

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

Camelia M. Kuhnen and Brian Knutson*

Abstract

Neuroeconomics research shows that brain areas that generate emotional states also process information about risk, rewards, and punishments, suggesting that emotions influence financial decisions in a predictable and parsimonious way. We find that positive emotional states such as excitement induce people to take risks and to be confident in their ability to evaluate investment options, while negative emotions such as anxiety have the opposite effects. Beliefs are updated so as to maintain a positive emotional state by ignoring information that contradicts individuals' prior choices. Marketplace features or outcomes of past choices may change emotions and thus influence future financial decisions.

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, Weber, Hsee, and Welch (2001)). The evidence gathered so far is supportive of this claim. For instance, environmental factors that have been shown to influence people's moods 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 influence behavior by changing risk preferences, or the belief formation process, or both.

*Kuhnen, c-kuhnen@kellogg.northwestern.edu, Kellogg School of Management, Northwestern University, 2001 Sheridan Rd., Evanston, IL 60208; Knutson, knutson@psych.stanford.edu, Department of Psychology, Stanford University, 470 Jordan Hall, Stanford, CA 94305. We are grateful to Colin Camerer, Michael Jensen (the referee), Paul Malatesta (the editor), Paola Sapienza, and Devin Shanthikumar; to participants at the Fall 2007 National Bureau of Economic Research Corporate Finance meeting, the 2007 European Summer Symposium in Financial Markets in Gerzensee (Switzerland), the 2008 Western Finance Association, and the 2009 Royal Economic Society meeting; and to seminar participants at Stanford University, the University of California at Berkeley, and Northwestern University. We thank the Zell Center for Risk Research at the Kellogg School of Management for providing generous funding for this project. All remaining errors are ours.

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 (e.g., seeking food or avoiding predators) are also important for the processing of information about monetary rewards and punishments, as well as risk. These brain areas are present in all mammals, and their role is to allow the organism to react quickly to cues or changes in the environment that are important for survival.

Two components of the emotional brain have been found to be particularly important for 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, Aharon, Kahneman, Dale, and Shizgal (2001), Knutson, Adams, Fong, and Hommer (2001)). Activation in this structure has been associated with experiencing positive emotions such as excitement (Bjork, Knutson, Fong, Caggiano, Bennett, and Hommer (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 negative emotions such as anxiety (Chua, Krams, Toni, Passingham, and Dolan (1999), Simmons, Matthews, Stein, and Paulus (2004)). Moreover, these two structures track information about the risk entailed by available choices (Preuschoff, Bossaerts, and Quartz (2006)).

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 nonmonetary visual stimuli characterized by strong positive emotions) before a financial decision causes the subjects to make riskier investments (Knutson, Wimmer, Kuhnen, and Winkielman (2008)).

The goal of this paper is 2-fold. First, we would like to find out whether emotions change financial choices by modifying risk preferences, or beliefs, or both. In an experimental setting, we observe financial choices and also elicit subjective beliefs about the payoffs of the available investments. Changes in emotional (or 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 emotional brain 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 2 emotional states, excitement and anxiety, influence 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 emotional brain 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, Nystrom, Fissell, Noll, and Fiez (2000), Pessiglione, Seymour, Flandin, Dolan, and Frith (2006)).¹ Activation in the anterior insula increases when the outcome is worse than expected (Seymour, O’Doherty, Dayan, Koltzenburg, Jones, Dolan, Friston, and Frackowiak (2004), Pessiglione et al.), and when actions not chosen have larger payoffs than the chosen one (Kuhnen and Knutson (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 emotional brain.

We build on these prior findings by proposing that affective states 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. We first hypothesize that risk aversion is diminished by excitement, and increased by anxiety, consistent with evidence from Loewenstein et al. (2001), Kuhnen and Knutson (2005), Knutson et al. (2008), and Gilad and Kliger (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 2 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 excitement and decreased by anxiety, as predicted.

We also 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 emotional brain areas, and influence the encoding of the new information conveyed by the event. Since these brain areas allow the individual to engage in self-preservation behaviors

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, Kakade, and Dayan (2002)).

(e.g., approaching food or avoiding predators), it is possible that beliefs are formed with the same self-preservation motive. Specifically, the emotional brain may generate beliefs strategically, in order to maintain a positive emotional state or avoid a negative state. 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 seems to be at odds with their prior actions. We also find that events that generate positive affective states, 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 emotions 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 the absence of future feedback allows a more positive affective state prior to choice. There is also some evidence that emotions influence learning. Individuals sometimes deviate from Bayesian updating by following suboptimal strategies (broadly referred to as reinforcement learning) such as the "win-stay/lose-shift" heuristic. This 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 this rule would not consider the fact that he has already depleted that resource. In an experimental setting, Charness and Levin (2005) find that if individuals' pay does not depend on the outcome of their choices, they are more likely to adopt Bayesian updating and less likely to use this suboptimal heuristic. Their finding suggests that removing emotional input from feedback (i.e., not paying individuals 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), and Grinblatt and Keloharju (2009)), and between overconfidence and managerial decisions (Heaton (2002), Malmendier and Tate (2005), Gervais, Heaton, and Odean (2005), and Ben-David, Graham, and Harvey (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 influences 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 distribution of payoffs of the risky security. Twenty-eight subjects (14 male) were recruited on the Northwestern University campus and were given detailed instructions, as shown in the Appendix. The experiment lasted 90 minutes.

Each subject made 90 investment decisions, choosing each time between a risky security (stock) and a riskless security (bond). Figure 1 shows the sequence of events involved in a typical trial. On any trial, if subjects chose to invest in the bond, they received \$3 for sure at the end of the trial. If they chose to invest in the stock, they received a dividend, which was either \$10 or $-\$10$. The stock could be paying these dividends according to 1 of 2 probability distributions, one of which was better than the other, in the sense of 1st-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 5 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 they were drawn from the same distribution during all 5 trials in each block. Subjects had 4 seconds to choose an asset in each trial. Immediately following the choice, subjects observed for 2 seconds the dividend paid by the stock, whether they chose the stock or the bond in that particular trial. In the following 2 seconds they saw the total value of dividends accumulated in the task up to that time. At the end of the trial subjects were asked 2 questions and had 4 seconds to answer each question:

- i) “What you think is the probability that the stock is the good stock?”
- ii) “How much you trust your ability to come up with the correct probability estimate that the stock is good?”

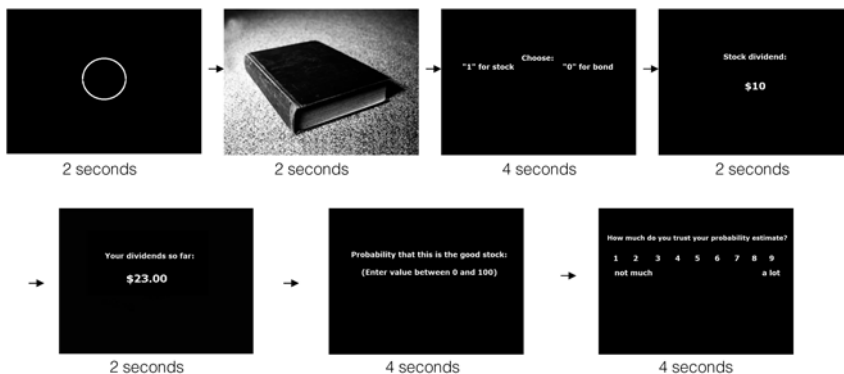
When the investment task was completed, we added \$1 to the subjects’ earnings for every time they provided us with a probability estimate that was within 5% of the correct value. For example, that happened if the correct probability was 80% and the subject’s answer was between 75% and 85%. During the task, however, no feedback was provided about the accuracy of the probability estimates elicited from participants. The final pay was calculated as the sum of a \$15 show-up fee and 1/20th of the task earnings, where task earnings included the dividends accumulated through investing in the 2 assets and the amount earned by providing accurate probability estimates.² The average final pay was \$24.91.

Before each investment decision, subjects saw a geometric shape followed by a picture for a total of 4 seconds. The shape and picture had no connection

²Subjects did not receive any payment for trials where they failed to make an asset choice or failed to answer our question about the value of their probability estimate.

FIGURE 1
Trial Structure of the Beliefs Task

In Figure 1, the sequence of events in a typical trial is presented.



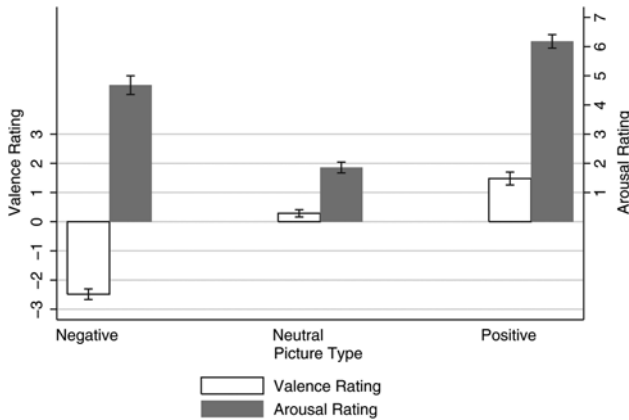
to the financial choice faced by the subject. Pictures were used to exogenously manipulate the emotional state of the subjects immediately prior to making each financial decision and belonged to 3 different categories: highly arousing and positive (e.g., erotic scenes), highly arousing and negative (e.g., 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 intensifies emotional reactions, since subjects know which type of image to expect before they actually see it. All information about the timing of each trial and the payoff structure was known by subjects before the experiment started.

After making all of the investment decisions, subjects were asked to rate the pictures on 2 dimensions that are widely used in the psychology literature to characterize emotional states: arousal and valence. Arousal measures how energized or activated people feel at a given time. Low arousal values indicate that people are feeling slow, still, or deenergized, whereas high arousal values indicate that individuals are feeling very activated, charged, or energized, physically or mentally. Valence indicates how positively or negatively people are feeling at a given time. Valence is negative when people are feeling unhappy, upset, irritated, frustrated, angry, sad, fearful, or depressed, and it is positive when they are feeling happy, pleased, satisfied, competent, proud, contented, or delighted. Arousal was measured on a scale from 1 to 9, while valence was measured on a scale from -4 to 4. 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 states of excitement (characterized by high arousal and positive valence) and anxiety (characterized by high arousal and negative valence). Moreover, we infer that the presentation of positive pictures

FIGURE 2

Valence and Arousal Ratings for Exogenous Affective Visual Cues

In Figure 2, subjects' ratings of negative, neutral, and positive pictures are shown along 2 dimensions: valence and arousal.



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, Slifstein, Broft, Malawi, Hwang, Huang, Cooper, Kegeles, Zarahn, Abi-Dargham, Haber, and Laruelle (2003)), and that anterior insula activation correlates with negative arousal ratings (Paulus, Rogalsky, Simmons, Feinstein, and Stein (2003), Simmons et al. (2004)).

III. Results

A. Affect and Choice

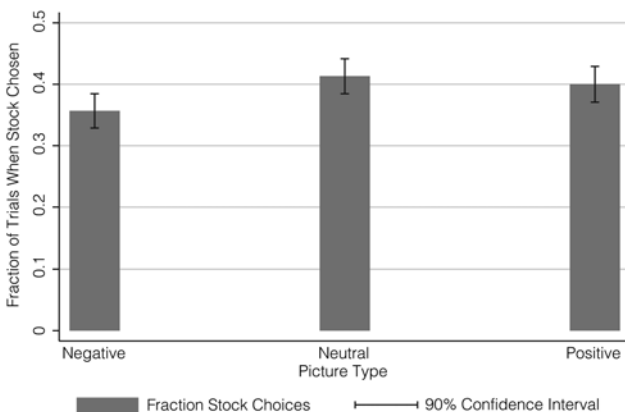
The evidence we find is consistent with the hypothesis that exogenous and endogenous affect cues can change risk-taking behavior, in line with our predictions. On average, after subjects saw a negative picture, they were less likely to choose the risky asset relative to trials when neutral pictures were 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 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 picture.³

Table 1 presents the results of a multivariate probit model of asset choice. The dependent variable, $STOCK_CHOICE_t$, is equal to 1 if the subject chose to

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.

FIGURE 3
Exogenous Affective Cues and Risky Investments

In Figure 3, the fractions of trials where subjects chose the risky asset after viewing a negative, neutral, or positive picture are shown.



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 stock 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 ($\text{OBJECTIVE_PROBABILITY}_{t-1}$). 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: $\text{Prob}\{\text{GOOD_STOCK} | k \text{ successes in } n \text{ trials}\} = 1 / (1 + 3^{n-2k})$. Since the subjects' posterior beliefs are not necessarily equal to the objective posterior, we estimate the probit model also using these elicited subjective probabilities ($\text{PROBABILITY_ESTIMATE}_{t-1}$), and also control for the confidence that subjects put in their probability estimates ($\text{CONFIDENCE_ESTIMATE}_{t-1}$).

Given the simple decision rule dictated by expected utility maximization, the probability of choosing the stock should only depend on whether the prior probability (objective 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 the 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 dummy variables POSITIVE_CUE_t and NEGATIVE_CUE_t , the omitted category being the neutral pictures), a dummy variable indicating whether the dividend paid by the stock on the prior trial was \$10 ($\text{HIGH_STOCK_DIVIDEND}_{t-1}$), and the past choice of the subject ($\text{STOCK_CHOICE}_{t-1}$). All these variables relate to the affective state of the individual at the time of choice,

TABLE 1

Probit Model for Choosing the Stock versus the Bond

The dependent variable, STOCK_CHOICE_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. POSITIVE_CUE_t and NEGATIVE_CUE_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 cue, respectively. HIGH_STOCK_DIVIDEND_{t-1} is an indicator variable equal to 1 if the dividend paid by the stock in trial *t* - 1 is \$10, and 0 if the dividend is -\$10. TOTAL_EARNINGS_{t-1} is the total amount of money earned by the subject up to and including trial *t* - 1. OBJECTIVE_PROBABILITY_{t-1} 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* - 1. PROBABILITY_ESTIMATE_{t-1} is the subjective probability that the stock is the good stock, elicited at the end of trial *t* - 1. CONFIDENCE_ESTIMATE_{t-1} 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 stock. Table 1 reports marginal effects. Standard errors are clustered by subject (*t*-statistics are shown in parentheses). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable STOCK_CHOICE _t	Previous Choice was the Bond (i.e., STOCK_CHOICE _{t-1} = 0)			Previous Choice was the Stock (i.e., STOCK_CHOICE _{t-1} = 1)			All Data		
	POSITIVE_CUE _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)
NEGATIVE_CUE _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)**
HIGH_STOCK_DIVIDEND _{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)**
TOTAL_EARNINGS _{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)**
OBJECTIVE_PROBABILITY _{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)***	
PROBABILITY_ESTIMATE _{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)**
CONFIDENCE_ESTIMATE _{t-1}			-0.03 (-2.13)**			-0.01 (-0.66)			-0.03 (-1.74)*
PROBABILITY_ESTIMATE _{t-1} × CONFIDENCE_ESTIMATE _{t-1}			0.06 (2.41)**			0.01 (0.16)			0.05 (1.41)
STOCK_CHOICE _{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	0.26	0.28	0.25	0.24	0.25	0.23	0.31	0.32	0.29
No. of obs.	1,443	1,443	1,443	970	970	970	2,413	2,413	2,413

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 ($TOTAL_EARNINGS_{t-1}$), and block number fixed effects. As in all models estimated in the paper, we include subject fixed effects and cluster the standard errors by subject.

The results in Table 1 indicate that subjects were 7% less likely to choose the risky asset after the presentation of negative pictures, relative to the presentation of neutral pictures, 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 stock, 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 $STOCK_CHOICE_t$. If subjects estimated that the stock they faced had a higher probability of being the good stock (as measured by $PROBABILITY_ESTIMATE_{t-1}$), 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 $CONFIDENCE_ESTIMATE_{t-1}$), their estimate was a stronger predictor of choosing the stock in the next trial. The objective, Bayesian probability that the stock is good ($OBJECTIVE_PROBABILITY_{t-1}$) 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 $PROBABILITY_ESTIMATE_{t-1}$ and $OBJECTIVE_PROBABILITY_{t-1}$ among the right-hand side variables. As can be seen in Table 1, 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 dummy variables, not reported here for space reasons). Also, especially after a bond choice, individuals with more money earned on the task up to that point (measured by the variable $TOTAL_EARNINGS_{t-1}$) were more likely to switch to the stock.

Hence, since the results in Table 1 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 (albeit at low significance levels) 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. The amount of confidence subjects had in their probability estimates was highest for extreme low and high probabilities, as shown in Figure 5.

FIGURE 4
Subjective Probability Estimates

In Figure 4, average subjective probability estimates are charted against objective probability estimates.

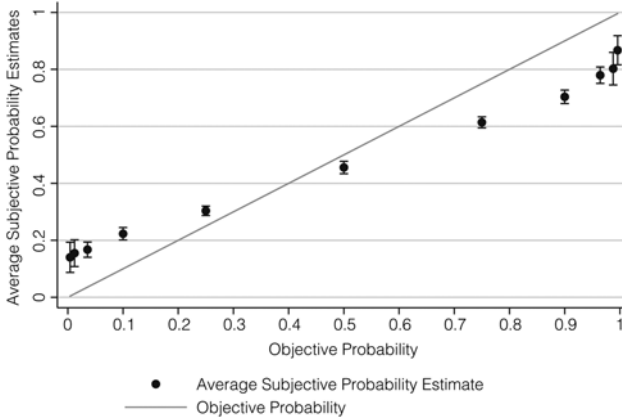
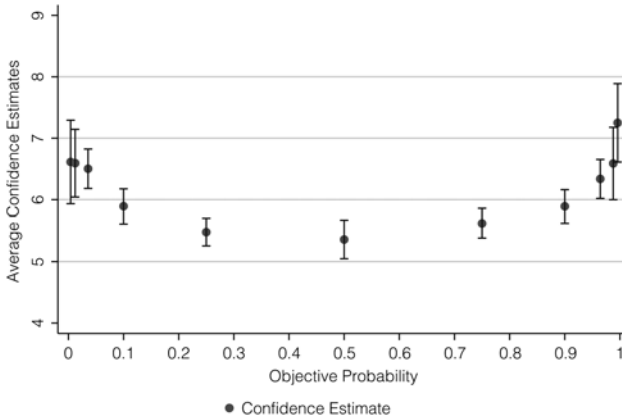


FIGURE 5
Confidence Estimates

In Figure 5, the amount of confidence subjects had in their probability estimates is shown at the different objective probability levels.



The multivariate regression results in Table 2 also show that the subjective beliefs ($PROBABILITY_ESTIMATE_t$) are highly positively correlated with the objective, Bayesian posterior ($OBJECTIVE_PROBABILITY_t$). If the latter

increases by 10%, the subjective posterior increases by 6.1%. The difference between these 2 probabilities can come from subjects using an updating formula different 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 effect, 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 several past dividends.⁴ For any given belief elicited from the subject in the prior trial ($\text{PROBABILITY_ESTIMATE}_{t-1}$), 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 $\text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t$. As can be seen from Table 2, 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 that they have high confidence in their probability estimates in trial $t - 1$, as would naturally be expected.⁵ However, since there is not a one-to-one mapping between the subjective and the

TABLE 2
Elicited Posterior Probability Estimates

In Table 2 the dependent variable, $\text{PROBABILITY_ESTIMATE}_t$, is the elicited, subjective posterior probability that the stock is the good stock. POSITIVE_CUE_t and NEGATIVE_CUE_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 cue, respectively. STOCK_CHOICE_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. $\text{OBJECTIVE_PROBABILITY}_t$ 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 . $\text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t$ 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. $\text{CONFIDENCE_ESTIMATE}_t$ 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 stock. Standard errors are clustered by subject (t-statistics are shown in parentheses). *** indicates significance at the 1% level.

Dependent Variable $\text{PROBABILITY_ESTIMATE}_t$	Panel A	Panel B	Panel C
POSITIVE_CUE_t	0.01 (0.55)	0.01 (0.67)	0.01 (0.40)
NEGATIVE_CUE_t	-0.01 (-0.81)	-0.00 (-0.26)	-0.00 (-0.35)
STOCK_CHOICE_t	0.07 (3.69)***	0.10 (5.69)***	0.09 (5.53)***
$\text{OBJECTIVE_PROBABILITY}_t$	0.61 (11.59)***		
$\text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t$		0.62 (12.48)***	0.41 (4.87)***
$\text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t$ $\times \text{CONFIDENCE_ESTIMATE}_{t-1}$			0.03 (3.09)***
$\text{CONFIDENCE_ESTIMATE}_{t-1}$			-0.02 (-3.67)***
Subject fixed effects	Yes	Yes	Yes
Adj. R^2	0.59	0.55	0.56
No. of obs.	2,413	2,413	2,413

⁴If the prior probability that the stock is good is p , then following a high (\$10) dividend, the posterior obtained using Bayes' rule is $(3p)/(2p + 1)$. The posterior after a low (-\$10) dividend is $p/(3 - 2p)$.

⁵Estimating a truncated regression model to account for the fact that the dependent variable in Table 2 is constrained to the interval [0,1] yields results of similar magnitude and statistical significance.

correct posteriors, Bayes' rule is a only rough predictor of the way subjects update beliefs, leaving room for other factors to influence 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 2 shows, relative to subjects who chose the bond in trial t , those individuals who chose the stock (i.e., those for whom $\text{STOCK_CHOICE}_t = 1$) declare a probability estimate that the stock is the good stock 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 1 indicates, subjects select assets based on their priors. The results in Table 2 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. All else being 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. This is not an instance of the well-known endowment effect (Thaler (1980)), whereby individuals who have ownership of certain assets value them more than the other individuals do. In our task subjects were not endowed with the stock or bond. They had to decide at every trial which of the 2 assets they wanted to hold, for that trial only. When beliefs are elicited in our experiment, subjects do not have any asset in their portfolio and are getting ready for the next round of portfolio selection.⁶

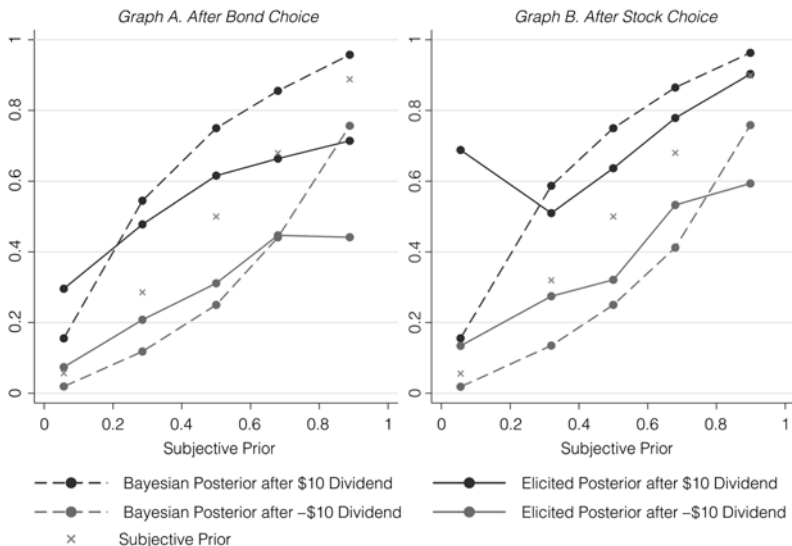
To further investigate the role of past choices on beliefs, we split the data based on whether the stock or the bond was chosen in a given trial, and we analyzed how subjects in these 2 conditions updated their beliefs upon the arrival of the dividend news. For any prior belief expressed by the subject, we calculated the correct Bayesian posterior and compared that to the posterior actually declared by the subject. Several patterns suggestive of biased learning emerge, as seen in Figure 6. Confirming the results in Table 2, stock holders provide higher posterior estimates than do bond holders, independent of their prior beliefs and the actual dividend paid by the stock. Also, 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 opposes 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

⁶Recent papers that relate to the endowment effect in the context of investment decisions are Hales (2007) and Ko and Hansch (2008). Importantly, belief updating (learning) is not considered by either of these papers.

FIGURE 6

Dependence of the Learning Process on Prior Actions and Beliefs

In Figure 6, the elicited subjective posterior beliefs and the objective posteriors obtained with Bayesian updating from subjective priors are shown as a function of the subjects' choice and the dividend revealed during the trial.



the summary statistics in Table 3. If the subject has an extremely 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 3). All of these effects are statistically significant ($p < 0.001$).

TABLE 3
Errors in Updating Extreme Priors

Table 3 presents the error in probability updating (BAYESIAN.POSTERIOR.USING.SUBJECTIVE.PRIOR_t - PROBABILITY.ESTIMATE_t) as a function of the stock dividend, for extreme low (i.e., PROBABILITY.ESTIMATE_{t-1} ∈ [0, 0.2]), extreme high (i.e., PROBABILITY.ESTIMATE_{t-1} ∈ [0.8, 1]) and moderate (i.e., PROBABILITY.ESTIMATE_{t-1} ∈ (0.2, 0.8)) subjective prior beliefs. *** indicates that the estimate is significantly different than 0 ($p < 0.001$).

	STOCK_DIVIDEND _t = -\$10	STOCK_DIVIDEND _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***

For moderate priors, however, 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 4 presents the summary statistics for these errors in updating. For stock holders and bond holders, the

TABLE 4
Conservatism Bias

Table 4 presents the error in probability updating ($\text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t - \text{PROBABILITY_ESTIMATE}_t$) as a function of the subject's choice and the stock dividend, for moderate subjective prior beliefs (i.e., $\text{PROBABILITY_ESTIMATE}_{t-1} \in (0.2, 0.8)$). STOCK_CHOICE_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 0 ($p < 0.001$).

	$\text{STOCK_DIVIDEND}_t = -\10	$\text{STOCK_DIVIDEND}_t = \10
$\text{STOCK_CHOICE}_t = 0$	-0.07***	0.13***
$\text{STOCK_CHOICE}_t = 1$	-0.10***	0.10***

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 adequately revise down their priors. Their elicited posterior probabilities are higher than they should be by 10% and 7%, respectively. These effects are also statistically significant ($p < 0.001$).

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 4 as a function of asset choice and dividend type. We formally test it by estimating an ordinary least squares (OLS) regression where the dependent variable is the absolute value of the learning error ($|\text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t - \text{PROBABILITY_ESTIMATE}_t|$), first for trials where the stock paid a low dividend, and then for those where the dividend was high. The independent variable of interest is the dummy variable indicating whether in the beginning of the trial the subject chose the stock or not (STOCK_CHOICE_t). We include subject fixed effects to take into account the possibility that some people make larger size errors in updating than others, and we cluster the standard errors by subject. The results are presented in Table 5. The size of the probability updating error is 3% higher for stock holders than for bond holders if the dividend is low, and it

TABLE 5
Strategic Belief Formation

The dependent variable in Table 5 is the absolute value of the error in probability updating ($|\text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t - \text{PROBABILITY_ESTIMATE}_t|$), for moderate subjective prior beliefs (i.e., $\text{PROBABILITY_ESTIMATE}_{t-1} \in (0.2, 0.8)$). STOCK_CHOICE_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). ** indicates significance at the 5% level.

Dependent Variable $ \text{BAYESIAN_POSTERIOR_USING_SUBJECTIVE_PRIOR}_t - \text{PROBABILITY_ESTIMATE}_t $	$\text{STOCK_DIVIDEND}_t = -\10 Coeff./ t	$\text{STOCK_DIVIDEND}_t = \10 Coeff./ t
STOCK_CHOICE_t	0.03 (2.54)**	-0.03 (-2.61)**
Subject fixed effects	Yes	Yes
Adj. R^2	0.22	0.16
No. of obs.	927	857

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. This effect is different from confirmation bias, 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 (Mynatt, Doherty, and Tweney (1977)). 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 6. 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 6 we include as controls the

TABLE 6
Subjective Confidence Estimates

In Table 6 the dependent variable, CONFIDENCE_ESTIMATE_{*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 stock. POSITIVE_CUE_{*t*} and NEGATIVE_CUE_{*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 cue, respectively. HIGH_STOCK_DIVIDEND_{*t*} is an indicator variable equal to 1 if the dividend paid by the stock in trial *t* is \$10, and 0 if the dividend is -\$10. STOCK_CHOICE_{*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. OBJECTIVE_PROBABILITY_{*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*. PROBABILITY_ESTIMATE_{*t*} is the subject's estimate of the probability that the stock is the good stock. Standard errors are clustered by subject (t-statistics are shown in parentheses). *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Dependent Variable CONFIDENCE_ESTIMATE _{<i>t</i>}	Panel A	Panel B
POSITIVE_CUE _{<i>t</i>}	0.17 (1.88)*	0.18 (1.92)*
NEGATIVE_CUE _{<i>t</i>}	0.13 (1.67)	0.12 (1.50)
HIGH_STOCK_DIVIDEND _{<i>t</i>}	-0.25 (-1.81)*	-0.44 (-2.65)**
HIGH_STOCK_DIVIDEND _{<i>t</i>} × STOCK_CHOICE _{<i>t</i>}	0.73 (3.42)***	0.74 (4.23)***
STOCK_CHOICE _{<i>t</i>}	-0.37 (-2.52)**	-0.51 (-3.08)***
OBJECTIVE_PROBABILITY _{<i>t</i>}	-3.56 (-2.55)**	
OBJECTIVE_PROBABILITY _{<i>t</i>} ²	3.64 (2.57)**	
PROBABILITY_ESTIMATE _{<i>t</i>}		-3.98 (-3.10)***
PROBABILITY_ESTIMATE _{<i>t</i>} ²		4.83 (3.58)***
Subject fixed effects	Yes	Yes
Adj. R ²	0.54	0.55
No. of obs.	2,402	2,402

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 $HIGH_STOCK_DIVIDEND_t$ and $STOCK_CHOICE_t$. 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 influences decision making under risk. Affect may matter whether it is induced exogenously by an experimenter, a policymaker, 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, people ignore new information that is opposed to their actions, and as a result, learning is flawed. These findings have implications for several areas of finance and economics.

A. Learning and Incentives for Exploration

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. This in turn implies that to induce agents to efficiently explore new strategies, incentives must be provided to mitigate the detrimental effect on learning caused by emotional reactions to new information. For example, agents may ignore bad news about a new project that they have chosen to start, in order to avoid experiencing negative affect, or may be overtly confident in their project selection abilities when new information matches their prior choices. Moreover, agents could simply decide not to explore at all, to avoid the possibility of being disappointed. One possible mechanism that can alleviate these inefficiencies is to offer pay that is not contingent on the outcomes of the agents' efforts in the early and most critical stages of exploration. The experimental evidence in Ederer and Manso (2009) is consistent with this hypothesis. They find that the combination of tolerance for early failure and reward for long-term success is an

effective mechanism to induce agents to discover novel business strategies. Field evidence also suggests that decoupling agents' payoffs from outcomes may improve learning and exploration. For instance, Acharya and Subramanian (2009) and Acharya, Baghai-Wadji, and Subramanian (2009) find that innovation is encouraged by debtor-friendly bankruptcy laws and by stringent labor laws that restrict the dismissal of employees.

B. Market and Institutional Design

Policymakers 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 excitement 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.

C. 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 that 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.

Existing empirical evidence on investor behavior is consistent with these conjectures. For example, Vissing-Jorgensen (2003) shows that young investors without prior exposure to market crashes had the highest stock market return expectations during the boom years in the late 1990s, while Kaustia and Knüpfer (2008) find that individual investors who experienced high returns by participating in past initial public offerings (IPOs) are more likely to subscribe to future IPOs. Malmendier and Nagel (2011) document that individuals who have experienced low stock market returns throughout their lives report lower willingness to take financial risk, are less likely to participate in the stock market, and, conditional

on participating, invest a lower fraction of their liquid assets in stocks. Therefore, by inducing emotional reactions that influence risk preferences and learning, past outcomes experienced by investors can change future financial choices and asset prices.

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 neuroeconomics 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 affect 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 a positive emotional state and avoiding a negative state. Specifically, subjects do not fully incorporate news that contradicts their prior choices, and they 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, which 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 a factor to be considered in optimal contracting problems and may also help predict fluctuations in investor participation and asset prices.

Appendix. 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 that 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 they 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 they 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 2 things: i) What you think the probability is that the stock is the good stock (the answer must be a number between 0 and 100; do not add the % sign, just type in the value), ii) 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 sometimes 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 + $(1/20) \times \text{TASK_EARNINGS}$, where the $\text{TASK_EARNINGS} = (\text{Dividends you accumulate through investing in the 2 assets PLUS money you earn by guessing correct probabilities})$. 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

- Acharya, V. V.; R. Baghai-Wadji; and K. Subramanian. "Labor Laws and Innovation." Working Paper, New York University (2009).
- Acharya, V. V., and K. V. Subramanian. "Bankruptcy Codes and Innovation." *Review of Financial Studies*, 22 (2009), 4949–4988.
- Akerlof, G. A., and W. T. Dickens. "The Economic Consequences of Cognitive Dissonance." *American Economic Review*, 72 (1982), 307–319.
- Barber, B. M., and T. Odean. "Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors." *Journal of Finance*, 55 (2000), 773–806.
- Benartzi, S., and R. H. Thaler. "Myopic Loss Aversion and the Equity Premium Puzzle." *Quarterly Journal of Economics*, 110 (1995), 73–92.
- Ben-David, I.; J. R. Graham; and C. R. Harvey. "Managerial Overconfidence and Corporate Policies." Working Paper, Duke University (2007).
- Bernheim, D. B., and A. Rangel. "Addiction and Cue-Triggered Decision Processes." *American Economic Review*, 94 (2004), 1558–1590.
- Bjork, J. M.; B. Knutson; G. W. Fong; D. M. Caggiano; S. M. Bennett; and D. W. Hommer. "Incentive-Elicited Brain Activation in Adolescents: Similarities and Differences from Young Adults." *Journal of Neuroscience*, 24 (2004), 1793–1802.
- Breiter, H. C.; I. Aharon; D. Kahneman; A. Dale; and P. Shizgal. "Functional Imaging of Neural Responses to Expectancy and Experience of Monetary Gains and Losses." *Neuron*, 30 (2001), 619–639.
- Caplin, A., and M. Dean. "Dopamine, Reward Prediction Error, and Economics." *Quarterly Journal of Economics*, 123 (2008), 663–701.
- Caplin, A., and J. Leahy. "Psychological Expected Utility Theory and Anticipatory Feelings." *Quarterly Journal of Economics*, 116 (2001), 55–79.
- Charness, G., and D. Levin. "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect." *American Economic Review*, 95 (2005), 1300–1309.
- Chua, P.; M. Krams; I. Toni; R. Passingham; and R. Dolan. "A Functional Anatomy of Anticipatory Anxiety." *NeuroImage*, 9 (1999), 563–571.
- Daw, N. D.; S. Kakade; and P. Dayan. "Opponent Interactions between Serotonin and Dopamine." *Neural Networks*, 15 (2002), 603–616.
- Delgado, M. R.; L. E. Nystrom; C. Fissell; D. C. Noll; and J. A. Fiez. "Tracking the Hemodynamic Responses to Reward and Punishment in the Striatum." *Journal of Neurophysiology*, 84 (2000), 3072–3077.
- Ederer, F. P., and G. Manso. "Is Pay-for-Performance Detrimental to Innovation?" Working Paper, University of California at Los Angeles (2009).
- Elster, J. "Emotions and Economic Theory." *Journal of Economic Literature*, 36 (1998), 47–74.
- Gervais, S.; J. B. Heaton; and T. Odean. "Overconfidence, Investment Policy and Manager Welfare." Working Paper, Duke University (2005).
- Gervais, S., and T. Odean. "Learning to Be Overconfident." *Review of Financial Studies*, 14 (2001), 1–27.
- Gilad, D., and D. Kliger. "Priming the Risk Attitudes of Professionals in Financial Decision Making." *Review of Finance*, 12 (2008), 567–586.
- Gneezy, U., and J. Potters. "An Experiment on Risk Taking and Evaluation Periods." *Quarterly Journal of Economics*, 112 (1997), 631–645.
- Grinblatt, M., and M. Keloharju. "Sensation Seeking, Overconfidence, and Trading Activity." *Journal of Finance*, 64 (2009), 549–578.
- Hales, J. "Directional Preferences, Information Processing, and Investors' Forecast of Earnings." *Journal of Accounting Research*, 45 (2007), 607–628.
- Heaton, J. B. "Managerial Optimism and Corporate Finance." *Financial Management*, 31 (2002), 33–45.
- Hirshleifer, D., and T. Shumway. "Good Day Sunshine: Stock Returns and the Weather." *Journal of Finance*, 58 (2003), 1009–1032.
- Kaustia, M., and S. Knüpfer. "Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions." *Journal of Finance*, 63 (2008), 2679–2702.

- Kluger, A. N., and A. DeNisi. "The Effects of Feedback Interventions on Performance: A Historical Review, a Meta-Analysis, and a Preliminary Feedback Intervention Theory." *Psychological Bulletin*, 119 (1996), 254–284.
- Knutson, B.; C. M. Adams; G. W. Fong; and D. Hommer. "Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens." *Journal of Neuroscience*, 21 (2001), 1–5.
- Knutson, B.; G. E. Wimmer; C. M. Kuhnen; and P. Winkielman. "Nucleus Accumbens Activation Mediates the Influence of Reward Cues on Financial Risk Taking." *NeuroReport*, 19 (2008), 509–513.
- Ko, K. J., and O. Hansch. "Persistence of Beliefs in an Investment Experiment." Working Paper, Securities and Exchange Commission (2008).
- Kuhnen, C. M., and B. Knutson. "The Neural Basis of Financial Risk Taking." *Neuron*, 47 (2005), 763–770.
- Lo, A. W., and D. V. Repin. "The Psychophysiology of Real-Time Financial Risk Processing." *Journal of Cognitive Neuroscience*, 14 (2002), 323–339.
- Loewenstein, G. F.; E. U. Weber; C. K. Hsee; and N. Welch. "Risk as Feelings." *Psychological Bulletin*, 127 (2001), 267–286.
- Malmendier, U., and S. Nagel. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *Quarterly Journal of Economics*, 126 (2011), 373–416.
- Malmendier, U., and G. Tate. "CEO Overconfidence and Corporate Investment." *Journal of Finance*, 60 (2005), 2661–2700.
- Martinez, D.; M. Slifstein; A. Broft; O. Malawi; D. R. Hwang; Y. Huang; T. Cooper; L. Kegeles; E. Zarahn; A. Abi-Dargham; S. N. Haber; and M. Laruelle. "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*, 23 (2003), 285–300.
- Mynatt, C. R.; M. E. Doherty; and R. D. Tweney. "Confirmation Bias in a Simulated Research Environment: An Experimental Study of Scientific Inference." *Quarterly Journal of Experimental Psychology*, 29 (1977), 85–95.
- Paulus, M. P.; C. Rogalsky; A. Simmons; J. S. Feinstein; and M. B. Stein. "Increased Activation in the Right Insula during Risk-Taking Decision Making Is Related to Harm Avoidance and Neuroticism." *NeuroImage*, 19 (2003), 1439–1448.
- Pessiglione, M.; B. Seymour; G. Flandin; R. J. Dolan; and C. D. Frith. "Dopamine-Dependent Prediction Errors Underpin Reward-Seeking Behavior in Humans." *Nature*, 442 (2006), 1042–1045.
- Phillips, L. D., and W. Edwards. "Conservatism in a Simple Probability Inference Task." *Journal of Experimental Psychology*, 72 (1966), 346–354.
- Preuschoff, K.; P. Bossaerts; and S. R. Quartz. "Neural Differentiation of Expected Reward and Risk in Human Subcortical Structures." *Neuron*, 51 (2006), 381–390.
- Saunders, E. M., Jr. "Stock Prices and Wall Street Weather." *American Economic Review*, 83 (1993), 1337–1345.
- Schultz, W. "Behavioral Dopamine Signals." *TRENDS in Neuroscience*, 30 (2007), 203–210.
- Seymour, B.; J. P. O'Doherty; P. Dayan; M. Koltzenburg; A. K. Jones; R. J. Dolan; K. J. Friston; and R. S. Frackowiak. "Temporal Difference Models Describe Higher-Order Learning in Humans." *Nature*, 429 (2004), 664–667.
- Simmons, A.; S. C. Matthews; M. B. Stein; and M. P. Paulus. "Anticipation of Emotionally Aversive Visual Stimuli Activates Right Insula." *NeuroReport*, 15 (2004), 2261–2265.
- Thaler, R. "Toward a Positive Theory of Consumer Choice." *Journal of Economic Behavior and Organization*, 1 (1980), 39–60.
- Vissing-Jorgensen, A. "Perspectives on Behavioral Finance: Does 'Irrationality' Disappear with Wealth? Evidence from Expectations and Actions." *NBER Macroeconomics Annual*, 18 (2003), 139–194.