Essays in Neurofinance

Thesis by

Cary D. Frydman

In Partial Fulfillment of the Requirements for the degree

of

Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY

Pasadena, California

2012

(Defended May 23, 2012)

ACKNOWLEDGEMENTS

Early in my second year of graduate school in 2008, I was taking a class in Machine Learning, where lectures were often interrupted by questions from an enthusiastic student on the opposite side of the room from where I sat. The student's questions were always interesting, frequent, provocative, and very loud. After class one day, I met this student walking back from our class to Baxter Hall where both of our offices were located. The student was Antonio Rangel, who became my thesis advisor shortly after he convinced me in forty-five seconds that he had an interesting project for me to work on. Every day since then, without fail, Antonio has pushed me and encouraged me to pursue my ideas and research in neurofinance. His commitment and devotion to developing my skills as a successful and thoughtful scientist are unparalleled, and without his contagious enthusiasm and undying support for neuroeconomics, this thesis would not be possible.

I owe an equal amount of gratitude to my three other advisors. Peter Bossaerts has always been available to talk about ideas in financial economics and I have learned an immense amount from him through working on a diverse set of papers on the topics of market microstructure (not included in this thesis), behavioral genetics, and neuroimaging. Colin Camerer has been an invaluable source of wisdom in learning behavioral and experimental economics, and his sharp wit about those who are skeptical of newer methodologies in the social sciences has made pursuing work in neurofinance a much easier task. Colin's encyclopedic knowledge of the literatures from nearly every discipline has taught me to always ask whether seemingly separate phenomena have a common foundation. Finally, I became interested in behavioral finance at the end of my second year of graduate school when I attended a summer school at Yale organized by Nick Barberis. It was there that I developed the main idea for this thesis, and I owe a great deal to Nick for his encouragement in applying methods from neuroscience to questions is behavioral finance. I also thank him for his willingness to learn neuroeconomics over the last few years and for his generosity in working very closely with me from afar on this thesis.

There are several others at Caltech who have been integral in developing the ideas in the pages that follow. I can't begin to list all of them, but an abbreviated list includes: Alec Smith, Vikram Chib, Ben Bushong, Jon Malmaud, Ian Krajbich, Todd Hare, Ben Gillen, Salvatore Nunnari, Luke Boosey, Mike Alvarez, Leeat Yariv, Pietro Ortoleva, Yutaka Kayaba, Marina Agranov, Guilherme de Freitas, Sang Mok Lee, John Clithero, Erie Boorman, Alison Harris, and Taisuke Imai. I also owe a great deal of thanks to Laurel Auchampaugh, Jenny Niese, and Tiffany Kim for always making sure research went smoothly on a daily basis.

Finally, I thank those who encouraged me to apply to graduate school, and who have been with me the whole way through: I thank my girlfriend, Katherine, who every day has been encouraging and supportive of my work, and has always made sure I wore shoes at Caltech. I owe an endless amount of thanks to my mother for her support and love throughout graduate school, and during the first phase of my life, Before Caltech. Finally, I thank my father and grandfather for instilling in me a deep sense of intellectual curiosity, which has truly made this thesis possible.

ABSTRACT

Economists have learned a great deal about investor behavior over the last two decades with the availability of large discount brokerage data sets. While this has given economists a better understanding of the trading patterns that characterize individual investor behavior, less success has been achieved in understanding what drives these trading patterns. Part of the difficulty in this endeavor is that it is sometimes difficult to test alternative theories of investor behavior using only data from the field. In particular, the two trading patterns we investigate in this thesis, the disposition effect and the repurchase effect, are unlikely driven by standard rational models of trading, and alternative theories of their causes are difficult to test using only data from the field, or data from behavioral laboratory experiments.

In order to better understand the causes of the disposition effect and the repurchase effect, we use *neural data*, data collected from functional magnetic resonance imaging (fMRI) along with trading data to construct empirical tests of different theories. Chapter 1 uses fMRI data to test a model of realization utility, which can readily predict a disposition effect. In our experiment, we find that subjects exhibit strong disposition effects, although they are suboptimal, and the neural data strongly supports the realization utility hypothesis. While Chapter 1 is concerned with the selling behavior, we focus on systematic violations of buying behavior in Chapter 2. We propose a model of regret to explain the repurchase effect in the buy-side trading data, for which we find strong support in the neural data. Chapters 3 and 4 study whether the suboptimal trading behavior we find in the first two chapters is stable, and we explore what the source of the heterogeneity is. Specifically, in Chapter 3 we find that exogenously manipulating the display of information on the trading screen can significantly reduce the size of the disposition effect. Chapter 4 uses an approach from behavioral genetics to identify candidate genes that can help explain the cross-sectional variation in choice behavior.

TABLE OF CONTENTS

INTRODUCTION

Economists have documented several robust empirical facts regarding the behavior of individual investors. Some of this behavior is puzzling from the standpoint of a purely rational model of trading, which has led researchers to propose alternative models of trading behavior based on different specifications of preferences and/or beliefs. In this thesis, we aim to provide financial economics with a new methodology that can be useful in testing these alternative theories of investor behavior. In particular, we use methods from cognitive neuroscience to collect data from the human brain that allows us to construct sharp empirical tests of theories of investor behavior. Chapter 1 shows how neural data collected via functional magnetic resonance imaging (fMRI) can be useful in understanding why investors exhibit a puzzling behavior known as the disposition effect. We design an experimental stock market and generate disposition effects in the laboratory that are strongly suboptimal. We then use a time series of neural data to test a specific theory of the disposition effect, called *realization utility,* for which we find strong support. One of the main contributions of Chapter 1 is to clearly illustrate the added value of neural data for our understanding of economic behavior.

While Chapter 1 is concerned with understanding systematic violations of the optimal trading strategy on the sell-side, in Chapter 2 we investigate the *buying behavior* of individual subjects. In this chapter, we seek to understand what drives the “repurchase effect”, which is the recently documented empirical fact that investors have a great tendency to repurchase stocks that have gone down in value since last sale more often than they repurchase stocks that have gone up in value since last sale. We propose that this effect is driven by a regret-devaluation mechanism, and we employ empirical tests of this mechanism that cannot be run without neural data. Our results provide support for the regret-devaluation mechanism as we find: 1) stronger regret signals predict a higher probability of a repurchase mistake and 2) average regret signals can explain a portion of the variation in the size of the repurchase effect across subjects. Like Chapter 1, this chapter highlights the added value of using neural data in testing theories of investor behavior.

After achieving a better understand of what drives systematic deviations from optimal trading behavior on both the sell-side and the buy-side, it is interesting to compare the two behaviors. Figure 4 in Chapter 2 shows that the disposition effect and repurchase effect are highly correlated across subjects and that there is substantial variation in the size of both effects. This suggests that there may be a stable and common underlying "factor" that drives buying and selling behavior, but that there is heterogeneity in this factor across the population of investors. We examine these two ideas in Chapter 3 and Chapter 4, respectively.

In particular, Chapter 3 examines whether the preference structure that drives the disposition effect is stable and fixed, or whether it is malleable and can be manipulated by exogenous changes in the display of information about an investor's portfolio. We find that by introducing cues on the trading screen that shift attention away from realization utility, we can manipulate the size of the disposition effect, and therefore, expected final wealth. We discuss our experimental results in this chapter in connection with a recently enacted US government legislation that effectively increases the saliency of the cost basis, and may have unintended consequences for investor behavior.

Chapter 4 is concerned with studying the underlying source of heterogeneity in suboptimal investor behavior that we document in the first two chapters. We use a different task structure in this chapter, where subjects are asked to make a series of choices between a risky lottery and a certain option. We then test whether genetic markers can help explain the variation in choice behavior across subjects. We find significant heterogeneity in choice behavior, and subjects with a specific genetic polymorphism, MAOA-L, tend to choose the risky option more often than those with MAOA-H. After estimating a *computational phenotype*, we find that MAOA-L does not affect behavior through preferences, but instead through a different channel, known in computational neuroscience as the choice comparator.

Chapter 1

Testing Theories of Investor Behavior Using Neural Data

Over the past twenty years, economists have accumulated a large amount of evidence on how individual investors manage their financial portfolios over time. Some of this evidence is puzzling, in the sense that it is hard to reconcile with the simplest models of rational trading (Barberis and Thaler (2003); Campbell (2006)). Theorists have responded to this challenge by constructing new models of investor behavior. Empiricists, in turn, have started testing these newly-developed models.

Most of the empirical work that tests theories of investor behavior uses field data (Barber and Odean (2000); Barber and Odean (2001); Choi et al. (2009); Grinblatt and Keloharju (2009)). A smaller set of studies uses data from laboratory experiments. The advantage of experimental data is that it gives researchers a large degree of control over the trading and information environment, which can make it easier to tease theories apart (Plott and Sunder (1988); Camerer and Weigelt (1991); Camerer and Weigelt (1993); Weber and Camerer (1998); Bossaerts and Plott (2004); Bossaerts et al. (2007)).

In this paper, we show that another kind of data, namely measures of *neural* activity taken using functional magnetic resonance imaging (fMRI) while subjects trade in an experimental stock market, can also be very useful in testing theories of investing behavior. In particular, we show that neural data can be used to test theories designed to explain the "disposition effect," the robust empirical fact that individual investors have a greater propensity to sell stocks trading at a gain relative to purchase price, rather than stocks trading at a loss[1].

The disposition effect has attracted considerable attention because it has proven challenging to explain using simple rational models of trading behavior. This impasse has motivated the development of multiple competing alternative theories, both rational and

---

[1] See for example, Shefrin and Statman (1985), Odean (1998), Genesove and Mayer (2001), Grinblatt and Keloharju (2001), Feng and Seasholes (2005), Frazzini (2006), Jin and Scherbina (2011).

[2] In this paper, we use the word "behavioral" in two different senses. Most of the time, as in the last sentence of this paragraph, we take it to mean "pertaining to behavior". Occasionally, we take it to mean "less than fully rational" or "psychological". It should be clear from the context which of the two meanings is intended.

behavioral (Shefrin and Statman (1985); Odean (1998); Barberis and Xiong (2009); Kaustia (2010)). One such theory, which is the focus of this paper, is the realization utility hypothesis (Shefrin and Statman (1985); Barberis and Xiong (2011)). According to this theory, in addition to deriving utility from consumption, investors also derive utility *directly* from *realizing* gains and losses on the sale of risky assets that they own. For example, if an investor realizes a gain (e.g., by buying a stock at $20 and selling it at $40), he receives a positive burst of utility proportional to the capital gain. In contrast, if he realizes a loss (e.g., by buying a stock at $20 and selling it at $10), he receives a negative burst of utility proportional to the size of the realized loss. The presence of realization utility is important because, in combination with a sufficiently high time discount rate, it leads investors to exhibit a disposition effect (Barberis and Xiong (2011)).

Testing among competing theories of phenomena like the disposition effect using field or experimental data is difficult because these theories often make similar predictions about behavior (Weber and Camerer (1998) is an exception). Furthermore, it is extremely difficult, using such data alone, to carry out direct tests of the mechanisms driving behavior (e.g., of whether or not people actually receive bursts of utility proportional to realized capital gains). On the other hand, a combination of neural measurement and careful experimental design allows for direct tests of the extent to which the computations made by the brain at the time of decision-making are consistent with the mechanisms posited by different models.

In this paper, we describe the results of an fMRI experiment designed to test the hypothesis that subjects experience realization utility while trading in an experimental stock market, and that this is associated with trading patterns consistent with the disposition effect. The experiment allows us to test several behavioral and neural predictions of the realization utility hypothesis.[2]

[2] In this paper, we use the word "behavioral" in two different senses. Most of the time, as in the last sentence of this paragraph, we take it to mean "pertaining to behavior". Occasionally, we take it to mean "less than fully rational" or "psychological". It should be clear from the context which of the two meanings is intended.

Behaviorally, we find that the average subject in our experiment exhibits a strong and significant disposition effect. This stands in sharp contrast to the prediction of a simple rational trading model in which subjects maximize the expected value of final earnings. In particular, our experimental design induces positive short-term autocorrelation in stock price changes, which implies that a risk-neutral rational trader would sell losing stocks more often than winning stocks, thereby exhibiting the opposite of the disposition effect. In contrast, the strong disposition effect displayed by our subjects is consistent with the existence of realization utility effects.

When taken literally as a model of the decision-making process, the realization utility model also makes several clear predictions about the pattern of neural activity that should be observed at different times in the experiment. We describe these predictions in detail in the main body of the paper, but summarize them briefly here.

First, the realization utility model predicts that, at the moment when a subject is making a decision as to whether to sell a stock, neural activity in areas of the brain that are known to encode the value of potential actions should be proportional to the capital gain that would be realized by the trade (i.e. to the difference between the sale price and the purchase price). This prediction follows from the fact that, for an individual who experiences realization utility, the value of selling a stock depends on the associated capital gain or loss. Brain regions that have been widely shown to correlate with the value of potential actions include the ventromedial prefrontal cortex (vmPFC) and the ventral striatum (vSt)[3].

Second, the realization utility model predicts that, across individuals, the strength of the disposition effect should be correlated with the strength of the realization utility signal in decision value areas such as the vmPFC or the vSt. This follows from the fact that a subject who is strongly influenced by realization utility should exhibit both a strong disposition effect *and* neural activity in decision value areas that is highly responsive to the associated capital gain.

[3] See Hsu et al. (2005), Kable and Glimcher (2007), Knutson et al. (2007), Hare et al. (2008), Kennerley et al. (2008), Chib et al. (2009), Hare et al. (2009), Hsu et al. (2009), Kang et al. (2009), Hare et al. (2010), Levy et al. (2010), Litt et al. (2010), Kang et al. (2011).

Third, the realization utility hypothesis predicts that neural activity in areas that have been associated with the encoding of experienced utility (sometimes called "instantaneous hedonics") should increase at the moment that a subject decides to realize a capital gain. Previous research in behavioral neuroscience has shown that activity in regions of the vmPFC and the vSt also correlates with the reported level of instantaneous experienced utility[4]. This prediction is particularly interesting because it provides the most direct test of the realization utility hypothesis, and thus best illustrates the value of neural data for testing theories of financial decision-making.

Our fMRI measurements reveal patterns of neural activity that are consistent with the three neural predictions. This provides novel and strong support for the mechanisms at work in the realization utility model, and to our knowledge, provides the first example of how neural evidence can be used to test economic models of financial decision-making. We emphasize that the results do not imply that realization utility provides a complete description of the forces driving investor behavior, even in the context of our experiment. However, the fact that activity in the decision-making circuitry corresponds to some of the computations hypothesized by the realization utility model provides novel evidence that realization utility plays a significant role in the decisions made by our experimental subjects. It further suggests that mechanisms of this kind might also be at work in the real-world transactions of individual investors.

Using neural data to test an economic model is an unusual exercise in the field of economics because a common view in the profession is that models make as-if predictions about behavior, and are not to be taken as literal descriptions of how decisions are actually made (Gul and Pesendorfer (2008); Bernheim (2009)). In contrast to this view, we adopt a neuroeconomic approach which is based on the idea that knowledge about the computational processes that the brain uses to make decisions should be of central interest to economists because, since these processes describe the actual determinants of observed behavior, they provide valuable insights

[4] See Blood and Zatorre (2001), De Araujo et al. (2003), Kringelbach et al. (2003), Rolls et al. (2003), Small et al. (2003), McClure et al. (2004), Plassmann et al. (2008).

into the drivers of economic behavior (Camerer et al. (2005); Camerer (2007); Fehr and Rangel (2011)).

Our study contributes to the nascent field of neurofinance, which seeks to characterize the computations undertaken by the brain to make financial decisions, and to understand how these computations map to behavior. Several early contributions are worth highlighting. Lo and Repin (2002) investigated the extent to which professional experience affects the emotional arousal of traders in stressful situations, where arousal was measured using skin conductance responses and changes in blood pressure. Kuhnen and Knutson (2005) measured neural responses using fMRI during a simple investment task and found that activity in brain regions previously associated with emotional processing, such as the nucleus accumbens and the insula, predicted subjects' subsequent willingness to take risks. Knutson et al. (2008) took these ideas further by showing that exogenous emotional cues (e.g., erotic pictures) could be used to affect investment behavior, and that these cues increased activity in the same areas that they identified in their previous study. More recently, Bruguier et al. (2010) have shown that neural fMRI measurements of the extent to which subjects activate brain areas associated with concrete cognitive skills, such as the ability to predict others' state of mind, might be useful in identifying which subjects would be successful traders.

Our paper contributes to this literature by showing, for the first time, that a combination of fMRI neural measurements and careful experimental design can be used to test the validity of specific economic theories of financial decision making. Our work also contributes more broadly to the rapidly growing field of neuroeconomics, which seeks to characterize the computations made by the brain in different types of decisions, ranging from simple choices to choices involving risk, self-control and complex social interactions. For recent reviews, see Fehr and Camerer (2007), Glimcher et al. (2008), Rangel et al. (2008), Bossaerts (2009), Kable and Glimcher (2009), Rangel and Hare (2010), and Fehr and Rangel (2011).

The paper is organized as follows. Section I presents some background information about the disposition effect and realization utility. Section II describes the experimental design and the predictions of the realization utility hypothesis. Section III provides a detailed description of how the neural predictions can be tested using fMRI. Section IV describes the results. Section V briefly concludes.

I. Background: The Disposition Effect and the Realization Utility Model

Using an argument based on Kahneman and Tversky's (1979) prospect theory, Shefrin and Statman (1985) predict that individual investors will have a greater propensity to sell stocks trading at a gain relative to purchase price, rather than stocks trading at a loss. They label this the "disposition effect" and provide some evidence for it using records of investor trading. More detailed evidence for the effect can be found in Odean (1998), who analyzes the trading activity, from 1987 to 1993, of 10,000 households with accounts at a large discount brokerage firm. The phenomenon has now been replicated in several other large databases of trading behavior.

It will be useful to explain Odean's (1998) methodology in more detail because we will adopt a similar methodology in our own analysis. For any day on which an investor in Odean's (1998) sample sells shares of a stock, each stock in his portfolio on that day is placed into one of four categories. A stock is counted as a "realized gain" ("realized loss") if it is sold on that day at a price that is higher (lower) than the average price at which the investor purchased the shares. A stock is counted as a "paper gain" ("paper loss") if its price is higher (lower) than its average purchase price, but it is *not* sold on that day. From the total number of realized gains and paper gains across all accounts over the entire sample, Odean (1998) computes the Proportion of Gains Realized (PGR):

$$PGR = \frac{\text{\# of realized gains}}{\text{\# of realized gains} + \text{\# of paper gains}}$$

In words, PGR computes the number of gains that were realized as a fraction of the total number of gains that could have been realized. A similar ratio, PLR, is computed for losses:

$$PLR = \frac{\text{\# of realized losses}}{\text{\# of realized losses} + \text{\# of paper losses}}$$

The disposition effect is the empirical fact that PGR is significantly greater than PLR. Odean (1998) reports PGR = 0.148 and PLR = 0.098.

While the disposition effect is a robust empirical phenomenon, its causes remain unclear. This is due, in large part, to the fact that standard rational models of trading have had trouble capturing important features of the data. Consider, for example, an information model in which investors sell stocks with paper gains because they have private information that these stocks will subsequently do poorly, and hold on to stocks with paper losses because they have private information that these stocks will rebound. This hypothesis is inconsistent with Odean's finding that the average return of the prior winners sold by investors is 3.4% *higher*, over the next year, than the average return of the prior losers they hold on to. Another natural model involves taking into account the favorable treatment of losses by the tax code. However, this model also fails to explain the disposition effect because tax-loss selling predicts a greater propensity to sell stocks associated with paper *losses*. Another model attributes the disposition effect to portfolio rebalancing of the kind predicted by a standard framework with power utility preferences and i.i.d. returns. However, under this hypothesis, rebalancing is the "smart" thing to do, which implies that we should observe a stronger disposition effect for more sophisticated investors. In contrast to this prediction, it is *less* sophisticated investors who exhibit a stronger disposition effect (Dhar and Zhu (2006)).

Early on, researchers proposed behavioral economics models of the disposition effect, which can potentially explain the stylized facts that the rational explanations just described cannot explain. One popular model assumes that investors have an irrational belief in mean-reversion (Odean (1998); Weber and Camerer (1998); Kaustia (2010)). If investors *believe* that stocks that have recently done well will subsequently do poorly, and that stocks that have recently done poorly will subsequently do well, their optimal trading strategy would lead to a disposition effect. We label such beliefs "irrational" because they are at odds with Odean's (1998) finding that the winner stocks investors sell subsequently do well, not poorly. While the mean-reversion hypothesis is appealing for its simplicity, and is consistent with some evidence from psychology on how people form beliefs[5], some studies cast doubt on its empirical validity. Weber and Camerer (1998) ask subjects to trade stocks in an experimental stock market, and find that they exhibit a disposition effect in their trading. In order to test the mean-reversion hypothesis, they add a condition in which subjects' holdings are exogenously liquidated at full value at random times, after which subjects are asked to reinvest the proceeds across stocks in any way they like. Note that if subjects are holding on to stocks with paper losses because of a belief in mean-reversion, we would expect them to re-establish their positions in these stocks, but in fact, they do not.[6]

Another popular behavioral economics model posits that the disposition effect results from prospect theoretic preferences (Kahneman and Tversky (1979)). Prospect theory is a prominent theory of decision-making under risk which assumes that individuals make decisions by computing the utility of potential gains and losses measured relative to a reference point that is often assumed to be the status quo, and that utility is concave over gains and convex over losses. At first sight, it appears that prospect theory preferences may be helpful for understanding the disposition effect. If an investor is holding a stock that has risen in value, he may think of it as trading at a gain. Moreover, if the concavity of the value function over gains induces risk aversion, this may lead him to sell the stock. Conversely, if the convexity of the value function over losses induces risk-seeking, he may be inclined to hold on to a stock that has dropped in value. Contrary to this intuition, Barberis and Xiong (2009) have recently shown that it is

---

[5] For a review, see Rabin (2002).

[6] Odean (1998) and Kaustia (2010) provide additional evidence that is inconsistent with the mean-reversion hypothesis.

surprisingly difficult to derive behavior consistent with the disposition effect using this model. In fact, they show that an investor who derives prospect theory utility from the annual trading profit on each stock that he owns will often exhibit the *opposite* of the disposition effect. Further theoretical arguments against this model have been provided by Kaustia (2010), who has shown that it predicts that investors' propensity to sell a stock depends on the magnitude of the embedded paper gain in a way that is inconsistent with the empirical evidence.

Another behavioral model of the disposition effect is based on the realization utility hypothesis (Shefrin and Statman (1985); Barberis and Xiong (2011)). The central assumption of this model is that investors derive *direct* utility from realizing capital gains and losses on risky assets that they own: they experience a positive burst of utility when they sell an asset at a gain relative to purchase price, where the amount of utility depends on the size of the realized gain; and a negative burst when they sell an asset at a loss relative to purchase price, where the amount of disutility again depends on the size of the loss realized. Importantly, this hypothesis states that trades have a direct utility impact on investors, not just an indirect one through their effect on lifetime wealth and consumption.[7] Barberis and Xiong (2011) show that linear realization utility, combined with a sufficiently high time discount rate, leads to a disposition effect. The intuition is simple. If an investor derives pleasure from realizing capital gains and, moreover, is impatient, he will be very keen to sell stocks at a gain. Conversely, if he finds it painful to sell stocks at a

---

[7] Barberis and Xiong (2011) speculate that realization utility might arise because of the way people think about their investing history. Under this view, some investors – in particular, less sophisticated investors -- do not think about their investing history in terms of overall portfolio return, but rather as a series of investing "episodes," each of which is characterized by three things: the identity of the asset, the purchase price, and the sale price. "I bought GE at $40 and sold it at $70" might be one such episode, for example. According to this view, an investor who sells a stock at a gain feels a burst of positive utility right then because, through the act of selling, he is creating a positive new investing episode. Similarly, if he sells a stock at a loss, he experiences a burst of disutility: by selling, he is creating a negative investing episode.

capital loss and also discounts future utility at a high rate, he will delay selling losing stocks for as long as possible.[8]

While the realization utility hypothesis makes predictions about behavior that are consistent with the disposition effect, as well as with other empirical patterns [9], it is based on assumptions that depart significantly from those of traditional models. In particular, its predictions rely on the assumption that utility depends not only on consumption, but also on capital gains and losses realized from the sale of specific assets. Given the unusual nature of this assumption, it seems especially important to carry out direct tests of the extent to which the hypothesized source of utility is actually computed by subjects and affects their decisions. In the rest of the paper we show how this can be done using a combination of fMRI measures of neural activity and careful experimental design.

II. Experimental Design and Predictions

In this section, we first describe the experimental stock market that we set up to test the realization utility model. We then lay out the specific behavioral and neural predictions of the theory that we test.

A. Design

The design of the experimental stock market builds directly on an earlier non-neural experiment conducted by Weber and Camerer (1998).

---

8 Time discounting is not a critical part of the realization utility hypothesis. The disposition effect also follows from realization utility combined with an S-shaped value function, as in prospect theory (Barberis and Xiong, 2009). Adopting this interpretation of the realization utility hypothesis would not significantly affect the analysis that follows.

9 Barberis and Xiong (2011) show that realization utility can shed light on many empirical phenomena, not just on the disposition effect. Some of the other applications they discuss are the poor trading performance of individual investors, the greater volume of trading in bull markets than in bear markets, the individual investor preference for volatile stocks, and the low average return of volatile stocks.

Subjects are given the opportunity to trade three stocks – stock A, stock B, and stock C – in an experimental market. The experiment consists of two identical sessions separated by a one-minute break. Each session lasts approximately 16 minutes and consists of 108 trials. We use *t* to index the trials within a session.[10]

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 $t > 9$ consists of two parts: a price update and a trading decision, each of which corresponds to a separate screen 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 value of buying or selling a stock, but do not need to update their beliefs about the price process 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

[10] We split our experiment into two sessions in order to avoid running the fMRI machine for too long without a break, as this could lead to potential medical risks for the subjects.

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 time. 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 buy 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 two-state Markov chain with a good state and a bad state. The Markov chain for each stock is independent of the Markov chains for the other two stocks. Suppose that, in trial *t*, there is a price update for stock *i*. If stock *i* is in the good state at that time, its price increases with probability 0.55 and decreases with probability 0.45. Conversely, if it is in the bad state at that time, 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 a state 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 this trial remains the same as in trial *t-1* with probability 0.8, but switches with probability 0.2. In mathematical terms, if $s_{i,t} \in \{\text{good}, \text{bad}\}$ 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 never revealed to the subjects: they have to infer them from the observed price paths. To ease comparison of trading performance across subjects, the same set of realized prices is used for all subjects.

A key aspect of our design is that, conditional on the information available to subjects, each of the stocks exhibits positive short-term autocorrelation in its price changes. If a stock performed well 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 positive.

At the end of each session, we liquidate subjects' holdings of the three stocks and record the cash value of their position. We give subjects a financial incentive to maximize the final value of their portfolio at the end of each session. Specifically, if the total value of a subject's cash and risky asset holdings at the end of session 1 is \$X, in experimental currency, and the total value of his cash and risky asset holdings at the end of session 2 is \$Y, again in experimental currency, then his take-home pay in *actual* dollars is 15 + (X+Y)/24.[11] Subjects' 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.

[11] In other words, we average X and Y to get (X+Y)/2, convert the experimental currency to actual dollars using a 12:1 exchange rate, and add a \$15 show-up fee.

N=28 Caltech subjects participated in the experiment (22 male, age range 18 – 60).[12] 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.

Finally, note that, in our experiment, there is a straightforward way to measure the extent to which a subject exhibits a disposition effect in his trading. We simply adapt Odean's (1998) methodology, described in Section I, in the following way. Every time a subject faces a decision about selling a stock, we classify his eventual action as a paper gain (loss) if the stock's current price is above (below) the purchase price and he chooses not to sell; and as a realized gain (loss) if the stock's current price is above (below) the purchase price and he chooses to sell. We then count up the number of paper gains, paper losses, realized gains, and realized losses over all selling decisions faced by the subject and compute the PGR and PLR measures described earlier. We assign the subject a disposition effect measure of PGR-PLR. When this measure is positive (negative), the subject exhibits (the opposite of) a disposition effect.

B. Optimal trading strategy

We now characterize the optimal trading strategy for a risk-neutral Bayesian investor who is maximizing the expected value of his take-home earnings – from now on, we refer to such an investor as an "expected value" investor. 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.

[12] 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, thereby interfering with the reliability of the neural measurements.

Formally, let $p_{i,t}$ be the price of stock *i* in trial *t*, after any price update about the stock, and let $q_{i,t} = \Pr(s_{i,t} = good \mid p_{i,t}, p_{i,t-1}, \ldots, p_{i,1})$ 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:

(1)

$$q_{i,t}(q_{i,t-1}, z_t) = \frac{\Pr(z_t \mid s_{i,t} = good) * \Pr(s_{i,t} = good \mid q_{i,t-1})}{q_{i,t-1}\Pr(z_t \mid s_{i,t-1} = good) + (1 - q_{i,t-1})\Pr(z_t \mid s_{i,t-1} = bad)}$$

$$= \frac{(0.5 + 0.05z_t) * [0.8 * q_{i,t-1} + 0.2 * (1 - q_{i,t-1})]}{q_{i,t-1}[0.8 * (0.5 + 0.05z_t) + 0.2 * (0.5 - 0.05z_t)] + (1 - q_{i,t-1})[0.2 * (0.5 + 0.05z_t) + 0.8 * (0.5 - 0.05z_t)]}$$

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.

Note that a trader who follows the optimal strategy described above will exhibit the opposite of the disposition 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 selling winner stocks relatively rarely, and losing stocks more often, thereby generating the reverse of the disposition effect.

Of course, it is difficult for subjects to do the exact calculation in equation (1) in real time during the experiment. However, it is relatively straightforward for subjects to approximate the optimal strategy: they need simply keep track of each stock's most recent price changes, and then hold on to stocks that have recently performed well while selling stocks that have recently performed poorly. The fact that a stock's purchase price is reported on the trading screen makes it particularly easy to follow an approximate strategy of this kind: subjects can simply use the difference between the current market price and the purchase price as a proxy for the stock's recent performance.[13]

C. Behavioral and neural predictions of the realization utility model

We now lay out the behavioral and neural predictions of the realization utility model, and contrast them with the predictions made by the expected value agent model.

Consider the behavioral predictions first. Since the experimental stock prices exhibit short-term momentum, an expected value investor will exhibit the opposite of the disposition effect: for the actual price process that our subjects see, the value of the PGR-PLR measure of the disposition effect under the optimal trading strategy for an expected value investor is -0.76. In other words, such an investor will have a much greater propensity to realize losses than to realize gains. By contrast, a trader who experiences bursts of realization utility and who discounts future utility at a high rate will sell winner stocks more often than the expected value trader and loser stocks less often. After all, he is keen to realize capital gains as soon as possible and to postpone realizing capital losses as long as possible. This leads to our first prediction.

<u>Prediction 1 (Behavioral):</u> For an expected value investor, the value of the PGR-PLR measure is given by -0.76. On the other hand, for a subject who experiences bursts of realization utility, the value of PGR-PLR is greater than -0.76.

We now turn to the neural predictions made by the two models. As noted earlier, a maintained assumption here is that the theories are not only making predictions about behavior,

[13] Our rational benchmark assumes risk-neutrality because the monetary risk in our experiment is small. We have also considered the case of risk aversion, however, and have concluded that its predictions do not differ significantly from those of risk neutrality. In some frameworks, risk aversion can generate a disposition effect through rebalancing motives. This is not the case in our experiment, however, because the volatility of stock price changes is independent of the level of the stock price. Furthermore, any rebalancing motives would be of second-order importance relative to time variation in the *mean* stock return.

but are also describing the key computations that subjects have to undertake in order to make decisions.

The first two neural predictions build on a basic finding from the field of decision neuroscience. A sizable number of studies have found evidence consistent with the idea that in simple choice situations the brain makes decisions by assigning values (often called "decision values") to the options under consideration, and then comparing them to make a choice[14]. These value signals are thought to reflect the relative value of taking the action or option under consideration (e.g., sell a stock) versus staying with the status quo (e.g., don't sell it) (De Martino et al. (2006); Hutcherson et al. (2011)). A significant body of work, using various neural measurement techniques, has shown that activity in regions of the ventromedial prefrontal cortex (vmPFC), and often also the ventral striatum (vSt), correlates with the decision values of options across a range of choices. For example, a recent study shows that, when subjects have to make purchasing decisions for goods such as monetary lotteries, foods, or DVDs, activity in the vmPFC correlates with behavioral measures of their willingness to pay (their "decision value") taken prior to the choice task (Chib et al. (2009)). See Rangel and Hare (2010) for an up-to-date review of the evidence.

Now consider the decision value signals that would be computed at the time of making a selling decision by an individual who makes choices according to the expected value model. In the context of our experiment, the decision value of selling a stock is given by the value of selling the stock minus the value of holding it. For the expected value investor, the value of selling the stock is zero: if he sells it, he will no longer own any shares of it, and so it can no longer generate any value for him. In contrast, the value of holding the stock can be approximated by the stock's expected price change on its next price update:

---

[14] See, for example, Hsu et al. (2005), Padoa-Schioppa and Assad (2006), Kable and Glimcher (2007), Knutson et al. (2007), Tom et al. (2007), Hare et al. (2008), Kennerley et al. (2008), Chib et al. (2009), Hare et al. (2009), Hsu et al. (2009), Hare et al. (2010), Levy et al. (2010), Litt et al. (2010), Rangel and Hare (2010).

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

It follows that the decision value signal at the time of making a selling decision is given by *0-0.6(*$2q_{i,t}$*-1),* or *0.6(1-*$2q_{i,t}$*)*; we will refer to this quantity throughout the paper as the *net expected value* of selling, or NEV. Note that this is only an approximation because the exact value of holding a stock is the stock's expected *cumulative* price change until the subject decides to sell it. However, this approximation has little effect on our later results because the value of holding a stock until its next price change is highly correlated with the value of holding the stock until it is actually optimal to sell it (the latter quantity can be computed by simulation).

Now consider the decision value signal that would be computed at the time of making a selling decision by an individual who makes choices according to the realization utility model. In particular, consider a simple form of the model in which subjects maximize the sum of expected discounted realized capital gains and losses. For such a trader, the value of selling is linearly proportional to the capital gain or loss, given by $p_{i,t} - c$, where *c* is the purchase price, or cost basis. However, the expected impact of holding the stock on realization utility is approximately zero, as long as the discount rate is sufficiently high. Thus, for such an investor, the decision value of selling should be linearly related to $p_{i,t} - c$.[15] This, together with the fact that decision

---

[15] We say that the value of holding a stock is "approximately" zero for a realization utility investor because, in principle, there is some value to holding, namely expected future realization utility flows. However, under the realization utility hypothesis, the trader is essentially myopic – he discounts future utility flows at a high rate. To a first approximation, then, the value of holding is zero. It may initially seem surprising that a subject would discount future utility at a high rate in the context of a 30-minute experiment. However, the literature on hyperbolic discounting suggests that discounting can be steep even over short intervals, perhaps because people distinguish sharply between rewards available right now, and rewards available at all future times. Furthermore, what may be important in our experiment is not so much calendar time, as transaction time. A subject who can trade stock B now may view the opportunity to trade it in the future as a very distant event -- one that is potentially dozens of screens away – and hence one that he discounts heavily. Finally, we note that discounting is not a critical part of our hypothesis. The disposition effect also follows from a model that combines realization utility with an S-shaped utility function, as in prospect theory. To a first approximation, this model would produce the same decision value as the discounting-based model. The reason is that, under an S-shaped utility function, the utility of selling a stock at a gain (loss) immediately is significantly higher (lower) than the expected utility of holding on to it.

value signals have been found to be reliably encoded in the vmPFC and the vSt, leads to the next prediction.

Prediction 2 (Neural): For expected value traders, activity in the regions of the vmPFC and the vSt associated with the computation of decision values should be linearly proportional to the NEV, $0.6(1-2q_{i,t})$, at the time of making selling decisions, and thus independent of the cost basis. In contrast, for subjects who experience realization utility proportional to realized capital gains and losses, activity in these areas of the vmPFC and the vSt should be linearly related to the realizable gain or loss, $p_{i,t} - c$.

The previous arguments predict that traders who place a large weight on realization utility when making decisions should exhibit neural activity in the vmPFC and the vSt that is more strongly correlated with the realizable capital gains or losses. At the same time, subjects with a larger weight on realization utility when making decisions should exhibit a stronger disposition effect. It follows that the degree to which vmPFC and vSt activity correlates with the realizable capital gain should be correlated, across subjects, with the strength of the disposition effect in their trading.

Prediction 3 (Neural): The degree to which vmPFC and vSt activity correlates with the realizable capital gain should be correlated, across subjects, with the strength of the disposition effect in their trading.

The final neural prediction is qualitatively different, in that it seeks to test *directly* if the subject experiences a burst of realization utility at the time of selling a stock that is proportional to the realized capital gain. As before, we can test this prediction using fMRI by building on previous work in neuroscience which has shown that activity in regions of the vSt and the vmPFC correlates reliably with reports of subjective pleasure generated by a wide variety of stimuli – including music, paintings, attractive faces, food, and wine.[16] It follows that, if realizing a capital gain generates a positive burst of experienced utility for the investor, it should increase neural activity in these areas precisely at the moment that the decision is made.

Prediction 4 (Neural): Under the realization utility hypothesis, neural activity in areas known to encode instantaneous experienced utility, such as the vSt or the vmPFC, should increase at the precise moment that individuals decide to realize a capital gain, and decrease at the moment they decide to realize a capital loss.

III. fMRI data collection and analysis

In this section, we provide a primer on how fMRI measures of neural activity are collected and analyzed. For more details, see Huettel et al. (2004), Ashby (2011), and Poldrack et al. (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 local inflows of oxygenated hemoglobin and outflows of de-oxygenated hemoglobin that occur when neurons fire. fMRI provides measures of the BOLD response of relatively small "neighborhoods" of brain tissue

[16] See, for example, Blood and Zatorre (2001), De Araujo et al. (2003), Kringelbach et al. (2003), Rolls et al. (2003), Small et al. (2003), McClure et al. (2004), Plassmann et al. (2008)

known as *voxels*, and is thought to measure the sum of the total amount of neural firing into that voxel as well as the amount of neuronal firing within the voxel. [17]

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, subsequent BOLD measurements are affected for up to 24 seconds. Figure 2A provides a more detailed illustration of the nature of the BOLD response. In particular, it shows the path of the BOLD signal in response to 1 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 elapsed time since the neural activity impulse, and has been shown to approximate well the pattern of BOLD responses for most subjects, brain areas, and tasks.

Fortunately, the BOLD response has been shown to combine linearly across multiple sources of neural activity (Boynton et al. (1996)). This property, along with a specific functional form of the HRF, allows us to construct a mapping from neural activity to BOLD response so that we can control for BOLD responses that are generated by neural activity over the previous 24 seconds. In particular, if the 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

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

which is the convolution between the HRF and the neural inputs. The integral can be interpreted in a straightforward way: it is simply a lagged sum of all the BOLD responses triggered by previous neural activity. This is illustrated in Fig. 2B, which depicts a hypothetical path of neural activity, together with the associated BOLD response.

[17] Note that the neural activity measured by fMRI in a 1-mm^3 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.

We acquire two types of MRI data during the experiment 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 with a voxel size of 3 mm^3. We acquire data for the entire brain (~ 100,000 voxels) every 2.75 seconds.[18] We also acquire high-resolution anatomical scans that we use mainly for realigning the brains across subjects and for localizing the brain activity identified by our analyses.[19]

B. fMRI data pre-processing

Before the BOLD data can be analyzed to test our hypotheses, it has to be converted into a usable format. This requires the following steps, which are fairly standard – see Huettel et al. (2004), Ashby (2011), & Poldrack et al. (2011) – and were implemented using a specialized but commonly used software package called SPM5 (Wellcome Department of Imaging Neuroscience, Institute of Neurology, London, UK).

First, images are corrected 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 non-linear interpolation, temporally 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 was always associated with the same brain location throughout the experiment.[20]

---

[18] More precisely, we acquired 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 acquired the images in an oblique orientation of 30° to the anterior commissure–posterior commissure line (Deichmann et al. (2003)). Each volume of images had 45 axial slices. A total of 692 volumes were collected over two sessions. The imaging parameters were as follows: echo time, 30 ms; field of view, 192 mm; in-plane resolution and slice thickness, 3mm; repetition time, 2.75 s.

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

[20] 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

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 and, as a result, a given spatial location maps to different brain regions in different subjects. Spatial normalization involves a non-linear re-shaping 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 suffices for the purposes of this study. Furthermore, any imperfections in the re-alignment 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, with the weights decreasing with distance.[21] This step is necessary to make sure that the error structure of the data conforms to the normality assumptions about the error structure of the regression models, described below, that we use to test our hypotheses.

Finally, we remove low-frequency signals that are unlikely to be associated with neuronal responses to individual trials.[22]

C. fMRI main data analyses

The key goal of our exercise is to identify 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

---

model of the head motion (3 parameters for center movement, and 3 parameters for rotation) for each volume, and then removing the motion using these parameters. For details, see Friston et al. (1996).

[21] Smoothing was performed using an 8 mm full-width half-maximum Gaussian kernel.

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

responses can be driven by multiple signals. Fortunately, the linear properties of the BOLD signal allow for the identification of the neural signals of interest using standard linear regression methods.

The general 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 that 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 inherent in the BOLD signal, but are not explicitly used in any of our tests. The control variables are not convolved with the HRF because they reflect parameters that affect the measured BOLD responses, and not neural activity that triggers a hemodynamic response.[23]

The linearity of the BOLD signal implies that the level of BOLD activity in any voxel v should be given by

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

where $\varepsilon(t)$ denotes AR(1) noise. This model is estimated independently in each of the brain's voxels using standard regression methods.

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.

[23] For example, linear trends are often included because the scanner heats up with continuous operation and this induces a linear change in the measured BOLD responses.

Two additional considerations apply to most fMRI studies, including the present one. First, we are interested in testing hypotheses about the distribution of the signal coefficients in the population, and not about individual coefficients. This requires estimating a random effects version of the linear model specified above which, given the size of a typical fMRI dataset, is computationally intensive. Fortunately, it has been shown that there is a straightforward 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, given that these tests are carried out in each of the ~100,000 voxels in the brain, there is a serious concern about false-positives, and multiple comparison corrections are necessary. Several approaches have been proposed in the fMRI literature to address this problem, many of which rely on the idea that purely random activations are unlikely to come in sizable clusters.[24] Here, we follow a common approach in the literature, which consists of combining a sizable statistical threshold for the test in each voxel, given by $p<0.001$ uncorrected, together with a minimal cluster size of 15 voxels. These two criteria, taken together, severely reduce the likelihood of false positives.

The analyses described so far involve searching for neural correlates of signals of interest across the entire brain and are therefore known as whole brain analyses. Another popular and very useful type of exercise, which we use here, is a "region of interest" (ROI) analysis. Put simply, this analysis differs from a whole-brain analysis because it first restricts the set of voxels that is being analyzed. The most common types of ROI analyses involve 1) the measurement of signal strength in a pre-specified ROI (in other words, in a pre-specified subset of voxels), 2) computing the correlation across subjects between measures of signal strength in a particular ROI and behavioral or psychological measures, and 3) characterizing the time course of BOLD responses in an ROI for a particular event (e.g, selling a stock.)

[24] As noted earlier, a cluster is a set of spatially contiguous voxels.

The measurement of signal strength in pre-specified ROIs is a straightforward extension of the whole brain analysis. In this case, a general linear model is estimated *only* for the voxels in the ROI, and then a response estimate for the signal of interest is computed for every subject by averaging over the estimated coefficients over all of the voxels in the ROI. The distribution of average estimates for the group can then be compared across signals of interest using t-tests.

The characterization of the time course of BOLD responses in specific ROIs and around particular events is a little more complicated, but is needed in order to conduct a test of Prediction 4. It requires the specification of a version of the GLM described above that uses a series of "event-locked" dummy variables. The nature of this model is most easily explained with a concrete example. Suppose that we are interested in characterizing the time course of changes in BOLD activity that follows the rapid presentation of two types of images to subjects, type A and B. We then define a series of dummy variables

$$d(t|x,n) = \begin{bmatrix} 1 & \text{if stimulus x was presented at } t-n \\ 0 & \text{otherwise} \end{bmatrix}$$

for $x=A,B$, $n=1,\ldots, T$. The general model is then specified as

$$b_v(t) = \text{constant} + \textstyle\sum_{x,n} \gamma_{x,n} d(t|x,n) + \varepsilon(t).$$

The estimate of the change in the BOLD response *n* seconds after the presentation of stimulus *x* is then given by $\gamma_{x,n}$.

IV. Results

A. Behavioral predictions

We begin our test of Prediction 1 by computing the strength of the disposition effect for each subject using the PGR-PLR measure described at the end of Section IIA. We find that the average PGR and PLR across subjects are .412 and .187, respectively. This implies an average PGR-PLR value of 0.225, which is significantly greater than 0 ($p<0.001$). In other words, not only is the average value of PGR-PLR significantly greater than the expected value benchmark of -0.76, but it is actually significantly positive. These results reject the hypothesis that our subjects are all expected value investors and are consistent with the idea that some of our subjects are affected by realization utility.

Figure 3 tests the prediction at the individual level. Each bar shows the value of PGR-PLR for a particular subject. The horizontal dashed line near the bottom of the figure marks the -0.76 value of PGR-PLR that an expected value investor would exhibit. The figure shows that *every* subject exhibits a disposition effect greater than -0.76. The hypothesis that the average disposition effect is not different from -0.76 is rejected with a *t*-statistic of 16.52.

The figure also shows that there is significant heterogeneity in the strength of the disposition effect across subjects: the value of PGR-PLR ranges from -0.41 to 0.83 and has a standard deviation of 0.32. This cross-individual variation is consistent with Dhar and Zhu (2006) who, using data on actual trading decisions, also find significant variation in the strength of the disposition effect across investors. Interestingly, we find that, while each of PGR and PLR varies a good deal across subjects, the two variables have a correlation of only 0.03: subjects who are slow to sell losing stocks are not necessarily also quick to sell winning stocks[25]. This independence between selling behavior in the gain domain and in the loss domain is also consistent with the empirical findings of Dhar and Zhu (2006).

---

[25] The low correlation between PGR and PLR is not inconsistent with realization utility; it simply suggests that realization utility is not the only factor driving subjects' trading. For example, if our subjects care to varying extents about realization utility but also differ in how much they enjoy trading in general, they may exhibit a near-zero correlation between PGR and PLR: the negative correlation between the two variables induced by realization utility will be offset by the positive correlation induced by the taste for trading.

Figure 4 provides additional insight into subjects' selling behavior by showing, for each of the four types of decisions that a subject could make – decisions to realize a gain, decisions to realize a loss, decisions not to realize a gain, and decisions not to realize a loss -- the fraction of the decisions that are optimal, where "optimal" is defined by the expected value benchmark. For example, the figure shows that there were a total of 495 occasions in which our subjects realized gains, and that most of these decisions were suboptimal. Given that stocks exhibit short-term price momentum in the experiment, it is generally better to hold on to a stock that has been performing well. This explains why most (77.9%) of subjects' decisions to hold on to winning stocks were optimal, and why most (67.5%) of subjects' decisions to sell winning stocks were suboptimal. Similarly, in the experiment, it is generally better to sell a stock that has been performing poorly. This explains why most (79.2%) of subjects' decisions to sell losing stocks were optimal, while most (80.3%) of their decisions to hold these stocks were suboptimal.

The disposition effect exhibited by our subjects is stronger than that found in empirical studies (Odean (1998); Frazzini (2006)). One possible reason for this is that the current price and the cost basis of a stock are both prominently displayed on the trading screen.[26] If, because of realization utility, a subject has a preference for realizing gains and for not realizing losses, the fact that we report the purchase price might make it particularly easy for him to cater to this preference, and hence to exhibit a disposition effect.[27]

---

[26] One natural question about our experiment is how much of the realization utility effect that we have found depends on the fact that we display the original purchase price on the trading screen in a highly salient way. It is important to emphasize that it is unlikely that the presence of a realization utility effect depends critically on this aspect of the design. In follow-up work we have carried out behavioral experiments to investigate the impact of the saliency with which the stock purchase price information is displayed (Frydman and Rangel (2011)). We find that eliminating the purchase price from the trading screen diminishes the size of the disposition effect, but that it is still well above the optimal level that an expected value investor would exhibit. This suggests that reporting the purchase price on the trading screen is not a critical aspect of our current design. Moreover, given that most investors in naturally occurring financial markets have at least a rough sense of the price at which they purchased a stock, displaying the cost basis on the trading screen is likely a better approximation of reality.

[27] At the same time, because our experimental design induces a negative correlation between the capital gain and the NEV of selling (r= - 0.55), the fact that we report the purchase price also makes it easy for an

In summary, the behavioral results indicate that all of our subjects exhibit a strong disposition effect, which is inconsistent with the expected value model, but is consistent with the realization utility model.

B. Neural Prediction 2

We now turn to Prediction 2, which states that, for individuals who experience realization utility at the time of selling assets, activity in areas of the brain associated with the computation of decision values, such as the vmPFC and the vSt, should be correlated with the capital gain variable $(p_t - c_t)$. By contrast, it states that, for expected value subjects, activity in these areas should correlate with the NEV variable, but not with the capital gain.

We test this hypothesis in two stages. First, we estimate the following general linear model (GLM) of BOLD activity for each individual:

$$b_v\ (t) = \text{constant} + \ \beta_1^v I_{dec}(t)(p_t - c_t) + \beta_2^v I_{dec}(t) NEV_t + \beta_3^v \text{controls} + \varepsilon(\text{t}).$$

Here, $b_v(t)$ denotes BOLD signal at time t in voxel v. $I_{dec,t}$ is an indicator function that equals one at the time when the subject is presented with the opportunity to *sell* a stock at time t. NEV_t denotes the net expected value from selling the stock being considered at time t, namely $0.6(1-2q_{i,t})$, and $(p_t - c_t)$ is the realizable capital gain. Finally, the *controls* vector includes regressors that control for physical movement inside the scanner, session-specific effects, and any changes in neural activity that might be due to information processing during the price update screens, which is not an activation of interest for the hypothesis being tested. As described in Section III, the regressors involving computations of interest (here, the non-constant regressors NEV and p-c)

expected value subject to trade in a way that is close to *his* optimal strategy, namely to hold a stock when it has a capital gain and to sell it when it has a capital loss. If a subject is an expected value investor, then, we do not think that presenting the purchase price on the trading screen should bias him towards exhibiting a disposition effect.

are convolved with the HRF[28]. Finally, inferences about the extent to which the signals of interest are encoded in a given voxel are made by carrying out a one-sided *t*-test of the individually estimated coefficients (i.e., β_1^v and β_2^v) against zero.

Although we can carry out these tests in all of the brain's voxels, here we limit our search to voxels that belong to pre-specified anatomical areas of the vmPFC and the vSt. These areas were identified using the AAL digital atlas of the human brain (Tzourio-Mazoyer et al. (2002)). Note that these restrictions make our significance threshold of $p<.001$ uncorrected, together with a minimum cluster size of 15 voxels, even less likely to generate false positives than in the standard whole brain analyses to which it is typically applied.

The results from these tests are consistent with the implications of realization utility noted in Prediction 2: we find a cluster of 67 voxels in the vmPFC where $\beta_1^v > 0$. However, no voxels within the vSt exhibit a correlation with the capital gain at our statistical threshold. The location of the vmPFC voxels is depicted in Figure 5. In contrast, there are no clusters that significantly relate to the NEV variable at our statistical threshold. In short, the neural data is consistent with subjects computing the decision value predicted by realization utility, rather than the decision value predicted by the expected value agent model.

Because of the high correlation between the NEV variable and the capital gain variable ($r= -0.55$), we run a robustness check to make sure that the above results are not driven by spurious collinearity issues. This is done by introducing a single change to the GLM: the capital gain variable is orthogonalized (prior to convolution) to the NEV variable, using a standard Gram-Schmidt algorithm (Strang (1988)). Note that this provides an even more stringent test of the realization utility hypothesis because any shared variance between the two variables is now

---

[28] The amount of the price change during the price update screen, which represents our control for information processing, is convolved with the HRF because this will generate a BOLD response. Controls for physical movement inside the scanner and session-specific effects are not convolved because they do not elicit a BOLD response.

allocated to the NEV. As before, we find a cluster of 67 voxels in vmPFC that satisfies the significance criterion described above.

We also carry out an ROI analysis designed to test the properties of the vmPFC realization utility signals further. The relevant ROI (i.e., the relevant subset of voxels within the vmPFC) is defined by estimating the simpler GLM,

$$b_v\ (t) = \text{constant} + \ \beta_1^v I_{dec}(t)(p_t - c_t) + \beta_2^v \text{controls} + \varepsilon(\text{t}),$$

and identifying clusters in the vmPFC that are significantly responsive to the capital gain regressor. For the rest of the paper we refer to the resulting ROI, which contains 154 voxels, as the $\text{vmPFC}_{\text{ROI}}$. Note that we define this ROI using this additional regression to side-step the estimation noise introduced by the high correlation between the capital gain and the NEV regressors.

We then carry out the ROI analysis by estimating the following GLM for each voxel in the newly defined ROI:

$$b_v\ (t) = \text{constant} + \ \beta_1^v I_{dec}(t) p_t + \beta_2^v I_{dec}(t) c_t + \beta_3^v \text{controls} + \varepsilon(\text{t}).$$

This model is interesting because it allows us to compare the strength of the average beta value in the ROI *separately* for the price and cost basis components of the capital gain. Within $\text{vmPFC}_{\text{ROI}}$, β_1=0.025 ($p<0.001$) and β_2=-0.023 ($p<0.01$) and the absolute values of the two coefficients are not significantly different ($p=0.79$). These results demonstrate that the correlation with capital gains that we found above is affected by both the price and cost basis components of the capital gain.

Finally, we carry out a similar ROI analysis to test if the strength of the capital gain signal in $\text{vmPFC}_{\text{ROI}}$ is of similar magnitude in capital gain and capital loss trials. The associated GLM is:

$b_v\ (t) =$

$$\text{constant} + \beta_1^v I_{dec}(t)\ I_{\text{cap.gain}}\ (p_t - c_t) + \beta_2^v I_{dec}(t)\ I_{\text{cap.loss}}\ (p_t - c_t) + \beta_3^v \text{controls} + \varepsilon(\text{t}),$$

where $I_{\text{cap.gain}}$ and $I_{\text{cap.loss}}$ are indicator variables for trials involving capital gains and capital losses, respectively. The average values of β_1 and β_2 are not significantly different (*p=0.69*).

C. Neural Prediction 3

We now test Prediction 3. Specifically, we check whether, as predicted by the realization utility hypothesis, subjects whose neural activity in the vmPFC at the time of a sell decision is particularly sensitive to the realizable capital gain exhibit a stronger disposition effect.

For every subject, we compute the maximum beta value within the $\text{vmPFC}_{\text{ROI}}$ for the capital gain and capital loss regressors[29]. Consistent with Prediction 3, we find that the correlation between β_1 and PGR is 0.78 ($p<0.001$), indicating that subjects who exhibit stronger vmPFC activation in response to a capital gain *do* have a greater propensity to realize gains. Figure 6, which is a scatterplot of PGR against β_1, illustrates this graphically.

In contrast, we do not find a significant correlation between β_2 and PLR ($p=0.18$). One potential post-hoc explanation is that there may be different physiological systems involved in making decisions that involve capital gains and capital losses. Consistent with this hypothesis, the cross-subject correlation between the vmPFC (maximum) sensitivities to capital gains and losses, β_1 and β_2, is only -0.01.

D. Neural Prediction 4

[29] We use a maximum statistic instead of the average statistic because $\text{vmPFC}_{\text{ROI}}$ is relatively large (154 voxels) and because of the heterogeneity in anatomical and functional structure of vmPFC across subjects. Since we are using this beta value to test for a correlation (instead of testing for a particular value of the mean), using the max statistic will not bias our results.

We now test Prediction 4, which constitutes the most direct test of the realization utility hypothesis, and the one that, in our view, showcases the value of the neural data most clearly. The realization utility hypothesis posits that people experience a positive (negative) hedonic impact when they sell a stock at a gain (loss). Since earlier research in neuroscience suggests that activity in the areas of the vSt and the vmPFC correlate with such measures of experienced utility, or hedonics, we can test the hypothesis by looking at changes in the activity in these two areas at the moment that a subject decides to sell a stock, and compare it to changes in the activity in these areas at the moment that a subject issues a command to hold a stock.

More concretely, we test the hypothesis by carrying out the following ROI analysis in specific sub-areas of the vSt and vmPFC that have been shown, in previous studies, to correlate with experienced utility. In particular, we define $vmPFC_{EU\text{-}ROI}$ as the set of voxels that are within 5mm of the voxel whose activity exhibited the highest correlation with experienced utility during consumption of wine in Plassman et al. (2008). Similarly, we define $vSt_{EU\text{-}ROI}$ as the set of voxels that are within 6mm of the two voxels (bilateral) that exhibited the highest correlation with stimulus salience in Zink et al. (2003). Note that the EU subscript in $vmPFC_{EU\text{-}ROI}$ emphasizes that this is a different ROI from the one described above in the analysis of decision values, as it involves a different area of the vmPFC, one that has been previously shown to correlate with hedonics.

The ROI analysis involves estimating the time course of responses in these two ROIs during sell trials involving a capital gain, as a function of whether or not the subject chooses to sell. Figure 7A depicts the results of the analysis for the $vSt_{EU\text{-}ROI}$. The red line indicates changes in the vSt BOLD response for trials in which subjects choose to sell; the blue line shows activity in trials in which subjects choose not to sell. Note that $t=0$ corresponds to the time at which subjects indicate their decision by pressing a button on a hand-held button box-- it is not the time at which the trading screen becomes visible. Interestingly, the figure shows there is no significant difference between the two time series until a decision is made. Afterwards, and consistent with the realization utility hypothesis, activity in the vSt is significantly larger for the next six seconds. The average capital gain for stocks that are held is not significantly different from the average capital gain for stocks that are sold (\$15.77 vs. \$18.35). The effect in Figure 7A is therefore not due to the fact that the stocks subjects sell have larger capital gains than the stocks they hold (p=0.09).

Contrary to our expectations, we did not find a similar result in the $vmPFC_{EU\text{-}ROI}$ (Figure 7B). One possible explanation is that there might be more heterogeneity across subjects in the anatomical and functional structure of the vmPFC, than in the organization of the vSt, which would perhaps mean that the region identified as $vmPFC_{EU\text{-}ROI}$ does not really reflect the areas where experienced utility is computed in our sample. We emphasize, however, that this is pure speculation.

E. Tests of the mean-reversion theory of the disposition effect

As discussed in Section I, one prominent alternative behavioral theory of the disposition effect is that investors believe that stock prices mean-revert (Odean (1998); Weber and Camerer (1998); Kaustia (2010)). In our setting, such a belief would be irrational: stock prices in our experiment exhibit short-term *positive* autocorrelation. Nonetheless, if our subjects, for some reason, *think* that the stock prices in our experiment are mean-reverting, this could potentially explain why they tend to sell stocks that have recently gone up while holding stocks that have recently gone down.

To investigate whether a belief in mean-reversion could be driving our subjects' behavior, we estimate the following mixed effects logistic model to test whether recent price changes can significantly predict subjects' decisions to sell or hold a stock:

$$sell_{i,t,s} = (\alpha + a_i) + \beta_1 NEV_{t,s} + \beta_2 (p_{t,s} - c_{t,s}) + \beta_3 \Delta^1 p_{t,s} + \beta_4 \Delta^2 p_{t,s} + \varepsilon_t \quad (6)$$

where $sell_{i,t,s}$ =1 if subject i sold stock s at time t and 0 if he held it, a_i denotes a subject-level fixed effect, and $\Delta^m p_{t.s}$ denotes the m^{th} most recent price change for stock s (these price changes may not have occurred in consecutive trials because price updates in the experiment take place at random times). We find that the capital gain is a significant predictor of the propensity to sell, (*t-stat=10.04*), but that none of the other variables are. In particular, neither β_3 nor β_4 is significantly different from zero (p=.164 and p=0.160, respectively). In other words, contrary to the mean-reversion hypothesis, recent price changes do not significantly predict the decision to sell.

The neural data can also be used to test some aspects of the mean-reversion hypothesis. In particular, we test if neural activity in either the vmPFC or the vSt is correlated with recent price changes. This is done by estimating the following GLM:

$$b_v\ (t) = \text{constant} + I_{dec}(t)[\beta_1^v(p_t - c_t) + \beta_2^v \Delta^1 p_{t,s} + \beta_3^v \Delta^2 p_{t,s}] + \beta_4^v \text{controls} + \varepsilon(\text{t})$$

Under the mean reversion hypothesis, the decision value of selling should be positively correlated with recent price changes because a recent price increase indicates a lower expected future return, leading to a higher decision value of selling. Neural activity in the vmPFC and vSt should therefore correlate positively with past price changes. Contrary to this hypothesis, we do not find any activity in the vmPFC that is significantly associated with these regressors.

In summary, then, both the behavioral and the neural analyses cast doubt on the mean-reversion hypothesis.

V. Final Remarks

In this paper, we show that neural data obtained through fMRI techniques can be useful in testing theories of investor behavior. Specifically, we use neural data gathered from subjects trading stocks in an experimental market to test a "realization utility" theory of investor trading. While this theory can potentially explain the disposition effect as well as several other financial phenomena, it relies on an unusual assumption: that people derive utility directly from realizing gains. We identify the neural predictions of realization utility and find broad support for them in our data. Perhaps most striking of all, we find that, at the moment a subject issues a command to sell a stock at a gain, there is a sharp rise in activity in the ventral striatum, an area of the brain that, based on recent research in cognitive neuroscience, is known to encode feelings of subjective pleasure.

We emphasize that the methods we present in this paper are hardly a substitute for traditional empirical methods in finance. On the contrary, brain imaging techniques are simply complementary tools that can be used to test assumptions about investor behavior that are difficult to evaluate using field data or experimental data alone. In particular, we see neural data as a valuable resource when studying the more psychological dimensions of individual investor behavior, precisely because these may derive from variables that are only observable at the neural level.

REFERENCES

Ashby, F. Gregory, 2011. *Statistical analysis of fmri data* (The MIT Press).

Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773-806.

Barber, Brad M., and Terrance Odean, 2001, Boys will be boys: Gender, overconfidence, and common stock investment, *Quarterly Journal of Economics* 116, 261-292.

Barberis, Nicholas, and Richard Thaler, 2003, A survey of behavioral finance, in *Handbook of the economics of finance* (Elsevier).

Barberis, Nicholas, and Wei Xiong, 2009, What drives the disposition effect? An analysis of a long-standing preference-based explanation, *Journal of Finance* 64, 751-784.

Barberis, Nicholas, and Wei Xiong, 2011, Realization utility, *Journal of Financial Economics* Forthcoming.

Bernheim, B. Douglas, 2009, On the potential of neuroeconomics: A critical (but hopeful) appraisal, *American Economic Journal: Microeconomics* 1, 1-41.

Blood, Anne J., and Robert J. Zatorre, 2001, Intensely pleasurable responses to music correlate with activity in brain regions implicated in reward and emotion, *Proceedings of the National Academy of Sciences of the United States of America* 98, 11818-11823.

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

Bossaerts, Peter, and Charles Plott, 2004, Basic principles of asset pricing theory: Evidence from large-scale experimental financial markets, *Review of Finance* 8, 135-169.

Bossaerts, Peter, Charles Plott, and William R. Zame, 2007, Prices and portfolio choices in financial markets: Theory, econometrics, experiments, *Econometrica* 75, 993-1038.

Boynton, Geoffrey M., Stephen A. Engel, Gary H. Glover, and David J. Heeger, 1996, Linear systems analysis of functional magnetic resonance imaging in human v1, *The Journal of Neuroscience* 16, 4207-4221.

Bruguier, Antoine J., Steven R. Quartz, and Peter Bossaerts, 2010, Exploring the nature of "trader intuition", *Journal of Finance* 65, 1703-1723.

Camerer, Colin F., 2007, Neuroeconomics: Using neuroscience to make economic predictions, *The Economic Journal* 117, C26-C42.

Camerer, Colin, George Loewenstein, and Drazen Prelec, 2005, Neuroeconomics: How neuroscience can inform economics, *Journal of Economic Literature* 43, 9-64.

Camerer, Colin, and Keith Weigelt, 1991, Information mirages in experimental asset markets, *Journal of Business* 64, 463-93.

Camerer, Colin, and Keith Weigelt, 1993, Convergence in experimental double auctions for stochastically lived assets, in Daniel Friedman, and John Rust,

eds.: *The double auction market: Theories, institutions and experimental evaluations* (Addison-Wesley, Redwood City, CA).
Campbell, John, 2006, Household finance, *Journal of Finance* 61, 1553-1604.
Chib, V.S., A. Rangel, S. Shimojo, and J. 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.
Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, 2009, Reinforcement learning and savings behavior, *Journal of Finance* 64, 2515-2534.
De Araujo, Ivan E. T., Edmund T. Rolls, Morten L. Kringelbach, Francis McGlone, and Nicola Phillips, 2003, Taste-olfactory convergence, and the representation of the pleasantness of flavour, in the human brain, *European Journal of Neuroscience* 18, 2059-2068.
De Martino, Benedetto, Dharshan Kumaran, Ben Seymour, and Raymond J. Dolan, 2006, Frames, biases, and rational decision-making in the human brain, *Science* 313, 684-687.
Deichmann, R., J. A. Gottfried, C. Hutton, and R. Turner, 2003, Optimized epi for fmri studies of the orbitofrontal cortex, *NeuroImage* 19, 430-441.
Dhar, Ravi, and Ning Zhu, 2006, Up close and personal: Investor sophistication and the disposition effect, *Management Science* 52, 726-740.
Fehr, Ernst, and Colin F. Camerer, 2007, Social neuroeconomics: The neural circuitry of social preferences, *Trends in Cognitive Sciences* 11, 419-427.
Fehr, Ernst, and Antonio Rangel, 2011, Neuroeconomics for economists: Why economists should care about what the brain does, *Working Paper*.
Feng, Lei, and Mark S. Seasholes, 2005, Do investor sophistication and trading experience eliminate behavioral biases in financial markets?, *Review of Finance* 9, 305-351.
Frazzini, Andrea, 2006, The disposition effect and underreaction to news, *Journal of Finance* 61, 2017-2046.
Frydman, Cary, and Antonio Rangel, 2011, Attention to realization utility modulates the disposition effect, *Caltech Working Paper*.
Genesove, David, and Christopher Mayer, 2001, Loss aversion and seller behavior: Evidence from the housing market, *Quarterly Journal of Economics* 116, 1233-1260.
Glimcher, Paul W., Ernst Fehr, Colin Camerer, and Russell A. Poldrack, 2008. *Neuroeconomics: Decision making and the brain* (Academic Press, Amsterdam).
Grinblatt, Mark, and Matti Keloharju, 2001, What makes investors trade?, *Journal of Finance* 56, 589-616.
Grinblatt, Mark, and Matti Keloharju, 2009, Sensation seeking, overconfidence, and trading activity, *Journal of Finance* 64, 549-578.
Gul, Farak, and Wolfgang Pesendorfer, 2008, The case for mindless economics, in Andrew Caplin, and Andrew Shotter, eds.: *The foundations of positive and normative economics* (Oxford University Press).
Hare, Todd , Colin F. Camerer, Daniel Knoepfle, John O'Doherty, and Antonio Rangel, 2010, Value computations in ventral medial prefrontal cortex during

charitable decision making incorporate input from regions involved in social cognition, *Journal of Neuroscience* 30, 583-590.

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.

Hare, Todd, John O'Doherty, Colin F. Camerer, Wolfram Schultz, 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.

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.

Hsu, Ming, Ian Krajbich, Chen Zhao, and Colin F. Camerer, 2009, Neural response to reward anticipation under risk is nonlinear in probabilities, *The Journal of Neuroscience* 29, 2231-2237.

Huettel, Scott, Allen Song, and Gregory McCarthy, 2004. *Functional magnetic resonance imaging* (Sinauer Associates).

Hutcherson, C.A., B. Bushong, M. Rabin, and Antonio Rangel, 2011, Pro-social behavior is sustained by dissociable neural mechanisms of different effectiveness, *Working paper*.

Jin, Li, and Anna Scherbina, 2011, Inheriting losers, *Review of Financial Studies* 24.

Kable, Joseph W., and Paul W. Glimcher, 2007, The neural correlates of subjective value during intertemporal choice, *Nat Neurosci* 10, 1625-1633.

Kable, Joseph W., and Paul W. Glimcher, 2009, The neurobiology of decision: Consensus and controversy, *Neuron* 63, 733-745.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263-291.

Kang, Min Jeong, Ming Hsu, Ian M. Krajbich, George Loewenstein, Samuel M. McClure, Joseph Tao-yi Wang, and Colin F. Camerer, 2009, The wick in the candle of learning, *Psychological Science* 20, 963-973.

Kang, Min Jeong, Antonio Rangel, Mickael Camus, and Colin F. Camerer, 2011, Hypothetical and real choice differentially activate common valuation areas, *The Journal of Neuroscience* 31, 461-468.

Kaustia, Markku, 2010, Prospect theory and the disposition effect, *Journal of Financial and Quantitative Analysis* 45, 1-36.

Kennerley, Steven W., Aspandiar F. Dahmubed, Antonio H. Lara, and Jonathan D. Wallis, 2008, Neurons in the frontal lobe encode the value of multiple decision variables, *Journal of Cognitive Neuroscience* 21, 1162-1178.

Knutson, Brian, Scott Rick, G. Elliott Wimmer, Drazen Prelec, and George Loewenstein, 2007, Neural predictors of purchases, *Neuron* 53, 147-156.

Knutson, Brian, G. Elliott Wimmer, Camelia M. Kuhnen, and Piotr Winkielman, 2008, Nucleus accumbens activation mediates the influence of reward cues on financial risk taking, *NeuroReport* 19, 509-513.

Kringelbach, M. L., J. O'Doherty, E. T. Rolls, and C. Andrews, 2003, Activation of the human orbitofrontal cortex to a liquid food stimulus is correlated with its subjective pleasantness, *Cerebral Cortex* 13, 1064-1071.

Kuhnen, Camelia M., and Brian Knutson, 2005, The neural basis of financial risk taking, *Neuron* 47, 763-770.
Levy, Ifat, Jason Snell, Amy J. Nelson, Aldo Rustichini, and Paul W. Glimcher, 2010, Neural representation of subjective value under risk and ambiguity, *Journal of Neurophysiology* 103, 1036-1047.
Litt, Ab, Hilke Plassmann, Baba Shiv, and Antonio Rangel, 2010, Dissociating valuation and saliency signals during decision-making, *Cerebral Cortex*.
Lo, Andrew W., and Dmitry V. Repin, 2002, The psychophysiology of real-time financial risk processing, *Journal of Cognitive Neuroscience* 14, 323-339.
McClure, Samuel M., Jian Li, Damon Tomlin, Kim S. Cypert, Latané M. Montague, and P. Read Montague, 2004, Neural correlates of behavioral preference for culturally familiar drinks, *Neuron* 44, 379-387.
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.
Penny, William, Karl Friston, John Ashburner, Stefan Kiebel, and Thomas Nichols, 2006. *Statistical parametric mapping: The analysis of functional brain images* (Academic Press).
Plassmann, Hilke, John O'Doherty, Baba Shiv, and Antonio Rangel, 2008, Marketing actions can modulate neural representations of experienced pleasantness, *Proceedings of the National Academy of Sciences* 105, 1050-1054.
Plott, Charles R., and Shyam Sunder, 1988, Rational expectations and the aggregation of diverse information in laboratory security markets, *Econometrica* 56, 1085-1118.
Poldrack, Russell A., Jeanette Mumford, and Thomas Nichols, 2011. *Handbook of functional mri data analysis* (Cambridge University Press).
Rabin, Matthew, 2002, Inference by believers in the law of small numbers, *Quarterly Journal of Economics* 117, 775-816.
Rangel, Antonio, Colin Camerer, and P. Read Montague, 2008, A framework for studying the neurobiology of value-based decision making, *Nat Rev Neurosci* 9, 545-556.
Rangel, Antonio, and Todd Hare, 2010, Neural computations associated with goal-directed choice, *Current Opinion in Neurobiology* 20, 262-270.
Rolls, Edmund T., Morten L. Kringelbach, and Ivan E. T. De Araujo, 2003, Different representations of pleasant and unpleasant odours in the human brain, *European Journal of Neuroscience* 18, 695-703.
Shefrin, Hersh, 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.
Small, Dana M., Michael D. Gregory, Y. Erica Mak, Darren Gitelman, M. Marsel Mesulam, and Todd Parrish, 2003, Dissociation of neural representation of intensity and affective valuation in human gustation, *Neuron* 39, 701-711.
Strang, Gilbert, 1988. *Linear algebra and its applications* (Harcourth Brace Jovanovich).

Tom, Sabrina M., Craig R. Fox, Christopher Trepel, and Russell A. Poldrack, 2007, The neural basis of loss aversion in decision-making under risk, *Science* 315, 515-518.

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 F. Camerer, 1998, The disposition effect in securities trading: An experimental analysis, *Journal of Economic Behavior & Organization* 33, 167-184.

Zink, Caroline F., Giuseppe Pagnoni, Megan E. Martin, Mukeshwar Dhamala, and Gregory S. Berns, 2003, Human striatal response to salient nonrewarding stimuli, *The Journal of Neuroscience* 23, 8092-8097.

Figure 1. **Sample screens from a typical trial in the fMRI experiment.** Subjects saw the *price update* screen for two seconds, followed by the *trading* screen for which they had up to three seconds to enter a decision (a blank screen was displayed in between in order to temporally separate neural activity associated with decision-making.) The screens shown below are for a trial in which the subject owns a unit of both stocks A and B. The screens were displayed while subjects were inside the fMRI scanner, and decisions were entered with a handheld device.

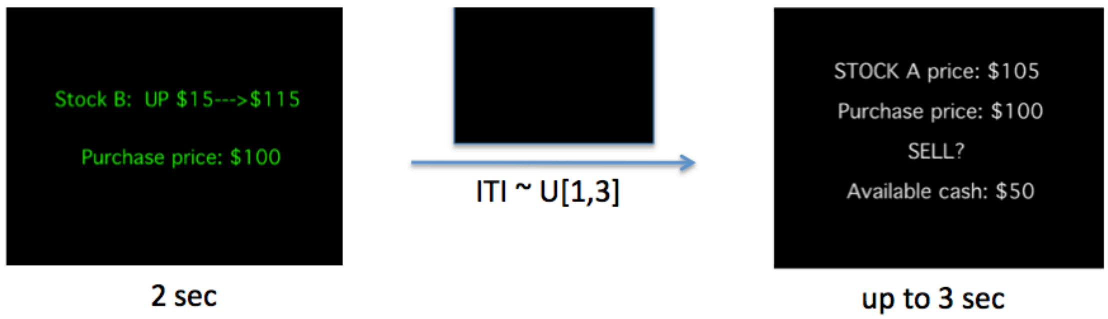

Figure 2. **BOLD measurements of neural activity.** (A) Canonical hemodynamic response function that approximates the BOLD response that follows one arbitrary unit of instantaneous neural activity at time 0. (B) Example of a path of neural activity together with the associated BOLD response.

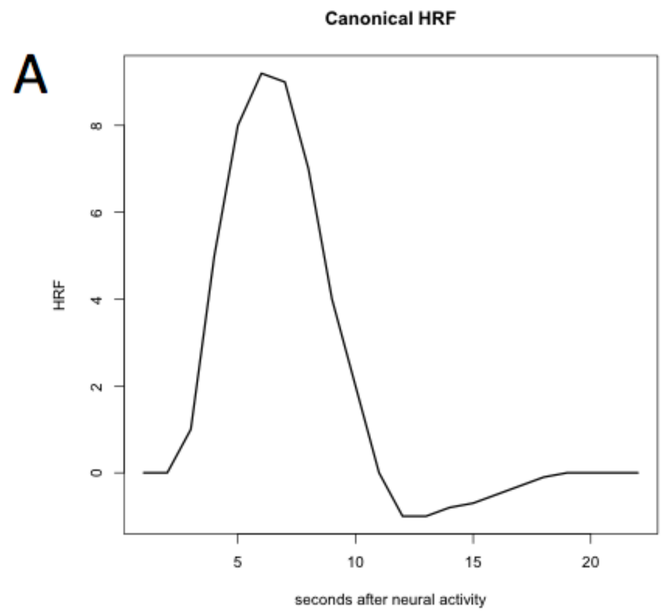

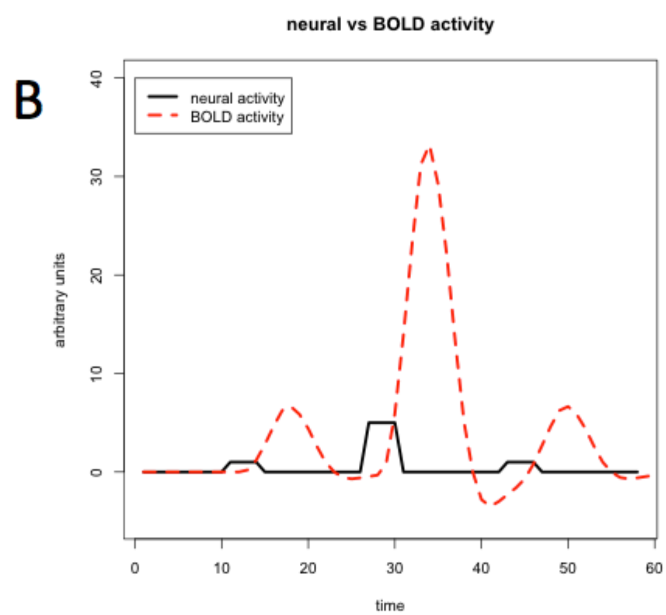

Figure 3. **Measures of the disposition effect (PGR-PLR) for each subject.** Standard error bars are computed as in Odean (1998) and the dotted line indicates the optimal level of the disposition effect, namely -0.76. All subjects exhibit a disposition effect greater than the optimal level and a majority of subjects have a disposition effect that is significantly *positive.* The figure indicates that there is significant heterogeneity in the size of the disposition effect across subjects (SD: 0.32).

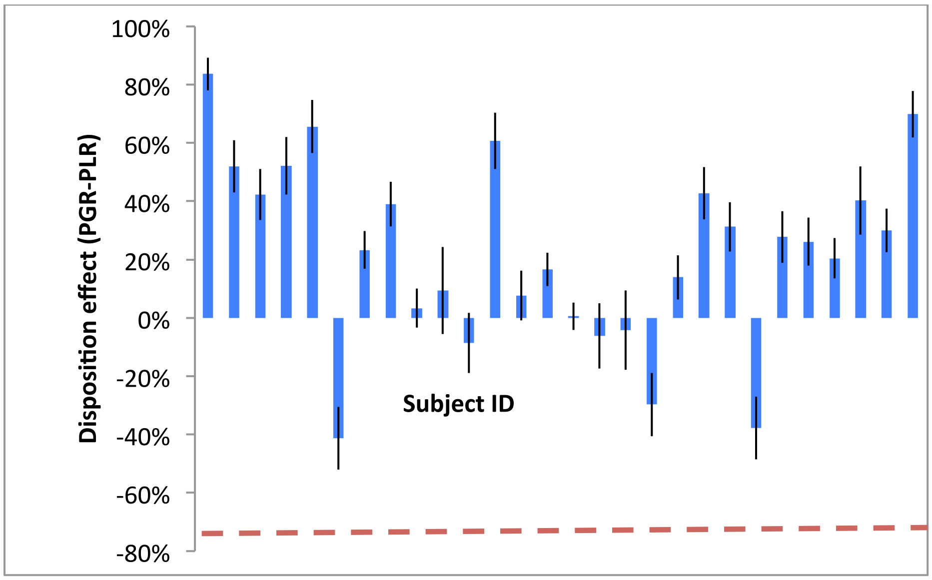

Figure 4. **Total number of sell decisions by decision type and optimality.** Realized gains and losses refer to decisions where subjects sold a stock trading at a gain (loss.) Paper gains (losses) refer to decisions where subjects decided to hold a stock trading at a gain (loss). The optimality measures show an important aspect of our design: selling winners and holding losers, which leads to a disposition effect, are typically suboptimal decisions. Decisions are pooled across all subjects.

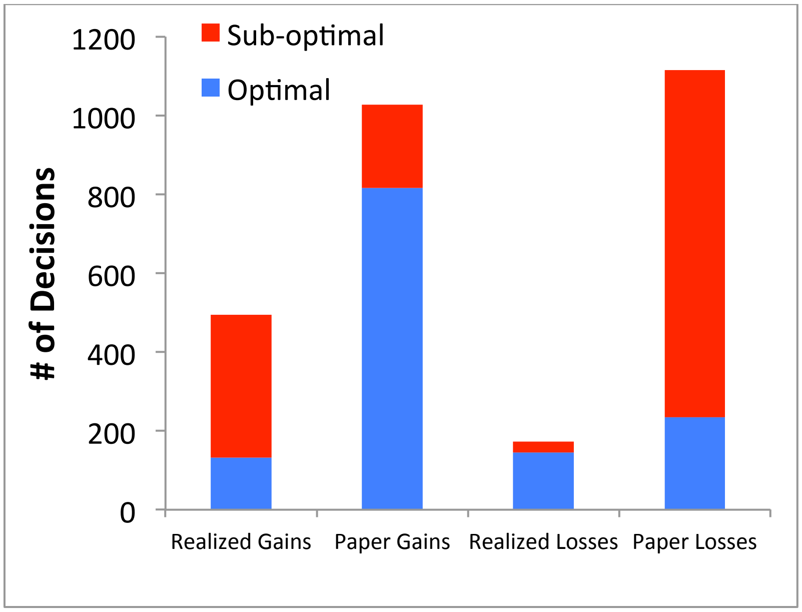

Figure 5. **vmPFC activity reflects realization utility.** Voxels that are shown in yellow all have a p-value less than 0.001, and only clusters of at least 15 significant voxels are shown. Color bar denotes t-statistics.

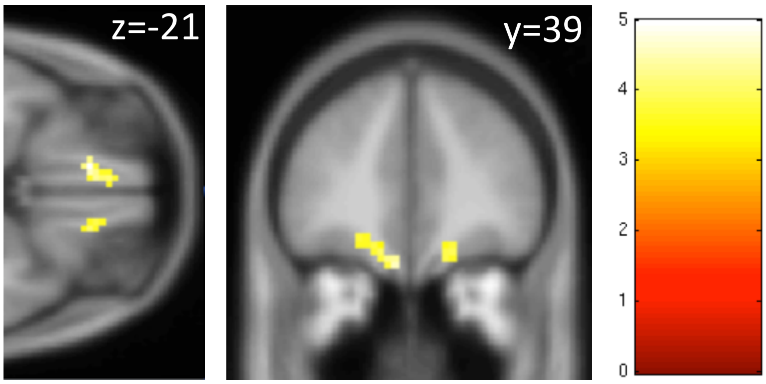

Figure 6. **Correlation between brain activity and measures of the disposition effect.** Each data point in the figure represents a single subject. We find that activity in the vmPFC at the time subjects are offered the opportunity to sell a capital gain is highly correlated with their propensity to realize gains. We do not find a similar correlation between vmPFC activity and the propensity to realize losses.

A)

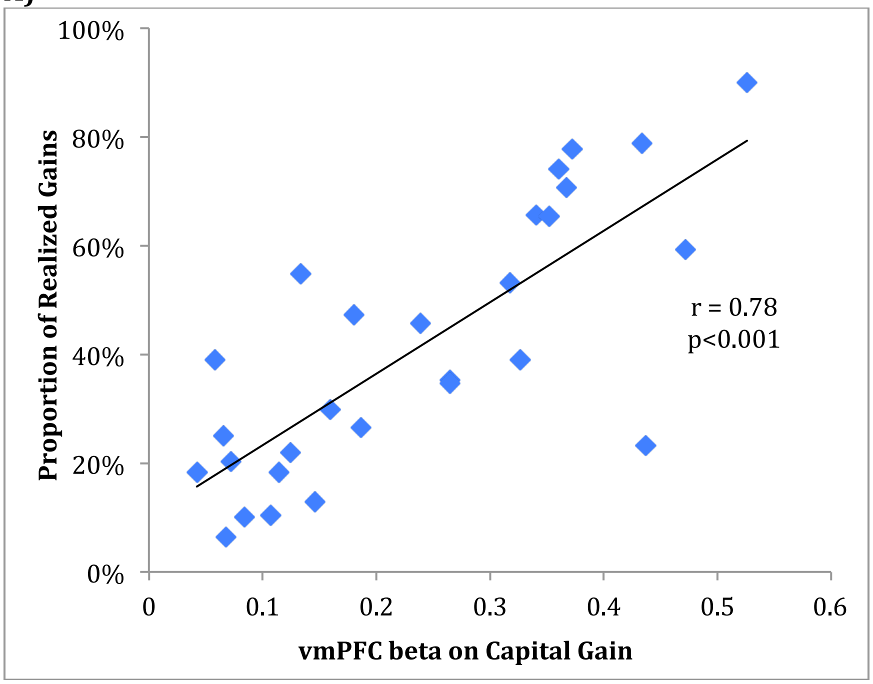

B)

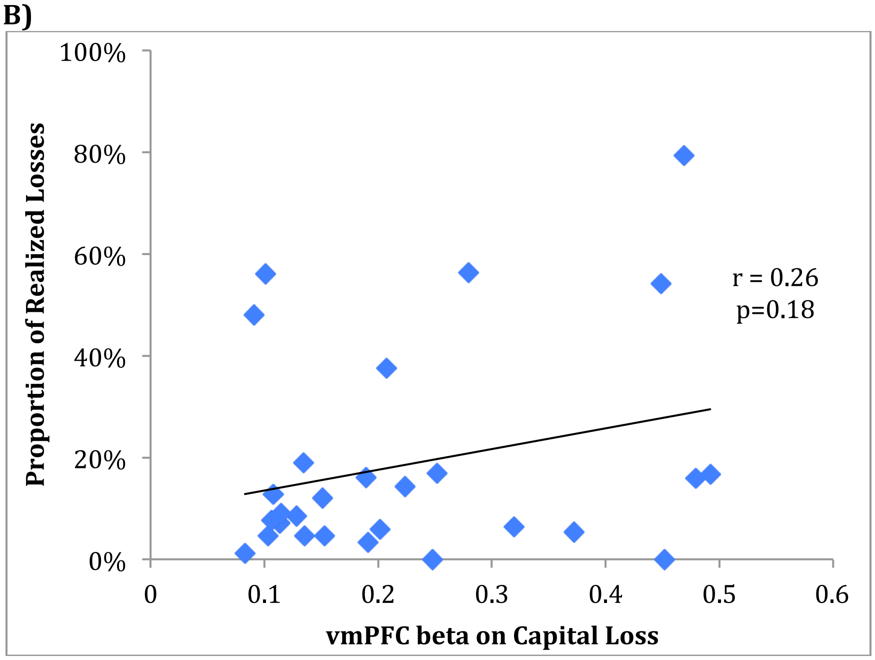

Figure 7. **Direct tests of the realization utility hypothesis.** Average activity in the vSt (Panel A) and vmPFC (Panel B) during trials when subjects were offered the opportunity to sell capital gains. The blue time series plots the average activity in trials where subjects realized capital gains, while the red time series plots the average activity in trials where subjects decided to hold capital gains. *** denotes p<0.001, ** denotes p<0.01, * denotes p<0.05 (paired t-test). *t=0* corresponds to the instant at which the subject enters his trading decision on a hand-held device.

A)

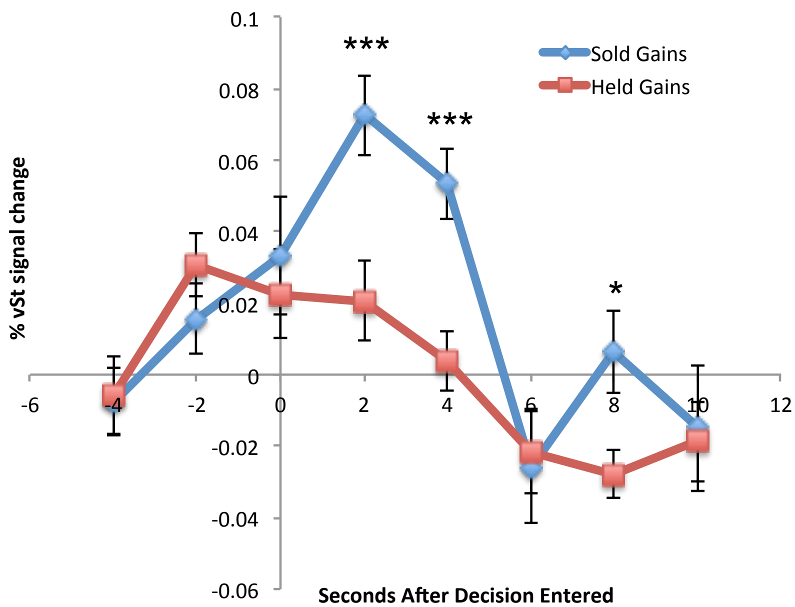

B)

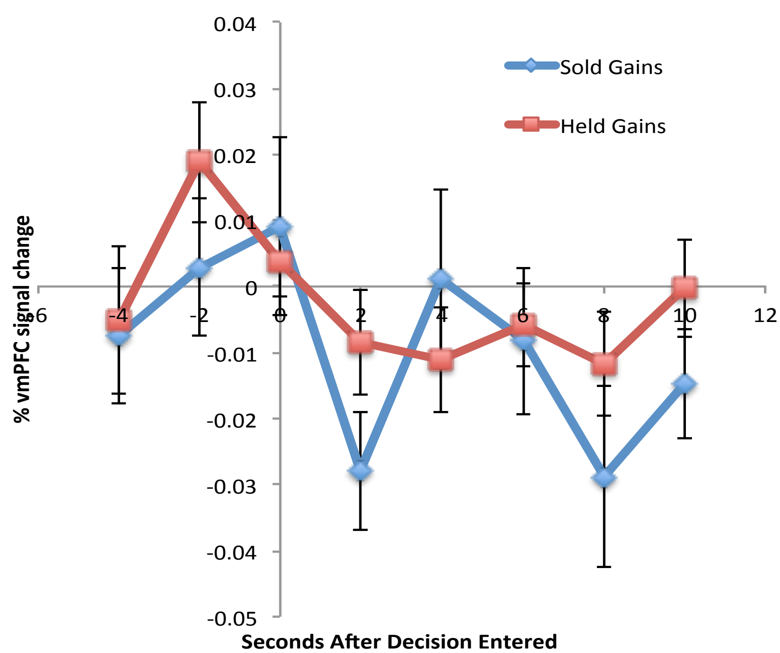

Chapter 2

Neural Measures of Regret and Repurchase Behavior

Economists have learned a great deal about the trading patterns of individual investors over the last two decades. While standard models of trading behavior have had some success in explaining the decision making of individual investors, there remains a set of stylized facts that are hard to reconcile with a standard rational model of trading. In this paper, we examine one of the most recently documented stylized facts about investor behavior, the repurchase effect, which belongs to this set of investor behaviors that are robust, yet difficult to understand using standard preferences and beliefs.

We propose that the repurchase effect, which is the empirical fact that investors tend to repurchase stocks that have gone *down* in value since last sold more often that stocks that have gone *up* in value since last sold, is driven by a negative affective state of regret (Barber et al. (2011)). . In particular, we present a model of a *regret-devaluation* mechanism that provides a link between the affective state induced by prior stock returns and the subjective valuation that the investor applies to subsequent investment opportunities. The main contribution of this paper is to test the regret-devaluation mechanism as the driver of the repurchase effect. However, it is difficult to directly measure regret using a standard empirical data set, and this presents an obstacle for constructing empirical tests of any regret theory. To circumvent this issue, we collect *neural data* that allows us to obtain measurements of regret, so that we can construct several empirical tests of our proposed regret-devaluation mechanism.

We are able to construct sharp empirical tests of this mechanism by relying on two key characteristics of our experimental design. First, we specify a price process in our experimental stock market that induces the optimal trading strategy to be the "opposite" of the repurchase effect, so that stronger repurchase effects lead to lower final earnings. Under this design, if subjects exhibit a repurchase effect, then we can be confident that behavior is due to some alternative belief or preference specification. The second key characteristic of our design is that we temporally separate the events where subjects see news about stock returns from the events where subjects have the ability to trade. This is critical because it allows us to isolate a potential regret signal in the brain, which we use to help predict the trading decision in the subsequent screen.

We find that subjects do exhibit repurchase effects that are higher than an agent who trades to maximize expected final wealth. We also find neural data that is consistent with the regret-devaluation mechanism. In particular, we find that subjects encode a regret signal at the time when they see positive stock returns for assets they positively elected not to purchase in the past. Stronger regret signals at this event predict a higher probability of a repurchase "mistake" during the next opportunity to buy the stock. Finally, we are able to explain a portion of the cross-subject variation of the repurchase effect using the average neural regret signal for each subject. Taken together, our behavioral and neural results provide strong evidence that the repurchase effect is driven by our proposed regret-devaluation mechanism.

The rest of this paper is organized as follows. Section II provides background on investor repurchase behavior and the psychology of regret and its impact on subsequent decision-making. Section III describes our experimental design and outlines the key behavioral and neural predictions of the regret-devaluation mechanism. We give a brief overview of the methods used to collect neural data in Section IV, and we continue with presenting our behavioral and neural results in Section V. We conclude in Section VI and discuss potential directions for future research.

I. Background: Repurchase Behavior, Regret, and Inaction Inertia

A. Repurchase behavior

One of the most recent empirical findings in the household finance literature is that individual investors exhibit a "repurchase effect:" they have a greater propensity to repurchase stocks that have decreased in price since last sale, relative to stocks that have increased in price since last sale (Barber et al. (2011); Weber and Welfens (2011)). This type of behavior is puzzling because stocks that have strong recent performance tend to continue performing well over the short term (Jegadeesh and Titman (1993)) and individual investors could earn higher returns by instead repurchasing stocks that have strong recent performance. Interestingly, it has been found that investors do typically purchase stocks that have strong recent performance, but this does not hold for stocks investors currently own (Odean (1998)). This suggests that past ownership of a specific stock may affect subsequent repurchase behavior.

Because the repurchase effect is not well-explained by a standard rational model of informed trading, (Barber et al. (2011)) suggest that regret may play a role in this type of behavior, noting however, that their *"field data do not enable us to determine definitely the psychological mechanisms that drive these trading patterns."* For example, one potential alternative theory that is consistent with this type of repurchase behavior is a belief in mean-reverting prices: if a stock goes down after an investor sells at a high price, he may want to repurchase it because he believes the price will revert back to the high price. Another theory is that the repurchase effect may arise from tax-motivated trading. In contrast to the field data in (Barber et al. (2011)), the combination of neural and trading data that we collect in our experiment allow us to test a specific theory of regret and its impact on subsequent behavior that can explain the repurchase effect.

B. Regret and Inaction Inertia

Research in psychology has documented an effect called *inaction inertia:* when bypassing an initial action opportunity decreases the likelihood that subsequent similar opportunities will be taken (Tykocinski et al. (1995); Tykocinski and Pittman (1998); Arkes et al. (2002); Kumar (2004)). A simple example involving consumer choice is useful in illustrating this effect. Suppose a consumer is faced with the opportunity to buy a pair of shoes on sale at $40 that are regularly priced at $80. Now suppose for an exogenous reason that the consumer does not buy the shoes on sale at $40, and the next day is offered the shoes at a smaller discount of $50. The inaction inertia effect occurs when the consumer has a lower propensity to buy the shoes at $50 on the 2nd day than he would had he not "missed" the initial, cheaper bargain. In other words, a forgone opportunity has an impact on subsequent behavior[30].

One theory that psychologists have used to explain inaction inertia is based on a *regret-devaluation* mechanism (Arkes et al. (2002)). Using the above example, this theory says that the consumer experiences regret when he is confronted the second opportunity to buy the shoes at the higher price of $50 because he realizes he has "wasted" a more attractive opportunity in the past. This experienced regret then triggers a devaluation of the shoes, and hence results in a lower likelihood of purchasing the shoes during the second opportunity to buy. (Arkes et al. (2002)) provide experimental evidence that is consistent with this theory, as they show that both (self-reported) measures of experienced regret and changes in valuation mediate the link between the price change of a good and the likelihood of buying a good.

We hypothesize that this regret-devaluation mechanism may in part be driving the repurchase effect that individual investors exhibit. To be clear, we now identify the three steps by which regret-devaluation could be applied in a financial market setting. Consider an investor who has a valuation v for a stock, and he has the opportunity to buy this stock at price $p_0 < v$ on day 0, but for an exogenous reason (e.g., liquidity concerns) the investors does not buy the stock. Because the investor has passed up this initial buying opportunity, he finds himself in a setting where the regret-devaluation mechanism applies as follows:

1) <u>*Experienced regret:*</u> Suppose on day 1 the investor sees that the stock price has gone up to $p_1 > p_0$ where $p_1 < v$. The investor will now experience regret in

[30] The inaction inertia effect may seem similar to the well-known cognitive bias of "anchoring," whereby humans have a tendency to rely too heavily, or "anchor," on one piece of information when making decisions. In our setting, the piece of information would be the first price offered. However, a critical difference between the two effects is that inaction inertia relies on a decision of inaction during the first buying opportunity to reduce the probability of buying on the second buying opportunity, whereas anchoring relies only on the presentation of the first buying opportunity, without any decision being made.

proportion to (p_1-p_0) because he has "missed out" on a profit of (p_1-p_0)>0 by not investing on day *0*.

2) *Devaluation:* This experienced regret triggers a change in the investor's valuation of the stock and this change is negatively correlated with (p_1-p_0). Since (p_1-p_0)>0, the investor's new valuation of the stock is $v^* < v$.

3) *Inaction:* This lower valuation, v^*, leads to a lower probability of buying the stock on day 1, compared to a scenario in which the investor did not experience regret from missing out on a profit. That is, $Pr(buy \mid \text{value}=v^*, \text{price}= p_1) < Pr(buy \mid \text{value}=v, \text{price}= p_1)$

This regret-devaluation theory makes several specific predictions about the neural activity we should observe in our experiment. We will develop these predictions in detail in the main text, but we briefly introduce them here.

First, an area of the brain called the ventral Striatum (vSt) has been shown to exhibit activity that correlates with regret or counterfactual signals (Coricelli et al. (2005); Lohrenz et al. (2007); Li and Daw (2011); Nicolle et al. (2011)). In particular, there is evidence that this brain structure computes a signal during a stochastic outcome that is proportional to the difference in the outcome that the subject received and the outcome that the subject *could have received* had he acted differently in the past. In our experiment, the vSt should encode the change in stock price when he is presented with this information, and importantly, this signal should be a function of the subject's previous action associated with the stock: if the subject owns the stock, vSt activity should positively correlate with the change in price, and if the subject does not own the stock, vSt should negatively correlate with the change in price.

Our second neural prediction is that the magnitude of regret we measure using vSt activity should be proportional to the devaluation the subject applies to the stock. A greater devaluation will lead to a lower probability that the subject buys on the next trading opportunity.

Therefore, we should observe, on a trial-by-trial basis, that a stronger regret signal should lead to a larger decrease in the probability of buying the stock.

Finally, those subjects who exhibit a stronger regret signal for a given price change should, on average, decrease their valuation of the stock more than subjects who experience a weaker regret signal for the same price change. Hence, a subject with a stronger average regret signal over the course of the experiment should exhibit a lower average probability of repurchasing a stock that has recently increased in price, and should therefore exhibit a stronger repurchase effect.

II. Experimental Design and Predictions

A. Design

Both the design and the data set used in our analysis were generated in (Frydman et al. (2011)). Subjects are given the opportunity to trade three stocks –- stock A, stock B, and stock C -– in an experimental market. The experiment consists of two identical sessions separated by a one-minute break. Each session lasts approximately 16 minutes and consists of 108 trials. We use t to index the trials within a session.[31]

At the beginning of each session, each subject is given \$350 in experimental currency and is required to buy one share of each stock so that all subsequent purchase decisions are in fact "repurchase" decisions. The initial share price for each stock is \$100; after the initial purchase, each subject is therefore left with \$50. Every trial $t > 9$ consists of two parts: a price update and a trading decision, each of which corresponds to a separate screen 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

[31] We split our experiment into two sessions in order to avoid running the fMRI machine for too long without a break, as this could lead to potential medical risks for the subjects.

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 value of buying or selling a stock, but do not need to update their beliefs about the price process 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 time. 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 buy 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 two-state Markov chain with a good state and a bad state. The Markov chain for each stock is independent of the Markov chains for the other two stocks. Suppose that, in trial t, there is a price update for stock i. If stock i is in the good state at that time, its price increases with probability 0.55 and decreases with probability 0.45. Conversely, if it is in the bad state at that time, 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 a state 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 this trial remains the same as in trial $t-1$ with probability 0.8, but switches with probability 0.2. In mathematical terms, if $s_{i,t} \in \{\text{good}, \text{bad}\}$ 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 never revealed to the subjects: they have to infer them from the observed price paths. To ease comparison of trading performance across subjects, the same set of realized prices is used for all subjects.

A key aspect of our design is that, conditional on the information available to subjects, each of the stocks exhibits positive short-term autocorrelation in its price changes. If a stock performed well 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 positive.

At the end of each session, we liquidate subjects' holdings of the three stocks and record the cash value of their position. We give subjects a financial incentive to maximize the final value of their portfolio at the end of each session. Specifically, if the total value of a subject's cash and risky asset holdings at the end of session 1 is \$X, in experimental currency, and the total value of his cash and risky asset holdings at the end of session 2 is \$Y, again in experimental currency, then his take-home pay in *actual* dollars is 15 + (X+Y)/24.[32] Subjects' 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.

N=28 Caltech subjects participated in the experiment (22 male, age range 18 – 60).[33] 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.

B. Optimal trading strategy

We now characterize the optimal trading strategy for a risk-neutral Bayesian investor who is maximizing the expected value of his take-home earnings – from now on, we refer to such an investor as an "expected value" investor. 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.

---

[32] In other words, we average X and Y to get (X+Y)/2, convert the experimental currency to actual dollars using a 12:1 exchange rate, and add a $15 show-up fee.
[33] One additional subject participated in the experiment but was excluded from further analyses because his head motion during the scanning exceeded a prespecified threshold, thereby interfering with the reliability of the neural measurements.

Formally, let $p_{i,t}$ be the price of stock *i* in trial *t*, after any price update about the stock, and let $q_{i,t} = \Pr(s_{i,t} = good \mid p_{i,t}, p_{i,t-1}, \ldots, p_{i,1})$ 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:

$$q_{i,t}(q_{i,t-1}, z_t) = \frac{\Pr(z_t \mid s_{i,t} = good) * \Pr(s_{i,t} = good \mid q_{i,t-1})}{q_{i,t-1}\Pr(z_t \mid s_{i,t-1} = good) + (1 - q_{i,t-1})\Pr(z_t \mid s_{i,t-1} = bad)}$$

$$= \frac{(0.5 + 0.05z_t) * [0.8 * q_{i,t-1} + 0.2 * (1 - q_{i,t-1})]}{q_{i,t-1}[0.8 * (0.5 + 0.05z_t) + 0.2 * (0.5 - 0.05z_t)] + (1 - q_{i,t-1})[0.2 * (0.5 + 0.05z_t) + 0.8 * (0.5 - 0.05z_t)]}$$

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

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

We refer to this quantity as the *net expected value (NEV)* of buying. Hence, a risk neutral subject will buy whenever NEV is positive ($q_{i,t}$ >0.5) and he will not buy 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 and neural predictions of regret-devaluation

We now lay out the predictions of the model whereby trading decisions are directly affected by the regret-devaluation mechanism. A useful variable we will use in this analysis and in our subsequent behavioral analyses is that of the *foregone capital gain.* This is defined as the difference between the current price of the stock and the price at which the subject last sold the stock. For example, if a subject sold Stock A in trial 19 at a price of $125, and now has the opportunity to repurchase the stock in trial 25 at a price of $105, the foregone capital gain is -$20. By design, this variable is always well defined (when the subject does not hold the stock) because we force subjects to buy all three stocks at the beginning of the experiment. The repurchase effect in our experiment is therefore the tendency of subjects to buy stocks with negative foregone capital gains more often than stocks with positive foregone capital gains. In order to quantify this effect, we use a methodology very similar to that of (Barber et al. (2011)). Each time a subject has an opportunity to repurchase a stock, we place this opportunity into one of four mutually exclusive categories based on the purchase decision (buy or no buy) and the sign of the foregone capital gain (positive or negative). We then compute the relative frequencies that a subject repurchases a stock based on the sign of the foregone capital gain.

In particular, we define two variables: the proportion of stocks that have increased in price since being sold that were repurchased (PDR) and the proportion of stocks that have decreased in price since being sold that were repurchased (PUR). We compute a measure of each of these two variables at the subject level by calculating the following ratios:

$$PDR = \frac{\text{\# of stocks down since being sold repurchased}}{\text{\# of opportunities to repurchase stocks down since being sold}}$$

$$PUR = \frac{\text{\# of stocks up since being sold repurchased}}{\text{\# of opportunities to repurchase stocks up since being sold}}$$

A repurchase effect arises when PDR-PUR>0. In our experiment, the momentum that we build into the stock prices makes the repurchase effect a costly mistake. To see why, consider a subject who sells a stock in period *t* at a price of *s*. Now suppose the subject is offered the opportunity to repurchase the stock several periods later, at a price $p<s$. The fact that the stock price has recently decreased over the last several period is likely because it is in the bad state, and since there is a high probability of the stock remaining in the bad state in the next period, the investor should not invest in this stock. Hence, the optimal strategy dictates buying stocks with strong recent performance, and not buying stocks with poor recent performance. This trading strategy will then result in the *opposite* of a repurchase effect. With the specific price path that is used in our experiment, the optimal strategy will result in a PDR-PUR measure of -0.75.

In contrast, if the regret-devaluation mechanism impacts a subject's decision-making circuitry, then a subject who sees a price increase for a stock he does not own (and therefore one that he either sold recently or chose not to buy) will experience regret. This regret will trigger a devaluation of the stock, leading to a lower propensity to buy the stock. Because this stock has just gone up and because of the positive autocorrelation in price changes, this stock has likely also risen in price since the subject last sold. Hence, the subjects will repurchase fewer stocks that have gone up since last sale than an expected value trader would. This leads to our first prediction.

<u>Prediction 1 (Behavioral):</u> For an expected value investor, the value of the PDR-PUR measure is given by -0.75. On the other hand, for a subject whose decision-making is affect by the regret-devaluation mechanism, the value of PUR-PDR>-0.75.

We now turn to the neural predictions made by the inaction inertia model. As described in section II, an area of the brain called the ventral Striatum (vSt) has been associated with computing *counterfactual* or *regret* signals. In particular, during the outcome of a stochastic event, the vSt has been shown to compute a signal that encodes the difference between what was received and what *could* have been received had the subject acted differently (Coricelli et al. (2005); Lohrenz et al. (2007); Li and Daw (2011); Nicolle et al. (2011)). During the price update screen in our experiment, we therefore expect to see a neural signal in the vSt which satisfies two key properties: 1) the magnitude of the signal should be proportional to the size of the change in price and 2) the sign of the signal should be a function of asset ownership. Formally, define an ownership function, *O(t)*, which takes on the value of 1 if the subject owns the stock being updated in trial *t,* and -1 if the subject does not own the stock in trial *t*. We then expect activity in the vSt to positively correlate with $O(t) \times \Delta p$.

<u>Prediction 2 (Neural):</u> Activity in the vSt during a price update screen should *positively* correlate with the price change if a subject owns the stock. Activity in the vSt during a price update should *negatively* correlate with the price change if the subject does not own the stock.

In the rest of the analysis to follow, it will be useful to distinguish between the counterfactual or regret signals that occur at the price update screen, based on whether the subject owned the stock or not. As the focus of this paper is on testing the regret-devaluation model, we will mainly be concerned with the counterfactual signal computed at the price update screen when the subject *does not* own the stock. For the rest of the paper, we will refer to this signal as the *inaction regret* signal.

Recall from Section II that under the regret-devaluation hypothesis, the amount of regret a subject experiences during a price update screen should be negatively correlated with the price change. Moreover, this experienced regret should be proportional to the devaluation the subject applies to the next trading opportunity. A key implication of this model is that a stronger experienced regret signal during the price update screen on trial *t* should induce a stronger devaluation of the stock on the subsequent trading screen. It follows that a stronger inaction regret signal should induce a lower probability of investing in the stock on the next screen.

Additionally, the link between the experienced regret and the devaluation of the subsequent repurchase opportunity should be asset specific; that is, experienced regret about a rising price in Stock X should affect the probability of repurchasing *only* Stock X on the following trading screen. In contrast, the regret-devaluation hypothesis makes no predictions about how experienced regret from Stock X should impact decision-making regarding repurchasing stock Y. In particular, the regret-devaluation hypothesis *does not* stipulate that experienced regret for one stock will lead to a devaluation for all stocks. We therefore formulate Prediction 3 as follows:

Prediction 3 (Neural): Under the regret-devaluation mechanism, stronger inaction regret signals generated by a price increase about Stock X should lead to a lower probability of repurchasing Stock X on the following trading screen. In contrast, stronger inaction regret signals generated by a price increase about Stock X should not affect the probability of repurchasing stock Y on the following trading screen.

The final prediction of the regret-devaluation hypothesis is aimed at explaining part of the cross-subject variation in the size of the repurchase effect. Recall that the impetus of the regret-devalution mechanism is the experienced regret that is generated when a subject is confronted with news about a price increase for a stock he does not own. If there is any cross-sectional variation in the amount of experienced regret that is generated from a given price increase, then this should help explain some of the cross-sectional variation we see in the size of the repurchase effect. This leads us to prediction 4:

Prediction 4 (Neural): Under the regret-devaluation hypothesis, a stronger inaction regret signal in the vSt should lead to a stronger repurchase effect, across subjects.

III. fMRI data collection and analysis

In this section, we provide a primer on how fMRI measures of neural activity are collected and analyzed. For more details, see Huettel et al. (2004), Ashby (2011), and Poldrack et al. (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 local inflows of oxygenated hemoglobin and outflows of de-oxygenated hemoglobin that occur when neurons fire. fMRI provides measures of the BOLD response of relatively small "neighborhoods" of brain tissue known as *voxels*, and is thought to measure the sum of the total amount of neural firing into that voxel as well as the amount of neuronal firing within the voxel. [34]

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, subsequent BOLD measurements are affected for up to 24 seconds.

[34] Note that the neural activity measured by fMRI in a 1-mm^3 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.

Figure 2 provides a more detailed illustration of the nature of the BOLD response. In particular, it shows the path of the BOLD signal in response to 1 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 elapsed time since the neural activity impulse, and has been shown to approximate well the pattern of BOLD responses for most subjects, brain areas, and tasks.

Fortunately, the BOLD response has been shown to combine linearly across multiple sources of neural activity (Boynton et al. (1996)). This property, along with a specific functional form of the HRF, allows us to construct a mapping from neural activity to BOLD response so that we can control for BOLD responses that are generated by neural activity over the previous 24 seconds. In particular, if the 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

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

which is the convolution between the HRF and the neural inputs. The integral can be interpreted in a straightforward way: it is simply a lagged sum of all the BOLD responses triggered by previous neural activity. This is illustrated in Figure 2B, which depicts a hypothetical path of neural activity, together with the associated BOLD response.

We acquire two types of MRI data during the experiment 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 with a voxel size of 3 mm^3. We acquire data for the entire brain (~ 100,000 voxels) every 2.75 seconds.[35] We also acquire high-resolution anatomical

---

[35] More precisely, we acquired 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 acquired the images in an oblique orientation of 30° to the anterior commissure–posterior commissure line ((Deichmann et al. (2003))). Each volume of images had 45 axial slices. A total of 692 volumes were

scans that we use mainly for realigning the brains across subjects and for localizing the brain activity identified by our analyses.[36]

B. fMRI data preprocessing

Before the BOLD data can be analyzed to test our hypotheses, it has to be converted into a usable format. This requires the following steps, which are fairly standard – see Huettel et al. (2004), Ashby (2011), & Poldrack et al. (2011) – and were implemented using a specialized but commonly used software package called SPM5 (Wellcome Department of Imaging Neuroscience, Institute of Neurology, London, UK).

First, images are corrected 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, temporally 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 was always associated with the same brain location throughout the experiment.[37]

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 and, as a result, a given spatial location maps to different brain regions in different subjects. Spatial normalization involves a nonlinear reshaping of the brain to maximize the match with a target

---

collected over two sessions. The imaging parameters were as follows: echo time, 30 ms; field of view, 192 mm; in-plane resolution and slice thickness, 3 mm; repetition time, 2.75 s.

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

[37] 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 the head motion (3 parameters for center movement, and 3 parameters for rotation) for each volume, and then removing the motion using these parameters. For details, see Friston et al. (1996).

template. Although the transformed data are not perfectly aligned across subjects due to remaining neuroanatomical heterogeneity, the process suffices for the purposes of this study. 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, with the weights decreasing with distance.[38] This step is necessary to make sure that the error structure of the data conforms to the normality assumptions about the error structure of the regression models, described below, that we use to test our hypotheses.

Finally, we remove low-frequency signals that are unlikely to be associated with neuronal responses to individual trials.[39]

C. fMRI main data analyses

The key goal of our exercise is to identify 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. This is complicated by the fact that, since every voxel contains thousands of neurons, the BOLD responses can be driven by multiple signals. Fortunately, the linear properties of the BOLD signal allow for the identification of the neural signals of interest using standard linear regression methods.

The general 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 that we are looking for (e.g., an inaction regret signal). They are specified by a time series $s_i(t)$ describing

[38] Smoothing was performed using an 8 mm full-width half-maximum Gaussian kernel.
[39] Specifically, we applied a high-pass temporal filter to the BOLD data with a cutoff of 128 seconds.

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 inherent in the BOLD signal, but are not explicitly used in any of our tests. The control variables are not convolved with the HRF because they reflect parameters that affect the measured BOLD responses, and not neural activity that triggers a hemodynamic response.[40]

The linearity of the BOLD signal implies that the level of BOLD activity in any voxel v should be given by

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

where $\varepsilon(t)$ denotes AR(1) noise. This model is estimated independently in each of the brain's voxels using standard regression methods.

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.

Two additional considerations apply to most fMRI studies, including the present one. First, we are interested in testing hypotheses about the distribution of the signal coefficients in the population, and not about individual coefficients. This requires estimating a random effects version of the linear model specified above which, given the size of a typical fMRI dataset, is computationally intensive. Fortunately, it has been shown that there is a straightforward 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.

[40] For example, linear trends are often included because the scanner heats up with continuous operation and this induces a linear change in the measured BOLD responses.

Second, given that these tests are carried out in each of the ~100,000 voxels in the brain, there is a serious concern about false-positives, and multiple comparison corrections are necessary. Several approaches have been proposed in the fMRI literature to address this problem, many of which rely on the idea that purely random activations are unlikely to come in sizable clusters.[41] Here, we follow a common approach in the literature, which consists of combining a sizable statistical threshold for the test in each voxel, given by $p<0.001$ uncorrected, together with a minimal cluster size of 15 voxels. These two criteria, taken together, severely reduce the likelihood of false positives.

The analyses described so far involve searching for neural correlates of signals of interest across the entire brain and are therefore known as whole brain analyses. Another popular and very useful type of exercise, which we use here, is a “region of interest” (ROI) analysis. Put simply, this analysis differs from a whole-brain analysis because it first restricts the set of voxels that is being analyzed. The most common types of ROI analyses involve 1) the measurement of signal strength in a prespecified ROI (in other words, in a prespecified subset of voxels), 2) computing the correlation across subjects between measures of signal strength in a particular ROI and behavioral or psychological measures, and 3) characterizing the time course of BOLD responses in an ROI for a particular event (e.g, seeing a price update screen.)

The measurement of signal strength in prespecified ROIs is a straightforward extension of the whole brain analysis. In this case, a general linear model is estimated *only* for the voxels in the ROI, and then a response estimate for the signal of interest is computed for every subject by averaging over the estimated coefficients over all of the voxels in the ROI. The distribution of average estimates for the group can then be compared across signals of interest using t-tests.

[41] As noted earlier, a cluster is a set of spatially contiguous voxels.

IV. Results

A. Behavioral predictions

We begin our test of Prediction 1 by computing the strength of the repurchase effect for each subject using the PDR-PUR measure described earlier. We find that the average PDR and PUR across subjects are .301 and .337, respectively. This implies an average PDR-PUR value of -0.029, which is significantly higher than the optimal level of PDR-PUR=-0.75 ($p<0.001$). At the individual subject level, all but 2 of 28 subjects in our experiment exhibit a repurchase effect that is greater than the optimal level, which is consistent with some of our subjects being affected by the regret-devaluation mechanism. Moreover, Figure 3 shows there is significant variation in the size of the repurchase effect across subjects and suggests there are different types of traders. 6 traders exhibit a significantly positive repurchase effect, 6 traders exhibit a significantly negative repurchase effect, while the remaining 16 subjects exhibit a repurchase effect which is not significantly different form zero.

One potential reason that we do not see more subjects exhibiting a measure of PDR-PUR above zero is because it may be difficult for subjects to recall the price at which they last sold a stock that they are considering repurchasing. If this is the case, we may still see a preference for smaller forgone capital gains (the accumulation of price changes since last sale), without seeing a sharp discontinuity in preferences for foregone capital gains at 0. We can test this by running a logistic regression of the buy decision on the foregone capital and the NEV of buying.

Table 1 shows these regression results, and model 1 indicates that subjects do exhibit a preference for buying stocks with smaller foregone capital gains ($p=0.002$). However, when we run another logistic regression decomposing the foregone capital gain into its two components (model 2), the current price and the price at which the subject last sold, we find that only the current price is a significant predictor of the purchase decision ($p<0.001$). The coefficient on the price at which the subject last sold the stock is positive, but not significantly greater than 0 ($p=0.24$). This suggests that since only the current price is displayed on the trading screen, subjects may have a difficult time calculating the sign of the foregone capital without perfect recall of the price at which they last sold the stock. Note also that the coefficient on NEV is positive in both regression specifications, suggesting that subjects are also (partially) tracking the optimal strategy.

While the focus of this paper is on buying behavior, it is interesting to examine whether there is any relationship to the sell-side behavior of this same set of subjects. In previous work with the same data set (Frydman et al. (2011)), we find that subjects exhibit significant disposition effects, which is the tendency of an investor to *sell* a stock with a capital gain more often than she sells a stock with a capital loss. This effect is suboptimal in our experiment for the same reason that the repurchase effect is suboptimal in our experiment, namely, because of the short-term positive autocorrelation in price changes. Figure 4 shows that there is indeed a correlation between buy-side and sell-side behavior, as subjects who have stronger repurchase effects tend to have stronger disposition effects. In other words, subjects who have a high propensity to sell stocks with capital gains tend to have a low propensity to repurchase stocks with strong recent performance. Such a subject in our experiment is therefore selling stocks with high expected returns and buying stocks with low expected returns.

B. Neural Predictions

We now turn to Prediction 2 which states that under the regret-devaluation mechanism, we should observe a neural signal in the vSt at the price update screen which correlates with the size of the price change. Moreover, this signal should positively correlate with the price change when the subject owns the stock, and it should negatively correlate with the price change when the subject does not own the stock. In other words, the vSt should encode a counterfactual signal representing the difference between what the subjects earned and what he *could have earned* had he acted differently on the previous asset-specific trading screen.

We test this hypothesis by first estimating a general linear model of BOLD activity for each subject:

$$(1) \qquad b_v\ (t) = \text{constant} + I_{pupdate}(t)(\Delta p_t)[\ \beta_1^v I_{own}(t) + \beta_2^v I_{no\ own}(t)] + \beta_3^v \text{controls} + \varepsilon(\text{t}).$$

Here, $b_v\,(t)$ denotes BOLD signal at time t in voxel v. $I_{pupdate}$(t) is an indicator function that equals one at the time when the subject is presented with a price update screen at time t. Δp_t denotes the price change at time t, and I_{own}(t) and $I_{no\ own}$(t) are indicator functions that equal one if the subject owns or doesn't own the stock at the price update screen at time t, respectively. Finally, the *controls* vector includes regressors that control for physical movement inside the scanner, session-specific effects, and any changes in neural activity that might be due to the decision processes from previous trading screens, which are not activations of interest for the hypothesis being tested. Finally, inferences about the extent to which the signals of interest are encoded in a given voxel are made by carrying out a one-sided *t*-test of the individually estimated coefficients (i.e., β_1^v and β_2^v) against zero.

Although we can carry out these tests in all of the brain's voxels, here we limit our search to voxels that belong to the prespecified anatomical area of the vSt. This area was identified using the AAL digital atlas of the human brain (Tzourio-Mazoyer et al. (2002)). Note that this restrictions makes our significance threshold of $p<.001$ uncorrected, together with a minimum cluster size of 15 voxels, even less likely to generate false positives than in the standard whole brain analyses to which it is typically applied.

Consistent with our hypothesis, we find that the vSt exhibits activity at the time of a price update screen which correlates with the size of the price change. In particular, Figure 5a shows areas of the brain where $\beta_1^v > 0$ and Figure 5b shows areas of the brain were $\beta_2^v < 0$. We find a cluster of 36 voxels in the left vSt and a cluster of 41 voxels in the right vSt where *both* $\beta_1^v >$

$0\ and\ \beta_2^v < 0$. Within this set of 77 voxels, the average β_1^v =0.090 and the average β_2^v =-0.058, and the magnitude of the counterfactual signal during the price update screen is significantly higher when subjects own the stock as compared to when they don't own the stock (p=0.02). The fact that $\beta_2^v < 0$ implies that when subjects are presented with a price update for a stock they do not own, the vSt encodes a signal which negatively correlates with the size of the price change.

We now turn to Prediction 3, which examines the key implication of the regret-devaluation mechanism: experienced regret causes a devaluation of the subsequent repurchase opportunity, which then leads to a lower probability of repurchase. Until now, this mechanism has only been tested using self-reported measures of experienced regret (Arkes et al. (2002)); here, we introduce a method which highlights the value of the neural data in *directly* measuring experienced regret so that we can assess its impact on subsequent trading decisions.

Our design allows us to construct a novel test of this mechanism by exploiting exogenous variation in the ordering of specific stocks in the price update and trading screens. Recall from the experimental design that there is a 1/3 probability that each stock is displayed on the price update screen, and there is a 1/3 probability each stock is displayed on the trading screen. This allows us to perform tests about whether experienced regret from stock X (generated during the price update screen) affects the repurchase decision of stock X and *only* stock X. We construct this test in two steps. First, we partition trials where the subject sees a price update for stock X (an asset she doesn't own) into three categories based on the trading opportunity in the subsequent trading screen: 1) trading screens where the subject has an opportunity to buy Stock X, 2) trading screens where the subject has an opportunity to buy stock Y≠X, and 3) trading screens where the subject does not have the opportunity to buy any stock (i.e., only has the opportunity to sell). The second step is to estimate a GLM containing regressors that model these three types of trials:

(2)

$$b_v\ (t) = \text{constant} + I_{no\_own\_pupdate}(t)(\Delta p_t)\left[\ \beta_1^v I_{same}(t) + \beta_2^v I_{different}(t) + \beta_3^v I_{no\_opp}(t)\right] +$$

$$I_{no\_own\_update}(t) I_{Buy}(t)\left[\ \beta_4^v I_{same}(t) + \beta_5^v I_{different}(t)\right] + \beta_6^v \text{controls} + \varepsilon(\text{t}).$$

With the above GLM, we are able to test the precise implications of prediction 3, by examining whether inaction regret signals in the vSt predict the stock-specific repurchase decision. We perform an ROI analysis in the vSt[42] to test whether activity in this area of the brain can predict the stock-specific repurchase decision. Consistent with *both* parts of prediction 3, we find that the left vSt does indeed predict the subsequent repurchase decision, *only* when the trading screen asset is the same as the price update screen asset[43]. Specifically, Figure 6 shows that within our left vSt ROI, $\beta_4^v = 2.21$ ($p = 0.003$) and $\beta_5^v = 0.11$ ($p = 0.83$) and that $\beta_4^v > \beta_5^v$ ($p < 0.003$). We do not find a similar effect in the right vSt, but the signal in the left vSt passes a Bonferroni test with two comparisons.

It is important to note that GLM (2) above includes a control for the change in price, Δp_t, and we are still able to predict the stock-specific repurchase decision from the inaction regret signal. This is important because it shows that there is additional predictive information in this signal beyond the price change itself, and hence provides evidence of the marginal value of the neural data. It is also consistent with previous work in psychology (Arkes et al. (2002)) that experienced regret is a mediating factor of the inaction inertia effect.

Our final neural prediction is aimed at using the inaction regret signal to explain a portion of the cross-subject variation in the repurchase effect (Figure 3). Building on our result from prediction 3 that activity in the left vSt is a significant predictor of the repurchase decision, we expect to see a correlation between the strength of the inaction regret signal in the left vSt and the

---

[42] We define the vSt ROI as the set of voxels that are within 6 mm of the two voxels (bilateral) that exhibit the highest correlation with stimulus salience in (Zink et al. (2003)).

[43] One subject was excluded from this analysis because there were no trials where the subject had an opportunity to buy stock Y≠X, after seeing a price update for stock X.

size of the repurchase effect across subjects. Because the association between left vSt activity and the repurchase decision is present only for trials in which the price update asset was the same as the trading screen asset, we first compute a "restricted" measure of (PDR-PUR) for each subject *i* using only those trials where the price update screen asset is the same as the trading screen asset[44]. We then run a regression of this restricted measure of (PDR-PUR) on the inaction regret signal, which is defined as the estimated β_2 coefficient from GLM (1):

$$(PDR - PUR)_i = \alpha + \gamma\beta_{2_i} + \epsilon$$

After weighting each observation inversely proportional to the standard error of $(PDR - PUR)_i$, we find that $\hat{\gamma}$=-1.81 (*t-stat: -2.09, p=0.047)* which is consistent with prediction 4 that subjects with stronger inaction regret signals exhibit a stronger repurchase bias. Figure 7 displays a scatterplot of the inaction regret signal vs. (PDR-PUR). It is interesting to note that the two subjects who are nearly tracking the optimal level of (PDR-PUR)=-0.75 have an inaction regret signal of almost 0 or slightly positive. This suggests that the few subjects who are repurchasing optimally have very little sensitivity to news about returns on foregone investments.

V. Final Remarks

In this paper we provide a model of investor behavior that yields a repurchase effect, which is a systematic trading pattern that has recently been documented in the class of individual investors {Barber, 2011 #130}. Our model is based on a regret-devaluation mechanism, which is difficult to test using standard empirical data sets because regret is not directly observable. We sidestep this issue by collecting neural data that allows us to directly measure regret and test for its effect on subsequent decision-making. The neural data strongly supports the regret-devaluation

[44] The unrestricted and restricted measures of (PDR-PUR) have a correlation of 0.46.

mechanism as the driver of the repurchase effect. While our concern in this paper is with a very specific behavioral implication of regret, the repurchase effect, there is a vast literature on models of regret in economics which may benefit from a similar style of empirical testing by using neural data to directly observe regret signals.

REFERENCES

Arkes, Hal R., Yi-Han Kung, and Laura Hutzel, 2002, Regret, valuation, and inaction inertia, *Organizational Behavior and Human Decision Processes* 87, 371-385.

Ashby, F. Gregory, 2011. *Statistical analysis of fmri data* (The MIT Press).

Barber, Brad M., Terrance Odean, and Michal Strahilevtz, 2011, Once burned, twice shy: How naive learning, conterfactuals, and regret affect the repurchase of stocks previously sold, *Journal of Marketing Research* 48, 102-120.

Boynton, Geoffrey M., Stephen A. Engel, Gary H. Glover, and David J. Heeger, 1996, Linear systems analysis of functional magnetic resonance imaging in human v1, *The Journal of Neuroscience* 16, 4207-4221.

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, *Nat Neurosci* 8, 1255-1262.

Deichmann, R., J. A. Gottfried, C. Hutton, and R. Turner, 2003, Optimized epi for fmri studies of the orbitofrontal cortex, *NeuroImage* 19, 430-441.

Frydman, Cary, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel, 2011, Testing theories of investor behavior using neural data, *Caltech Working Paper*.

Huettel, Scott, Allen Song, and Gregory McCarthy, 2004. *Functional magnetic resonance imaging* (Sinauer Associates).

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65-91.

Kumar, Piyush, 2004, The effects of social comparison on inaction inertia, *Organizational Behavior and Human Decision Processes* 95, 175-185.

Li, Jian, and Nathaniel D. Daw, 2011, Signals in human striatum are appropriate for policy update rather than value prediction, *The Journal of Neuroscience* 31, 5504-5511.

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.

Nicolle, Antoinette, Stephen M. Fleming, Dominik R. Bach, Jon Driver, and Raymond J. Dolan, 2011, A regret-induced status quo bias, *The Journal of Neuroscience* 31, 3320-3327.

Odean, Terrance, 1998, Are investors reluctant to realize their losses?, *Journal of Finance* 53, 1775-1798.

Penny, William, Karl Friston, John Ashburner, Stefan Kiebel, and Thomas Nichols, 2006. *Statistical parametric mapping: The analysis of functional brain images* (Academic Press).

Poldrack, Russell A., Jeanette Mumford, and Thomas Nichols, 2011. *Handbook of functional mri data analysis* (Cambridge University Press).

Tykocinski, O., and T. Pittman, 1998, The consequences of doing nothing: Inaction inertia as avoidance of anticpated counterfactual regret, *Journal of Personality and Social Psychology* 75.

Tykocinski, O., T. Pittman, and E. Tuttle, 1995, Inaction inertia: Foregoing future benefits as a result of an iniitial failure to act, *Journal of Personality and Social Psychology* 68.

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 Frank Welfens, 2011, The follow-on purchase and repurchase behavior of individual investors: An experimental investigation, *Die Betriebswirtschaft* 71, 139-154.

Zink, Caroline F., Giuseppe Pagnoni, Megan E. Martin, Mukeshwar Dhamala, and Gregory S. Berns, 2003, Human striatal response to salient nonrewarding stimuli, *The Journal of Neuroscience* 23, 8092-8097.

Table 1. **Determinants of Propensity to Buy.** Logistic regression where dependent variable equals 1 if the subject bought on the trial and 0 if the subject did not buy (conditional on the opportunity to buy). Forgone 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.

	Model 1		Model 2	
	coefficient	*p-val*	*coefficient*	*p-val*
NEV	**3.26**	0.087	**4.42**	0.025*
Foregone Capital Gain	**-0.014**	0.002**		
Current Price			**-0.022**	0.001***
Last Sale Price			**0.005**	0.237
constant	**-0.96**	0.001***	**1.06**	0.031*
# of obs	2410		2410	

Figure 1. **Sample screens from a typical trial in the fMRI experiment.** Subjects saw the *price update* screen for two seconds, followed by the *trading* screen for which they had up to three seconds to enter a decision (a blank screen was displayed in between in order to temporally separate neural activity associated with decision-making.) The screens shown below are for a trial in which the subject owns a unit of both stocks A and B. The screens were displayed while subjects were inside the fMRI scanner, and decisions were entered with a handheld device.

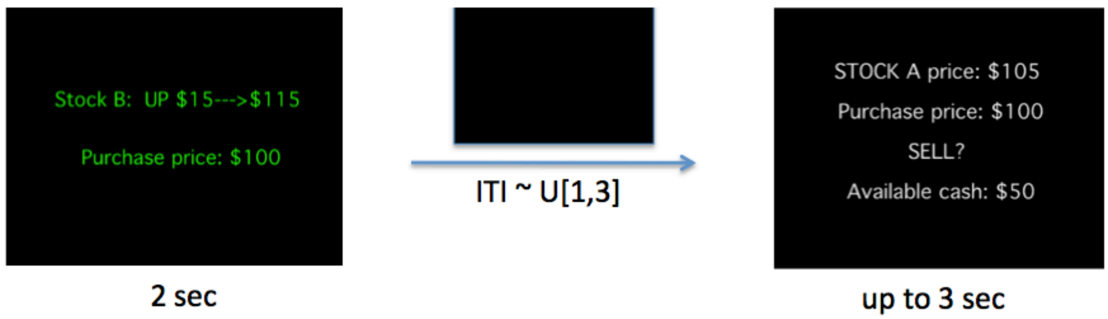

Figure 2. **BOLD measurements of neural activity.** (A) Canonical hemodynamic response function that approximates the BOLD response that follows one arbitrary unit of instantaneous neural activity at time 0. (B) Example of a path of neural activity together with the associated BOLD response.

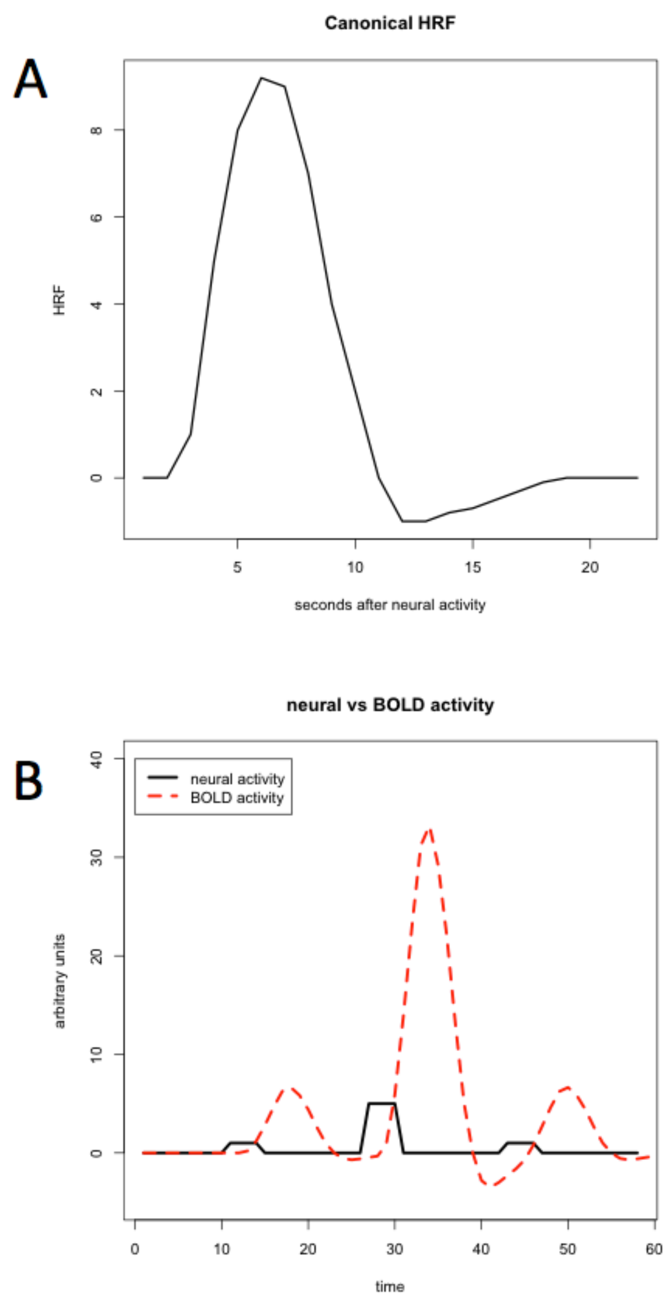

Figure 3. **Measure of (PDR-PUR) for each of the 28 subjects.** Each black line represents 2 standard errors. All but 2 subjects exhibit a measure of (PDR-PUR) that is significantly greater than the optimal level of -0.75 (denoted by the dashed horizontal line.) The figure indicates that there is significant heterogeneity in the size of the repurchase effect across subjects (SD: 0.30).

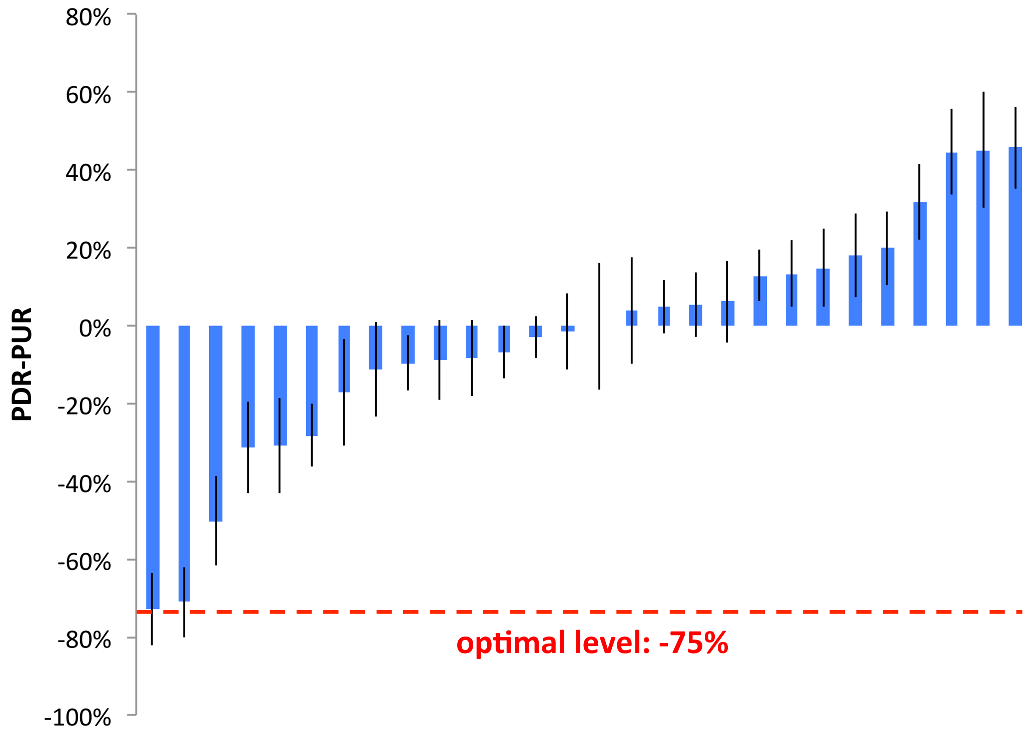

Figure 4. **Relationship between buying and selling behavior.** Each blue point represents a subject. The horizontal axis measures the repurchase effect and the vertical axis measures the disposition effect. Similar to the repurchase effect, a positive disposition effect is also suboptimal in our experiment, where the optimal level of the disposition effect is -0.76.

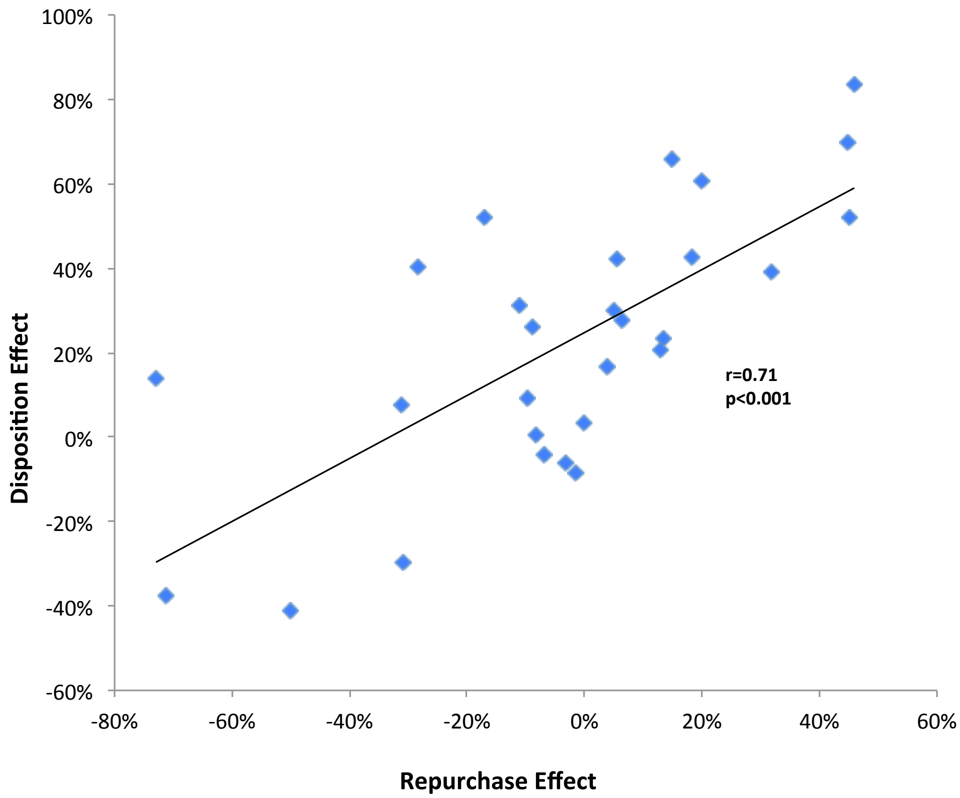

Figure 5. **vSt encodes a counterfactual signal.** A) Voxels in the brain which exhibit activity that positive correlates with the price change at the time when subjects receive news about an asset they own. B) Voxels in the brain which exhibit activity that negatively correlates with the price change at the time when subjects receive news about an asset they do not own.

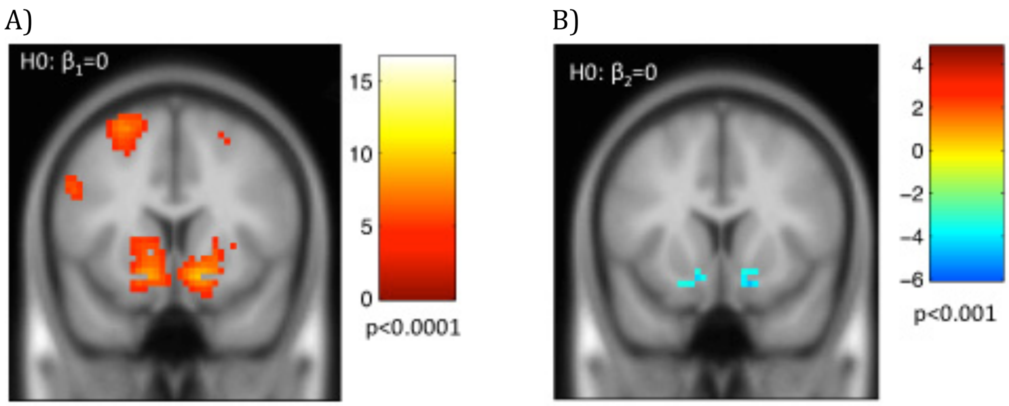

Figure 6. **Activity in vSt predicts trial-by-trial repurchase decision.** Average activity in the left vSt during price update screen when asset is not owned. Stronger inaction regret signals are seen preceding failure to repurchase relative to repurchase, on those trials where price update asset is the same as the trading screen asset (RIGHT). There is no difference in inaction regret signals when assets are different on price update and trading screen (LEFT).

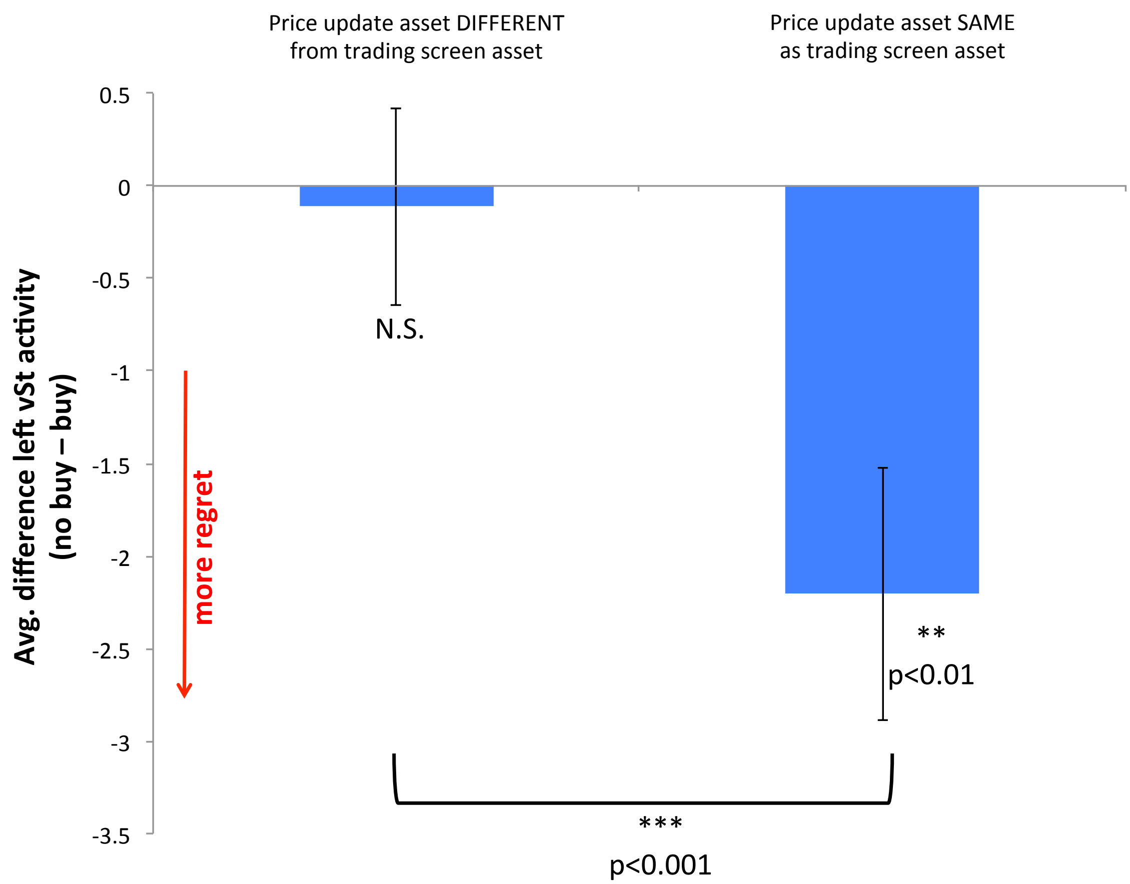

Figure 7. **Inaction regret signal in vSt correlates with repurchase effect across subjects.** The sensitivity of the left ventral Striatum to price changes when a subjects does not own an asset (inaction regret signal) is negatively correlated with the size of the repurchase effect. Inaction regret signals are defined as the estimate of the β_2 coefficient from GLM (1) in the text. Stronger inaction regret signals (more negative vSt sensitivity to price update) are therefore associated with stronger repurchase effects.

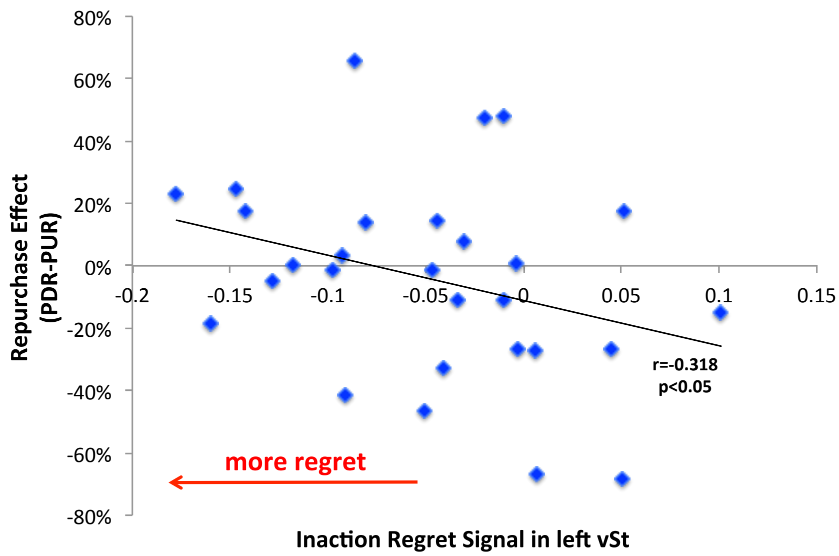

Chapter 3

Attention to Realization Utility Modulates the Disposition Effect

Individual investors face a daunting task when constructing a portfolio of stocks because of the vast amount of information available to them. Attention is a scarce resource, and it is inevitable that investors will eventually deplete their stock of attention and disregard valuable information when trading. A growing body of literature in finance suggests that these constraints on attention can have a systematic effect on investor behavior and asset prices (Hong and Stein (1999); Hirshleifer and Teoh (2003); Peng and Xiong (2006); Barber and Odean (2008); Cohen and Frazzini (2008); Dellavigna and Pollet (2009); Duffie (2010); Da et al. (2011)). While much of this literature focuses on underreaction to firm news or strategic disclosure of information to investors, another implication of inattention is the manner in which an investor perceives the "attributes" of his own portfolio. For example, in the case that an investor experiences a liquidity shock and is forced to sell a position in his portfolio, what attributes does the investor attend to when formulating the decision of which security to sell?

In this paper, we assume the investor allocates his attention to two general attributes of each security that he holds in his portfolio: the expected return and the past performance. We focus on these two general attributes because standard theory predicts that investors should trade based on expected returns, but there is substantial evidence that investors also trade based on a stock's past performance (Grinblatt and Keloharju (2001); Kaustia (2010)). Our main question in this paper is whether an investor's allocation of attention to each of these two attributes is fixed, or whether it can be modulated by the saliency of expected returns or past performance. In essence, are investor preferences malleable to be more or less "forward-looking?"

This question is important because if investors allocate attention to the most salient items on financial statements (Libby et al. (2002); Hirshleifer and Teoh (2003)), then the design of financial statements can have a direct impact on investor behavior. To study this attention allocation procedure, we have subjects trade in an experimental stock market and we manipulate the saliency of the expected return and the past performance of a stock. We use the capital gain on the stock as the measure of the stock's past performance, and as we describe in the main text, we use an optimal (Bayesian) prediction of the stock's next period price change as a measure of expected return.

We use a stock's capital gain as its measure of past performance because there is a large empirical literature documenting that individual investors trade based on how well the stock has done since the investor purchased it (Grinblatt and Keloharju (2001); Kaustia (2010)). While this trading behavior does not typically arise from a standard model (eg, power utility with i.i.d. returns), a recently developed preference specification, called realization utility, provides a theoretical foundation for the empirical link between the capital gain and propensity to sell (Shefrin and Statman (1985); Barberis and Xiong (2012)). Realization utility is based on the assumption that investors derive positive (negative) utility *directly* from the act of selling capital gains (capital losses) on risky assets that they own. Contemporaneous experimental and neurobiological research provides evidence that is consistent with the idea that investors do have a direct preference for selling stocks with capital gains (Frydman et al. (2011)). If investors have realization utility preferences in addition to standard preferences over consumption, it seems plausible that trading behavior should be driven by an investor's attention to *both* capital gains and expected returns; that is, trading behavior should be affected by the saliency of the past and future.

One particular implication of realization utility that we will focus on in this paper is the *disposition effect*, which is the empirical fact that investors have a greater propensity to sell stocks that have gone up in price since purchase, relative to stocks that have gone down in price since purchase[45]. We focus on this effect for two reasons: 1) it provides a measure of the strength of

---

[45] See for example (Shefrin and Statman (1985); Odean (1998); Genesove and Mayer (2001); Feng and Seasholes (2005); Frazzini (2006); Jin and Scherbina (2011)).

realization utility preferences (Frydman et al. (2011)) and 2) it is a robust behavior that has been found to lower overall trading performance (Odean (1998); Weber and Camerer (1998)). In particular, (Odean (1998)) finds that, on average, the winning stocks investors sell tend to outperform the losing stocks they hold by 3.4%. This is consistent with the momentum found in equity returns, because winning stocks will on average continue to have higher returns than losing stocks over the short-term (Jegadeesh and Titman (1993)). It follows that if attention to realization utility can be manipulated, then we should be able to affect the size of the disposition effect and therefore the trading performance of the investor.

Indeed, we are able to establish a causal link between attention to realization utility and the size of the disposition effect. Our results can be useful for the optimal design of financial statements. For example, brokerage houses can design regular financial statements that highlight expected returns, or they can decrease the salience of past trading performance. One particular method of implementing the latter would be to decrease the saliency of the cost basis, so that less relative attention is focused on the capital gain of the stock. In our experimental stock market, we implement this treatment and find that the disposition effect is significantly attenuated, and overall earnings are increased. This is interesting because recent legislation enacted by the US government has mandated that the cost basis should effectively become *more* salient when investors are in the decision phase of selling an asset[46]. Congress passed this cost basis legislation in order to increase tax compliance with capital gain laws, but a potential unintended consequence of such a law could be that investors now focus more on the cost basis when deciding whether to sell. According to our experimental results, this will induce a higher disposition effect, and will lower overall average trading performance. Interestingly, this should work to increase government tax revenue even more, since an investor's optimal tax policy is to realize losses immediately and defer realizing gains far into the future in order to minimize the present value of the capital gains tax burden (Constantinides (1983)).

---

[46] See Emergency Economic Stabilization Act of 2008

Our results can also inform specific channels through which effective financial education can be implemented. Many current forms of financial education, for example, disclaimers on mutual fund investments, stress that past performance should not be used to form beliefs about future returns. However, our results show that besides this information channel, there is another potential mechanism through which past performance may influence trading decisions, namely, via a direct preference for realizing capital gains. This additional preference-based mechanism should therefore also be taken into consideration when educating investors about the potential disadvantages of using past performance when trading.

I. Theory and Experimental Design

A. Design

We now describe an experiment that we use to examine the attention that subjects pay to past performance and future returns, which is based on a design used in (Frydman et al. (2011)). All subjects are given the opportunity to trade three stocks – stock A, stock B, and stock C – in an experimental stock market. The experiment consists of two sessions separated by a two-minute break. Each session lasts approximately 16 minutes and consists of 108 trials. In this section, we use t to index the trials. The first session consists of trials t=1 through t=108, and the second, of trials t=109 through 216. We now describe the structure of the first session; the structure of the second session is identical to that of the first.

Before trial 1, 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 this transaction, each subject is therefore left with \$50. The majority of the trials – specifically, trials 10 through

108 – consist of two parts: a price update and a trading decision, each of which corresponds to a separate screen that the subject sees (see 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 decision 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 the trading decision part.

Trials 1 through 9 consist only of a price update part; subjects are not given the opportunity to buy or sell during these trials. The idea behind this is to let subjects 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 time. 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 buy a stock (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. We now explain how this price is determined.

The price path of each stock is governed by a two-state Markov chain with a good state and a bad state. The Markov chain for each stock is independent of the Markov chains for the other two stocks. Suppose that, in trial t, there is a price update about stock i. If stock i is in the good state at that time, its price increases with probability 0.7 and decreases with probability 0.3. Conversely, if it is in the bad state at that time, its price increases with probability 0.3 and decreases with probability 0.7. 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 a state 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 the trial remains the same as in trial t-1 with probability 0.8, but switches with probability 0.2. In mathematical terms, if $s_{i,t} \in \{\text{good}, \text{bad}\}$ 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 three stocks are not revealed to the subjects: they have to infer them from the observed price paths. To ease comparison of trading performance across subjects, the same set of realized prices was used for all subjects.

A key aspect of our design is that, conditional on the information available to subjects, each of the stocks exhibits short-term price continuation. If a stock performed well 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 positive.

At the end of the first session, we liquidate subjects' holdings of the three stocks and record the cash value of their position. As noted above, the second session is identical in structure to the first. Before the start of the second session – in other words, just before trial 109 – we again randomly assign a state to each of the three stocks, reset each share price to \$100, give each subject a fresh \$350 and require them to immediately buy a share of each of the three stocks. The first nine trials of the second session consist only of a price update, while the remaining trials have both a price update part and a trading decision part. At the end of the second session, we again liquidate subjects' asset holdings and record the value of the cash proceeds.

We give subjects a financial incentive to maximize the final value of their portfolio at the end of each session Specifically, if the total value of a subject's cash and risky asset holdings at the end of session 1 is \$X, in experimental currency, and the total value of his cash and risky asset holdings at the end of session 2 is \$Y, again in experimental currency, then his take-home pay in *actual* dollars is 15 + (X+Y)/24. In other words, we average X and Y to get (X+Y)/2, convert the experimental currency to actual dollars using a 12:1 exchange rate, and add a \$5 show-up fee. Average total earnings were \$30.20. Earnings (not including the show-up fee) ranged from \$19.14 to \$29.14 and the standard deviation of earnings was \$2.58.

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.

Fifty-eight Caltech subjects participated in the experiment, and each subject was randomly assigned to one of three different experimental conditions. Twenty subjects were assigned to the *control* condition where the price and cost basis (if the asset was owned) was displayed on both the price update and trading screens. Eighteen subjects were assigned to the *cost basis* treatment, which was identical to the *control* condition, except that the cost basis was removed from the price update and trading screens. Finally, twenty subjects were assigned to the *forecast* treatment, which was identical to the *control* condition, except that the expected next period price change, (as computed by a Bayesian agent) was displayed on both the price update and trading screens. We refer to this variable as the *optimal forecast,* and the derivation of this variable is given in the next section. Figure 1 depicts each of the two screens for the *control* condition and the two treatment conditions.

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 beginning the first sessions, subjects engaged in a practice session of 25 trials to familiarize themselves with the market software.

B. Expected Value Investors

In the environment we have just described, the optimal strategy for a risk-neutral Bayesian investor who is maximizing the expected value of his take-home earnings – from now on, we refer to such an investor as an "expected value" 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. An expected value investor will therefore only be concerned with the expected return on a stock, and will not pay attention to previous performance.

Formally, let $p_{i,t}$ be the price of stock i in trial t, after any price update about the stock, and let $q_{i,t} = Pr(s_{i,t} = good \mid p_{i,t}, p_{i,t-1}, p_{i,t,\ldots}, p_{i,1})$ 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 indicated a price increase for the stock in question; and -1 if the price update indicated 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:

$$q_{i,t}(q_{i,t-1}, z_t) = \frac{\Pr(z_t \mid s_{i,t} = good) * \Pr(s_{i,t} = good \mid q_{i,t-1})}{q_{i,t-1}\Pr(z_t \mid s_{i,t-1} = good) + (1 - q_{i,t-1})\Pr(z_t \mid s_{i,t-1} = bad)}$$

$$= \frac{(0.5 + 0.2z_t) * [0.8 * q_{i,t-1} + 0.2 * (1 - q_{i,t-1})]}{q_{i,t-1}[0.8 * (0.5 + 0.2z_t) + 0.2 * (0.5 - 0.2z_t)] + (1 - q_{i,t-1})[0.2 * (0.5 + 0.2z_t) + 0.8 * (0.5 - 0.2z_t)]}$$

As noted above, 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. We can also compute the expected next period price change as a function of $q_{i,t}$, which is given by

$$E_t\left[\Delta p_{i,t+1} \middle| q_{i,t}, \Delta p_{i,t+1} \neq 0\right] = 0.6\left(2q_{i,t} - 1\right)$$

This variable is the *optimal forecast.* For an expected value investor, the expected utility of selling stock *i* is then $0\text{-}0.6(2q_{i,t}\text{-}1) = 0.6(1\text{-}2q_{i,t})$. We refer to this quantity as the net present value of selling stock *i* in trial *t*, $NEV_{i,t.}$

C. Realization Utility Investors

As described in the introduction, several studies haves documented the empirical regularity that investors have a greater propensity to sell stocks that have risen in price since purchase, relative to those that have fallen in price since purchase; in other words, investors exhibit a disposition effect (Shefrin and Statman (1985); Odean (1998); Frazzini (2006)). While this type of behavior is hard to reconcile with a rational model of informed trading, a recent theoretical model based on the assumption that investors derive utility directly from the *act of selling* stocks, in addition to standard sources of utility based on consumption, can help explain the disposition effect. This theory, called realization utility, assumes that investors derive a positive flow of utility precisely at the moment when selling a stock that has increased in value since purchase, and this utility is proportional to the capital gain on the stock (Barberis and Xiong (2009); Barberis and Xiong (2012)). The psychological intuition behind this theory is that investors may think about investing in stocks as a series of episodes characterized by the name of the stock, the purchase price, and the sale price. When investors sell a stock at a gain, they are creating a positive memory of their investment experience, and they feel good about this. Conversely, selling a stock at a loss creates a negative episode, one which may conjure negative emotions of regret and low self-esteem. Because the variable that generates realization utility is the capital gain, investors who trade to maximize realization utility will pay attention to past performance via the capital gain.

Recent experimental work using brain imaging has confirmed the presence of such realization utility signals (Frydman et al. (2011)). In particular, these authors allowed subjects to trade in an experimental stock market while their brain activity was monitored. Subjects were found to have a strong disposition effect, and an area of the brain known to hedonic impact exhibited a spike in neural activity precisely at the moment when they realized a stock with a capital gain.

If subjects in our experiment have realization utility preferences and they discount the future at a high rate, then we expect them to sell stocks with capital gains, and hold stocks with capital losses, more often than an expected value investor will. In essence, subjects with "pure" realization utility preferences will only focus attention on past performance through the capital gain, and they will not pay attention to expected returns at the time of the sell/hold decisions. In contrast, an expected value investor will only focus on expected returns, and will not pay any attention to his capital gains.

D. Hypotheses

Our first hypothesis is that investors are not simply one type or the other, but in fact have a *hybrid* preference structure over both realization utility and expected utility of future returns. That is, investors have preferences over realizing capital gains and they also have standard preferences over final consumption. Such an investor will compute the sell/hold decision as a function of both the capital gain and the NEV. One way to model the probability of selling is to assume it is a logistic function of both variables:

$$\Pr\left(Sell_{i,t}\right) = \alpha + \beta_1 NEV_{i,t} + \beta_2\left(p_{i,t} - c\right) + \varepsilon$$

Our hypothesis is that investors pay attention to both past performance and expected returns when computing the sell/hold decision, which we formulate as follows:

Hypothesis 1: Investors use both past performance and expected future returns when deciding whether to sell, which leads to: $\beta_1 > 0$ and $\beta_2 > 0$.

The purpose of our experiment is not only to test this hypothesis, but moreover, to test whether β_1 and β_2 can be manipulated by introducing cues which trigger the desire to experience realization utility. In particular, we run two treatments in which we test whether $\frac{\beta_1}{\beta_2}$ increases relative to the control condition. In the *cost basis* treatment we remove the cost basis from the decision screen, which decreases the salience of the net capital gain, thus shifting attention away from realization utility, and thus inducing a higher $\frac{\beta_1}{\beta_2}$. In the *forecast* treatment, we prominently display the optimal forecast of the price change, which should shift attention towards expected returns, and thus induces a higher $\frac{\beta_1}{\beta_2}$.

Hypothesis 2a: Removing the cost basis from the trading screen will increase the relative weight that subjects attach to expected returns: $\frac{\beta_1}{\beta_2}_{control} < \frac{\beta_1}{\beta_2}_{cost\ basis\ treatment}$.

Hypothesis 2b: Displaying the optimal forecast on the trading screen will increase the relative weight that subjects attach to expected returns: $\frac{\beta_1}{\beta_2}_{control} < \frac{\beta_1}{\beta_2}_{forecast\ treatment}$

II. Results

A. Treatment effects on disposition effect size

We now describe our method for calculating the disposition effect for each subject and then test for differences in this statistic across conditions. We follow a similar methodology to that of (Odean (1998)), where every time a subject is offered the opportunity to sell a stock, we classify this decision into one of four mutually exclusive categories: "realized gains", "realized losses", "paper gains" or "paper losses". A decision classified as a realized gain (realized loss) if the market price of the stock is above (below) the purchase price, and the subject decides to sell the stock. A decision is classified as a paper gain (paper loss) if the market price of the stock is above (below) the purchase price, and the subject decides not to sell the stock. For each subject, we count the number of realized gains, realized losses, paper gains, and paper losses over the course of both experimental sessions and compute the Proportion of Gains Realized (PGR) and the Proportion of Losses Realized (PLR):

$$PGR = \frac{\text{\# of realized gains}}{\text{\# of realized gains + \# of paper gains}}$$

$$PLR = \frac{\text{\# of realized losses}}{\text{\# of realized losses + \# of paper losses}}$$

For each subject, we measure the disposition effect using the difference between PGR and PLR. When the difference is positive (negative), the subject exhibits (the opposite of) a disposition effect.

Figure 2 shows that subjects exhibit an average disposition effect that is significantly greater than the optimal level of -55% in all three conditions (p<0.001 for each condition.) Both the *forecast* treatment and *cost basis* treatment significantly reduce the disposition effect relative to the control condition. By removing the cost basis from the price update screen and the decision screen *(cost basis treatment)*, the average disposition effect decreases from 11.9% to -8.7% (p=0.01). By displaying the optimal forecast on both screens (*forecast treatment)*, the average disposition effect decreases from 11.9% to -16.9% (p=0.014). While both treatment conditions reduce the severity of the disposition effect, the disposition effect is still significantly greater than the optimal level of -55% in both treatments.

Figure 3 shows the number of decisions aggregated across conditions when the subject either sells a stock at a gain, or holds a stock at loss. A realized gain is optimal when the NEV is negative, as the expected return is negative and so the investor should sell. A paper loss is optimal when the NEV is positive, as the expected return is positive and so the investor should hold on to the stock. The blue bars denote those decisions where the subject behaves according to the optimal strategy, and the red bars show those decisions where the subject fails to follow the optimal strategy. It is clear that both the forecast treatment and cost basis treatment substantially reduce the number of paper losses, suggesting that subjects may be paying less attention to the capital loss they have accrued, and instead paying more attention to the negative expected return from holding a stock with a capital loss. It is also interesting to note that in the forecast treatment, over half of all decisions to sell capital gains or hold capital losses are suboptimal, despite the optimal forecast being expliclty displayed on the trading screen.

Figure 4 shows the disaggregated subject data by condition. In the control condition, 45% of subjects (9 of 20) exhibited a disposition effect significantly above 0%. In contrast, only 11% of subjects in the cost basis treatment (2 of 18) exhibited a disposition effect significantly above 0%. Similarly, only 10% of subjects in the forecast treatment (2 of 20) exhibited a disposition effect significantly above 0%[47]. Figure 4 also highlights an important aspect of the experimental design: the disposition effect is a suboptimal behavior, indicated by the negative correlation between earnings and the disposition effect (p<0.001).

---

[47] Within subject standard errors of the disposition effect are computed as in (Odean (1998)).

B. Effect of *cost basis* treatment on sell decision

Our goal in this section is to investigate the mechanism through which the absence of the cost basis leads to a smaller disposition effect, and hence, a higher rate of optimal behavior. Our model says that when subjects are offered the opportunity to sell a stock, the decision is computed by assessing the tradeoff between the accrued capital gain that generates realization utility (backwards looking) and the NPV of selling (forward looking). It follows that if this causes subjects to place a higher relative weight on the NPV of selling, which we manipulate by decreasing the saliency of the cost basis, then the disposition effect should decrease. We test this by estimating a logistic regression of the sell decision on the capital gain and the NPV of selling, allowing for different sensitivities to each of these variables in the *cost basis* condition.

Panel A of table 1 indicates that the effect of the capital gain in the control condition is significantly positive (p=0.049), but the effect of the NPV is not significantly different from zero (p=0.32). There is also no significant marginal effect of the *cost basis* treatment on the capital gain in computing the sell decision (p=0.54). Critically, we find that the NPV is significantly more predictive of the sell decision in the *cost basis* treatment than in the control condition (p=0.05). This suggests that subjects pay more relative attention to expected returns, not because they pay less attention to the capital gain, but instead because they pay more attention to the NPV of selling.

This is important because it rules out the following alternative hypothesis: subjects would like to cater to their realization utility preferences, but are unable to do so because of imperfect recall of the cost basis. If this alternative theory held, then the effect of the treatment on the disposition effect would occur only through a diminished weight on the cost basis. Panel B of table 1 provides an alternative specification of the model that splits the capital gain into its price and cost components, and indicates that there is no differential effect of the cost basis in the treatment condition. Instead, our main results are robust to this specification, and the key interaction effect is stronger (p=0.035).

We also find that the cost basis does not significantly predict the sell decision in the control condition, although the effect is in the right direction. This is likely attributed to a noisier estimate of this coefficient because of small variation in the cost basis over time. For a given holding period, the price will, on average, change every third trial, but the cost basis will remain the same. Therefore, we have variation in the cost basis only across different *holding periods*, whereas we have variation in the price also across trials.

C. Effect of *forecast* treatment on sell decision

We now report results from an analysis designed to understand the mechanism through which the *forecast* treatment lowers the magnitude of the disposition effect. In essence, we replicate the regression in the above section, but use data only from the control group and from those participants in the *forecast* treatment. We hypothesize that displaying the optimal forecast should shift attention away from realization utility preferences and towards the forward-looking expected returns. Panel A of table 2 confirms that in the control condition, the capital gain has a significant effect on the propensity to sell (p=0.049), but the NPV of selling does not (p=0.32). As hypothesized, we find that displaying the optimal forecast on the trading screen significantly increases the weight of the NPV (p=0.002), but there is no significant marginal effect on the capital gain (p=0.50). Panel B of table 2 displays an alternative specification where we split the capital gain into its price and cost components, and we find that the price is a significant predictor of the propensity to sell (p=0.03), but the cost basis is not (p=0.30).

III. Discussion

We find that subjects exhibit a preference for realizing capital gains and losses on risky assets they own and they also exhibit a preference for final consumption. Additionally, the weights on these two components of a subject's hybrid preference structure can be manipulated by the saliency of the capital gain and expected return. This is exhibited most starkly in Figure 2, which shows that the disposition effect is much weaker when subjects are exposed to trading screens where the capital gain has a low degree of saliency. This manipulation of behavior is likely due to a shift of attention away from realization utility and towards expected returns, and it is unlikely because subjects cannot readily compute the capital gain due to lack of information about the cost basis. In other words, the effect on investor behavior we see is due to the manner in which information is *displayed,* and not due to the information that is *known* to the subject.

Our ability to change the size of the disposition effect through two simple manipulations of the display of information is interesting because the disposition effect has been found to decrease average trading performance (Odean (1998)). In order to model this empirical fact, we design our experiment so that stronger disposition effects explicitly lead to lower overall earnings (Figure 4). Together, this shows that we are able to systematically manipulate the trading performance of a subject through different types of portfolio displays.

To conclude, we highlight two main implications of our results, one is theoretical, and the other is practical. A theoretical implication of our results concerns the malleability of preferences. We show that saliency and attention to specific objects in a portfolio has the potential to change the relative weighting of preferences and therefore impact investor behavior (Libby et al. (2002); Hirshleifer and Teoh (2003)). In particular, we show that realization utility preferences may not be fixed and "hard-wired" but are subject to environmental cues that can trigger a desire to realize stocks with capital gains. Hence, a richer theoretical analysis of investor behavior incorporating attention as a state variable may prove useful in future research.

On the practical side, regulators may use our results as a policy tool in order to influence investor behavior. For example, investors with realization utility preferences will trade to realize capital gains and hold stocks with capital losses, which leads to a disposition effect that can lower overall trading performance. If regulators stipulate that brokerage houses decrease the saliency of the capital gain by removing the cost basis from the regular financial statements (as in our cost basis treatment), this would likely attenuate the disposition effect, and could increase individual investor trading performance.

Interestingly, the US government recently enacted a new cost basis legislation in January 2011 which effectively makes the cost basis *more* salient and may have the unintended effect of shifting investors' attention towards realization utility. Specifically, this legislation mandates that investors must decide, at the time of trading, which cost basis method they will use when reporting capital gains for tax purposes. Previously, investors decided this method *after* the trading decision, and so this legislation effectively increases the saliency of the cost basis during the decision phase. Our experimental results suggest that this may have a systematic and detrimental effect on the trading behavior of some individual investors.

This legislation introduces a natural experiment to test the effect of modulating attention to past performance on trading behavior. In particular, a testable implication of our theory of attention modulation is that the new cost basis legislation should induce investors to pay more attention to past performance, and it should lead to a higher average disposition effect among individual investors starting in January 2011. We leave this to future empirical research of the effect of government policy on individual investor behavior.

References

Barber, Brad M., and Terrance Odean, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *Review of Financial Studies* 21, 785-818.

Barberis, Nicholas, and Wei Xiong, 2009, What drives the disposition effect? An analysis of a long-standing preference-based explanation, *Journal of Finance* 64, 751-784.

Barberis, Nicholas, and Wei Xiong, 2012, Realization utility, *Journal of Financial Economics* 104.

Cohen, Lauren, and Andrea Frazzini, 2008, Economic links and predictable returns, *The Journal of Finance* 63, 1977-2011.

Constantinides, George M., 1983, Capital market equilibrium with personal tax, *Econometrica* 51.

Da, Z., J. Engelberg, and P. Gao, 2011, In search of attention, *Journal of Finance* 66.

Dellavigna, Stefano, and Joshua M. Pollet, 2009, Investor inattention and friday earnings announcements, *The Journal of Finance* 64, 709-749.

Duffie, Darrell, 2010, Presidential address: Asset price dynamics with slow-moving capital, *The Journal of Finance* 65, 1237-1267.

Feng, Lei, and Mark S. Seasholes, 2005, Do investor sophistication and trading experience eliminate behavioral biases in financial markets?, *Review of Finance* 9, 305-351.

Frazzini, Andrea, 2006, The disposition effect and underreaction to news, *Journal of Finance* 61, 2017-2046.

Frydman, Cary, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel, 2011, Testing theories of investor behavior using neural data, *Caltech Working Paper*.

Genesove, David, and Christopher Mayer, 2001, Loss aversion and seller behavior: Evidence from the housing market, *Quarterly Journal of Economics* 116, 1233-1260.

Grinblatt, Mark, and Matti Keloharju, 2001, What makes investors trade?, *Journal of Finance* 56, 589-616.

Hirshleifer, David, and Siew Hong Teoh, 2003, Limited attention, information disclosure, and financial reporting, *Journal of Accounting and Economics* 36, 337-386.

Hong, Harrison, and Jeremy C. Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *The Journal of Finance* 54, 2143-2184.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65-91.

Jin, Li, and Anna Scherbina, 2011, Inheriting losers, *Review of Financial Studies* 24.

Kaustia, Markku, 2010, Prospect theory and the disposition effect, *Journal of Financial and Quantitative Analysis* 45, 1-36.

Libby, Robert, Robert Bloomfield, and Mark W. Nelson, 2002, Experimental research in financial accounting, *Accounting, Organizations and Society* 27, 775-810.

Odean, Terrance, 1998, Are investors reluctant to realize their losses?, *Journal of Finance* 53, 1775-1798.

Peng, Lin, and Wei Xiong, 2006, Investor attention, overconfidence and category learning, *Journal of Financial Economics* 80, 563-602.

Shefrin, Hersh, 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.

Weber, Martin, and Colin F. Camerer, 1998, The disposition effect in securities trading: An experimental analysis, *Journal of Economic Behavior & Organization* 33, 167-184.

Table 1. Panel A shows results from a logistic regression of the sell decision on the capital gain and NPV of selling. C2 denotes a dummy variable indicating that the subject was a participant in the *cost basis* treatment. Only subjects from the control condition and *cost basis* treatment are included in the model, and all standard errors are clustered at the subject level. Panel B provides another specification of the model that splits the capital gain into its price and cost components.

	Panel A		Panel B	
	coefficient	*p-val*	*coefficient*	*p-val*
NEV	**0.157**	0.319	**0.168**	0.299
Capital Gain	**0.006**	0.049*		
Price			**0.006**	0.033*
Cost			**-0.004**	0.297
C2	**0.621**	0.002**	**0.024**	0.95
NEV * C2	**0.523**	0.05*	**0.57**	0.035*
Capital Gain * C2	**-0.003**	0.535		
Price * C2			**-0.002**	0.664
Cost * C2			**0.007**	0.226
constant	**-2.02**	0.001***	**-2.275**	0.001***
# of obs	4612		4612	

standard errors clustered at subj. level

Table 2. Panel A shows results from a logistic regression of the sell decision on the capital gain and NPV of selling. C3 denotes a dummy variable indicating that the subject was a participant in the *forecast* treatment. Panel B provides another specification of the model which splits the capital gain into its price and cost components.

	Panel A		Panel B	
	coefficient	*p-val*	*coefficient*	*p-val*
NEV	**0.157**	0.319	**0.168**	0.299
Capital Gain	**0.006**	0.049*		
Price			**0.007**	0.032*
Cost			**-0.004**	0.296
C3	**0.809**	0.001***	**0.468**	0.303
NEV * C3	**0.826**	0.002**	**0.858**	0.002**
Capital Gain * C3	**-0.003**	0.501		
Price * C3			**-0.002**	0.565
Cost * C3			**0.006**	0.345
constant	**-2.02**	0.001***	**-2.275**	0.001***
# of obs	4743		4743	

standard errors clustered at subj. level

Figure 1. **Sample screens from control condition and treatments.** Each trial in every condition consisted of a "price update" screen (2 seconds), followed by a 1 second ITI, then followed by a "trading" screen (3 seconds). A) Control condition. Price and cost displayed on both screens. B) *Cost basis* treatment. Cost is removed on both screens. C) *Forecast* treatment. The optimal forecast of the asset-specific next period price change is displayed on both screens.

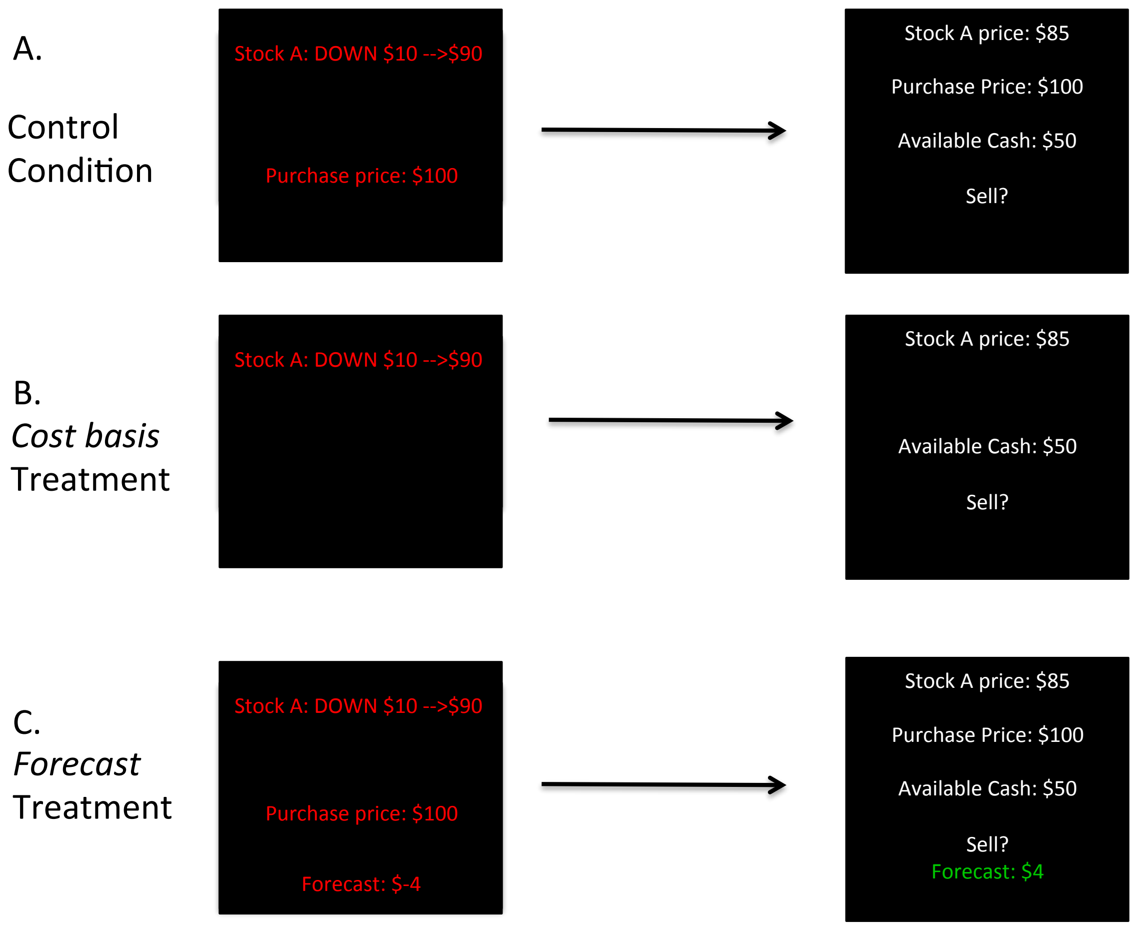

Figure 2. **Average Disposition Effects (PGR-PLR).** Average disposition effects are displayed for the control condition and the two treatment conditions. Both of the treatments significantly reduce the size of the disposition effect relative to the control condition. The optimal level of the disposition effect for an expected value maximizing agent is given in red is -55%, which is significantly lower than the average in any of the three conditions.

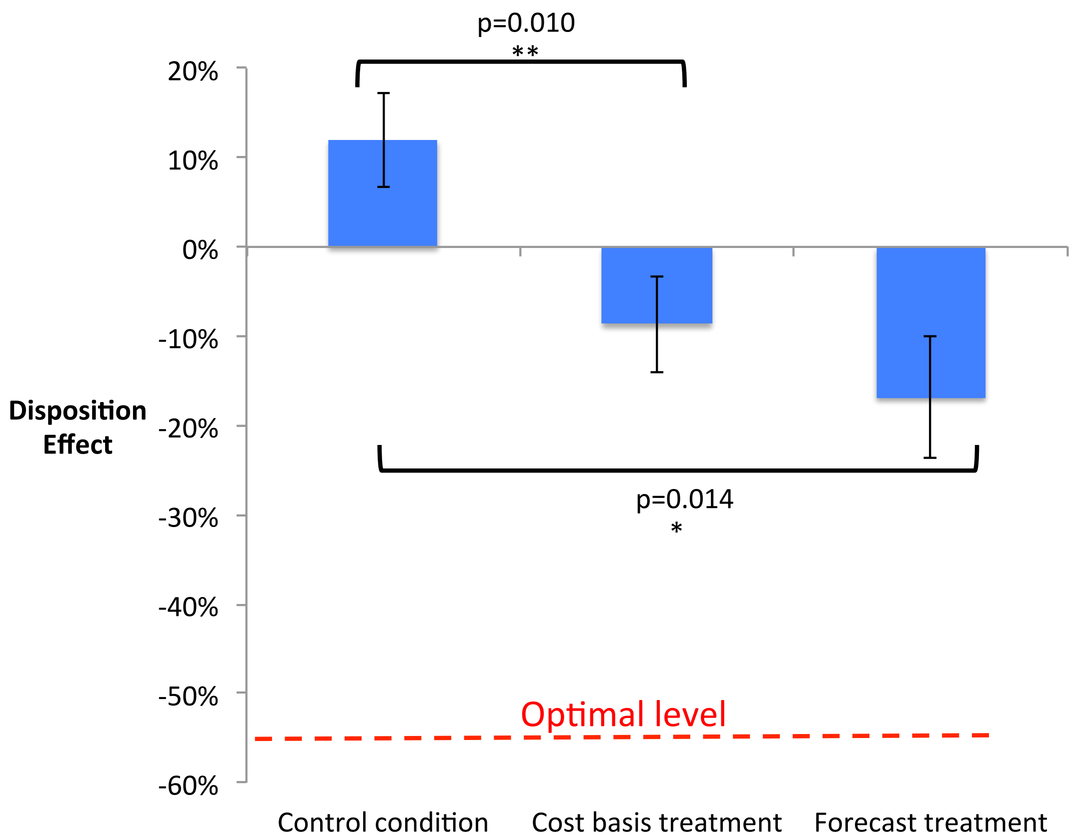

Figure 3. **Decisions to Realize Gains and Hold Losses.** The figure displays the total number of trials, across treatments, where subjects sold a capital gain or held a capital loss. Data are color-coded by whether the decision was optimal or sub-optimal. Realized gains are optimal when the NPV is negative, and paper losses are optimal when the NPV is positive. In each condition, including the forecast treatment where the NPV is displayed on the screen, over half the decisions to realize gains and hold losses are suboptimal.

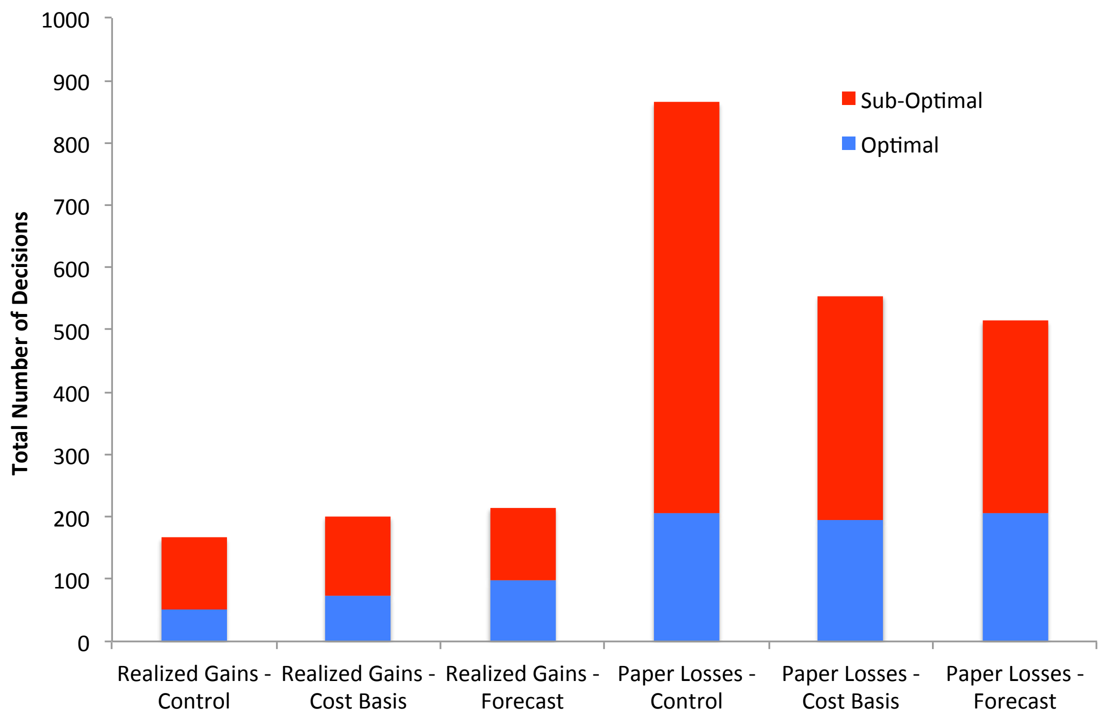

Figure 4. **Disposition effect and earnings.** Each point represents a single subject, and data are presented by condition. The negative relationship between the disposition effect and earnings highlights an important aspect of our experimental design, which is that the disposition effect is a suboptimal behavior. This relationship arises because of the momentum, or positive short-term autocorrelation, in the stock price changes.

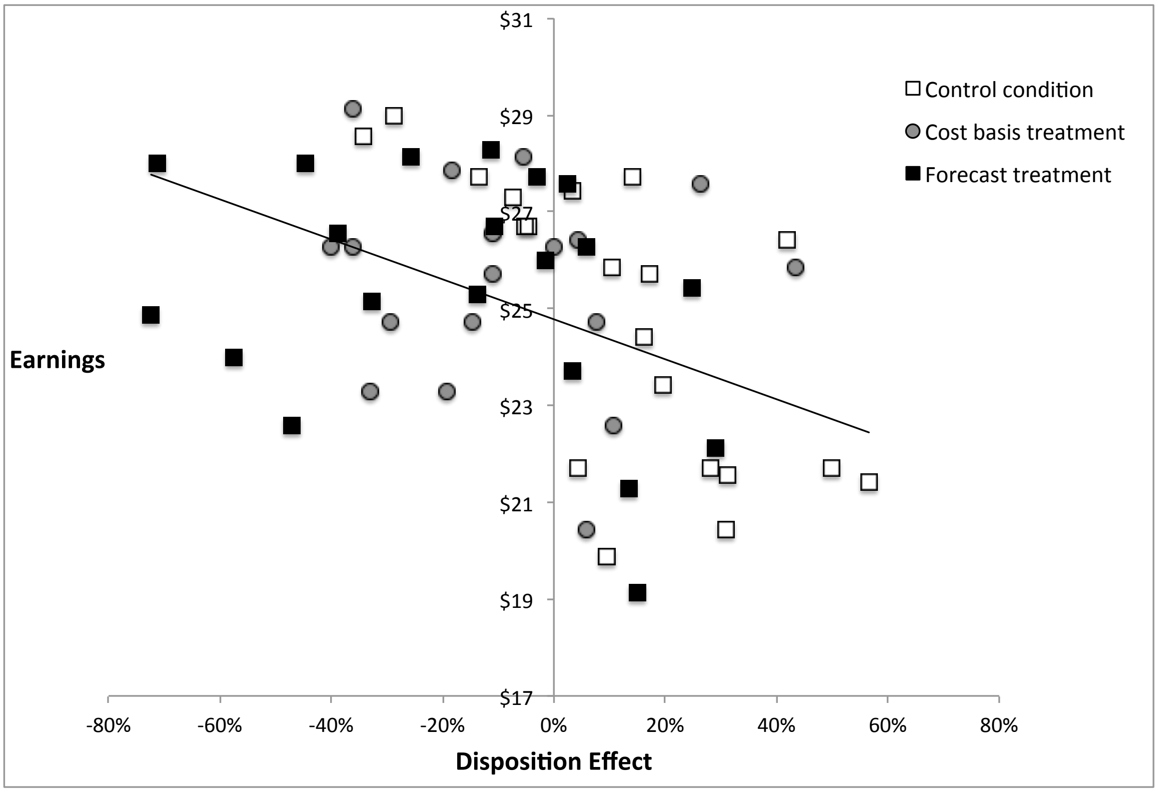

Chapter 4

MAOA-L Carriers are better at making optimal financial decisions under risk

I. INTRODUCTION

Recent research using twin-genetic studies has shown that some of the variation across people in their willingness to take risks can be attributed to heritability [1-3]. Additional work has begun to show associations between specific genes and financial risk-taking behavior [4-7]. Although these studies have been very valuable in identifying particular genes that might be associated with risk-taking behavior, they have not been able to identify the neurocomputational mechanisms that mediate the impact of genes on behavior. This is an important shortcoming since a growing body of work in behavioral neuroscience and neuroeconomics has shown that there might be multiple mechanisms through which genes could affect risk-taking behavior. In particular, genes may affect behavior by changing the value assigned to different risky options [8, 9], or they may affect behavior by changing the way in which the brain adjudicates between the options based on their values [10].

In this paper we shed some light on how genes affect the psychological processes associated with risk-taking behavior by combining tools from behavioral genetics, neuroeconomics, and experimental economics. In particular, we use experimental choice data to estimate well parameterized computational models of financial behavior under risk that allow us to test for the impact of the genes encoding for monoamine oxidase-A (*MAOA*), the serotonin transporter (*5-HTTLPR*), and the dopamine D4 receptor (*DRD4*) have on the two computations described above. Employing a computational model allows us to isolate the underlying psychological mechanisms that contribute to choice heterogeneity across these genes. Consistent with previous results, we find that a specific polymorphism of the *MAOA* gene is associated with an increased propensity to take financial risk. Our computational modeling approach also allowed us to identify the specific mechanism responsible for this increased appetite for risk, which allows for an improved interpretation of previous behavioral genetic results.

We focus on these three genes because they have been the subject of various previous behavioral genetic studies, and because much of the behavioral neuroscience literature points to an important role of the serotonergic, dopaminergic and noradrenergic systems in decision-making [11-15].

Monoamine oxidase-A is an enzyme that regulates the catabolism of monoamines including serotonin, dopamine, norepinphrine, and epinephrine. These monoamines function as neurotransmitters in the central nervous system. Expression of monoamine oxidase-A in the brain has been shown to be influenced by the variable number of tandem repeats (VNTR) in the *MAOA* gene [16]. In particular, carriers of the 3.5 or 4 repeats (MAOA-H) allele exhibit higher expression of the enzyme, whereas carriers of the 2, 3, or 5 repeats (MAOA-L) allele are associated with lower enzymatic expression. The low-activity variant of the *MAOA* gene has been shown to contribute to aggressive and impulsive behavior in mice and humans [17, 18]. At the neuroanatomical level, MAOA-L carriers show lower activity in regulatory prefrontal areas and increased functional connectivity between vmPFC and amygdala regions [19, 20]. In addition, genetic variation in the *MAOA* gene has also been linked to a susceptibility to psychiatric diseases including pathological gambling [21].

The serotonin transporter (*5-HTTLPR*) encodes a protein responsible for the reuptake of serotonin at the synaptic cleft. A short variant has been associated with lower transcriptional efficiency of the gene promoter and higher levels of anxiety, harm-avoidance, and financial risk-aversion [5, 22, 23]. A long variant of the gene is associated with higher transcriptional efficiency and thus higher reuptake of serotonin into the presynaptic neuron. The serotonin transporter is also the target of many anti-depressant drugs that act to inhibit the reuptake of 5HT in order to prolong neuronal firing and serotonin transmission.

The *DRD4* gene influences the function of dopamine D4 receptors, for which a particular repeat sequence leads to functional differences in ligand binding. This gene contains a 48-bp VNTR in exon III which contains 2-11 repeats; carriers of the 7-repeat allele have been shown to require higher levels of dopamine to produce a response of similar magnitude to those without the 7-repeat allele [24]. Behaviorally, carriers of the 7-repeat allele score higher on novelty-seeking personality tests and also exhibit higher rates of pathological gambling [25, 26]. A recent study of an all-male population has also shown that carriers of the 7-repeat allele are willing to take more financial risk in an investment experiment [4], and a similar mixed-gender experiment confirms these results [5].

II. RESULTS

90 male subjects were asked to make choices between 140 different pairs of monetary gambles. Each pair contained a certain option (CO) involving a payout of $\$c$ with 100% probability, and a risky option (RO) involving a gain $\$g$ and a loss $\$l$ with equal probability (see methods for details). Subjects cared about their choices because at the end of the experiment one trial was selected at random and the payouts associated with the selected option were implemented. We failed to obtain successful genotyping on 6 subjects, and 1 additional subject was excluded because ex-post debriefing showed that he did not understand the instructions. As a result, our effective sample size is N=83.

Basic behavioral results. We compared the frequency with which the risky option was chosen by the different genetic groups (Fig. 1). MAOA-L carriers accepted the RO in 41% of trials while MAOA-H carriers accepted 36% of the risky options ($n=83$, $t=1.70$, $p<.046$, one-tailed). However, the *DRD4* and *5-HTTLPR* polymorphisms were not associated with differences in the propensity to accept the risky option. DRD4 7R+

carriers accepted 39% of the risky options, while non-carriers accepted 38% of the time (n=83, t=0.60, p=0.27, one-tailed). Subjects that were homozygous for the short allele of the *5-HTTLPR* gene accepted 37% of risky options, while carriers of the long variant had a 39% acceptance rate (n=83, t=-0.66, p=0.26, one-tailed). We also examined a different categorization of the *5-HTTLPR* genotype and found that carriers of the short allele and those that were homozygous for the long allele both accepted the risky option 38% of the time (n=83, p=0.58, one-tailed). One-tailed test statistics were used for these basic behavioral tests because multiple previous studies have shown that these polymorphisms are associated with increased risk-taking.

Basic computational phenotype. In order to investigate the psychological mechanisms through which the *MAOA* gene affects the propensity to take financial risks we estimated the parameters of a linear prospect theoretic model for each of the subjects based on their choices. The use of this model is justified by the fact that a growing body of behavioral and neuroimaging evidence suggests that most individuals make risky choices by first assigning a value to the different lotteries according to the rules of Prospect Theory (PT) [8, 27-29], and then comparing those values to make a choice.

We assumed that subjects evaluated the gambles using a simple linear version of prospect theory in which the utility of taking the RO is given by

$$U(RO)=pg - \lambda(1-p)l,$$

and the utility of taking the CO is given by

$$U(CO)=CO.$$

Here g denotes the gain associated with the risky option, l denotes the loss, p denotes the probability of the positive payoff, and is a parameter measuring the relative value that the individual assigns to gains and losses. Note that most PT models also assume that probabilities are weighted non-linearly [27, 29]. However, since our study only considers 50-50 gambles, and previous studies have found that the probability distortion at p =0.5

is small [30], we ignore this aspect of the theory.

We assume that the choices are a stochastic function of values that is described by the softmax function:

$$Pr(accept\ RO)=(1+exp(-a\ (U(RO) - U(CO)))^{-1},$$

where a is the inverse-temperature parameter that controls the quality of the decision-making process: when $a=0$ subjects choose both options with equal probability regardless of their associated underlying values; as a increases the probability of choosing the option associated with the largest value increases.

The model has two free parameters for each subject, and a, which we estimated using maximum likelihood (see methods for details). For technical reasons described in the Supplementary Methods, we were able to successfully estimate all of the parameters for the two computational models discussed in the paper for 64 of the 83 subjects. For this reason all of the computational results here and below are limited to this smaller sample. The estimate of λ was 1.520.11 (meanse). The estimate of a was 3.060.60. See Table S2 for a full list of the individual estimates. 41 subjects were significantly loss averse ($\lambda>1$), 21 subjects were loss neutral, and 2 subjects were significantly loss seeking ($\lambda<1$). The average level of λ is similar to that found in other behavioral studies of loss aversion [30, 31].

Note the relationship between the parameters of the model and the psychological processes that affect choice. The coefficient of loss aversion measures the relative value placed on potential gains and losses. When is low, subjects engage in risk-taking behavior by overvaluing gains relative to losses. The opposite is true for high . As a result, this coefficient is a good indicator of the impact that the valuation process has on risk-taking. In contrast, the coefficient a measures the facility with which subjects are able to choose the option with the highest value. Thus, *a* is a good measure of the performance of the comparison or choice processes.

Genetic effects on the valuation process. We first investigated the extent to which the impact of the genetic polymorphisms on risk-taking behavior can be explained by changes in the valuation process. We did this by first regressing the individually estimated loss aversion parameters on each gene polymorphism, including controls for ethnicity and school attended. We found no significant results for *MAOA* (n=64, t=-1.33, p=.19, two-tailed), *DRD4* (n=64, t=-.09, p=.93, two-tailed) or *5-HTTLPR* (n=64, t=.22 p=.82, two-tailed). We then ran a multivariate specification by including all three gene polymorphisms in the same model (which also included the ethnicity and schooling controls) and found no significant effects for *MAOA* (t=-1.33, p=.19, two-tailed), *DRD4* (t=-.03, p=.98, two-tailed), or *5-HTTLPR* (t=.31, p=.76, two-tailed). Finally, because the multiplicative nature of the loss aversion parameter may bias the estimate of the mean, as a robustness check we also ran a version of the multivariate specification using log(λ) as the independent variable. Again, we found no significant effects for *MAOA* (t=-.92, p=.36, two-tailed), *DRD4* (t=-.06, p=.95, two-tailed), or *5-HTTLPR* (t=.43, p=.67, two-tailed).

Genetic effects on the choice process. We then investigated the extent to which the impact of the genetic polymorphisms on risk-taking behavior can be explained by changes in the comparison process. We did this by regressing the individual-fit of the inverse temperature parameters on each gene polymorphism controlling for ethnicity and school population. We found no significant effects for *MAOA* (n=64, t=-.47, p=.64, two-tailed), *DRD4* (n=64, t=.66, p=.51, two-tailed), or *5-HTTLPR* (n=64, t=-.84, p=.40, two-tailed). The multivariate specification including all three polymorphisms and controls also failed to find significant effects for *MAOA* (t=-.40, p=.69, two-tailed), *DRD4* (t=.44, p=.66, two-tailed), and *5-HTTLPR* (t=-.64, p=.53, two-tailed).

Advanced computational phenotype. Since our basic behavioral results show a statistically significant difference in risk-taking behavior for *MAOA* that could not be explained by the simple computational model described above, we decided to complicate the model slightly by allowing for the inverse-temperature parameter to differ for choices in which the gamble had positive net expected utility and those that had negative net expected utility. In particular, we used the same equations to model the valuation of the RO and CO, but we allowed for an asymmetric stochastic choice function:

$$Pr(accept\ RO)=(1+exp(-a^{+}(U(RO)-U(CO)))^{-1}, \quad if \quad U(RO)-U(CO)0$$

$$Pr(accept\ RO)=(1+exp(-a^{-}(U(RO)-U(CO)))^{-1}, \quad if \quad U(RO)-U(CO)< 0$$

The selection of this specification was motivated by the fact that, under the estimates of the basic model, on average subjects rejected a higher percentage (93.3%) of gambles among those with negative expected utility than they accepted (85.4%) with positive expected utility ($n=64$, $t=-4.43$, $p<0.001$, two-tailed). This suggested that subjects might be using a different comparison process when making choices between these two types of risks.

As before, we estimated the individual model parameters using maximum likelihood (see Table S3 and Fig. S3 for a description of the individual fits). The estimate of λ was 1.49.10. The median estimates of a^{+} and a^{-} were 1.97 and 2.25, respectively (see the Supplementary Materials for discussion on technical issues related to the estimation of these two inverse-temperature parameters and a justification for the use of the median statistic to describe the population.) The Bayesian Information Criterion (BIC) value for the unconstrained and constrained ($a^{-}=a^{+}$) models were 5022 and 5058, respectively, indicating a better fit when allowing for asymmetric temperature parameters.

Genetic effects on the comparator process of advantageous and disadvantageous risks. Fig. 2 displays logistic fits to the average group choices for each *MAOA* group, allowing for different slopes in the positive and negative EU domains. The net utility for each gamble was computed using the model fits from the advanced computational model. Note that the logit curve summarizes the performance of the comparator process by relating the net utility of the risky option to the probability that it is chosen as steeper slopes of the logistic curve correspond to higher rates of optimal decision making. The figure shows that there were no differences when the RO had a lower value than the CO, but that there was a systematic difference when this was not the case: MAOA-L carriers chose the optimal action more often than MAOA-H carriers when faced with advantageous risk. A formal statistical test confirmed the difference between the two groups (Fig. 3): MAOA-L carriers accepted the risky option 6.4% more often than MAOA-H carriers ($n=64$, $t=2.49$ $p=0.015$, two-tailed) when the risky option had a positive net expected utility, but there was no significant difference in acceptance rates over the negative EU domain ($n=64$, $t=.51$, $p=0.62$, two tailed), in both cases controlling for ethnicity. Because the latter conclusion was justified by rejection of the null, we subsequently ran a more conservative statistical test by estimating the interaction effect between the *MAOA* genotype and a dummy for positive EU gambles; we found a positive coefficient on the interaction term, consistent with our previous test, although the result was slightly weaker ($n=64$, $t=1.95$, $p=0.056$, two tailed). Note that our statistical test was constructed by integrating under each of the choice curves in the positive and negative EU domains, and thus acts as a non-parametric test of group differences between the a^+ and a^- parameters. Similar analyses for *DRD4* and *5-HTT* did not reveal any significant differences in choice behavior.

III. DISCUSSION

The computational approach used in the paper allowed us to conclude that MAOA-L carriers are more likely to take a financial risk than their MAOA-H counterparts, but only when it is advantageous to do so given their preferences over risk. For disadvantageous gambles there was no difference between the two groups. This suggests that MAOA-L carriers perform better in the case of risky financial decision making since they exhibit an improved ability to select the optimal response when it is advantageous. Contrary to previous findings in the literature [4, 5], we found no significant differences in either gambling tendencies or the computations associated with valuation or choice for the *5-HTTLPR* and *DRD4* genes.

Our results for *MAOA* are consistent with previous related behavioral genetic studies, although our computational approach provides novel insights about the mechanism through which this gene influences risky financial choice. Previous studies have found that MAOA-L carriers are more likely to exhibit aggressive and risky behavior [6, 18, 19, 21]. Contrary to previous discussion in the literature [6, 19, 21], our results show that these behavioral patterns are not necessarily counterproductive [19, 21], since in the case of financial choice these subjects engage in more risky behavior only when it is advantageous to do so. This provides a cautionary tale on the interpretation of previous behavioral results related to *MAOA*, and on the common practice in the literature of relating genes to behavior without specifying and estimating a computational phenotype.

The fact that the *MAOA* gene influences the catabolism of monoamines (such as serotonin, dopamine, norepinphrine, and epinephrine) also allow us to connect our findings with various other strands of the literature. Previous neuroscience studies have shown that humans with higher levels of norepinephrine typically choose the action carrying the highest immediate reward [11, 15]. Our results are consistent with this claim as monoamine oxidase is responsible for the catabolism of norepinephrine, and low activity carriers of *MAOA* will tend to have lower enzymatic activity and thus increased levels of norepinephrine. A recent study which examined the cognitive effects of norepinephrine in mice found that pharmacologically manipulating norepineprhine levels downward resulted in decreased "immediate performance accuracy" [32], which is also consistent with our finding that *MAOA* affects the temperature parameters that control the accuracy of choices.

Monoamine oxidase also plays a role in breaking down dopamine. Therefore, dopaminergic transmission might also play a role in the computational phenotype identified here. Consistent with our findings, a recent study in which in vitro dopaminergic levels were experimentally manipulated through L-DOPA and the impact on optimal choice behavior was measured [33] found that increased dopamine levels lead to more optimal choices in a simple learning task.

The fact that we failed to find behavioral or computational differences between the *5-HTTLPR* and *DRD4* genotypes is also consistent with the previous literature. Some recent studies have found significant effects of both of these genes on financial risk-taking behavior [4, 5], but other studies have failed to replicate these results. For example, a recent fMRI study found a significant effect of *5-HTTLPR* on the framing induced choice biases, but it failed to find a link between *the 5-HTTLPR* polymorphism and financial risk taking [34]. Another study also failed to find any *5-HTTLPR* associations between risk attitudes over the gain and loss domains [35]. The DRD4 gene has also been implicated in impulsive behavior and novelty-seeking in a variety of studies [25, 26], but these results have also not been consistently replicated [36, 37]. In particular, a larger meta-analysis also does not find a significant association between the DRD4 polymorphism and impulsive or risky behavior [38]. One potential reason for our failure to identify a significant effect of *5-HTTLPR* on financial risk taking is limited statistical power for this gene: the distribution of the key polymorphism was unbalanced in our subject population, with only 27% being homozygous for the short allele.

As with any behavioral genetic study, it is also important to pay close attention to the behavioral specificity of the phenotype we define. It is possible that the phenotypic difference we find for the MAOA-L polymorphism may arise from a more general cognitive effect, such as intelligence or numerical ability. We do not have a sufficient battery of controls that can definitively rule out these broader psychological mechanisms nor do we have controls for potential environmental variables (eg, income) that could interact with the MAOA gene to produce the effect. However, one advantage of estimating a computational phenotype is that it allows us to precisely identify the parameter that is driving the heterogeneity in choice within the model. If this heterogeneity were driven by a more general cognitive or environmental variable, then this mechanism should also mediate choice behavior in a manner consistent with our asymmetric result on optimal action selection.

Our results suggest several natural directions for further research. First, future studies should investigate the neurochemical basis of decision-making to understand the quantitative relationship between norepineprhine, dopamine, monoamine oxidase, and optimal choice. Our results provide support for the hypothesis that higher levels of norepinephrine and dopamine correspond to a greater level of action selection optimality, but further research must be conducted to fully understand this relationship [11]. Second, our results indicate the need for future genetic studies to specify a computational phenotype that separates the valuation and choice processes, as subjects with similar preferences might still make very different choices.

IV. MATERIALS AND METHODS

Subjects. 90 male subjects, ages 19-27, participated in the study. Subjects were students at Caltech (59) or at a nearby community college. We restricted our population to males to avoid gender as a confounding factor and to avoid difficulties in the analysis of the MAOA gene (males carry only one allele while females carry two). Subjects self-reported ethnicity was as follows: 53 Caucasian, 13 Latin/Hispanic, 9 Indian, 3 African-American, 3 Asian, and 9 other. However, we failed to obtain successful genotyping on 6 subjects, and 1 additional subject was excluded because ex-post debriefing showed that he did understand the instructions. As a result, our effective sample size is N=83. The study was approved by Caltech's Human Subjects Committee.

Behavioral task. Subjects received twenty-five dollars for participating in the study. They were allowed to risk part of these funds during the following decision-making task. In each trial they were shown a pair of gambles and had to choose one of them. One option involved certain nonnegative payoffs (e.g., gain $0 with probability 100%). We refer to it as the certain option (CO). The other option involved a 50%-50% gamble between a gain and a loss (e.g., winning $7 and losing $4 with equal probability). We refer to it as the risky option (RO). Subjects made decisions in 140 different trials without feedback on a private computer. The order of the choices was randomized within subjects. Table S1 lists the entire set of payoffs used.

Both options were displayed simultaneously on the screen until the subject made a decision. Subjects made a decision using a five point scale: 1=strongly reject the risk option, 2=weakly reject the risky option, 3=indifferent between both options, 4=accept the risky option, 5=strongly accept the risky option. For the purpose of the computational analysis, the responses were collapsed into a binary response (with 5 & 4 coded as accept, 1 and 2 coded as reject, and 3's allocated randomly to the two conditions). To make sure that we did not lose information when collapsing the choice data into binary responses, we estimated an ordered logistic regression and found that 95% confidence intervals for the interior cutpoints (responses 2, 3 and 4) overlapped. This suggests that using the 5-point scale would not add significant information to the behavioral and genetic analyses performed in the paper. Subjects failed to enter a response in 4% of the trials, which were excluded from further analyses. Subjects cared about the choices because one trial was selected at random at the end of the experiment and his choice for that trial was implemented. Average earnings were $28.

Genotyping. Genetic data was collected from each subject using an Oragene DNA OG-500 saliva collection kit. Six subjects were unsuccessfully genotyped for one or more genes and were dropped from all genetic analyses.

5-HTTLPR was identified as follows. The forward primer was labeled with 6FAM-5'-GGC GTTGCC GCT CTG AAT GC-3', the reverse primer was unlabelled 5'-GAG GGA CTGAGC TGG ACA ACC AC-3', which yielded 484-bp (short) and 527-bp (long) fragments. Polymerase chain reaction (PCR) was performed in a total volume of 25 µL, containing 50 ng of DNA; 1µl of each primer(10µM stock); 1.5µl of (25mM)MgCl2; 2%

DMSO (v/v); 2.5 U Amplitaq Gold DNA polymerase (Applied Biosystems, Foster City, California); 2ul of Deaza dNTP (2mM each dATP,dCTP,dTTP,1mMdGTP,1mM deaza dGTP). Cycling conditions consisted of (1) an initial 12 min denaturation at 94°C; (2) 8 cycles with denaturation for 30 sec at 94°C, varied annealing temperatures consisting of 30 sec at 66°C (2 cycles), then 65°C (3 cycles), then 64°C (3 cycles), followed by hybridization for 1 min at 72°C; (3) 35 cycles with an annealing temperature of 63°C and the same denaturation and hybridization parameters; and (4) a final extension for 20 min at 72°C.

MAOA was identified as follows. The forward primer was labeled with VIC-5'-ACAGCCTGACCGTGGAGAAG -3', the reverse primer was unlabelled 5'-GAACGGACGCTCCATTCGGA -3'. Polymerase chain reaction (PCR) was performed in a total volume of 10 µL, containing 25 ng of DNA; 0.5µl of each primer(10µM stock);10x PCR buffer 0.8µl, dNTP 0.8µl,DMSO 0.8 µl,25mM MgCl2 0.8 µl, 0.064 µl of Amplitaq Gold (AppliedBiosystems). Cycling conditions consisted of (1) an initial 12 min denaturation at 95°C; (2) 35cycles of 94°C for 30sec, 59°C for 30sec, 72°C for 2min.

DRD4 was identified as follows. The forward primer was labeled with VIC-5'-AGG ACC CTC ATG GCC TTG -3', the reverse primer was unlabelled 5'-GCG ACT ACG TGG TCT ACT CG -3'. Polymerase chain reaction (PCR) was performed in a total volume of 10 µL, containing 25 ng of DNA; 0.5µl of each primer(10µM stock);Takara LA Taq 0.1 µl,5µl 2x GC Buffer II ,1.6 µl dNTP. Cycling conditions consisted of (1) an initial 1 min denaturation at 95°C; (2) 30 cycles of 94°C for 30sec, 62°C for 30sec, 72°C for 2min; (3) 72°C for 5min. In all cases the PCR products were electrophoresed on an ABI 3730 DNA analyzer (Applied Biosystems) with a LIZ1200 size standard (AppliedBiosystems), and Data collection and analysis used the Genemapper software (Applied Biosystems).

Genotype equilibrium. Allele and genotype frequencies are given in Tables S4 – S7. A Pearson-Chi squared test failed to reject the null hypothesis that the *5-HTTLPR* gene was in Hardy-Weinberg Equilibrium (HWE) in our subject pool (^{2}=0.98, df=1, p>0.32). Since males possess only one allele of the *MAOA* gene, HWE is trivially satisfied. Finally, because of its multiple allele structure [39], we used a Markov Chain Monte Carlo (MCMC) method to test the null hypothesis that *DRD4* was in HWE. The test failed to reject the null hypothesis (p=0.689).

Computational phenotype. The parameters for the two computational models described in the results section were estimated by optimizing the non-linear likelihood function using the Nelder-Mead Simplex Method [40], as implemented in Matlab 2008b. We computed standard errors for the estimated parameters using parametric bootstrapping with a re-sampling size of 500. For each subject we estimated individual parameters from the choice data and then used the estimates to generate a set of 500 pseudosamples of choice data. We then used the same MLE procedure described above to estimate the parameters in each of the pseudosamples. The standard error of the parameter estimate was then estimated by the standard deviation of this set of samples.

We assessed the model fit of the unconstrained computational model by computing the percent of choices correctly predicted for each subject at individually fitted parameter values, which was 88.8% on average.

For technical reasons explained in the Supplementary Materials, we failed to estimate one or more model parameters for 19 out of 83 subjects. The computational results described in the paper only apply to the 64 subjects for which all parameters were estimated successfully.

A potentially important simplification used in the computational models is the linearity of the value function. We tested the robustness of this assumption by estimating a non-linear version of the simple prospect theoretic model given by the following three equations:

1. *(1)* $U(RO) = pg^{\rho} - \lambda(1-p)l^{\rho}$
2. *(2)* $U(CO) = CO^{\rho}$
3. *(3)* $Pr(\text{accept } RO) = (1+\exp(-a\,(U(RO) - U(CO)))^{-1}$.

This model contains an additional parameter that allows for the possibility that value might be a non-linear function of the payoffs. We estimated the model using the same MLE procedure described above. However, due to insufficient concavity of the likelihood function, we failed to successfully estimate parameters for 5 additional subjects that we were able to estimate the model for under the constraint $\rho=1$. Of the remaining subjects, average estimates of λ, ρ and *a* were 1.51, 1.03, and 2.95, respectively. We ran a likelihood ratio test for each individual under the null hypothesis that $\rho=1$, and determined that we could reject a linear value function in 46 of 65 subjects at the 5% significance level. Furthermore, a t-*test* on the distribution of the unconstrained estimates of ρ did not reject the null hypothesis that the average value of ρ in the population is 1 ($p=.29$) Because of the lack of heterogeneity in ρ, and because including this extra parameter did not significantly improve the model fit, we focused the analysis in the paper on the simple and advanced versions of the linear prospect theory model.

Estimation of advanced computational model. The advanced computational model is described by the following four equations:

1) $U(RO)=pg - \lambda(1-p)l$,
2) $U(CO)=CO$.
3) $Pr(\text{accept } RO)=(1+\exp(-a^{+}(U(RO) - U(CO)))^{-1}$, if $U(RO)-U(CO)>=0$
4) $Pr(\text{accept } RO)=(1+\exp(-a^{-}(U(RO) - U(CO)))^{-1}$, if $U(RO)-U(CO)< 0$

The first two equations describe the valuation process and the second pair of equations describes the probability with which subjects choose the option with the highest net expected utility. We used maximum likelihood to estimate the parameter vector θ=(λ, a^+, a^-) for each subject. This required maximizing the following likelihood function:

$$l(\theta \mid y,p) = \sum_{i=1}^{140} y_i \log(F(p,\theta)) + (1-y_i)\log(1-F(p,\theta))$$

where

$$F(p, \theta)=(1+\exp(-a^+(U(RO)-U(CO)))^{-1}, \quad \text{if } U(RO)-U(CO)>=0$$
$$F(p, \theta)=(1+\exp(-a^-(U(RO)-U(CO)))^{-1}, \quad \text{if } U(RO)-U(CO)<0,$$

i indexes the trial number, *y* indicates the response, *p* describes the design matrix of the behavioral task, and θ indicates the parameter vector to be estimated. We used the Nelder-Mead Simplex Method as implemented in Matlab 2008b to obtain point estimates for each parameter.

As described in the methods section, we failed to successfully estimate at least one parameters for 19 out of 83 subjects that comprise our effective sample size. 9 subjects were dropped due to insufficient variation in responses, which makes estimation impossible. 8 subjects were excluded because their behavior was random, in the sense of being unresponsive to the underlying valuations options. 2 were excluded for failing to satisfy the basic "rationality" constraint that when the expected utility of the risky option is higher than the certain option, the risky option should always be accepted. Table S8 describes the sample sizes and explanations for all analyses in the main text.

Estimation problems for randomless choice behavior. We failed to estimate parameters of the advanced computational phenotype for 9 subjects (6 MAOA-H and 3 MAOA-L) due to lack of variation in observed choices. This type of complication arises when subjects always choose the highest value option without any noise, which corresponds to the case $a=\infty$. This makes parameter estimation impossible since the resulting choice behavior can be generated using any sufficiently large temperature parameter *a*. This leads to a flat likelihood function in this range of the temperature parameters that makes maximization of the likelihood function over this range infeasible. This type of complication is more severe for the advanced computational phenotype because estimation will fail if subjects respond without noise in either the positive *or* negative EU domain.

Estimation problems for random choice behavior. On the other end of the spectrum from noiseless choice performance is random behavior. We failed to estimate the parameters of 8 subjects (5 MAOA-H and 3 MAOA-L) due to this problem. In particular, for these subjects the maximum likelihood procedure generated a nonsensical negative estimate for either a^+ or a^-. This problem can arise when subjects' parameters induce valuations that lead to a positive net value for the risky option in only a small fraction of the 140 choice pairs, for which the subject responds sub-optimally within this small set of trials.

About the identification of the temperature parameters: a^+ and a^- **.** The introduction of two temperature parameters in the advanced model makes the estimation problem more difficult than in the basic model. The fundamental problem is illustrated in Fig. S1, which shows that the fraction of trials in which the risky option has a positive net utility decreases rapidly with λ, which makes it difficult to obtain precise estimates of a^+ and a^- . Intuitively, the econometric difficulty arises because the sample size of trials in the positive and negative EU domain is endogenously determined by λ. When λ takes on an extremely high (low) value, the sample size of the positive (negative) EU domain becomes very small, which induces highly imprecise estimates of all computational model parameters. This estimation problem is intensified when subjects respond using either random or purely randomless behavior, as described above.

Example of estimation problems. Here we show that the estimation problems described above can arise even with simulated data in which we know that the underlying computational model applies. Consider a hypothetical subject with λ=2.5, in which case only 23% of trials (33 of 140 trials) will have a positive net RO, $a^-=3$, and a^+ very large, so that she responds with noiseless choice performance in the positive EU domain. We simulated choice data from this hypothetical subject and attempted to estimate parameters using both the basic and advanced computational model. For the basic computational model, the parameters are estimated correctly and the likelihood function is concave in a (Fig S4). However, when estimating the advanced computational phenotype, the maximization algorithm does not converge and terminates the search procedure prematurely at: λ=2.38, $a^+=386$, $a^-=208$. Fig S5 shows that this is because the likelihood function is not concave in a^+; the likelihood surface is flat in the a^+ dimension, and there is a continuum of parameter values that fit the data equally well. This leads to a failed maximization procedure, and an inability to estimate the advanced computational model.

To compare this function with data generated from a subject who does not respond with noiselss choice performance, we generated a data set from another hypothetical subject with λ=2.5 and $a^{+}=a^{-}=3$. Because this hypothetical subject does not respond with noiseless choice performance in the positive EU domain, we are able to successfully estimate the advanced computational model. The likelihood function for this subject is plotted in Fig S6, which shows the function is concave in both dimensions, allowing for successful maximization.

Fig. 1. Choice of the risky option by genetic group. MAOA-L carriers accepted the risky optional significantly more often that MAOA-H carriers: 41.2% vs. 36.3% (p=0.046). Differences in the acceptance rates for the *5-HTTLPR* and *DRD4* polymorphisms were not significant. 46% of the risky options in our design had a positive net expected value.

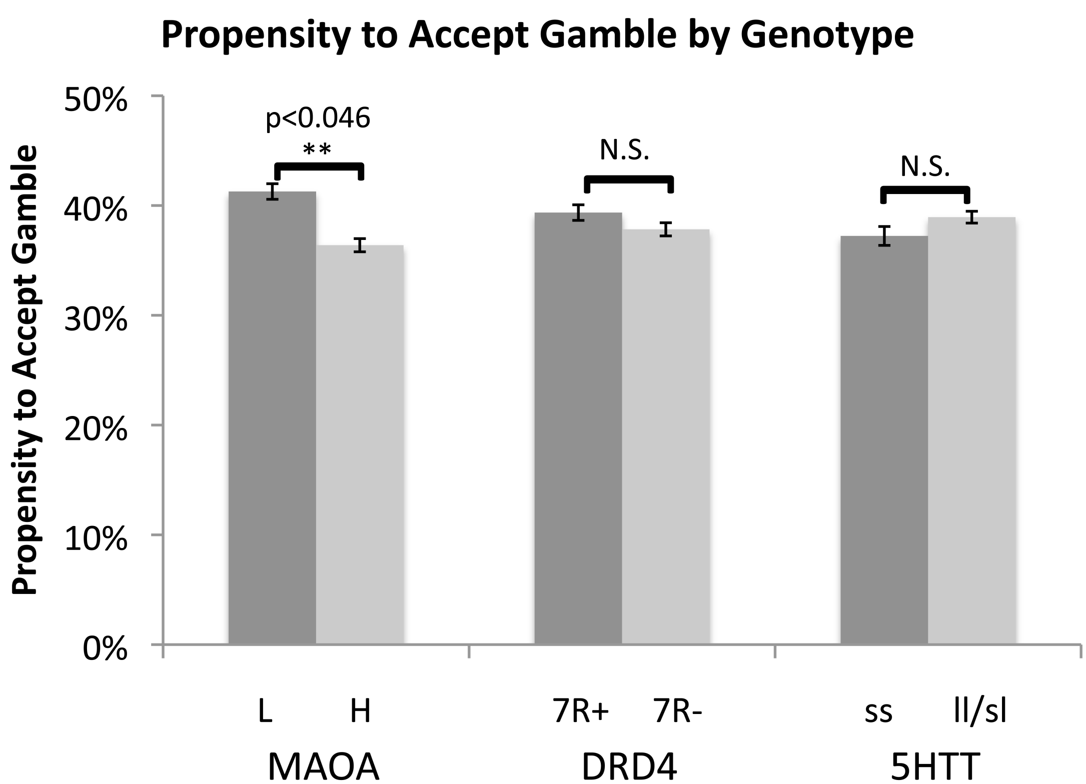

Fig. 2. Propensity to choose the RO as a function of its net expected utility (using individually fitted PT parameters). The single solid black curve in the negative net EU domain indicates there was no difference in acceptance rates across the *MAOA* polymorphism when the net EU was negative. However, the two dashed curves in the positive net EU domain show there was a systematic difference in the propensity to accept the RO: MAOA-L carriers (black) accepted the risky offer significantly more often than MAOA-H carriers (grey). Net EU is partitioned into bins of length 0.5 and the average group acceptance rate within each bin is displayed for MAOA-L and MAOA-H.

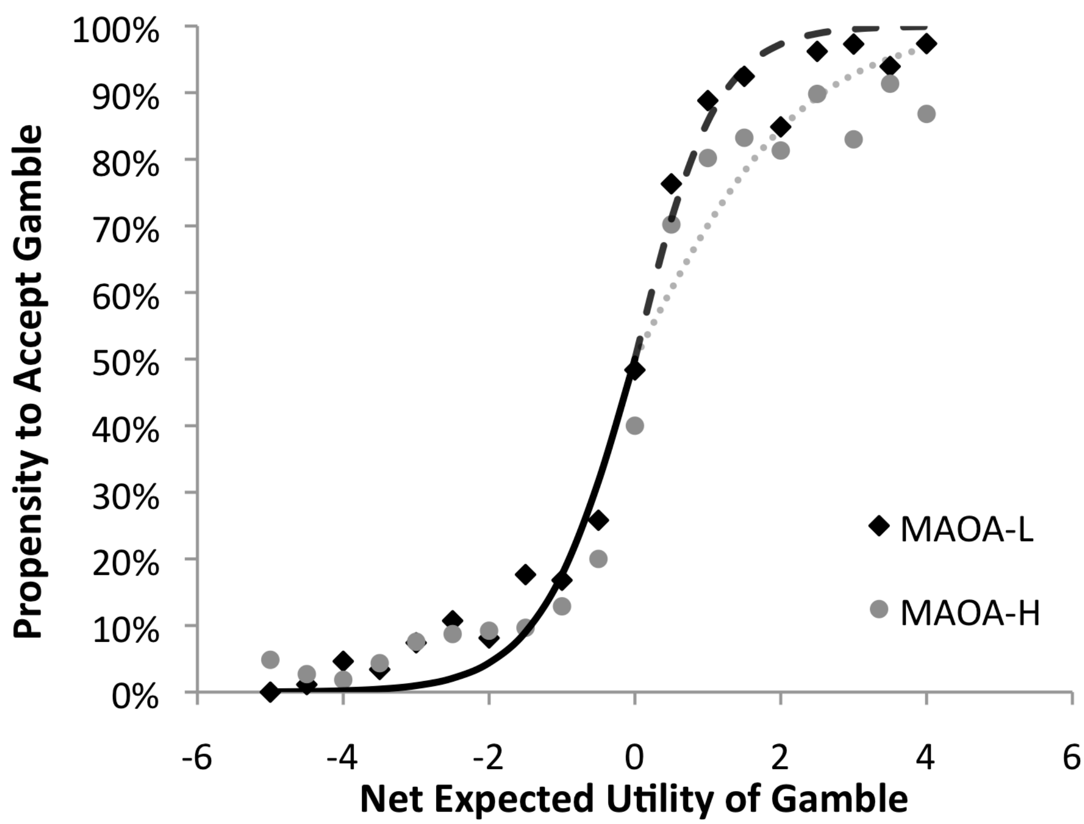

Fig. 3. Propensity to choose the option with highest expected utility as a function of the *MAOA* polymorphism.

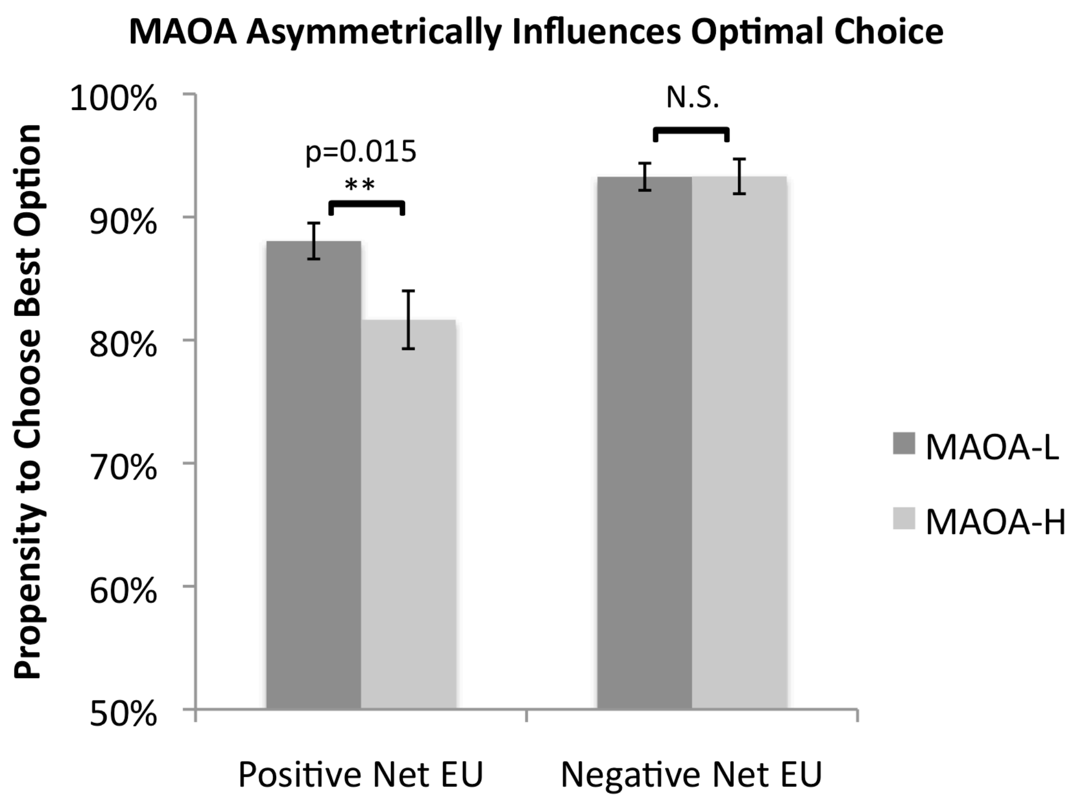

Fig. S1. Number of experimental trials in which the RO had positive and negative net EU as a function of the underlying loss aversion parameter.

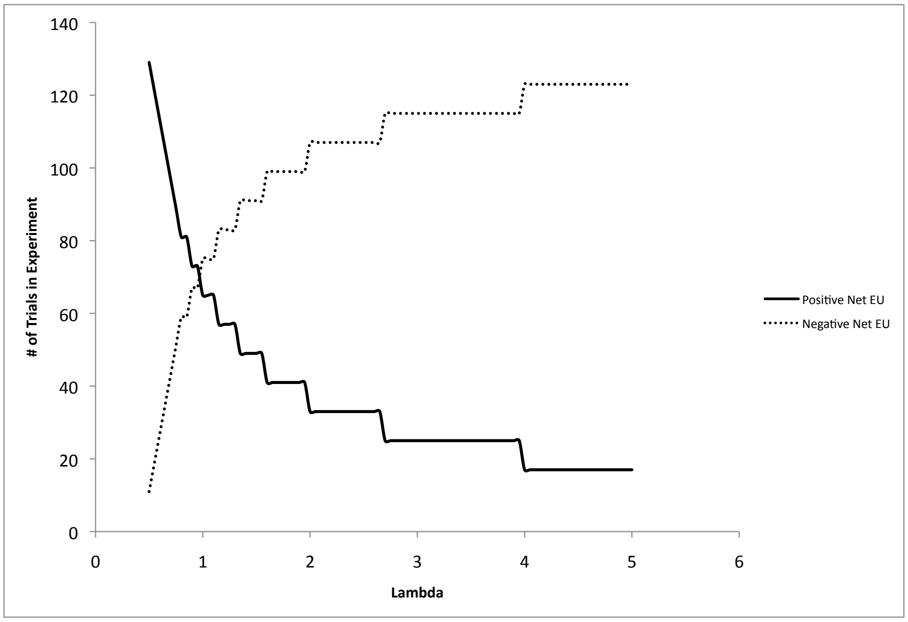

Fig. S2. MLE estimates of the loss aversion parameter under the basic and advanced computational models.

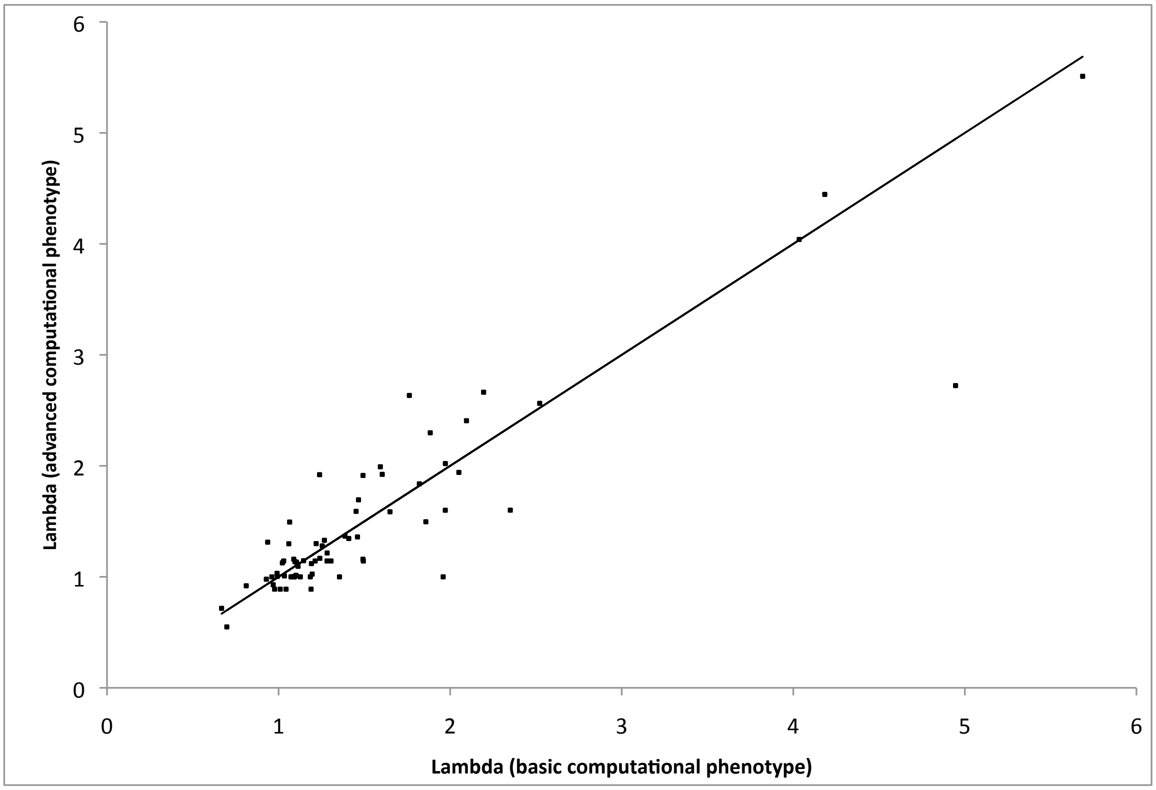

Fig. S3. Distribution of individual loss aversion estimates under the advanced computational phenotype.

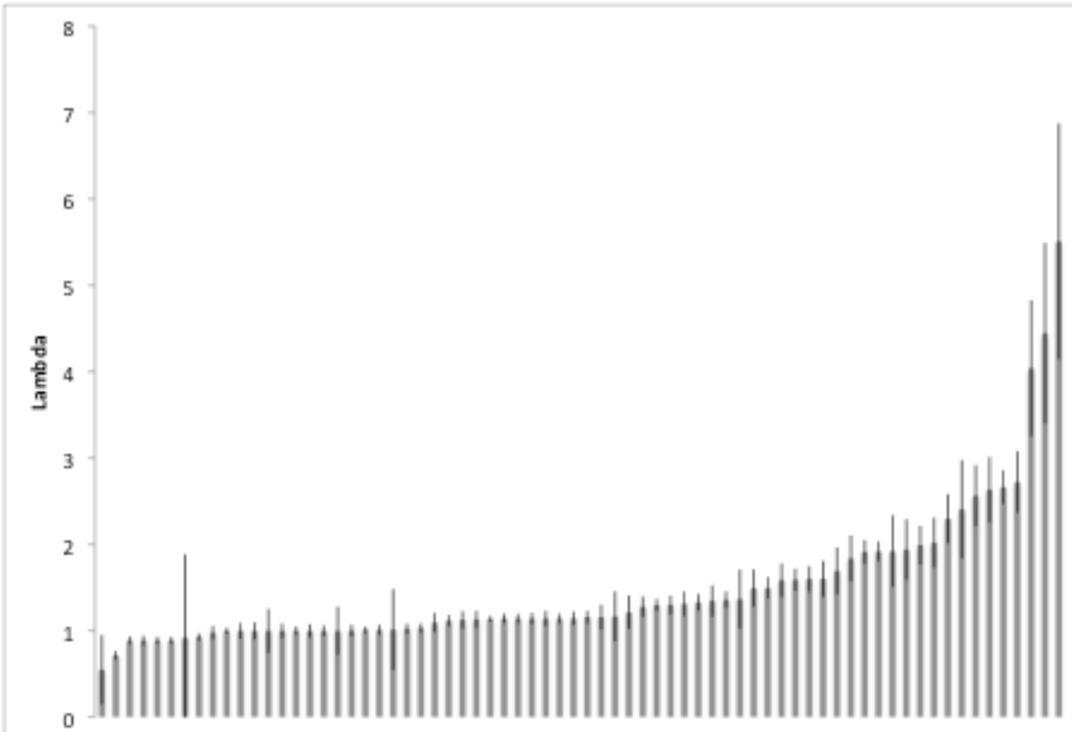

Fig. S4. Log-likelihood function of basic computational model. Choice data is simulated from hypothetical subject with λ=*2.5*, a^{-}=*3*, and noiseless choice performance in positive EU domain. The likelihood function is plotted at λ =*2.5*. The function is concave, allowing for maximization and successful estimation of the basic computational model.

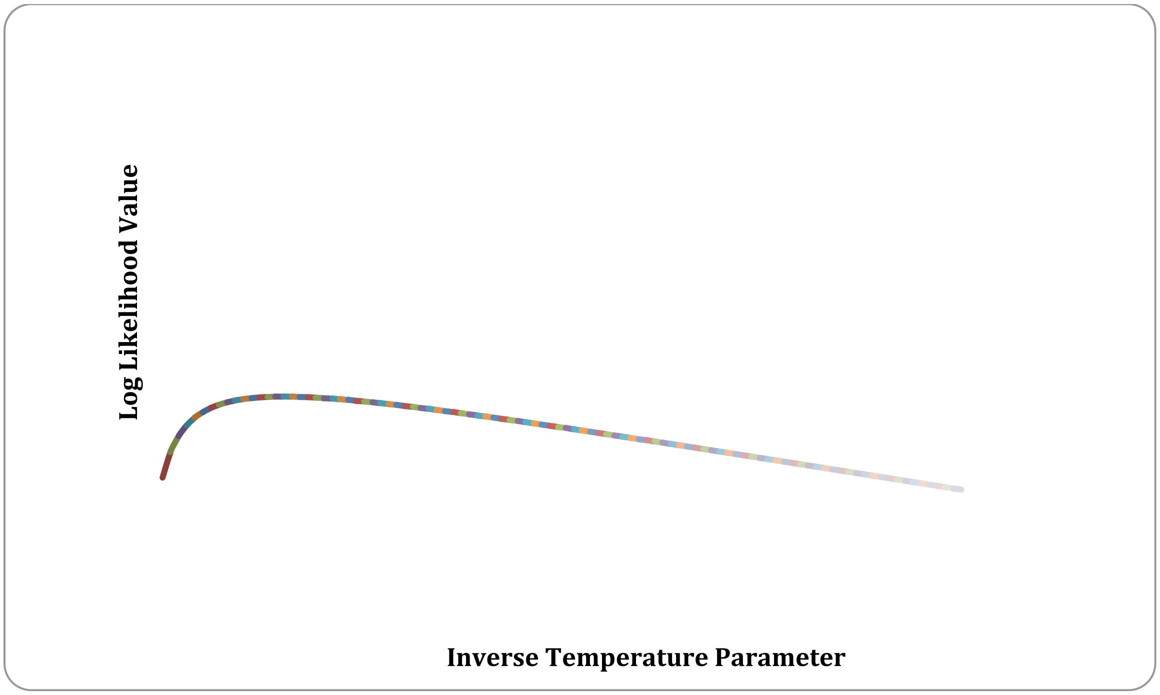

Fig. S5. Log-likelihood function of advanced computational model. Choice data is simulated from hypothetical subject with λ=*2.5*, a^{-}=*3*, and noiseless choice performance in positive EU domain. The likelihood function is plotted at λ =*2.5*. The function is flat in the a^{+} dimension, because of the noiseless choice performance in the positive EU domain. This causes the maximization of the log-likelihood function to fail, and leads to estimation problems for the a^{-} parameter as well.

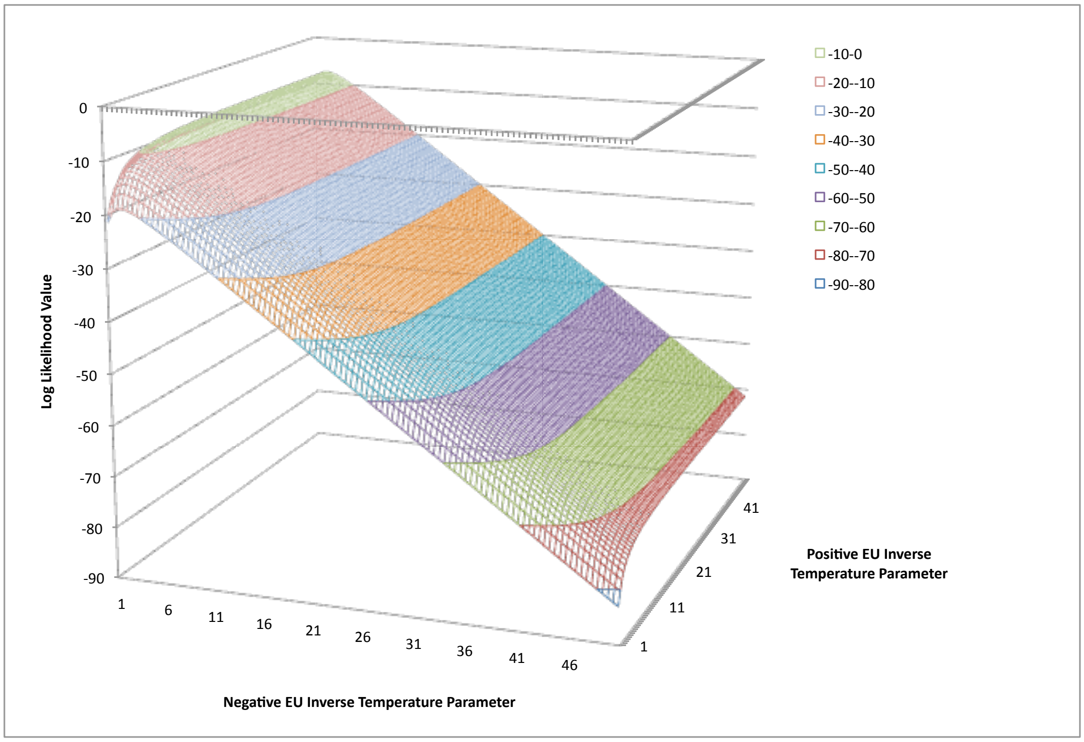

Fig. S6. Log-likelihood function of advanced computational model. Choice data is simulated from hypothetical subject with λ=2.5, a^{+}=a^{-}=3. The likelihood function is plotted at λ =2.5. The function is concave in both the $a^{+}$and a^{-} dimensions, which allows for successful estimation. The region that maximizes the log-likelihood function is depicted in orange.

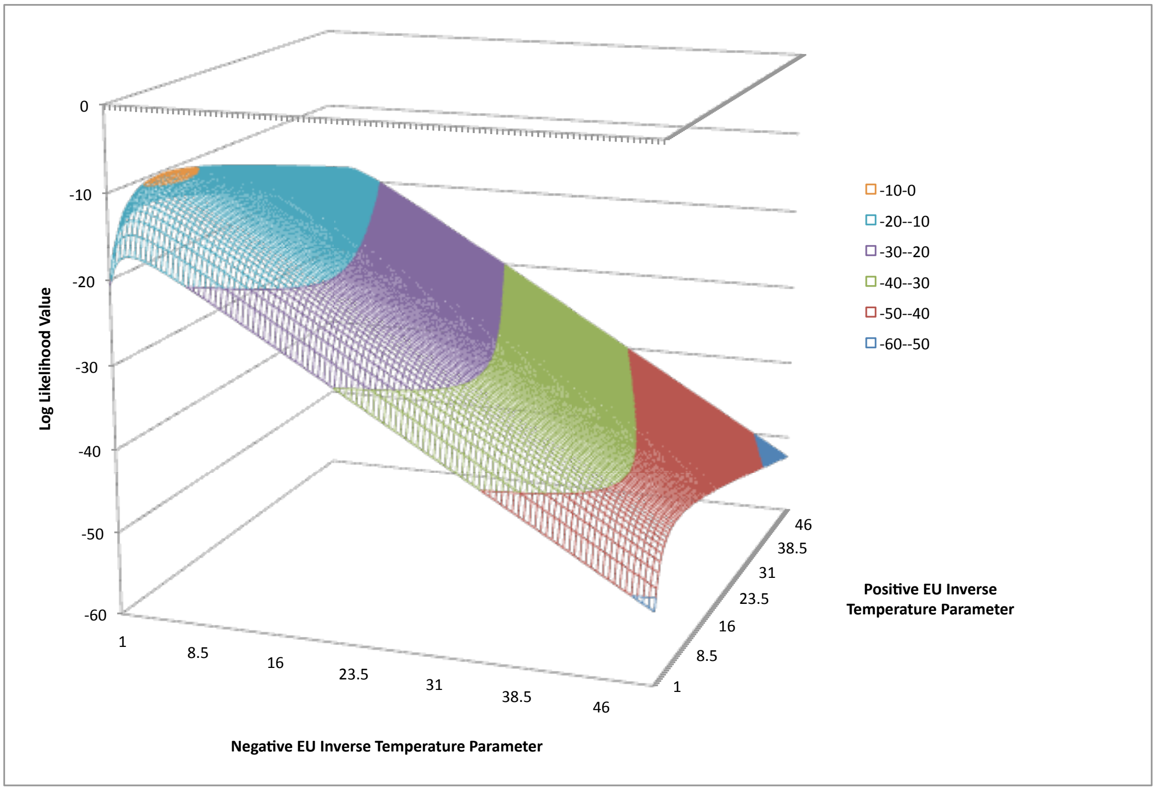

Table S1. Binary choices used in the experiment. CO indicates certain option.

Gain	Loss	CO	Gain	Loss	CO	Gain	Loss	CO
$12.00	-$24.00	$0.00	$2.00	-$3.50	$0.00	$4.00	-$1.50	$0.00
$12.00	-$22.50	$0.00	$4.00	-$5.50	$0.00	$5.00	-$2.50	$0.00
$10.00	-$20.00	$0.00	$6.00	-$7.50	$0.00	$10.00	-$7.50	$0.00
$9.00	-$18.00	$0.00	$12.00	-$13.50	$0.00	$4.00	-$1.00	$0.00
$12.00	-$21.00	$0.00	$2.00	-$3.25	$0.00	$6.00	-$3.00	$0.00
$10.00	-$18.75	$0.00	$5.00	-$6.25	$0.00	$8.00	-$5.00	$0.00
$8.00	-$16.00	$0.00	$10.00	-$11.25	$0.00	$12.00	-$9.00	$0.00
$9.00	-$16.88	$0.00	$9.00	-$10.13	$0.00	$5.00	-$1.88	$0.00
$10.00	-$17.50	$0.00	$2.00	-$3.00	$0.00	$9.00	-$5.63	$0.00
$12.00	-$19.50	$0.00	$4.00	-$5.00	$0.00	$5.00	-$1.25	$0.00
$8.00	-$15.00	$0.00	$8.00	-$9.00	$0.00	$6.00	-$2.25	$0.00
$9.00	-$15.75	$0.00	$2.00	-$2.75	$0.00	$10.00	-$6.25	$0.00
$10.00	-$16.25	$0.00	$6.00	-$6.75	$0.00	$8.00	-$4.00	$0.00
$6.00	-$12.00	$0.00	$5.00	-$5.63	$0.00	$6.00	-$1.50	$0.00
$8.00	-$14.00	$0.00	$2.00	-$2.50	$0.00	$9.00	-$4.50	$0.00
$12.00	-$18.00	$0.00	$4.00	-$4.50	$0.00	$12.00	-$7.50	$0.00
$9.00	-$14.63	$0.00	$2.00	-$2.25	$0.00	$8.00	-$3.00	$0.00
$6.00	-$11.25	$0.00	$2.00	-$2.00	$0.00	$10.00	-$5.00	$0.00
$5.00	-$10.00	$0.00	$4.00	-$4.00	$0.00	$9.00	-$3.38	$0.00
$8.00	-$13.00	$0.00	$5.00	-$5.00	$0.00	$8.00	-$2.00	$0.00
$10.00	-$15.00	$0.00	$6.00	-$6.00	$0.00	$12.00	-$6.00	$0.00
$6.00	-$10.50	$0.00	$8.00	-$8.00	$0.00	$10.00	-$3.75	$0.00
$9.00	-$13.50	$0.00	$9.00	-$9.00	$0.00	$9.00	-$2.25	$0.00
$12.00	-$16.50	$0.00	$10.00	-$10.00	$0.00	$10.00	-$2.50	$0.00
$5.00	-$9.38	$0.00	$12.00	-$12.00	$0.00	$12.00	-$4.50	$0.00
$4.00	-$8.00	$0.00	$2.00	-$1.75	$0.00	$12.00	-$3.00	$0.00
$8.00	-$12.00	$0.00	$2.00	-$1.50	$0.00	$2.00	$0.00	$1.00
$5.00	-$8.75	$0.00	$4.00	-$3.50	$0.00	$3.00	$0.00	$1.00
$6.00	-$9.75	$0.00	$5.00	-$4.38	$0.00	$4.00	$0.00	$2.00
$10.00	-$13.75	$0.00	$2.00	-$1.25	$0.00	$5.00	$0.00	$2.00
$4.00	-$7.50	$0.00	$6.00	-$5.25	$0.00	$7.00	$0.00	$3.00
$9.00	-$12.38	$0.00	$2.00	-$1.00	$0.00	$8.00	$0.00	$3.00
$5.00	-$8.13	$0.00	$4.00	-$3.00	$0.00	$12.00	$0.00	$6.00
$4.00	-$7.00	$0.00	$8.00	-$7.00	$0.00	$12.00	$0.00	$5.00
$6.00	-$9.00	$0.00	$9.00	-$7.88	$0.00	$12.00	$0.00	$4.00
$8.00	-$11.00	$0.00	$2.00	-$0.75	$0.00	$13.00	$0.00	$5.00
$12.00	-$15.00	$0.00	$5.00	-$3.75	$0.00	$13.00	$0.00	$6.00
$4.00	-$6.50	$0.00	$10.00	-$8.75	$0.00	$19.00	$0.00	$8.00
$5.00	-$7.50	$0.00	$2.00	-$0.50	$0.00	$22.00	$0.00	$10.00
$10.00	-$12.50	$0.00	$4.00	-$2.50	$0.00	$23.00	$0.00	$10.00
$6.00	-$8.25	$0.00	$6.00	-$4.50	$0.00	$25.00	$0.00	$9.00
$9.00	-$11.25	$0.00	$12.00	-$10.50	$0.00	$25.00	$0.00	$10.00
$2.00	-$4.00	$0.00	$5.00	-$3.13	$0.00	$26.00	$0.00	$10.00
$4.00	-$6.00	$0.00	$4.00	-$2.00	$0.00	$26.00	$0.00	$12.00
$8.00	-$10.00	$0.00	$8.00	-$6.00	$0.00	$28.00	$0.00	$13.00
$5.00	-$6.88	$0.00	$6.00	-$3.75	$0.00	$30.00	$0.00	$12.00
$2.00	-$3.75	$0.00	$9.00	-$6.75	$0.00			

Table S2. Individual parameter estimates in the basic computational model.

ID	λ	a
1	2.35 ±0.37	0.72 ±0.23
2	1.76 ±0.16	1.16 ±0.29
3	0.99 ±0.04	2.65 ±0.66
4	4.94 ±1.49	0.72 ±0.44
5	1 ±0.04	2.77 ±0.72
6	0.96 ±0.03	6.51 ±2.45
7	0.93 ±0.05	1.49 ±0.3
8	1.08 ±0.03	6.18 ±2.26
9	1.1 ±0.06	1.67 ±0.39
10	0.92 ±0.05	1.74 ±0.34
11	1.04 ±0.06	1.67 ±0.38
12	1.18 ±0.08	1.2 ±0.25
13	5.68 ±1.42	1.16 ±3.39
14	1.24 ±0.11	0.82 ±0.19
15	1.09 ±0.04	7.58 ±2.96
16	1.49 ±0.1	1.5 ±0.37
17	1.08 ±0.07	1.1 ±0.23
18	1.6 ±0.18	0.7 ±0.17
19	1.06 ±0.08	1.04 ±0.22
20	1.02 ±0.06	1.67 ±0.35
21	1.03 ±0.03	10.22 ±3.89
22	1.12 ±0.04	3.01 ±0.83
23	1.28 ±0.07	2.02 ±0.52
24	1.14 ±0.05	2.95 ±0.8
25	1.4 ±0.1	1.18 ±0.28
26	1.45 ±0.07	2.38 ±0.67
27	1.23 ±0.15	0.21 ±0.09
28	0.66 ±0.02	2.49 ±0.52
29	1.3 ±0.06	2.98 ±0.78
30	1.1 ±0.03	6.73 ±2.63
31	1.97 ±0.13	1.98 ±0.62
32	1.19 ±0.06	2.51 ±0.66
33	0.99 ±0.03	5.69 ±2.01
34	2.05 ±0.22	0.97 ±0.24
35	0.97 ±0.04	2.65 ±0.64
36	1.26 ±0.05	3.44 ±1.23
37	1.46 ±0.12	1.09 ±0.24
38	4.18 ±0.62	3.73 ±18.9
39	1.21 ±0.04	3.89 ±1.26
40	1.49 ±0.09	2.07 ±0.54
41	1.85 ±0.16	1.24 ±0.33
42	2.19 ±0.17	1.94 ±0.63
43	1.45 ±0.07	2.77 ±0.89
44	1.09 ±0.06	7.7 ±3.12
45	1.49 ±0.12	1.11 ±0.25
46	1.97 ±0.18	1.37 ±0.38
47	1.07 ±0.03	6.7 ±2.49
48	2.52 ±0.22	1.62 ±0.55
49	1.21 ±0.05	2.71 ±0.68
50	1.05 ±0.05	1.89 ±0.45
51	1.28 ±0.11	0.94 ±0.21
52	1.35 ±0.11	1 ±0.23
53	0.81 ±0.2	0.3 ±0.12
54	1.38 ±0.15	0.67 ±0.16
55	1.11 ±0.07	1.48 ±0.32
56	1.1 ±0.05	0.7 ±0.22
57	1.95 ±0.36	0.48 ±0.15
58	0.69 ±0.01	0.86 ±3.69
59	2.09 ±0.13	0.65 ±0.13
60	1.02 ±0	4.39 ±0.77
61	1.88 ±1.44	1.03 ±0.36
62	1.08 ±0	14.13 ±0
63	1.25 ±0.15	1.74 ±0.42
64	0.96 ±0.29	2.53 ±7.34

Table S3. Individual parameter estimates for advanced computational phenotype

ID	λ	a^+	a^-
1	1.6 ±0.2	0.27 ±0.22	>100 ±14.25
2	2.63 ±0.37	15.61 ±9.13	0.43 ±0.15
3	1 ±0.06	2.92 ±4.17	2.35 ±3.61
4	2.72 ±0.35	0.31 ±0.34	13.16 ±10.33
5	0.88 ±0.04	1.86 ±0.77	>100 ±2.13
6	0.99 ±0.03	>100 ±15.19	4.16 ±6.37
7	1.31 ±0.14	5.51 ±10.79	0.47 ±1.1
8	1 ±0.08	3.99 ±2.03	>100 ±0.4
9	1.13 ±0.09	2.06 ±3.32	1.35 ±1.6
10	0.97 ±0.07	2.38 ±0.98	1.22 ±1.99
11	0.88 ±0.04	1.12 ±0.33	>100 ±4.21
12	0.88 ±0.05	0.7 ±0.24	>100 ±6.15
13	5.5 ±1.36	1.14 ±0.71	1.26 ±15.26
14	1.16 ±0.29	0.63 ±3.41	1.16 ±0.7
15	1.13 ±0.03	33.74 ±15.31	5.1 ±4.2
16	1.15 ±0.07	0.9 ±0.34	36.11 ±13.49
17	1.15 ±0.14	1.35 ±1.83	0.84 ±2.23
18	1.92 ±0.41	1.25 ±1.58	0.45 ±0.43
19	1.49 ±0.21	3.04 ±6.95	0.39 ±1.08
20	1.12 ±0.1	2.77 ±4.76	1.02 ±0.74
21	1 ±0.04	8.25 ±4.7	34.44 ±13.85
22	1 ±0.05	2.05 ±0.87	>100 ±2.96
23	1.14 ±0.05	1.38 ±0.49	>100 ±13.24
24	1.14 ±0.07	2.95 ±4.85	2.96 ±5.77
25	1.34 ±0.18	0.95 ±0.49	1.6 ±6.79
26	1.35 ±0.09	1.9 ±1.16	4.71 ±9.33
27	1.92 ±0.11	0.75 ±3.41	0.09 ±0.7
28	0.71 ±0.04	3.58 ±3.22	1.3 ±1.4
29	1.14 ±0.06	1.99 ±0.79	>100 ±4.69
30	1 ±0.05	4.02 ±1.41	>100 ±5.66
31	1.6 ±0.14	1.25 ±0.54	>100 ±5.21
32	1.02 ±0.06	1.62 ±0.57	>100 ±6.48
33	1.03 ±0.05	>100 ±8.21	3.4 ±2.35
34	1.94 ±0.35	0.77 ±0.45	1.24 ±6.18
35	0.88 ±0.04	1.86 ±0.73	>100 ±3.16
36	1.32 ±0.09	5.34 ±8.46	2.36 ±3.67
37	1.69 ±0.27	1.44 ±1.39	0.73 ±2.62
38	4.44 ±1.04	71.41 ±10.07	1.9 ±17.73
39	1.14 ±0.05	2.91 ±1.24	15.11 ±11.66
40	1.91 ±0.13	11.62 ±8.5	0.84 ±1.37
41	1.49 ±0.12	0.74 ±0.3	>100 ±12.6
42	2.66 ±0.19	>100 ±8.19	0.96 ±1.52
43	1.58 ±0.12	6.77 ±10.87	1.63 ±1.44
44	1 ±0.27	4.62 ±1.87	>100 ±0.78
45	1.14 ±0.08	0.64 ±0.25	>100 ±8.72
46	2.02 ±0.29	1.6 ±2.26	1.19 ±3.79
47	1 ±0.07	4.62 ±1.73	>100 ±2.17
48	2.56 ±0.35	1.97 ±1.96	1.39 ±3.65
49	1.29 ±0.1	3.81 ±7.11	1.89 ±3.58
50	1.29 ±0.07	12.41 ±15.4	0.94 ±0.35
51	1.21 ±0.19	0.69 ±0.33	1.38 ±4.64
52	1 ±0.09	0.63 ±0.2	3.16 ±8.03
53	0.91 ±0.96	0.37 ±2.7	0.16 ±6.06
54	1.36 ±0.34	0.65 ±0.91	0.7 ±2.44
55	1.09 ±0.11	1.39 ±0.53	1.64 ±4.11
56	1.01 ±0.47	0.55 ±1.81	1.04 ±5.21
57	1 ±0.25	0.14 ±0.16	2.49 ±5.78
58	0.54 ±0.4	0.65 ±1.69	9.25 ±8.63
59	2.4 ±0.57	0.99 ±0.91	0.47 ±1.63
60	1.14 ±0.05	>100 ±3.67	2.16 ±1.3
61	2.29 ±0.28	5.55 ±7.16	0.56 ±0.27
62	1 ±0.06	7.94 ±3.27	>100 ±0.61
63	1.27 ±0.11	1.96 ±1.86	1.53 ±3.34
64	0.92 ±0.04	2.93 ±2.73	6.43 ±8.76

Tables S4. Allelic and genotype frequencies for 5HTT and MAOA for sample used in basic behavioral results (N=83). "s" indicates the short allele of the 5HTT gene. MAOA-L (MAOA-H) indicates the low (high) variant of the MAOA gene.

A)

5HTT	**N**	**%**
Allele		
S	83	50.00%
L	83	50.00%
Genotype		
s/s	23	27.38%
s/l	37	44.05%
l/l	23	27.38%

B)

MAOA	**N**	**%**
Allele (bp-repeats)		
3	35	42.17%
3.5	1	1.20%
4	46	55.42%
5	1	1.20%
Genotype		
MAOA-L	36	43.37%
MAOA-H	47	56.63%

Table S5. Allelic and genotype frequencies for DRD4 sample used in basic behavioral results (N=83). 7+ denotes a carrier of the 7-repeat allele.

DRD4		
Allele	N	%
2	17	10.24%
3	7	4.22%
4	97	58.43%
5	4	2.41%
7	40	24.10%
8	1	0.60%
Genotype		
2/2	1	1.20%
2/3	1	1.20%
2/4	8	9.64%
2/7	6	7.23%
3/4	3	3.61%
3/7	3	3.61%
4/4	30	36.14%
4/5	4	4.82%
4/7	21	25.30%
4/8	1	1.20%
7/7	5	6.02%
7+	**35**	**42.17%**
7-	**48**	**57.83%**

Tables S6. Allelic and genotype frequencies for 5HTT and MAOA for sample used in basic and advanced computational phenotype analyses (N=64). "s" indicates the short allele of the 5HTT gene. MAOA-L (MAOA-H) indicates the low (high) variant of the MAOA gene.

A)

5HTT	**N**	**%**
Allele		
S	66	51.56%
L	62	48.44%
Genotype		
s/s	18	28.13%
s/l	30	46.88%
l/l	16	25.00%

B)

MAOA	**N**	**%**
Allele (bp-repeats)		
3	28	43.75%
3.5	0	0.00%
4	35	54.69%
5	1	1.56%
Genotype		
MAOA-L	29	45.31%
MAOA-H	35	54.69%

Table S7. Allelic and genotype frequencies for DRD4 for sample used in basic and advanced computational phenotype analyses (N=64). 7+ denotes a carrier of the 7-repeat allele.

DRD4		
Allele	N	%
2	14	10.94%
3	5	3.91%
4	73	57.03%
5	3	2.34%
7	32	25.00%
8	1	0.78%
Genotype		
2/2	1	1.56%
2/3	1	1.56%
2/4	6	9.38%
2/7	5	7.81%
3/4	2	3.13%
3/7	2	3.13%
4/4	23	35.94%
4/5	3	4.69%
4/7	15	23.44%
4/8	1	1.56%
7/7	5	7.81%
7+	**27**	**42.19%**
7-	**37**	**57.81%**

Table S8. Summary of sample sizes. "Basic behavioral results" refers to results from analysis shown in Fig 1. "Basic & advanced computational phenotype" refers to results from all behavioral and genetic analyses using either the basic (2 parameter) or advanced (3 parameter) computational phenotype.

Initial Sample Size	**90**
Failure to genotype	6
Failure to understand instructions	1
Basic behavioral results sample size	**83**
Randomless Choice Behavior	9
Random Choice Behavior	8
Violate rationality constraint	2
Failure to estimate basic or advanced computational phenotype	19
Basic & Advanced computational phenotype sample size	**64**

REFERENCES

1. Cesarini, D., et al., Genetic Variation in Financial Decision-Making. Journal of Finance, 2010. **65**(5): p. 1725-1754.
2. Cesarini, D., et al., Genetic Variation in Preferences for Giving and Risk Taking*. Quarterly Journal of Economics, 2009. **124**(2): p. 809-842.
3. Zhong, S., et al., The heritability of attitude toward economic risk. Twin Res Hum Genet, 2009. **12**(1): p. 103-7.
4. Dreber, A., et al., The 7R polymorphism in the dopamine receptor D4 gene (DRD4) is associated with financial risk taking in men. Evolution and Human Behavior, 2009. **30**(2): p. 85-92.
5. Kuhnen, C.M. and J.Y. Chiao, Genetic determinants of financial risk taking. PLoS ONE, 2009. **4**(2): p. e4362.
6. Zhong, S., et al., Monoamine oxidase A gene (MAOA) associated with attitude towards longshot risks. PLoS One, 2009. **4**(12): p. e8516.
7. Roe, B.E., et al., Financial and psychological risk attitudes associated with two single nucleotide polymorphisms in the nicotine receptor (CHRNA4) gene. PLoS One, 2009. **4**(8): p. e6704.
8. Tom, S.M., et al., The Neural Basis of Loss Aversion in Decision-Making Under Risk. Science, 2007. **315**(5811): p. 515-518.
9. Tobler, P.N., et al., Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. J Neurophysiol, 2007. **97**(2): p. 1621-32.
10. Rangel, A., C. Camerer, and P.R. Montague, A framework for studying the neurobiology of value-based decision making. Nat Rev Neurosci, 2008. **9**(7): p. 545-56.
11. Doya, K., Metalearning and neuromodulation. Neural Netw, 2002. **15**(4-6): p. 495-506.
12. Schultz, W., Behavioral dopamine signals. Trends Neurosci, 2007. **30**(5): p. 203-10.
13. Daw, N.D., S. Kakade, and P. Dayan, Opponent interactions between serotonin and dopamine. Neural Netw, 2002. **15**(4-6): p. 603-16.
14. Aston-Jones, G. and J.D. Cohen, An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. Annu Rev Neurosci, 2005. **28**: p. 403-50.
15. Ishii, S., W. Yoshida, and J. Yoshimoto, Control of exploitation-exploration meta-parameter in reinforcement learning. Neural Netw, 2002. **15**(4-6): p. 665-87.
16. Sabol, S.Z., S. Hu, and D. Hamer, A functional polymorphism in the monoamine oxidase A gene promoter. Human Genetics, 1998. **103**(3): p. 273-279.
17. Shih, J.C. and K. Chen, MAO-A and -B gene knock-out mice exhibit

distinctly different behavior. Neurobiology (Bp), 1999. **7**(2): p. 235-46.

18. McDermott, R., et al., Monoamine oxidase A gene (MAOA) predicts behavioral aggression following provocation. Proc Natl Acad Sci U S A, 2009. **106**(7): p. 2118-23.
19. Meyer-Lindenberg, A., et al., Neural mechanisms of genetic risk for impulsivity and violence in humans. Proc Natl Acad Sci U S A, 2006. **103**(16): p. 6269-74.
20. Buckholtz, J.W., et al., Genetic variation in MAOA modulates ventromedial prefrontal circuitry mediating individual differences in human personality. Mol Psychiatry, 2008. **13**(3): p. 313-24.
21. Ibanez, A., et al., Pathological gambling and DNA polymorphic markers at MAO-A and MAO-B genes. Mol Psychiatry, 2000. **5**(1): p. 105-9.
22. Lesch, K.P., et al., Association of anxiety-related traits with a polymorphism in the serotonin transporter gene regulatory region. Science, 1996. **274**(5292): p. 1527-31.
23. Canli, T. and K.P. Lesch, Long story short: the serotonin transporter in emotion regulation and social cognition. Nat Neurosci, 2007. **10**(9): p. 1103-9.
24. Asghari, V., et al., Modulation of Intracellular Cyclic AMP Levels by Different Human Dopamine D4 Receptor Variants. Journal of Neurochemistry, 1995. **65**(3): p. 1157-1165.
25. Ebstein, R.P., et al., Dopamine D4 receptor (D4DR) exon III polymorphism associated with the human personality trait of Novelty Seeking. Nat Genet, 1996. **12**(1): p. 78-80.
26. Perez de Castro, I., et al., Genetic association study between pathological gambling and a functional DNA polymorphism at the D4 receptor gene. Pharmacogenetics, 1997. **7**(5): p. 345-8.
27. Kahneman, D. and A. Tversky, Prospect Theory: An Analysis of Decision under Risk. Econometrica, 1979. **47**(2): p. 263-291.
28. Hsu, M., et al., Neural response to reward anticipation under risk is nonlinear in probabilities. J Neurosci, 2009. **29**(7): p. 2231-7.
29. Tversky, A. and D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty, 1992. **5**(4): p. 297-323.
30. Bateman, I., et al., Testing competing models of loss aversion: an adversarial collaboration. Journal of Public Economics, 2005. **89**(8): p. 1561-1580.
31. Sokol-Hessner, P., et al., Thinking like a trader selectively reduces individuals' loss aversion. Proc Natl Acad Sci U S A, 2009. **106**(13): p. 5035-40.
32. Luksys, G., W. Gerstner, and C. Sandi, Stress, genotype and norepinephrine in the prediction of mouse behavior using reinforcement learning. Nat Neurosci, 2009. **12**(9): p. 1180-1186.

33. Pessiglione, M., et al., Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. Nature, 2006. **442**(7106): p. 1042-5.
34. Roiser, J.P., et al., A genetically mediated bias in decision making driven by failure of amygdala control. J Neurosci, 2009. **29**(18): p. 5985-91.
35. Zhong, S., et al., A neurochemical approach to valuation sensitivity over gains and losses. Proceedings of the Royal Society B: Biological Sciences, 2009. **276**(1676): p. 4181-4188.
36. Lusher, J.M., C. Chandler, and D. Ball, Dopamine D4 receptor gene (DRD4) is associated with Novelty Seeking (NS) and substance abuse: the saga continues. Mol Psychiatry, 2001. **6**(5): p. 497-9.
37. Kreek, M.J., et al., Genetic influences on impulsivity, risk taking, stress responsivity and vulnerability to drug abuse and addiction. Nat Neurosci, 2005. **8**(11): p. 1450-7.
38. Schinka, J.A., E.A. Letsch, and F.C. Crawford, DRD4 and novelty seeking: results of meta-analyses. Am J Med Genet, 2002. **114**(6): p. 643-8.
39. Guo, S.W. and E.A. Thompson, Performing the exact test of Hardy-Weinberg proportion for multiple alleles. Biometrics, 1992. **48**(2): p. 361-72.
40. Lagarias, J.C., et al., Convergence Properties of the Nelder--Mead Simplex Method in Low Dimensions. SIAM J. on Optimization, 1998. **9**(1): p. 112-147.

Appendix: 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

	Good state	*Bad state*
+	55%	45%
-	45%	55%

State changes

	Good state today	*Bad state today*
Good state tomorrow	80%	20%
Bad state tomorrow	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.

Earnings=cash + price A(Hold A) + Price B*(Hold B) + Price C*(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.