In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

_____ _____

Andrew M. Brooks Date

The Neural Coding of Loss

By

Andrew M. Brooks

Doctor of Philosophy

Graduate Division of Biological and Biomedical Sciences
Neuroscience

_____
Gregory Berns
Advisor

_____
Jan Barton
Committee Member

_____
C. Monica Capra
Committee Member

_____
Helen Mayberg
Committee Member

_____
Read Montague
Committee Member

_____
James Rilling
Committee Member

Accepted:

_____

Lisa A. Tedesco, Ph.D. Dean of the James T. Laney School of Graduate Studies

_____ Date

The Neural Coding of Loss

By

Andrew M. Brooks

B.S., Mercer University, 2008

Advisor**:** Gregory Berns, MD, Ph.D.

An abstract of

a dissertation submitted to the Faculty of the

James T. Laney School of Graduate Studies of Emory University

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in Graduate Division of Biological and Biomedical Science, Neuroscience

2012

Abstract

The Neural Coding of Loss

By Andrew M. Brooks

After decades of research, we know that the mesolimbic dopamine pathway is heavily involved in nearly all aspects of financial decision-making. Its role has been defined by its involvement in reward-related learning as it relates to monetary gains. There is little evidence implicating it in aversive or loss-related learning. This dissertation was designed to advance our understanding of this systems involvement in aversive and loss related decision making. To do this, we use functional magnetic resonance imaging while human participants engage in experimental economic tasks that deal with financial decisions over gains and losses. We show that, 1) during decision-making, the ventral striatum tracks the expected values of gambles that are entirely aversive, 2) heterogeneity in loss-holding behavior in the stock market can be explained via activity within the ventral striatum, and 3) the ventral striatal processes earnings announcements which lead to financial loss, and are correlated with subsequent changes in stock price. We find that the BOLD response in the ventral striatum to loss outcomes in all three of our experiments fit with the prediction error hypothesis of dopamine activity. Finally, we discuss future research in financial loss, and its implications for sub-optimal behavior, such as gambling addiction.

The Neural Coding of Loss

By

Andrew M. Brooks

B.S., Mercer University, 2008

Advisor: Gregory Berns, MD, Ph.D.

A dissertation submitted to the Faculty of the

James T. Laney School of Graduate Studies of Emory University

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in Graduate Division of Biological and Biomedical Science, Neuroscience

2012

Contents

Chapter 1

Introduction

Goals of the Research Project

If any single conclusion can be drawn from the past decades of research on the neurobiology of human decision-making, it is that the mesolimbic dopamine system is the hub of reward-related learning. Consisting of the ventral tegmental area, the ventral striatum, and regions of the pre-frontal cortex, the mesolimbic dopamine system has been shown, first by electrophysiological experiments in the 1950s (Olds and Milner, 1954), and more recently by human neuroimaging (Breiter et al., 1997; Samanez-Larkin et al., 2008), to be intricately involved in signaling reward-related cues that reinforce behavior. Neuroeconomics is a relatively recent field dedicated to studying human decision using neuroscience tools but with the rigor and incentives of economic paradigms. Research in neuroeconomics has found that when decisions are made over money, structures in the mesolimbic dopamine system encodes various aspects of decision, including value.

The majority of studies in neuroeconomics have focused on how the human brain encodes monetary gains, but only a small portion of the neuroeconomic literature has come to any consensus on how the brain encodes monetary loss. This is likely due to a multitude of reasons, not the least of which is the difficulty of having participants experience true monetary loss in an experimental paradigm. This is because it is only ethical to make sure that participants receive compensation for their time. Even with the experimental research that has been done on monetary loss, there is still no conclusion yet on whether the mesolimbic dopamine system encodes both sides of the monetary coin. Experiments both supporting (Delgado et al., 2003; Seymour et al., 2007; Tom et

al., 2007) and refuting (Yacubian et al., 2006; Levin et al., 2012) this single-system idea exist in the literature. Nor is there any conclusive evidence to suggest alternative loss-value encoding structures, though candidate regions do exist. My dissertation tries to advance our understanding of how the brain encodes financial loss, using functional magnetic resonance imaging (fMRI). Furthermore, my research aims to use current theories on reinforcement learning to help explain both individual investor and market behavior as it relates to loss. The primary goals assigned to each chapter are as follows:

1. To explore whether economic paradigms used to study monetary gains and losses elicit the same neural responses when all of the outcomes are unpleasant.

2. To help explain loss-holding behavior in the stock market by using fMRI to distinguish between several proposed theories of the behavior.

3. To test whether the ventral striatal processes company's earnings announcements when they miss or beat expectations, and whether these responses are correlated with market return.

Background

The Mesolimbic Dopamine System

Through decades of research, the mesolimbic dopamine system has been identified as the primary reward-related pathway. More recent research has expanded its role to a more general valuation system that works across both primary and secondary

rewards (Zink, 2005; Kable and Glimcher, 2007; Haber and Knutson, 2010). This system and its structures are the focus of my dissertation research, and thus an understanding of the system's neuroanatomy is important before proceeding further. **Figure 1-1** outlines the three most prominent structures of the mesolimbic dopamine system as they relate to this dissertation. As the name implies, the primary neurotransmitter of the mesolimbic dopamine system is dopamine. The synthesis of dopamine occurs in two midbrain regions, the ventral tegmental area (VTA) and the substantia nigra (SN). These regions contain tyrosine hydroxylase-positive neurons, a necessary enzyme for the synthesis of dopamine from its precursors. Efferent projections from the VTA and SN terminate in the striatum, both dorsally and ventrally. In the rat, VTA projections end primarily in the ventral striatum, which includes the nucleus accumbens (NAcc) and olfactory tubercule (Moore and Bloom, 1978). SN neurons on the other hand terminate primarily in the dorsal striatum (Smith and Kieval, 2000). The NAcc lies within the ventral part of the striatum, and thus the term 'ventral striatum' is often used interchangeably with NAcc, despite the NAcc only being a subset of the ventral striatum. The VTA indirectly connects to the prefrontal cortices via cholinergic neurons in the basal forebrain (Gray, 1999).

The ventral striatum lies in a position well suited for integration of sensory input and ultimately calculation of value. Besides the VTA, the ventral striatum also receives excitatory afferents from nearly all the prefrontal association cortices, the basolateral amygdala, and the hippocampus via mostly glutamatergic neurons. Thus it is in a prime location for integration of past experience, current sensory input, and higher-order executive output, all of which are requirements for the computation of value. Other areas

of the striatum receive input mainly from motor cortices and the thalamus. The output targets of the ventral striatum include the globus palladus, hypothalamus, and SN, which projects back to the thalamus (Nicola et al., 2000). The ventral striatum is thus in a position to influence motor action via globus palladus and thalamic projections.

Within the ventral striatum, the majority of the neurons are inhibitory GABAergic medium spiny neurons. Afferents from the VTA terminate primarily on the neck of the dendrites of spiny neurons in the ventral striatum, while glutamatergic input from the cortical regions, hippocampus, and amygdala, terminate on the head of the same dendrites. The close proximity of these two inputs allows dopamine to serve as a neuromodulator of the glutamatergic input (Pisani et al., 2001).

Dopamine's modulatory role in the ventral striatum is quite complex. Its influence on glutamatergic input depends on which receptor subtype it is bound to (D1-like vs. D2-like dopamine receptors), which itself is governed by both the concentration of the receptors and of dopamine. Dopamine neuron firing can also take on two different temporal states – either tonic firing or phasic firing. Tonic firing of dopamine neurons determines the baseline steady-state concentrations of dopamine within the synaptic cleft and extracellular space. Phasic dopamine firing on the other hand occurs in bursts and release much higher concentrations of dopamine than tonic firing. Reward prediction error, a behaviorally-relevant learning signal discussed throughout this dissertation, is thought to be signaled via phasic dopamine firing (West et al., 2003; Glimcher, 2011).

Reward Prediction Error

Though we know that the mesolimbic dopamine pathway is involved in all types of reward behavior; the application of reinforcement learning theories from artificial

intelligence has helped shed light on exactly *what* dopamine might be signaling. The current consensus is that dopamine signals reward prediction error (RPE), and thus far the model best able to explain empirical findings from electrophysiological and fMRI experiments, is the temporal-difference model of learning. RPE in its simplest form is the difference between what reward is expected, and what reward is actually received. Unexpected rewards elicit positive RPEs, while fully-predicted rewards elicit no RPE. The first manifestation of RPE was in a mathematical formulation by Bush and Mosteller (1953). This formula was based on Pavlov's dogs, who learned to associate a bell that preceded feeding with food, evidenced through the dog's salivation at the bell. They formulated the following equation to describe the bell-food association:

$$A_{next\_trial} = A_{last\_trial} + \alpha(R_{current\_trial} - A_{last\_trial})$$

What the equation describes is the probability that a dog will salivate to a bell on the next trial (A_{next_trial}) based upon the probability of the last trial (A_{last_trial}), plus a correction term. This correction is the difference between the reward received in the current trial ($R_{current\_trial}$) and what was expected in the previous trial (A_{last_trial}), which forms the basis of RPE. The rate at which new information from the correction term updates expectations is represented by α, which ranges from 0 to 1. Over multiple trials, A_{next_trial} becomes a weighted sum of previous rewards, with the weights being determined by α. Thus if α = 0.5, the probability of salivating rests 50% with $R_{current\_trial}$, 25% with R_{t-1}, 12.5% with R_{t-2}, and so on. The equation does a sufficient job of modeling the most basic of classical conditioning.

All subsequent models of reinforcement learning that are prominent in helping explain dopamine activity build upon Bush and Mosteller's equation, which did an

accurate job at representing basic associations between a single stimulus and subsequent reward. Rescorla and Wagner (1972) built on this model by allowing multiple conditioned stimuli to predict a single reward, each vying for association with the reward. The Rescorla-Wagner model of reinforcement learning is arguably the most well-known of the reinforcement learning theories (Wynne and Staddon, 1998; Glimcher, 2011). However, both the Bush and Mosteller and Rescorla-Wagner models are limited in their scope – they allow an agent to learn the association between a stimulus and reward based on past rewards. Modeling a more complex decision-maker requires the ability to update values and actions in real-time, and not only by trials. Furthermore, complex decisions require the association not just of stimuli to reward, but of stimuli to other stimuli which predict reward.

Sutton and Bartow (1998) created the temporal difference (TD) model of learning to account for the shortcomings of previous models. In the TD model of learning, the ultimate goal of an agent is to build the most accurate prediction of future rewards. The RPE in a TD model is the difference between what rewards one expects to receive on into the future and what rewards/signals they currently are consuming. Also, unlike the Bush and Mosteller (1953) or Rescorla-Wagner (1972) models, the TD model updates expectations not on a trial-wise basis but at discrete points in time. To illustrate this, take the expectation of future reward an agent might have at time point *t*. In the TD model, this is given by the following equation:

$$(1) \quad V_t = r_t + \gamma^1 r_{t+1} + \gamma^2 r_{t+2} + \cdots + \gamma^n r_{t+n}$$

where the expectation of future reward at time *t* is given by V_t. Reward or average reward at a given time point is given by *r*, and is discounted into the future by a factor of γ (because humans discount rewards received in the future).

Clearly there is no way to predict what future rewards will be received with complete accuracy, and thus an estimation of V_t must be made. Because of the recursive nature of the value function, future expectations of reward can be estimated by adjusting the current expectation based on new information about future reward. This estimated value of Vt, denoted $\hat{V}_t$, is given by:

$$(2)\quad \hat{V}_t = \hat{V}_t + \alpha\delta_t$$

$$(3)\quad \delta_t = r_t + \gamma\hat{V}_{t+1} - \hat{V}_t$$

In equation (2), α is the rate of learning (similar to α in the Bush and Mosteller model), and δ_t is the RPE term. Unlike the Bush and Mosteller model though, the current prediction is updated not only by the reward received at the current state in time (r_t), but also by the difference between the current expectation ($\hat{V}_{t+1}$) and the previous expectation ($\hat{V}_t$). One can see that by substituting equation (3) in for RPE in equation (2) and ignoring the learning rate, yields the following equation which is an estimate of V_t given in equation (1):

$$(4)\quad \hat{V}_t = r_t + \gamma\hat{V}_{t+1}$$

The most important aspect of the TD model as it relates to dopamine signaling, is that the RPE back-propagates to the first piece of information which predicts future reward. That is, δ is 0 when expectations are fully met. Schultz et al. (1993; 1997) showed that during a classical conditioning task in monkeys, before training, midbrain dopamine neurons fired to a reward received after an unconditioned cue. After learning

the association between the cue and the reward, midbrain dopamine neurons would fire exclusively to the cue, consistent with a TD-model RPE. A large literature now supports the role of dopamine in signaling RPE in humans (Pagnoni et al., 2002; O'Doherty et al., 2004), monkeys (Mirenowicz and Schultz, 1996), honey bees (Hellstern et al., 1998), and rats (Roesch et al., 2007). The ventral striatum has a very dense innervation of dopaminergic afferents. Thus, the interpretation of signals gleaned from the ventral striatum during human decision making can best be interpreted through the view of RPE and related reinforcement learning theory models.

Functional Magnetic Resonance Imaging

MRI Physics

Much of the contemporary decision making research in the field of neuroscience utilizes functional magnetic resonance imaging (fMRI). Because fMRI is a non-invasive and relatively safe method that allows the indirect measurement of neural activity in awake and behaving human participants, it is the primary method of this dissertation research. Most of the literature supporting mesolimbic structures in signaling RPE and other aspects of reinforcement learning models in humans mentioned above have done so using fMRI. As fMRI is a relatively complex technique in regards to both the physics behind it, and the analysis of the data, a brief overview is presented below.

As the name implies, functional *magnetic resonance* imaging relies on a property of protons known as nuclear magnetic resonance. Protons are abundant in the human body. When subjected to a strong magnetic field, proton nuclei will precess parallel to the magnetic field, a low-energy state. If these protons are then subjected to energy at their Larmor frequency (the frequency at which protons precess around the longitudinal

axis), they will absorb the energy and 'flip' into a high-energy state that is anti-parallel to the magnetic field; this is known as excitation. After the Larmor frequency energy is no longer present, these nuclei will return to the low-energy parallel state, emitting a photon equivalent to the difference between the low-energy and high-energy state. This release of energy upon relaxation is what constitutes the MR signal that MRI scanners measure.

The MR signal changes over time, based on 1) longitudinal recovery (T1*), and 2) transverse magnetization decay (T2). Longitudinal relaxation occurs after a proton absorbs Larmor frequency energy, and begins to relax back into the low-energy state. The rate at which the proton relaxes is the T1-recovery time constant, usually a few hundred milliseconds to a few seconds. Transverse magnetization decay refers to the fact that initially, all protons subjected to Larmor frequency energy will be in-phase in the transverse plane, but over time, due to spin-spin interactions with surrounding nuclei, they fall out-of-phase with each other, resulting in decay of the MR signal. The time it takes for this to happen is governed by the T2-time constant, which ranges from 10ms – 200ms depending on the tissue type. Other factors, such as magnetic field inhomogeneities, also affect the time it takes for excited nuclei to go out of phase. Both T1 and T2 rates differ based on tissue type, thus providing contrast. T2* rates are what provide the signal for functional MRI, as the rate of T2* depends not just on interactions with surrounding nuclei, but also on magnetic inhomogeneities – including those created from the flow of oxygenated blood.

Cerebral Blood Flow and the BOLD Response

The flow of oxygenated blood to a brain region, and the resulting change in T2* effects provide the contrast for fMRI. Oxygenated blood is diamagnetic, which means it has no unpaired electrons and no magnetic moment. Deoxygenated blood is paramagnetic, as it has unpaired electrons and a magnetic moment. Because it is paramagnetic, it contributes to a faster transverse decay (a shorter T2* time). Thus more MR signal will be picked up in regions with higher amounts of oxygenated blood (Ogawa, 1990). This forms the basis of the blood-oxygenation-level-dependent response (BOLD) – the MR signal captured in a brain region where activity has increased. Taking a single snapshot of the brain from T2*-weighted images would not provide any information regarding cognitive processes that occur over time. fMRI is able to capture activity related to cognitive processes by capturing T2*-weighted MR signal over time. A typical fMRI experiment will capture the BOLD response over time with a temporal resolution of around 2 seconds, though this time varies across experiments (Huettel et al., 2009).

Using the BOLD response as an indirect measure of neuronal activity relies on the observation that as synaptic activity increases in a region of the brain, cerebral blood flow (CBF) also increases to that region. Neurons and astrocytes require ATP for their energy. These cells utilize pathways known as cellular respiration to replenish ATP stores, including glycolysis, aerobic respiration, and anaerobic respiration. The cellular respiration process breaks down glucose using oxygen (if present), and the downstream result is ATP. When a neuron or astrocyte requires more glucose or oxygen for ATP synthesis, local vasoactive substances help dilate surrounding blooding vessels resulting in higher CBF to that region. Still-undetermined regional messengers tasked with

signaling the need for more CBF are also released. One of these candidates is nitric oxide, which can act as both a local and distant vasodilator.

It is important to note that an increase in CBF does not necessarily mean neurons in that region are increasing in their action potential frequency. Work by Logothetis et al. (2001; 2002) has demonstrated that the BOLD response is mainly correlated with local field potentials, as opposed to action potentials from neurons in that region. Local field potentials are driven by all dendritic synaptic activity (excitatory or inhibitory), and thus the BOLD response is likely to reflect not only synaptic input into a region, but also intra-regional processing.

Data Preprocessing

Before fMRI data can be analyzed, it must go through a series of processing steps. These steps help with two aspects of the data – they help improve the signal-to-noise ratio (SNR), and help prepare data for analysis. The first step that is typically done is slice-timing correction. Because each slice in a 3D volume of the brain is captured at slightly different points in time, and because adjacent slices are often not captured in sequential order (even slices are often captured first, odds second, or visa-versa), slice timing correction is often applied so that each slice's data can be 'assumed' to be the same temporally. This is done using interpolation, where data from previous and succeeding slices are used to estimate the current slices data. To do this, a reference slice is given, and all other slices are interpolated with that slice.

The second step of standard data pre-processing is motion correction. Head motion is highly detrimental to fMRI data. An important assumption for data analysis is

that each voxel remains in the same spatial position over time. Movement violates this assumption, and can result in spurious findings. Prevention of head motion using foam blocks and participant training is often done. To some degree, motion can be corrected for after data acquisition by using algorithms which find the best way to translate and rotate volumes to match a reference volume. These algorithms utilize cost functions, which are measures of how similar two images are from one another. The goal of motion correction is to minimize the cost function between the reference volume and other volumes in the data set. Most neuroimaging tools adhere to this type of motion correction, however the types of cost functions that are used differ based on the neuroimaging packages used. Another step that can be taken to account for mild to moderate movement, is to include motion correction parameters, which describe the degree to which each volume has to be moved to minimize the cost function. This helps absorb variance in the data attributed to movement, though it can never fully account for all motion related artifact (Johnstone et al., 2006).

For group analyses, a third step in fMRI data preprocessing is spatial normalization. As the adult human brain can differ by up to 42% (Allen et al., 2002), each individual's fMRI data must be transformed into the same coordinate system in order for group-level statistics to be accurate. Furthermore, transformation to templates allows for cross-study comparison of brain activation. The Montreal Neurological Institute template or Talairach templates are the most commonly used templates (Talairach and Tournoux, 1988; Evans et al., 1993). Similar to motion correction, spatial normalization procedures attempt to find the best transformation that minimizes the differences between two images, based on landmarks that are manually labeled or on

voxel intensities. Spatial normalization can be done in a linear fashion (12-parameter affine transformation), non-linear (which allows warping), or both (Huettel et al., 2009).

Finally, fMRI data typically undergoes spatial smoothing to help increase the probability of signal detection. Spatial smoothing works to improve the signal-to-noise ratio if the size of the kernel is similar to the size of the signal of interest, a principle known as matched filters (Friman et al., 2003). Spatial smoothing is accomplished by convolving the data with a Gaussian function, such that neighboring data points are averaged. The distance and weights with which neighbors are averaged are determined by the kernel size, usually given in terms of full width at half maximum (FWHM). FWHM describes the width of the Gaussian kernel at half its maximum height. A larger FWHM will smooth the data over more voxels, whereas a smaller FWHM will limit the smoothing to a smaller region. fMRI studies will typically employ kernel sizes of 4 - 8mm at FWHM. Generally, the kernel size is chosen based on the size of the regions expected *a priori* to be active during the task along with the resolution of the data.

Data Analysis

The most common type of fMRI statistical analysis for human cognition studies is the general linear model (GLM). A GLM model takes the form:

$$Y = \beta X + \varepsilon$$

Where Y is a matrix of dependent variables, X is a matrix of independent variables, β is a matrix of estimates of the relationship between the independent and dependent variables, and ε is an error term. In terms of an fMRI task, Y is the pre-processed BOLD signal, and X is a design matrix where the number of columns represent events of interest (and their interactions) convolved with a hemodynamic response function (HRF). A simple

task where a participant stares at a screen which either flashes a checkerboard (ON) or is blank (OFF), might model each individual voxel the following way:

$$BOLD_i = \beta_0 + \beta_1(ON_i \otimes HRF) + \varepsilon_i$$

The index of scans is represented by i, β_0 is the baseline composed of all non-modeled events, and β_1 is an estimate of the effect of the checkerboard on the BOLD response in that voxel. The BOLD response from a single event takes on a characteristic shape, the HRF, which peaks between 4-6 seconds and lasts around 10 seconds, and is followed by a slight undershoot (**Figure 1-2**). Thus the regressors which represent events in the task must first be convolved ($\otimes$) with an HRF before being regressed onto the BOLD signal. This regression is run for each voxel in the brain and for each participant in the study before moving onto group-level statistics.

The majority of group-level statistics use random-effects analyses, where variance for each subject is calculated, as opposed to fixed-effects models where variance is estimated scan-by-scan (Friston et al., 1999). Random-effects analyses allow for inferences about the population from which the sample was pulled, whereas fixed-effects models allow only inferences about the sample itself. To do a random-effects analysis, contrast maps for each participant are created. For example, a contrast map representing the effect of the checkerboard would be a 3D volume where each voxel represents β_1 from the GLM. A random-effects analysis takes the contrast map representing β_1 for each participant and runs a one-sample t-test on every voxel in the brain to test whether β_1 is equal to zero or not. All group-level analyses in this dissertation represent random-effects models.

The statistical thresholding of the group-level t-maps has been the focus of much research interest (Forman et al., 1995; Genovese et al., 2002; Loring et al., 2002; Lieberman and Cunningham, 2009). Because a t-test is performed for each voxel in the brain, and there can be as many as 10,000 voxels, at a threshold of $\alpha = 0.05$, 500 voxels can be expected to show significance due to random chance. This summarizes the multiple comparisons problem, which is inherent in the analysis of any type of large data set, including genome-wide association studies (Moskvina and Schmidt, 2008). Fortunately, there are many methods to account for multiple comparisons without being unnecessarily strict. A Bonferroni correction at $\alpha=0.05$ would yield a voxel-wise α of 5 x 10-6, assuming 10,000 voxels in the brain, which is unreasonably strict given the SNR of fMRI data. Alternatives such as cluster-thresholding, false-discovery rate correction, AFNI's AlphaSim estimation, and family-wise error rate correction through Gaussian random field theory all provide more reasonable ways of multiple comparisons correction. All results presented in this dissertation include some method of multiple comparisons correction.

Gains and Losses in the Brain

The bulk of the research in how the brain processes financial gains and losses comes from the field of neuroeconomics. Neuroeconomics emerged in the late 1990s (Loewenstein et al., 2008). Neuroeconomics merges the more rigid theory-based approach of economists, with psychology that focuses on the subjective experience of humans in their decision making, along with neuroscience that seeks to understand the mechanics behind our decision making process (Glimcher and Rustichini, 2004). This

powerful approach allows us to discover where and how decision variables such as risk, expected value, losses and gains are represented in the brain.

Early neuroeconomic studies sought to locate where expected utility was represented in the brain. Utility is a term that describes the subjective value a person places on something. The term comes from expected utility theory, an economic theory of human choice under uncertainty (Von Neumann and Morgenstern, 1947). Before the use of fMRI, utility was only observable via human choice. That is, if a person picks one item over another, it is said that the chosen item confers more utility to the person than the non-chosen option. However, if the brain ultimately perceives, processes, and carries out decisions, utility might be able to be captured using fMRI. Thus the field of neuroeconomics was born.

An early series of studies done by Knutson et al. (2000, 2001a) captured the brain's response to outcomes of differing monetary value, and presumably differing utility. Participants in these studies completed what is known as the monetary incentive delay (MID) task. The MID task was modeled after a reward anticipation task that Schultz et al. (1997) used to demonstrate TD-model-like RPE in VTA dopamine neurons. For each trial in the MID, the participant sees a cue that indicates what the potential outcome will be (a reward, punishment, or control trial; for Knutson et al. (2001a) the cue indicated the amount of potential reward). Following a brief delay, participants were required to make a button press response within a brief window of time in order to win money or avoid loss of money from an endowed amount. This window of time was adjusted for each participant, such that 60% of the time participants were successful. Both dorsal and ventral (among a host of other regions) showed an increased

BOLD response to potential monetary reward. In a follow-up study, the BOLD response scaled to the size of the potential reward – thus larger rewards evoked a larger BOLD response than smaller rewards.

A host of subsequent studies have shown that the mesolimbic structures signal much greater complexities of value than monetary tasks with single potential outcomes – most notably expected value, which combines both probability and magnitude into a single term. Expected value is a fundamental concept in economics, probability theory, and a host of other disciplines. It describes the weighted average that a random variable can take on, calculated by:

$$E[X] = x_1 p_1 + x_2 p_2 + \cdots + x_n p_n$$

where E[X] is the expectation of random variable X, and where outcome x_i occurs with probability p_i. Take a simple 50/50 gamble of \$10 and \$0. If this gamble is played a large number of times, on average, the player can be expected to win \$5. From the above equation:

$$E(gamble) = (\$10) \bullet (0.5) + (\$0) \bullet (0.5)$$

$$E(gamble) = \$5$$

Objectively, a gamble with a higher expected value earns a player more money over the long term than one with a lower expected value. In fMRI studies, the BOLD response in the ventral striatum increases with expected value during both decision-making and anticipation, suggesting that the striatum is representing expected value, or some correlate of it (Kable and Glimcher, 2007; Staudinger et al., 2009). It is unclear whether probability and magnitude, both required for expected value calculations, are represented in the ventral striatum, or are calculated in spatially different regions or separated

temporally (Knutson et al., 2005; Preuschoff et al., 2006; Yacubian et al., 2007; Berns and Bell, 2012). In either case, the observed BOLD response in the ventral striatum is consistent with a TD-model RPE (Montague and Berns, 2002).

There is strong evidence for a common neural 'currency' when it comes to positive-valence stimuli. Social rewards (Rilling et al., 2002; Izuma et al., 2008, 2010), food rewards (Kelley and Berridge, 2002; Smeets et al., 2006), juice reward (Bowman et al., 1996; Pagnoni et al., 2002), and monetary reward are all correlated with increases in the BOLD response in the ventral striatum. When it comes to monetary loss, the picture becomes less clear. Some studies have shown a decrease in BOLD response to signals of a monetary loss, as would be expected from a RPE standpoint (Delgado et al., 2000, 2008). Others have not (D'Ardenne et al., 2008). Tom et al. (2007) demonstrated that the ventral striatum increased in BOLD response for higher potential gains and decreased in response to higher potential losses. They also showed that the slope of the response to losses was twice that of gains – reflecting the observation that people weight financial losses twice as much as financial gains (Kahneman and Tversky, 1979). Others find no striatal response to monetary loss in the ventral striatum, but increases to monetary loss in other regions such as the amygdala and insula (Yacubian et al., 2006).

If monetary loss is akin to a physically aversive stimulus, then there is lots of support for the idea that loss can be tracked in the ventral striatum. Separate sub-populations of dopaminergic neurons have been shown to respond to aversive stimuli by an increase in firing rate, and other neurons by an inhibition of firing (Coizet et al., 2006; Matsumoto and Hikosaka, 2009). The anticipation of pain has been shown to increase the BOLD response in the ventral striatum (Becerra et al., 2001; Jensen et al., 2003).

Appetitive RPEs that signal better-than-expected outcomes as well as aversive RPEs that signal worse-than-expected outcomes have both been found in the ventral striatum, but separated spatially (Seymour et al., 2004, 2005). Just as primary rewards activate similar regions as secondary rewards such as monetary gains, we might expect primary punishments or aversive stimuli to activate similar regions as monetary loss. It is unclear whether this should manifest itself as a decrease in activation or increase in activation to monetary loss, as neuronal populations exist for both appetitive and aversive RPE.

Experiment Overview

To better understand how the brain encodes loss, I setup three experiments to test different aspects of human decision making. The **first experiment (chapter 2)** was designed to tackle a problem inherent in the majority of research studies using money as the primary outcome. Research study participants must be compensated for their time, thus participants can't leave with less money than they arrived with. In most experiments dealing with financial loss, participants are endowed with a sum of money at the beginning of the experiment, and can gain or lose from that amount. Whether this prompts participants to behave as if the money was their own is unclear, though it is known that behaviorally, people behave more risky when given a windfall gain – known as the 'house money effect' (Thaler and Johnson, 1990). To attempt to get around this problem, researchers will often modify their incentive structure to make participants feel more like the money is their own. One way of doing this is to pay participants upon signing up for the experiment, and have them bring that money with them when they return for the actual task several weeks later (Tom et al., 2007). Though likely to be the best way to have participants experience a financial loss, it is difficult to implement. An

alternative is to have participants work for their endowment prior to completing the actual task.

To bypass this issue entirely, we take an economic paradigm commonly used to study what brain structures track monetary value, but instead use a medium unpleasant for most people – electric shocks. In this paradigm, participants make a decision to accept or reject mixed gambles of electric shocks while undergoing fMRI. Mixed gambles are gambles that consist of both a gain and a loss. If traditional value-related structures in the mesolimbic dopamine system, namely the ventral striatum, track aversive stimuli – we should expect to see an increase in BOLD response in these regions for better gambles.

In **second experiment (chapter 3)**, I use fMRI to help differentiate theories of the disposition effect, a behavior which has a large loss-related component. We focus on ventral striatal activity, and interpret it through the TD-model RPE framework. The disposition effect is a phenomenon first observed in the stock market where investors hold onto losing assets longer than they hold onto gaining assets (Shefrin and Statman, 1985). There are a multitude of theories behind what drives the behavior, but no conclusive way to differentiate them based on behavior alone. Focusing on participant behavior over financial losses, and instituting an incentive structure designed to make financial losses more salient, we test whether activity in the ventral striatum can help explain which theories are more likely to be correct.

In the **third experiment (chapter 4)**, we test hypotheses related to how investors process company earnings announcements, and its relation to subsequent changes in stock price. In this experiment, second-year MBA students serve as proxies for investors.

They take a position (long or short) in a firm's stock based on past earnings per share information, and past and current analyst consensus forecasts. We then capture the ventral striatum's BOLD response to the actual earning's announcement for the current quarter. Participants' monetary compensation is based on the subsequent change in stock price. I focus first on whether the ventral striatum processes both positive and negative earnings announcements, akin to a monetary gain and loss. Second, I test whether an aggregated BOLD response to both the positive and negative earnings announcements from the ventral striatum can predict the subsequent change in market price. Third, I test whether the position (long or short) affects the aforementioned striatal responses.

Figures and Tables

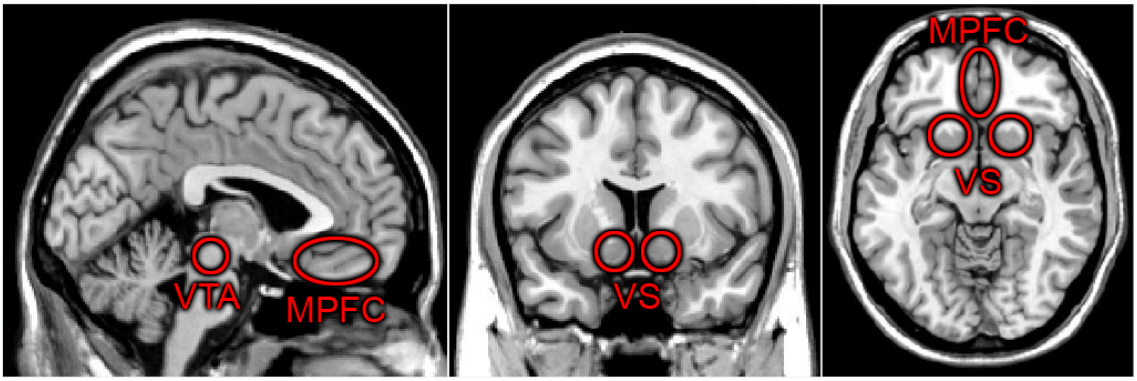

Figure 1-1: Relevant anatomical structures of the mesocorticolimbic dopamine system to neuroeconomic literature. Synthesis of dopamine in the mesolimbic dopamine system occurs in the ventral tegmental area (VTA), and extensive axonal projections from this region terminate in the ventral striatum (VS). Cortical regions, including the medial pre-frontal cortex (mPFC) have projections back to the VS, making the VS an ideal area for the integration of information related to value. The magnitude of unexpected financial gains and losses have been shown to correlate with synaptic activity in the VS and mPFC.

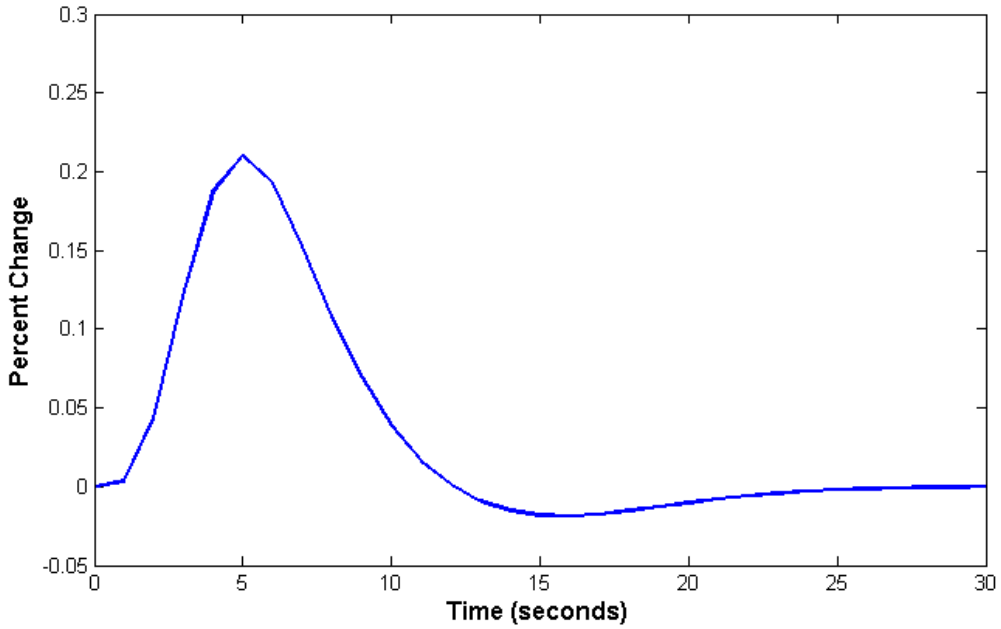

Figure 1-2. Canonical hemodynamic response function (HRF) from the SPM8 neuroimaging package. The response to a single stimulus evokes a bloody-oxygenation-level-dependent (BOLD) response which peaks between 4-6 seconds, lasts 10 seconds, followed by an undershoot before returning to baseline. In the SPM neuroimaging package (Wellcome Department of Imaging Neuroscience, University College of London), a double-gamma function is used to model the peak and undershoot of the HRF.

Chapter 2

From Bad to Worse: Striatal Coding of the Relative Value of Painful Decisions

Introduction

Many real world decisions involve the possibility of both good and bad outcomes, but sometimes the choices are between bad and worse. Consider, for example, an individual who purchases a cell phone plan only to realize that the reception with that carrier is terrible. The individual is then faced with the decision to either stay with the carrier and suffer bad reception, or pay an exorbitant cancellation fee. In either case, the outcome is bad. The recent advent of neuroeconomics has brought new methods of analysis to the study of human decision-making, but the vast majority of these studies have focused on decisions in which all possible outcomes are non-negative (Knutson et al., 2001b; McClure et al., 2004a; Padoa-Schioppa and Assad, 2006; Preuschoff et al., 2006; Tobler et al., 2007). But because relatively few studies have examined decisions made entirely in the domain of losses, it is not clear how the brain gauges relative value when all of the outcomes are bad. One hypothesis regarding valuation in the brain suggests that the utility of positive outcomes is evaluated by a separate neural system from that of negative outcomes. In its simplest form, the dual-systems hypothesis associates the ventral striatum and the orbitofrontal cortex (OFC) exclusively to the evaluation of gains (Mirenowicz and Schultz, 1996), and the amygdala and insula exclusively to the evaluation of losses (Yacubian et al., 2006).

There is some evidence that the striatum and other orbitostriatal structures are involved in both gain and loss processing (Delgado et al., 2003; Seymour et al., 2007; Tom et al., 2007). However, most of these studies pitted a potential gain against a loss, used a medium that is generally rewarding (money), or focused solely on the anticipation of the gain or loss, and thus it is not clear that when all outcomes are negative, whether the striatal system would still be engaged or whether a separate system would perform the decision-making processing. Some research suggests that anticipation and experience of aversive stimuli activate the striatum (LaBar et al., 1998; Becerra et al., 2001; Jensen et al., 2003; Seymour et al., 2004). Indeed, there are populations of dopaminergic neurons that respond to aversive stimuli (Coizet et al., 2006; Matsumoto and Hikosaka, 2009). Thus, we hypothesized that the striatal system also processes the value of non-rewarding stimuli during the decision-making process itself, as opposed to solely the anticipation of the stimuli. For painful outcomes (electric shocks), we predicted that fewer electric shocks from a reference amount would be viewed as a "gain" and more electric shocks as a "loss." Furthermore, we predicted that the ventral striatum would be involved in processing "gains" (fewer electric shocks) even though the overall outcome medium was always unpleasant.

To test these hypotheses, we used fMRI along with a gambling task involving electric shocks. In a manner similar to that in the task used by Tom et al. (2007), participants were asked to accept or reject a 50/50 gamble of "more" or "fewer" electric shocks compared to a reference amount that they received at the beginning of each trial. If participants rejected the gamble, they received the reference amount of shocks. If they accepted the gamble, they either received "more" or "fewer" shocks from the reference

amount. Using this task, we tested whether participants' choice behavior was consistent with an adaptation of their status quo to the reference level of shocks. Our analysis of the neuroimaging data focused on the period in which participants decided whether to accept or reject these lotteries. This allowed us to identify regions involved specifically in decision-making as opposed to the anticipation of the outcomes.

Materials and Methods

Participants

Thirty-six participants (18 female, 18 male; 18–45 years) were recruited from the Emory University campus. All participants were right-handed, reported no psychiatric or neurological disorders, or other characteristics that might preclude them from safely undergoing fMRI, and provided informed consent to experimental procedures approved by the Emory University Institutional Review Board. Participants received a base pay of $40.

Experimental Procedures

A Biopac STM100C stimulator module with a STMISOC isolation unit (Biopac Systems, Inc., CA, USA) was used to deliver electric shocks cutaneously to the dorsum of the left foot through shielded, gold electrodes placed 2–4 cm apart. The STMISOC unit controlled current output to the electrodes, with each pulse lasting 15 ms. The stimulator module was connected via a serial-interface to a laptop which controlled the timing and delivery of the shocks.

Prior to scanning, shock intensity was calibrated by finding each participant's "maximum shock intensity", I_{max}. Participants were told that their maximum shock

intensity would be set to the highest intensity that they could bear. For the calibration procedure, each trial consisted of 18 shocks over 340 ms (the maximum number per trial in the subsequent experiment). The current was slowly increased until participants notified the experimenter that they couldn't bear it anymore, and this current level was set as their I_{max}. The current level for all shocks throughout the experiment was set at 90% of I_{max}.

To gain familiarity with the different numbers of shock outcomes, participants were passively exposed to all possible outcomes. An attempt to induce a status quo of 10 shocks was made by subjecting participants to the 10 shocks at the beginning of each trial. On each outcome, the number of shocks (SN) was evenly spaced in time over 340 ms, yielding an inter-pulse interval of 340/(SN-1). This was done to avoid confounding the number of shocks with the total duration of shocks. The number of shocks, SN, within a trial was 2, 3, 4, 5, 6, 8, 10, 12, 15, or 18. These numbers were determined based on previous literature that suggests that the Weber fraction for many stimuli range from 0.01 to 0.10, meaning that a difference of at least 1–10% between stimuli is needed in order to be distinguishable from each other (Teghtsoonian, 1971; Lavoie and Grondin, 2004). To insure that participants could distinguish between different numbers of shocks, a difference in number of at least 25% between shocks was used. The status quo was set at 10 shocks, and so a relative gain was framed as "2, 4, 5, 6, 7, or 8 less" and a relative loss as "2, 5, or 8 more."

Following the calibration phase of the experiment, participants entered the scanner to begin the experimental phase, which was modeled after a monetary gambling paradigm used by Tom et al. (2007). Each trial began with a status quo (10 shocks),

which was indicated by the presentation of a circle with the text "Sure Thing" centered in the middle (**Figure 2-1**). After 2 s, this circle turned yellow, indicating the impending onset of the shocks, which occurred after a further 2 s. Following an interstimulus interval (ISI) of 3 s, a 50/50 gamble appeared with the words "Accept" and "Reject" below it. This gamble consisted of two possible outcomes, indicated by separate, equally sized slices of the circle, where the left side was always more potential shocks and the right side always fewer potential shocks. The number of shocks more and less than the reference amount varied between trials, such that every possible combination of shocks was presented.

Two seconds after presentation of the gamble, participants were allowed to "Accept" or "Reject" the gamble by using a button box in the scanner. If participants accepted the gamble, a pink ball flipped between options for a varying amount of time between 3 and 6 s, landing with a 50/50 chance on the more shocks or fewer shocks outcome. The side on which the ball landed turned yellow, indicating the outcome of the gamble and impending shocks, which occurred 4.7 s after the outcome was revealed. If participants rejected the gamble, an identical presentation including the ball-flip and outcome selection occurred. However, in this case the reference shocks were the only possible outcome. After the shocks were administered, the outcome remained on screen for 3 s, and was followed by an inter-trial interval (ITI) of 3 s. The experimental phase consisted of three runs with 18 trials per run (54 trials in total). Trials were randomly ordered for each run within-subjects, but remained the same between-subjects. COGENT 2000 (FIL, University College London) was used for stimulus presentation and response acquisition for this phase.

To confirm that participants could distinguish between the different numbers of shocks and that increasing shocks were increasingly averse, participants rated all possible sets of shocks relative to the reference shocks (after the above procedure but while still in the scanner). A visual analog scale (VAS) was presented on screen, with a white arrow in the center labeled as "reference shocks." Participants were given the reference shocks, and then were given another set of shocks, blinded to the number. They were asked to rate "How much better or worse it is from your reference," by moving the arrow on screen either left ("better") or right ("worse"). All possible sets of shocks were given three times each for a total of 30 data points.

fMRI Measurements

Functional imaging was performed with a Siemens 3 T Trio whole-body scanner. T1-weighted images (TR = 2300 ms, TE = 3.04 ms, flip angle = 8,192 × 146 matrix, 176 sagittal slices, 1 mm cubic voxel size) were acquired for each subject prior to the three experimental runs. For each experimental run, T2*-weighted images using an echo-planar imaging sequence were acquired, which show blood oxygen level-dependent (BOLD) responses (echo-planar imaging, TR = 2350 ms, TE = 30 ms, flip angle = 90, FOV = 192 mm × 192 mm, 64 × 64 matrix, 35 3-mm thick axial slices, and 3 mm^3 voxels).

fMRI Analysis

fMRI data were analyzed using SPM5 (Wellcome Department of Imaging Neuroscience, University College London) using a standard 2-stage random-effects

regression model. Data were subjected to standard preprocessing, including motion correction, slice timing correction, normalization to an MNI template brain and smoothing using an isotropic Gaussian kernel (full-width half-maximum = 8 mm).

Four main regressors were included in the first-level models. (1) The status quo shock at the beginning of each trial was modeled as an impulse function. (2) The "decision" period, during which a decision to accept or reject the gamble was required, was modeled from the onset of gamble presentation until button press. The expected value of the gamble was also included as a parametric modulator for this period. (3) The "ball" period, in which the gamble outcome was resolved over a varying period of time between 3 and 6 s, was modeled as a variable duration function. (4) The "wait" period was modeled from the display of the gamble outcome to the receipt of the shocks. For this period, the number of shocks received was included as a parametric modulator. Subject motion parameters were also included as regressors. All regressors were convolved to the standard HRF function.

Because we were interested in investigating the neural basis of decision parameters that affect choice, the second-level analysis focused on the decision period (#2 above). To identify regions involved in valuation during choice, we first identified regions showing correlations with the expected value of the gamble. We assumed shocks are "bad" and have negative value; for example, the reference shocks would have an expected value (EV) of −10. We calculated the expected value of the gambles with the equation: $EV_{gamble} = -10 + (\text{number of shocks less} - \text{number of shocks more})/2$. EV_{gamble} ranged from −7 for the best gamble, and −13 for the worst gamble. This parameter was expected to directly affect choice, because a less negative EV_{gamble} would

indicate a better gamble and a more negative EV_{gamble} a worse gamble, assuming individuals find electric shocks unpleasant. To further analyze the interaction between potential outcomes with less or more shocks within identified regions, we performed an ROI analysis using beta estimates from a different first-level model in which the number of shocks less and the number of shocks more than the reference amount were modeled by separate parametric modulators. This allowed us to identify the extent to which better and worse potential outcomes separately contributed to EV_{gamble}.

Finally, another first-level model was constructed in order to extract BOLD responses for each individual gamble type during the decision period. Instead of a single lottery period modulated by the number of shocks less and number of shocks more than the reference amount, this model included each lottery period associated with a different gamble as a separate regressor, such that there were 18 columns in the design matrix for the decision period, along with the remaining regressors that appeared in the primary first-level model described above. This allowed the average BOLD activity during the decision period for each separate gamble to be extracted. These values were then used to create "heat maps" of activation which give snapshots of how a particular region responds to all possible gambles.

Results

Behavioral

For monetary payments, if an individual prefers to receive a certain payment rather than a gamble with the same expected value, he is said to be risk averse. If he instead prefers the gamble, he is said to be risk seeking. Prior research with monetary

payments shows that on average, individuals are risk-averse (risk-seeking) for positive (negative) payoffs. We consider whether the shock quantities in our experiment are treated in the same way. To consider the issue, participant behavior in symmetric lotteries was analyzed. Symmetric lotteries were lotteries with the same amount of shocks less and more than the reference amount, and therefore had the same expected value as the reference shocks. Averaged across all runs for all participants, the symmetric lotteries were chosen over the reference shocks 74% of the time, which suggests risk-seeking behavior. For the individual symmetric lotteries of 8/8, 5/5, and 2/2, participants chose the lottery 56%, 78%, and 89% of the time, respectively. Interestingly, this was significantly different between the three symmetric lottery types ($F(2,105) = 10.21$; $p < 0.0001$).

As another indicator of overall risk-preference, the average indifference point across participants was determined by graphing the probability of choosing the lottery as a function of the expected value of the gambles. A sigmoidal curve, shown in **figure 2-2**, was fit to the data using a logistic function to determine the average indifference point. If participants on average were risk neutral, their indifference point would equal the expected value of the reference shocks (−10). If participants were risk-seeking, their indifference point would be less than −10. The average indifference point was −10.94 shocks ($f(-10.94) = 0.500 \pm 0.218$), indicating risk-seeking behavior. The reference point of −10 did not lie within the 95% confidence interval of the logistic fit ($f(-10) = 0.720 \pm 0.214$), and therefore it is likely that this observed indifference point was significantly different from risk-neutrality.

To determine individual risk-preference, the curvature of the utility function, $u(x) = x^{\alpha}$ was estimated for each participant using a non-linear least-squares regression. Participant values were not normally distributed nor were they lognormal, and therefore non-parametric statistics were used to test for significance. A Wilcoxon signed-rank test indicated that, on average, participant α values (median $\alpha = 0.934$, SD = 0.309) were significantly different from one ($p = 0.0381$). Due to the method of estimation (where a larger expected value is a more unfavorable gamble), an $\alpha < 1$ indicates convexity over losses and therefore a preference for risk-seeking behavior, whereas an $\alpha = 1$ indicates a risk-neutral preference. In addition, average VAS ratings for each possible outcome in the study were computed and normalized to the reference shock ratings. When plotted, these ratings revealed a convex function resembling a value function over losses (see **figure 2-3**). The slope of the VAS rating over more and less potential shocks were computed for each participant, using linear regression. A paired-samples *t*-test revealed that the slope for less potential shocks ($M = 1.974$, SD = 0.644) was significantly greater than the slope for more potential shocks ($M = 0.920$, SD = 0.548), $p < 0.001$, consistent with a convex value function.

fMRI

The expected value of the gambles, EV_{gamble}, was used to identify brain regions involved in the valuation of gambles during the decision period (see **figure 2-4**). Used as a parametric modulator, this allowed for identification of regions of the brain whose BOLD signal correlated with the objective gamble value. Positive correlations between EV_{gamble} and BOLD activity were found in the visual cortex, intraparietal sulcus, frontal

eye fields, and the left ventral striatum, among other areas (see upper portion of **table 2-1**). A less negative EV_{gamble} indicated a better gamble, which demonstrates that these regions responded in a graded manner to comparatively better possible outcomes – even though all outcomes were still painful. Given that all outcomes were aversive, it is interesting that ventral striatum activity increased for relatively "less bad" outcomes. Regions with negative EV_{gamble} correlations, or a greater response for worse outcomes (more expected shocks), included the posterior cingulate, anterior cingulate (ACC), inferior parietal lobule, insula, and the lateral OFC (lower portion of **table 2-1**).

To determine how these regions responded to the individual components of the gambles (less or more potential shocks), beta values for the *lottery ×number of shocks less* and *lottery × number of shocks more* condition were extracted from regions identified in the *lottery* × EV_{gamble} contrast mentioned above. The left ventral striatum showed significant positive and negative correlations with the number of potential shocks less (better) and number of potential shocks more (worse), respectively. Other areas identified in the positive EV_{gamble} contrast revealed the same relationship: a significant positive correlation with the number of potential shocks less and negative correlation with the number of potential shocks more. The opposite trend was seen for several regions identified in the negative EV_{gamble} contrast: significant positive correlations with the number of potential shocks more and negative correlations with the number of potential shocks less were observed in the insula, intraparietal sulcus, and dorsomedial prefrontal cortex (DMPFC). To visualize activity to each individual gamble type, we extracted beta values from ROIs in the *lottery*× EV_{gamble} contrast for each gamble type. In the left ventral striatum, gambles with a higher EV_{gamble} were associated with less

deactivation, and gambles with a lower EV_{gamble} were associated with more deactivation, as revealed in a heat map (see **figure 2-5**). A heat map of beta values from the DMPFC for each gamble revealed less activation to gambles with a higher EV_{gamble}, and more activation to gambles with a lower EV_{gamble} (see **figure 2-5**). In other words, more potential shocks elicited above-baseline BOLD activity in these regions. Similar activity was observed in the genual ACC (see **figure 2-5**), with less deactivation for gambles with a lower EV_{gamble}.

Discussion

Contrary to the simplest form of the dual-systems view, which would predict no response from the ventral striatum to gambles consisting solely of losses, our results indicate that the ventral striatum encodes information regarding value irrespective of the type of outcome (e.g., "more" or "less" shocks) and whether the outcomes are globally "good" or "bad" (e.g., appetitive or aversive). In particular, the positive correlation of left ventral striatal activity with the expected value of the shock lotteries supports its role in valuation and extends this to include the relative valuations of "bads." While previous neuroimaging studies have demonstrated the role of the striatum in integrating the value of rewards with a variety of costs (Tom et al., 2007; Croxson et al., 2009; Talmi et al., 2009), our results extend these findings to the domains of pain and loss even when there is no possibility of gain.

That these decisions were viewed as occurring in the loss domain is reinforced by the fact that, despite being exposed to the reference shocks for each trial, participants viewed every outcome as a "loss." This was evidenced by consistent risk-seeking

behavior over the full range of lotteries and a larger slope for less shocks than more shocks relative to the status quo for the VAS ratings. These results are consistent with past research showing risk-seeking behavior over hypothetically painful outcomes (Eraker and Sox, 1981). Interestingly, this risk-seeking behavior cannot explain the changes in striatal activation as others have suggested (Fiorillo et al., 2003; Preuschoff et al., 2006) because the variance of the best and worst lotteries is the same in our task. One possible reason for this lack of status quo inducement is the transient nature of the reference shocks. Although participants were presented with reference shocks between each trial, the majority of the time participants were not experiencing painful stimuli. It is possible that a constant painful stimulus, such as would arise with the use of capsaicin to induce a constant state of pain which can then be attenuated or exacerbated with temperature, might be more effective in inducing a status quo (Seymour et al., 2005).

It is important to distinguish between the loss of something desirable, which has been investigated in a considerable number of prior studies, and the receipt of something undesirable, which has received less attention. Previous neuroeconomic studies of loss aversion have shown that the ventral striatum deactivates to the prospect of monetary loss (Tom et al., 2007). Similarly, striatal deactivation has been observed with increased effort and pain to obtain a monetary gain (Croxson et al., 2009; Talmi et al., 2009). These results point to the integrative role of the striatum in determining net value for monetary rewards but do not directly address its role in the relative valuation of things that are universally bad. Evidence exists, however, that the striatum dynamically scales for relative coding of value (Seymour and McClure, 2008). In a similar manner, dopamine neurons have been observed to adaptively code reward value (Tobler et al.,

2005), so it is plausible that the striatum could exhibit adaptive signaling even in the realm of painful outcomes – for which we find strong evidence here.

Beyond the striatum's adaptive coding of value, its more general role in pain processing has been hotly debated (Leknes and Tracey, 2008). Some studies have shown ventral striatal activity during the anticipation of painful stimuli (Becerra et al., 2001; Jensen et al., 2003), a finding echoed by PET evidence of dopamine release to pain (Scott et al., 2006), while others have argued this activity merely reflects the anticipated relief (Baliki et al., 2010). Still others have suggested that the ventral striatum functions more generally in motivated behavior (Horvitz, 2000; Zink et al., 2003, 2004; Delgado et al., 2004; Nicola et al., 2004; Leknes and Tracey, 2008). Our results showed increased ventral striatal activity in anticipation of fewer shocks, which suggests that the striatum is not simply functioning to prime the system to avoid pain – i.e., an analgesic effect (Scott et al., 2006, 2007; Wood and Holman, 2009). If that were the case, we would expect to see increased striatal activity to more potential shocks. Instead, we observed the opposite trend, precluding an analgesic explanation.

Although the aforementioned discussion pertains to the role of the striatum in relative valuation, we also find evidence for such signals in cortical regions classically associated with pain and punishment evaluation (Bechara et al., 1998; O'Doherty et al., 2001; Koyama et al., 2005; Kringelbach, 2005; Raij et al., 2005; Seymour et al., 2005). These regions appear to signal valuation in an inverse manner from the striatum, with both systems operating in synchrony during the decision period. Indeed, evidence for the co-existence of both appetitive-valuation and aversive-valuation signals in the brain exists, with the aversive-valuation signals residing in some of the same regions that we

observe, namely in the lateral OFC and genual ACC (O'Doherty et al., 2001; Small et al., 2001; Gottfried et al., 2002; Seymour et al., 2005; Nitschke et al., 2006). In our study, the lateral OFC and genual ACC convey this valuation information during decision-making itself over painful stimuli, as opposed to only during passive learning tasks, which build on prior evidence for these structures roles in signaling bad outcomes, perhaps to facilitate reversal-learning or changes in action, as has been suggested (Kringelbach and Rolls, 2003; Seymour et al., 2005). Furthermore, given past research showing lateral OFC and genual ACC activation to non-painful but aversive stimuli, such as monetary loss (O'Doherty et al., 2001; Liu et al., 2007) and unpleasant odors (Gottfried et al., 2002; Rolls et al., 2003), this information is likely coded in a "common currency", as has been suggested of activity in the orbitofrontal-striatal system (Montague and Berns, 2002; Murray et al., 2007).

In addition to valuation, the increase in BOLD response to worse gambles that we observed could be related to attention or cognitive control in general, which refers to the process by which attention, memory, and other cognitive abilities are shifted to accomplish a variety of goals. In addition to the lateral OFC and ACC, we found that the DMPFC signaled worse gambles with above-baseline activation in a location that has been recently implicated in decision-related control (Venkatraman et al., 2009b), and that has been shown to be more active for more difficult decisions and for decisions that run counter to overall behavioral strategy (Paulus et al., 2002; Zysset et al., 2006; Hampton and O'Doherty, 2007). Though our experimental design does not allow us to separate these functions from valuation or vice-versa, it is likely the case that there exists a complex interplay between regions signaling aversive valuation, such as the lateral OFC,

and higher-level decision-control regions which integrate these signals, possibly the DMPFC. Much like the striatum, activity in the DMPFC has been demonstrated during decision-making over a variety of stimuli, suggesting that its role might be independent of the type of outcome that is being decided on (Rushworth et al., 2005; Pochon et al., 2008; Venkatraman et al., 2009a).

The current body of research in decision-making points to the idea of a universal valuation system that signals how "good" or "bad" a potential outcome is, relative to some reference point. Structures that were originally thought to be involved solely in reward processing during decision-making are increasingly being shown to be involved in the processing of punishing stimuli as well. Similar activity in these orbitofrontal-striatal regions is observed between more abstract punishments (e.g., monetary losses) and painful stimuli as we have shown here, much like the similar activation patterns for a wide variety of rewarding stimulus modalities. Future research might focus on how a baseline is determined for this valuation activity and whether it is directly related to the status quo, and whether loss aversion can be observed for non-monetary outcomes once a status quo has been set.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Acknowledgments

Supported by grants from NIDA (R01 DA016434 and R01 DA025045).

Figures and References

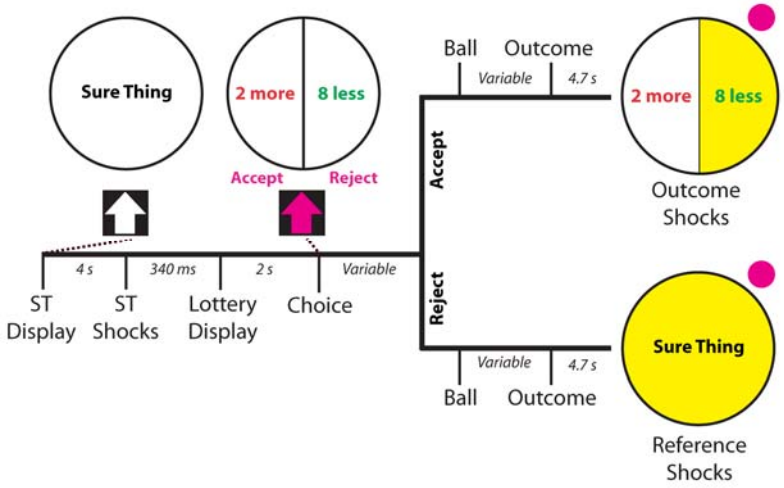

Figure 2-1: Schematic and timing of experimental task. Participants were given a status quo number of shocks (10 shocks) indicated by a circle with the text "Sure Thing", followed by the presentation of a 50/50 gamble of more/ less electric shocks from a reference amount. After 2 s of presentation of the gamble, participants could accept or reject the gamble. If they accepted the gamble, a pink ball flipped between outcomes for a period varying between 3 and 6 s, and landed with a 50/50 probability on either outcome, which turned yellow upon selection. If they rejected the gamble, the same presentation appeared. However the outcome that appeared in that case was always the reference quantity of shocks. After 4.7 s (two scanner repetition times), the outcome shocks were administered followed by a further 3 s of the outcome display. The ITI remained constant throughout the experiment at 3 s.

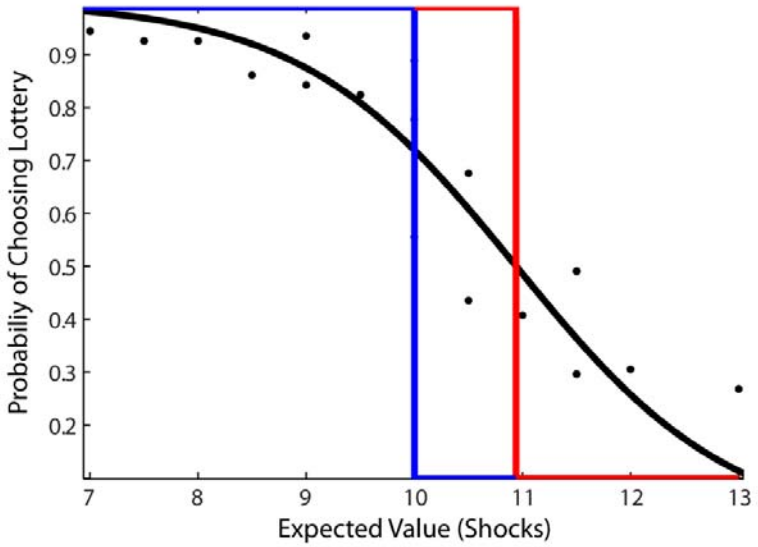

Figure 2-2: Risk-seeking behavior. To analyze risk attitude, a sigmoid curve was fit to the lottery choice data using a logistic function. An indifference point at the status quo characterizes risk-neutral behavior. Participants' actual estimated indifference point was at a lower expected value than the status quo, indicated by the red correspondence, which demonstrates risk-seeking behavior, typical in the realm of losses (Kahneman and Tversky, 1979). The reference point of 10 shocks (−10) did not lie within the 95% confidence interval of the logistic fit ($f(-10) = 0.720 \pm 0.214$), suggesting that this observed indifference point was significantly different from risk-neutrality.

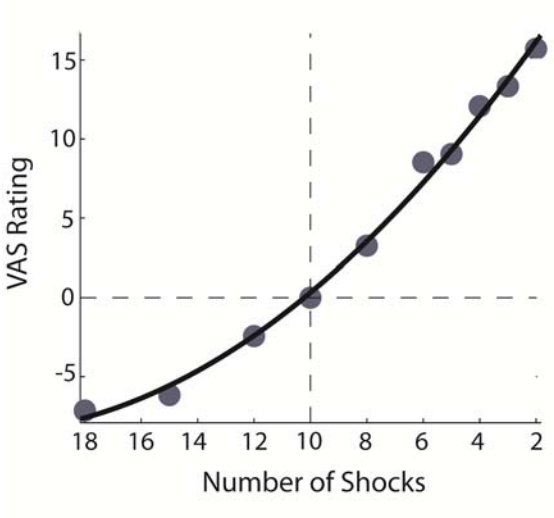

Figure 2-3: Average normalized VAS ratings as a function of the number of shocks received. A second-order polynomial function was fitted to the data in order to demonstrate the convexity of the observed ratings ($R^2 = 0.995$). The reference shocks (10 shocks) are indicated by the lines at the origin. A paired-samples *t*-test revealed that the slope for less electric shocks ($M = 1.974$, SD = 0.644) was significantly greater than the slope for more electric shocks ($M = 0.920$, $SD = 0.548$), $p < 0.001$, consistent with a convex value function.

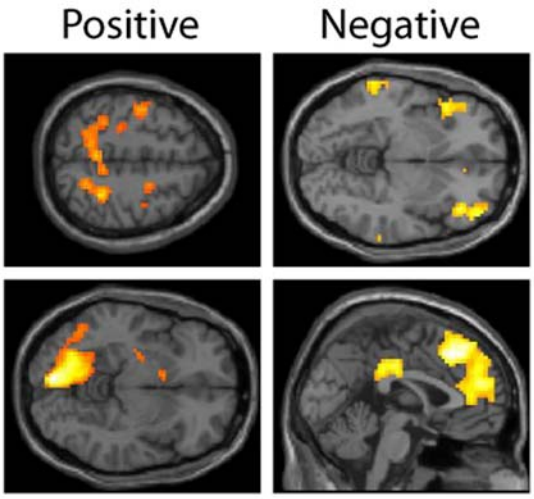

Figure 2-4: Lottery × EV_{gamble} response. Regions whose BOLD activity responded in an increasing manner to less negative outcomes are shown in the left column, and regions whose BOLD activity responded increasingly to more negative outcomes are shown in the right column. Regions shown in the positive contrast include the frontal eye fields, intraparietal sulcus, left visual area, and ventral striatum. Regions shown in the negative contrast included the left OFC, DMPFC, genual ACC, and posterior cingulate cortex. The contrasts were thresholded at $p < 0.05$, FDR-corrected (effectively a voxel-wise alpha level of 0.002), and with a cluster-threshold of 10.

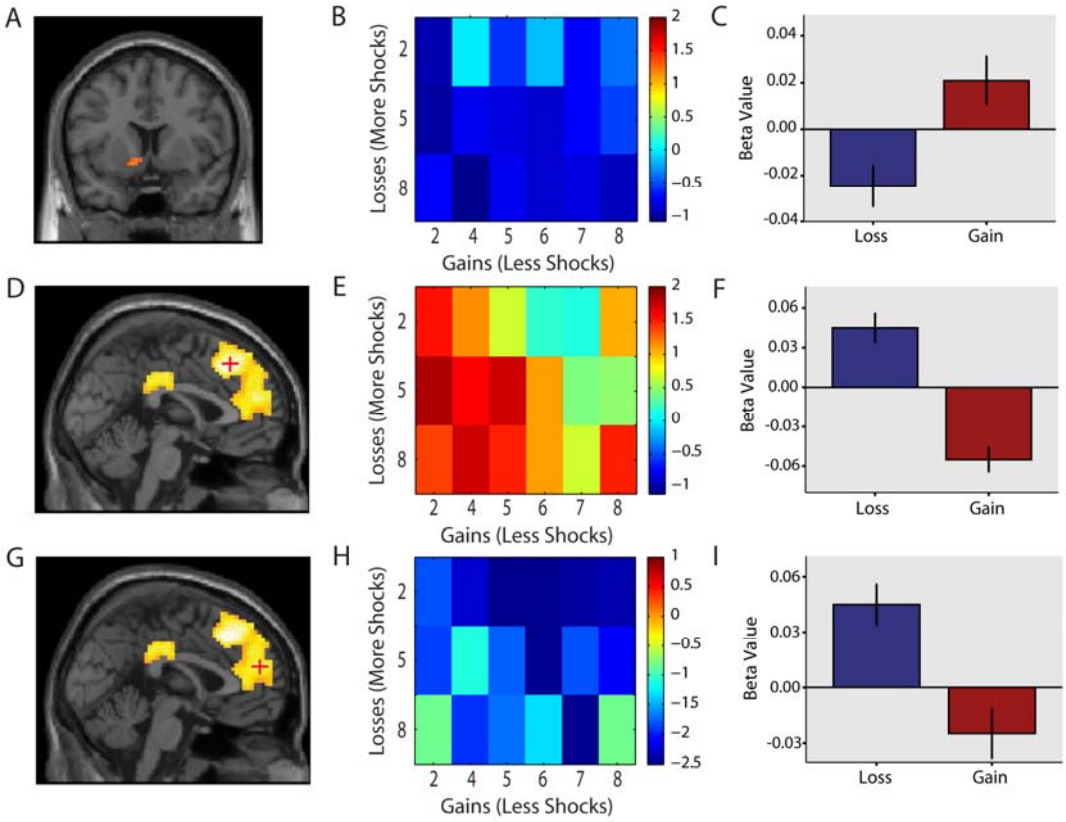

Figure 2-5: Ventral striatum, DMPFC, and genual ACC activity during the lottery period. The ventral striatum showed less deactivation to better gambles (less negative EV_{gamble}), as seen in the whole-brain EV_{gamble}analysis at $p < 0.05$, FDR (effectively a voxel-wise alpha level of 0.002), with a cluster threshold of 10 voxels **(A)**. A heat map of ventral striatum activity for each gamble type was generated by taking the ventral striatum ROI (14 voxels) from A and extracting BOLD estimates for each gamble type **(B)**. Activity in the ventral striatum showed a significant positive correlation with the number of shocks less than the reference amount and significant negative correlation with the number of shocks more than the reference amount **(C)**. The DMPFC and genual ACC showed increasing activity to worse gambles (more negative EV_{gamble}), as seen by the whole-brain EV_{gamble} analysis at $p < 0.05$, FDR, with a cluster threshold of 10

voxels **(D**,**G)**. A heat map for activity in the DMPFC and genual ACC is shown in **(E)** and **(H)**. Activity in the DMPFC and genual ACC showed a significant positive correlation with the number of shocks more than the reference amount **(F,I)**, and a significant negative correlation with the number of shocks less than the reference amount **(F)**. Beta values for the DMPFC and genual ACC were extracted from an 8-mm sphere ROI centered on the peak voxel for that cluster (marked with a red cross).

Structure	L/R	Voxels	X	Y	Z	Max T
EV_{gamble} Positive						
Primary Visual	L	1250	-6	-84	-6	8.59
Pre-Motor	L	143	-36	-9	48	5.47
Intraparietal Sulcus	R	410	30	-45	54	5.31
Pre-Motor	R	82	42	-3	48	4.1
Frontal Eye Fields	R		15	-3	66	3.37
Cuneus	R	28	18	-90	18	4.67
Frontal Eye Fields	L	29	-24	-6	63	4.34
Cerebellum	R	15	24	-51	-18	4.33
Insula	R	16	39	-21	21	3.92
Ventral Striatum	L	14	-12	9	-9	3.85
Superior Temporal Gyrus	L	19	-45	-39	12	3.79
Mid-Cingulate Gyrus	R	13	12	-18	48	3.7
EV_{gamble} Negative						
Anterior Cingulate/Middle Frontal		1467	0	30	45	6.07
Anterior Insula/Inferior Orbital	R	152	-42	24	-12	5.94
Posterior Cingulate		103	3	-24	36	4.87
Angular Gyrus	R	229	39	-66	54	4.82
Angular Gyrus	L	38	-39	-60	57	4.44
Middle Temporal Gyrus	L	43	-66	-42	-6	4.34
Middle Frontal Gyrus	R	40	39	15	57	4.12
Angular Gyrus	L	16	-33	-75	48	4.05
Precuneus	R	20	6	-63	39	3.91
Middle Frontal Gyrus	R	11	45	36	30	3.52

Table 1. Regions showing BOLD activity that correlates with the expected value of the gamble (EV_{gamble}) during the lottery period.

Chapter 3

Neural Insensitivity to Upticks in Value is Associated with the Disposition Effect

Introduction

The disposition effect is a behavioral phenomenon in which investors "sell winners too early and ride losers too long" (Shefrin and Statman, 1985). There is robust field evidence that investors sell shares of stock significantly more often after an increase in the value of the shares, than they do after a decrease in value (Ferris et al., 1988; Odean, 1998; Grinblatt and Keloharju, 2001; Choe and Eom, 2009). This occurs despite the behavior being suboptimal in terms of profit earning (Odean, 1998; Choe and Eom, 2009). In experimental settings, where complex market forces are not present and participants are not experienced investors, the disposition effect still exist (Weber and Camerer, 1998; Da Costa et al., 2008; Vlcek and Wang, 2008). Even when decision-makers are aware that losing assets are more likely to continue falling, they hold onto these assets more than they hold onto rising assets (Weber and Camerer, 1998). These observations suggest that the disposition effect is a common property of decision-making and not simply a product of the stock market or limited to investors.

Several explanations for the disposition effect have been suggested. The most common is preference-based: namely, that people are risk-averse over gains and risk-seeking over losses (Kahneman and Tversky, 1979; Shefrin and Statman, 1985; Lakshminarayanan et al., 2010). Assuming that the purchase price of an asset is the reference point upon which investors judge gains or losses, then risk-aversion over gains

would lead to the selling of assets when the value of the shares rise (the less risky option), while risk-seeking over losses would lead to holding assets when the value of the shares fall (the more risky option). This explanation is based on prospect theory, which assumes a value function that is concave over gains, but convex over losses (Kahneman and Tversky, 1979). Although risk-preference is often suggested as the driving force behind the disposition effect, several theoretical papers and experimental studies have questioned whether a preference-based explanation can fully explain the disposition effect (Vlcek and Wang, 2008; Barberis and Xiong, 2009; Kaustia, 2010). Alternatively, the realization utility hypothesis suggests that investors receive utility from the act of realizing a gain or loss, driving investors to sell gains to receive positive utility and hold losses to avoid negative utility (Barberis and Xiong, 2008). Another explanation suggests that investors have an irrational belief in mean reversion (Barberis and Thaler, 2003). This explanation predicts that an investor would hold onto a losing asset with the expectation that it would rise and sell a gaining asset with the expectation that it would fall. Despite it being an important phenomenon in finance, there is no consensus in the behavioral economics literature that mean reversion drives the disposition effect, although some experimental studies have found evidence for such a relationship (Andreassen, 1988; Hung and Yu, 2006). Hence, the question remains as to whether the disposition effect is driven by asymmetric risk-preferences over gains and losses, realization utility, or by belief-related mechanisms such as mean reversion.

In this paper, we utilize functional magnetic resonance imaging (fMRI) during an asset-trading task to test these alternative theories of the disposition effect. We measured the blood oxygenation level dependent (BOLD) response in valuation regions of the brain

during decisions to keep or to sell an asset that followed a random walk in price. A preference-based explanation for the disposition effect would predict correlations between individuals' risk-preferences, the magnitude of their disposition effect, and activation in valuation structures of the brain. A realization utility explanation would predict differential responses in valuation regions during the decision to sell versus keep an asset that correlate with the magnitude of the disposition effect. Finally, if participants believe that the asset price will eventually revert to the mean, we would predict an attenuated ventral striatum response to upticks in value below the purchase price and a greater response to upticks in value above the purchase price for individuals with a disposition effect. The ventral striatum has been shown to signal prediction error, and thus an expectation of a rise in asset price followed by an uptick should result in lessened striatal activity. Conversely, if an expectation of a fall in asset price is met by an uptick, striatal activity should increase. Of these three theories, the fMRI data were most consistent with an irrational belief in mean reversion.

Materials and Methods

Thirty-eight participants (18 female; 18–51 years) were recruited from the Emory University campus and completed the asset-trading task. Of this group, thirty-three participants were scanned using fMRI (17 female; 18–51 years). Of the thirty-three scanned participants, we excluded one participant from our imaging model because of excessive motion and five participants because they lacked observations for the regressors of interest. This was due to a high variability in participant behavior that was tied to the number of observations in each of the regressors in our model. All participants

were right-handed, reported no psychiatric or neurological disorders, or other characteristics that might preclude them from safely undergoing fMRI. All participants provided informed consent to experimental procedures approved by the Emory University Institutional Review Board. Prior to beginning the experiment, participants were told that they would be paid $20 for showing up, and $30 for completing three questionnaires, including the BIS/BAS, EPQR, and a risk-preference worksheet. The total amount of $50 had to be used in the subsequent asset-trading task, where participants could earn an additional $50, or lose $50. Therefore the total possible compensation ranged from $0 to $100 (actual $30–$75).

Asset trading task

Participants completed 40 trials of an asset-trading task while undergoing fMRI. The main screen of the asset-trading task consisted of a graph with relative value (in dollars) of an asset on the *y*-axis, and time (in periods) on the *x*-axis. For each trial, participants were initially forced to purchase an asset worth $50 using all of the money they earned prior to entering the scanner ($50). A gray circle at relative value zero and period 10 indicated the purchase point. To make sure that the current decision period was always centered on the screen, each asset had a 10-period history prior to purchase that was different for all 40 assets. Each history was generated in the same manner as the forward asset price trajectory—by generating a random walk beginning at the purchase price, but in this case going backwards for ten periods. After a button press, the asset subsequently increased or decreased in value by $5 with equal probability.

Participants were then given the choice of keeping the asset for another period, or selling the asset for its worth in that period. If participants chose to keep the asset, it again increased or decreased by $5 with equal probability after a 3 second delay. A single decision period is shown in **figure 3-1** (panel A). To avoid influencing participant decision-making by having a finite number of periods where participants potentially behave differently towards the end of the trial, we implemented an infinite horizon. This was accomplished by using a 'soft' ending, where the trial ended and asset force-sold with a 5% probability each time the participant chose to keep the asset (Camerer and Weigelt, 1993; Noussair and Matheny, 2000). Participants could therefore keep the asset as long as they wanted, keeping in mind that each trial had a 5% probability of ending each time they kept the asset. If participants chose to sell the asset, they earned what the asset was worth in that period. Each sell period was followed by five periods where they saw what trajectory the price would have followed had they of kept the asset. Each time the asset was sold or force-sold, participants saw an outcome screen stating the relative amount that their asset was sold for. The maximum relative value that the asset could reach was + $50, and minimum of −$50.

Participants were given full information regarding the determination of asset price and the infinite horizon prior (see **supplement 1**). Participants were verbally tested on the probabilities and independence of the asset pricing, and completed two practice trials prior to the actual asset-trading task. Note that because the size and probability of an increase or decrease in asset value was always equal and independent across periods, the expected value of keeping the asset was the same as selling the asset at every decision point. Furthermore, because of the random-walk nature of the asset price, there is no

optimal selling strategy that maximizes earnings in the task. Therefore, no matter what strategy is taken, participants on average earn a relative value of $0 across trials ($50).

Behavioral data analysis

To measure the disposition effect for each individual, we computed the integral of asset value (in dollars) relative to purchase price over time and averaged across all trials. If each asset was held regardless of value and never sold, the average integral would be close to zero because the price of the asset followed a random walk (**figure 3-1;** panel B). If assets below the purchase price were held longer than assets above the purchase price, then the average integral would be negative (**figure 3-1**; panel C).

We estimated risk-preference and loss aversion parameters using a method developed by (Tanaka et al., 2010). Participants were given three series of paired lotteries. In each series, participants were asked to choose the point at which they would switch from lottery A to lottery B. The expected value of lottery B increased downward in each series. By solving a system of inequalities in which constant relative risk aversion is assumed ($U(x) = x^{\alpha}$), a unique set of risk-preference and loss aversion parameters were estimated for each participant.

To understand the factors that were most important in driving participants' decisions to keep or sell, and thus create an imaging model most relevant to participant behavior, we estimated a mixed-effects logistic regression with “keep” or “sell” for each period as the outcome variable (1 for sell, 0 for keep), five fixed-effects regressors which described local asset price characteristics, and a subject-wise regressor describing the

magnitude of a participant's disposition effect. Subject was included as a random-effects factor. The regression took the following form:

$$\text{logit}(\text{Sell}_t) = \beta_0 + \beta_1 \text{Value}_t + \beta_2 \text{Value}'_t + \beta_3 \text{Value}''_t + \beta_4 \text{DE} + \text{interactions}$$

Value_t was the value of the asset at time t (which ranged from − 50 to 50 in multiples of 5). Value'_t was the difference in value from time t to $t - 1$ (which could carry values − 5 or 5). We subsequently refer to value'_t as "delta" throughout this paper. Value''_t was the change in price direction, calculated the following way: $(\text{value}_t - \text{value}_{t-1}) - (\text{value}_{t-1} - \text{value}_{t-2})$. Thus, value''_t can carry values of 10, 0, or − 10. DE was the average integral across trials for each participant, which was a measure of the individual's disposition effect. We also included interaction terms between DE and the main effects in order to probe where people with a disposition effect behaved differently.

fMRI data acquisition and analysis

Functional imaging was performed with a Siemens 3 T Trio whole-body scanner. T1-weighted images (TR = 2600 ms, TE = 3.02 ms, flip angle = 8°, 240 × 256 matrix, 176 sagittal slices, 1 mm^3 voxel size) were acquired for each subject prior to the four experimental runs. For each experimental run, T2*-weighted images using an echo-planar imaging sequence were acquired, which show blood oxygen level-dependent (BOLD) responses (echo-planar imaging, TR = 2000 ms, TE = 30 ms, flip angle = 73°, FOV = 192 mm × 192 mm, 64 × 64 matrix, 33 3.5-mm thick axial slices, and 3 × 3 × 3.5 mm voxels).

fMRI data were analyzed using SPM8 (Wellcome Department of Imaging Neuroscience, University College London) using a standard 2-stage random-effects

regression model. Data were subjected to standard preprocessing, including motion correction, slice timing correction, normalization to an MNI template brain and smoothing using an isotropic Gaussian kernel (full-width half-maximum = 8 mm). Task regressors were modeled using a standard canonical hemodynamic response function. Our model focused on the decision period, which encompassed the time at which the outcome of the asset increase or decrease was shown, to the time in which a participant made a response. To attempt to capture the response only to the asset price increase or decrease outcome, we also created an alternative model using an impulse function at the moment the asset increase or decrease was shown, leaving the rest of the decision period in the model baseline. In both models, we separated these periods into five separate regressors: the first period in the trial where participants hit the 'continue' button to see whether the asset increased or decreased in value, decisions above the purchase price, decisions below the purchase price, decisions at the purchase price, and the periods after an asset was sold. Decisions above and below the purchase price were modulated by choice (keep vs. sell), value_t, and delta ($\text{value}_t - \text{value}_{t-1}$). We also included a regressor for the outcome screen. The 3-second interval where the participant's choice was highlighted was randomly assigned to four regressors of no-interest to prevent correlation with the constant. Because individual behavior varied greatly, and was reflected in the number of observations in each regressor, we concatenated all runs in our model together, and included $n - 1$ dummy variables to account for run differences. This allowed us to avoid complex contrast weights for each participant, given the high likelihood of missing regressors for multiple runs. Six motion parameters and a single constant were also included for each subject.

Statistical thresholds were determined based on the smoothness of the second-level contrasts. We used AlphaSim, a routine in AFNI, to estimate the combination of height and extent thresholds that yielded a whole-brain FDR < 0.05 (10,000) iterations. To do this, we used 3dFWHMx to estimate image smoothness from the square root of the masked SPM-generated residual (ResMS) image. We then used AlphaSim with a voxel-level threshold of $p < 0.001$ to find the cluster threshold at which the whole-brain FDR would be < 0.05, and which ranged from 42 to 57 voxels, depending on the specific contrast.

Results

We found evidence for a robust disposition effect in our task. As a measure of the disposition effect, we took the integral of asset price averaged across trials. Across all participants (38 including behavioral and scanned), the average integral was − 27.7 $*periods ($\sigma = 31.4$, range = − 113.6–18.9), which was significantly different from zero [$t(37) = -5.4$, $p < 0.001$]. This indicates that participants spent significantly more time below the purchase price than above the purchase price across trials. The majority of participants displayed the disposition effect (31 of 38 had integrals less than zero). As another measure of the disposition effect, participants sold the asset when it was below the purchase price significantly less often (mean = 6.5 trials, $\sigma = 5.6$) than when the asset was above the purchase price (mean = 14.4 trials, $\sigma = 5.5$; $t(37) = 5.4$,$p < 0.001$]. Furthermore, the average number of periods the asset was kept above the purchase price (mean = 6.3 periods, $\sigma = 3.0$) was significantly lower than the average number of periods below the purchase price [mean = 9.3, $\sigma = 5.2$; $t(37) = -4.9$, $p < 0.001$].

The magnitude of the disposition effect varied widely across participants, ranging from a large disposition effect (integral of − 113.6 $*periods) to an anti-disposition effect (integral of 18.9 $*periods). The heterogeneity in the disposition effect could be driven primarily by individual differences in gain-selling behavior, loss-holding behavior, or both. The mean integral was significantly correlated with the average period in which an asset was sold or force-sold below the purchase price ($r = -0.73$, $p < 0.001$), but not the average period in which an asset was sold or force-sold above the purchase price ($r = -0.13$, $p = 0.43$). This suggests that the heterogeneity in the disposition effect was driven primarily by the length of loss holding. For this reason we focused the imaging analysis on decision-making below the purchase price, i.e. in the loss domain.

A preference-based explanation of the disposition effect would suggest that risk-preference should be correlated with the magnitude of the disposition effect. Overall, our risk-preference questionnaire revealed that participants were slightly risk-averse, with a mean $\alpha = 0.91$ ($\sigma = 0.34$), which was consistent with past studies on risk preference (Tversky and Kahneman, 1992). The mean loss aversion parameter, however, was higher ($\lambda = 3.8$, $\sigma = 3.6$) than is typically reported in the literature ($\lambda \sim 2$). This was driven primarily by several outliers with unusually high estimates for the loss aversion coefficient. The median value of $\lambda = 2.7$ was closer to that found in the literature. We found no correlation between the risk-preference parameters, the loss aversion coefficients, and the disposition effect, providing further behavioral evidence against a preference-based explanation of the disposition effect (Vlcek and Wang, 2008).

Participants were not asked in a structured manner about their beliefs in the task, however casual conversation during debriefing revealed that a large portion of

participants expressed a 'hope' that the asset would return to the purchase price when at a loss, despite being aware of the random-walk nature of the price. It is important to note that self-reporting of the motivation behind investment decisions might not accurately reflect the underlying process driving decisions. More formally, psychological aspects of participant behavior were gathered through BIS/BAS and EPQR questionnaires. We found no significant correlations between these personality measurements and the disposition effect. A negative correlation between BAS reward and the disposition effect approached significance ($r = -0.32$, $p = 0.06$), however using Bonferroni correction for multiple comparisons makes the trend much less compelling.

In order to build an imaging model that best reflects participant behavior, we ran a logistic regression with the decision to sell or keep as the outcome variable, and characteristics of the asset price as regressors. We found that for any given period there was a low probability of selling (**table 3-1**). The largely negative constant represents the fact that for each trial, the participant could sell on only one period (what would be the final one). Consequently, the default prediction, with all other values set to zero, was that a participant would keep the asset for any particular period. We found that three regressors significantly ($p < 0.05$) predicted participant choice behavior: the value of the asset, the change in value of the asset from the previous period, and the change in price direction. These three characteristics of the asset price were subsequently included in our imaging model.

Imaging

Our neuroimaging model was based on two behavioral observations: 1) that heterogeneity in the disposition effect was mostly due to decisions below the purchase price; and 2) that the regressors in our behavioral model that described local asset price characteristics ($value_t$, $value_t'$, $value_t''$), predicted participant keep/sell behavior. We therefore created a model focusing on the decision period, which was divided into decisions above the purchase price (gains), and decisions below the purchase price (losses). These decision periods were modulated by participant choice (keep vs. sell), value, and delta. We report coordinates of activation and their corresponding statistics in **table 3-2**. As would be expected, a simple contrast between decisions over gains versus losses revealed robust ventral striatal activation. Ventral striatal activity was not significantly correlated with overall asset price relative to the purchase (neither in the gain nor loss domain). Instead, we found that the ventral striatum responded to upticks in delta both above and below the purchase price, which was the difference between the current asset price, and the price in the previous period (**figure 3-2**). This suggests that a change in asset price from the previous period best explains variation in the BOLD response during the decision-making period within the ventral striatum, as opposed to the total present value of the asset.

Because of the evidence implicating the striatum in the process of valuation during decision-making, and its role in explaining individual heterogeneity in decision-making strategies (Venkatraman et al., 2009a), we extracted beta values from the left and right ventral striatum and putamen clusters from the aforementioned delta contrast for decisions below the purchase price (right panel; **figure 3-2**). These beta values describe

the slope of the relationship between the BOLD response and delta and were correlated with the magnitude of the disposition effect; both correlations approached significance ($r = 0.367$, $p = 0.055$; $r = 0.360$, $p = 0.057$, left and right respectively). However, in an alternative model where an impulse function was used to model the immediate response to an increase or decrease in asset value (as opposed to the entire decision period), we find significant correlations in the left and right ventral striatum ($r = 0.384$, $p = 0.044$ and $r = 0.395$, $p = 0.037$ for left and right, respectively).

This suggests that our results are likely driven by immediate outcome response as opposed to deliberation over keeping vs. selling the asset. Participants who did not have a disposition effect displayed positive striatal responses to upticks in value, but participants who had large disposition effects showed little striatal response (**figure 3-3**; left ventral striatum/putamen cluster). This relationship was not seen for decisions above the purchase price, where the BOLD response was significantly modulated by delta but not correlated with the magnitude of the disposition effect ($r = 0.252$,$p = 0.187$). We found no correlation between the measured utility curvature (indicative of risk-preference) of individual participants and overall activity in the left and right ventral striatum/putamen below the purchase price ($r = 0.09$, $p = 0.64$ and $r = 0.10$, $p = 0.60$ for both, respectively), as well as above the purchase price ($r = -0.18$, $p = 0.36$). Whole-brain correlation analysis revealed no brain regions whose response to delta correlated with utility curvature, above or below the purchase price.

To test whether the disposition effect was driven by a differential response to realized (sold) versus paper (kept) gains and losses, we contrasted sell versus keep decisions (averaged across gains and losses). We found a significant difference for this

contrast in bilateral insula, DLPFC, thalamus and ACC, among other regions (**figure 3-4**; **table 3-1**), areas consistently activated for risk-averse choices and decisions to quit loss-chasing behavior (Paulus et al., 2003; Campbell-Meiklejohn et al., 2008; Preuschoff et al., 2008). However, none of these regions were significantly correlated with the magnitude of the disposition effect. We found no activation in typical reward-related regions, such as the ventral striatum or medial prefrontal cortex (MPFC) that was greater for sell than keep decisions (even at a liberal threshold of $p < 0.05$, $k \geq 10$). We also found no regions significantly more active for keep than sell decisions at our statistical threshold ($p < 0.001$,$k \geq 57$). However, at a more liberal threshold ($p < 0.005$, $k \geq 10$) we found greater activation in the ventral striatum and MPFC for keep versus sell (Kuhnen and Knutson, 2005; Campbell-Meiklejohn et al., 2008). Neither the ventral striatum nor MPFC activation from the [keep − sell] contrast correlated with the magnitude of the disposition effect ($r = -0.23$, $p = 0.26$, and $r = 0.24$, $p = 0.24$, for ventral striatum and MPFC, respectively). Even when separated for decisions above and below the purchase price, whole-brain correlation analysis on [keep − sell] and [sell − keep] contrasts revealed no brain regions whose activity significantly correlated with the disposition effect or with participant risk-preference.

Discussion

The disposition effect is a robust behavioral phenomenon in which individuals hold onto losing assets longer than they hold onto gaining assets. Here, we tested three theories on why people display a disposition effect using an asset-trading task and fMRI: a risk-preference, realized utility, and an irrational belief in mean reversion. We found

three pieces of evidence suggesting that a preference-based explanation of the disposition effect does not fully explain the phenomenon. First, we found that a behavioral measure of risk-preference did not correlate with the magnitude of the disposition effect. Second, we found no correlation between measured risk-preference parameters and BOLD response in valuation regions. Third, we found no correlation between the disposition effect and activation in regions previously shown to be correlated with risk attitude, namely the MPFC. Similarly, we did not find a relationship between the disposition effect and differences in keep/sell activations in reward-related brain regions, as would be predicted by differences in realized utility. Instead, our results are more consistent with a mean reversion hypothesis. We show that for decisions below the purchase price, a greater disposition effect is correlated with a blunted ventral striatum response to upticks in value in some individuals. Given the established role of the ventral striatum in signaling reward-prediction errors (RPE) (Pagnoni et al., 2002; Schultz, 2010), this blunted response is consistent with meeting an expectation of an uptick towards the mean.

Behaviorally, our data did not point towards a risk-preference view of the disposition effect. According to this explanation, the disposition effect reflects an asymmetry in risk-preference across gains and losses (Kahneman and Tversky, 1979; Lakshminarayanan et al., 2010). Risk-preference is often measured using choice-based paradigms that estimate utility function curvature (Binswanger, 1980; Tversky and Kahneman, 1992; Holt and Laury, 2002). Often these measurements are tied to individual variation in aspects of financial and non-financial decision-making, suggesting validity in their ability to predict real-world decision-making (Barsky et al., 1997;

Anderson and Mellor, 2008). We used a similar choice paradigm that estimated both utility function curvature (risk-preference) and loss aversion (Tanaka et al., 2010). Consistent with past literature, we found that our participants were both risk-averse and loss-averse on average (Tversky and Kahneman, 1992). However, neither curvature of the utility function nor loss aversion correlated with the disposition effect in our task, which builds on evidence that risk attitude is not driving the disposition effect (Vlcek and Wang, 2008).

Our imaging data further suggest that risk-preference does not fully explain the disposition effect. A large body of evidence has demonstrated that a network of brain regions, including the ventral striatum, MPFC, and orbitofrontal cortex (OFC) is closely associated with the economic concept of expected utility and the behavioral concept of reward (Knutson et al., 2001a; Montague and Berns, 2002; Berridge and Robinson, 2003; Delgado, 2007; Plassmann et al., 2007). Activity in these regions scale with both the expected and subjective value of stimuli (Knutson et al., 2001a; Kable and Glimcher, 2007). Moreover, the ventral striatum and MPFC are not limited to the processing of gains, but they also process expected loss during decision-making (Tom et al., 2007; Brooks et al., 2010). Individual risk-preference is directly related to the subjective valuation (i.e. utility) of potential outcomes during choice. Thus, if risk attitude played a role in our task, we should have found that activity in the aforementioned brain regions was correlated with risk-preference parameters. However, none of the relative activations between keep versus sell decisions were correlated with subject risk-preference in these regions. In addition, there were no brain regions where the BOLD response to delta was correlated with risk-preference. This suggests that subject

heterogeneity in risk-preference did not significantly contribute to decisions in our asset-trading task.

In addition to finding no valuation regions that correlated with measures of risk-preference, we found no significant correlation between the magnitude of the disposition effect and activity in these same regions. For example, activity in the ventral MPFC has been found to positively correlate with individual measures of risk seeking when risky decisions are made, while activity in the lateral OFC and dorsal MPFC positively correlates with risk-aversion when less-risky decisions are made (Tobler et al., 2007; Christopoulos et al., 2009; Xue et al., 2009). When we compared risky (keep) versus riskless (sell) decisions below the purchase price, we found no correlation between ventral MPFC, lateral OFC, or dorsal MPFC activity and the disposition effect. We found the same absence of correlation when comparing riskless (sell) versus risky (keep) decisions.

The realization utility theory suggests that investors get a jolt of positive utility when they sell a stock at a gain, and a jolt of negative utility when they sell a stock at a loss (Barberis and Xiong, 2008). In this theory, people want to avoid the pain of negative utility so they hold onto losses, and because they want to receive the pleasure of positive utility they sell gains. This hypothesis would suggest that individual differences in the disposition effect should be correlated with the magnitude of the BOLD response in valuation-related brain regions during the decision to sell versus keep an asset, as others have found (Frydman et al., 2011). In our study, we did not find such correlations. However, if the receipt of realization utility requires the act of completing a successful investing episode (including purchase and sale), it is possible that our experimental

design did not capture this because our participants did not have a choice about purchasing the assets (only selling).

Though the aforementioned lack of results supporting alternative hypothesis is not definitive, our results are more consistent with an irrational belief in mean reversion. We found a negative correlation between the magnitude of the disposition effect and the ventral striatum response to upticks in value when the asset price was below the purchase price. This finding is consistent with models of reward prediction error (RPE) in the striatum (Schultz, 1997; Hollerman and Schultz, 1998; O'Doherty et al., 2003) which suggest that the striatum responds only to unexpected upticks in value. In its simplest form, the RPE is computed by taking the difference between the reward received and the reward that was expected. At each period in our task there was a 50/50 probability of an increase or decrease in asset value. Thus, the objective expected value of the asset was always $0. Any uptick in value should therefore result in a positive RPE, and consequently a ventral striatum response, regardless of the total present value of the asset. Although one might predict a decreasing marginal striatal response as the total asset value increases, one would not predict a blunted response to upticks in value when the asset value was negative. Thus, the lack of a striatal response would suggest no RPE, even though the asset had increased in value. Previous neuroimaging studies have shown that the striatal response to reward disappears when the subject can fully predict the reward (Berns et al., 2001; Pagnoni et al., 2002; McClure et al., 2003; O'Doherty et al., 2003). Thus, the lack of striatal response to upticks in value suggests that some subjects with a large disposition effect predicted the uptick when the asset was in the loss domain, implying that they were expecting a return to the mean. Our results build on previous

behavioral research that supports an irrational belief in mean reversion (Andreassen, 1988; Weber and Camerer, 1998; Vlcek and Wang, 2008).

Acknowledgments

We thank C. N. Noussair for experimental design input and manuscript feedback and S. E. Moore for help with programming the experiment. Funding was provided by a grant from the National Institute on Drug Abuse (DA024045) through the American Recovery and Reinvestment Act (ARRA).

Figures and Tables

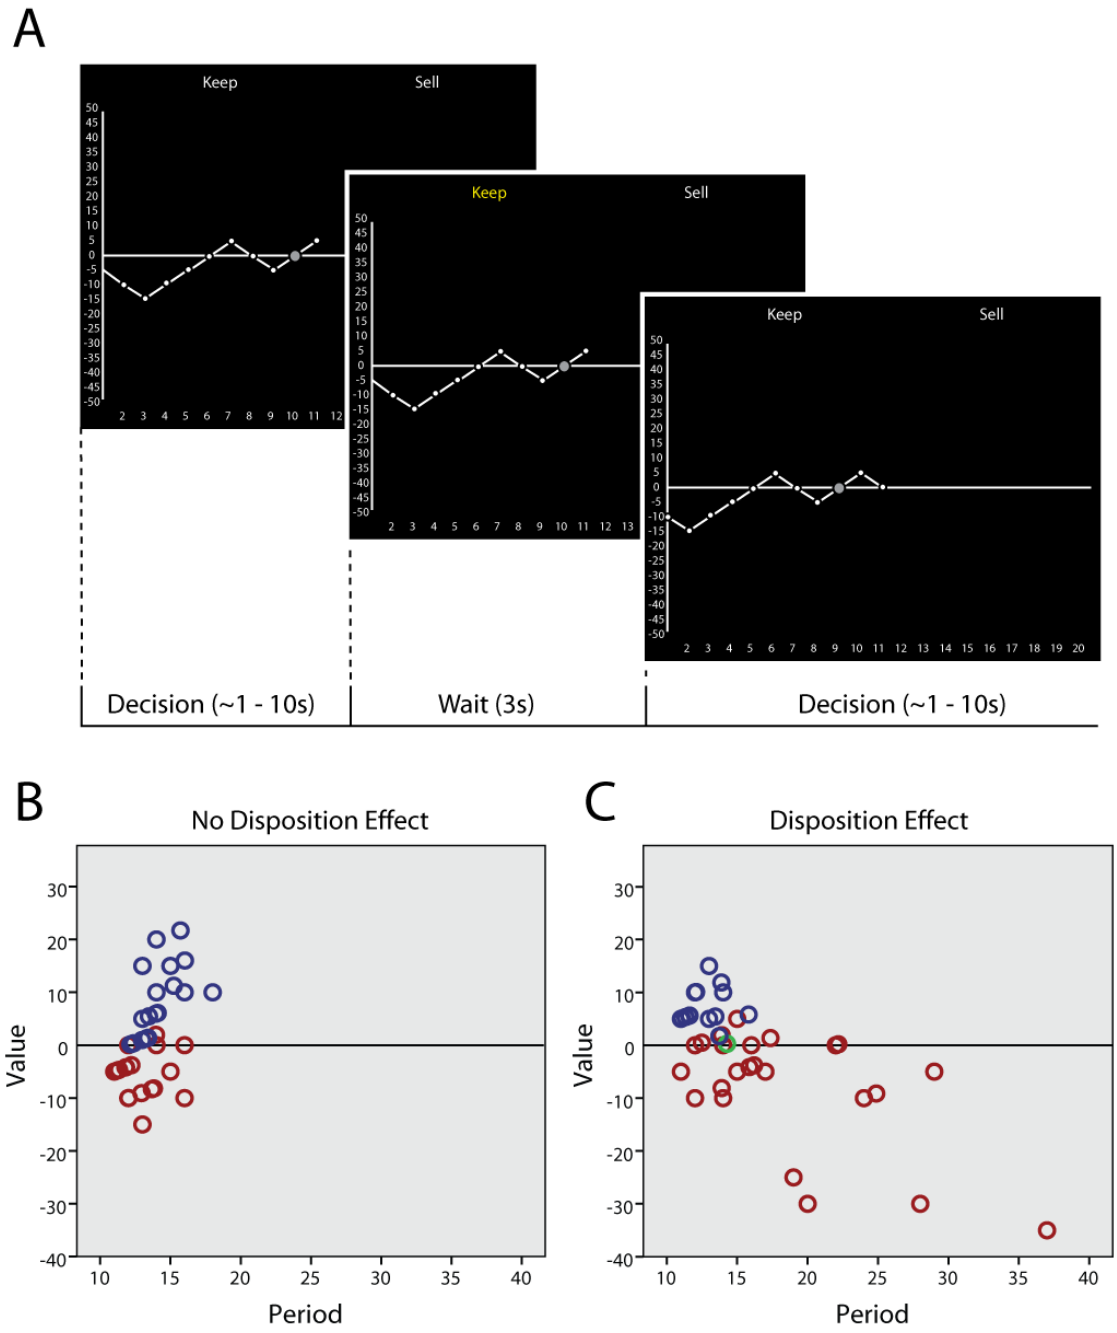

Figure 3-1: Asset trading task overview. (A) Participants were forced to purchase an asset for $50 (purchase point marked by a gray circle). The asset price then followed a random walk, always increasing or decreasing in increments of $5 with equal probability.

At each period, participants chose to keep or sell the asset (*first panel*). After the participant made a choice, the choice was highlighted for 3 seconds (*second panel*). The next period outcome was then shown and the participant could immediately make a decision (third panel). An infinite horizon was implemented, where there was a 5% probability of the trial ending (and asset being force-sold). Therefore, participants were able to keep the asset as many periods as they liked, with the caveat that each time they kept the asset, there was a 5% chance of the trial ending. (B) and (C) show data from subjects without a disposition effect (*left*; average integral of 10.4 $*periods) and with a disposition effect (*right*; average integral of − 46.8 $*periods). Each data point represents the period in which an asset was sold or force-sold (assets that were held until the trial ended). Trials in which the overall integral of asset price through time was positive are shown in blue, negative in red, and zero in green. Forty trials are shown for each participant.

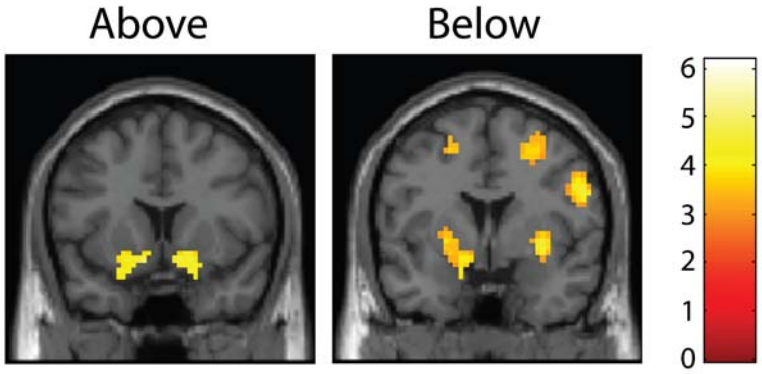

Figure 3-2: Whole-brain analysis for identifying brain regions responding to positive delta. Regions are shown whose BOLD response increased with an increase in asset price from the previous period, above (*left*) and below (*right*) the purchase price. We find ventral striatal activity both above and below the purchase price in response to upticks in value. Below the asset price, we also find DLPFC and motor area activation. Coronal sections are shown here (MNI $y = 11$ above, and $y = 2$ below). Statistical thresholding was set at $p < 0.001$, $k \geq 53$ (*above*), and $k \geq 42$ (*below*), resulting in a whole-brain FDR < 0.05 for both contrasts.

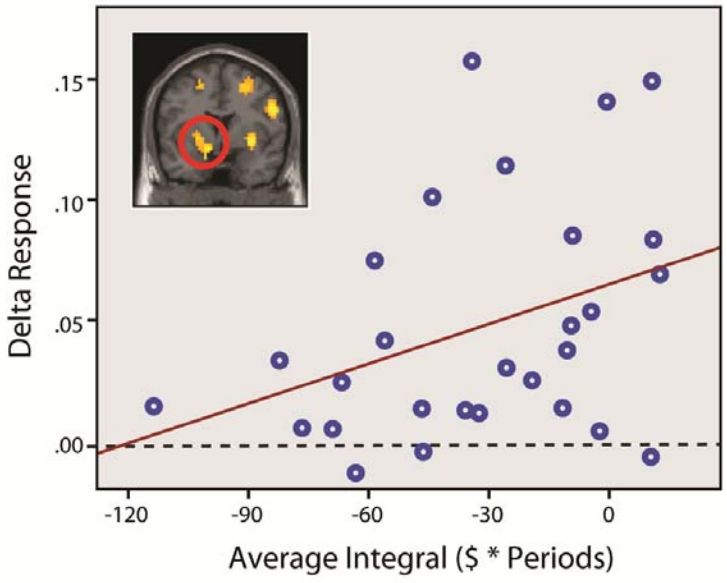

Figure 3-3: Relationship between striatal response to upticks in value below the purchase price and the magnitude of the disposition effect. Beta values in the left ventral striatum during decisions below the purchase price, modulated by the change in asset price from the previous period (delta), correlated with the magnitude of the disposition effect, as measured by the integral ($r = 0.367$, $p = 0.055$). A more negative integral is associated with a larger disposition effect. Participants who showed a larger disposition effect were more likely to show little or no response in the ventral striatum to upticks in value.

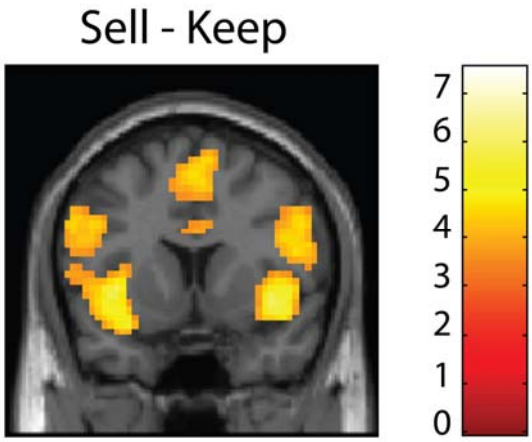

Figure 3-4: Whole-brain contrast between decisions to sell versus keep. A greater BOLD response was observed in the bilateral insula, DLPFC, and ACC during decisions to sell versus keep an asset (MNI $y = 11$). This is consistent with past research where participants quit chasing losses, but inconsistent with a realization utility explanation of the disposition effect. Statistical thresholding was set at $p < 0.001$, $k \geq 57$, resulting in a whole-brain FDR < 0.05.

		Coefficient	Std. Error	z	P > z	[95% Conf.	Interval]
$Value_t$		0.036	0.005	6.95	0.000	0.026	0.047
$Value_t'$		-0.091	0.016	-5.79	0.000	-0.122	-0.060
$Value_t''$		0.027	0.010	2.78	0.005	0.008	0.047
Integral (DE)		0.021	0.005	4.68	0.000	0.0126	0.031
$Value_t$ x Integral		-0.001	0.000	-7.78	0.000	-0.001	-0.001
$Value_t'$ x Integral		-0.004	0.000	-8.06	0.000	-0.005	-0.003
$Value_t''$ x Integral		0.001	0.000	3.78	0.000	0.001	0.002
Constant		-1.90	0.186	-10.21	0.000	-2.259	-1.531

Table 3-1: Results of keep/sell decisions regressed onto local asset price characteristics using a logistic regression. Participants had a greater propensity to sell when the asset price increased in value ($value_t$), decreased from the previous period ($value_t$'), and had a positive change in price direction (value*t*"). Value*t*" was calculated by taking the difference between (value$t$ − valuet_{-1}) and (valuet_{-1} − valuet_{-2}). A larger average integral (more negative for a greater disposition effect; DE) increased the probability of making a sell decision. We found significant interactions between the disposition effect and local asset price characteristics.

Structure	L / R	Maximum T	Cluster Size	Z	P (unc)	X	Y	Z
Decision Gain - Decision Loss								
Middle Occipital	R	10.97	174	6.84	3.9E-12	18	-94	9.5
Ventral Striatum	L	7.40	710	5.50	1.89E-08	-18	14	-8
Ventral Striatum	R	6.88	495	5.26	7.33E-08	15	14	-8
Medial Pre Frontal	R	5.83	555	4.70	1.29E-06	24	47	-8
Thalamus	L, R	5.74	78	4.65	1.64E-06	0	-7	6
Superior Parietal Lobule	L	4.64	382	3.98	3.47E-05	-33	-55	45
Superior Parietal Lobule	R	4.61	308	3.96	3.71E-05	21	-61	48
Below PP x Delta (Positve)								
Middle Occipital	R	7.12	87	5.30	5.94E-08	24	-97	9.5
Putamen/Ventral Striatum	L	5.75	243	4.61	2.06E-06	-30	-13	2.5
Superior Parietal	R	5.64	510	4.55	2.74E-06	21	-64	59
Putamen	R	5.51	135	4.47	3.89E-06	33	-7	-1
Dorsolateral PFC	R	5.28	119	4.34	7.23E-06	57	2	27
Pre-Motor	R	5.03	193	4.19	1.4E-05	27	11	55
Pre-Motor	L	4.40	57	3.78	7.72E-05	-21	-4	59
Superior Parietal	L	4.36	131	3.76	8.54E-05	-15	-70	59
Above PP x Delta (Positve)								
Cuneus	L	5.93	466	4.73	1.1E-06	-6	-82	17
Ventral Striatum	L, R	5.48	236	4.48	3.78E-06	-9	5	-12
Lingual Gyrus	R	4.23	94	3.69	0.000114	12	-82	-1
Sell - Keep								
Anterior Cingulate	L, R	7.52	1097	5.39	3.55E-08	9	29	27
Cuneus	L, R	6.73	1353	5.04	2.38E-07	12	-85	13
Insula	R	6.18	1290	4.77	9.17E-07	42	14	-12
Insula	L	5.83	360	4.59	2.22E-06	-36	11	-1
Thalamus	L, R	5.67	354	4.51	3.31E-06	6	-19	2.5
Inferior Parietal	R	5.34	226	4.32	7.77E-06	54	-40	59
Dorsolateral PFC	R	5.19	157	4.23	1.15E-05	-45	5	31
Posterior Cingulate	L, R	5.08	228	4.17	1.53E-05	6	-22	31
Inferior Parietal	L	4.94	301	4.09	2.19E-05	-54	-34	55
Cerebellum	R	4.68	150	3.93	4.30E-05	24	-52	-26

Table 3-2: Whole-brain analysis results for selected contrasts.

Supplementary Material

Asset Trading Task Instructions

This is an experiment in decision-making. If you follow the instructions carefully and make good decisions, you can earn a considerable amount of money. There are 40 trials in the experiment. However, only one of the trials will actually count toward your final payment. A random draw will decide which of the 40 trials will count.

1. For each trial, you initially see the history of an asset you have purchased. Your purchase point is represented by the grey circle at period 10. On the y-axis is the change in value of the asset, and on the x-axis is the period in which the asset is in. At each period, there is a 50/50 chance that the asset increases or decreases in value by $5. Each period is completely independent of one another: an increase or decrease in a previous period does not change the probability of an increase or decrease in the following period.

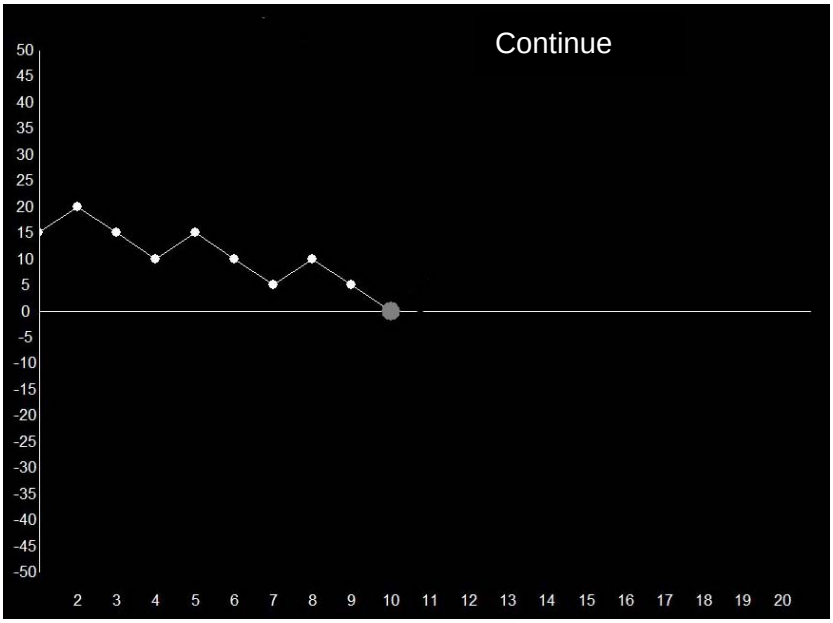

2. After pressing the left or right key, your asset will initially increase or decrease in value by $5 with a 50/50 chance.

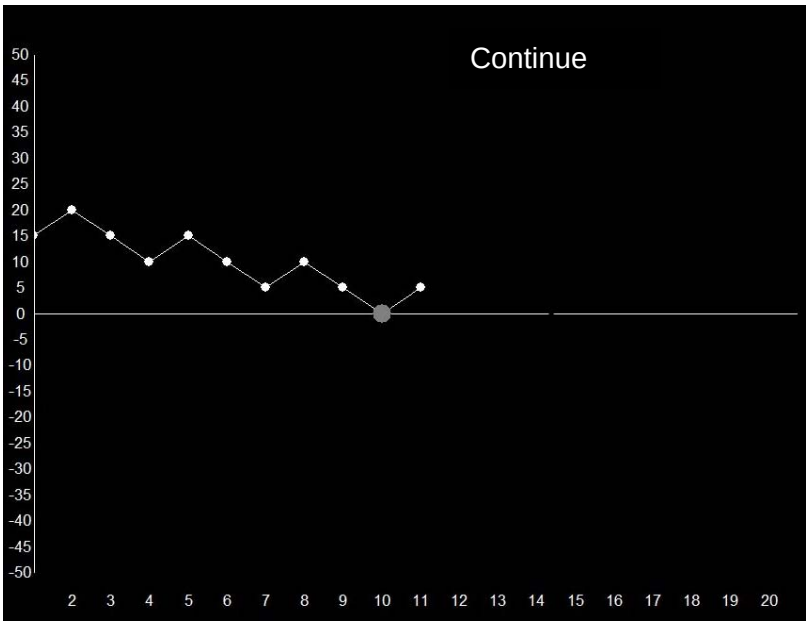

3. After the asset has increased or decreased in value, you will be asked if you want to keep or sell your asset. Press the *left arrow* key to keep the asset and the *right arrow* key to sell the asset:

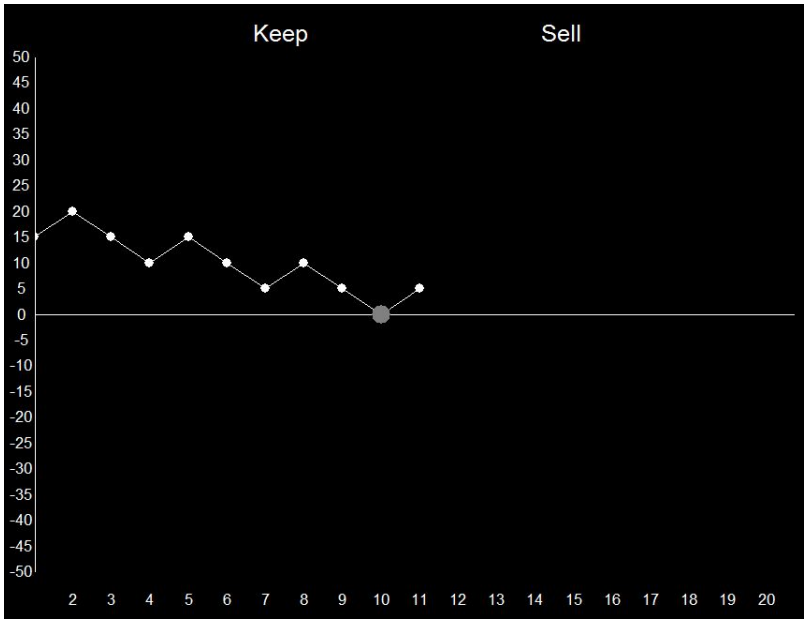

4. If you choose to KEEP your asset, there is a 5% chance that the trial will end, and you will receive what your asset is worth at that period. If the trial continues, there is a 50/50 chance that the asset will increase or decrease in value by $5. You can keep as many periods as you like, however keep in mind that there is a 5% chance that the trial will end every time you choose to keep the asset, and that the asset will increase or decrease by $5 with a 50/50 chance if the trial continues.

   In this example, the asset was kept and the trial continued on to the next period, where the asset decreased in value by $5.

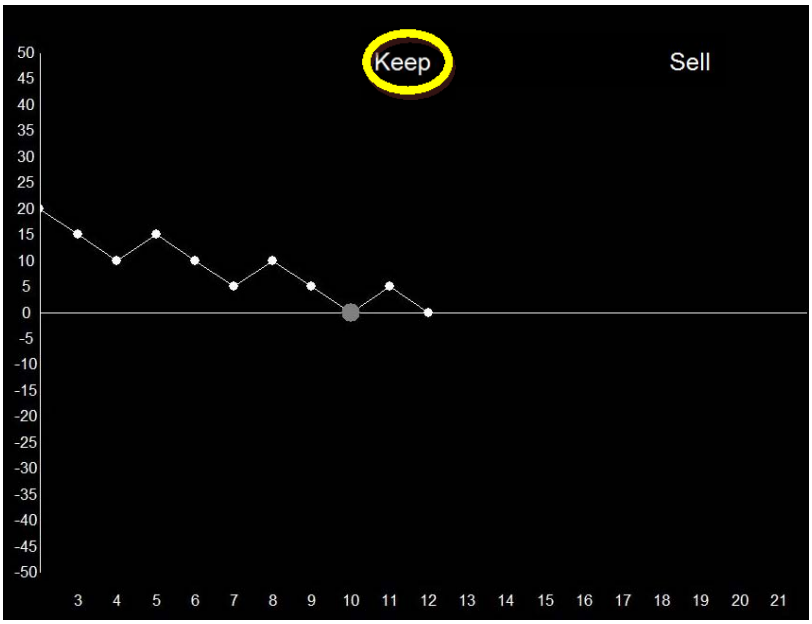

5. If you choose to keep the asset and the trial ends, you will receive what the asset is worth for that trial:

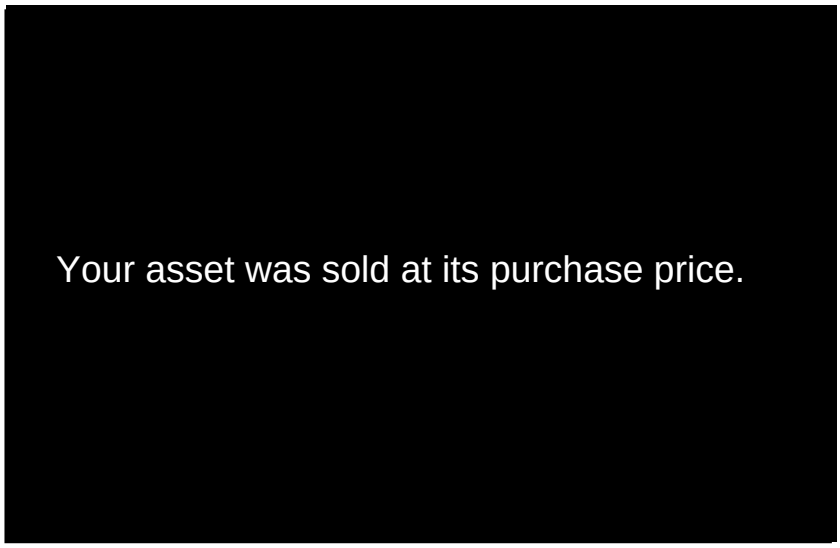

6. If you choose to SELL your asset, you will receive what the asset is worth for that period. This period is marked by a yellow circle, as seen below.

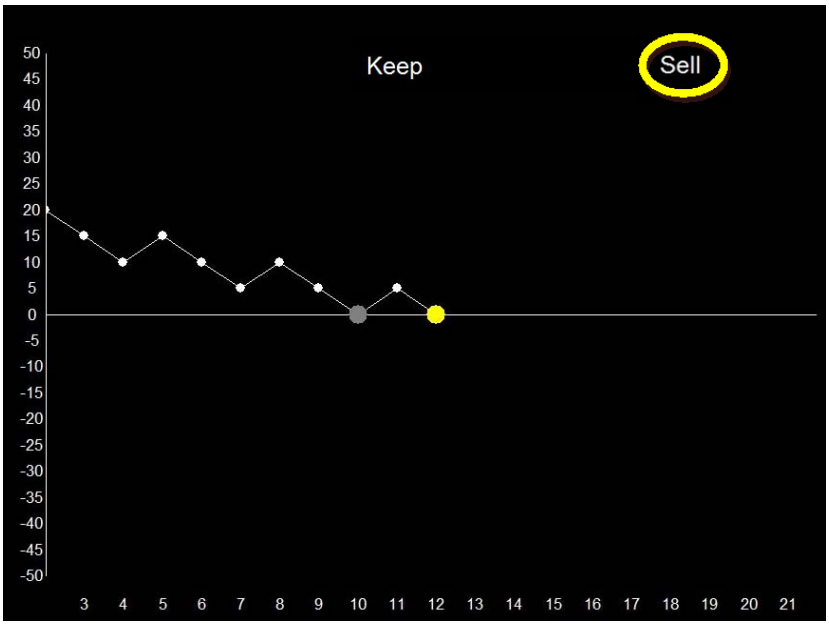

7. After you sell your asset, you may press the left or right arrow key to view what would have happened for five more periods if you had kept it. For example, assume you sell your asset in period 12. After pressing the left or right arrow key to continue, you will see what the asset would have done for each of the five periods after selling.

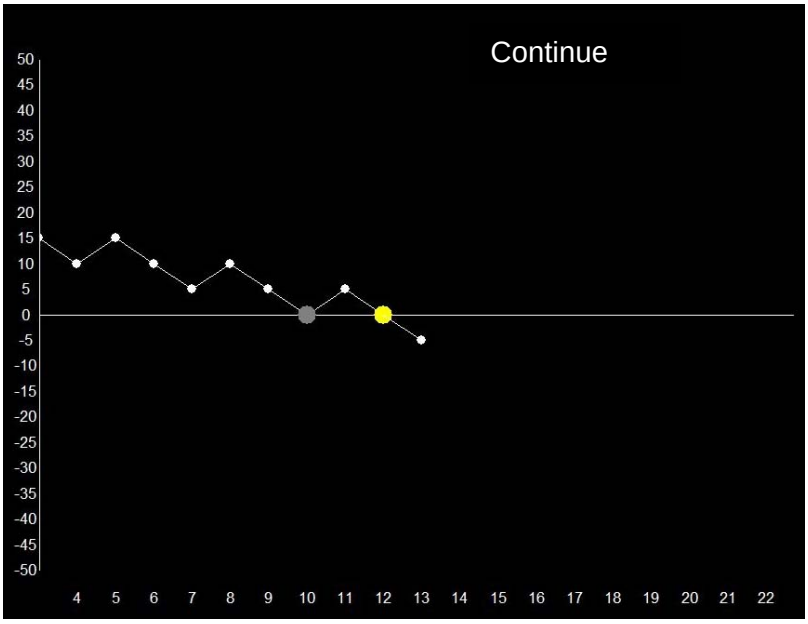

8. After the trial in which you sold your asset ends, you will see a screen that tells you for how much you sold your asset for in that trial.

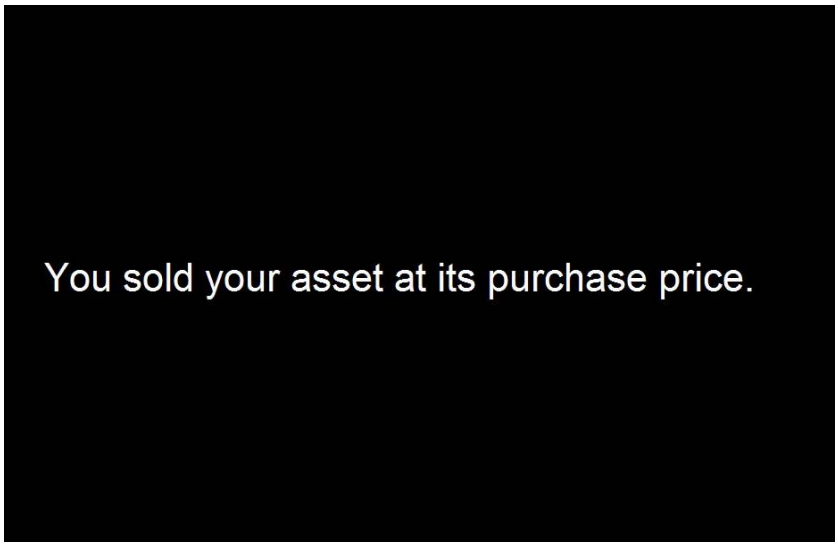

9. After all of the trials are over, you will role dice to determine a trial for which you will be paid.

Chapter 4

The Neuroscience Behind the Stock Market's Reaction to Corporate Earnings News

Introduction

As many investors can attest to, when corporate earnings announcements beat analyst's forecasts, the stock price of that company generally rises. Likewise, if corporate earnings miss expectations even by a small amount, the stock price tends to fall, usually by a larger amount than the rise from beating expectations (Skinner and Sloan, 2002). Analyst forecasts can be thought of as a measurable proxy for actual investor expectations, though the relationship between analysts and the companies they follow can be complicated. As stock price is a direct link to shareholder wealth, expectations that investors have on earnings announcements carry tremendous weight for the managers of a company. How investors generate their expectations of company earnings based on publicly available financials information is still not entirely clear, though a large amount of literature has been generated on this topic (Ball and Brown, 1968; Beaver, 1968; Dugar and Nathan, 1995; Calegari and Fargher, 1997; Libby et al., 2002; Bradshaw, 2004; Kaniel et al., 2012).

Ultimately, the trading behavior of investors is what drives the market reaction to earnings announcement beats or misses. But why do earnings misses lead to larger declines in stock price compared to equivalent gains? How do investors process the most basic of earnings information, their historical numbers? It is technically unfeasible to measure the expectation of each and every individual investor who plays a role in the

market; and analyst forecasts are not necessarily an accurate representation of individual investor expectations. Even if every investor could be contacted, investor expectations of a company's performance (and thus earnings surprise) can encompass more than just a numerical forecast or factors easily measured through self-report. Neuroimaing holds promise in being able to capture individual investor expectations, which should include factors not easily measurable.

Most of the research in human financial decision making has centered on the ventral striatum, a sub-cortical structure in the brain that has been implicated largely in reward-processing, and reward-related learning (Delgado et al., 2000; Knutson et al., 2001a; Rilling et al., 2002; McClure et al., 2004b). It has been shown to encode the subjective value of prospective options, and because it receives input from many cortical and sub-cortical regions of the brain, it is a prime structure for the integration of past and present sensory information necessary for the generation of expectation (Smith and Kieval, 2000). The ventral striatum has been shown to signal a reward prediction-error (RPE), a term which is negative when expectations about present and future reward are too low, zero when expectations are met, and positive when expectations are too low (Schultz, 1997; Pagnoni et al., 2002; McClure et al., 2003; O'Doherty et al., 2004; Abler et al., 2006; Hare et al., 2008).

We hypothesize that like other forms of monetary reward and punishment, the magnitude of earning surprise is signaled in the ventral striatum as a prediction error. Furthermore, we hypothesize that the ventral striatum will process both positive and negative earnings surprise. Finally, we hypothesize that an aggregate reward prediction error signal across subjects will correlate with subsequent changes in the market price of

the stock. To test these hypotheses, we recruited 36 second-year MBA students to serve as proxies for investors. The participants forecast 60 firm's earnings per share (EPS) numbers using actual EPS figures from three years prior, as well as financial analysts' consensus forecasts. Based on their own EPS forecasts, participants then take either a long or short position in the firm's stock. Finally, while undergoing fMRI, the participants learn the actual EPS, and the subsequent change in stock market price. We focus on the BOLD response in the ventral striatum during the phase in which the actual EPS is revealed, which we expect will capture the difference between the actual earning's numbers participant expectations.

Materials and Methods

Participants

Thirty-six full-time MBA students (8 female; 28 male) were recruited from Emory University's Goizueta Business School. The average age of participants was 28. Full-time MBA students were recruited because of their basic knowledge of accounting and investing, which are required for an understanding of the forecasting task. Full-time MBA students have been shown to be a good proxy for informed nonprofessional investors, the population we hope to make inferences regarding (Libby et al., 2002; Elliott et al., 2006). All participants reported no psychiatric, neurological, or other disorders which would preclude them from safely undergoing fMRI. All participants provided informed consent to experimental procedures approved by the Emory University Institutional Review Board. Participants rolled dice at the end of the experiment, which determined the trial participants would be compensated for.

Compensation was based on the change in stock price from $50 after the earnings announcement for that trial. Participants earned between $33 and $67.

Prior to the forecasting task, we assessed each participant's knowledge of accounting and investing concepts. To do this, participants completed a 25-question test covering topics such as the meaning of long and short positions, earnings announcements, and other task-related topics. We required a minimum of an 80% correct response rate. One individual did not make this cut-off, and therefore was excluded from analysis. The majority of participants reported having used financial reports and following and investing in the stock market. Academically, 65.7% of participants are second-year MBA students, 54.3% are concentrating in finance, and 60% hold undergraduate business degrees. All participants completed core MBA accounting and finance classes, and 80% completed at least two courses in finance, accounting, economics, and statistics.

Forecasting Task

The forecasting task was divided into two parts. In the first part, participants entered forecasts for 60 firms based on historical information, and then took a position (long or short) in each of the companies. In the second part, while undergoing fMRI, participants viewed the actual EPS of the companies, and the subsequent change in stock price surrounding the announcement. The first portion of the experiment outside the scanner consisted of three screens (**figure 4-1; panel A**). In the first screen, participants were given three years of historical information – EPS and analyst consensus forecasts, as well as the difference between these numbers. Participants were also given the current year's analysts' consensus forecast. After reviewing the historical information,

participants had to enter an EPS forecast. After this, they were asked if they wanted to go long or short in the company's stock.

The second portion of the experiment took place while participants were undergoing fMRI. Participants viewed three screens for each company they made forecasts for. On the first screen (**figure 4-1; panel B**), participants reviewed historical information, their forecast, and their position in the company's stock. In the second screen, the actual EPS was revealed. In the third screen, participants learned of the change in stock price and were told whether won or lost money based on the direction of the stock price change and their position. The interval between each screen and between trials was jittered using a negative exponential function with a mean of three seconds. A minimum of two seconds was required before participants could advance the screen.

Firm Selection

Each participant forecasted the EPS of 60 publicly traded companies. These companies were chosen based on several factors. First, the sample period was restricted to the time period between 2000-2009 because of well-documented changes in reported earnings, earnings surprise, and earnings thresholds during the last half of the century (Francis and Schipper, 1999; Givoly and Hayn, 2000). Second, we restricted the sample to firms with earnings announcement-period stock returns that have the same sign as the earnings surprise. Thus, if a company beats analyst expectations (positive surprise), the change in stock price would be positive. Likewise, if a company misses expectations, the change in stock price would be negative. Meeting expectations could lead to either a positive or negative surprise. This was done in order to prevent information other than what our participants viewed, such as CEO explanations as to why the company suffered

a loss, from playing a major role in the subsequent change in stock price. Of the firms we selected based on the aforementioned criteria, 45% beat the analyst consensus forecast, 40% missed the forecast, and 15% met the forecast.

fMRI data acquisition and analysis

Functional imaging was performed with a Siemens 3 T Trio whole-body scanner. T1-weighted images (TR = 2600 ms, TE = 3.02 ms, flip angle = 8°, 240 × 256 matrix, 176 sagittal interleaved slices, 1 mm^3 voxel size) were acquired for each subject prior to the experimental run. For the experimental run, where participants viewed actual EPS, T2*-weighted images using an echo-planar imaging sequence were acquired, which show blood oxygen level-dependent (BOLD) responses (echo-planar imaging, TR = 2000 ms, TE = 30 ms, flip angle = 90°, FOV = 192 mm × 192 mm, 64 × 64 matrix, 33 3.5-mm thick interleaved axial slices, and 3 × 3 × 3.5 mm voxels).

fMRI data were analyzed using SPM8 (Wellcome Department of Imaging Neuroscience, University College London) using a standard GLM model. Data were subjected to standard preprocessing, including motion correction, slice-timing correction, normalization to an MNI template brain, and smoothing using an isotropic Gaussian kernel (full-width half-maximum = 8 mm). Task regressors were modeled using a standard canonical hemodynamic response (HRF) function, which combines two gamma functions- one to model the peak HRF response, and one to model the undershoot. Further analysis of individual participant beta coefficients from the regressions was completed in Matlab (Mathworks) and STATA (StataCorp).

To estimate an average ventral striatal BOLD response to the actual earnings announcement for each company across individuals, we estimated the within-subjects BOLD response by running the following regression:

$$BOLD_{vt} = \beta_{0v} + \sum_j \beta_{1vj}[Screen1_t \otimes HRF_{(t)}] + \sum_j \beta_{2vj}[Screen2_t \otimes HRF_{(t)}] + \sum_j \beta_{3vj}[Screen3_t \otimes HRF_{(t)}] + \mathcal{E}_{vt}$$

where BOLD is the preprocessed signal in voxel *v* and time *t*, and *j* represents the trial/firm. Additional regressors controlling for separate runs and participant motion were included. Screen1 is the time period in which participants reviewed their position and the firm's historical information, Screen2 is when the earnings of the firm for the current period is displayed, and Screen3 is the stock price change. Because of our *a priori* hypothesis regarding the ventral striatum and earnings surprise, we limited the regression to the anatomically defined ventral striatum (**figure 4-2**). To obtain an average neural surprise response for each company, we simply summate each β_{2v} over all *v* in the ventral striatum ROI and over all participants for each firm *j*, to obtain 60 average BOLD responses. Some observations were excluded if the participant took the wrong position (72 observations), or there were scanner-related problems preventing the capture of a response (11 observations).

To test whether the ventral striatal response is correlated with the change in stock price after the company's earnings announcement, we ran the following linear regression:

$$CRR_j = \beta_1 + \beta_2 VS\_BOLD_j + \mathcal{E}_j$$

where CRR represents the cumulative raw return during the three day period centered on the earnings announcement, VS_BOLD is the average BOLD response across

participants in the ventral striatum, and *j* represents the firm. To aid in the interpretation of the beta coefficient β_2, we flip the sign of the BOLD signal for short positions such that they can be interpreted in the same way as long positions. Additionally, we compared a model with the only regressor being earnings surprise (actual EPS – analyst forecast) with a model containing both surprise and VS_BOLD described above. To test whether the ventral striatum reflects the larger effect that earnings misses has on stock price, we run a model where VS_BOLD is split into beat, meet, and miss categories, regressed against CRR. All regressions used robust standard errors, which help deal with heteroscedasticity in the data (though the regression results do not change substantially when using regular standard errors).

Results

Participants' Behavior

The participant estimates of EPS (mean = 2.70, μ = 1.95) were not statistically different from the consensus forecast (mean = 2.67, μ = 1.95; $t(59) = 0.70$), and were strongly correlated ($r = 0.998$, $p < 0.001$). The average surprise for each firm, defined as the difference between the actual EPS and the forecasts (participant or analysts) averaged over firms, significantly differed between participants (mean = 0.02, μ = 0.48) and analysts (mean = -0.01, μ = 0.44; $t(59) = 2.17$) . Forecasting ability was defined as the proportional mean absolute forecast error (Clement et al., 2007), and was not correlated with experience, education, or quiz scores when corrected for multiple comparisons using the Bonferroni method.

BOLD Response Analysis

To get rid of generic task-related effects, we subtract the BOLD response to the first screen (review) from the screen of interest (surprise). We find that the direction of striatal response to earnings surprise in all categories is as expected – positive when companies beat (miss) analyst expectations and participants are long (short), and negative when companies miss (beat) expectations and analysts are long (short). Against what we might expect, the BOLD response in the ventral striatum is not significantly different from zero in all conditions (**table 4-1**). Only in conditions where the participant would lose money based on their position, or where it's unclear if they will lose or gain (when the company meets expectations), was the striatal BOLD response significantly different from zero at a threshold of $\alpha = 0.05$. The absolute response to earnings announcements that would subsequently lead to monetary loss for the participant (miss-long and beat-short), were larger than the absolute response to announcements that would lead to monetary gain.

Across participants, we find no significant correlation between the magnitude of the earnings surprise announcement (actual EPS – analyst consensus) and the BOLD response in the striatum during the surprise screen for long ($r = 0.02$; $p = 0.52$) or short ($r = -0.06$; $p = 0.13$) positions. Calculating earnings surprise using the participant forecasts (actual EPS – participant forecast), we also find no correlation with the BOLD response in the ventral striatum for long ($r = 0.02$, $p = 0.58$) or short ($r = -0.05$; $p = 0.18$) positions.

To test whether an aggregate BOLD response for each firm would correlate with subsequent raw market return, we regressed the BOLD response to company's earnings announcements against raw stock return. To make beta coefficients more easily

interpretable, we flip the sign of the BOLD response on short positions to transform them to equivalent long positions. We find that the average ventral striatal BOLD response for each firm significantly correlates with raw return during the announcement period (**table 4-2; panel A**). To test whether the participant BOLD response provides any additional information over earnings surprise, we compared two models with earnings surprise, one with and without the BOLD data. As expected, a simple linear regression model with analyst surprise as a regressor was significant, with an adjusted R^2 of 0.16 (**table 4-2; panel B**). Adding the average BOLD response in the ventral striatum increases the adjusted R^2 of the model from 0.16 to 0.41, suggesting that the BOLD data adds additional explanatory power to the model (**table 4-2; panel C**).

Earnings announcement misses lead to a larger decline in stock price than the rise from beating expectations. To test whether the ventral striatal response to company earnings misses has a larger effect on market price, we divided the average ventral striatal BOLD signal into beat, meet, and miss regressors, took their absolute value, and regress them against raw return. In both the meet and beat conditions, the BOLD response significantly relates to the market return (**table 4-3**). The BOLD response to meeting expectations did not significantly correlate with market return. The absolute value of the beta coefficients for beat vs. miss were similar in size ($\beta = 0.272$ vs. $\beta = -0.243$ for beat and miss, respectively).

Discussion

Company earnings announcements have a large effect on shareholder wealth. When companies beat analyst expectations, the stock price generally rises, and when they

miss analyst expectations, the stock price generally falls. Little is known about how individual investors process earnings information. It is not feasible to measure every investor's expectation of earnings in the market. A large body of literature within the field of neuroeconomics has demonstrated that certain brain structures signal expectations of future reward, namely the ventral striatum (Hollerman and Schultz, 1998; McClure et al., 2003; O'Doherty et al., 2004). Here we utilize fMRI as a means to measure brain activity in this structure while a sample of proxy-investors complete a task where they view earnings announcements for companies they have a stake in. Their compensation for the task was tied to the performance of the company's stock post-announcement. This allows us to answer several questions – is earnings surprise tracked in the ventral striatum of investors? Are both positive and negative earnings surprises, which have differential effects in the magnitude of price change in the market, reflected in the ventral striatum? Finally, can the ventral striatum activity of a subset of investors be correlated with overall market activity?

We found that the BOLD response in the anatomically-defined ventral striatum to company earnings surprise versus a baseline screen was significant, but only for negative surprise when participants were long, and positive surprise when participants were short – both conditions in which participants would lose money due to market price change. Many neuroimaging studies find a significant activation to monetary gain in the ventral striatum, but not monetary loss (Knutson et al., 2001a; Yacubian et al., 2006). This is possibly attributed to the lack salience of monetary losses in experiments, because of the fact that participants have to be compensated for their time, and thus always leave with more money than when they started. However, in this case we find that a cue signaling

monetary loss results in significant less ventral striatal activity than a baseline information screen. In the market, earnings misses tend to result in more drastic market changes, which could be reflected in our sample by the larger absolute value for the loss-conditions. Other groups have found similar results regarding steeper responses to monetary loss-predicting stimuli than monetary-gain predicting ones (Tom et al., 2007).

Most surprisingly, we find that the ventral striatal response to company earnings surprise correlates with subsequent changes in stock price. Earnings surprise, as defined by the difference in actual earnings per share and analyst consensus forecasts, is positively correlated with the change in stock price. Investor expectations and subsequent trading behavior is likely to be partially responsible for this change. We find that ventral striatal activity, which has been shown to signal expectations of reward (Schultz, 1997; Kable and Glimcher, 2007), correlates with changes in stock price around the announcement period in a group of proxy investors. This suggests that the ventral striatal response might be capturing individual expectations, which could potentially help explain some market variation in terms of raw return around the announcement period. We find that the BOLD response adds additional explanatory power above and beyond earnings surprise itself, increasing the adjusted R^2 of the regression model.

We show that fMRI can potentially be a useful tool in helping explain investor expectations, and subsequent market reaction. In our sample, the ventral striatum captures information from individual proxy-investors that is above and beyond the earnings surprise numbers themselves. Determining what information the ventral striatum is actually conveying beyond the reaction to the earnings surprise could be done in future experiments. Neuroimaging can help explain the relationship between the

market, accounting numbers, and investor expectations through carefully controlled experiments. Here we take a first step in this direction. Whether our results are applicable to the market as a whole can only be determined by future experiments.

Figures and Tables

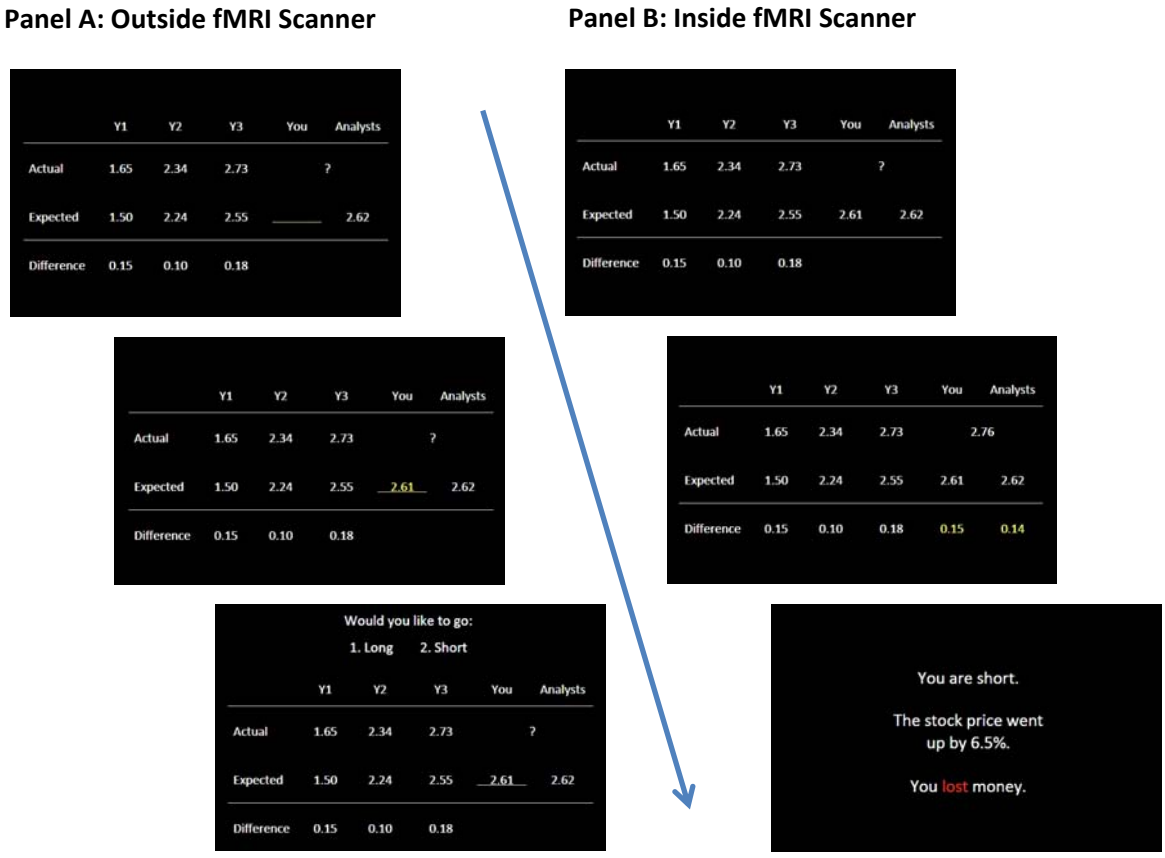

Figure 4-1: Experimental design. Each trial consisted to two components. The first component shown in panel A, took place outside of the MRI scanner. Participants were asked to enter a forecast of the EPS for the current period based on given historical information. This historical information included the past three years of actual earnings numbers, as well as analyst consensus numbers. They then had to take either a long or short position in the firm's stock. The second component (panel B), which was completed inside the MRI scanner, consisted of three screens – the first was a review of the participant's position, historical information, and the entered earnings estimate. The

second screen revealed the actual earnings per share of the current period, as well as the difference between actual and analyst consensus, and actual and participant forecast numbers. The third screen displayed the stock price change during the three-day period centered on the earnings announcement, and whether the participant won or lost money based on their position.

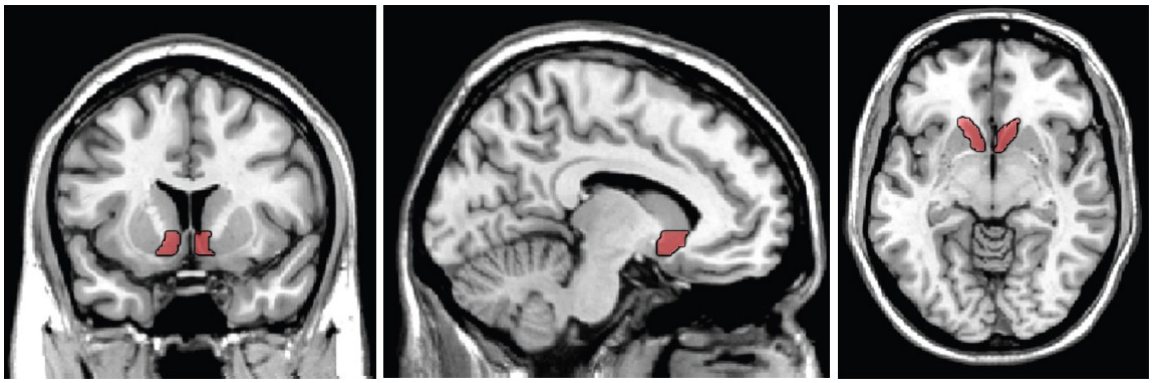

Figure 4-2: Anatomical ventral striatum ROI. Our *a priori* hypothesis was that the ventral striatum would process earnings surprise information. We therefore limited our analysis to this region (shaded in red).

	Beat	**Meet**	**Miss**
Long	*M* = 0.03 *SD* = 1.20 *t* = 0.68	*M* = -0.15 *SD* = 1.11 *t* = -1.92	*M* = -0.40 *SD* = 1.20 *t* = -7.60**
Short	*M* = -0.44 *SD* = 1.13 *t* = -6.46**	*M* = -0.41 *SD* = 1.61 *t* = -2.67*	*M* = 0.08 *SD* = 1.26 *t* = 1.07

Table 4-1: Average BOLD response in the anatomical ventral striatum to earnings surprise. ** = $p < 0.001$, * = $p < 0.05$.

Panel A.

CRR = β_0 + β_1*VS_BOLD + ε

	Coefficient	SE	t	P>t	[95% Conf.	Interval]
VS_BOLD	0.210	0.042	5.050	0.000	0.127	0.293
Constant	0.006	0.013	0.500	0.620	-0.019	0.031

R^2 = 0.306

Adj. R^2 = 0.294

Panel B.

CRR = β_0 + β_1*Surprise + ε

	Coefficient	SE	t	P>t	[95% Conf.	Interval]
Surprise	0.110	0.031	3.510	0.001	0.047	0.173
Constant	-0.002	0.014	-0.130	0.893	-0.029	0.025

R^2 = 0.175

Adj. R^2 = 0.161

Panel C.

CRR = β_0 + β_1*Surprise + β_2*VS_BOLD + ε

	Coefficient	SE	t	P>t	[95% Conf.	Interval]
Surprise	0.093	0.027	3.490	0.001	0.040	0.146
VS_BOLD	0.193	0.038	5.020	0.000	0.116	0.269
Constant	0.005	0.011	0.450	0.653	-0.018	0.028

R^2 = 0.428

Adj. R^2 = 0.408

Table 4-2: Results of market regressions. To test whether the BOLD response in the ventral striatum relates to the change in market price, we regress surprise and the BOLD response (VS) separately onto the cumulative raw return (CRR). VS_BOLD alone significantly correlates with CRR (panel A); however earnings surprise has already been shown to correlate strongly with market return. We test to see whether the ventral striatal response provides additional information on top of surprise alone by comparing a model with only surprise (panel B) with a model containing both surprise and VS_BOLD (panel C). The adjusted ^{R2} increases with the addition of VS_BOLD, suggesting an increase in explanatory power over surprise with the ventral striatal BOLD response.

Chapter 5

Summary of Findings

The ventral striatum is thought to signal reward prediction error (RPE) that theoretically should account for both financial gain and loss. A cue predicting a monetary gain should elicit a positive RPE, and a cue predicting a monetary loss might be expected to elicit a negative RPE. Cues signaling monetary gain consistently activate the ventral striatum. However, the case for ventral striatum signaling of RPE for financial losses is less clear. Some research studies show an increase in BOLD response to financial gains, and a decrease to financial losses (Tom et al., 2007), while others demonstrate that separate regions respond to gains and losses (Yacubian et al., 2006). To advance our knowledge of both ventral striatum loss processing, as well as help explain real-world phenomenon as they relate to financial loss using neuroimaging, we carried out a series of three experiments described in **chapters 2**, **3**, and **4**. We summarize the results and conclusions of these experiments below.

From Bad to Worse: Striatal Coding of the Relative Value of Painful Decisions

1. When participants made decisions over gambles consisting of electric shocks, they behaved in a risk-seeking manner. This is consistent with the view that participants treated the lotteries in a similar manner as financial losses.

2. Even in this loss realm with entirely aversive outcomes, we find that cues predicting better potential outcomes elicited a BOLD response in the ventral striatum.

3. The magnitude of the BOLD response in the ventral striatum correlated with the overall expected value of the lotteries.

4. The BOLD response correlated with each component of the lotteries – positively with the potential shocks less than a status quo amount, and negatively with the potential shocks less.

Neural Insensitivity to Upticks in Value is Associated with the Disposition Effect

5. In an asset-trading task where price followed a random walk and there was no optimal trading behavior, participants held onto losing assets much longer than gaining assets. This behavior is known as the disposition effect.

6. Risk-preference measured outside of the scanner did not correlate with the magnitude of the disposition effect. This suggests that whatever governs risk-seeking behavior over losses and risk-averseness over gains is not likely the same mechanism governing the disposition effect.

7. Heterogeneity in the disposition effect was driven by behavior when the asset was at a loss, but not when the asset was at a gain.

8. The slope of the BOLD response to an uptick in asset value when participants were below the purchase price was negatively correlated with the disposition effect. In other words, participants who had a larger disposition effect had a blunted response in the ventral striatum to upticks in value when at a loss. Under an RPE interpretation of the ventral striatum response, the blunted response in participants who had a larger disposition effect might signify that their expectation that the losing asset would begin to return to the mean was met. This is consistent with a mean reversion hypothesis.

The Neuroscience Behind the Stock Market's Reaction to Corporate Earnings News

9. Participant forecasts for the current period earnings per share was significantly correlated with analyst consensus.

10. The BOLD response in the ventral striatum decreases for company earnings announcements which lead to monetary loss, but surprisingly did not increase for announcements that lead to monetary gain.

11. The overall BOLD response did not correlate with earnings surprise (the difference between the actual earnings and the analyst consensus).

12. The BOLD response in the ventral striatum to earnings announcements significantly correlated with the cumulative market return.

13. The BOLD response added significant explanatory power, demonstrated by an increase in the adjusted R^2, when moving from a model with only earnings surprise to a model with both earnings surprise and the BOLD response in the ventral striatum.

Significance

Basic Knowledge

The overarching conclusion that can be drawn from the above three studies is that the ventral striatum does indeed encode predictions of financial loss through decreases in activity. We find that the BOLD response in the ventral striatum correlates with the expected value of entirely aversive stimuli that participants treat as losses (**chapter 2**). We also find that loss-chasing behavior in the disposition effect is correlated with a striatal response while participants are at a loss (**chapter 3**). Finally, company earnings announcements which lead to stock price changes that result in participant financial loss are coded within the ventral striatum (**chapter 4**). From a basic research perspective, this is important for understanding the range of stimuli with which the ventral striatum signals RPE. Modern theories of decision-making are tied to concepts of reward, utility, and RPE, all of which neuroimaging finds is represented within the mesolimbic dopamine system (Hollerman and Schultz, 1998; Pagnoni et al., 2002; Kable and Glimcher, 2007; Berridge and Kringelbach, 2008). Models of decision-making associated with these concepts can account for both sides of the financial gain-loss coin, and might help explain what happens when decision making goes wrong (Montague et al., 2004, 2012;

Sharp et al., 2012). Whether they actually provide benefit to human health hinges on whether the brain behaves in a manner that these models prescribe, and having an understanding of how the parameters of these models are represented (O'Doherty et al., 2007; Daw et al., 2011). My dissertation adds to the emerging literature on the brain's representation of reinforcement learning models by showing that negative RPE can be represented in the ventral striatum.

Clinical Relevance

From a health-oriented standpoint, understanding the brain's neuroanatomy and mechanism of dealing with financial loss is important for developing methods of treatment for mental disorders characterized by aberrant financial decision-making. Pathological gamblers incur large financial losses, yet continue to gamble despite this. Cognitive distortions like the gambler's fallacy are known to be present in pathological gamblers (Toneatto, 1999; Raylu and Oei, 2002). The gambler's fallacy is a belief that every segment of a sequence of random events must reflect the true proportion – that is, the belief that if a coin flip results in heads twice in a row, on the third flip the coin will be more likely to land tails to keep the 50/50 true proportion (Tversky and Kahneman, 1971). In cases of gambling or investment, a belief in the gambler's fallacy would lead to the chasing of loses due to the belief that multiple losses in a row would increase the likelihood of a win. In **chapter 3**, I show that the disposition effect is likely driven by an irrational belief in mean reversion. Mean reversion is similar to the gambler's fallacy in that they both represent a belief in a random variable returning to its mean in short segments of time. We showed that a healthy sample of MBA students exhibited the disposition effect, which mirrors the literature showing that the effect is present in a

variety of populations, including undergraduates and professional investors (Odean, 1998; Weber and Camerer, 1998). Thus one aspect of pathological gambling, the chasing of losses, could be a general human trait that is exacerbated in pathological cases. Interestingly, dopamine agonists used to treat Parkinson's disease symptoms have been shown to increase the likelihood of pathological gambling behavior, of which loss-chasing is a large component (Driver-Dunckley et al., 2003; Dodd et al., 2005). It is possible that their influence on dopaminergic transmission within the ventral striatum might drive the increase in loss-chasing behavior.

Applications

Beyond models of decision-making and health-related significance, some of my research has suggested that neuroimaging of sub-samples of a population can help provide additional information about population behavior than behavioral metrics alone. In **chapter 4**, the BOLD response to company earnings announcements in the ventral striatum correlated with the subsequent change in market price. I show that this ventral striatum response, which can be interpreted as the difference between actual earnings and the proxy-investor expectations, helps explain additional variance in market return above and beyond numerical earnings surprise. As with most scientific literature, caution must be taken with extrapolating these results. The firms in **chapter 4** were chosen based a clear pattern of positive-surprise → increased market return, negative surprise → decreased market return. In the actual market, these patterns are not always met. The R^2 for a model with many more firms not meeting these criteria will unquestionably be lower. Despite this, fMRI holds some potential in capturing expectations that otherwise couldn't be elicited from self-report.

fMRI is expensive relative to surveys and questionnaires. However, if it can allow one to use sub- sample of a population to add additional explanatory power over surveys and questionnaires, then in the long-term it could be a useful tool for both marketing purposes and understanding some market phenomena. Indeed, the field of neuromarketing is already burgeoning (Plassmann et al., 2008; Ariely and Berns, 2010). My research adds to the field by showing that even extremely complex market behavior can be partially explained using neuroimaging responses in a sub-sample.

Cellular Speculation

The technique of fMRI takes a coarse spatial and temporal look at brain activity compared to cellular-level techniques such as electrophysiology. As with any rapidly adopted technique, there has been some caution leveled in the interpretation of fMRI results in terms of the underlying brain processes governing the signal (Attwell and Iadecola, 2002; Logothetis, 2008). Research published in the last decade combining electrophysiology with fMRI has demonstrated that the BOLD signal reflects local field potential more so than action potentials (Logothetis et al., 2001; Logothetis, 2002). The local field potential principally consists of all synaptic activity on the dendrites of neurons within a region, not necessarily the action potentials generated by neurons in that region. The interpretation of a BOLD signal in a region thus has to be tempered by the fact that it could be reflecting excitatory or inhibitory input from outside regions, or intra-regional processing. Ultimately the conclusions drawn about the mesolimbic dopamine system signaling both financial gain and loss from fMRI studies remain the same; and a substantial amount of literature at the cellular-level reflects these results.

The neuronal architecture to support the mesolimbic dopamine systems involvement in signaling financial gain and loss is there. GABAergic neurons within the ventral striatum increase in their firing rate prior to delivery of a reward (Schultz et al., 1992). Dopaminergic neurons in the VTA signal expectations of reward in a RPE-manner (Hollerman and Schultz, 1998). Dopaminergic neurons in the substantia nigra of humans have been shown to increase in firing rate to unexpected gains, and decrease to unexpected losses (Zaghloul et al., 2009). Some populations of midbrain dopamine neurons increase in firing rate to both appetitive and aversive stimuli, potentially signaling saliency. Other populations are excited by appetitive stimuli and inhibited by aversive stimuli (Matsumoto and Hikosaka, 2009). And yet others only fire for positive RPEs (Bayer and Glimcher, 2005). Further downstream, neurons in the lateral habenula signal negative reward which is thought to drive dopaminergic prediction error signals (Matsumoto and Hikosaka, 2007). Financial loss could be regarded as an aversive stimulus, or it could be regarded as less of an appetitive stimulus. In either case, the type of neuron required to carry these signals are indeed there.

The characteristics of dopamine neurons place a limit on the range of RPE that can be signaled. The baseline action potential frequency of midbrain dopaminergic neurons is roughly 5 Hz (Schultz, 1997). Dopaminergic neurons can be excited up to 100 Hz for a brief period of time, but can only be suppressed down to 0 Hz (Hyland et al., 2002; Pan et al., 2005). Thus if the baseline firing rate is an RPE of 0, the positive RPE could be orders of magnitude larger than negative RPEs. Financial gains and losses are reference dependent, as demonstrated behaviorally by Kahneman & Tversky (1979), and evidence suggests that signals from the ventral striatum are relative to the person's

reference point (De Martino et al., 2009). Thus the baseline firing rate of dopaminergic neurons might be re-set to represent a different reference point in order to accommodate largely negative RPEs. Alternatively, aversive RPEs that are positive for worse-than-expected outcomes might carry the financial loss side of the coin. In the above three experiments described in my dissertation, increasing potential loss always resulted in a lessened BOLD response in the ventral striatum, so the former explanation is more likely.

Future Directions

The debate on whether a single system processes financial gains and losses, or whether positive and negative valence are separated spatially or temporally will likely continue. I provide evidence for the former argument by showing that the ventral striatum is involved in loss-related processes in three separate experiments. However the problem of having participant experience ecologically-valid financial loss in an experimental setting can't be ignored. Future research might compare incentive structures and brain activation to financial loss to determine whether conflicting results in the financial loss literature are driven by differences in incentive structure. This could have wide-reaching impact on the future study of financial loss in neuroeconomic studies.

Future research might also focus on loss-holding behavior in mental health behaviors and in aging. In **chapter 3**, I use a simple paradigm to study the disposition effect, where participants choose to keep or sell an asset whose price follows a random walk. In a healthy population, we show that the disposition effect is clearly present, and that based on activation within the ventral striatum, is likely driven by an irrational belief in mean reversion. One could hypothesize that pathological gamblers would be more likely to show a large disposition effect compared to a healthy population. Using

neuroimaging, we might be able to determine whether this exacerbated loss-chasing is due to a stronger belief in mean reversion (or the gambler's fallacy) based on striatal activation or whether some other belief is driving it. Follow-up studies could then be able to test training routines tailored to the findings in order to decrease loss-chasing behavior.

Future research might utilize higher field strength MRI with localizers on the ventral striatum to increase both the spatial and temporal resolution of the signal. By imaging a smaller region of the brain, both spatial and temporal signal can be increased (D'Ardenne et al., 2008; Salas et al., 2010). Furthermore, the restriction of MRI data-capture to pre-defined regions would increase the likelihood of finding significant small-cluster activations because extent thresholds for multiple comparison correction would be not be as stringent over a smaller volume. Given the electrophysiological evidence that there are separate sub-populations of dopaminergic neurons that are excited by and inhibited by aversive stimuli, we might expect to see these distinct appetitive and aversive RPEs spatially separated in the ventral striatum for monetary gains and losses (Seymour et al., 2005).

Final Words

In closing, the findings presented here help tackle several aspects of human decision-making. We have framed the above experiments primarily in terms of the neural coding of financial loss, but they each contribute to a host of other areas of research that span from psychology, to economics, to behavioral finance. Our findings have focused primarily on the ventral striatum as a substrate of valuation, but other brain regions both within and outside the mesolimbic dopamine system have been shown to contribute to human decision-making. From the large amount of neuroimaging research that is ongoing, we are hopeful that a complete picture of where and how the brain makes decisions can emerge.

References

Abler B, Walter H, Erk S, Kammerer H, Spitzer M (2006) Prediction error as a linear function of reward probability is coded in human nucleus accumbens. NeuroImage 31:790–795.

Allen JS, Damasio H, Grabowski TJ (2002) Normal neuroanatomical variation in the human brain: an MRI-volumetric study. American Journal of Physical Anthropology 118:341–358.

Anderson LR, Mellor JM (2008) Predicting health behaviors with an experimental measure of risk preference. Journal of Health Economics 27:1260–1274.

Andreassen PB (1988) Explaining the price-volume relationship: the difference between price changes and changing prices. Organizational Behavior and Human Decision Processes 41:371–389.

Ariely D, Berns GS (2010) Neuromarketing: the hope and hype of neuroimaging in business. Nature Reviews Neuroscience 11:284–292.

Attwell D, Iadecola C (2002) The neural basis of functional brain imaging signals. Trends in Neurosciences 25:621–625.

Baliki MN, Geha PY, Fields HL, Apkarian AV (2010) Predicting value of pain and analgesia: nucleus accumbens response to noxious stimuli changes in the presence of chronic pain. Neuron 66:149–160.

Ball R, Brown P (1968) An empirical evaluation of accounting income numbers. Journal of Accounting Research 6:159–178.

Barberis N, Thaler R (2003) A Survey of Behavioral Finance. In: Handbook of the Economics of Finance, pp 1053–1128. Elsevier.

Barberis N, Xiong W (2008) Realization Utility. NBER Working Papers.

Barberis N, Xiong W (2009) What drives the disposition effect? An analysis of a long-standing preference-based explanation. Journal of Finance 64.

Barsky RB, Juster FT, Kimball MS, Shapiro MD (1997) Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study. The Quarterly Journal of Economics 112:537–579.

Bayer HM, Glimcher PW (2005) Midbrain dopamine neurons encode a quantitative reward prediction error signal. Neuron 47:129–141.

Beaver WH (1968) The Information Content of Annual Earnings Annoucements. Journal of Accounting Research 6:67–92.

Becerra L, Breiter HC, Wise R, Gonzalez RG, Borsook D (2001) Reward circuitry activation by noxious thermal stimuli. Neuron 32:927–946.

Bechara A, Damasio H, Tranel D, Anderson SW (1998) Dissociation Of Working Memory from Decision Making within the Human Prefrontal Cortex. Journal of Neuroscience 18:428–437.

Berns GS, Bell E (2012) Striatal topography of probability and magnitude information for decisions under uncertainty. NeuroImage 59:3166–3172.

Berns GS, McClure SM, Pagnoni G, Montague PR (2001) Predictability modulates human brain response to reward. Journal of Neuroscience 21:2793–2798.

Berridge KC, Kringelbach ML (2008) Affective neuroscience of pleasure: reward in humans and animals. Psychopharmacology 199:457–480.

Berridge KC, Robinson TE (2003) Parsing reward. Trends in Neurosciences 26:507–513.

Binswanger HP (1980) Attitudes toward risk: experimental measurement in rural India. American Journal of Agricultural Economics 62:395–407.

Bowman EM, Aigner TG, Richmond BJ (1996) Neural signals in the monkey ventral striatum related to motivation for juice and cocaine rewards. Journal of Neurophysiology 75:1061–1073.

Bradshaw MT (2004) How Do Analysts Use Their Earnings Forecasts in Generating Stock Recommendations? The Accounting Review 79:25–50.

Breiter HC, Gollub RL, Weisskoff RM, Kennedy DN, Makris N, Berke JD, Goodman JM, Kantor HL, Gastfriend DR, Riorden JP, Mathew RT, Rosen BR, Hyman SE (1997) Acute Effects of Cocaine on Human Brain Activity and Emotion. Neuron 19:591–611.

Brooks AM, Pammi VSC, Noussair C, Capra CM, Engelmann JB, Berns GS (2010) From bad to worse: striatal coding of the relative value of painful decisions. Frontiers in Neuroscience 4:176.

Bush RR, Mosteller F (1953) A Stochastic Model with Applications to Learning. The Annals of Mathematical Statistics 24:559–585.

Calegari M, Fargher NL (1997) Evidence that Prices Do Not Fully Reflect the Implications of Current Earnings for Future Earnings: An Experimental Markets Approach. Contemporary Accounting Research 14:397–433.

Camerer CF, Weigelt K (1993) Convergence in experimental double auctions for stochastically live assets. In: Proceedings Volume in the Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley.

Campbell-Meiklejohn DK, Woolrich MW, Passingham RE, Rogers RD (2008) Knowing when to stop: the brain mechanisms of chasing losses. Biological Psychiatry 63:293–300.

Choe H, Eom Y (2009) The disposition effect and investment performance in the futures market. Futures 29:496–522.

Christopoulos GI, Tobler PN, Bossaerts PL, Dolan RJ, Schultz W (2009) Neural correlates of value, risk, and risk aversion contributing to decision making under risk. Journal of Neuroscience 29:12574–12583.

Clement MB, Koonce L, Lopez TJ (2007) The roles of task-specific forecasting experience and innate ability in understanding analyst forecasting performance. Journal of Accounting and Economics 44:378–398.

Coizet V, Dommett EJ, Redgrave P, Overton PG (2006) Nociceptive responses of midbrain dopaminergic neurones are modulated by the superior colliculus in the rat. Neuroscience 139:1479–1493.

Croxson PL, Walton ME, O'Reilly JX, Behrens TEJ, Rushworth MFS (2009) Effort-based cost-benefit valuation and the human brain. Journal of Neuroscience 29:4531–4541.

Da Costa N, Mineto C, Da Silva S (2008) Disposition effect and gender. Applied Economics 15:411–416.

Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ (2011) Model-Based Influences on Humans' Choices and Striatal Prediction Errors. Neuron 69:1204–1215.

De Martino B, Kumaran D, Holt B, Dolan RJ (2009) The neurobiology of reference-dependent value computation. Journal of Neuroscience 29:3833–3842.

Delgado MR (2007) Reward-related responses in the human striatum. Annals of the New York Academy of Sciences 1104:70–88.

Delgado MR, Li J, Schiller D, Phelps EA (2008) The role of the striatum in aversive learning and aversive prediction errors. Philosophical Transactions of the Royal Society of London 363:3787–3800.

Delgado MR, Locke HM, Stenger VA, Fiez JA (2003) Dorsal striatum responses to reward and punishment: effects of valence and magnitude manipulations. Cognitive, Affective, & Behavioral Neuroscience 3:27–38.

Delgado MR, Nystrom LE, Fissell C, Noll DC, Fiez JA (2000) Tracking the Hemodynamic Responses to Reward and Punishment in the Striatum. Journal of Neurophysiology 84:3072–3077.

Delgado MR, Stenger VA, Fiez JA (2004) Motivation-dependent responses in the human caudate nucleus. Cerebral Cortex 14:1022–1030.

Dodd ML, Klos KJ, Bower JH, Geda YE, Josephs KA, Ahlskog JE (2005) Pathological gambling caused by drugs used to treat Parkinson disease. Archives of Neurology 62:1377–1381.

Driver-Dunckley E, Samanta J, Stacy M (2003) Pathological gambling associated with dopamine agonist therapy in Parkinson's disease. Neurology 61:422–423.

Dugar A, Nathan S (1995) The Effect of Investment Banking Relationships on Financial Analysts' Earnings Forecasts and Investment Recommendations. Contemporary Accounting Research 12:131–160.

D'Ardenne K, McClure SM, Nystrom LE, Cohen JD (2008) BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. Science (New York, NY) 319:1264–1267.

Elliott WB, Hodge FD, Kennedy SJ, Pronk M (2006) Are MBA Students a Good Proxy for Non-professional Investors? SSRN Electronic Journal.

Eraker SA, Sox HC (1981) Assessment of patients' preferences for therapeutic outcomes. Medical Decision Making 1:29–39.

Evans AC, Collins DL, Mills SR, Brown ED, Kelly RL, Peters TM (1993) 3D statistical neuroanatomical models from 305 MRI volumes. In: 1993 IEEE Conference Record Nuclear Science Symposium and Medical Imaging Conference, pp 1813–1817. IEEE.

Ferris SP, Haugen RA, Makhija AK (1988) Predicting contemporary volume with historic volume at differential price levels: evidence supporting the disposition effect. Journal of Finance 43:677–697.

Fiorillo CD, Tobler PN, Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299:1898–1902.

Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC (1995) Improved Assessment of Significant Activation in Functional Magnetic Resonance Imaging (fMRI): Use of a Cluster-Size Threshold. Magnetic Resonance in Medicine 33:636–647.

Francis J, Schipper K (1999) Have Financial Statements Lost Their Relevance? Journal of Accounting Research 37:319–352.

Friman O, Borga M, Lundberg P, Knutsson H (2003) Adaptive analysis of fMRI data. NeuroImage 19:837–845.

Friston KJ, Holmes AP, Price CJ, Büchel C, Worsley KJ (1999) Multisubject fMRI studies and conjunction analyses. NeuroImage 10:385–396.

Frydman C, Barberis N, Camerer CF, Bossaerts PL, Rangel A (2011) Testing Theories of Investor Behavior Using Neural Data. SSRN Electronic Journal.

Genovese CR, Lazar NA, Nichols TE (2002) Thresholding of statistical maps in functional neuroimaging using the false discovery rate. NeuroImage 15:870–878.

Givoly D, Hayn C (2000) The changing time-series properties of earnings, cash flows and accruals: Has financial reporting become more conservative? Journal of Accounting and Economics 29:287–320.

Glimcher PW (2011) Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. Proceedings of the National Academy of Sciences 108:15647–15654.

Glimcher PW, Rustichini A (2004) Neuroeconomics: the consilience of brain and decision. Science 306:447–452.

Gottfried JA, O'Doherty JP, Dolan RJ (2002) Appetitive and aversive olfactory learning in humans studied using event-related functional magnetic resonance imaging. Journal of Neuroscience 22:10829–10837.

Gray TS (1999) Functional and Anatomical Relationships among the Amygdala, Basal Forebrain, Ventral Striatum, and Cortex: An Integrative Discussion. Annals of the New York Academy of Sciences 877:439–444.

Grinblatt M, Keloharju M (2001) What makes investors trade? Journal of Finance 56:589–616.

Haber SN, Knutson B (2010) The reward circuit: linking primate anatomy and human imaging. Neuropsychopharmacology 35:4–26.

Hampton AN, O'Doherty JP (2007) Decoding the neural substrates of reward-related decision making with functional MRI. Proceedings of the National Academy of Sciences 104:1377–1382.

Hare TA, O'Doherty JP, Camerer CF, Schultz W, Rangel A (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.

Hellstern F, Malaka R, Hammer M (1998) Backward inhibitory learning in honeybees: a behavioral analysis of reinforcement processing. Learning & Memory 4:429–444.

Hollerman JR, Schultz W (1998) Dopamine neurons report an error in the temporal prediction of reward during learning. Nature Neuroscience 1:304–309.

Holt CA, Laury SK (2002) Risk aversion and incentive effects. The American Economic Review 92:1644 – 1655.

Horvitz J (2000) Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. Neuroscience 96:651–656.

Huettel SA, Song AW, McCarthy G (2009) Functional Magnetic Resonance Imaging. Sinauer Associates.

Hung M-W, Yu H-Y (2006) A heterogeneous model of disposition effect. Applied Economics 38:2147–2157.

Hyland B., Reynolds JN., Hay J, Perk C., Miller R (2002) Firing modes of midbrain dopamine cells in the freely moving rat. Neuroscience 114:475–492.

Izuma K, Saito DN, Sadato N (2008) Processing of social and monetary rewards in the human striatum. Neuron 58:284–294.

Izuma K, Saito DN, Sadato N (2010) Processing of the incentive for social approval in the ventral striatum during charitable donation. Journal of Cognitive Neuroscience 22:621–631.

Jensen J, McIntosh AR, Crawley AP, Mikulis DJ, Remington G, Kapur S (2003) Direct activation of the ventral striatum in anticipation of aversive stimuli. Neuron 40:1251–1257.

Johnstone T, Ores Walsh KS, Greischar LL, Alexander AL, Fox AS, Davidson RJ, Oakes TR (2006) Motion correction and the use of motion covariates in multiple-subject fMRI analysis. Human Brain Mapping 27:779–788.

Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. Nature Neuroscience 10:1625–1633.

Kahneman D, Tversky A (1979) Prospect Theory: An Analysis of Decision under Risk. Econometrica 47:263–291.

Kaniel R, Liu S, Saar G, Titman S (2012) Individual Investor Trading and Return Patterns around Earnings Announcements. The Journal of Finance 67:639–680.

Kaustia M (2010) Prospect theory and the disposition effect. Journal of Financial and Quantitative Analysis 45:791–812.

Kelley AE, Berridge KC (2002) The Neuroscience of Natural Rewards: Relevance to Addictive Drugs. Journal of Neuroscience 22:3306–3311.

Knutson B, Adams CM, Fong GW, Hommer D (2001a) Anticipation of increasing monetary reward selectively recruits nucleus accumbens. Journal of Neuroscience 21:RC159.

Knutson B, Fong G, Adams C, Varner J, Hommer D (2001b) Dissociation of reward anticipation and outcome with event-related fMRI. Neuroreport 12:3683–3687.

Knutson B, Taylor J, Kaufman M, Peterson R, Glover GH (2005) Distributed neural representation of expected value. Journal of Neuroscience 25:4806–4812.

Knutson B, Westdorp A, Kaiser E, Hommer D (2000) FMRI visualization of brain activity during a monetary incentive delay task. NeuroImage 12:20–27.

Koyama T, McHaffie JG, Laurienti PJ, Coghill RC (2005) The subjective experience of pain: where expectations become reality. Proceedings of the National Academy of Sciences 102:12950–12955.

Kringelbach ML (2005) The human orbitofrontal cortex: linking reward to hedonic experience. Nature Reviews Neuroscience 6:691–702.

Kringelbach ML, Rolls ET (2003) Neural correlates of rapid reversal learning in a simple model of human social interaction. NeuroImage 20:1371–1383.

Kuhnen CM, Knutson B (2005) The neural basis of financial risk taking. Neuron 47:763–770.

LaBar KS, Gatenby JC, Gore JC, LeDoux JE, Phelps EA (1998) Human Amygdala Activation during Conditioned Fear Acquisition and Extinction: a Mixed-Trial fMRI Study. Neuron 20:937–945.

Lakshminarayanan VR, Chen MK, Santos LR (2010) The evolution of decision-making under risk: framing effects in monkey risk preferences. Journal of Experimental Social Psychology 47:693–689.

Lavoie P, Grondin S (2004) Information processing limitations as revealed by temporal discrimination. Brain and Cognition 54:198–200.

Leknes S, Tracey I (2008) A common neurobiology for pain and pleasure. Nature Reviews Neuroscience 9:314–320.

Levin IP, Xue G, Weller JA, Reimann M, Lauriola M, Bechara A (2012) A neuropsychological approach to understanding risk-taking for potential gains and losses. Frontiers in Neuroscience 6:15.

Libby R, Bloomfield R, Nelson MW (2002) Experimental research in financial accounting. Accounting, Organizations and Society 27:775–810.

Lieberman MD, Cunningham WA (2009) Type I and Type II error concerns in fMRI research: re-balancing the scale. Social Cognitive and Affective Neuroscience 4:423–428.

Liu X, Powell DK, Wang H, Gold BT, Corbly CR, Joseph JE (2007) Functional dissociation in frontal and striatal areas for processing of positive and negative reward information. Journal of Neuroscience 27:4587–4597.

Loewenstein G, Rick S, Cohen JD (2008) Neuroeconomics. Annual Review of Psychology 59:647–672.

Logothetis NK (2002) The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. Philosophical Transactions of the Royal Society of London 357:1003–1037.

Logothetis NK (2008) What we can do and what we cannot do with fMRI. Nature 453:869–878.

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. Nature 412:150–157.

Loring D., Meador K., Allison J., Pillai J., Lavin T, Lee G., Balan A, Dave V (2002) Now you see it, now you don't: statistical and methodological considerations in fMRI. Epilepsy & Behavior 3:539–547.

Matsumoto M, Hikosaka O (2007) Lateral habenula as a source of negative reward signals in dopamine neurons. Nature 447:1111–1115.

Matsumoto M, Hikosaka O (2009) Two types of dopamine neuron distinctly convey positive and negative motivational signals. Nature 459:837–841.

McClure SM, Berns GS, Montague PR (2003) Temporal prediction errors in a passive learning task activate human striatum. Neuron 38:339–346.

McClure SM, Laibson DI, Loewenstein G, Cohen JD (2004a) Separate neural systems value immediate and delayed monetary rewards. Science 306:503–507.

McClure SM, York MK, Montague PR (2004b) The neural substrates of reward processing in humans: the modern role of FMRI. The Neuroscientist 10:260–268.

Mirenowicz J, Schultz W (1996) Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. Nature 379:449–451.

Montague PR, Berns GS (2002) Neural economics and the biological substrates of valuation. Neuron 36:265–284.

Montague PR, Dolan RJ, Friston KJ, Dayan P (2012) Computational psychiatry. Trends in cognitive sciences 16:72–80.

Montague PR, Hyman SE, Cohen JD (2004) Computational roles for dopamine in behavioural control. Nature 431:760–767.

Moore RY, Bloom FE (1978) Central catecholamine neuron systems: anatomy and physiology of the dopamine systems. Annual review of neuroscience 1:129–169.

Moskvina V, Schmidt KM (2008) On multiple-testing correction in genome-wide association studies. Genetic epidemiology 32:567–573.

Murray EA, O'Doherty JP, Schoenbaum G (2007) What we know and do not know about the functions of the orbitofrontal cortex after 20 years of cross-species studies. Journal of Neuroscience 27:8166–8169.

Nicola SM, Surmeier J, Malenka RC (2000) Dopaminergic modulation of neuronal excitability in the striatum and nucleus accumbens. Annual Review of Neuroscience 23:185–215.

Nicola SM, Yun IA, Wakabayashi KT, Fields HL (2004) Cue-evoked firing of nucleus accumbens neurons encodes motivational significance during a discriminative stimulus task. Journal of Neurophysiology 91:1840–1865.

Nitschke JB, Sarinopoulos I, Mackiewicz KL, Schaefer HS, Davidson RJ (2006) Functional neuroanatomy of aversion and its anticipation. NeuroImage 29:106–116.

Noussair C, Matheny K (2000) An experimental study of decisions in dynamic optimization problems. Economic Theory 15:389–419.

Odean T (1998) Are investors reluctant to realize their losses? Journal of Finance 53:1775–1798.

Ogawa S (1990) Brain Magnetic Resonance Imaging with Contrast Dependent on Blood Oxygenation. Proceedings of the National Academy of Sciences 87:9868–9872.

Olds J, Milner P (1954) Positive Reinforcement Produced By Electrical Stimulation of Septal Area And Other Regions of Rat Brain. Journal of Comparative and Physiological Psychology 47:419–427.

O'Doherty JP, Dayan P, Friston KJ, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. Neuron 38:329–337.

O'Doherty JP, Dayan P, Schultz J, Deichmann R, Friston KJ, Dolan RJ (2004) Dissociable roles of ventral and dorsal striatum in instrumental conditioning. Science 304:452–454.

O'Doherty JP, Hampton AN, Kim H (2007) Model-based fMRI and its application to reward learning and decision making. Annals of the New York Academy of Sciences 1104:35–53.

O'Doherty JP, Kringelbach ML, Rolls ET, Hornak J, Andrews C (2001) Abstract reward and punishment representations in the human orbitofrontal cortex. Nature Neuroscience 4:95–102.

Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. Nature 441:223–226.

Pagnoni G, Zink CF, Montague PR, Berns GS (2002) Activity in human ventral striatum locked to errors of reward prediction. Nature Neuroscience 5:97–98.

Pan W-X, Schmidt R, Wickens JR, Hyland BI (2005) Dopamine cells respond to predicted events during classical conditioning: evidence for eligibility traces in the reward-learning network. Journal of Neuroscience 25:6235–6242.

Paulus MP, Hozack N, Frank L, Brown GG (2002) Error rate and outcome predictability affect neural activation in prefrontal cortex and anterior cingulate during decision-making. NeuroImage 15:836–846.

Paulus MP, Rogalsky C, Simmons A, Feinstein JS, Stein MB (2003) Increased activation in the right insula during risk-taking decision making is related to harm avoidance and neuroticism. NeuroImage 19:1439–1448.

Pisani A, Bonsi P, Picconi B, Tolu M, Giacomini P, Scarnati E (2001) Role of tonically-active neurons in the control of striatal function: Cellular mechanisms and behavioral correlates. Biological Psychiatry 25:211–230.

Plassmann H, O'Doherty JP, Rangel A (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. Journal of Neuroscience 27:9984–9988.

Plassmann H, O'Doherty JP, Shiv B, Rangel A (2008) Marketing actions can modulate neural representations of experienced pleasantness. Proceedings of the National Academy of Sciences 105:1050–1054.

Pochon J-B, Riis J, Sanfey AG, Nystrom LE, Cohen JD (2008) Functional imaging of decision conflict. Journal of Neuroscience 28:3468–3473.

Preuschoff K, Bossaerts PL, Quartz SR (2006) Neural differentiation of expected reward and risk in human subcortical structures. Neuron 51:381–390.

Preuschoff K, Quartz SR, Bossaerts PL (2008) Human insula activation reflects risk prediction errors as well as risk. Journal of Neuroscience 28:2745–2752.

Raij TT, Numminen J, Närvänen S, Hiltunen J, Hari R (2005) Brain correlates of subjective reality of physically and psychologically induced pain. Proceedings of the National Academy of Sciences 102:2147–2151.

Raylu N, Oei TP. (2002) Pathological gambling. Clinical Psychology Review 22:1009–1061.

Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Classical Conditioning II: Current Research and Theory (Black A, Prokasy W, eds), pp 64–99. New Yorks: Appleton Century Crofts.

Rilling JK, Gutman D, Zeh T, Pagnoni G, Berns GS, Kilts C (2002) A neural basis for social cooperation. Neuron 35:395–405.

Roesch MR, Calu DJ, Schoenbaum G (2007) Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. Nature Neuroscience 10:1615–1624.

Rolls ET, Kringelbach ML, de Araujo IET (2003) Different representations of pleasant and unpleasant odours in the human brain. European Journal of Neuroscience 18:695–703.

Rushworth MFS, Kennerley SW, Walton ME (2005) Cognitive neuroscience: resolving conflict in and over the medial frontal cortex. Current Biology 15:R54–6.

Salas R, Baldwin P, de Biasi M, Montague PR (2010) BOLD Responses to Negative Reward Prediction Errors in Human Habenula. Frontiers in Human Neuroscience 4:36.

Samanez-Larkin GR, Hollon NG, Carstensen LL, Knutson B (2008) Individual differences in insular sensitivity during loss anticipation predict avoidance learning. Psychological Science 19:320–323.

Schultz W (1997) A Neural Substrate of Prediction and Reward. Science 275:1593–1599.

Schultz W (2010) Dopamine signals for reward value and risk: basic and recent data. Behavioral and Brain Functions 6:24.

Schultz W, Apicella P, Ljungberg T (1993) Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. Journal of Neuroscience 13:900–913.

Schultz W, Apicella P, Scarnati E, Ljungberg T (1992) Neuronal activity in monkey ventral striatum related to the expectation of reward. Journal of Neuroscience 12:4595–4610.

Scott DJ, Heitzeg MM, Koeppe RA, Stohler CS, Zubieta J-K (2006) Variations in the human pain stress experience mediated by ventral and dorsal basal ganglia dopamine activity. Journal of Neuroscience 26:10789–10795.

Scott DJ, Stohler CS, Egnatuk CM, Wang H, Koeppe RA, Zubieta J-K (2007) Individual differences in reward responding explain placebo-induced expectations and effects. Neuron 55:325–336.

Seymour B, Daw ND, Dayan P, Singer T, Dolan RJ (2007) Differential encoding of losses and gains in the human striatum. Journal of Neuroscience 27:4826–4831.

Seymour B, McClure SM (2008) Anchors, scales and the relative coding of value in the brain. Current Opinion in Neurobiology 18:173–178.

Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS (2004) Temporal difference models describe higher-order learning in humans. Nature 429:664–667.

Seymour B, O'Doherty JP, Koltzenburg M, Wiech K, Frackowiak R, Friston KJ, Dolan RJ (2005) Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. Nature Neuroscience 8:1234–1240.

Sharp C, Monterosso J, Montague PR (2012) Neuroeconomics: a bridge for translational research. Biological Psychiatry 72:87–92.

Shefrin H, Statman M (1985) The disposition to sell winners too early and ride losers too long: theory and evidence. Journal of Finance 40:777–790.

Skinner DJ, Sloan RG (2002) Earnings Surprises, Growth Expectations, and Stock Returns or Don't Let an Earnings Torpedo Sink Your Portfolio. Review of Accounting Studies 7:289–312.

Small DM, Zatorre RJ, Dagher A, Evans AC, Jones-Gotman M (2001) Changes in brain activity related to eating chocolate: from pleasure to aversion. Brain 124:1720–1733.

Smeets PAM, de Graaf C, Stafleu A, van Osch MJP, Nievelstein RAJ, van der Grond J (2006) Effect of satiety on brain activation during chocolate tasting in men and women. The American Journal of Clinical Nutrition 83:1297–1305.

Smith Y, Kieval JZ (2000) Anatomy of the dopamine system in the basal ganglia. Trends in Neurosciences 23:S28–S33.

Staudinger MR, Erk S, Abler B, Walter H (2009) Cognitive reappraisal modulates expected value and prediction error encoding in the ventral striatum. NeuroImage 47:713–721.

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. IEEE Transactions on Neural Networks 9:1054.

Talairach J, Tournoux P (1988) Co-planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System - an Approach to Cerebral Imaging. New York, NY: Thieme Medical Publishers.

Talmi D, Dayan P, Kiebel SJ, Frith CD, Dolan RJ (2009) How humans integrate the prospects of pain and reward during choice. Journal of Neuroscience 29:14617–14626.

Tanaka T, Camerer CF, Nguyen Q (2010) Risk and time preferences: linking experimental and household survey data from Vietnam. American Economic Review 100:557–571.

Teghtsoonian R (1971) On the exponents in Stevens' law and the constant in Ekman's law. Psychological Review 78:71–80.

Thaler RH, Johnson EJ (1990) Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice. Management Science 36:643–660.

Tobler PN, Fiorillo CD, Schultz W (2005) Adaptive coding of reward value by dopamine neurons. Science 307:1642–1645.

Tobler PN, O'Doherty JP, Dolan RJ, Schultz W (2007) Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. Journal of Neurophysiology 97:1621–1632.

Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision-making under risk. Science 315:515–518.

Toneatto T (1999) Cognitive psychopathology of problem gambling. Substance Use & Misuse 34:1593–1604.

Tversky A, Kahneman D (1971) Belief in the law of small numbers. Psychological Bulletin 76:105–110.

Tversky A, Kahneman D (1992) Advances in prospect theory: cumulative representation of uncertainty. Journal of Risk and Uncertainty 5:297–323.

Venkatraman V, Payne JW, Bettman JR, Luce MF, Huettel SA (2009a) Separate neural mechanisms underlie choices and strategic preferences in risky decision making. Neuron 62:593–602.

Venkatraman V, Rosati AG, Taren AA, Huettel SA (2009b) Resolving response, decision, and strategic control: evidence for a functional topography in dorsomedial prefrontal cortex. Journal of Neuroscience 29:13158–13164.

Vlcek M, Wang M (2008) Speculating in Gains, Waiting in Losses: An Experimental Study on the Disposition Effect. Available at SSRN 1101627.

Von Neumann J, Morgenstern O (1947) Theory of games and economic behavior. Princeton: Princeton Univ. Press.

Weber M, Camerer CF (1998) The disposition effect in securities trading: an experimental analysis. Journal of Economic Behavior & Organization 33:167–184.

West AR, Floresco SB, Charara A, Rosenkranz JA, Grace AA (2003) Electrophysiological Interactions between Striatal Glutamatergic and Dopaminergic Systems. Annals of the New York Academy of Sciences 1003:53–74.

Wood PB, Holman AJ (2009) An elephant among us: the role of dopamine in the pathophysiology of fibromyalgia. Journal of Rheumatology 36:221–224.

Wynne CDL, Staddon JER (1998) Models of Action: Mechanisms for Adaptive Behavior. Routledge.

Xue G, Lu Z, Levin IP, Weller JA, Li X, Bechara A (2009) Functional dissociations of risk and reward processing in the medial prefrontal cortex. Cerebral Cortex 19:1019–1027.

Yacubian J, Gläscher J, Schroeder K, Sommer T, Braus DF, Büchel C (2006) Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. Journal of Neuroscience 26:9530–9537.

Yacubian J, Sommer T, Schroeder K, Gläscher J, Braus DF, Büchel C (2007) Subregions of the ventral striatum show preferential coding of reward magnitude and probability. NeuroImage 38:557–563.

Zaghloul KA, Blanco JA, Weidemann CT, McGill K, Jaggi JL, Baltuch GH, Kahana MJ (2009) Human substantia nigra neurons encode unexpected financial rewards. Science 323:1496–1499.

Zink CF (2005) Human striatal activity : A signal for saliency rather than reward. :132.

Zink CF, Pagnoni G, Martin ME, Dhamala M, Berns GS (2003) Human striatal response to salient nonrewarding stimuli. Journal of Neuroscience 23:8092–8097.

Zink CF, Pagnoni G, Martin-Skurski ME, Chappelow JC, Berns GS (2004) Human Striatal Responses to Monetary Reward Depend On Saliency. Neuron 42:509–517.

Zysset S, Wendt CS, Volz KG, Neumann J, Huber O, von Cramon DY (2006) The neural implementation of multi-attribute decision making: a parametric fMRI study with human subjects. NeuroImage 31:1380–1388.