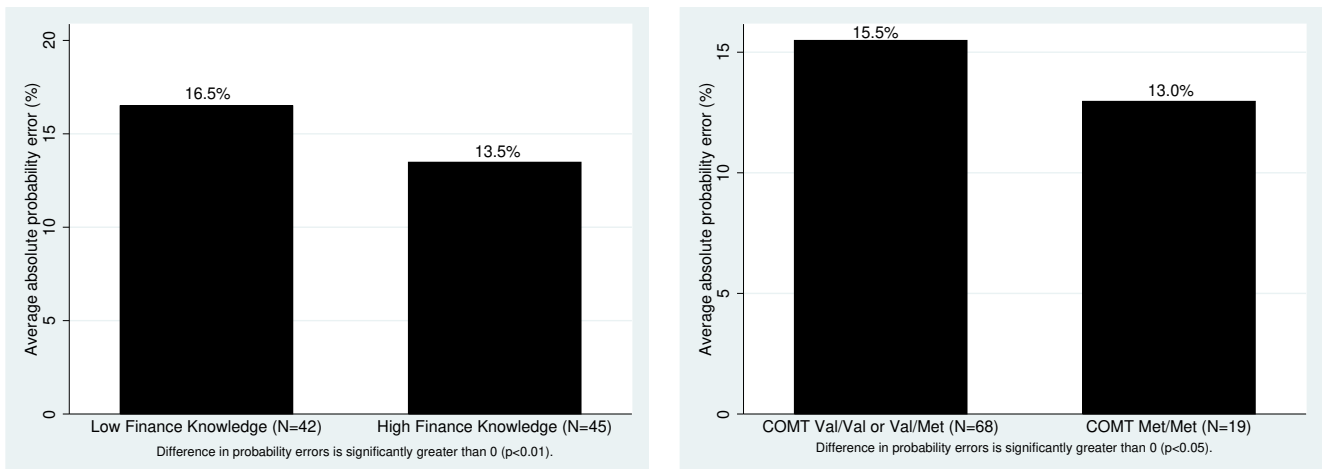Figure 7: Learning across domains: within-subject correlations.

Figure 8: Left: Average probability estimation errors are significantly higher for the 42 participants with a low finance knowledge score (i.e., those who made at least one mistake when answering the portfolio return question asked after the experimental task), relative to the 45 participants with a high finance knowledge score (i.e., those who made no mistake answering the portfolio return question). Right: Average probability estimation errors are significantly higher for the 68 participants with other variants of the $COMT$ genotype, relative to the 19 participants who have the Met/Met genotype. These differences are significant at conventional levels ($p < 0.05$ or better).
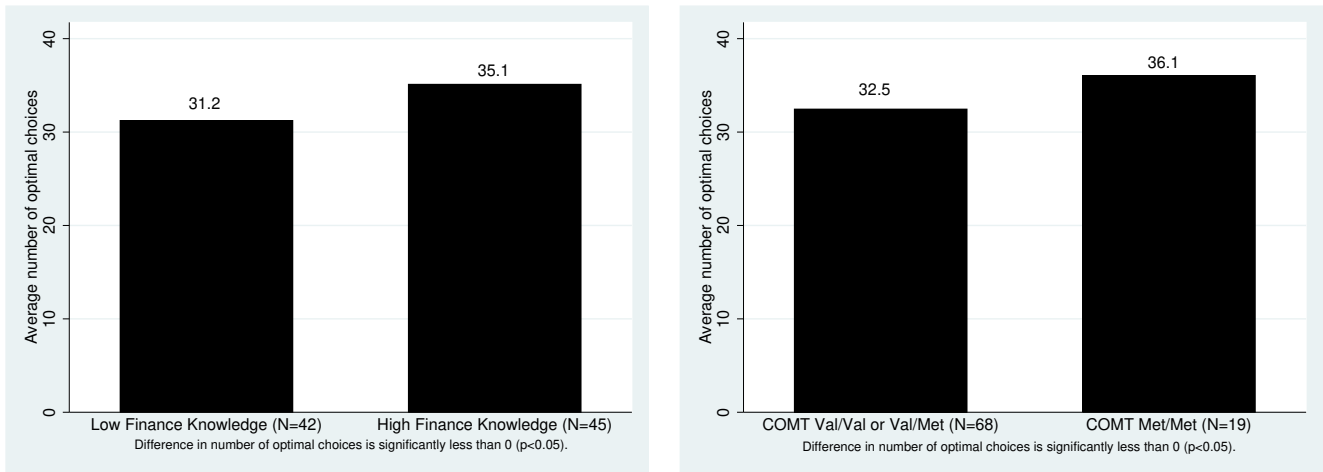
Figure 9: Left: The average number of optimal choices is significantly lower for the 42 participants with a low finance knowledge score (i.e., those who made at least one mistake when answering the portfolio return question asked after the experimental task), relative to the 45 participants with a high finance knowledge score (i.e., those who made no mistake answering the portfolio return question). Right: The average number of optimal choices is significantly lower for the 68 participants with other variants of the $COMT$ genotype, relative to the 19 participants who have the Met/Met genotype. These differences are significant at conventional levels ($p < 0.05$).

Table 1: Summary statistics.

Gender	Age (years)	Ethnicity	Financial knowledge	$COMT$ genotype
Male: 37	Mean: 20	Caucasian:44	Score of 3: 45	Met/Met: 19
Female: 50	St. Dev.: 1.6	Asian: 20	Score of 2: 24	Val/Met: 34
	Range: 18-29	African-American: 8	Score of 1: 17	Val/Val: 34
		Indian: 8	Score of 0: 1	
		Hispanic: 7		
$N = 87$ participants				

Table 2: Probability estimation errors and their dependence on context and type of information received. The dependent variable in these OLS regressions is the absolute value of the probability estimation error made by participant i in a given trial t (i.e., the absolute value of the difference between the subjective probability estimate and the objective Bayesian probability that the stock pays dividends from the better distribution, given the information seen so far by the participant). Constant term included, omitted here for brevity. Robust standard errors are clustered by subject.

Dependent variable				$Absolute\ Probability\ Error_{it}$			
	All trials	Active trials only	Passive trials only	All trials	Gain trials only	Loss trials only	Trials with objective posteriors $\geq 50\%$
$Loss\ trial_{it}$	1.86	2.56	1.16				4.31
	$(3.88)^{***}$	$(3.95)^{***}$	$(1.83)^{*}$				$(6.71)^{***}$
$Passive\ trial_{it}$				−1.06	−0.36	−1.76	
				$(−1.72)^{*}$	$(−0.51)$	$(−2.23)^{**}$	
$Adj.\ R^2$	0.004	0.007	0.001	0.001	0.000	0.003	0.02
Observations	10377	5177	5200	10377	5193	5184	5938

Table 3: Distribution of learner types. Each participant is assigned to one of the four types below, based on whether their average of the absolute values of probability errors committed is higher in the Loss condition vs. the Gain condition, or in the Active vs. Passive trials. For example, the table shows that 36 out of 87 participants (i.e., 41%), learn worse during Loss and during Active trials, relative to the other types of trials.

	Number of participants		
	Worse learning in Loss trials	Worse learning in Gain trials	
Worse learning in Active trials	36 (41%)	18 (21%)	Total: 54
Worse learning in Passive trials	21 (24%)	12 (14%)	Total: 33
	Total: 57	Total: 30	Overall: 87

Table 4: Heterogeneity of learning performance across different conditions.

Standard Deviation of Absolute Probability Errors			
Loss condition	15.85%	Active trials	15.45%
Gain condition	13.99%	Passive trials	14.39%
Difference	1.86%***	Difference	1.06%**

Table 5: Predictors of learning errors. The dependent variable in these OLS regressions is the absolute value of the probability estimation error (expressed as percentage points) made by participant i in a given trial t. Ethnicity fixed effects are included. Standard errors are robust to heteroskedasticity and clustered by participant.

Dependent variable	Absolute Probability Error$_{it}$						
	All trials	All trials	Active trials only	Passive trials only	Gain trials only	Loss trials only	Loss, high prior ($\geq 50\%$) trials only
COMT Met/Met$_i$		−2.34	−1.87	−2.85	−2.19	−2.61	−4.88
		(−2.01)**	(−1.66)*	(−1.98)*	(−2.00)**	(−1.91)*	(−2.40)**
Finance Knowledge$_i$	−1.83	−1.67	−1.41	−1.96	−1.46	−1.85	−1.39
	(−2.43)**	(−2.24)**	(−2.13)**	(−1.99)**	(−2.19)**	(−1.91)*	(−1.23)
Male$_i$	−1.49	−1.62	−2.60	−0.57	−1.59	−2.06	−4.36
	(−1.35)	(−1.48)	(−2.27)**	(−0.41)	(−1.43)	(−1.61)	(−2.24)**
Age$_i$	0.28	0.26	0.12	0.39	0.29	0.27	0.77
	(0.78)	(0.75)	(0.36)	(0.87)	(0.87)	(0.66)	(1.10)
ObjectiveProbability$_{i,t-1}$	0.04	0.04	0.06	0.02	−0.01	0.09	0.29
	(3.15)***	(3.20)***	(3.77)***	(1.32)	(−0.62)	(5.08)***	(5.25)***
Ethnicity Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj R^2	0.030	0.034	0.036	0.039	0.031	0.050	0.069
Observations	10377	10377	5177	5200	5193	5184	1848