

The Influence of Affect on Beliefs, Preferences and Financial Decisions*

Camelia M. Kuhnen[†]

Brian Knutson[‡]

September 2008

Abstract

Recent research in neuroeconomics suggests that the same brain areas that generate emotional states are also involved in the processing of information about risk, rewards and punishments. These findings imply that emotions may influence financial decisions in a predictable and parsimonious way. Our evidence suggests that affect – generated either by exogenous manipulations, or endogenously by outcomes of prior actions – indeed matters for financial risk taking, and that it does so by changing preferences as well as the belief formation process. Positive and arousing emotional states such as excitement induce people to take more risk, and to be more confident in their ability to evaluate the available investment options, relative to neutral states, while negative emotions such as anxiety have the opposite effects. Moreover, beliefs are updated in a way that is consistent with the self-preservation motive of maintaining positive affect and avoiding negative affect, by not fully taking into account new information that is at odds with the individuals' prior choices. Therefore, characteristics of markets, economic policies or organization design that have an impact on emotional brain circuits may influence decision making and affect important outcomes at the individual and aggregate level.

*The authors are grateful to Paola Sapienza and to participants at the Fall 2007 NBER Corporate Finance meeting, the 2007 European Summer Symposium in Financial Markets in Gerzensee, Switzerland, and at the Managerial Economics & Decision Science seminar at the Kellogg School of Management for helpful comments. All remaining errors are ours.

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

[‡]Department of Psychology, Stanford University, 470 Jordan Hall, Stanford, CA 94305-2130, knutson@psych.stanford.edu

I. Introduction

Recent work in economics (Elster (1998), Caplin and Leahy (2001)) has proposed that emotions may play a role in decision making under risk, following a significant stream of work in the psychology literature (Loewenstein et al. (2001)). The evidence gathered so far is supportive of this claim. For instance, environmental factors that have been shown to influence people's mood seem to correlate with stock market returns (Saunders (1993), Hirshleifer and Shumway (2003)). Moreover, investing itself is an activity that induces strong emotional responses, even when the individuals involved are professional traders (Lo and Repin (2002)). These empirical findings, however, do not allow us to distinguish whether emotions impact behavior by changing risk preferences, or the belief formation process, or both.

A microfoundation for these effects has started to emerge from the neuroscience literature. It has been shown that parts of the brain that generate emotional states and help guide behavior in primitive circumstances (such as seeking food or avoiding predators) are also important for the processing of information about monetary rewards, punishments, as well as risk. These brain areas are present in all mammals, and are referred to as the brain's limbic system. The role of this old (in an evolutionary sense) system is to allow the organism to react quickly to cues or changes in the environment that are important for survival.

Two components of this system in particular have been found to be engaged in decision making under risk: the nucleus accumbens and the anterior insula. The nucleus accumbens is crucial for the processing of information about gains or rewards, and for motivating the individual to approach potentially rewarding cues in the environment (Breiter et al. (2001), Knutson et al. (2001)). Activation in this structure has been associated with experiencing a state of excitement, also referred to as positive arousal (Bjork et al. (2004)). The anterior insula has been implicated in the processing of information about losses or punishments, and in the avoidance of aversive cues, and its activation leads to experiencing states of anxiety, also referred to as negative arousal (Chua et al. (1999), Simmons et al. (2004)). Moreover, activation in these two structures encodes information about the risk entailed by available choices (Preuschoff et al. (2006), Preuschoff et al. (2008)).

Prior work has shown that activation in these two areas helps predict whether individuals choose risky or riskless investment opportunities, even after controlling for the effect of informational and wealth-related variables that should be the determinants of this choice. Specifically, higher activation in the nucleus accumbens is associated with

a higher likelihood of switching to risky assets, while higher activation in the anterior insula is associated with switching to holding riskless assets. Excessive activation in these brain areas is followed by taking on too much or too little risk relative to the optimal choice of an expected utility maximizing, Bayesian updating agent (Kuhnen and Knutson (2005)). The relationship between activation in these emotional areas of the brain and the riskiness of the chosen investments seems to be causal. Exogenously increasing the nucleus accumbens activation (by presenting non-monetary visual stimuli characterized by strong positive affect) before a financial decision causes the subjects to make riskier investments (Knutson et al. (2008)).

The goal of this paper is two-fold. First, we would like to find whether emotions change financial choices by modifying risk preferences, or beliefs, or both. In our experimental setting, we observe financial choices and also elicit subjective beliefs about the payoffs of the available investments. Changes in affective states are either induced exogenously, or are the result of the investing activity itself. Second, we propose a parsimonious framework for explaining multiple types of deviations from rational choice, or behavioral biases. We build this framework on the existing results regarding the role of the brain's limbic system on decision making.

Our main insight is that emotions matter for decision making under risk, whether they are caused by exogenous factors or induced by outcomes of past choices. In particular, we argue that two emotional states, excitement (or positive arousal) and anxiety (or negative arousal), impact our risk preferences and the way we learn.

The psychology literature has shown that feedback about outcomes of past choices generates strong affective reactions, which in turn “have automatic and pervasive effects on performance even on tasks other than the one that induced the affect” (Kluger and DeNisi (1996)). Moreover, neuroscience findings point to the limbic system as being the source of these reactions. Specifically in the context of feedback about decisions under risk, it has been shown that activation in the nucleus accumbens increases when we learn that the outcome of a past choice was better than expected (Delgado et al. (2000), Pessiglione et al. (2006)).¹ Activation in the anterior insula increases when the outcome is worse than expected (Seymour et al. (2004), Pessiglione et al. (2006)), and when actions not chosen have larger payoffs than the chosen one (Kuhnen and Knutson

¹The neurotransmitter responsible for these reactions to rewards is dopamine. See Schultz (2007) for a review of the neuroscience literature on the role of dopamine in learning, and Bernheim and Rangel (2004) and Caplin and Dean (2008) for theoretical models of choice based on the properties of the dopamine system. Whether or not another neurotransmitter is responsible with learning from punishments (or unexpected losses) is still debated. One possible candidate for this role is serotonin (Daw et al. (2002)).

(2005)). Importantly, the increase in activation in these areas is positively correlated with the size of the surprise associated with the outcome, also referred to as the prediction error, suggesting that learning is dependent on the brain's limbic system.

We build on these prior findings by proposing that emotions such as excitement and anxiety – or equivalently, changes in activation of the nucleus accumbens and of the anterior insula – modify risk preferences, as well as the learning process itself.

First, we hypothesize that risk aversion is diminished by excitement, and increased by anxiety, consistent with evidence from Loewenstein et al. (2001), Kuhnen and Knutson (2005) and Knutson et al. (2008). The results in Knutson et al. (2008) suggest that nucleus accumbens activation modulates risk preferences. In that experiment, subjects knew the probability distributions of the outcomes of the two risky lotteries they could choose from, and thus their beliefs were fixed. Hence, shifts in risk taking behavior that followed the exogenous affect manipulation – switching from the low risk to the high risk lottery — could not be attributed to changes in beliefs about the probability distribution of outcomes, but rather to a change in risk preferences. In the current paper we find that, controlling for beliefs, the propensity to take risks is increased by positive affect and decreased by negative affect, as predicted.

Second, we hypothesize that beliefs are formed in accordance with the current emotional state, which is generated either by exogenous manipulations, or by outcomes of past choices. For instance, the learning process will differ if an individual experiences a high or a low outcome, one that is better or worse than that of other investments (or actions) not chosen, or if the outcome is confirming or disconfirming of prior choices. These events generate either excitement or anxiety, respectively, by triggering activation in the limbic system, and influence the encoding of the new information conveyed by the event. Since the limbic system allows the individual to engage in self-preservation behaviors (such as approaching food, or avoiding predators) it is possible that beliefs are formed with the same self-preservation motive. Specifically, the limbic system may generate beliefs strategically, in order to maintain positive affect or avoid negative affect. This is akin to selecting beliefs in order to eliminate cognitive dissonance (see Akerlof and Dickens (1982)). We find that the belief updating process indeed depends on how new information matches with prior beliefs and choices. People do not fully incorporate news about investment options that seem to be at odds with their prior actions. We also find that events that generate positive affect, such as instances where new information matches the subjects' prior choices, lead them to be more confident in their ability to identify the quality of the available investment options.

These findings suggest that affect can influence financial risk taking when feedback

and updating are involved. The influence of feedback has been suggested by prior findings related to the phenomenon of myopic loss aversion (Benartzi and Thaler (1995)), in which people take less financial risk if they know that they will receive more frequent feedback about their investment outcomes. Gneezy and Potters (1997) validate this phenomenon in an experimental setting and conjecture that when investors expect less frequent feedback they anticipate fewer opportunities to experience losses, and thus are more willing to make risky investments in advance. This conjecture is consistent with the notion that emotion plays a role in risk taking, and that an absence of future feedback allows a more positive affective state prior to choice.

Affect has also been linked to learning, or the updating of beliefs. Charness and Levin (2005) argue that in situations when reinforcement strategies (i.e. rules-of-thumb such as the “win-stay / lose-shift” heuristic²) compete with Bayesian updating strategies, people are more likely to make choices according to reinforcement strategies. They test this hypothesis in an experimental setting, and find strong evidence to support it. Importantly, the authors also find that if the subjects’ pay does not depend on the outcome of their choices, they are more likely to adopt Bayesian updating strategies and less likely to use the rule-of-thumb reinforcement strategies. One can interpret these findings as indicating that removing affective input from feedback (i.e. not paying the subjects based on the outcome of their choice) helps people behave in a more rational way.

The present findings also contribute to the literature focusing on the link between mood and stock returns (Saunders (1993), Hirshleifer and Shumway (2003)), between overconfidence and trading (Barber and Odean (2000), Gervais and Odean (2001), Grinblatt and Keloharju (2006)), and between overconfidence and managerial decisions (Heaton (2002), Malmendier and Tate (2005), Gervais et al. (2005), Ben-David et al. (2007)) by suggesting that affect (exogenous or generated by past outcomes) may be the source of patterns in financial choices that have been documented in this prior work.

II. Experimental Design

To examine whether exogenous and endogenous affect impacts beliefs, preferences or both during decision making under risk, we designed the Beliefs Task, in which subjects repeatedly chose to invest in a risky or a riskless security while learning about the

²This heuristic refers to the tendency to repeat actions that have been successful and avoid those that were not, irrespective of the implications of past outcomes for the optimal strategy going forward. If, for instance, past choices are rewarding but at the same time they deplete the source of the reward (e.g. picking fruit from a tree), an individual acting according to reinforcement learning would not even consider the fact that he has already depleted that resource.

distribution of payoffs of the risky security.

Before each choice, subjects saw a geometric shape followed by a picture. The shape and picture had no connection to the investment choices faced by the subject. Pictures belonged to three different categories: highly arousing and positive (e.g. erotic scenes), highly arousing and negative (e.g. rotten food), and neutral (e.g. a book sitting on the floor). Subjects were instructed to pay attention to the pictures because they would be asked questions about them after the investment task was over. Each geometric shape was associated with one type of picture. A square was always followed by a positive picture, a circle was always followed by a neutral picture, and a triangle was always followed by a negative picture. Presenting these geometric shapes before the actual pictures increases the anticipatory affect, since subjects know which type of image to expect before they actually see it.

After making all the investment decisions, subjects were asked to rate the pictures on arousal and valence dimensions. Arousal was measured on a scale from 1 to 9, while valence was measured on a scale from -4 to 4. The arousal scale measures how aroused or activated people feel at a given time. At the far right on this scale, individuals are feeling very aroused, activated, charged or energized, physically or mentally. At the far left, people are feeling completely unaroused, slow, still, deenergized, physically and mentally. The valence scale measures how positive or negative people are feeling. At the far left people are feeling very negative: unhappy, upset, irritated, frustrated, angry, sad, fearful or depressed. At the far right people are feeling very positive: happy, pleased, satisfied, competent, proud, content or delighted. In the middle of this scale people are feeling completely neutral, neither positive or negative.

Figure 2 shows that subjects rated positive pictures as being more positive and arousing than neutral pictures, and negative pictures as being more negative and arousing than neutral pictures ($p < 0.001$). Thus, the exogenous visual cues induced the desired affective state.

Moreover, we infer that the presentation of positive pictures increased nucleus accumbens activation, while the presentation of negative pictures increased insula activation, based on prior evidence that nucleus accumbens activation correlates with the positive arousal rating that subjects assign to cues (Bjork et al. (2004), Martinez et al. (2003)), and that anterior insula activation correlates with negative arousal ratings (Paulus et al. (2003), Simmons et al. (2004)).

28 subjects (14 male) were recruited via email announcements to the Kellogg School of Management subject pool and by posting fliers on the Northwestern University campus. The instructions given to subjects are contained in the appendix. All participants signed

informed consent forms before doing the experiments, as required by the Institutional Review Board at Northwestern University.

The experiment took 1.5 hours to complete and the average pay per subject was \$24.91. The task contained 90 trials, during which subjects had to decide whether to invest in a risky security (stock) or in a riskless security (bond). The timeline of one such trial is shown in Figure 1.

On any trial, if subjects chose to invest in the bond, they would get \$3 for sure at the end of the trial. If they chose to invest in the stock, they would receive a dividend which was either \$10 or -\$10. The stock could be paying these dividends according to one of two probability distributions, one of which was better than the other, in the sense of first-order stochastic dominance. If the stock was “good” then the probability of receiving the \$10 dividend was 75% and the probability of receiving the -\$10 dividend was 25%. If the stock was “bad” then the probability of receiving the \$10 dividend was 25% and the probability of receiving the -\$10 dividend was 75%. At the beginning of each block of five trials, the computer selected at random which type of stock the subject would face. The dividends paid by the stock were independent from trial to trial, but were drawn from the same distribution during all five trials in each block.

Subjects saw the dividend paid by the stock, whether they chose the stock or the bond in a particular trial. After making the choice between the two assets and learning the dividend paid by the stock, at the end of each trial subjects were asked two questions:

(1) “What you think is the probability that the stock is the good one?”

(2) “How much you trust your ability to come up with the correct probability estimate that the stock is good?”

At the end of the investment task, we added \$1 to the subjects’ earnings for every time they provided us with a probability estimate that was within 5% of the correct value (for example, that happened if the correct probability was 80% and the subject’s answer was between 75% and 85%). During the task, however, no feedback was provided about the accuracy of the probability estimates elicited from participants.

Throughout the task, subjects were informed how much money they had accumulated through dividends paid by the assets chosen up to that point. The final pay was calculated as the sum of a \$15 show-up fee and one twentieth of the task earnings, where task earnings included the dividends accumulated through investing in the two assets and the amount earned by providing accurate probability estimates.

All this information was known by subjects before the experiment started. The only thing that participants did not know was which type of stock they would encounter at the beginning of each block of five trials.

III. Results

A. Affect and choice

The evidence we find is consistent with the hypothesis that exogenous and endogenous affect cues can change risk taking behavior, in line with our predictions.

On average, after subjects saw a negative picture, they were less likely to choose the risky asset relative to trials when neutral pictures were shown (see Figure 3). The fraction of trials when the risky asset was chosen was 0.36 when negative cues were presented, 0.41 when the neutral cues were presented, and 0.40 when positive cues were presented. Mean comparison tests between the fraction of times people made the risky investment when exposed to negative versus positive cues, or negative versus neutral cues, indicate a significantly reduced likelihood of choosing the stock in the negative cue condition ($p < 0.05$). However, overall, subjects were not more likely to choose the risky asset after seeing a positive picture relative to seeing a neutral one.³

Table I presents the results of a multivariate probit model of asset choice. The dependent variable, $StockChoice_t$, is equal to 1 if the subject chose to invest in the stock on trial t , and 0 otherwise.

An important determinant of choice is the probability that the stock is paying from the good dividend distribution. If our subjects are risk-neutral for the small stakes involved in each asset choice, and maximize the expected value of the dividend to be received at the end of each trial, they should only pick the stock if the probability that it is the good one is at least 80%. When this condition is not met, the bond should be chosen. We therefore include as an independent variable the objective, Bayesian value of this probability based on the history of dividends paid up to and including the prior trial ($ObjectiveProbability_{t-1}$). The objective probability that the stock is good after observing k dividend payments of \$10 in the past n trials in the block (and thus $(n - k)$ dividend payments of -\$10) is: $Prob\{GoodStock|k \text{ successes in } n \text{ trials}\} = \frac{1}{1+3^{n-2k}}$.

Since the subjects' posterior beliefs are not necessarily equal to the objective posterior, we estimate the probit model also using these elicited subjective probabilities ($ProbabilityEstimate_{t-1}$), and also control for the confidence that subjects put in their probability estimates ($ConfidenceEstimate_{t-1}$).

Given the simple decision rule dictated by expected utility maximization, the probability of choosing the stock should only depend on whether the prior probability (objec-

³While negative pictures reduced risk taking in this study, in a prior experiment, positive pictures increased risk taking (Knutson et al. (2008)). This difference may be due to the use of more potent negative stimuli in the present experiment.

tive or subjective) that the stock is good is greater than 80%. We find, however, that we can explain significantly more variation in subjects' choices if we allow the probability of choosing the stock to be a linear function of prior beliefs, instead of being a step-function. Hence, in interests of space, in the analysis reported here we use the better-fit specification.

As independent variables for stock versus bond choice we also include picture type (captured by dummies $PositiveCue_t$ and $NegativeCue_t$, the omitted category being the neutral pictures), a dummy indicating whether the dividend paid by the stock on the prior trial was \$10 ($HighStockDividend_{t-1}$), and the past choice of the subject ($StockChoice_{t-1}$). All these variables relate to the affective state of the individual at the time of choice, as indicated by prior neuroeconomics work.

To control for wealth and time effects on choice we include in the model the overall earnings in the experiment up to that trial ($TotalEarnings_{t-1}$), and block number fixed-effects. We include subject fixed effects to take into account the possibility that some people are more risk seeking than others, and cluster the standard errors by subject.

The results in Table I indicate that subjects were 7% less likely to choose the risky asset after the presentation of negative pictures, relative to the presentation of neutral ones, especially if their prior choice was the bond. Positive pictures increased the likelihood that the stock was chosen by 6%, if the prior choice was the stock (this effect, however, is only marginally significant).

Moreover, the dividend paid by the stock in the prior trial influenced the asset choice. If the stock paid the high (\$10) dividend in the prior trial, controlling for the subjective belief that the stock was the good one, subjects were 11% more likely to switch from holding the bond to holding the stock, relative to trials where the prior dividend was low (-\$10). Interestingly, the effect of the prior stock dividend on the current choice is driven out of the model when the objective probability that the stock is good is used as a control variable, instead of the subjective posterior.

As expected, the objective and subjective beliefs that the stock was good given the dividend history were positive and significant predictors of $StockChoice_t$. If subjects estimated that the stock they faced had a higher probability of being the good stock (as measured by $ProbabilityEstimate_{t-1}$), they were more likely to select that stock in the subsequent trial. Additionally, if subjects were more confident in their probability estimate (as measured by $ConfidenceEstimate_{t-1}$), their estimate was a stronger predictor of choosing the stock in the next trial. The objective, Bayesian probability that the stock is good ($ObjectiveProbability_{t-1}$) also robustly predicted subsequent stock choice.

It is important to establish, though, that subjective beliefs influence choice on their

own, and not simply because they may correlate with the objective belief. To check whether this is the case, we estimate the probit model of stock choice by simultaneously including $ProbabilityEstimate_{t-1}$ and $ObjectiveProbability_{t-1}$ among the right-hand side variables. As can be seen in Table I, the subjective probability estimate continues to be a positive and significant predictor of choosing the stock, even when controlling for the objective probability.

The results also show the existence of time and wealth effects. Subjects were less likely to choose a stock in later blocks of the task (as indicated by the coefficients on the block fixed-effects dummies, not reported here for space reasons). Also, especially after a bond choice, individuals with more money earned on the task up to that point were more likely to switch to the stock.

Hence, since the results in Table I indicate that affect-inducing stimuli change choices even after controlling for beliefs, we infer that they must change risk preferences, with negative affect leading to higher risk aversion and positive affect leading to more risk seeking.

B. Affect and beliefs

We find that subjective beliefs tracked objective probabilities relatively well, indicating that subjects understood and paid attention to the task. Although subjects tended to overestimate small probabilities and underestimate large probabilities, their estimates closely approximated the correct values, as can be seen in Figure 4. Interestingly, however, the amount of confidence subjects had in their probability estimates was highest for extreme low and high probabilities, as shown in Figure 5.

The multivariate regression⁴ results in Table II also show that the subjective beliefs ($ProbabilityEstimate_t$) are highly positively correlated with the objective, Bayesian posterior ($ObjectiveProbability_t$). If the latter increases by 10%, the subjective posterior increases by 6.1%. The difference between these two probabilities can come from subjects using a different updating formula than Bayes' rule, or can be caused by imperfect recall of past outcomes. The task design allows us to disentangle the learning effect from the memory effect⁵, since the value of the Bayesian posterior that the stock is good on trial t only depends on the value of the prior held in trial $t - 1$ and on the stock dividend paid in trial t , and therefore can be obtained without having to recall

⁴As in all models estimated in the paper, here we include subject fixed effects and cluster the standard errors by subject.

⁵Recent evidence indicates that affect may also impact memory formation, but the underlying mechanisms are not yet fully known.

several past dividends.⁶ For any given belief elicited from the subject in the prior trial ($ProbabilityEstimate_{t-1}$), we can compute the correct Bayesian posterior using that prior and the news about the dividend paid by the stock in the current trial t . We refer to this variable as $BayesianPosteriorUsingSubjectivePrior_t$. As can be seen from Table II, this quantity is also a positive and significant predictor of the actual posterior belief declared by the subject. Its predictive power is greatest if subjects declared to have high confidence in their probability estimates in trial $t - 1$, as would naturally be expected.

However, since there is not a one-to-one mapping between the subjective and the correct posteriors, Bayes' rule is a only rough predictor of the way subjects update beliefs, leaving room for other factors to impact learning.

While we do not find a significant impact of positive or negative pictures on the probability estimates reported by subjects, we document that feedback about past choices has effects on these subjective beliefs that are consistent with our predictions.

As Table II shows, relative to subjects who chose the bond in trial t , those individuals who chose the stock (i.e. those for whom $StockChoice_t = 1$) declare a probability estimate that the stock is the good one that is about 10% higher. This effect is present even after controlling for the objective belief that this is the good stock, or for the Bayesian posterior formed using the subject's own probability estimate elicited at the end of the prior trial. As the probit model of stock choice in Table I indicates, subjects select assets based on their priors. The results in Table II suggest that the fact that a subject chose a particular asset will make that asset even more desirable to the person *after* the choice, even when controlling for what the subject believed about the asset right before he decided to add it to his portfolio.

This suggests that people may choose their beliefs strategically, to maintain positive affect or avoid negative affect.⁷ All else equal, after learning the dividend paid by the stock those individuals who chose the stock in the beginning of a particular trial will be more optimistic that the stock is good, relative to people who chose the bond, in order to validate their choice and feel good about it.

To further investigate the role of past choices on beliefs, we split the data based on whether the stock or the bond were chosen in a given trial, and analyzed how subjects in these two conditions updated their beliefs upon the arrival of the dividend news. For

⁶If the prior probability that that stock is good is p , then following a high (\$10) dividend, the posterior obtained using Bayes' rule is $\frac{3p}{2p+1}$. The posterior after a low (-\$10) dividend is $\frac{p}{3-2p}$.

⁷Arguably, this can be seen an instance of the endowment effect, but in our task subjects were not endowed with the stock or bond. They had to decide every trial which of the two assets they wanted to hold, for that trial only.

any prior belief expressed by the subject we calculated the correct Bayesian posterior, and compared that to the posterior actually declared by the subject. Several patterns emerge, as seen in Figure 6.

First, as also indicated by the results in Table II, stock holders provide higher posterior estimates than bond holders, independent of their prior beliefs and the actual dividend paid by the stock.

Second, the results in Figure 6 suggest that updating extreme priors (very low or very high subjective probability estimates that the stock is good) is different from updating moderate priors. There is a significant overreaction to news that oppose the extreme prior. For example, if a subject believes that the probability the stock is good is low (less than 20%), and then a high dividend is realized, he will form a much higher posterior that this is the good stock than is warranted by the news. The elicited posterior is 22% higher than it should be, as indicated by the summary statistics in Table III. If the subject has an extreme high prior belief (greater than 80%) that the stock is good, and then the stock pays a low dividend, the subjective posterior is 20% lower than what it should be conditional on the prior, and on the new dividend (see Table III). All these effects are statistically significant ($p < 0.001$).

Third, for moderate priors, subjective posterior beliefs are conservative. That is, the information about the new dividend is not fully incorporated to generate a posterior that matches that produced by Bayes' rule when applied to the subject's prior belief. This can be a manifestation of the well-known conservatism bias (Phillips and Edwards (1966)), whereby people do not update their prior beliefs enough when presented with new information. Table IV presents the summary statistics for these errors in updating. For stock holders and bond holders, the elicited posteriors are lower by 10% and 13%, respectively, than the Bayesian posterior, if the stock pays a high dividend. If the stock pays a low dividend, stock and bond holders do not revise down their priors enough. Their elicited posterior probabilities are higher than they should be by 10% and 7%, respectively. These effects are also statistically significant ($p < 0.001$).

Fourth, the results in Figure 6 indicate that this conservatism in updating is asymmetric, in that it depends on the choice the subject made in the beginning of the trial. This is a novel result and is not an implication of conservatism bias per se. Specifically, stock holders especially ignore news about low dividends, and bond holders especially ignore news about high dividends. This asymmetry is also supported by the evidence on the size of the learning errors presented in Table IV as a function of asset choice and dividend type. We formally test it by estimating an OLS regression where the dependent variable is the absolute value of the learning error ($|BayesianPosteriorUsingSubjectivePrior_t -$

$ProbabilityEstimate_t$), first for trials where the stock paid a low dividend, and then for those where the the dividend was high. The independent variable of interest is the dummy indicating whether in the beginning of the trial the subject chose the stock or not ($StockChoice_t$). We include subject fixed effects to take into account the possibility that some people make larger size errors in updating than others, and cluster the standard errors by subject. The results are presented in Table V. The size of the probability updating error is 3% higher for stock holders than for bond holders if the dividend is low, and it is 3% higher for bond holders relative to stock holders if the dividend is high. These differences are statistically significant ($p < 0.05$).

This asymmetry is consistent with our hypothesis that subjects choose beliefs that match their prior choices in order to avoid the negative affect caused by the admission of making a mistake. Importantly, this effect is different from confirmation bias (Mynatt et al. (1977)), which refers to people’s tendency to choose to consult information sources that can help confirm a particular hypothesis, instead of seeking those that would help reject it. People acting in accordance with confirmation bias simply have a preference for how to explore information sources, but when they obtain explicit falsifying information, they use it to reject incorrect hypotheses. In contrast to confirmation bias, our results speak to an error in the process of rejecting incorrect hypotheses. Specifically, we show that individuals underweight or ignore information that is in disagreement with their past choice.

We also find that affect changes the confidence in one’s estimates, and not just the probability estimates themselves. We report these effects in the OLS regression in Table VI. Since the unconditional results in Figure 5 revealed a U-shaped relationship between confidence in one’s probability estimate and the probability itself, in the model of confidence in Table VI we include as controls the subjective probability estimate and its square (Panel A) and the objective probability and its square (Panel B). The model includes subject fixed effects, and standard errors are clustered by subject.

Confidence levels, declared using a scale from 1 to 9, are 0.17 higher ($p < 0.07$) in trials where an exogenous positive picture was presented, relative to those where neutral pictures were used. Negative cues are also positively associated with confidence, but the effect is not statistically significant. In other words, exogenously-induced excitement leads to more confidence in one’s probability estimation ability.

Moreover, confidence is significantly higher (by 0.2 to 0.7, depending on specification, $p < 0.01$) in trials where the dividend paid by the stock matches the choice just made by the subject, than in trials where it does not, as indicated by the positive coefficient on the interaction term between $HighStockDividend_t$ and $StockChoice_t$. These instances

occur when the subject chooses the stock and then the dividend paid is high, or when the bond is chosen and then the dividend paid is low. As discussed earlier, such events are likely to be characterized by a state of endogenous positive affect, as the subject's choice is validated by the new dividend information.

Hence, the results indicate that positive affect induces the subjects to be more confident in their ability to identify the quality of the available risky investments.

IV. Discussion

The evidence presented here suggests that emotion impacts decision making under risk. Affect may matter whether it is induced exogenously by an experimenter, a policy maker, or by institutional features of markets and the environment in which the choice is made, or whether it is generated by the outcomes of prior decisions. Positive affect makes us more risk seeking, and more confident in our beliefs. To maintain positive affect and avoid negative affect we ignore new information that is opposed to our actions, and as a result, our learning is flawed. These findings have implications for several areas of finance and economics.

A. Modeling learning

The learning model used in virtually all economic theories is Bayes' rule. However, as suggested by our results, and those in Charness and Levin (2005), exogenous factors and outcomes of past choices may change the way people update their beliefs. We propose that this influence is mediated by activation in the brain's limbic system, in a predictable and parsimonious way. Related to this idea is the work of Caplin and Dean (2008) on the role of the neurotransmitter dopamine in economic choice. This neurotransmitter is the source of activation in the nucleus accumbens, the brain's reward area, and has been shown to encode the prediction error generated by outcomes, which is defined as the difference between the actual outcomes and their ex-ante expected value. Caplin and Dean (2008) develop the axiomatic foundations of economic choice that would be necessary for dopamine to indeed encode prediction errors. However, the approach in that paper explicitly abstracts from a learning environment and thus does not provide a model of belief updating per se.

B. Incentives for exploration

Since affect induces errors in our learning process, the optimal contract offered by a principal to an agent to induce efficient exploration of new strategies may need to take emotions into account. For example, agents may ignore bad news about a new project that they have chosen to start, in order to avoid experiencing negative affect. This error in learning may not be there, for instance, if the news is perceived by the agent to only speak to the quality of the project, but not to the quality of the agent himself. In general, contracts should undo the agents' tendency towards ignoring information that seems to contradict their prior actions, as well as their tendency to be more confident in their project selection abilities when the new information matches their prior choices.

C. Individual differences

Individuals vary in their reaction to affect cues. Some are influenced more by positive cues in the environment, while others are affected more by negative cues, depending in part on the reactivity of their limbic system. These effects are likely to work on chief executive officers (CEOs) and other top corporate decision makers, just as on any other individuals, and therefore may have a significant impact on firm value. We can speculate that some CEOs might be likely to become overconfident after getting good news (which may be unrelated to the firm's operations) but not react as much to bad news. Other CEOs might react more to bad news (again, perhaps orthogonal to the firm) relative to good news.

Hence, firm policies regarding investment or capital structure may depend on the environment (i.e. the type of cues available to the top decision maker, such as information about business conditions, or information unrelated to the firm) as well as on the personality of the CEO, or their emotional reactivity. Prior studies have shown the importance of CEO fixed-effects (Bertrand and Schoar (2003), Malmendier and Tate (2005)) on corporate policies. Our work can shed light on what these fixed-effects capture. They might provide a summary statistic of the CEO's risk aversion, his sensitivity to gains and to losses, or the responsiveness of his beliefs to events that generate affect. Understanding the importance of these components of CEO fixed-effects helps to optimally allocate managers to firms, since these managerial personal characteristics influences the quality of the CEO-firm match. One can envision equilibrium settings where firms with different growth opportunities are optimally matched with managers with different levels of risk aversion or different emotional reactivity to good or bad news. More broadly, our results suggest that it is important for top corporate decision makers to be aware of the role of

affect on financial choices, and of the types of information, or stimuli in the environment that may change one's affect and ultimately drive decisions that affect the entire firm.

D. Market and institutional design

Policy makers may wish to encourage people to save more for retirement by investing more in the stock market, casino owners may want to have their visitors gamble more money, and insurance companies may want to have their clients be more risk averse. All these desired behaviors can be induced with the appropriate affective manipulation. For instance, our results suggest that presenting ads, information or other types of stimuli that induce a state of positive arousal will cause people to take more risk. In the context of casinos, such institutional design features are already being employed. Visitors to casinos floors are surrounded by sights of potential rewards, such as free food and drinks offered by attractive individuals, that have nothing to do with the odds of winning at the roulette table. As shown by prior research, the presence of these cues triggers the brain's reward area and induces a state of positive arousal. According to our results, this will cause people to gamble more.

E. Asset bubbles and crashes

The relationship between affect and risk taking that we propose here suggests a possible explanation for asset bubbles and crashes. Positive returns in financial markets may induce a positive affective state and make investors more willing to invest in stocks, and more confident that they have chosen the right portfolio, which will lead to increased buying pressure and future positive returns. This effect would be even stronger when more individuals are already investing in the stock market, since we show that simply adding a risky asset to one's portfolio makes people be more favorable about that asset's future payoffs. After losses in the financial markets, investors may experience a state of negative affect which will reduce their willingness to take on more risk, and their confidence in their ability to choose stocks. This will create selling pressure and further negative stock returns. Such feedback effects can therefore lead to bubbles or crashes, depending on the starting point in these chains of events.

V. Conclusion

Recent theories have proposed that emotion can influence decision making under risk, and advances in technology enable experiments that can shed light on the mechanisms

responsible for these effects. Here, we build on existing neuroeconomic evidence that shows that the same brain areas that generate emotional states are also involved in the processing of information about risk, rewards and punishments.

We examine whether emotion indeed influences risk taking, and whether it does so by changing beliefs, preferences or both. In an investment selection task that requires subjects to update their beliefs about the return distribution of a risky asset, we find that events associated with positive and arousing emotions such as excitement lead to riskier choices, while those associated with negative and arousing emotions such as anxiety lead to more risk averse choices. Moreover, affect influences the belief formation process. Positive arousal increases the subjects' confidence in their ability to evaluate the risky investments they are faced with. Beliefs about these investments are updated in a way that is consistent with the self-preservation motive of maintaining positive affect and avoiding negative affect. Specifically, subjects do not fully incorporate news that contradict their prior choices, and form incorrect posterior beliefs.

Several important caveats limit the implications of our experimental findings. First, the effects documented here only speak directly to the influence of affect on fast decisions in an experimental setting, whereas outside the laboratory people have much more time to deliberate and make financial choices. Therefore, we can only speculate that our results apply to real life financial decisions. Moreover, we propose here a framework for understanding how affect changes risk preferences and beliefs that is based on neuroscience evidence on how our brain works during the process of economic choice. We do not, however, measure brain activation during our experiment, both because it is very costly to acquire such data, and also, because in doing so we would only replicate existing findings regarding activation patterns in the brain's emotional areas. We therefore rely on prior neuroeconomics findings to infer how the exogenous manipulations in the experiment, and the endogenous outcomes of past choices influence the brain mechanisms that we propose to be important in financial decision making.

These caveats notwithstanding, the evidence provided here suggests that characteristics of markets, economic policies or organization design that have an impact on emotional brain circuits may influence decision making under risk by changing both risk preferences, and the learning process. A more realistic theory of learning needs to account for the effects of emotions on beliefs. Reactivity to emotional events may be the source of individual differences observed in economic behavior, may be a factor to be considered in optimal contracting problems, and may also help predict fluctuations in asset prices.

Appendix A: Instructions for the Beliefs task

You will be able to make 90 investments in a risky asset (a stock) and in a riskless asset (a bond, or a savings account). On any trial, if you choose to invest in the bond, you get \$3 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 -\$10. 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 75% and the probability of receiving the -\$10 dividend is 25%. 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 75%, and the odds of it being -\$10 are 25%.

If the stock is bad then the probability of receiving the \$10 dividend is 25% and the probability of receiving the -\$10 dividend is 75%. 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 25%, and the odds of it being -\$10 are 75%.

At the beginning of each block of 5 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 safe asset and add \$3 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. The 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.

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%) then we will add \$1 to your task earnings at the end of the task.

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.

PAY: Your final pay for being in our experiment will be: Show-up fee + $1/20$ * Task Earnings, where the Task Earnings = (Dividends you accumulate through investing in the two assets PLUS money you earn by guessing correct probabilities). The show-up fee is \$15.

PICTURES: During each trial you will see a geometric shape and a picture before you make the investment decision for that trial. The shape and picture have no connection to the investment choice you are facing. However, we would like you to pay attention to them because we will ask you questions about them after the investment task is over.

References

- Akerlof, George, and William T. Dickens, 1982, The economic consequences of cognitive dissonance, *American Economic Review* 72, 307–319.
- Barber, Brad, and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773–806.
- Ben-David, Itzhak, John R. Graham, and R. Harvey, Campbell, 2007, Managerial overconfidence and corporate policies, *Working paper* .
- Benartzi, Shlomo, and Richard Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *The Quarterly Journal of Economics* .
- Bernheim, Douglas B., and Antonio Rangel, 2004, Addiction and cue-conditioned cognitive processes, *American Economic Review* 94.
- Bertrand, Marianne, and Antoinette Schoar, 2003, Managing with style: The effect of managers on firm policies, *Quarterly Journal of Economics* 118, 1169:1208.
- Bjork, J. M., B. Knutson, G. W. Fong, D. M. Caggiano, S. M. Bennett, and D. W. Hommer, 2004, Incentive-elicited brain activation in adolescents: Similarities and differences from young adults, *Journal of Neuroscience* 24, 1793–1802.
- Breiter, H. C., I. Aharon, D. Kahneman, A. Dale, and P. Shizgal, 2001, Functional imaging of neural responses to expectancy and experience of monetary gains and losses, *Neuron* 30, 619–639.
- Caplin, A., and J. Leahy, 2001, Psychological expected utility theory and anticipatory feelings, *The Quarterly Journal of Economics* 116, 55–79.
- Caplin, Andrew, and Mark Dean, 2008, Dopamine, reward prediction error, and economics, *Quarterly Journal of Economics* 663:701.
- Charness, Gary, and Dan Levin, 2005, When optimal choices feel wrong: A laboratory study of Bayesian updating, complexity, and affect, *American Economic Review* 95, 1300–1309.
- Chua, P., M. Krams, I. Toni, R. Passingham, and R. Dolan, 1999, A functional anatomy of anticipatory anxiety, *NeuroImage* 9.

- Daw, Nathaniel D., Sham Kakade, and Peter Dayan, 2002, Opponent interactions between serotonin and dopamine, *Neural Networks* 15, 603–616.
- Delgado, M.R., L.E. Nystrom, C. Fissell, D.C. Noll, and J.A. Fiez, 2000, Tracking the hemodynamic responses to reward and punishment in the striatum, *Journal of Neurophysiology* 84, 3072–3077.
- Elster, J., 1998, Emotions and economic theory, *Journal of Economic Literature* 36, 47–74.
- Gervais, Simon, J.B. Heaton, and Terrance Odean, 2005, Overconfidence, investment policy and managerial welfare, *Working paper* .
- Gervais, Simon, and Terrance Odean, 2001, Learning to be overconfident, *Review of Financial Studies* 14, 1–27.
- Gneezy, Uri, and Jan Potters, 1997, An experiment on risk taking and evaluation periods, *The Quarterly Journal of Economics* 112, 631–645.
- Grinblatt, Mark, and Matti Keloharju, 2006, Sensation seeking, overconfidence and trading activity, *NBER working paper* .
- Heaton, J. B., 2002, Managerial optimism and corporate finance, *Financial Management* 33–45.
- Hirshleifer, David, and Tyler Shumway, 2003, Good day sunshine: Stock returns and the weather, *Journal of Finance* 58.
- Kluger, Avraham N., and Angelo DeNisi, 1996, The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory, *Psychological Bulletin* 119, 254–284.
- Knutson, B., C. M. Adams, G. W. Fong, and D. Hommer, 2001, Anticipation of increasing monetary reward selectively recruits nucleus accumbens, *Journal of Neuroscience* 21.
- Knutson, Brian, Elliott Wimmer, Camelia M. Kuhnen, and Piotr Winkielman, 2008, Nucleus accumbens activation mediates reward cues’ influence on financial risk-taking, *NeuroReport* 19, 509–513.
- Kuhnen, Camelia M., and Brian Knutson, 2005, The neural basis of financial risk taking, *Neuron* 47, 763–770.

- Lo, Andrew W., and Dmitry V. Repin, 2002, The psychophysiology of real-time financial risk processing, *Journal of Cognitive Neuroscience* 14, 323–339.
- Loewenstein, G., E.U. Weber, and C.K. Hsee, 2001, Risk as feelings, *Psychological Bulletin* 267–86.
- Malmendier, Ulrike, and Geoffrey Tate, 2005, CEO overconfidence and corporate investment, *Journal of Finance* 60, 2661–2700.
- Martinez, D., M. Slifstein, A. Broft, O. Malawi, D. R. Hwang, Y. Huang, T. Cooper, L. Kegeles, E. Zarahn, and A. Abi-Dargham, 2003, Imaging human mesolimbic dopamine transmission with positron emission tomography. Part II: Amphetamine-induced dopamine release in the functional subdivisions of the striatum, *Journal of Cerebral Blood Flow and Metabolism* 28, 285–300.
- Mynatt, C.R., M.E. Doherty, and R.D. Tweney, 1977, Confirmation bias in a simulated research environment: An experimental study of scientific inference, *The Quarterly Journal of Experimental Psychology* 29, 85–95.
- Paulus, M. P., C. Rogalsky, and A. Simmons, 2003, Increased activation in the right insula during risk-taking decision making is related to harm avoidance and neuroticism, *NeuroImage* 19, 1439–1448.
- Pessiglione, Mathias, Ben Seymour, Guillaume Flandin, Raymond J. Dolan, and Chris D. Frith, 2006, Dopamine-dependent prediction errors underpin reward-seeking behavior in humans, *Nature* 442, 1042–1045.
- Phillips, L., and W. Edwards, 1966, Conservatism in a simple probability inference task, *Journal of Experimental Psychology* 72, 346–354.
- Preuschoff, Kerstin, Peter Bossaerts, and Steven R. Quartz, 2006, Neural differentiation of expected reward and risk in human subcortical structures, *Neuron* 51, 381–390.
- Preuschoff, Kerstin, Steven R. Quartz, and Peter. Bossaerts, 2008, Human insula activation reflects risk prediction errors as well as risk, *Journal of Neuroscience* 28, 2745–2752.
- Saunders, Edward M. Jr, 1993, Stock prices and Wall Street weather, *American Economic Review* 83, 1337–1345.
- Schultz, Wolfram, 2007, Behavioral dopamine signals, *TRENDS in Neuroscience* 30.

Seymour, Ben, John P. O’Doherty, Peter Dayan, Martin Koltzenburg, Anthony K. Jones, Raymond J. Dolan, Karl J. Friston, and Richard S. Frackowiak, 2004, Temporal difference models describe higher order learning in humans, *Nature* 429.

Simmons, A., S. C. Matthews, M. B. Stein, and M. P. Paulus, 2004, Anticipation of emotionally aversive visual stimuli activates right insula, *NeuroReport* 15, 2261–2265.

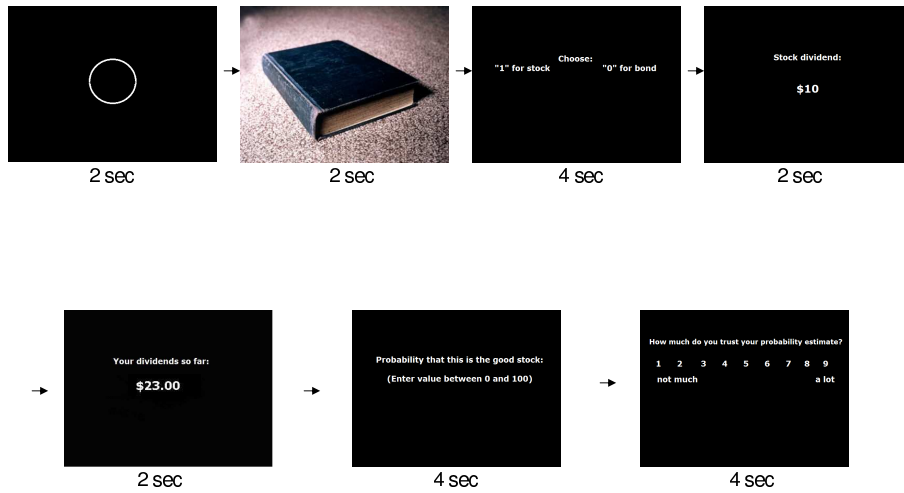


Figure 1
Trial structure of the Beliefs Task.

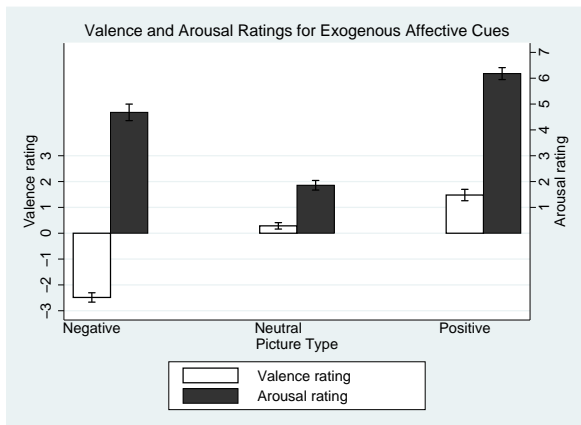


Figure 2
Valence and arousal ratings for exogenous affective visual cues

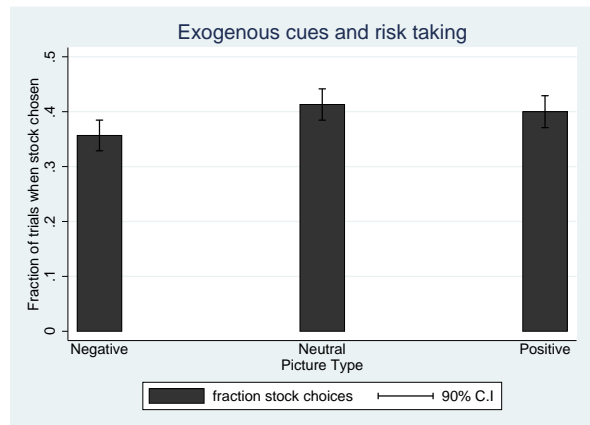


Figure 3
Exogenous affective cues and risky investments

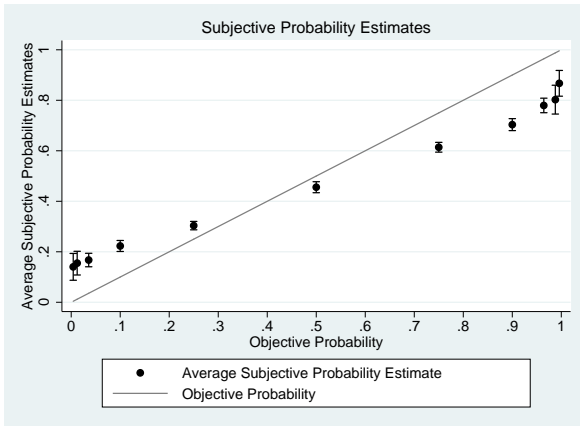


Figure 4
Subjective probability estimates

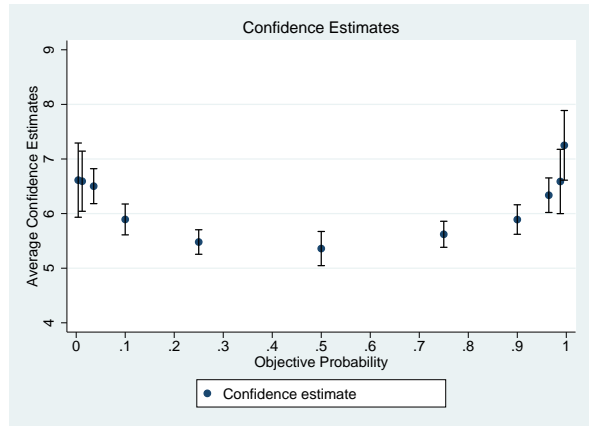


Figure 5
Confidence estimates

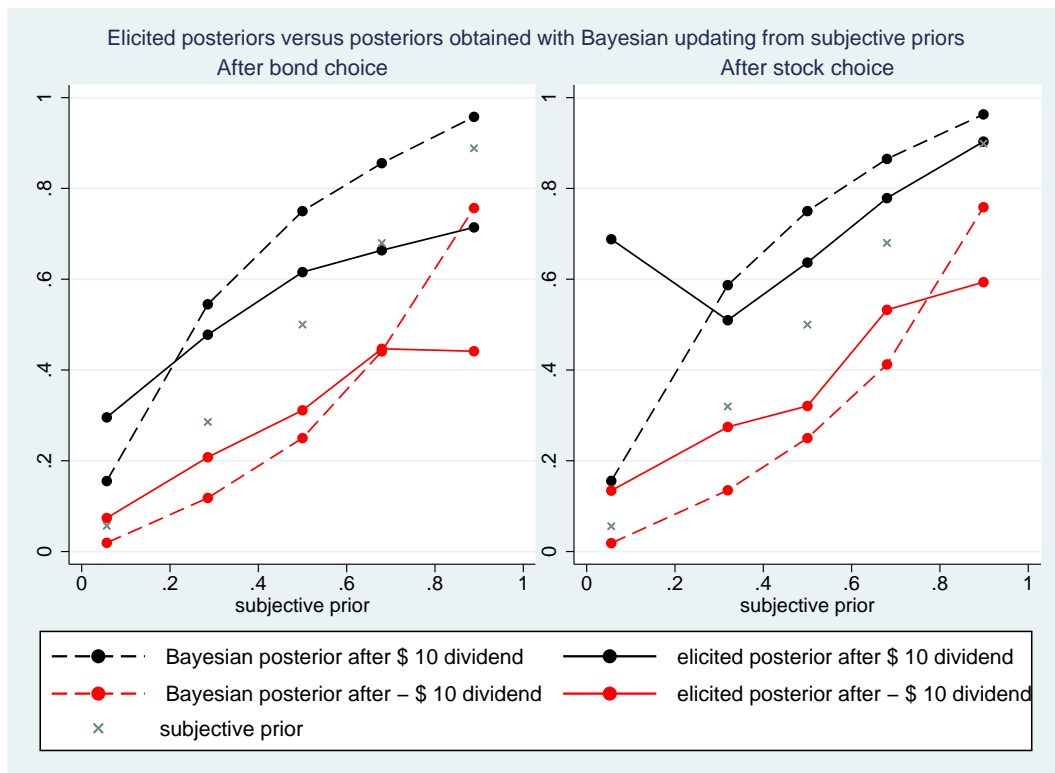


Figure 6
Dependence of the learning process on prior actions and beliefs

Table I

Probit model for choosing the stock versus the bond

The dependent variable, $StockChoice_t$, is an indicator variable equal to 1 if the subject chose to hold the stock in trial t , and 0 if the bond was chosen. $PositiveCue_t$ and $NegativeCue_t$ are indicator variables equal to 1 if the exogenous visual cue presented at the beginning of the trial was a positive or a negative one, respectively. $HighStockDividend_t$ is an indicator variable equal to 1 if the dividend paid by the stock in trial t is \$10, and zero if the dividend is -\$10. $TotalEarnings_t$ is the total amount of money earned by the subject up to and including trial t . $ProbabilityEstimate_t$ is the subjective probability that the stock is the good one, elicited at the end of trial t . $ObjectiveProbability_t$ is the Bayesian posterior belief that the stock is good, conditional on the history of dividends paid by the stock up to and including trial t . $ConfidenceEstimate_t$ is the subjective confidence (1 through 9, with 1 being the lowest confidence) in the elicited posterior probability estimate that the stock is the good one. The table reports marginal effects. Standard errors are clustered by subject (t-statistics are shown in parentheses).

Dependent variable $StockChoice_t$	Previous choice was the bond (i.e. $StockChoice_{t-1} = 0$)			Previous choice was the stock (i.e. $StockChoice_{t-1} = 1$)			All data		
$PositiveCue_t$	-0.04 (-1.25)	-0.04 (-1.24)	-0.05 (-1.32)	0.07 (1.90)*	0.06 (1.78)*	0.06 (1.55)	-0.02 (-0.62)	-0.02 (-0.68)	-0.02 (-0.66)
$NegativeCue_t$	-0.07 (-2.24)**	-0.07 (-2.15)**	-0.07 (-2.30)**	-0.06 (-1.03)	-0.06 (-1.13)	-0.07 (-1.21)	-0.09 (-2.08)**	-0.09 (-2.10)**	-0.09 (-2.23)**
$HighStockDividend_{t-1}$	0.02 (0.45)	0.01 (0.15)	0.11 (2.99)***	-0.05 (-0.89)	-0.07 (-1.12)	0.05 (0.84)	0.00 (0.01)	-0.02 (-0.40)	0.11 (2.13)**
$TotalEarnings_{t-1}$	0.03 (2.68)***	0.03 (2.78)***	0.03 (3.08)***	0.02 (1.38)	0.02 (1.33)	0.02 (1.23)	0.03 (2.04)**	0.03 (2.10)**	0.03 (2.29)**
$ObjectiveProbability_{t-1}$	0.61 (7.12)***	0.46 (5.12)***		0.69 (4.59)***	0.56 (3.79)***		0.77 (6.23)***	0.60 (5.06)***	
$ProbabilityEstimate_{t-1}$		0.28 (3.54)***	0.15 (0.91)		0.26 (1.41)	0.52 (2.29)**		0.34 (2.56)**	0.38 (2.14)**
$ConfidenceEstimate_{t-1}$			-0.03 (-2.13)**			-0.01 (-0.66)			-0.03 (-1.74)*
$ProbabilityEstimate_{t-1} \times$ $ConfidenceEstimate_{t-1}$			0.06 (2.41)**			0.01 (0.16)			0.05 (1.41)
$StockChoice_{t-1}$							0.16 (3.94)***	0.14 (3.68)***	0.17 (5.15)***
Subject fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Block fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R^2	0.263	0.275	0.246	0.244	0.249	0.225	0.307	0.315	0.291
Observations	1443	1443	1443	970	970	970	2413	2413	2413

Table II
Elicited posterior probability estimates

The dependent variable, $ProbabilityEstimate_t$, is the elicited, subjective posterior probability that the stock is the good one. $PositiveCue_t$ and $NegativeCue_t$ are indicator variables equal to 1 if the exogenous visual cue presented at the beginning of the trial was a positive or a negative one, respectively. $StockChoice_t$ is an indicator variable equal to 1 if the subject chose to hold the stock in trial t , and 0 if the bond was chosen. $ObjectiveProbability_t$ is the Bayesian posterior belief that the stock is good, conditional on the history of dividends paid by the stock up to and including trial t . $BayesianPosteriorUsingSubjectivePrior_t$ is the value of the posterior belief that the stock is good obtained by updating the subject's prior belief using the new dividend information according to Bayes' rule. $ConfidenceEstimate_t$ is the subjective confidence (1 through 9, with 1 being the lowest confidence) in the elicited posterior probability estimate that the stock is the good one. Standard errors are clustered by subject (t-statistics are shown in parentheses).

Dependent Variable:	Panel A	Panel B	Panel C
$ProbabilityEstimate_t$			
$PositiveCue_t$.01 (.55)	.01 (.67)	.01 (.40)
$NegativeCue_t$	-.01 (-.81)	-.00 (-.26)	-.00 (-.35)
$StockChoice_t$.07 (3.69)***	.10 (5.69)***	.09 (5.53)***
$ObjectiveProbability_t$.61 (11.59)***		
$BayesianPosteriorUsingSubjectivePrior_t$.62 (12.48)***	.41 (4.87)***
$BayesianPosteriorUsingSubjectivePrior_t$ x $ConfidenceEstimate_{t-1}$.03 (3.09)***
$ConfidenceEstimate_{t-1}$			-.02 (-3.67)***
Subject fixed effects	Yes	Yes	Yes
Adj. R^2	.585	.550	.558
Observations	2413	2413	2413

Table III
Errors in updating extreme priors

The error in probability updating ($BayesianPosteriorUsingSubjectivePrior_t - ProbabilityEstimate_t$) as a function of the stock dividend, for extreme low (i.e. $ProbabilityEstimate_{t-1} \in [0, 0.2]$), extreme high (i.e. $ProbabilityEstimate_{t-1} \in [0.8, 1)$) and moderate (i.e. $ProbabilityEstimate_{t-1} \in (0.2, 0.8)$) subjective prior beliefs. *** indicates that the estimate is significantly different than zero ($p < 0.001$).

	$StockDividend_t = -\$10$	$StockDividend_t = \$10$
Extreme low subjective prior	-0.06***	-0.22 ***
Moderate subjective prior	-0.08***	0.11 ***
Extreme high subjective prior	0.20***	0.10 ***

Table IV
Conservatism bias

The error in probability updating ($BayesianPosteriorUsingSubjectivePrior_t - ProbabilityEstimate_t$) as a function of the subject's choice and the stock dividend, for moderate subjective prior beliefs (i.e. $ProbabilityEstimate_{t-1} \in (0.2, 0.8)$). $StockChoice_t$ is an indicator variable equal to 1 if the subject chose to hold the stock in trial t , and 0 if the bond was chosen. *** indicates that the estimate is significantly different than zero ($p < 0.001$).

	$StockDividend_t = -\$10$	$StockDividend_t = \$10$
$StockChoice_t = 0$	-0.07***	0.13***
$StockChoice_t = 1$	-0.10***	0.10***

Table V
Strategic belief formation

The dependent variable is the absolute value of the error in probability updating ($|BayesianPosteriorUsingSubjectivePrior_t - ProbabilityEstimate_t|$), for moderate subjective prior beliefs (i.e. $ProbabilityEstimate_{t-1} \in (0.2, 0.8)$). $StockChoice_t$ is an indicator variable equal to 1 if the subject chose to hold the stock in trial t , and 0 if the bond was chosen. Standard errors are clustered by subject (t-statistics are shown in parentheses).

Dependent Variable	$StockDividend_t = -\$10$	$StockDividend_t = \$10$
$ BayesianPosteriorUsingSubjectivePrior_t - ProbabilityEstimate_t $	Coef./t	Coef./t
$StockChoice_t$.03 (2.54)**	-.03 (-2.61)**
Subject fixed effects	Yes	Yes
Adj. R^2	.219	.164
Observations	927	857

Table VI
Subjective confidence estimates

The dependent variable, $ConfidenceEstimate_t$, is the subjective confidence (1 through 9, with 1 being the lowest confidence) in the elicited posterior probability estimate that the stock is the good one. $PositiveCue_t$ and $NegativeCue_t$ are indicator variables equal to 1 if the exogenous visual cue presented at the beginning of the trial was a positive or a negative one, respectively. $HighStockDividend_t$ is an indicator variable equal to 1 if the dividend paid by the stock in trial t is \$10, and zero if the dividend is -\$10. $StockChoice_t$ is an indicator variable equal to 1 if the subject chose to hold the stock in trial t , and 0 if the bond was chosen. $ObjectiveProbability_t$ is the Bayesian posterior belief that the stock is good, conditional on the history of dividends paid by the stock up to and including trial t . $ProbabilityEstimate_t$ is the subject's estimate of the probability that the stock is the good one. Standard errors are clustered by subject (t-statistics are shown in parentheses).

Dependent variable:	Panel A	Panel B
$ConfidenceEstimate_t$		
$PositiveCue_t$.17 (1.88)*	.18 (1.92)*
$NegativeCue_t$.13 (1.67)	.12 (1.50)
$HighStockDividend_{t-1}$	-.25 (-1.81)*	-.44 (-2.65)**
$HighStockDividend_t \times StockChoice_t$.73 (3.42)***	.74 (4.23)***
$StockChoice_t$	-.37 (-2.52)**	-.51 (-3.08)***
$ObjectiveProbability_t$	-3.56 (-2.55)**	
$ObjectiveProbability_t^2$	3.64 (2.57)**	
$ProbabilityEstimate_t$		-3.98 (-3.10)***
$ProbabilityEstimate_t^2$		4.83 (3.58)***
Subject fixed effects	Yes	Yes
Adj. R^2	.539	.548
Observations	2402	2402