usc Marshall

USC Marshall School of Business

**Marshall Research Paper Series
Working Paper FBE 02.16**

**Neural Evidence of Regret and its Implications for Investor
Behavior**

Cary Frydman
University of Southern California, Marshall School of Business

Colin Camerer
Caltech, Division of Humanities and Social Sciences, and
Computational and Neural Systems

Neural Evidence of Regret and its Implications for Investor Behavior

Cary Frydman and Colin Camerer¹

ABSTRACT: We use neural data collected from an experimental asset market to measure regret preferences while subjects trade stocks. When subjects observe a positive return for a stock they chose not to purchase, a regret signal is observed in an area of the brain that is commonly active during reward processing. Subjects are unwilling to repurchase stocks that have recently increased in price, even though this is suboptimal in our experiment. The strength of stock repurchasing mistakes is correlated with the neural measures of regret. Subjects with high rates of repurchasing mistakes also exhibit large disposition effects.

JEL Codes: G02, G11

Corresponding Author:

Cary Frydman
3670 Trousdale Parkway
Los Angeles, CA 90089
(213)-821-5586
cfrydman@marshall.usc.edu

¹ Frydman is at the Marshall School of Business at University of Southern California, cfrydman@marshall.usc.edu. Camerer is at the Division of Humanities and Social Sciences, and Computational and Neural Systems at Caltech, camerer@hss.caltech.edu. We are grateful for comments from Nicholas Barberis, Phillip Bond, Peter Bossaerts, Campbell Harvey, Milica Mormann, Antonio Rangel and from seminar participants at Caltech, the University of Southern California, the University of Washington, the University of Western Ontario, the 2013 Miami Behavioral Finance Conference, and the 2013 Boulder Consumer Financial Decision-Making Conference.

Most models of investor behavior assume that investors maximize a utility function defined over a stream of consumption. In these models, consumption is typically the *only* carrier of utility. Such a simplified specification of preferences is parsimonious. However, some facts about investor behavior remain difficult to explain when consumption is the only carrier of utility (Barberis and Thaler (2003), Barber and Odean (2013)). One general approach to explain the anomalous facts is to assume there are additional sources of utility. However, to maintain a disciplined and accurate approach to modeling investor behavior, any additional assumptions about the source of utility must be tested empirically.

In this paper, the additional source of utility that we hypothesize and measure is regret, the difference between actual outcomes and the foregone (or counterfactual) outcomes that would have resulted from different decisions. Regret theory appears prominently in the early work on behavioral finance (Shefrin and Statman (1984, 1985)) and has continued to receive attention (Barberis, Huang, and Santos (2001), Strahilevitz, Barber, and Odean (2011)). Although the idea that investors experience regret is plausible, evidence of investor regret has not been documented, perhaps because of the difficulty in obtaining ideal data measuring both regret and trading decisions.

Our main contribution is to provide a novel measure of regret that can be used to test whether investors exhibit trading behavior that is consistent with the experienced regret hypothesis. The novel measure of regret is neural data collected using functional magnetic resonance imaging (fMRI). When combined with our experimental design, the neural data is helpful because it can provide evidence that regret is generated at the precise moment when *ex-post* suboptimal trading outcomes are revealed, but before trading decisions are made. We can then test whether the regret generated upon news of a stock return is correlated with subsequent trading behavior. Testing this hypothesis using trading data alone would typically be very difficult. Our data therefore allow us to provide, for the first time, evidence that an apparent

regret signal is encoded at the time of a stock's price change, and that this signal can explain cross-sectional variation in subsequent trading behavior.

Shefrin and Statman (1984) provide one of the earliest discussions of regret in finance by proposing that firms pay dividends because investors are regret-averse. Another early application of regret is proposed by the same authors who argue that the disposition effect – the fact that investors are more likely to sell winning stocks compared to losing stocks – can be explained by regret aversion (Shefrin and Statman (1985)). Here, the idea is that if investors have a stock with a paper loss, then selling the stock will trigger regret and thus investors will be reluctant to realize losses compared to gains².

Other theory (Barberis, Huang, Thaler (2006)) shows that if investors engage in narrow framing – that is, they evaluate risk in isolation rather than first merging it with other pre-existing risks – then this can make the stock market look excessively risky, which will lead to low stock market participation. One interpretation of narrow framing that these authors offer is based on regret: although an investor should evaluate any new gamble by first combining it with all other sources of background risk, accepting the new gamble is linked to a specific decision. This specific decision can generate regret if the gamble turns out badly³.

More recently, using data from a large discount brokerage, Strahilevitz, Odean, and Barber (2011) show that when an investor sells a stock, he is less likely to repurchase this same stock if the price has increased since the sale, compared to when the price has decreased since the sale. The authors call this new empirical fact the “repurchase effect.” Similar to the logic in

² Recent theoretical work by Barberis and Xiong (2012) and Ingersoll and Jin (2013) presents a related idea where investors derive “realization utility” from the act of selling risky assets. One interpretation of this preference structure is that the disutility from realizing a loss is driven by regret. Later in the paper, we provide some neural evidence suggesting there is a link between regret and realization utility.

³ Regret has also been used as a microfoundation for behavioral models of the aggregate stock market. Barberis, Huang, and Santos (2001) propose a model in which investors derive utility from both consumption and fluctuations in financial wealth. This second utility component is the key to delivering predictions that can match statistical properties of the data on stock returns. The authors interpret this key source of utility as deriving, in part, from feelings of regret from investing in a falling stock market.

Shefrin and Statman (1984), Strahilevitz et al. (2011) argue that if an investor observes a stock price increase after selling the stock, this induces regret because the investor would have been better off not selling in the first place. In order to avoid being reminded over this poor premature selling decision, the investor chooses to mentally distance himself from the stock by not repurchasing it. Weber and Welfens (2011) were the first to show a repurchase effect in a laboratory setting.

Our experimental design is optimized to closely examine this potential influence of regret on the repurchasing behavior of individual investors. Strahilevitz et al. (2011) argue that alternative explanations besides regret have difficulty explaining the full extent of the repurchase effect. First, the behavior is not tax-motivated because the repurchase effect is found in both taxable and tax-deferred accounts. Second, the behavior is unlikely driven by private information because investors do not consistently earn higher returns when exhibiting a repurchase effect. Finally, a general belief in mean-reverting prices cannot explain the data as investors who exhibit a repurchase effect also purchase winning stocks that they have previously not owned. On the other hand, regret theory is a plausible candidate to explain the repurchase effect, but it has not yet been tested. In section I.C. we provide a basic model of how regret can generate a repurchase effect, which is based on the “regret-devaluation” mechanism from social psychology (Arkes et al. (2002)).

The details of our experimental design are discussed at length later in the paper, but we briefly introduce the two key design aspects here. Subjects in our experiment are endowed with cash and the opportunity to trade in three separate stocks. The stock price changes for each stock are mutually independent and positively autocorrelated, and subjects are given this information at the beginning of the experiment. This autocorrelation is the first key aspect of our design, and it allows us to perform a strong test of the repurchase effect. To see why, consider a subject who sells a stock and subsequently observes an increase in the stock price. Because of the positive

autocorrelation, this stock is likely to continue to rise in price, and hence the subject should purchase it to capitalize on the future price increase. Conversely, consider a subject who sells a stock and observes that the stock price subsequently decreases in price. The positive autocorrelation in this setting implies that the stock should continue to do poorly, and the subject should avoid repurchasing it. Taken together, subjects who trade to maximize their final portfolio value should repurchase stocks that have recently gone up in price and should avoid purchasing stocks that have recently decreased in price. In other words, the optimal trading strategy in this setting predicts that subjects will exhibit *the opposite* of a repurchase effect.

The second key aspect of our experimental design involves temporally separating the event where a subject receives information about a stock price change from the event where a subject is asked to make a trading decision. This temporal separation is critical because it allows us to identify the neural signal that is generated in response to news about a stock price change. Many regret applications assume that investors experience regret at the moment when they receive news about returns; without the temporal separation built into the experimental design, it would be difficult to test this assumption.

The main results from our experiment can be summarized as follows. First, 92% of our subjects exhibit a repurchase effect that is significantly greater than what would be exhibited by a Bayesian agent following the optimal trading strategy. Second, at the moment when a subject observes a price increase for a recently sold stock, we observe brain activity that is consistent with a regret signal. Moreover, those subjects who exhibit stronger neural regret signals are the same subjects who exhibit a greater repurchase effect. In other words, the neural measures of regret we collect using fMRI can explain a significant component of the cross-sectional variation in trading behavior. Third, we test for a systematic relationship between the repurchase effect and the disposition effect, which is also a suboptimal behavior in our experiment. We find that these

two trading behaviors are highly correlated across subjects, suggesting that there may be a common psychological mechanism that generates buying and selling mistakes.

Although prices in our experiment are not mean-reverting, it is possible that subjects may incorrectly believe that prices mean-revert. This incorrect belief specification will readily predict both a repurchase effect and a disposition effect; it will also predict that these two behaviors are highly correlated. To help rule out this belief-based explanation, we run an additional experiment (outside the fMRI scanner) in which the expected price change that would be computed by a Bayesian trader is explicitly displayed. While this extra information pushes subjects closer to optimal behavior, we still observe substantial disposition effects and repurchase effects, indicating that incorrect beliefs cannot fully explain our data.

Our study also contributes to the literature in cognitive neuroscience on regret. Early work by Camille et al. (2004) shows that when subjects with damage to a brain area called the orbitofrontal cortex (OFC) choose between two lotteries, they exhibit a blunted emotional reaction to counterfactual payoffs. This suggests that the OFC plays a causal role in generating regret – though it does not imply that regret is *exclusively* encoded in the OFC (since the lesions only damaged OFC). In a follow-up study by the same group, Coricelli et al. (2005) employ the same experimental paradigm and use whole-brain imaging to explore all possible neural activity encoding regret. The whole-brain analysis finds activity in the OFC, but also in another area of the brain called the striatum.

Lohrenz et al. (2007) were the first to show an fMRI response in the ventral striatum (vSt) to “fictive learning” – that is a response to the reward of an action that was not chosen, though they specifically avoid tying their result to regret. In their study, subjects were asked to allocate a fraction of their wealth to a risky asset over a series of trials. While their study also uses an investment task to investigate neural correlates of counterfactual payoffs, our study differs on

at least one crucial dimension: we provide subjects with explicit information about the risky asset’s return generating process⁴. Li and Daw (2011) find that the vSt responds to foregone rewards when choosing between one of two slot machines. Moreover, the strength of this neural signal explains variation in the weight that subjects attach to foregone rewards, as inferred through their choice of slot machines. Finally, Brassen et al. (2012) investigate whether there are differences in the neural response to regret across the lifespan. They find a neural regret signal in the vSt for young subjects and depressed older subjects, but this neural regret signal is absent among healthy older subjects.

Our contribution to this literature is to provide a detailed characterization of the regret mechanism that is deployed during stock trading. We document a regret signal that is generated upon news of a stock price change and we then show that this neural signal is encoded at the time of a trading decision. Perhaps most novel of all, we find that the strength of the repurchase effect is correlated across subjects with the size of the neural regret signal. This is important because it demonstrates, for the first time, that a neural signal generated in response to price changes is associated with subsequent trading behavior. The neural regret signals also provide key motivation to generate new testable hypotheses about a systematic relationship between buy-side and sell-side trading biases.

I. Experimental Design and Predictions

A. Design

The design and data used in our analysis come from Frydman et al. (2014). The focus in that paper was exclusively on *selling* behavior and testing the realization utility explanation of the

⁴ As we describe in the main text, this allows us derive a sharp prediction about optimal behavior, which we can then compare to the trading predictions based on regret theory.

disposition effect. Here, the focus is on *buying* behavior. We test the hypothesis that neural measures interpretable as regret are associated with a repurchase effect.

Subjects have the opportunity to trade three stocks — denoted stock A, stock B, and stock C — in an experimental asset market. The experiment has two identical sessions separated by a one-minute break. Each session lasts approximately 16 minutes and consists of 108 trials. The value t indexes the trials within a session.⁵

At the beginning of each session, each subject is given \$350 in experimental currency and is required to buy one share of each stock. The initial share price for each stock is \$100; after the initial purchase, each subject is therefore left with \$50. Every trial $t > 9$ has two parts: a price update, and a trading decision. The price update and subsequent trading decision are represented by two different screen displays that the subject sees (Figure 1).

In the price update part, one of the three stocks is chosen at random and the subject is shown a price change for this stock. Note that stock prices only evolve during the price update screens; as a result, subjects see the entire price path for each stock. In the trading part, one of the three stocks is again chosen at random and the subject is asked whether he wants to trade the stock. Note that no new information is revealed during this part.

We split each trial into two parts so as to temporally separate different computations associated with decision-making. At the price update screen, subjects are provided with information about a change in the price of one of the three stocks, but do not have to compute the value of buying or selling the stock, both because they are not allowed to make decisions at this stage, and also because they do not know which of the three assets will be selected for trading in the next screen. At the trading screen the opposite situation holds: subjects need to compute the

⁵ We split our experiment into two sessions in order to avoid running the fMRI machine for too long without a break. This can lead to potential medical risks for the subjects. It also provides a simple way to test for persistence of individual differences in trading patterns by correlating behavior across the two sessions (see Table 1).

value of buying or selling a stock, but do not need to update their beliefs since no new information about prices is provided.

Trials 1 through 9 consist only of a price update stage; i.e., subjects are not given the opportunity to buy or sell during these trials. We designed the experiment in this way so that subjects can accumulate some information about the three stocks before having to make any trading decisions.

Each subject is allowed to hold a maximum of one share and a minimum of zero shares of each stock at any point in time⁶. In particular, short-selling is not allowed. The trading decision is therefore reduced to deciding whether to sell a stock (conditional on holding it), or deciding whether to repurchase it (conditional on not holding it). The price at which a subject can buy or sell a stock is given by the current market price of the stock.

The price path of each stock is governed by a hidden-state Markov process with a good state and a bad state. The states and transitions of the three stocks are independent from each other. If stock i is in the good state, its price in a trial increases with probability 0.55 and decreases with probability 0.45. Conversely, if it is in the bad state, its price increases with probability 0.45 and decreases with probability 0.55. The magnitude of the price change is drawn uniformly from $\{\$5, \$10, \$15\}$, independently of the direction of the price change.

The state of each stock changes over time in the following way. Before trial 1, we randomly assign one of the two equally-likely states to each stock. If the price update in trial $t > 1$ is *not* about stock i , then the state of stock i in trial t remains the same as its state in the previous trial, $t-1$. If the price update in trial $t > 1$ *is* about stock i , then the state of stock i in trial t stays the same as in trial $t-1$ with probability 0.8, and switches with probability 0.2. In mathematical terms,

⁶ Combined with the design feature that subjects are required to buy one share of each stock at the beginning of the experiment, this implies that all subsequent purchase decisions are in fact “repurchase” decisions.

if $s_{i,t} \in \{good, bad\}$ is the state of stock i in trial t , then $s_{i,t} = s_{i,t-1}$ if the time t price update is not about stock i , whereas if the time t price update is about stock i , the state switches as follows:

	$s_{i,t+1}=good$	$s_{i,t+1}=bad$
$s_{i,t}=good$	0.8	0.2
$s_{i,t}=bad$	0.2	0.8

The states of the stocks are always hidden to the subjects; but of course, subjects can make Bayesian inferences about the state from observed price changes.

To reduce error in comparison of trading performance across subjects, the same set of realized prices is used for all subjects (as in Weber and Camerer (1998)). There are certainly potential disadvantages from using a single price path for all subjects (e.g., the inability to control for order effects through random variation across subjects). However, we thought that the benefits of easily comparing trading performance across subjects outweighed the costs of potential order effects. We also show below that statistical properties of the realized price sequence are very close to the expected properties of the sequence, conditional on the parameters known to subjects.

Note that in expectation, the stocks exhibit positive short-term autocorrelation in their price changes. If a stock increased on the last price update, it was probably in a good state for that price update. Since it is highly likely (probability 0.8) to remain in the same state for its next price update, its next price change is likely to also be positive⁷.

At the end of each of the two sessions, subjects' holdings of the three stocks are liquidated and added to their cash value to determine a final cash balance. To motivate subjects to

⁷ One potential issue with using a single realized price path for all subjects is that the degree of positive autocorrelation (which generates the short-term price momentum) may be substantially different from the expected degree of momentum due to simulation error. To test this, we computed a momentum statistic as the degree to which the sign of the price change in period $t-1$ predicts the sign of the price change in period t . Using only those parameters given to subjects before the experiment, we ran simulations to compute this probability: $Pr(\text{sign}(P_t - P_{t-1}) = \text{sign}(P_{t-1} - P_{t-2})) = 0.5029$. This represents the *expected* degree of momentum. For the single realized price path that was used, the analogous empirical probability is computed to be 0.5047. In other words, the ex-post momentum is only very slightly larger than the expected degree of momentum.

make careful decisions, their US dollar earnings are equal to \$15 (for participation) plus the total of the two sessions' final cash balance divided by 24. Their earnings ranged from \$43.05 to \$57.33 with a mean of \$52.57 and a standard deviation of \$3.35.

In order to avoid liquidity constraints, we allow subjects to carry a negative cash balance in order to purchase a stock if they do not have sufficient cash to do so at the time of a decision. If a subject ends the experiment with a negative cash balance, this amount is subtracted from the terminal value of his portfolio. The large cash endowment, together with the constraint that subjects can hold at most one unit of each stock at any moment, was sufficient to guarantee that no one ended the experiment with a negative portfolio value, or was unable to buy a stock because of a shortage of cash during the experiment. Moreover, while subjects are allowed to carry a negative cash balance throughout the experiment, the stock price exceeded a subject's cash balance on less than 2% of trials where there was an opportunity to repurchase.

N=28 Caltech subjects participated in the experiment (22 male, age range 18 – 60).⁸ All subjects were right-handed and had no history of psychiatric illness, and none were taking medications that interfere with fMRI. The exact instructions given to subjects at the beginning of the experiment are included in the Appendix. The instructions carefully describe the stochastic structure of the price process, as well as all other details of the experiment. Before entering the scanner, the subjects underwent a practice session of 25 trials to ensure familiarity with the market software.

⁸ One additional subject participated in the experiment but was excluded from further analyses because his head motion during the scanning exceeded a pre-specified threshold, making neural measurement of BOLD signal inaccurate.

B. Optimal trading strategy

We now characterize the optimal trading strategy for an “expected value investor,” defined as a risk-neutral Bayesian investor who is maximizing the expected value of take-home earnings. The optimal strategy of such an investor is to sell (or not buy) a stock when he believes that it is more likely to be in the bad state than in the good state; and to buy (or hold) the stock when he believes that it is more likely to be in the good state.

Formally, let $p_{i,t}$ be the price of stock i in trial t , after any price update about the stock, and let $q_{i,t} = \Pr(s_{i,t} = \text{good} \mid p_{i,t}, p_{i,t-1}, \dots, p_{i,1})$ be the probability that a Bayesian investor, after seeing the price update in trial t , would assign to stock i being in the good state in trial t . Also, let z_t take the value 1 if the price update in trial t indicates a price increase for the stock in question; and -1 if the price update indicates a price decrease. Then $q_{i,t} = q_{i,t-1}$ if the price update in trial t was *not* about stock i ; but if the price update in trial t was about stock i , then:

$$\begin{aligned}
 q_{i,t}(q_{i,t-1}, z_t) &= \Pr(s_{i,t} = \text{good} \mid q_{i,t-1}, z_t) \\
 &= \frac{\Pr(z_t \mid s_{i,t} = \text{good})\Pr(s_{i,t} = \text{good} \mid q_{i,t-1})}{\Pr(z_t)} \\
 &= \frac{\Pr(z_t \mid s_{i,t} = \text{good})\Pr(s_{i,t} = \text{good} \mid q_{i,t-1})}{\Pr(z_t \mid s_{i,t} = \text{good})\Pr(s_{i,t} = \text{good} \mid q_{i,t-1}) + \Pr(z_t \mid s_{i,t} = \text{bad})\Pr(s_{i,t} = \text{bad} \mid q_{i,t-1})} \\
 &= \frac{(0.5 + 0.05z_t)(0.8q_{i,t-1} + 0.2(1 - q_{i,t-1}))}{(0.5 + 0.05z_t)(0.8q_{i,t-1} + 0.2(1 - q_{i,t-1})) + (0.5 - 0.05z_t)(0.2q_{i,t-1} + 0.8(1 - q_{i,t-1}))}.
 \end{aligned} \tag{1}$$

The optimal strategy for an expected value investor is to sell (if holding) or not buy (if not holding) stock i in trial t when $q_{i,t} < 0.5$; and to hold or buy it otherwise. This is because the expected price change on the next price update is given by

$$E_t[\Delta p_{i,t+1} \mid q_{i,t}, \Delta p_{i,t+1} \neq 0] = 0.6(2q_{i,t} - 1). \tag{2}$$

The left quantity in (2) is called the *net expected value (NEV)* of buying. A risk neutral subject will buy whenever NEV is positive ($q_{i,t} > 0.5$) and will not buy (or will sell) whenever NEV is negative ($q_{i,t} < 0.5$).

Note that a trader who follows the optimal strategy described above will exhibit the opposite of the repurchase effect. If a stock performed well on the last price update, it was probably in a good state for that price update. Since it is very likely to remain in the same state for its next price update, its next price change is likely to also be positive. The optimal strategy therefore involves buying stocks that have recently increased in price, and not buying stocks that have recently decreased in price, hence generating the opposite of a repurchase effect.

C. Behavioral predictions of the regret-devaluation model

We now lay out the predictions of the model where trading decisions are also driven, in part, by regret. While the psychology literature discusses several different regret mechanisms, we focus on a specific one called the regret-devaluation mechanism (Arkes et al. (2002)). The core idea of this theory is that after selling a stock, the subsequent price changes impact the investor's affective state: after selling a stock and seeing the price rise, the investor experiences regret. Repurchasing the stock would prolong the regret, creating an undeniable fact that the stock was sold at a low price and then bought again at a higher price. The subjective sense that this *ex-post* trading mistake is regretful is encoded as a disutility from repurchase, so the investor decreases his valuation (expected utility) of repurchasing the stock⁹ (Arkes et al. (2002)). In what follows, we assume this mechanism is symmetric¹⁰, which implies that subjects who observe the price fall

⁹ The hypothesized regret devaluation is similar to the Aesop fable of the “sour grapes” – a valued object (grapes) is mentally devalued because it cannot be reached.

¹⁰ See Lin et al. (2006) for a discussion on the symmetry assumption.

after selling will experience a positive affective response – often called “rejoicing” – and this increases the expected utility of repurchasing the stock.

Formally, if a subject sells stock i in period s at p_i^s , and sees the price change to $p_i^{s+1} > p_i^s$ in period $s+1$, he experiences regret that is proportional to $p_i^{s+1} - p_i^s$ at the moment the new price is revealed. This negative affective response changes the expected utility of repurchasing the stock by $\lambda(p_i^{s+1} - p_i^s)$, for some $\lambda < 0$. Suppose now that the subject decides not to repurchase the stock in period $s+1$, and the price subsequently increases to p_i^{s+2} ; in this case, the subject experiences another burst of regret and the expected utility of repurchasing the stock changes further by $\lambda(p_i^{s+2} - p_i^{s+1})$. In general, for any period $T > s$, the accumulated net regret from period s through T will cause a series of changes in expected utility of repurchasing the stock. Assuming that periodic regrets can be summed, this total change in expected utility of repurchasing the stock can be written as, $\sum_{t=s}^T \lambda(p_i^{t+1} - p_i^t) = \lambda(p_i^T - p_i^s)$. Hence, a sufficient statistic for the change in expected utility induced by the regret devaluation mechanism is simply the difference between the current price and last sale price, multiplied by a negative constant¹¹. From now on, we will refer to the variable, $(p_i^T - p_i^s)$, as the *foregone capital gain*.

An agent who is influenced by the regret-devaluation mechanism (in addition to the expected return on the stock), will compute the expected utility of repurchasing stock i as $EU(\text{repurchase}_i) = NEV_i + \lambda(p_i^T - p_i^s)$ for $\lambda < 0$. Note that the optimal trading strategy discussed above is a special case of this expected utility specification when $\lambda = 0$. When $\lambda < 0$,

¹¹ The key idea in the two-period model of Arkes et al. (2002) is that the initial mistake of omission (not buying an item on sale) makes a participant less likely to purchase the same item during a smaller sale. In the current setting, the first mistake is an error of commission (selling too early), and a subsequent price increase will lower the propensity to repurchase the asset. However, once the subject declines the opportunity to repurchase the stock, this represents an “inaction” mistake and all subsequent opportunities that the subject turns down to repurchase the stock can be traced to the most recent decision not to purchase.

the expected utility of repurchasing recent winners will decrease and the expected utility of repurchasing recent losers will increase, relative to the optimal strategy.

As a behavioral measure of the regret-devaluation mechanism, we use the repurchase effect as defined by Strahilevitz et al. (2011). Following their methodology, the repurchase effect is computed as the difference between the proportion of stocks that are repurchased after going down since the previous sale (PDR) and the proportion of stocks that are repurchased after going up since the previous sale (PUR):

$$PDR = \frac{\text{\# of stocks down since last sale that are repurchased}}{\text{\# of opportunities to repurchase stocks down since last sale}} \quad (3)$$

$$PUR = \frac{\text{\# of stocks up since last sale that are repurchased}}{\text{\# of opportunities to repurchase stocks up since last sale}} \quad (4)$$

The size of the repurchase effect is defined as the difference between these two frequencies, $(PDR - PUR)$. Note that because the repurchase effect is defined as a difference in ratios, the effect itself should be independent of the number of opportunities the subject has to repurchase stocks with foregone capital gains or with foregone capital losses. For the particular price process used in our experiment, a trader following the optimal strategy will exhibit a measure of $PDR - PUR = -0.75$. In contrast, for a trader who is also influenced by the regret devaluation mechanism, we expect $\lambda < 0$, and as discussed above, this will cause the repurchase effect to be greater than it would be under the optimal strategy. This leads us to make the following predictions:

Prediction 1A (Behavioral): If subjects are influenced by the regret-devaluation mechanism, then they will exhibit a repurchase effect that is larger than that predicted by the optimal strategy. In particular, $PDR-PUR > -0.75$.

Prediction 1B (Behavioral): If subjects are influenced by the regret-devaluation mechanism, then the foregone capital gain should be a negative predictor of the repurchase decision ($\lambda < 0$).

D. Neural predictions of the regret-devaluation model

We now turn to the neural predictions. Based on substantial prior research in neuroeconomics, we develop two hypotheses which correspond to the neural signals that are generated from the regret-devaluation mechanism. The first neural signal is the prediction error at the time at which stock prices are updated (but before trading decisions are made). The second neural signal is the decision value of repurchasing. We now develop each of these two hypotheses in detail.

The price update screen gives new information about the state of the updated stock. In neuroscience terms, this is a kind of *reward prediction error* (RPE). In general, an RPE measures the change in expected present value of utility induced by the new information. In our design, the RPE is the change in the expected utility of repurchasing the stock whose price is updated. Crucially, if the regret devaluation mechanism is responsible for the repurchase effect, then the regret generated upon a price increase should be detected in the RPE, since the regret induces a change in the utility of the updated stock.

A large amount is known about RPEs in neuroeconomics. RPEs are very commonly found in the area near the center of the brain called the ventral striatum (vSt). vSt is part of the basal ganglia, an evolutionarily-conserved part of the brain (common to mammalian and other

species). Evidence for RPEs in this area comes from two main sources: recording of single-neuron firing rates from implanted electrodes and human fMRI. The evidence from single-neuron studies is particularly strong and broad: Romo and Schultz (1990) showed early on that neurons in the midbrain, which are known to project to vSt, fire when unexpected rewards are presented, and scale with the size of the reward (Schultz et al. 1997). There is also evidence that the RPE is distinct from a simple response to reward (Hare et al. 2008.) Many other studies of human fMRI have also shown clear signals of RPE in vSt¹².

Central to the novelty of our research is the hypothesis that when a stock is not owned, the news from a price increase will generate an RPE that reflects, in part, the regret from selling too early. The neural evidence for this hypothesis is much less extensive and clear than the voluminous literature on RPE encoding in general. Lohrenz et al. (2007) were the first to show an fMRI response in vSt to “fictive learning” – that is a response to the reward of an action that was not chosen¹³, though they specifically avoid tying their result to regret.

Several other studies have examined regret, defined decision-theoretically as the difference between the payoff from an unchosen action and the received payoff. Early regret studies focused attention mostly on activity in medial orbitofrontal cortex (mOFC). We will explain why that selective focus on mOFC came about, and describe why many later studies clearly show that vSt also (and perhaps more robustly) encodes regret.

¹² See, for example, McClure et al. (2003), O’Doherty et al. (2003), Bayer and Glimcher (2005), Abler et al. (2006), Pessiglione et al (2006), Glascher et al. (2010), Daw et al. (2010), Caplin et al. (2010), Rutledge et al. (2010), and Lin et al. (2012).

¹³ Their design is different than ours and shows a different result. Their subjects invest a fraction of their wealth, b_t , in each period in a single stock index (drawn from historical data). The “fictive error” is the difference between the return r_t and $b_t * r_t$. They show positive encoding of this signal in vSt. Behaviorally, subjects respond positively to the difference $(r_t - b_t * r_t)$ for positive returns, showing the opposite of regret devaluation (and show a weak rejoicing-valuation response for negative returns). We think their result is due to the fact that subjects are typically partly invested at all times, so that there is no sharp distinction between owning and not owning a stock, which reduces the effect of regret devaluation.

Camille et al. (2004) used a task in which subjects chose between two circular spinners. Each spinner had two shaded regions indicating the amounts of points won or lost if the spinner arrow, when it stopped after spinning, pointed to those regions. In a partial feedback condition, subjects only saw the spinner arrow for the choice that they made. In a complete feedback condition, subjects saw both spinner arrows (so they can see exactly what counterfactual payoff they would have gotten if they made the foregone choice). Unlike neurotypical control subjects, their small sample ($n=5$) of patients with damage to mOFC showed a blunted emotional reaction to counterfactual payoffs, with emotional reaction (conventionally) measured by both self-report and skin conductance. However, note that since this is a lesion study measuring only behavior (not whole-brain activity), the result does not imply that mOFC is exclusively sensitive to regret. Ventral striatum could have been activated by regret. Indeed, given strong fiber connectivity between areas of frontal and striatal cortex, it is quite possible that ventral striatal areas would have shown activity if they had imaged the participants.

Coricelli et al. (2005) is a sequel study by the same group that did use whole-brain imaging to explore all possible neural activity encoding regret. Using the same two-spinner choice task, they did find activity in putamen (-14,0,6), as well as lateral OFC (not medial) and in two other regions (ACC, IPL). Putamen is part of the dorsal striatum (it is higher up toward the top of the brain or “dorsal” to the lower ventral region). However, they do not report how large the active area is, so it is possible that while the peak activity is in putamen, the entire active area includes ventral striatum¹⁴.

¹⁴ In further analysis, they do not revisit the putamen activity and instead concentrate on mOFC. We think their intention was to most clearly link the evidence from damage to mOFC with evidence imaged from mOFC (while not denying parallel activity in vSt). Furthermore, they find a much stronger difference between win and loss encoding in vSt when the subject makes their own choice, compared to a neutral condition in which subjects merely watch a computer-generated choice played out. This stronger contrast is exactly what one would expect if regret was revaluing wins and losses based on personal regret.

Eryilmaz et al. (2014) used a risky choice paradigm and focused on regret-related activation during a 90-second period after learning about regret or rejoicing outcome. This is an apt comparison to our design, in which information that potentially generates regret is first given (the price update) and then a choice is made a few seconds later. They find positive activity in vSt in response to rejoicing (the opposite of regret) during post-update regret. They also report a negative correlation between the strength of rejoicing encoding in vSt and measured levels of depression. Further evidence of regret signals in the vSt can be found in Camille et al. (2010), Buchel et al. (2011) and Brassen et al. (2012).

Thus, while the early neuroscience literature emphasized the role of mOFC in processing regret, *every* study we know of that has looked at the whole brain (beginning with Coricelli et al. (2005)) has found evidence that regret signals are encoded in vSt or in nearby striatal areas. We therefore focus on measures of neural activity in this area of the brain at the time news is revealed to the subject; in particular, at the time of the price update screen.

Recall that a subject who is influenced by both the optimal trading strategy and the regret-devaluation mechanism will compute the expected utility of repurchasing the stock as: $EU(\text{repurchase}_i) = NEV_i + \lambda(p_i^T - p_i^S)$. Owing to the specific Markov chain that governs price changes in our experiment, when a subject sees a positive price update, the probability that the stock is in the good state increases, and this leads to an increase in the NEV. However, a price increase will also lead to a greater value of the regretted foregone capital gain, $(p_i^T - p_i^S)$. Therefore, the net effect of the price increase on the expected utility of repurchase depends on the value of λ . If λ is strongly negative, then the change in expected utility of repurchasing the stock, which is equivalent to the prediction error signal, should negatively correlate with the price change. Because prediction errors are encoded in the vSt, we therefore expect that the vSt should

negatively correlate with the price change at the price update screen for a stock that is not owned.

This leads to our second prediction:

Prediction 2 (Neural): If subjects are sufficiently influenced by the regret-devaluation mechanism, then neural activity in the vSt should negatively correlate with the price change at the price update screen for a stock that is not held.

The next prediction relates to the heterogeneity in the strength of the prediction errors and repurchase effects across subjects. Subjects who are influenced by the regret-devaluation mechanism to a large degree will have strongly negative values of λ . Behaviorally, this means that these subjects will have very low (high) expected utilities of repurchasing winners (losers), and therefore they will exhibit large repurchase effects. Neurally, it predicts that these subjects will generate prediction errors that are strongly negative (positive) at the time a positive (negative) price update is revealed for a stock that is not owned. This leads us to our third prediction:

Prediction 3 (Neural): The strength of the prediction error generated at the time of the price update should explain a portion of the cross-subject heterogeneity in the size of the repurchase effect. The repurchase effect should negatively correlate with the vSt response to price updates for a stock that is not held.

Predictions 2 and 3 relate to the prediction error signals that are encoded at the time of price updating. However, the regret-devaluation mechanism predicts that these signals are also used to update the subjective expected utility, or “decision value,” of repurchasing the stock.

Such decision values are thought to reflect the relative value of obtaining the option under consideration versus staying with the status quo. Decision values are important for guiding choice and have been robustly found in an area of the brain called the ventromedial prefrontal cortex (vmPFC) (Hsu et al. (2005), Padoa-Schioppa and Assad (2006), Kable and Glimcher (2007), Knutson et al. (2007), Chib et al. (2009), and Hare, Camerer and Rangel (2009)). Two recent meta-analyses statistically integrating hundreds of studies have shown strong evidence across decisions involving money and primary rewards (such as foods and erotic pictures) that vmPFC encodes decisions values (Bartra et al. (2013) and Clithero and Rangel (2014)).

In our particular setting, the time of choice occurs at the trading screen, and the option under consideration is a stock repurchase. The decision value of repurchasing the stock is therefore equal to the expected utility from repurchasing the stock minus the expected utility from not repurchasing. Because the expected utility from not repurchasing is zero, the decision value of repurchasing reduces to the expected utility of repurchasing.

For a subject who is sufficiently influenced by the regret-devaluation mechanism, this decision value at the time of choice is well-approximated by the opposite of the foregone capital gain. To see this, note that the expected utility of repurchasing a stock under the regret-devaluation mechanism is given by: $EU(repurchase_i) = NEV_i + \lambda(p_i^T - p_i^S)$. When λ is sufficiently negative – as is the case when a subject is strongly influenced by the regret-devaluation mechanism – the second term will swamp the first term and the overall expected utility is well-approximated by $\lambda(p_i^T - p_i^S)$. This leads to Prediction 4:

Prediction 4 (Neural): When a subject who is influenced by the regret-devaluation mechanism is presented with the opportunity to repurchase a stock at the trading screen, activity in the vmPFC should negatively correlate with the foregone capital gain.

II. Results

A. Test of Behavioral Predictions

We begin our analyses of the trading data by computing the size of the repurchase effect for each subject. We find that the average PDR and PUR across subjects are .301 and .337, respectively. Consistent with Prediction 1A, the average PDR-PUR value is -0.029, which is much less negative than the optimal level of PDR-PUR = -0.75 ($p < 0.001$). Figure 2 shows there is large variation in the size of the repurchase effect across subjects and that the repurchase effect is greater than the level displayed by an optimal trader for 26 of the 28 subjects.

Table 1 provides data on each component of the repurchase effect, disaggregated by subject and by experimental session. These data allow us to test whether subjects exhibit learning between the two sessions. The mean repurchase effect is smaller in the second session compared to the first session (0.01 versus -0.07, a reduction of 12.4% relative to the optimal level of -0.75), although the difference is not statistically significant ($p = 0.34$). There is also a strong 0.62 ($p < 0.001$) correlation of the individual-level repurchase effects across sessions, showing that the repurchase effect is persistent within subjects.

We next test Prediction 1B by running a logistic regression of the repurchase decision on the NEV and the foregone capital gain. Under the regret-devaluation mechanism, we expect the foregone capital gain to be a negative predictor of the repurchase decision. Column 1 in Table 2 shows the estimation results from this regression, clustering the standard errors by subject. Consistent with Prediction 1B, the foregone capital gain is a negative and significant predictor of the decision to repurchase¹⁵. Column 2 displays results when the foregone capital gain is broken

¹⁵ Two other results in column (1) of Table 1 deserve further explanation. First, the estimated coefficients on the NEV and foregone capital gain differ by several orders of magnitude. This is because the ranges of those two variables are quite different: the standard deviations are 0.073 for NEV and 18.45 for the foregone capital gain. Second, because our experimental design induces positive autocorrelation in the

into its two components, p_i^T and p_i^S ; these results indicate that the negative relationship between the foregone capital gain and the decision to repurchase is driven mainly by the current price. One reason for this asymmetry may be that the current price is prominently displayed on the decision screen, while the last sale price is not. Columns (3) and (4) re-estimate the regressions using a linear probability model and yield similar results to the logistic regressions. Note that the statistical strength of NEV, as a correlate of repurchase, is only significant at conventional levels in the logit specifications. This is because repurchase effects and NEV-buying cannot comfortably co-exist: repurchasing losers is correlated with buying assets with negative NEV.

One concern with the models in columns (1) – (4) is that the assumption of independence of observations within a subject and over time may not hold. In particular, the decision to repurchase a stock on trial t is dependent on the subject not purchasing the stock on trial $t-1$. To address this concern, we run a survival analysis using a Cox proportional hazard model, similar to the analysis done in Strahilevitz, Odean and Barber (2011). In particular, we estimate the hazard rate of repurchasing a stock on trial t , which is the probability of repurchasing the stock on trial t conditional on 1) not repurchasing through trial $t-1$ and on 2) a set of time-varying covariates.

In column (5), we find that the hazard ratio of the foregone capital gain is 0.994, which indicates that a \$1 increase in the foregone capital gain yields a hazard rate of 99.4% relative to the stock's hazard rate before \$1 increase in foregone capital gain, although this result is not statistically significant. The results in column (6) show that a \$1 increase in the current price yields a hazard rate that is 98.7% of the stock's hazard rate before the \$1 price increase. This hazard rate, while close to one, is significantly less than one at the 1% level. In summary, we find that the average subject exhibits a strong repurchase effect, and while the foregone capital gain is

price changes, the NEV and foregone capital gain variables are positively correlated. This collinearity causes a potential concern about identification, so as a robustness check we re-estimate the logistic regression at the individual subject level to examine the joint distribution of the two coefficients. We find that the NEV and foregone capital gain coefficients are not correlated ($p=0.392$), suggesting that there are no major identification problems arising from the correlated regressors.

a negative predictor of the repurchase decision, this effect is mainly driven by the current price component.

B. Test of Neural Prediction 2

We now turn to Prediction 2, which examines the sign of the prediction error generated at the revelation of a price update for an asset not held. In the appendix, we provide a brief primer for economists on fMRI methodologies that can aid in understanding the following neural results.

If subjects are sufficiently influenced by the regret devaluation mechanism, then a positive price update should lower the overall expected utility of repurchasing the asset, even after taking the increased NEV into account. Therefore, under the regret devaluation model, activity in the vSt should negatively correlate with the price change at the time of a price update screen for an asset not held.

The pre-specified region of the vSt in which we test for prediction errors is taken from Frydman et al. (2014) and is the area colored in yellow and orange in Panel C of Figure 3 (68 voxels)¹⁶. To carry out the main test of Prediction 2, we estimate the following general linear model (GLM) for BOLD activity in every subject and voxel:

$$b^v(t) = \alpha + \beta_1^v(\Delta p_t)I_{own\ update}(t) + \beta_2^v(\Delta p_t)I_{no\ own\ update}(t) \quad (5)$$

$$+ \beta_3^v(p_t - p_s)I_{repurchase\ opp}(t) + \beta_4^v(p_t - c)I_{sell\ opp}(t) + \beta_5^v controls + \varepsilon(t)$$

¹⁶ The pre-specified vSt region we use is constructed by taking a sphere with 15 mm radius around the coordinates (MNI-space x=-15, y=6, z=-12) found to exhibit peak correlation with prediction errors in Lin, Adolphs and Rangel (2012). We then intersect this sphere with an anatomical mask of the vSt identified using the Talairach Daemon atlas within the PickAtlas software (Maldjian et al. (2003)), which includes the nucleus accumbens and the ventral putamen.

where $b_v(t)$ denotes the BOLD signal at time t in voxel v . $I_{own\ update}$ ($I_{no\ own\ update}$) is an indicator function that equals one if, at time t , the subject is presented with a price update screen for an owned (not owned) stock. The third nonconstant regressor is an interaction between the foregone capital gain and an indicator that equals one if at time, t , the subject is presented with a repurchase opportunity. Similarly, the fourth nonconstant regressor is an interaction between the holding period capital gain and an indicator that equals one if at time, t , the subject is presented with a selling opportunity.

The “controls” vector includes the following variables: (1) an indicator function denoting the onset of a price update screen for an owned stock, (2) an indicator function denoting the onset of a price update screen for a stock not owned, (3) an indicator function denoting the onset of a repurchase opportunity, (4) an indicator function denoting the onset of a repurchase opportunity interacted with the NEV of buying, (5) an indicator function denoting the onset of a selling opportunity, (6) an indicator function denoting the onset of a selling opportunity interacted with the NEV, (7) regressors controlling for physical movement inside the scanner, and (8) session indicator variables. Controls 1 to 6 are convolved with the HRF, whereas control 7 and 8 are not. Finally, inferences about whether the variables of interest are encoded within a voxel are made by carrying out a one-sided t -test against zero of the average of the individually estimated coefficients, and by correcting for multiple comparisons within the pre-specified vSt region of interest.

As shown in Panel A of Figure 3, the results are consistent with the regret-devaluation model. Within the pre-specified 68-voxel vSt target region that has been shown to correlate with prediction errors, we find a cluster of 19 voxels in which β_2^v , averaged across subjects, is significantly negative. Panel B shows that there is a cluster of 34 voxels in which β_1^v , averaged across subjects, is significantly positive. Panel C shows the conjunction of voxels (colored red

and orange) in the pre-specified vSt target region (colored yellow and orange) for which β_1^v is significantly positive and β_2^v is significantly negative. Finally, Panel D provides the average β_1^v and β_2^v , averaged across subjects and averaged across voxels within the pre-specified vSt target region. This figure shows that within subjects, the vSt response to price changes is a function of whether the subject owns the stock or not.

One potential concern in interpreting the vSt signal at the price update screen is that it may be driven by a “preparatory” decision value signal rather than a prediction error. In other words, even though subjects do not have the ability to execute a trading decision at the price update screen, they may nonetheless begin thinking about their trading decision, should they be given a trading opportunity. To test for this possibility, we re-estimated the GLM in equation (5), but also included two additional regressors at the time of the price update screen when a subject does not hold the stock. Specifically, we included 1) the NEV conditional on the new price and 2) the foregone capital gain conditional on the new price. These two variables represent the decision values under the optimal trading strategy and regret devaluation mechanism, respectively. Therefore, if the vSt signal that we originally detect is driven by a “preparatory” decision value computation, we expect to see activity in the vSt significantly correlate with these variables.

Contrary to this hypothesis, we find no voxels in the vSt for which the coefficient on either of the two additional regressors are different from zero at our omnibus threshold of $p < 0.05$ SVC. Because decision values are reliably found in the vmPFC, we also tested for signals in this region, and again find no voxels that exhibit activity that correlated with either decision value at our omnibus threshold. Instead, we find that the original 19 voxels in the vSt that significantly negatively correlate with the price change still exhibit this correlation even after including the foregone capital gain and NEV regressors. Taken together, these results suggest that the vSt

signal that we detect is more likely to be driven by a prediction error than a decision value computation.

C. Test of Neural Prediction 3

We now test Prediction 3, which builds on the GLM and analysis we used to conduct our tests in the previous section. Recall that if a subject is heavily influenced by the regret-devaluation mechanism, we should observe a strong negative (positive) prediction error when the subject sees a positive (negative) price update for a stock he doesn't own. In other words, a subject who is heavily influenced by regret will exhibit a strongly negative measure of β_2^v in the vSt target region. This same subject should also exhibit a large repurchase effect, as he avoids repurchasing winners and is keen to repurchase losers. Hence, we test for a negative correlation between the repurchase effect and β_2^v . For each subject, we first compute the average β_2^v across the 68 voxels in the vSt target region. Consistent with Prediction 3, Figure 4 shows that the correlation, across subjects, between β_2 and the PDR-PUR measure is -0.399 ($p=0.037$).

D. Test of Neural Prediction 4

The results from neural predictions 2 and 3 are so far consistent with the regret-devaluation model. In particular, these results suggest that upon receipt of news about a stock price change, the vSt computes a prediction error, which can be used to update the value of repurchasing the stock. This value updating procedure should lead to a decision value that is used to guide choice when presented with an opportunity to repurchase a stock. Prediction 4 states that this decision value should equal the opposite of the foregone capital gain, and it should be reflected in vmPFC activity at the time of choice.

We test for the presence of this decision value in a pre-specified region of the vmPFC, which consists of 429 voxels and has been shown to reflect the computation of decision values in previous studies¹⁷. We use the estimation results from regression (5) and perform hypothesis tests about β_3^v , the coefficient on the foregone capital gain. As shown in Panel A of Figure 5, the results from this test are consistent with the regret-devaluation model as we find a cluster of 22 voxels in which β_3^v , when averaged across subjects, is significantly negative.

It is also worth noting that this cluster of 22 voxels exhibits an overlap with the voxels that were previously found to encode the capital gain at the time of a selling decision, which represents the decision value of selling under the realization utility model of trading (Frydman et al. (2014)). In particular, Panel B of Figure 5 shows there is a cluster of 71 voxels in the pre-specified region of vmPFC for which β_4^v , when averaged across subjects, is significantly positive. Panel C shows there is an overlap of 14 voxels that encode both the decision value of selling under the realization utility model and the decision value of repurchasing under the regret-devaluation model. The next section further investigates this link between buying and selling decisions.

E. Additional results on the link between buying and selling behavior

Motivated by the result that the same sub-region of the vmPFC is responsible for computing the decision value of selling and purchasing decisions, we investigate whether these two types of decisions are themselves correlated. Previous work using the same data set finds that subjects exhibit a significant disposition effect; that is, they exhibit a greater tendency to sell

¹⁷ The pre-specified vmPFC region we use is constructed by taking a sphere with 15 mm radius around the coordinates (MNI-space $x=3$, $y=36$, $z=-18$) found to exhibit peak correlation with decision values in (Plassmann, O'Doherty and Rangel (2010)). We then intersect this sphere with an anatomical mask of the vmPFC that was identified using the AAL digital atlas of the human brain (Tzourio-Mazoyer et al. (2002)), and includes the rectus, the orbital part of the superior frontal gyrus, and the orbital part of the middle frontal gyrus.

stocks at a gain compared to stocks at a loss (Frydman et al. 2014). This effect is costly in our experiment for the same intuition that the repurchase effect is costly. If a stock is trading at a capital gain then it is likely to have accrued this gain because of the stocks's strong recent performance. Because our experimental design induces positive short-term autocorrelation in price changes, the strong recent performance predicts, on average, prices will continue to rise. Therefore, selling a winning stock is, on average, a costly mistake. The same logic shows that holding losers is also a costly mistake.

Following Odean (1998), we compute the disposition effect as the difference between a subject's proportion of gains realized (PGR) and his proportion of losses realized (PLR), where each ratio is defined as follows:

$$PGR = \frac{\# \text{ of realized gains}}{\# \text{ of realized gains} + \# \text{ of paper gains}} \quad (6)$$

$$PLR = \frac{\# \text{ of realized losses}}{\# \text{ of realized losses} + \# \text{ of paper losses}} \quad (7)$$

In our experiment, an optimal Bayesian trader will exhibit a $PGR - PLR$ measure of -0.76, but we find the average subject exhibits a $PGR - PLR$ measure of 0.23. There is also significant variation in the disposition effect, and Figure 6A shows that there is a strong positive correlation, across subjects, between the disposition effect and repurchase effect ($r=0.71$, $p<0.001$). In other words, subjects who exhibit purchasing mistakes are also more likely to exhibit selling mistakes, suggesting that there may be a common psychological mechanism that governs both types of trading mistakes.

To investigate this further, we test whether the sub-components of each effect are also correlated. In particular, previous work documents that PGR and PLR are statistically independent (Dhar and Zhu (2006), Weber and Welfens (2008), Frydman et al. (2014)), suggesting that traders who are quick to realize gains are not necessarily slow to realize losses.

We therefore tested for the analogous results in the buying data. Indeed we find that PDR is independent of PUR ($p=0.29$, under the null that the two variables are uncorrelated), suggesting that traders who are quick to repurchase stocks that have gone down, are not necessarily slow to repurchase stocks that have gone up.

However, traders who are slow to realize losses are also slow to repurchase stocks that have gone up since last sale, as demonstrated by the positive correlation between PLR and PUR in Figure 6B ($r=0.58$, $p<0.001$). Similarly traders who are quick to realize gains are also quick to repurchase stocks that have gone down since last sale, as demonstrated by the positive correlation between PGR and PDR in Figure 6C ($r=0.60$, $p<0.001$). Taken together, these results suggest that the two components within each behavioral effect are likely governed by different psychological mechanisms; yet, the analogous components *across* behavioral effects are likely to be governed by the same psychological mechanism.

F. Ruling out incorrect beliefs

An alternative theory that can potentially explain the disposition effect, the repurchase effect, *and* their high correlation is driven by subjects' beliefs. Although subjects are given full information about the stochastic process that generates the price changes, it is possible that subjects may not form beliefs in a Bayesian manner. In particular, if subjects (irrationally) believe that prices exhibit mean reversion, then they will purchase recent losers and sell recent winners. This will induce subjects to exhibit both a disposition effect and a repurchase effect, and moreover, the two effects will be positively correlated across subjects.

To rule out this alternative belief-based theory for our data, we run an additional laboratory experiment (outside of the fMRI scanner). We reasoned that, if belief in mean-reversion is driving the results, then explicitly providing subjects with the NEV should reduce the

repurchase effect and disposition effects to the level that an expected value trader would exhibit. We recruited an additional N=30 subjects to participate in an “NEV treatment” that is identical to the experiment used in the fMRI scanner, except for two key differences: (i) the NEV is explicitly displayed on both the price update and trading screen in every trial and (ii) the parameters of the stochastic process are changed to $\Pr(\text{price increase} \mid \text{good state}) = \Pr(\text{price decrease} \mid \text{bad state}) = 0.7$. By explicitly displaying the NEV, any deviations from expected value behavior are unlikely to be driven by non-Bayesian belief updating. By increasing $\Pr(\text{price increase} \mid \text{good state}) = \Pr(\text{price decrease} \mid \text{bad state})$ from 0.55 to 0.7, each price change becomes a more informative signal of the hidden Markov state, thus making both the repurchase effect and disposition effect even more costly.

Because the incentive structure has changed, we cannot compare behavior among subjects in the “NEV treatment” to behavior of subjects in our fMRI experiment. We therefore rely on data from a control condition comprised of N=33 subjects that were collected and previously analyzed in Frydman and Rangel (2014). The incentive structure in this control condition is identical to that used the “NEV treatment.” Therefore, the only difference between this control condition and the “NEV treatment” is the explicit display of the NEV on both the price update and trading screens.

Figure 7 shows that explicitly displaying the NEV does increase the average percentage of decisions that are optimal compared to the control condition. However, the average percentage of decisions that are optimal in all four decision types – sell, hold, repurchase, and no repurchase – is still significantly below the optimal level ($p < 0.001$ for each decision category)¹⁸. This evidence suggests that while a portion of the behavioral and neural data we observe in the fMRI

¹⁸ We also note that that even after increasing a subject’s incentive to implement the optimal strategy (by increasing $\Pr(\text{price increase} \mid \text{good state}) = \Pr(\text{price decrease} \mid \text{bad state})$ from 0.55 to 0.7), we still find that a substantial portion of trading decisions are suboptimal. This shows that the overall disposition effect and repurchase effects are unlikely to be driven by the incentive structure that we impose, which is governed by the degree of predictability in price changes.

data set may be driven by irrational beliefs, there remains a substantial portion of the data that an irrational belief in mean reversion cannot explain¹⁹.

Another important piece of evidence against the mistaken beliefs explanation is from the brain activity: recall that given the same event—e.g., a stock goes up—there is a positive vSt signal when a subject owns the stock and a negative signal when a subject does not own the stock (Figure 3D). In other words, the prediction error, which measures the change in expected utility generated by news contained in the price change, is a function of stock ownership. However, the belief in mean-reversion theory predicts that the direction of belief updating should be independent of stock ownership. Thus, the brain activity at time of price update also casts doubt on the belief in mean-reverting prices hypothesis.

III. Discussion

In this paper we use neural data collected while subjects trade in an experimental stock market to directly test whether regret can have an impact on investor behavior. The neural data is helpful for conducting empirical tests on regret because it allows us to obtain inferred measures of regret arising from behavior. Our approach linking only behavior and neural activity, however, means that we did not collect any other measures (e.g., self-reports), which can be informative about the subjective or emotional content of regret. Further experiments could usefully do so.

¹⁹ Two recent studies on the disposition effect also cast doubt on the incorrect beliefs hypothesis. Fischbacher, Hoffmann, and Schudy (2015) allow subjects to pre-commit to selling losers, by placing stop-loss orders. Many subjects do implement this stop-loss commitment and they exhibit smaller disposition effects than control subjects who cannot place stop-loss orders. Investors who mistakenly believe in mean-reversion will not place stop-loss orders, so the frequency of those orders is evidence against belief mistakes. Furthermore, the causal effect of commitment devices in reducing the disposition effect suggests that there may be a dynamic inconsistency involving self-control that contributes to the disposition effect: it appears that some subjects understand they will be averse to realizing a loss in the future, so they pre-commit to doing so at the time of purchase via a stop-loss order. In recent work using field data, Frydman, Hartzmark, and Solomon (2015) show that the disposition effect is unlikely to be driven by a belief-explanation as the purchase price of a *previously* sold stock affects the propensity to sell a currently owned stock.

The main results of our neural tests are consistent with the idea that investors do experience regret when receiving information that indicates a trading decision is ex-post suboptimal. In particular, at the moment when a subject observes a price increase for a recently sold stock, we observe a decrease in brain activity in the vSt. The location of this activity is consistent with other studies that have shown vSt response to prediction errors and regret-related signals. Moreover, the strength of this neural regret signal correlates with the size of the repurchase effect across our subject pool. It is important to note that our interpretation of vSt activity at the price update screen as a regret signal is based on the combination of two facts: (i) the large body of evidence showing that the vSt encodes reward prediction errors (see section I.D.) and (ii) regret should be detected in the reward prediction error signal, if the regret-devaluation mechanism is responsible for the observed trading behavior.

We also find that neural activity in the vmPFC negatively correlates with the foregone capital gain at the moment when a subject is presented with the opportunity to repurchase a stock. This result is consistent with a large body of evidence from decision neuroscience that finds decision values are encoded in the vmPFC. This result provides the critical motivation for our subsequent empirical tests that examine the link between buying and selling behavior. In particular, because the specific subarea of the vmPFC that we find to encode the foregone capital gain at a buying opportunity is nearly the exact subarea that we find to encode the capital gain at the time of a selling decision, this leads us to conduct tests on the relationship between buying behavior (repurchase effect) and selling behavior (disposition effect). Our results indicate that these two behaviors are highly correlated, suggesting there may be a single psychological mechanism that generates both trading biases. One interpretation is that regret itself can provide a microfoundation for realization utility. Another interpretation is that regret is closely related to the psychological phenomenon of cognitive dissonance, which has recently been proposed to explain the disposition effect (Chang et al. (2014)).

However, yet another interpretation is that both effects are driven by a belief-based explanation. Our follow-up behavioral experiment is designed precisely to investigate this hypothesis. We find there are still a significant number of suboptimal trading decisions even when the Bayesian expected stock price change is displayed (although the percentage of both optimal buying and selling decisions does increase). This casts doubt on the hypothesis that the disposition effect and repurchase effect are completely driven by mistaken beliefs.

While the neural data do allow us to perform biologically-direct tests for the existence of regret signals, one limitation of our study is that we conduct these tests under a single model of regret: the so called “regret devaluation” model. There are other models of regret in the psychology literature that would make different predictions in our setting. For example, Shiller (2005) proposes that when an investor regrets missing out on investing early in a rising stock market, such regret will induce the investor to take action and buy stocks. Note that the effect Shiller describes is not a *repurchase* effect, and Shiller also emphasizes the likely role of social contagion, and observing returns of others, as emotional fuel for this effect. In Shiller’s example, regret is generated because a neighbor has done well in the stock market and the investor regrets his failure to participate in the market in comparison to his neighbor. Therefore an important open question is how the presence of social influence interacts with regret to impact investor behavior. Fortunately, this question can also be explored using neural data too, as shown by other recent studies in neurofinance that collect neural data from multiple subjects simultaneously (Lohrenz et al. (2013), Smith et al. (2014); Frydman (2015)).

REFERENCES

- Abler, Birgit, Henrik Walter, Susanne Erk, Hannes Kammerer, and Manfred Spitzer, 2006, Prediction error as a linear function of reward probability is coded in human nucleus accumbens, *Neuroimage* 31, 790-795.
- Arkes, Hal R., Yi-Han Kung, and Laura Hutzel, 2002, Regret, valuation, and inaction inertia, *Organizational Behavior and Human Decision Processes* 87, 371-385.
- Barber, Brad M., and Terrance Odean, 2013, The behavior of individual investors, in George Constantinides, Milton Harris, and Rene Stulz, eds.: *Handbook of the economics of finance* (Elsevier, Amsterdam).
- Barberis, Nicholas, Ming Huang and Tano Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1-53.
- Barberis, Nicholas, Ming Huang and Richard Thaler, 2006, Individual preferences, monetary gambles, and stock market participation: A case for narrow framing, *American Economic Review* 96, 1069-1090.
- Barberis, Nicholas, and Richard Thaler, 2003, A survey of behavioral finance, in George Constantinides, Milton Harris, and Rene Stulz, eds.: *Handbook of the economics of finance* (Elsevier, Amsterdam).
- Barberis, Nicholas, and Wei Xiong, 2012, Realization utility, *Journal of Financial Economics* 104, 251-271.
- Bartra, Oscar, Joseph McGuire, and Joseph Kable, 2013, The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value, *Neuroimage* 76, 412-427.
- Bayer, Hannah, and Paul Glimcher, 2005, Midbrain dopamine neurons encode a quantitative reward prediction error signal, *Neuron* 47, 129-141.
- Brassen, Stefanie, Matthias Gamer, Jan Peters, Sebastian Gluth and Christian Buchel, 2012, Don't look back in anger! Responsiveness to missed chances in successful and unsuccessful aging, *Science*, 336, 612-614.
- Buchel, Christian, Stefanie Brassen, Juliana Yacubian, Raffael Kalisch, and Tobias Sommer, 2011, Ventral striatal signal changes represent missed opportunities and predict future choice, *Neuroimage* 57, 1124-1130.
- Camille, Nathalie, Giorgio Coricelli, Jerome Sallet, Pascale Pradata-Diehl, Jean-Rene Duhamel and Angela Sirigu, 2004, The involvement of the orbitofrontal cortex in the experience of regret, *Science* 304, 1167-1170.
- Camille, Nathalie, V. A. Pironti, C. M. Dodds, M. R. F. Aitken, T. W. Robbins, and L. Clark, 2010, Striatal sensitivity to personal responsibility in a regret-based decision-making task, *Cognitive, Affective, & Behavioral Neuroscience* 10, 460-469.
- Campbell, John, 2006, Household finance, *Journal of Finance* 61, 1553-1604.
- Caplin, Andrew, Mark Dean, Paul Glimcher, and Robb Rutledge, 2010, Measuring beliefs and rewards: a neuroeconomic approach, *Quarterly Journal of Economics* 125, 923-960.

- Chib, Vikram, Antonio Rangel, Shin Shimojo, and John O'Doherty, 2009, Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex, *Journal of Neuroscience* 29, 12315-12320.
- Chang, Tom, David Solomon, and Mark Westerfield, 2014, Looking for someone to blame: Delegation, Cognitive Dissonance and the Disposition Effect, *Journal of Finance*, Forthcoming
- Clithero, John and Antonio Rangel, 2014, Informatic parcellation of the network involved in the computation of subjective value, *Social, Cognitive and Affective Neuroscience* 9, 1289-1302.
- Coricelli, Giorgio, Hugo D. Critchley, Mateus Joffily, John P. O'Doherty, Angela Sirigu, and Raymond J. Dolan, 2005, Regret and its avoidance: A neuroimaging study of choice behavior, *Nature Neuroscience* 8, 1255-1262.
- Daw, Nathaniel, Samuel Gershman, Ben Seymour, Peter Dayan, and Raymond Dolan, 2011, Model based influences on humans' choices and striatal prediction errors, *Neuron* 69, 1204-1215.
- Dhar, Ravi, and Ning Zhu, 2006, Up close and personal: Investor sophistication and the disposition effect, *Management Science* 52, 726-740.
- Eryilmaz, Hamdi, Dimitri Van De Ville, Sophie Schwartz, and Patrik Vuilleumier, 2014, Lasting Impact of Regret and Gratification on Resting Brain Activity and Its Relation to Depressive Traits, *Journal of Neuroscience*, 34, 7825 – 7835.
- Fischbacher, Urs, Gerson Hoffmann, and Simeon Schudy, 2015, The causal effect of stop-loss and take-gain orders on the disposition effect, *Working paper*.
- Frydman, Cary, 2015, What drives peer effects in financial decision-making? Neural and behavioral evidence, *Working paper*.
- Frydman, Cary, Samuel Hartzmark, and David Solomon, 2015, Rolling mental accounts, *Working paper*.
- Frydman, Cary and Antonio Rangel, 2014, Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price, *Journal of Economic Behavior & Organization* 107, 541-552.
- Frydman, Cary, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel, 2014, Using neural data to test a theory of investor behavior: An application to realization utility, *Journal of Finance* 69, 907-946.
- Glascher, Jan, Nathaniel Daw, Peter Dayan and John O'Doherty, 2010, States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning, *Neuron* 66, 585-595.
- Hare, Todd, John O'Doherty, Colin Camerer, and Antonio Rangel, 2008, Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors, *Journal of Neuroscience* 28, 5623-5630.
- Hare, Todd, Colin F. Camerer, and Antonio Rangel, 2009, Self-control in decision-making involves modulation of the vmPFC valuation system, *Science* 324, 646-648.

- Hsu, Ming, Meghana Bhatt, Ralph Adolphs, Daniel Tranel, and Colin F. Camerer, 2005, Neural systems responding to degrees of uncertainty in human decision-making, *Science* 310, 1680-1683.
- Ingersoll, Jonathan and Lawrence Jin, 2013, Realization utility with reference-dependent preferences, *Review of Financial Studies* 26, 723-767.
- Kable, Joseph W., and Paul W. Glimcher, 2007, The neural correlates of subjective value during intertemporal choice, *Nature Neuroscience* 10, 1625-1633.
- Knutson, Brian, Scott Rick, G. Elliott Wimmer, Drazen Prelec, and George Loewenstein, 2007, Neural predictors of purchases, *Neuron* 53, 147-156.
- Li, Jian, and Nathaniel Daw, 2011, Signals in human striatum are appropriate for policy update rather than value prediction, *Journal of Neuroscience*, 31, 5504-5511.
- Lin, Alice, Ralph Adolphs, and Antonio Rangel, 2012, Social and monetary reward learning engage overlapping neural substrates, *Social Cognitive and Affective Neuroscience* 7, 274-281.
- Lin, Chien-Huang, Wen-Hsien Huang, and Marcel Zeelenberg, 2006, Multiple reference points in investor regret, *Journal of Economic Psychology*, 27, 781-792.
- Lohrenz, Terry, Kevin McCabe, Colin F. Camerer, and P. Read Montague, 2007, Neural signature of fictive learning signals in a sequential investment task, *Proceedings of the National Academy of Sciences* 104, 9493-9498.
- Lohrenz, Terry, Meghana Bhatt, Nathan Apple, and P. Read Montague, 2013, Keeping up with Joneses: Interpersonal prediction errors and the correlation of behavior in a tandem sequential task, *PLoS computational biology* 9.
- Maldjian, J.A., P.J. Laurienti, J.B. Burdette, and R.A. Kraft, 2003, An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets, *Neuroimage* 19, 1233-1239.
- McClure, Samuel, Gregory Berns, and P. Read Montague, 2003, Temporal prediction errors in a passive learning task activate human striatum, *Neuron* 38, 339-346.
- O'Doherty, John, Peter Dayan, Karl Friston, Hugo Critchley, and Ray Dolan, 2003, Temporal difference models and reward-related learning in the human striatum, *Neuron*, 329-337.
- Odean, Terrance, 1998, Are investors reluctant to realize their losses?, *Journal of Finance* 53, 1775-1798.
- Padoa-Schioppa, Camillo, and John A. Assad, 2006, Neurons in the orbitofrontal cortex encode economic value, *Nature* 441, 223-226.
- Pessiglione, Mathias, Ben Seymour, Guillaume Flandin, Raymond Dolan, and Christopher Frith, 2006, Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans, *Nature* 442, 1042-1045.
- Plassmann, Hilke, John P. O'Doherty, and Antonio Rangel, 2010, Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making, *The Journal of Neuroscience* 30, 10799-10808.

- Romo, Ranulfo and Wolfram Schultz, 1990, Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements, *Journal of Neurophysiology* 63, 592-606.
- Rutledge, Robb, Mark Dean, Andrew Caplin, and Paul Glimcher, 2010, Testing the reward prediction error hypothesis with an axiomatic model, *Journal of Neuroscience* 30, 13525-13536.
- Schultz, Wolfram, Peter Dayan and P. Read Montague, 1997, A neural substrate of prediction and reward, *Science* 275, 1593-1599.
- Shefrin, Hersch, and Meir Statman, 1984, Explaining investor preference for cash dividends, *Journal of Financial Economics* 13, 253-282.
- Shefrin, Hersch, and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long: Theory and evidence, *Journal of Finance* 40, 777-790.
- Shiller, Robert, 2005, *Irrational Exuberance*.
- Smith, Alec, Terry Lohrenz, Justin King, P. Read Montague, and Colin Camerer, 2014, Irrational exuberance and neural crash warning signals during endogenous experimental market bubbles, *PNAS* 111, 10503-10508.
- Strahilevitz, Michal, Brad M. Barber, and Terrance Odean, 2011, Once burned, twice shy: How naive learning, counterfactuals, and regret affect the repurchase of stocks previously sold, *Journal of Marketing Research* 48, 102-120.
- Tzourio-Mazoyer, N., B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, B. Mazoyer, and M. Joliot, 2002, Automated anatomical labeling of activations in spm using a macroscopic anatomical parcellation of the mni mri single-subject brain, *NeuroImage* 15, 273-289.
- Weber, Martin, and Colin Camerer, 1998, The disposition effect in securities trading: an experimental analysis, *Journal of Economic Behavior & Organization* 33, 167-184.
- Weber, Martin, and Frank Welfens, 2011, The follow-on purchase and repurchase behavior of individual investors: An experimental investigation, *Die Betriebswirtschaft* 71, 139-154.
- Weber, Martin, and Frank Welfens, 2008, Splitting the Disposition Effect: Asymmetric Reactions Towards 'Selling Winners' and 'Holding Losers', *Working Paper*.

Table 1. Summary Statistics. The following table provides summary statistics on the repurchase effect, disaggregated by individual subject and over the course of the two experimental sessions. “Average Episode Length” is the average number of trials that elapse between selling a stock and repurchasing it. “Repurchase Rate” is the fraction of trials on which a subject repurchases a stock given an opportunity to do so. “PDR” and “PUR” are the proportion of repurchases as defined by equations (3) and (4) in the main text.

Subject ID	Overall Sample					Session 1				Session 2			
	Average Episode Length	Repurchase Rate	PDR	PUR	Repurchase Effect	Repurchase Rate	PDR	PUR	Repurchase Effect	Repurchase Rate	PDR	PUR	Repurchase Effect
1	1.74	55%	67%	42%	25%	61%	73%	47%	25%	49%	60%	35%	25%
2	2.90	33%	29%	42%	-13%	43%	31%	71%	-40%	27%	28%	25%	3%
3	4.64	15%	17%	12%	5%	19%	15%	25%	-10%	12%	19%	4%	15%
4	3.79	24%	58%	13%	45%	15%	44%	8%	36%	35%	70%	21%	49%
5	3.20	28%	38%	20%	18%	24%	39%	13%	27%	36%	38%	35%	2%
6	2.71	34%	20%	64%	-43%	39%	18%	63%	-45%	29%	22%	67%	-45%
7	5.22	10%	21%	5%	16%	6%	25%	0%	25%	14%	18%	12%	6%
8	5.09	11%	28%	1%	26%	8%	29%	3%	26%	13%	28%	0%	28%
9	1.56	64%	64%	65%	-1%	76%	88%	67%	21%	55%	50%	63%	-13%
10	3.67	19%	17%	23%	-6%	13%	12%	14%	-2%	27%	20%	54%	-33%
11	1.37	72%	74%	70%	4%	81%	83%	78%	6%	66%	67%	65%	1%
12	4.26	19%	33%	11%	22%	13%	21%	9%	13%	25%	41%	14%	27%
13	2.82	34%	26%	54%	-28%	33%	27%	41%	-14%	35%	25%	86%	-61%
14	2.81	34%	35%	33%	2%	48%	67%	33%	33%	24%	23%	33%	-11%
15	4.86	14%	11%	17%	-7%	15%	22%	11%	11%	13%	5%	40%	-35%
16	9.13	8%	8%	8%	0%	8%	8%	8%	0%	8%	8%	8%	0%
17	9.00	9%	7%	13%	-5%	8%	6%	14%	-9%	10%	9%	11%	-2%
18	2.59	37%	29%	59%	-31%	39%	32%	50%	-18%	36%	27%	75%	-48%
19	6.71	44%	31%	93%	-62%	57%	38%	88%	-49%	38%	28%	100%	-72%
20	2.18	30%	33%	26%	8%	38%	56%	23%	33%	24%	21%	28%	-7%
21	2.83	45%	40%	54%	-14%	53%	53%	53%	0%	38%	32%	56%	-23%
22	2.24	39%	23%	89%	-66%	42%	17%	85%	-67%	37%	27%	100%	-73%
23	2.45	17%	16%	18%	-1%	23%	10%	27%	-17%	13%	19%	7%	12%
24	5.40	30%	27%	33%	-6%	27%	35%	20%	15%	33%	21%	64%	-42%
25	2.85	9%	14%	4%	10%	11%	25%	6%	19%	8%	11%	0%	11%
26	5.89	23%	14%	38%	-23%	17%	12%	26%	-14%	29%	18%	50%	-33%
27	4.17	14%	16%	10%	6%	20%	19%	21%	-2%	9%	14%	0%	14%
28	3.93	40%	56%	27%	28%	40%	53%	29%	24%	40%	59%	26%	33%
Mean	3.93	29%	30%	34%	-3%	31%	34%	33%	1%	28%	29%	39%	-10%
Std. Dev.	1.98	17%	18%	26%	26%	21%	23%	26%	27%	15%	18%	30%	32%

Table 2. **Determinants of Propensity to Repurchase.** Dependent variable equals 1 if the subject repurchases and equals 0 if the subject does not repurchase (conditional on the opportunity to buy). Foregone Capital Gain is the difference between current price and last sale price. NEV is the expected future price change of the stock conditional on all previous information. Standard errors are clustered at the subject level for all model specifications. *, **, and *** denotes significance at 10%, 5%, 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	Logit	Logit	Linear Probability Model	Linear Probability Model	Hazard Ratio	Hazard Ratio
Foregone Capital Gain	-0.014 (-3.10)***	x	-0.002 (-2.53)***	x	0.994 (-1.48)	x
Price		-0.022 (4.04)***		-0.003 (4.48)***		.987 (3.11)***
Last Sale Price		0.005 (1.18)		-0.001 (-0.09)		1.002 (0.75)
NEV	3.267 (1.71)*	4.420 (2.25)**	0.534 (1.30)	0.670 (1.59)	5.94 (1.07)	17.82 (1.77)*
Constant	-2.930 (-2.49)**	1.068 (2.16)**	0.275 (237.92)***	0.6459 (7.70)***	x	x
Observations	2410	2410	2410	2410	2410	2410
R ² (pseudo R ²)	0.008	0.03	0.008	0.032	x	x
Subj. Fixed Effects	N	N	Y	Y	Y	Y

Figure 1. **Sample screens from a typical trial in the fMRI experiment.** For trials 10 to 108, subjects see a “price update” screen for two seconds, followed by a “trading” screen for which they have up to three seconds to enter a decision. The screens shown above are for a trial in which the subject does not own stock B. If the subject owned stock B, then he would be given the opportunity to sell the stock instead of the opportunity to purchase the stock. For trials 1 to 9, subjects see only the price update screen; this allows them to accumulate information about price changes before having to make any decisions.

Figure 2. **Measures of the repurchase effect (PDR-PUR) for each subject.** Each vertical column corresponds to a specific subject in our experiment. The dashed line indicates the level of the repurchase effect that an expected value trader would exhibit, namely -0.75. All but two subjects exhibit a repurchase effect greater than this benchmark level. The figure also shows there is significant heterogeneity in the size of the repurchase effects across subjects.

Figure 3. **vSt activity reflects prediction errors at the price update screen.** The figure presents estimation results from the following equation:

$$b_v(t) = \alpha + \beta_1^v(\Delta p_t)I_{\text{own update}}(t) + \beta_2^v(\Delta p_t)I_{\text{no own update}}(t) + \beta_3^v(p_t - p_s)I_{\text{repurchase opp}}(t) + \beta_4^v(p_t - c)I_{\text{sell opp}}(t) + \beta_5^v \text{controls} + \varepsilon(t)$$

Panel A: Results from hypothesis test that β_2^v is significantly different from zero. Voxels are colored according to the associated t -statistic as indicated on the color bar on the right. For illustration purposes, the significance threshold for the image is $p < 0.001$ with a 20-voxel extent threshold, but all statistics reported in the paper are also small volume corrected at $p < 0.05$ using familywise error (FWE). Panel B: Results from hypothesis test that β_1^v is significantly different from zero. Panel C: Yellow and orange voxels represent the vSt target region (68 voxels) whereas red and orange voxels represent the conjunction of voxels for which β_1^v is significantly positive and β_2^v is significantly negative. Panel D: Beta values averaged across the 68 voxels in the vSt target region.

Figure 4. **Correlation between brain activity at price update screen and measures of the repurchase effect.** Each data point in the figure represents a single subject. For each subject, we compute the average β_2^v from regression (5) across all voxels in the vSt target region. We find that, across subjects, the degree to which the vSt negatively correlates with a price update when a stock is not held is itself correlated with the repurchase effect.

Figure 5. **vmPFC reflects decision values under both the regret-devaluation mechanism and realization utility.** The figures present estimation results from equation (5). Panel A: Results from hypothesis test that β_3^v is significantly different from zero. Voxels are colored according to the associated t -statistic as indicated on the color bar on the right. For illustration purposes, the significance threshold for the image is $p < 0.005$ with a 20-voxel extent threshold, but all statistics reported in the paper are also small volume corrected at $p < 0.05$ using familywise error (FWE). Panel B: Results from hypothesis test that β_4^v is significantly different from zero. Panel C: Conjunction of voxels for which β_3^v is significantly negative and β_4^v is significantly positive (orange) within the vmPFC target region of 429 voxels (yellow and orange).

Figure 6. **Buying and selling decisions are highly correlated.** Each point in the figure represents a single subject. Panel A: The x-axis plots the disposition effect (selling behavior) and the y-axis plots the repurchase effect (buying behavior.) Panel B: The x-axis plots the proportion of stocks that are repurchased after going up since the previous sale (PUR) and the y-axis plots the proportion of loss realized (PLR). Panel C: The x-axis plots the proportion of stocks that are repurchased after going down since the previous sale (PDR) and the y-axis plots the proportion of gains realized (PGR).

Figure 7. **Suboptimal trading decisions are not fully driven by beliefs.** Results from an additional behavioral experiment designed to test whether subjects' beliefs could explain the repurchase effect and disposition effect. The control condition is identical to the fMRI experiment except we increase the incentives to follow the optimal strategy. The NEV treatment is identical to the control condition except we explicitly display the NEV on both the price update screen and trading screen. The figure shows that even after explicitly providing subjects with the NEV, which is a sufficient statistic to implement the optimal strategy, more than 30% of all decisions were still suboptimal.

Appendix I. Experimental Instructions

Buying your stock

In this experiment you will be given 350 experimental dollars to invest in three different stocks. Your job is to choose when to buy and sell each stock, so that you earn the most money by the end of the experiment. Throughout the experiment, you will see the price of each stock changing (more detail below), and you will use this information to decide when to buy and sell. When you sell a stock, you receive an amount of cash equal to the price of the stock. When you buy a stock, you receive one unit of the stock, but you must give up an amount of cash equal to the current price of the stock.

The three stocks you can buy or sell are simply called Stock A, Stock B, and Stock C. To begin the experiment you *MUST* buy all three stocks, where each stock costs \$100. Therefore, after you buy the three stocks, you will own one unit of each stock and have a total of \$50 remaining. For the remainder of the experiment, you are only allowed to hold a maximum of 1 unit of each stock, and you cannot hold negative units (no short selling.) However, you can carry a *negative cash balance* by buying a stock for more money than you have, but any negative cash balances will be deducted from your final earnings.

Structure of the market

In the experiment, you will see two types of screens, a *price update screen* and an *action screen*. In the price update screen, one stock will be randomly selected and you will be told if the selected stock price has gone up or down, and by how much. Note that you will only see an update for *one* stock at a time. You will not be asked to do anything during this screen, you will simply see information about the change in price.

Following the price update screen, another stock will be randomly chosen (it may be the same one you just saw) and you will be asked to take an action. If you currently hold a unit of the stock, you will be asked if you would like to sell the stock at the current price. If you do not currently own a unit of the stock, you will be asked if you would like to buy a unit at the current price.

The experiment will start out with 9 consecutive price update screens, and then you will have the opportunity to buy or sell after each subsequent price update screen.

How the stock prices change

Each stock changes price according to the exact same rule. Each stock is either in a good state or in a bad state. In the good state, the stock goes up with 55% chance, and it goes down with 45% chance. In the bad state, the stock goes down with 55% chance and it goes up with 45% chance.

Once it is determined whether the price will go up or down, the *size* of the change is always random, and will either be \$5, \$10, or \$15. For example, in the bad state, the stock will go down with 55% chance, and the amount it goes down by is \$5, \$10, or \$15 with equal chance. Similarly, the good stock will go up with 55% chance, and the amount it goes up by will either be \$5, \$10, or \$15.

The stocks will all randomly start in either the good state or bad state, and after each price update, there is a 20% chance the stock switches state.

Stock price changes

	<i>Good state</i>	<i>Bad state</i>
+	55%	45%
-	45%	55%

State changes

	<i>Good state today</i>	<i>Bad state today</i>
<i>Good state tomorrow</i>	80%	20%
<i>Bad state tomorrow</i>	20%	80%

Earnings and payout

You will play this market game **TWO SEPARATE TIMES** in the scanner. Each game will last approximately 15 minutes, and each game is independent from the previous one. This means when you start the second game, you will have to buy the three stocks at \$100 again, and the stocks will start randomly in each state again.

Your earnings at the end of the experiment will be equal to the amount of cash you accrued over the two scanning sessions from buying and selling stocks, plus the current price of any stocks that you own.

$$\text{Earnings} = \text{cash} + \text{price } A * (\text{Hold } A) + \text{Price } B * (\text{Hold } B) + \text{Price } C * (\text{Hold } C)$$

Finally, your earnings will be converted using an exchange rate of 12:1. That means we divide your earnings by 12, and pay you this amount plus the \$15 show up fee.

Button presses

During the Action screens, you will either be given the option to “Buy?” or “Sell?” depending on whether you hold the stock or not. The LEFT (blue) button indicates “YES”. And the RIGHT (yellow) button indicates “NO.” You have three seconds to enter your response, otherwise the computer will randomly select a response for you.

Appendix II. fMRI Data Collection and Analysis

In this section, we describe how the fMRI measures of neural activity were collected and analyzed. The goal of this section, which is primarily taken from Frydman et al. (2014), is to provide enough information to serve as a brief primer on the subject for readers who are unfamiliar with fMRI. For a more detailed discussion, see (Huettel, Song and McCarthy (2004); Ashby (2011); Poldrack, Mumford and Nichols (2011)).

A. fMRI Data Collection and Measurement

We collected measures of neural activity over the entire brain using BOLD-fMRI, which stands for blood-oxygenated level dependent functional magnetic resonance imaging. BOLD-fMRI measures changes in local magnetic fields that result from the local inflows of oxygenated hemoglobin and outflows of de-oxygenated hemoglobin that occur when neurons fire. In particular, fMRI provides measures of the BOLD response in small “neighborhoods” of brain tissue called *voxels*, and is thought to measure the sum of the total amount of neuronal firing into that voxel and the total amount of neuronal firing within the voxel.²⁰

One important complication is that the hemoglobin responses measured by BOLD-fMRI are slower than the associated neuronal responses. Specifically, although the bulk of the neuronal response takes place quickly, BOLD measurements are affected for up to 24 seconds thereafter. Panel A of Figure A1 provides a more detailed illustration of the nature of the BOLD response. It depicts the path of the BOLD signal in response to one (arbitrary) unit of neural activity of infinitesimal duration at time zero. The function plotted here is called the canonical hemodynamic response function (HRF). It is denoted by $h(\tau)$, where τ is the amount of time elapsed since the

²⁰ The neural activity measured by fMRI in a 1 mm³ cube (about the size of a grain of salt) represents the joint activity of between 5,000 to 40,000 neurons, depending on the area of the brain.

neural activity impulse, and has been shown to approximate well the pattern of BOLD responses for most subjects, brain areas, and tasks.

Fortunately, there is a standard way of dealing with the complication described in the previous paragraph. In particular, the BOLD response has been shown to combine linearly across multiple sources of neural activity (Boynton et al. (1996)). This property, along with knowledge of the specific functional form of the HRF, allows us to construct a mapping from predicted neural activity to predicted BOLD responses. Specifically, if the predicted level of neural activity at any particular time is given by $a(t)$, then the level of BOLD activity at any instant t is well approximated by

$$b(t) = \int_0^{\infty} h(u)a(t-u)du, \quad (A1)$$

which is the convolution between the HRF and the neural inputs. This integral has a straightforward interpretation: it is a lagged sum of all the BOLD responses triggered by previous neural activity. Panel B of Figure A1 illustrates the connection between neural activity and BOLD responses; it depicts a hypothetical path of neural activity (the solid line), together with the associated BOLD response (the dashed line).

During our experiment, we acquire two types of MRI data in a 3.0 Siemens Tesla Trio MRI scanner with an eight-channel phased array coil. First, we acquire BOLD-fMRI data while the subjects perform the experimental task. We use a voxel size of 3 mm^3 , and collect these data for the entire brain ($\sim 100,000$ voxels) every 2.75 seconds.²¹ We also acquire high-resolution

²¹ More precisely, we acquire gradient echo T2*-weighted echoplanar (EPI) images with BOLD contrast. To optimize functional sensitivity in the orbitofrontal cortex (OFC), a key region of interest, we acquire the images in an oblique orientation of 30° to the anterior commissure–posterior commissure line [Deichmann, 2003 #16]. Each volume of images has 45 axial slices. A total of 692 volumes were collected over two sessions. The imaging parameters are as follows: echo time, 30 ms; field of view, 192 mm; in-plane resolution and slice thickness, 3 mm; repetition time, 2.75 s.

anatomical scans that we use mainly for realigning the brains across subjects and for localizing the brain activity identified by our analyses.²²

B. fMRI Data Pre-processing

Before the BOLD data can be analyzed to test our hypotheses, they have to be converted into a usable format. This requires the following steps, which are fairly standard – see Huettel, Song, and McCarthy (2004), Ashby (2011), and Poldrack, Mumford, and Nichols (2011) – and which are implemented by way of a specialized but commonly-used software package called SPM5 (Wellcome Department of Imaging Neuroscience, Institute of Neurology, London, UK).

First, we correct for slice acquisition time within each voxel. This is necessary because the scanner does not collect data on all brain voxels simultaneously. This simple step, which involves a nonlinear interpolation, realigns the data across all voxels.

Second, we correct for head motion to ensure that the time series of BOLD measurements recorded at a specific spatial location within the scanner is always associated with the same brain location throughout the experiment.²³

Third, we realign the BOLD responses for each individual into a common neuroanatomical frame (the standard Montreal Neurological Institute EPI template). This step, called spatial normalization, is necessary because brains come in different shapes and sizes; as a result, a given spatial location maps to different brain regions in different subjects. Spatial

²² More precisely, we acquire high-resolution T1-weighted structural scans (1 x 1 x 1 mm) for each subject. These are coregistered with their mean EPI images and averaged across subjects to permit anatomical localization of the functional activations at the group level.

²³ BOLD measurements were corrected for head motion by aligning them to the first full brain scan and normalizing to the Montreal Neurological Institute's EPI template. This entails estimating a six-parameter model of head motion for each volume (three parameters for center movement, and three parameters for rotation), and then removing the effect of the motion using these parameters. For details, see (Friston et al. (1996)).

normalization involves a nonlinear reshaping of the brain to maximize the match with a target template. Although the transformed data are not perfectly aligned across subjects due to remaining neuroanatomical heterogeneity, the process is sufficiently accurate for the purposes of most studies. Furthermore, any imperfections in the realignment process introduce noise that reduces our ability to detect neural activity of interest.

Fourth, we also spatially smooth the BOLD data for each subject by making BOLD responses for each voxel a weighted sum of the responses in neighboring voxels, where the weights decrease with distance.²⁴ This step ensures that the error structure of the data conforms to the normality assumptions on the error structure of the regression models that we will use to test our hypotheses (Huettel et al. (2004); Poldrack et al. (2011)).

Finally, we remove low-frequency signals that are unlikely to be associated with neuronal responses to individual trials.²⁵ An example of such a signal is the effect of a continuous head movement over the course of the experiment that is not fully removed by the second correction step described above.

C. fMRI Main Data Analyses

The key goals of our analysis are to test if the region of the vSt that has been repeatedly shown to encode prediction errors is consistent with Predictions 2 and 3. To do this, we run statistical tests to see if there are areas within these regions of the brain, given by collections of spatially contiguous voxels called *clusters*, where the BOLD response reflects neural activity that implements the computations of interest (e.g., realization utility computations). This is complicated by the fact that, since every voxel contains thousands of neurons, the BOLD

²⁴ Spatial smoothing was performed using an 8 mm full-width half-maximum Gaussian kernel. Essentially, this step entails replacing every measurement at every voxel with a weighted sum of the measurements in a neighborhood centered on the voxel, using weights that are given by the Gaussian kernel.

²⁵ Specifically, we applied a high-pass temporal filter to the BOLD data with a cut-off of 128 seconds.

responses in a voxel can be driven by multiple signals. Fortunately, the linear properties of the BOLD signal allow the neural signals of interest to be identified using standard linear regression methods.

The general statistical procedure is straightforward, and will be familiar to most economists. The analysis begins by specifying two types of variables that might affect the BOLD response: target computations and additional controls. The target computations reflect the signals we are looking for (e.g., a realization utility signal at the time of selling a stock). They are specified by a time series $s_i(t)$ describing each signal of interest. For each of these signals, let $S_i(t)$ denote the time series that results from convolving the signal $s_i(t)$ with the HRF, as described above. The additional controls, denoted by $c_j(t)$, are other variables that might affect the BOLD time series (e.g., residual head movement or time trends). These are introduced to further clean up the noise in the BOLD signal, but are not explicitly used in any of our tests. The control variables are not convolved with the HRF because, while they affect the measured BOLD responses, they do not reflect neural activity which triggers a hemodynamic response.²⁶

The linearity of the BOLD signal implies that the level of BOLD activity $b^v(t)$ in any voxel v at time t should be given by

$$b^v(t) = \text{constant} + \sum_i \beta_i^v S_i(t) + \sum_j \alpha_j^v c_j(t) + \varepsilon(t), \quad (\text{A2})$$

where $\varepsilon(t)$ denotes AR(1) noise. This model is estimated independently in each of the voxels that fall within the relevant region of interest (the vSt). Our hypotheses can then be restated as tests about the coefficients of this regression model: signal i is said to be associated with activity in voxel v only if β_i^v is significantly different from zero.

²⁶ For example, linear trends are often included as controls because the scanner heats up with continuous operation, inducing a linear change in the measured BOLD responses.

Two additional considerations apply to most fMRI studies, including this one. First, we are interested in testing hypotheses about the distribution of the signal coefficients in the population of subjects, not hypotheses about individual subject coefficients. This would normally require estimating a mixed effects version of the linear model specified above, which, given the size of a typical fMRI dataset, would be computationally intensive. Fortunately, there is a shortcut that provides a good approximation to the full mixed effects analysis (Penny et al. (2006)). It involves estimating the parameters separately for each individual subject, averaging them across subjects, and then performing t -tests. This is the approach we follow here.

Second, since our tests are carried out in each of the voxels in the relevant regions of interest (68 for the vSt), there is a concern about false-positives. To address this problem, we correct for multiple comparisons within the relevant region of interest, a procedure known in the fMRI literature as a small volume correction (SVC). We report results as significant if they pass SVC correction at a level of $p < 0.05$.²⁷

As noted earlier, we conduct our tests in an area of the vSt that has been linked to the computation of prediction errors. Specifically, we construct a sphere with a 15 mm radius around the coordinates (MNI-space, $x = -15$, $y = 6$, $z = -12$) that were found to exhibit peak correlation with prediction errors in (Lin, Adolphs and Rangel (2012)), and then intersect this sphere with an anatomical mask of the vSt.

²⁷ Specifically, we report results as significant if voxels within the pre-specified region of interest pass $p < 0.005$ uncorrected with a 20-voxel extent threshold and if they pass SVC with a family-wise error rate of less than 0.05.

Figure A1. **BOLD measurements of neural activity.** Panel A: Because fMRI measures the blood-oxygenated level dependent (BOLD) response, and not neural activity itself, we need a mapping from neural activity to BOLD response to make inferences about changes in neural activity. This mapping is known as the canonical hemodynamic response function and is shown here as the result of one unit of instantaneous neural activity at time 0. Panel B: This figure shows the BOLD response that results from three sequential sources of neural activity. The BOLD response combines linearly across multiple sources of neural activity.

REFERENCES

1. S. Huettel, A. Song, G. McCarthy, *Functional Magnetic Resonance Imaging*. (Sinauer Associates, 2004).
2. F. G. Ashby, *Statistical Analysis of fMRI Data*. (The MIT Press, 2011).
3. R. A. Poldrack, J. Mumford, T. Nichols, *Handbook of Functional MRI Data Analysis*. (Cambridge University Press, 2011).
4. G. M. Boynton, S. A. Engel, G. H. Glover, D. J. Heeger, Linear Systems Analysis of Functional Magnetic Resonance Imaging in Human V1. *The Journal of Neuroscience* **16**, 4207-4221 (1996).
5. K. J. Friston, A. Holmes, J.-B. Polin, C. J. Price, C. D. Frith, Detecting activations in PET and fMRI: Levels of inference and power. *Neuroimage* **40**, 223-235 (1996).
6. W. Penny, K. Friston, J. Ashburner, S. Kiebel, T. Nichols, *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. (Academic Press, 2006).
7. A. Lin, R. Adolphs, A. Rangel, Social and monetary reward learning engage overlapping neural substrates. *Social Cognitive and Affective Neuroscience* **7**, 274-281 (2012).