THE NEURAL MECHANISMS OF INTERTEMPORAL AND RISKY CHOICE

By

BRIAN G. ESSEX

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Psychology

December, 2011

Nashville, Tennessee

Approved:

David H. Zald

Paige Marta Skiba

Stephen D. Benning

René Marois

To my parents and brother

ACKNOWLEGDGEMENTS

I would first like to thank my advisor David Zald for all of the guidance and encouragement he has offered throughout my years as a doctoral student. He always provided me with guidance when I needed it, and also provided me with the independence to pursue my interests. I have learned a great deal from David, and am extremely grateful for his support.

I am also grateful for the invaluable advice and suggestions of the members of my dissertation committee: Stephen Benning, René Marois, and Paige Marta Skiba. My work has greatly benefitted from their expertise. I would additionally like to thank Martin Gallagher for providing consultation on medical issues, Warren Lambert for providing statistical consultation, and Jacob Sagi for providing consultation for my dissertation grant application to the NSF.

Special thanks are due to my friends and colleagues in the Affective Neuroscience Laboratory at Vanderbilt, especially Josh Buckholtz, Suzie Dukic, Joseph Kim, Maureen McHugo, Greg Samanez-Larkin, and Michael Treadway. Each of them has provided contributions to my development as a graduate student. Finally, I would not have completed this work without the help of four undergraduate students who assisted in collecting and managing the data: Sarah Clinton, Rebecca Qian, Tawny Spinelli, and Lucas Wonderley. I am grateful for each of their contributions.

This work was supported by a doctoral dissertation improvement grant from the National Science Foundation (SES-1061913).

TABLE OF CONTENTS

LIST OF TABLES

Table Page

LIST OF FIGURES

CHAPTER I

INTRODUCTION AND BACKGROUND

Introduction

In everyday life, people must frequently make intertemporal choices, which involve selecting between options that lead to costs or benefits at different points in time. Examples of such choices range from deciding whether to eat a healthy meal or a tastier, unhealthy meal to deciding whether to save money for retirement or spend it on something sooner. In both of these cases, an individual must decide whether they would prefer an option that has larger long-term benefits, but fewer short-term or immediate benefits than the alternative.

Many choices also involve another element: risk. Any option that leads to an outcome with less than 100% probability is risky, and it becomes more risky as the probability that it leads to a beneficial outcome decreases and as the probability that it leads to a non-beneficial or negative outcome increases. In situations involving risk, people often are confronted with a choice between a risky option that could lead to a better outcome than a safer alternative, but that also has a higher probability of leading to a worse outcome. Examples of such situations range from deciding whether or not to a buy a stock or a bond to deciding whether to have one of two medical procedures.

This dissertation examines the neural mechanisms of intertemporal choice and choice involving risk. In particular, it focuses on the role of the Dorsolateral Prefrontal

Cortex (DLPFC) in both types of choices, and on the role of the Posterior Parietal Cortex (PPC) in intertemporal choice. While both Experiments 1 and 2 examine the neural mechanisms of economic choice, Experiment 3 has a different focus. This third experiment examines the relationships between preferences on intertemporal and risky choice tasks, and between these preferences and impulsivity.

Evidence reveals that the Dorsolateral Prefrontal Cortex (DLPFC) has an important role in intertemporal choice. Disruption of this region with low frequency repetitive Transcranial Magnetic Stimulation (rTMS) leads to greater selection of an option for monetary gain that has a better immediate, but a worse long-term value than the alternative (Figner et al., 2010). Experiment 1 of this dissertation examines the neural mechanisms of intertemporal choice and addresses the following main questions:

1. Does the Posterior Parietal Cortex (PPC) have a similar role in intertemporal choice as the DLPFC?
2. Do the DLPFC and PPC have a similar role in intertemporal choices involving gains and in intertemporal choices involving losses?

The DLPFC also has an important role in choices involving risk. Disruption of the DLPFC with low frequency rTMS leads to greater selection of a potentially larger risky option that has a lower probability of reward, and a higher probability of loss, than the alternative (Knoch et al., 2006). Experiment 2 of this dissertation examines the role of the DLPFC in risky choice and addresses the following main questions:

1. Does the role of the DLPFC in risky choice depend on the level of risk?
2. Does the DLPFC have a similar role in risky choices that are limited to the gain domain (i.e. that do not contain the possibility of a loss)?

People tend to discount the value of delayed monetary incentives relative to immediate incentives (Frederick, Loewenstein, & O'Donoghue, 2004; Murphy, Vuchinich, & Simpson, 2001), and tend to discount the value of risky incentives with respect to certain incentives (Rachlin, Raineri, & Cross, 1991). People, however, differ in these tendencies. Experiment 3 focuses on these individual differences and addresses the following main questions:

1. Are there associations between types of monetary discounting (i.e. delayed gains, delayed losses, risky gains) and impulsivity?
2. Are the tendencies to discount the value of different types of incentives related?

This dissertation is organized into five chapters with the following content:

Chapter I reviews important background information for the experiments performed in later chapters. The next sections of Chapter I focus on the following topics:

- Economic theories of intertemporal and risky choice, and behavioral economics findings of how individuals value delayed and risky incentives
- The brain processes involved in the prediction of positive value, in the prediction of negative value, and in the components of predicted value (i.e. probability, magnitude, and expected value)
- The brain processes involved in intertemporal and risky choice. This section describes the key experiments that led to the research questions for Experiments 1 and 2.
- The functions of the DLPFC and of the lateral Prefrontal Cortex more generally (with a focus on specific functions relevant to economic choice)

- The functions of the PPC (with a focus on specific functions relevant to economic choice)
- Potential Roles of the DLPFC and PPC in intertemporal and risky choice
- Summary of the Research Plan of the experiments

Chapters II, III, and IV each focus on one experiment (i.e. Experiments 1, 2, and 3, respectively). Each of these chapters contains a full research report, including an introduction, methods, results, and discussion. While the background for Experiments 2 and 3 is presented in Chapter I, much of the background for Experiment 3 is presented in Chapter IV.

Chapter V concludes the dissertation with a broad discussion of some of the key findings of the experiments and of other important questions related to intertemporal and risky choice.

Economic Theory and Behavioral Economics Research

Choice involving Risk

In standard economic theory, each good in the marketplace is said to have a specific utility (i.e. subjective value) for an individual (Pindyck & Rubinfeld, 2006). A rational individual who is a utility maximizer makes choices in order to attain the bundle of goods that maximizes their own utility. It is traditionally thought that when making decisions amongst monetary options, the utility of each choice option is calculated as a function of the probability and magnitude of each potential monetary outcome associated with that option. This idea that both the probability and magnitude of reward influence

decisions is a central idea in economics that goes back to work done in the 17th century by Blaise Pascal, Antoine Arnauld, and Pierre Nicole (Glimcher, 2003). Their insights led to the creation of Expected Value Theory, which held that individuals should make choices for options that offer the highest expected value. The expected value of a choice option is calculated as the sum of the product of the probability (p) and magnitude (v) of each possible monetary outcome of that choice option (i.e. Expected Value = $\sum_i p_i * v_i$).

Later, Expected Value Theory was modified by Daniel Bernoulli in the eighteenth century (Bernoulli, 1738; 1954) who asserted that individuals choose options with the highest expected utility rather than with the highest expected value. It can be mathematically stated: Expected Utility = $\sum_i p_i * u_i$, meaning that the expected utility of any option is equal to the sum across outcomes of the option of the utility of each outcome (u) multiplied by its probability (p). Although the utility of an option is a function of its magnitude (e.g. as the magnitude of a potential monetary gain increases, its utility increases as well) the two terms are not synonymous. Unlike magnitude, utility was designed to be a function of an individual's total level of wealth.

The utility function designed by Bernoulli was concave, indicating that as the level of a person's wealth increases, the increase in utility from obtaining the same magnitude of monetary gain diminishes. The concave curvature of the utility function also assumes that individuals are risk averse. A risky choice option can be defined as an option available with a less than certain probability that is known to the individual making the choice. A risk-averse individual is someone who prefers a certain option to a risky option of equal expected value (http://www.econport.org/content/handbook.html). For example, suppose a risk-averse individual were to make a choice between $x for sure

and $2x with a probability of 50%. Since their utility function of money is concave, this individual is presumed to get more utility from an increase from 0 to $x than from $x to $2x. Because of this, the individual is more likely to choose the certain option, since the expected utility is greater than that of the risky option, even though they both have the same expected value. In contrast to this prediction of Expected Utility Theory, Expected Value Theory would hold that because both of the options have the same expected value, this individual would be indifferent between the two options.

Expected Utility Theory has been an extremely useful theory in economics and was an improvement over the older Expected Value Theory, but it still does not adequately describe a great deal of human behavior. In 1979, Kahneman and Tversky developed an alternative to Expected Utility Theory, Prospect Theory, in an attempt to provide a more realistic picture of how individuals make decisions and account for behavioral deviations from Expected Utility Theory (Kahneman & Tversky, 1979). Prospect theory holds that there are two phases in the process of making choices between two different options for monetary gambles: an editing phase and an evaluation phase. In the editing phase, individuals simplify and organize the information associated with each choice option, such as by cancelling out similarities between the two options and by rounding values. These editing operations are a departure from the assumption of Expected Utility Theory that individuals consider options exactly as they are presented. Because of editing operations, a person's preferences over the same choice set may change if they use different editing operations in different situations. During the second phase, evaluation, individuals compare the value of the two options, and then choose the one with the highest value. The value of a prospect (i.e. an option offering chances of

winning and/or losing money) is computed as a function of a value function and of a probability weighting function.

The value function in Prospect Theory is concave for gains and convex for losses and steeper for losses than gains. The differential curvature of the gain and loss functions indicates that people tend to be risk averse in choices involving sure gains, but risk seeking in choices involving sure losses, while the differential steepness accounts for loss aversion. Another concept in Prospect Theory is that the value of each outcome is multiplied by a decision weight. These weights are a function of the probability of the outcome, but are not identical to it. Decision weights represent the impact that probability has on the desirability of an option, and are able to capture findings that people tend to underweight outcomes that are merely probable in comparison to those that are certain, and tend to overweight outcomes of very low probability.

While the original formulations of Expected Utility Theory and Prospect Theory both assumed that individuals are risk averse when monetary gains are involved, not all individuals have similar risk preferences. As stated earlier, in Expected Utility Theory, utility functions are concave to reflect risk-aversion. However, those with different relationships to risk have utility functions with different forms of curvature (Weber & Johnson, 2009; http://www.econport.org/content/handbook.html). For a risk-seeking individual, that is someone who prefers a risky option over a certain option of equal expected value, the utility curve is convex, rather than concave. In contrast, a risk-neutral individual who does not strictly prefer either of the two options has a linear utility function that exhibits no curvature. While there are three categorically different relationships to risk, the degree of risk aversion or risk seeking differs among individuals.

The level of a person's risk preference can be assessed by calculating that individual's certainty equivalent for a risky option (http://www.econport.org/content/handbook.html), which is the amount of money available for certain that a person values subjectively equally to a risky amount of money. The difference between the certainty equivalent of a risky option and the expected value of the risky option indicates how far that individual's preferences are from risk-neutrality.

Not all risky options are risky to the same degree, and this is reflected in terms that measure the level of riskiness of a monetary option. In finance, the level of risk of an option is taken to be a function of the spread of outcomes associated with it (Rothschild & Stiglitz, 1970), and is often defined specifically as the variance or standard deviation of the different outcomes (Glimcher, 2008; Weber & Johnson, 2009). By this definition, an option that offers $20 with 50% probability and $30 with 50% probability is riskier than an option of equal expected value that offers $22.50 with 50% probability and $27.50 with 50% probability, because the variance of the outcomes of the first option is larger than that of the second.

Another definition of the level of risk, however, will be used throughout the rest of this dissertation. Unless stated otherwise, the level of risk will be taken as a direct function of the probability of a negative outcome or non-rewarding outcome on a task. In choices involving monetary gains, the choice that offers a gain with a smaller probability (i.e. greater probability of no gain) will be termed riskier. In choices involving monetary losses, the choice that offers a loss with a higher probability will be termed riskier. Thus, greater risk will be taken as lesser probability of gain or greater probability of a loss. This definition can only easily be applied to choice options that have two possible

outcomes, and for which one possibility is either a monetary loss or no monetary reward. Although risk level is not typically defined this way, I will use this definition because it can be used to help make sense of much of the findings in neuroeconomics and because it provides a parsimonious definition of risk level for choice scenarios it applies to. Importantly, it can easily be applied to the two experiments of this dissertation that investigated risky choice, which both employed tasks that involved choices between a certain option and a risky option with known probability.

Intertemporal Choice: Discounted Utility Theory and Anomalies

Expected Utility Theory and Prospect Theory apply to choices for monetary incentives available immediately. However, there is another type of choice problem, called intertemporal choices, in which choice options lead to outcomes available at different points in time. While the dominant economic model for choices involving risk has been Expected Utility Theory, the dominant economic model for intertemporal choices has been the Discounted Utility Theory articulated by Samuelson (1937). Like Expected Utility Theory, Discounted Utility Theory holds that individuals make decisions to maximize their utility, which is a function of the magnitudes of monetary gains and losses.

According to Discounted Utility Theory, the utility associated with future consumption is discounted by a constant value as an exponential function of time. It can be mathematically stated that the utility (U) obtained from choosing a specific consumption profile is: $U(c_o,\ldots, c_T) = \sum_{t=0}^{T} \delta^t * u(c_t)$, where t is time, $u(c_t)$ is the utility of an item of consumption at time t, and δ is the person's discount rate. In other words, the

amount of utility obtained from choosing a consumption bundle is equal to a sum of the utilities from all the options with one caveat: future consumption is discounted as an exponential function of the discount rate, δ, and the delay to consumption, t. The appeal of this model is in its simplicity and that it models an important observation. Individuals do tend to discount the value of money that could be received in the future compared to money available immediately, and tend to discount the value of money more as the time of receiving it becomes more delayed (Frederick et al., 2004).

Despite the appeal of Discounted Utility Theory, there is an important pattern of behavior that it does not adequately capture. It has been observed that individuals' preferences for monetary rewards often reverse when delays to both rewards are incremented by a constant amount (Loewenstein & Prelec, 1992). For example an individual may prefer receiving $1 today to $2 tomorrow, but prefer receiving $2 in 8 days to $1 in 7 days. This behavioral pattern violates Discounted Utility Theory, which holds that the preference between two options is dependent only on the amount of time between them and thus should not switch at any time point (Koopmans, 1960). Preference reversals reveal that individuals do not treat all similar intervals between rewards equally, and that an exponential model of discounting is not the best fit to actual behavior. Rather, discounting tends to be greater between the present and a future time point than it is between two future time points separated by the same temporal interval. Individuals are thus more impatient than the standard model of discounted utility predicts.

By identifying how individuals value money available at different time points, many studies have been able to assess individuals' actual levels of delay discounting

(Kirby & Marakovic, 1995; Madden, Begotka, Raiff, & Kastern, 2003; Murphy et al., 2001; Myerson & Green, 1995; Rachlin et al., 1991; Thaler, 1991). One way to assess this is to ask individuals to make several choices between amounts of money available immediately and amounts available in the future, until one uncovers the specific immediate monetary values (i.e. immediate equivalents) that the individual values equally to specific delayed monetary values. These specific delayed and immediate values that are subjectively valued equivalently can be used to construct indifference curves which map out amounts of money available at particular points in the future that are subjectively valued equally (i.e. have a similar level of expected utility) to specific amounts of money available now. A number of studies show that intertemporal indifference curves are better modeled by hyperbolic or quasi-hyperbolic functions than exponential ones (Kirby & Marakovic, 1995; Madden et al., 2003; Murphy et al., 2001; Myerson & Green, 1995; Rachlin et al., 1991), which are steeper than exponential functions in short time frames, but shallower than exponential functions between time points far in the future (see Figure 1). Critically, the shape of hyperbolic and quasi-hyperbolic functions allows these models to account for preference reversals, which are not accounted for by the exponential discounting model of Discounted Utility Theory.

Just as individuals differ in their risk preferences, individuals differ in their time preferences. Since individuals tend to discount delayed rewards (i.e. are delay aversive), as mentioned earlier, these differences tend to be differences in levels of discounting rather than qualitatively different time preferences (i.e. delay averse vs. delay seeking). The steepness of an individual's indifference curve is a function of that person's delay

discount rate. Steeper decreasing curves reflect a higher discount rate (i.e. greater preference for incentives available sooner) than do shallower decreasing curves.

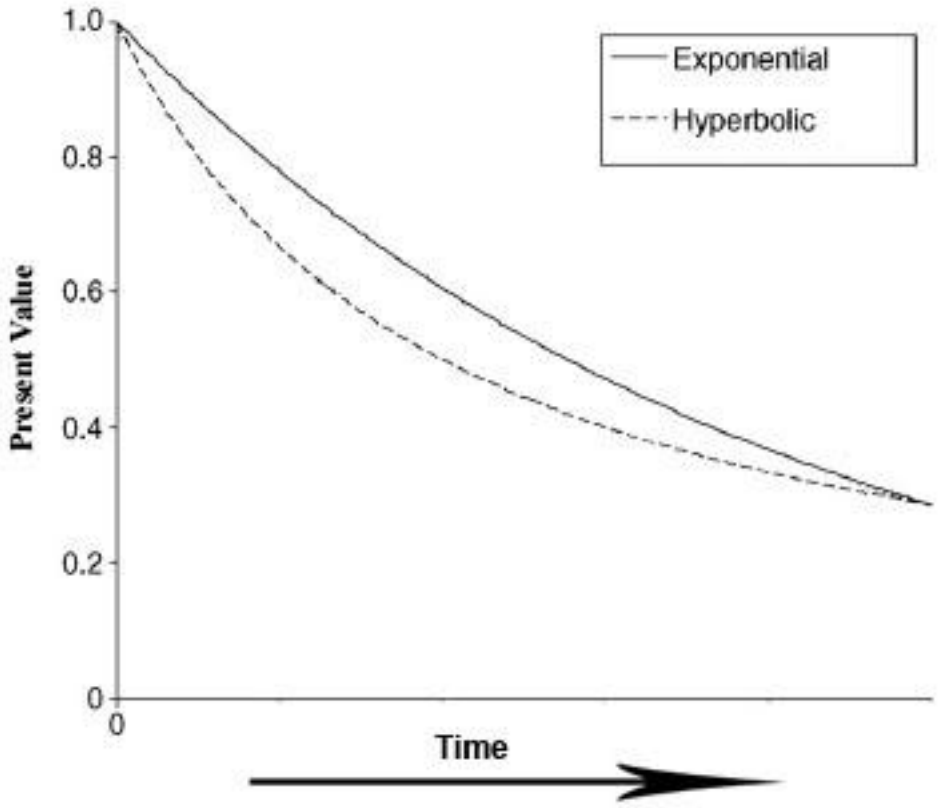

Figure 1: Intertemporal Choice Indifference Curves. Discount functions showing the present value of money that is subjectively equivalent to a constant amount of money available at various points in the future. Exponential discounting assumes constant discounting over time while hyperbolic discounting assumes faster discounting at shorter time scales than at longer time scales. Here, exponential discounting is described by a function of this form: Present Value = Ve^(-KD). In contrast, hyperbolic discounting is described by a function of this form: Present Value = V/(1+KD). K is the discount rate, and is individual specific since different people discount the value of delayed incentives by varying degrees. Larger values of K indicate greater value placed on immediate relative to delayed incentives. V is the objective value of money, and here V = 1. D is the Delay until payment, and is usually expressed in days. Figure modified from Berns, Laibson, & Loewenstein (2007).

Although the model of Discounted Utility Theory does not specify that there should be different discount rates for different types of delayed outcomes, the level of delay discounting is a function of both the valence and magnitude of a monetary outcome (Frederick et al., 2004). As with gains, individuals tend to discount the value of losses in the future relative to immediate losses. They typically prefer losing an equivalent amount of money in the future to losing the same amount immediately. Additionally, like gain discounting, loss discounting is well modeled by hyperbolic or quasi-hyperbolic discounting functions (Estle, Green, Myerson, & Holt, 2006; Murphy et al., 2001), and an individual's levels of delay discounting for monetary gains and losses are positively correlated, suggesting common underlying mechanisms (Murphy et al., 2001). However, numerous studies have reported that individuals discount delayed losses to a lesser extent than delayed gains (Baker, Johnson, & Bickel, 2003; Benzion, Rapoport, & Yagil, 1989; Estle et al., 2006; Frederick et al., 2004; Murphy et al., 2001). Furthermore, it has been shown that some individuals actually prefer to suffer a monetary loss or experience other aversive events sooner rather than later (Frederick et al., 2004).

The level of delay discounting also depends on the magnitude of the outcome. A large amount of data indicates that the rate of delay discounting is greater for smaller than larger monetary rewards (Benzion et al., 1989; Estle et al., 2006; Green, Myerson, & McFadden, 1997; Green, Myerson, & Ostaszewski, 1999; Kirby, 1997; Kirby, Petry, & Bickel, 1999; Thaler, 1991). However, the results of several studies suggest that the magnitude of an outcome may have less of an effect on the level of delay discounting for monetary losses (Baker et al., 2003; Estle et al., 2006).

Similarities and Differences between Intertemporal Choice and Choice involving Risk

While intertemporal choice and choices involving risk have typically been treated separately in economics, there are similarities between these two types of choices suggesting that both types of choices may rely on similar underlying processes. Across individuals, preferences for certainty and immediacy are positively correlated (Mitchell, 1999; Myerson, Green, Hanson, Holt, & Estle, 2003; Richards, Zhang, Mitchell, & de Wit, 1999). Additionally, just as individuals tend to discount delayed rewards relative to immediate rewards of the same magnitude, individuals tend to discount the maximum value of risky rewards relative to certain rewards of the same magnitude. Furthermore, like delay discounting, probability discounting (i.e. discounting of risky rewards) is well fit by a hyperbolic discount functions. More specifically, when probability is converted to odds against reward (Odds Against Reward = (1/probability of reward) – 1), the value of a risky reward decreases as the odds against reward increases, and this decrement in value can be modeled by a hyperbolic curve (Myerson et al., 2003; Rachlin et al., 1991) (see Figure 1 for example of a hyperbolic indifference curve). The steepness of this indifference curve reflects the level of an individual's probability discounting: a steeper curve corresponds to greater probability discounting.

These findings reveal a striking parallel. In choices involving time, individuals tend to place increased value on immediate relative to delayed rewards, while in choices involving risk they tend to place increased value on certain relative to risky rewards. And in both types of choices, preference reversals can occur. In intertemporal choices, an initial preference for a small immediate monetary gain over a larger delayed monetary gain will change to a preference for the larger delayed gain if a common delay is added to

both options, (Frederick et al., 2004). Similarly, an initial preference for a small certain monetary gain over a larger risky monetary gain will shift to a preference for the larger riskier gain if the probabilities of both options are reduced by a common ratio (Allais, 1953; Kahneman & Tversky, 1979).

It has in fact been theorized that individuals may discount the future precisely because it is risky (Keren & Roelofsma, 1995). Delayed rewards are always somewhat ambiguous, an individual never knows with certainty what the future will bring, and does not know whether or not they will be able to receive a reward in the future. If people discount rewards available in the future because they are less certain than immediate rewards, then the special weight given to immediacy could be due to a preference for certainty.

Despite the similarities between choices involving risk and intertemporal choices, there are also differences between both types of discounting, which have led others to argue against a single process account of probability and delay discounting (Christensen, Parker, Silberberg, & Hursh, 1998; Green & Myerson, 2004). Different magnitudes of monetary gain have divergent effects on each type of discounting. While levels of delay discounting become lower as the amount of the monetary gains increase, levels of probability discounting become higher (Christensen et al., 1998; Estle et al., 2006; Myerson et al., 2003). Additionally, there is some evidence that impulsivity, which can be defined as the tendencies to give into urges and to respond quickly without planning (Buss & Plomin, 1975), may have divergent relationships with each type of discounting. Many studies have found positive correlations between impulsivity and levels of delay discounting of monetary gains (Kirby et al., 1999; Madden, Petry, Badger, & Bickel,

1997; Mitchell, Fields, D'Esposito, & Boettiger, 2005; Petry, 2001, 2002; Reynolds, Richards, & de Wit, 2006). In contrast, a study that has investigated the relationship between impulsivity and probability discounting of monetary gains found a negative correlation (Mitchell, 1999). The relationships between each type of discounting and impulsivity will be more fully explained in Chapter IV.

Brain Regions involved in Value Prediction

Discussion of Neuroeconomics

The interdisciplinary field of neuroeconomics, located at the crossroads of psychology, economics, and neuroscience, is broadly concerned with uncovering how individuals make value-based decisions. The answer provided by standard economic theory is straightforward; people make choices that maximize their level of utility. According to this view, utility is not seen as something that can be directly measured, but is rather something that can only be inferred by the choices that someone makes. A very different view taken by those studying neuroeconomics is that it should be possible to measure the predicted level of utility associated with different options by recording the activity of brain regions that respond to the subjective value of the options. Knowing how the brain codes for subjective value should help predict what choices individuals will make, if the assumption holds that options that are valued highly will be chosen.

The search for the neural substrates of utility is not only concerned with how the brain codes for the subjective value of monetary options, but is also concerned with how the brain codes for the value of other types of rewards and aversive events, such as

receiving food and being attacked by a predator. Humans and other animals make choices to approach objects that are positively valued, and make choices to avoid objects that are negatively valued. Brain regions that represent subjective value probably developed to represent the value of rewards and aversive events broadly in order for an organism to make choices that increased their inclusive fitness.

In this section, I will discuss which brain regions have been associated with positive and negative valuation generally, by focusing on which regions are associated with coding for the prediction of rewards, such as money and food, and which are associated with coding for the prediction of aversive events, such as losing money or being exposed to an electric shock. I will largely limit the discussion in this section to the neural correlates of value prediction in non choice situations; how subjective value is represented in the brain in choice situations will be covered in the section entitled "Neuroeconomics of Choice".

Predicted Positive Value

Research has demonstrated that midbrain dopamine neurons have an essential role in the prediction of positive value. Midbrain dopamine neurons in the substantia nigra and ventral tegmental area in nonhuman primates phasically fire upon exposure to conditioned stimuli (i.e. sensory cues) that predict rewards (Schultz, 2006; Schultz, Dayan, & Montague, 1997). One interesting property of these neurons is that in addition to coding for reward prediction, they also code for whether or not the prediction was accurate (i.e. prediction errors), consistent with computational models of temporal difference learning (Schultz et al., 1997). Initially, before an animal has learned that a

particular conditioned stimulus (CS) leads to reward, midbrain dopamine neurons respond to the unexpected reward as it occurs. The initial firing to delivery of an unexpected reward appears to reflect a large positive prediction error, since this is a neural response in an environmental state that has a higher value than was predicted (i.e. there was a reward, but no reward was predicted). If the reward is repeatedly paired with a CS preceding it, over time the dopamine neurons will shift their phasic firing from the time of the unconditioned stimulus (US) to the time of the CS. This neural firing at the time of the CS appears to represent a prediction of the reward to come, while the subsequent lack of phasic firing at the US appears to indicate that there is no prediction error; the reward was fully predicted. There is also some evidence that dopamine neurons code for negative prediction errors; the activity of midbrain dopamine neurons is depressed if no reward is available at the predicted time (Hollerman & Schultz, 1998; Schultz et al., 1997).

Midbrain dopamine neurons not only are sensitive to whether a reward is predicted, they are sensitive to other important parameters of reward prediction. Dopamine neurons respond to the predicted magnitude of rewards; they fire more for predicted rewards of greater than of lesser magnitude (Tobler, Fiorillo, & Schultz, 2005). They also respond to the predicted probability of rewards. Phasic activity of midbrain dopamine neurons to a CS predictive of a risky reward is positively correlated with the probability that the reward will occur (Fiorillo, Tobler, & Schultz, 2003; Tobler et al., 2005). Importantly, among dopamine neurons, there is a positive correlation between the responsiveness to reward magnitude and the responsiveness to reward probability (Tobler et al., 2005). This suggests that dopamine neurons code for the expected reward value

(i.e. magnitude x probability), since individual neurons show activity that positively scales with both parameters of expected value.

In addition to responding to the magnitude and probability of predicted rewards, midbrain dopamine neurons respond to the predicted delay to rewards. Midbrain dopamine neurons phasically fire more for conditioned stimuli that predict rewards which will occur sooner than those which will occur later (Fiorillo, Newsome, & Schultz, 2008; Kobayashi & Schultz, 2008). Additionally, reductions in the firing rate of midbrain dopamine neurons as the predicted delay to reward increases on a reward prediction task are well fit by a hyperbolic discount function (Kobayashi & Schultz, 2008). Critically, these reductions may underlie hyperbolic discounting of delayed rewards, since they are seen in the neurons of monkeys that exhibit hyperbolic discounting during intertemporal choice (Kobayashi & Schultz, 2008). However, it is unclear whether dopamine neurons would discount the value, or even would be involved in predicting the value of delayed rewards available at the types of long delays (i.e. days, weeks, or years) that are typical in human intertemporal choice tasks, as the paradigm used by Kobayashi & Shultz only had delays on the order of seconds.

In recent years, a number of fMRI studies have looked at the neural correlates of predicted rewards in humans. One paradigm that has proven useful in this regard is the Monetary Incentive Delay (MID) task (Knutson, Westdorp, Kaiser, & Hommer, 2000). On this task, individuals are presented with a cue that indicates whether they will have the opportunity to respond to win an amount of money, to respond to avoid losing an amount of money, or to respond for no monetary outcome. Following each cue and a waiting period, individuals then make a response and receive feedback regarding their

performance; success on trials occurs only if subjects respond quickly enough. Because success does not always occur, cues indicate the potential to win or lose on the trial, rather than indicating there will be a certain outcome. On trials involving possibilities of monetary gain or loss, successful performance on a trial leads to feedback that they won or avoided losing the cued amount of money, respectively. Since subjects must quickly respond on the MID task to receive reward, activations during reward prediction in studies utilizing the MID task may reflect the enhancement of preparatory motor responses based on value prediction, rather than a passive signal of predicted positive value per se.

Many fMRI studies utilizing the MID task have looked at activations associated with predicted value, which is associated with the time of waiting following cue (i.e. CS) presentation. An activation likelihood estimate meta-analysis (ALE: Laird et al., 2005) of 20 studies using the MID or other cued-response tasks reveals that activation is greater for the prediction of a potential monetary gain than for prediction of a neutral outcome (i.e. no predicted monetary reward or loss) in various regions, including the ventral striatum, thalamus, and insula (Knutson & Greer, 2008). Although not found in the meta-analysis, there is also consistently more activation for the prediction of a potential monetary gain than for prediction of a neutral outcome on the MID in another region of the striatum, the caudate (Bjork et al., 2004; Juckel et al., 2006; Knutson, Adams, Fong, & Hommer, 2001; Knutson, Bhanji, Cooney, Atlas, & Gotlib, 2008; Knutson et al., 2004; Knutson, Fong, Adams, Varner, & Hommer, 2001; Knutson, Fong, Bennett, Adams, & Hommer, 2003; Samanez-Larkin et al., 2007; Wrase et al., 2007). Together this suggests

that in humans, the ventral striatum, the dorsal striatum (i.e. caudate), the thalamus, and insula are all involved in the prediction of rewards.

Predicted Negative Value

Understanding how the brain codes for the prediction of stimuli with positive value may not tell us about how the brain codes for the prediction of stimuli with negative value. Positive and negative incentives might prime behavior in different directions, and the neural substrates of these different types of behaviors may be different. Prediction of items with positive value might generate both positive feelings and high arousal (i.e. positive activation) which encourages one to approach them. In contrast, prediction of items with negative value might generate both negative feelings and high arousal (i.e. negative activation) which encourages one to avoid them (Panksepp, Knutson, & Burgdorf, 2002). Self-report data has shown that the emotional states of positive and negative activation are orthogonal (Russell, 1980), suggesting that each of these states has different neural substrates. Because each type of incentive may be primarily linked with one of these states, the neural substrates of value prediction might differ depending upon on whether the predicted item is a positive or a negative incentive.

While some neural substrates of positive and negative value might be different, others might be the same. Some regions of the brain could respond similarly to predicted value across valence, which could have a role in increasing arousal and in general priming of physical processes towards action. Conversely, some brain regions could respond to both positive and negative value but in different directions; these regions

could be particularly important for economic decisions that involve comparing positive and negative incentives. To have a full understanding of how the brain codes for predicted value, one must understand both how the brain codes for items with positive value and for items with negative value.

Although dopamine neurons play a distinct role in the prediction of rewards, their role in the prediction of aversive stimuli is less clear. In contrast to the general responsiveness of dopamine neurons to rewards, different dopamine neurons show very different responses to predictions of aversive stimuli. In an electrophysiological study, Matsumoto and Hikosaka (2009b) found that some midbrain dopamine neurons were excited by a CS predicting an aversive stimulus, while others were inhibited by it. These effects depended on the probability of the prediction, becoming greater as the probability associated with the predicted item increased. While there were two very different patterns of dopamine neurons to predictions of aversive stimuli, these neurons tended to show only one pattern for the prediction of positive stimuli: excitation. This study suggests that some dopamine neurons code for value on a single scale; by increasing their firing when rewards are predicted and by decreasing it when aversive stimuli are predicted. These neurons could prime behavior towards approach of rewards. Other dopamine neurons, however, that are excited by predictions of both positive and negative stimuli may instead have a role in increasing arousal.

In contrast to the divergent responses of dopamine neurons to stimuli predicting negative value, neurons in a region of the lateral habenula have been shown to exhibit more homogenous responses to these stimuli. On tasks involving reward, habenula neurons are excited by conditioned stimuli which predict no reward, and on tasks

involving aversive stimuli, habenula neurons are excited by predictions of these aversive stimuli (Matsumoto & Hikosaka, 2009a). Like some dopamine neurons, habenula neurons appear to code for value on a single scale. Contrary to the activity patterns of these dopamine neurons, however, activity in lateral habenula neurons is greater for predictions of more negative stimuli, raising the possibility that neurons of the lateral habenula may prime behavior towards avoidance of items with negative value.

Researchers have utilized fMRI to look for the neural correlates of predicted negative value just as they have to look for the neural correlates of predicted positive value. In their ALE meta-analysis of activation patterns during the MID and other cued-response tasks, Knutson and Greer (2008) found greater activation for prediction of a potential monetary loss than for prediction of a neutral outcome in the dorsal striatum (both caudate and putamen), insula, and thalamus. Notably, activation in the caudate, insula, and thalamus was present in their analysis of predicted positive value as well, suggesting a role for all three of these regions in prediction of monetary gains and losses. The putamen could have a larger role in prediction of negative value than of positive value, given that it was only significant in the monetary loss prediction analysis.

One should, however, be careful in interpreting these meta-analytic results as regions involved in the prediction of negative value. Many of the studies in the analysis used the MID. As the goal on this task is to respond quickly enough to avoid a loss, activations associated with the anticipation period on loss trials might represent predictions of no loss rather than predictions of loss per se.

One region notably absent from Knutson and Greer's (2008) meta-analysis of prediction of monetary loss is the ventral striatum. In fact, in their meta-analysis there

was greater activation associated with the prediction of potential monetary gain than with the prediction of potential monetary loss in the ventral striatum, suggesting that this region has a larger role in the prediction of positive than of negative value. Larger responses to predictions of positive value could be due to the behavior of dopamine neurons, as dopamine neurons project to this region (Haber, 2003). Since more midbrain dopamine neurons respond to the predicted value of rewards than of aversive stimuli, this could lead to greater activation increases associated with the prediction of positive than of negative value in the ventral striatum. This idea is consistent with the suggestion of Knutson and Gibbs that the fMRI BOLD response associated with the prediction of reward in the ventral striatum is a function of the release of dopamine in this region (Knutson & Gibbs, 2007).

Predicted Expected Value

One approach for investigating which regions in the brain are involved in valuation of monetary incentives is to look for brain activity that scales with the expected value, (i.e. magnitude x probability) of a predicted reward or aversive event. I have already discussed that dopamine neurons respond to both components of the expected value of rewards. Here I will focus on research findings that have looked at responses in the human brain to expected value, or its two components (i.e. magnitude and probability).

Several fMRI studies have looked at how brain activation scales with the expected value of monetary rewards in humans, and have revealed a particularly important role for the striatum. During passive monetary reward tasks, activation in the striatum positively

scales with the expected value of a predicted reward (Tobler, Christopoulos, O'Doherty, Dolan, & Schultz, 2009; Tobler, O'Doherty, Dolan, & Schultz, 2007). Other evidence suggests that striatal activation may also bias selection of options with a high expected value. In choice tasks for monetary rewards available with various probabilities, there is more activation in the striatum on trials in which high expected value options are chosen than on trials where options with lower expected value are chosen (Christopoulos, Tobler, Bossaerts, Dolan, & Schultz, 2009; Hsu, Bhatt, Adolphs, Tranel, & Camerer, 2005; Tobler et al., 2009).

Tobler and colleagues (2009; 2007) have also shown that activation in areas of the lateral prefrontal cortex (LPFC) scales linearly with the expected value of a predicted reward on passive monetary reward tasks, but leave unclear whether a specific portion of the LPFC is particularly responsive to expected value. In one of the fMRI studies by Tobler et al. (2007), activation that positively scaled with expected value was seen in the DLPFC, but not in the ventrolateral prefrontal cortex (VLPFC). In contrast, in another fMRI study by Tobler and colleagues (2009), it was only seen in the VLPFC.

Activation in the striatum not only positively scales with the expected value of rewards, but positively scales with both of its components: magnitude and probability. Tobler and colleagues have repeatedly shown that the probability of receiving a monetary reward as predicted by a cue is positively correlated with activation in the striatum (Tobler, Christopoulos, O'Doherty, Dolan, & Schultz, 2008; Tobler et al., 2007). Findings from Knuston et al. (2001) reveal that activation in the striatum also responds to predicted reward magnitude. On the MID, there is greater activation in both the ventral and dorsal striatum, and in two other regions (i.e. thalamus and dorsal MPFC), for

prediction of larger than of smaller rewards. The increased activation seen in the striatum as the expected value of predicted rewards increase or as the components of expected value increase reveals that the striatum has an especially important role in predicting the value of rewards.

While research shows that activation in the striatum scales linearly with the probability of predicted reward on passive reward tasks, activation patterns in the DLPFC associated with the probability of predicted reward appear to be more complex. In an fMRI study, Tobler and colleagues (2008) looked for activation in the brain that showed increased activation as an inverse S-function of probability, to see if responses in certain regions of the brain responded to probability in a manner consistent with the probability weighting function of Prospect Theory (Kahneman & Tversky, 1979). Compared to a linear function, this function overweights low probabilities and underweights high probabilities. Certain options (i.e. options with a probability of 0% or 100%), however, are not weighted differently than they would be by a linear function. They found that activation in the DLPFC responded to probability in a distorted manner, and was fit by the inverse S-function. In contrast, activation in the striatum increased linearly as probability increased.

No studies have investigated how activation in the human brain responds to the expected value of stimuli with negative values. A few studies, however, have examined how activation scales with the components of expected value (i.e. magnitude and probability). Activation in the striatum positively scales with the probability of a predicted aversive event (Bach, Seymour, & Dolan, 2009). This suggests that the

striatum has a similar role in signaling the probability of positive and negative stimuli, responding more as predicted stimuli are more likely to occur.

It is unresolved, however, whether regions of the brain respond to the predicted magnitude of positive and negative stimuli in a similar way. Knutson et al. (2001) found similar responses to the predicted magnitude of potential gains and losses on the MID in several regions. Activation in the dorsal striatum, thalamus, and dorsal MPFC were greater for both predicted gains and losses of larger magnitudes than of smaller magnitudes. In contrast, on a choice task between two risky options that could both lead to monetary gain or loss, Tom and colleagues (2007) found regions that had increased activation as the magnitude of a potential gain increased, but found no regions that had increased activation as the predicted magnitude of a potential loss increased. Interestingly, activation in the ventromedial prefrontal cortex (VMPFC) and striatum (i.e. ventral striatum and caudate) was shown to respond divergently to the predicted magnitudes of each type of incentive; activation increased in these regions as the predicted magnitude of a gain increased, but decreased as the predicted magnitude of a loss increased. This divergent response profile makes these regions good candidates for signaling economic value, because their patterns of response suggest they code for positive and negative predicted value on a single scale. The conflicting findings across studies, however, leave unresolved how the brain responds to the magnitude of negative stimuli.

Section Summary

This section has reviewed findings of how the brain responds to predictions of positive and negative value. This has revealed the importance of midbrain dopamine neurons, particularly in signaling positive value. It has also shown that a key region responsive to value is the striatum. This region responds to predictions of positive value, and has responses which directly scale with the expected value, magnitude, and probability of rewards. Activation in the striatum also responds to predictions of negative value and directly scales with the probability of predicted aversive events. Less is known, however, about how the brain responds to the magnitude and expected value of negatively valued stimuli.

There are some suggestions that the DLPFC is involved in the prediction of value as well, but it does not appear to have a central role in such processes. An interesting possibility suggested by one of the studies reviewed in this section is that probability might be coded differently in the DLPFC than in other regions responsive to value, such as in the striatum. It could respond differently to other components of prediction as well, such as delay. In contrast, there is little evidence that the PPC is responsive to basic predictions of value. The next section will focus specifically on the results of neuroeconomic studies of choice and will illustrate findings suggesting that the DLPFC and PPC do both have important roles during choice.

Subjective Value during Choice

Studying how the brain responds to predictions of rewarding or aversive stimuli in general can tell us much about the brain processes involved in economic decision making, but it has two important limitations. First, some regions that respond to value during choice situations may be different from those that respond to value during non-choice tasks. Second, not everyone values different incentives the same; individuals have different preferences. Because of this, value related responses in the brain that lead to choices might be better captured by looking for activity that scales with behavioral preferences (i.e. with subjective values) than with general properties of stimuli. Activity correlated with an individual's subjective valuations of items in risky choice or intertemporal choice tasks would be expected to directly lead to the risky or intertemporal choices they make.

Several fMRI studies have taken this approach by looking for activity that scales with subjective value on intertemporal choice or risky choice tasks involving positive incentives. To capture how each individual subjectively values items on these tasks, investigators have fit models (e.g. hyperbolic models of discounting) to each individual's choices. The results of studies using this approach have revealed a remarkable similarity in the regions that are responsive to the subjective value of risky monetary gains and delayed monetary gains. In intertemporal choice tasks, it has been consistently shown that activation in the striatum, MPFC, and posterior cingulate is positively correlated with the subjective value of delayed rewards (Kable & Glimcher, 2007, 2010; Peters &

Buchel, 2009, 2010). In risky choice tasks as well, several studies have shown that activation in these regions scales with the subjective value of risky rewards (Levy, Snell, Nelson, Rustichini, & Glimcher, 2010; Peters & Buchel, 2009). Together with the findings in the last section, this reveals that the striatum is both responsive to prediction of rewards in general and to the subjective value of rewards on choice tasks.

There is some evidence that the MPFC and posterior cingulate also play a role in signaling the subjective value of negative incentives during choice. FitzGerald and colleagues (2009) found that activation in both of these regions was associated with the difference in subjective value between two positive incentives and between two negative incentives on a basic choice task that did not involve intertemporal or risk components. However, it is unclear whether the MPFC and posterior cingulate would respond to the subjective value of loss options on intertemporal or risky choice tasks, as no prior studies have investigated how the brain responds to the subjective value of negative incentives during intertemporal or risky choice.

What about the DLPFC and PPC? There are some clues that the DLPFC may play a role in valuation on intertemporal and risky choice tasks. Research has shown that activation in the DLPFC scales with subjective value in a similar manner as does activation in the MPFC (i.e. VMPFC) in tasks that involve bidding for positive or negative incentives (Plassmann, O'Doherty, & Rangel, 2007; Plassmann, O'Doherty, & Rangel, 2010), and on a basic choice task (Hare, Camerer, & Rangel, 2009). Furthermore disruption of the DLPFC with repetitive Transcranial Magnetic Stimulation (rTMS) affects subjective valuation as measured on a bidding task (Camus et al., 2009). It could be that activation associated with subjective valuation in the DLPFC has not adequately

been captured by the models used to capture subjective value in past risky and intertemporal choice studies. However, another possibility is that the DLPFC is not involved in valuation on these tasks. Both the DLPFC and PPC could be involved in other functions during intertemporal and risky choice.

Intertemporal Choice

So far, I have argued how studying regions responsive to subjective value can inform us about what choices people make. While it should, there are other approaches to studying the neural mechanisms of choice. For example, investigators can look for activity that is associated with different types of choices or can see how choices are different when the functions of a brain region are disrupted. Studies using these approaches have revealed the importance of the DLPFC and PPC in intertemporal choices, and the DLPFC in choices involving risk.

The first clue that the DLPFC and PPC have an important role in intertemporal choice comes from an fMRI study by McClure and colleagues (McClure, Laibson, Loewenstein, & Cohen, 2004). They revealed that two sets of brain regions, which the authors labeled Beta and Delta regions, had different patterns of activation in intertemporal choices between smaller sooner monetary rewards and larger more delayed monetary rewards. The terms for these regions were taken from a theory of delay discounting that contains two components– a Beta variable that captures immediacy bias in choice and a Delta variable that discounts at a constant rate as a function of time (Phelps & Pollak, 1968). McClure and colleagues identified Beta regions as regions that were preferentially activated for intertemporal choices in which one of the two choice

options was immediate. These regions included the ventral striatum, MPFC, and posterior cingulate cortex. Notably, these are the same regions that have consistently shown activation patterns that scale with the subjective value of options in intertemporal choice tasks. The Delta regions, however, were similarly activated on intertemporal choices with and without an immediate option, and included the right DLPFC and bilateral PPC. As these regions are not typically associated with subjective valuation of options in intertemporal choice tasks, it is unclear what the functions of these regions were during choice. They did appear to be involved in the choice process itself, as they were preferentially activated by more difficult choices.

Using a similar design as McClure and colleagues (2004), several studies have replicated the effects, both with intertemporal choice for monetary rewards and for juice rewards (McClure, Ericson, Laibson, Loewenstein, & Cohen, 2007; Xu, Liang, Wang, Li, & Jiang, 2009). As in the original study by McClure et al. (2004), these studies found that the same two sets of regions were differentially activated depending on whether a choice was immediate. These consistent effects indicate that the DLPFC and PPC both have an important role in intertemporal choices, and suggest it is different from that of the regions that have typically been responsive to the subjective values of items during intertemporal choice.

Importantly, fMRI data suggests that the DLPFC and PPC bias intertemporal choice behavior differently than do the ventral striatum, MPFC, and posterior cingulate. The DLPFC and PPC may bias choice towards options of better long-term value, as there is greater activation in these regions when individuals choose larger more delayed rewards in intertemporal choice tasks (Weber & Huettel, 2008; Wittmann, Leland, &

Paulus, 2007). In contrast, the ventral striatum, MPFC, and posterior cingulate may bias choice towards options of better immediate value, as activation in this set of regions predicts a greater chance of choosing a small immediate over a larger delayed reward (McClure et al., 2007). Similarly, there is greater activation in the striatum for choice of small immediate than for choice of larger delayed options (Wittmann, Lovero, Lane, & Paulus, 2010). The relative activity in these two sets of regions as well, may have an important role in choice. McClure et al. (2004) observed that when subjects chose the larger later reward rather than the smaller immediate reward, there was more activation in the Delta regions (including the DLPFC and PPC) than in the Beta regions (i.e. the ventral striatum, MPFC, and posterior cingulate).

If the DLPFC and PPC bias intertemporal choices towards choice of the better long-term option, then it might be predicted that individuals with lesions to these regions would be more likely to choose immediate rewards in an intertemporal choice task than healthy controls. Fellows & Farah (2005), however, did not find any differences between the intertemporal choice behavior of individuals with lesions to either the left or right lateral PFC and that of individuals without lesions. There were some limitations to the study, however, that make it hard to assess the findings. Lesions to different individuals largely did not overlap; some individuals had lesions to the left PFC while others had them to the right, and different individuals had lesions to different regions of the lateral PFC. These differences make it possible that many of these individuals did not have lesions to the DLPFC regions important for intertemporal choices. Alternatively, effects were not seen because of limitations of the lesion approach; following brain damage, intact brain areas can sometimes compensate for the functions of the damaged ones.

An alternative approach to see whether or not a brain region has a necessary role in intertemporal choice behavior is to use low frequency rTMS to temporarily suppress functioning in a brain region. This method has several important advantages over the lesion approach. With rTMS, one can target the same region in different individuals, and the transient effects of the stimulation make it unlikely that a different brain region will be able to compensate for the disrupted functions of the targeted region.

Recently, using low frequency rTMS, Figner and colleagues (2010) were able to demonstrate that the DLPFC has an essential role in helping individuals make choices for better long-term, but worse immediate rewards. Following disruption of the left DLPFC, but not of the right DLPFC, individuals were more likely to choose a small immediate reward (and less likely to choose a larger delayed reward) on an intertemporal choice task than were individuals who had sham stimulation.

There were several important patterns of the rTMS effects. First, rTMS did not affect intertemporal choices that involved only delayed reward options; it only affected choices that involved both an immediate and a delayed option. This suggests that DLPFC functions may only be needed to choose larger delayed rewards when the other option is immediate. Second, the effects of rTMS were greatest on difficult trials (i.e. where there was an intermediate difference in monetary value between both options). Individuals were more likely to choose the immediate option following real rTMS administered to the left DLPFC on these trials, suggesting that the DLPFC has a greater role in helping individuals choose delayed rewards when the two options are close in subjective value. However, as stimuli were not matched for subjective value, it is unclear whether the options on the "difficult trials" were closer in subjective value than the

options on other trials. Third, rTMS affected choice but did not affect valuation of the items on the choice task, suggesting that the DLPFC does not bias choices because it evaluates the desirability of items differently than do other regions. However, this last suggestion is only tentative, as the valuation task used by Figner et al. had few trials and thus may have had low power to find effects.

The study by Figner and colleagues has revealed that the DLPFC has a critical function in intertemporal choices. But it leaves several unanswered questions. One question is whether only the left DLPFC is necessary for making better long-term choices. As only disruption of the left DLPFC affected choice in their study, one might conclude that the right DLPFC does not help individuals make such choices. However, other studies have found activation associated with intertemporal choices in the right DLPFC or in nearby regions of the right LPFC (McClure et al., 2007; McClure et al., 2004; Xu et al., 2009), suggesting that both sides may have a role in helping individuals make better long-term choices. It is possible that Figner and colleagues were not able to find an effect following disruption of the right DLPFC due to a lack of statistical power or due to specific components of the intertemporal choice task used.

An even more critical question that has currently been unaddressed is whether disruption of the PPC would similarly affect choice. As reviewed in this section, patterns of activation in the PPC and DLPFC are similar in intertemporal choice tasks (McClure et al., 2007; McClure et al., 2004; Xu et al., 2009), and findings show that activation in the PPC predicts choice of larger delayed rewards (Wittmann et al., 2007). This suggests that disruption of the PPC would also lead to greater choice of small immediate over larger delayed rewards, and that both the PPC and DLPFC may perform similar functions

during intertemporal choice. But as no studies have looked at how intertemporal choices are affected by neuromodulation or brain lesions to the PPC, this is only a prediction. Disruption of both regions could affect choice in a similar direction (e.g. towards greater choice of better long-term rewards), but for different reasons. The findings discussed implicating the DLPFC, but not the PPC in valuation, suggest that the DLPFC is more likely than the PPC to be involved in valuation processes during intertemporal choice.

Another critical question raised is whether disruption of the DLPFC or PPC would have similar effects in intertemporal choices involving losses. Delay discounting of gains and losses are positively correlated (Murphy et al., 2001) and are fit by the same type of hyperbolic function (Baker et al., 2003; Kirby & Marakovic, 1995; Murphy et al., 2001; Rachlin et al., 1991), suggesting that some processes in the brain might optimize choices of immediate vs. long-term value in both domains. However, the discussion of valuation processes revealed that brain responses to predictions of positive and negative value are not identical; while some regions have increased activation to the prediction of both positive and negative incentives (e.g. dorsal striatum), others are more likely to have increased activation for prediction of incentives of one domain than the other (e.g. greater responses to gain in the ventral striatum). This raises the possibility that the DLPFC and PPC might only have a critical role in choices involving gains.

The results of an fMRI study by Xu and colleagues (2009) suggest that the functions of the DLPFC and PPC may be similar in both intertemporal choices involving gains and those involving losses. Like McClure and colleagues (2004), Xu et al. analyzed their fMRI data to look for regions that preferentially responded on trials with an immediate option (Beta regions) and for regions that responded similarly on all trials

(Delta regions). Importantly, they did this separately for intertemporal choices involving gains and those involving losses. This analysis revealed similarities between patterns of activation for both types of choices. Although not all regions were identical across analyses, for both types of intertemporal choices, the Delta regions included the DLPFC and PPC, and the Beta regions included the MPFC and posterior cingulate.

The study by Xu and colleagues (2009) is only suggestive that the DLPFC and PPC have a similar role in intertemporal choices for gains and losses. Currently, there are no findings linking the DLPFC or PPC to specific choices on intertemporal choice tasks involving losses. If the functions of these regions help one select an option that has a better long-term value when another option has a better immediate value, one might expect disruption of these regions to lead to greater selection of larger delayed over smaller immediate losses.

Choice Involving Risk

Investigators have also looked for activation associated with different choice patterns in risky choice tasks, particularly activation that is associated with choosing more risky options (i.e. less probable) but potentially more rewarding options and activation that is associated with choosing smaller less risky (i.e. more certain) options. Importantly, this has suggested that the DLPFC has a key role in risky choice tasks.

A study by Matthews and colleagues has shown that choosing less risky rewards is associated with activation in the DLPFC (Matthews, Simmons, Lane, & Paulus, 2004). They found that there is more activation in the left DLPFC before selection of a certain monetary reward than before selection of a larger risky monetary reward. In contrast,

activation in the striatum, MPFC, and posterior cingulate has been associated with choosing riskier reward options (i.e. less probable) on choice tasks (Ernst et al., 2004; Weber & Huettel, 2008). Notably these are the same regions that have been implicated in subjective valuation of risky rewards (Levy et al., 2010; Peters & Buchel, 2009). One of these regions in particular, the striatum, appears to have a key role in biasing selection of risky options. Activation in the striatum, particularly in the ventral striatum, is not only associated with making risky choices, but has repeatedly been shown to predict whether or not a risky choice will be made. Matthews and colleagues (2004) found that there was more activation in the MPFC, ventral striatum, and caudate before selection of a risky option which could lead to a large gain or loss than before selection of a certain option for a small monetary gain. Similarly, Kuhnen and Knutson (2005) found that activation in the ventral striatum increased the chances of choosing a risky monetary option after having chosen a safe option on a financial decision making task. The ventral striatum plays a role in biasing selection of risky options in rodents as well. Cardinal and Howes (2005) found that rats with lesions to part of the nucleus accumbens chose more small certain reward options and fewer potentially larger risky reward options than did control rats.

Together these findings suggest that the DLPFC has a different role in choices involving risk than do the MPFC, striatum, and posterior cingulate. Specifically, the latter regions appear to bias behavior towards greater risk seeking, while the DLPFC appears to bias behavior towards greater risk aversion. However, the fMRI study reviewed implicating the DLPFC in such processes is only suggestive of such a role for

the DLPFC, since results from fMRI cannot be used to infer whether a brain region has an essential role in a behavior.

Studies using neuromodulation methods, however, have revealed that the DLPFC, particularly the right DLPFC, does have an essential role in helping individuals make choices for less risky options. Knoch and colleagues (2006) found that disruption of the right DLPFC with low frequency rTMS led to greater selection of riskier options than did sham stimulation on Roger's Risk Task (Rogers et al., 1999). Following disruption of the right DLPFC, individuals were more likely to choose the less probable, but potentially more rewarding riskier options (and were less likely to choose the safer options). Fecteau and colleagues (2007) have also shown that the right DLPFC has an important role in helping individuals avoid risk. They found that increasing activity in the right relative to the left DLPFC with transcranial Direct Current Stimulation (tDCS) increased choice of the riskier options on Roger's Risk Task. Together, these findings suggest that the right DLPFC biases individuals towards choosing less risky options. One should not, however, rule out a similar role for the left DLPFC, because some evidence suggests that both sides of the DLPFC may play a similar role in risky choice. One tDCS study found similar changes in preference on Roger's Risk Task both after increasing activity in the left relative to the right DLPFC and after increasing activity in the right relative to the left DLPFC (Boggio et al., 2010).

These neuromodulation studies have revealed that the DLPFC helps individuals avoid taking risks, but leave several critical questions unanswered. First, it is unclear whether the DLPFC biases behavior towards risk aversion for all levels of risk (i.e. all probabilities of reward), or only for some levels of risk. It may be that the DLPFC

becomes increasingly important in preventing choice of risky rewards as the probability of obtaining them decreases. Or the relationship between risk level and choice could be more complex. Either way, one might suppose that the DLPFC biases behavior differently depending on the risk level, given that activation in the DLPFC has been shown to scale differently with the probability of rewards than does activation in a region that biases behavior towards choice of risky options, the striatum (Tobler et al., 2008).

A second important question is whether the DLPFC helps individuals avoid taking risks for choices that do not involve the possibility of a loss. It could be that different brain regions bias behavior towards risk aversion in choices involving gains and in choices involving losses, and that the DLPFC is only involved in helping people avoid taking risks that can lead to losses. This possibility cannot be ruled out, since prior neuromodulation studies of risky choice have used Roger's Risk Task, which requires subjects to make a choice between two options that can both lead to a gain or a loss.

Section Summary

This section has revealed that the DLPFC has an important role in both intertemporal choice and choice involving risk. Disruption of the DLPFC increases choice of small immediate rewards during intertemporal choice and increases choice of riskier rewards during risky choice. This suggests that a similar process of the DLPFC might help individuals choose less risky monetary rewards and large rewards that are delayed. The PPC as well, appears to have an important role in intertemporal choice, similar to that of the DLPFC. An important question is what functions are performed by the DLPFC and PPC during economic choice. In order to better understand what they

may be, the next sections of this chapter will review research on the functions of these two regions of the brain.

Functions of the DLPFC

Divisions of the Prefrontal Cortex

The prefrontal cortex (PFC) is the most anterior part of the frontal cortex. In humans and other primates, it consists of all cortical regions anterior to the motor and premotor regions of the frontal lobe. It can be divided into three large subregions: lateral prefrontal cortex (LPFC), orbitofrontal cortex (OFC), and medial prefrontal cortex (MPFC). Based on cytoarchitectonic features, the PFC can be further divided into a number of smaller subregions. The most frequently used system for humans is that of Korbinian Brodmann (1909), and using this terminology the human PFC consists of Brodmann Areas (BA) 8-13, 32, 46, and 47 (Fuster, 1997).

Each subregion of the PFC is broad and is involved in many diverse brain functions. However, it has been generalized that while the OFC and MPFC are involved in emotional functions, the LPFC is involved in other, more cognitive functions (Fuster, 2001). Here I will focus on the role of the LPFC, with an emphasis on that of the dorsal section (i.e. the DLPFC).

The LPFC can be subdivided into two regions, the DLPFC lying along the middle and superior frontal gyrii, containing BA 9,46, and 8, and the VLPFC lying along the inferior frontal gyrus, made up of BA 45, 44, and the lateral aspect of BA 47 (Petrides, 2005). Rostral to both regions is BA 10, the frontopolar cortex, which contains lateral,

medial, and orbital aspects. This section will predominantly focus on the role of a specific portion of the DLPFC, that of the mid-DLPFC which contains BA 9 and 46 (or alternatively areas 9, 46, and 9/46 (Petrides, 1998b), following the PFC parcellation of Petrides and Pandya (Petrides & Pandya, 1994). Discussion is focused on this portion of the DLPFC because Experiments 1 and 2 of this dissertation both involve disruption of this area, given its known involvement in both intertemporal choice and choice involving risk (Fecteau et al., 2007; Figner et al., 2010; Knoch et al., 2006; Matthews et al., 2004; McClure et al., 2007; McClure et al., 2004; Xu et al., 2009) (see Figure 2 for location of regions).

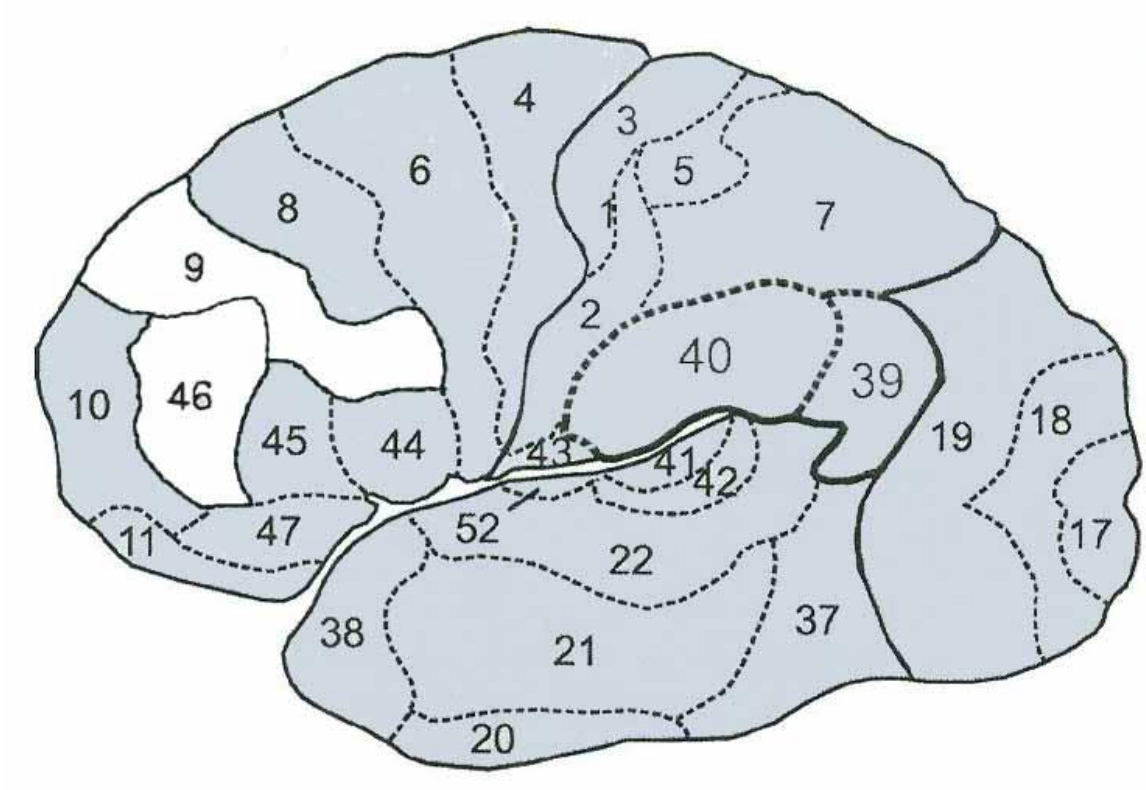

Figure 2: Location of the Mid-DLPFC. Figure shows Left Lateral view of the human brain with Brodmann areas (Brodmann, 1909) listed. Areas 9 and 46 of the Mid-DLPFC are highlighted in white. Figure modified from Caspers et al. (2006).

Theories of lateral PFC function with a focus on functions of the DLPFC

Even though the LPFC is involved in many functions, this section will largely focus on the role of this region in the functions of working memory and cognitive control, both which may be hypothesized to explain some of the contributions of the LPFC to economic decision making.

Working Memory

Working memory has been defined by Knight and Stuss (2002, p. 577) as the "ability to maintain information over a delay and to manipulate the contents of this short-term memory storage system." A large amount of research implicates the DLPFC, and the LPFC more broadly, in working memory processes. Some neurons in the mid-DLPFC of the primate fire during the delay period of spatial delayed response tasks, some have firing that is sensitive to the duration of delay, and some are unresponsive during this delay period when there is no memory load (Goldman-Rakic, 1987). All of these findings suggest that activity in the DLPFC is involved in maintaining information online. Additionally, findings from lesion studies more directly show that the DLPFC is important for working memory, since monkeys with lesions to this brain region perform poorly on working memory tasks (Butters & Pandya, 1969; Butters, Pandya, Sanders, & Dye, 1971; Goldman-Rakic, 1987). Other regions of the LPFC, however, are also implicated in working memory. Primate neurophysiological studies have demonstrated that neurons throughout the LPFC fire during the delay period of delayed response tasks (Bodner, Kroger, & Fuster, 1996; Funahashi, Bruce, & Goldman-Rakic, 1989; Fuster & Alexander, 1971; Miller, Erickson, & Desimone, 1996).

Research with humans also indicates that the DLPFC is involved in working memory. A meta-analysis of a large number of human Positron Emission Tomography (PET) and fMRI studies found bilateral activation in both the DLPFC (BA 9, 46) and VLPFC (BA 44) associated with working memory processes (Cabeza & Nyberg, 2000). Additionally, disruption of the DLPFC with rTMS has repeatedly been shown to impair performance on working memory tasks (Mottaghy, Doring, Muller-Gartner, Topper, & Krause, 2002; Mottaghy, Gangitano, Sparing, Krause, & Pascual-Leone, 2002).

Several different theories have been proposed to account for the role of different parts of the LPFC in working memory. Goldman-Rakic's strong view of working memory holds that the lateral PFC is organized along a dorsal-ventral division with different regions of the LPFC involved in similar working memory processes, but involved in working memory processes for content of different domains (Goldman-Rakic, 1987, 1998). According to this view the mid-DLPFC is involved in spatial working memory, and the mid-VLPFC is involved in non-spatial object working memory. Although some non-human primate lesion and electrophysiological studies show such regional distinctions for visuospatial working memory (Levy & Goldman-Rakic, 2000), a meta-analysis of human functional neuroimaging imaging studies of working memory does not support this distinction (Wager & Smith, 2003). Evidence that many prefrontal neurons show delay-period activity for both object and spatial information are also non-supportive (Rao, Rainer, & Miller, 1997).

A different view is that of Petrides, whose theory holds that the DLPFC and VLPFC are involved in different working memory processes, rather than in different working memory content domains (Petrides, 1998a, 2005). This view holds that the mid-

DLPFC (areas 46 and 9) is involved in monitoring of stimuli in working memory (i.e. keeping track of the current status of stimuli in relation to other stimuli in working memory (Champod & Petrides, 2007)), and in working memory processes that are thought to require monitoring, such as manipulation (i.e. rearrangement) of stimuli in working memory. In contrast, this view holds that the mid-VLPFC is involved in first-order executive processes involved in explicit encoding and retrieval. Several functional imaging studies support these claims (Champod & Petrides, 2007, 2010; Doyon, Owen, Petrides, Sziklas, & Evans, 1996; Petrides, Alivisatos, & Evans, 1995).

Some evidence suggests that the DLPFC is only necessary for working memory tasks when distracters are present. Monkeys with lesions of the DLPFC perform normally on delayed response tasks if visual distractions are absent during the delay period (i.e. when the lights of the room are turned off), but perform poorly when distractions are present (Malmo, 1942). This has lead to the suggestion that the DLPFC is involved in executive aspects (i.e. executive attention) of working memory, rather than in basic maintenance (Kane & Engle, 2002). Executive attention, or keeping memory representations active in the presence of interference, can be considered to rely on the function of cognitive control, which will be discussed shortly.

Meta-analytic findings also show a role for the DLPFC beyond the basic storage of information in working memory. A meta-analysis of human functional neuroimaging studies reveals that there is more activation in the bilateral mid-DLPFC and bilateral frontal pole for tasks that engage working memory storage and executive processes (i.e. manipulation, temporal order memory, or updating) than for those that only engage working memory storage (Wager, Jonides, & Reading, 2004). Of note, these findings are

consistent with the view of Petrides (1998a; 2005), since working memory tasks that engage executive processes can be said to require monitoring of information in working memory.

Cognitive Control

Many investigators have delineated a more executive role for the LPFC that goes beyond its role in working memory. One influential theory of PFC function is that it is important for cognitive control, which can be defined as top-down processing used to guide behavior towards internal intentions or states (Miller & Cohen, 2001). More specifically, according to this theory, the PFC is said to be involved in actively maintaining goals, attentional templates, and/or rules. These PFC activity patterns are said to provide bias signals throughout much of the brain so that an individual can perform a task a particular way. Cognitive control can be used to help individuals select weaker, task relevant responses over more frequently used ones that have more established neural pathways. This control is important for many processes including selective attention, behavioral inhibition, working memory, rule-based behavior, and goal-directed behavior.

Evidence for the cognitive control theory of PFC functions comes from human and monkey lesion studies showing that those with prefrontal cortex damage have problems with perseveration and distractibility, and perform poorly on tasks that require the ability to select a weaker, task relevant response over a more habitual, prepotent response (Chao & Knight, 1997; Milner, 1963; Mishkin, 1964; Perret, 1974; Vendrell et al., 1995). Human fMRI evidence suggests that the mid-DLPFC in particular appears to be involved in the implementation of cognitive control, as there is more activation in this

region during the instruction period of trials when subjects are told to make a less habitual response than when they are told to make a more habitual response (MacDonald, Cohen, Stenger, & Carter, 2000). Additional evidence that the LPFC may be important in cognitive control comes from findings that many LPFC neurons encode the current behavioral context, by exhibiting response patterns that are dependent on task rules (Asaad, Rainer, & Miller, 2000; Hoshi, Shima, & Tanji, 1998; White & Wise, 1999). By encoding task rules, these neurons may help bias brain activity to help an individual perform behavior consistent with those rules.

Several different theories hold that the LPFC is functionally organized in a rostral-caudal manner, with more rostral areas involved in more abstract control processes than more caudal areas (see Badre & D'Esposito, 2009; Fuster, 2002; O'Reilly, 2010; Petrides, 2005). Evidence from fMRI studies is supportive of this, revealing that the mid-DLPFC is involved in more abstract control than more posterior DLPFC regions (BA 8) and premotor regions (Badre & D'Esposito, 2007; Koechlin, Ody, & Kouneiher, 2003). Two theories of rostral-caudal function, the Cascade model of Koechlin and Summerfield (2007) and the policy abstraction model of Badre and D'Esposito (2007), also hold that organization in the LPFC is hierarchical, such that it is more common for rostral regions to influence more caudal regions, than for the reverse. The findings of an fMRI study by Koechlin and colleagues (2003) do suggest that the rostral-caudal axis of LPFC function is hierarchically organized, with rostral frontal regions able to account for variance in activation in more caudal regions, but not vice versa.

Dorsal vs. ventral divisions in cognitive control have also been postulated. Recently O'Reilly has proposed that the DLPFC provides cognitive control for the

brain's posterior dorsal "how" pathway, and the VLPFC provides such control for the posterior ventral "what" (i.e. object and verbal information) pathway (Goodale & Milner, 1992; O'Reilly, 2010). According to this framework, the DLPFC is postulated to provide cognitive control for transforming perception into action. This view builds off of Goldman-Rakic's (1987) content model of working memory, but in substituting "how" processing for "where" (i.e. spatial information) processing in the dorsal stream, it is also able to account for some of the process-related differences in the model of Petrides (2005). Consistent with O'Reilly's model is evidence showing that the DLPFC is involved in rule based response selection (Bunge, 2004).

Functions of the PPC

Divisions of the Posterior Parietal Cortex

The parietal lobe of the brain lies on the dorsal and lateral surface of the cerebral cortex posterior to the frontal lobe and superior to the temporal and occipital lobes. There are two major sections, the anterior section of the parietal lobe, which contains the somatosensory cortex, and the PPC. In humans, the PPC can be divided into three major sections. The superior parietal lobule (SPL) which contains BA areas 5 and 7, the inferior parietal lobule (IPL) made up of BA 40 (supramarginal gyrus) and 39 (angular gyrus), and the intraparietal sulcus (IPS) which separates the inferior and superior regions (Caminiti et al., 2010) (see Figure 3).

Comparison of data from human and monkey studies investigating the functions of the PPC are difficult, because the inferior parietal lobule has expanded a great deal in

humans (Hyvärinen, 1982). It has been suggested that the monkey inferior parietal lobule may actually be homologous to the human superior parietal lobule (Culham & Kanwisher, 2001). In monkeys, the superior parietal lobule consists of BA 5, and the inferior parietal lobule consists of BA 7; as in humans the IPS separates the regions. BA 39 and 40, however, were not identified in the non-human primate brain. Others, however, have parcellated the human and monkey parietal lobes differently, creating similar homologous divisions in both species: PE in the superior PPC, and areas PG and PF in the inferior PPC (Von Bonin & Bailey, 1947; Von Economo, 1929). In the human, PE roughly corresponds to BA 5 and 7, and PG and PF roughly correspond to BA 40 and 39, respectively (Hyvärinen, 1982).

Given the prevalence of the listing of Brodmann areas in the human functional imaging literature, I will utilize the labels provided by Brodmann in this section where available. Since there are discrepancies in which regions are considered homologous across species, this section will focus more on the functions of the human PPC than that of the monkey. This discussion will largely focus on the roles of BA 7 (excluding the medial aspect), 40, and the IPS, since Experiment 1 of this dissertation involves disruption of this part of the PPC, given its involvement in intertemporal choice (McClure et al., 2007; McClure et al., 2004; Xu et al., 2009).

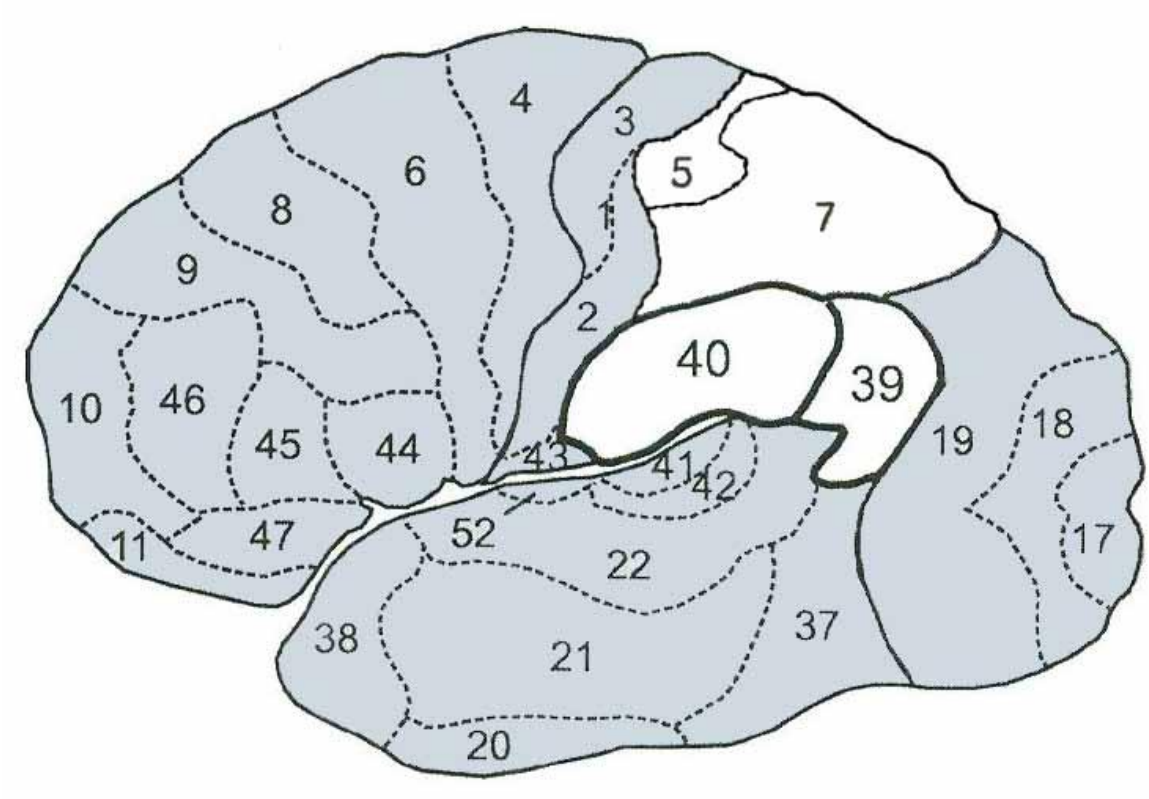

Figure 3: Location of the PPC. Figure shows Left Lateral view of the human brain with Brodmann areas (Brodmann, 1909) listed. Areas 5, 7, 39, and 40 of the PPC are highlighted in white. Figure modified from Caspers et al. (2006).

Roles of the PPC

PPC functions

One classic view of the PPC focuses on its role as part of the dorsal visual processing stream, which is involved in spatial perception (Ungerleider & Mishkin, 1982). The PPC is, however, neither structurally nor functionally monolithic. Lesions to the SPL produce different deficits than those to the IPL (Goldenberg & Karnath, 2006; Mort et al., 2003; Perenin & Vighetto, 1988; Vallar & Perani, 1987). These differences have led to the proposal that the SPL is involved in controlling action while the IPL is important in understanding actions and spatial perception (Rizzolatti & Matelli, 2003).

As will be discussed, the PPC plays an important role in two key functions that may be especially relevant to economic decision making: working memory and top-down attention.

Like the DLPFC, the PPC has a role in working memory, which is impaired following both lesions to the SPL or IPL and following disruption of the PPC with rTMS (Baldo & Dronkers, 2006; Koenigs, Barbey, Postle, & Grafman, 2009; Mottaghy, Doring et al., 2002). Similarly, meta-analytic findings of human functional imaging data indicate that regions of both the SPL and IPL (BA 7 and 40) are consistently implicated in working memory tasks, and that activations tend to be left-lateralized for verbal/numeric working memory tasks (Cabeza & Nyberg, 2000). Other meta-analytic findings reveal that in the SPL (BA 7) there is a preference for activation during simple storage in working memory of spatial, rather than object or verbal information (Wager & Smith, 2003). Furthermore, IPS activity is positively correlated with working memory load during both manipulation and maintenance tasks (Champod & Petrides, 2010; Todd & Marois, 2004).

Executive processing tasks in working memory (i.e. tasks involving updating, order, or manipulation in working memory) are also consistently associated with more activation in BA 7 and BA 40 than are simple storage tasks, and BA 7 is activated for executive working memory involving different storage types (Wager et al., 2004; Wager & Smith, 2003). However, unlike in the DLPFC where BOLD signal is primarily associated with working memory monitoring, rather than with manipulation, in the IPS, BOLD signal is more associated with working memory manipulation than with monitoring itself (Champod & Petrides, 2007, 2010).

A function related to working memory is attention, since information in working memory has also been attended to (Knudsen, 2007). Meta-analytic findings of human neuroimaging data show consistent activations during attention tasks in the PPC (Cabeza & Nyberg, 2000). Furthermore, identical regions of the bilateral IPS are activated in both attention and working memory tasks (LaBar, Gitelman, Parrish, & Mesulam, 1999).

Attention can be segmented into different sub-processes, and one of these, orienting, has particularly been associated with the PPC (Raz & Buhle, 2006). According to Raz and Buhle (2006, p. 372), orienting or attentional selection is the "ability to select specific information from among multiple sensory stimuli". It can be both top-down (i.e. goal directed) and bottom-up (i.e. driven by salient or unexpected stimuli in the environment). Top down and bottom up selection have been associated with different parts of the parietal lobe. The more dorsal aspects of the PPC, such as the SPL and IPS, are associated with top-down attention, while an inferior region of the IPL, the right temporoparietal cortex has a role in bottom-up attention (Corbetta, Kincade, Ollinger, McAvoy, & Shulman, 2000; Corbetta & Shulman, 2002; Friedrich, Egly, Rafal, & Beck, 1998). In humans, the left IPS in particular has an important role in top-down attention. Several studies have shown that disruption of activity in this region with TMS impairs performance on an identification task when highly salient distracters must be ignored and less salient items have to be selected (Mevorach, Hodsoll, Allen, Shalev, & Humphreys, 2010; Mevorach, Humphreys, & Shalev, 2006).

Numerical Processing and the IPS

The PPC is also involved in another function that is relevant to economic decision making: numerical processing. A meta-analysis of fMRI studies investigating number

processing reveals that the bilateral horizontal segment of the IPS (HIPS) is consistently activated during tasks involving numerical quantity manipulations, such as numerical comparisons, addition, and subtraction (Dehaene, Piazza, Pinel, & Cohen, 2003). Similarly, lesions to this region lead to deficits in performing calculations (Takayama, Sugishita, Akiguchi, & Kimura, 1994). These findings have led to the proposal by Deheane and colleagues (2003) that the HIPS represents the brain's core quantity system.

TMS studies particularly implicate the left IPS in the ability to perform numerical comparisons. Disruption of the left IPS using TMS or rTMS leads to impaired performance on numerical comparison tasks as demonstrated by increased reaction times (Andres, Seron, & Olivier, 2005; Cappelletti, Barth, Fregni, Spelke, & Pascual-Leone, 2007), and impairments occur both when quantities are expressed symbolically or non-symbolically (Cappelletti et al., 2007). Additionally, deficits following rTMS are greater when comparing numbers with similar magnitudes than with dissimilar magnitudes, consistent with findings that activity in IPS is sensitive to distances from specific quantities (Cappelletti et al., 2007; Nieder & Miller, 2004; Piazza, Izard, Pinel, Le Bihan, & Dehaene, 2004). Reaction time deficits on numerical comparison tasks have also been observed following rTMS disruption of the left IPL (BA 39 or 40), but these deficits occur less consistently than do those following disruption of the left IPS (Cappelletti et al., 2007; Gobel, Walsh, & Rushworth, 2001; Sandrini, Rossini, & Miniussi, 2004). In contrast to the effects seen following disruption of left PPC regions, disruption of the IPS or the adjacent IPL in the right hemisphere does not lead to significant impairments on numerical comparison tasks (Andres et al., 2005; Cappelletti et al., 2007; Gobel et al., 2001; Sandrini et al., 2004), and has even led to increased performance in one study

(Cappelletti et al., 2007). This reveals that the right PPC is less important for numerical comparisons than is the left PPC.

Connectivity between the DLPFC and PPC

The parietal and frontal cortices are largely interconnected through the superior longitudinal fasciculus (SLF). The connections are bidirectional; that is they contain both neurons projecting from the parietal to the frontal lobe, and neurons projecting from the frontal to the parietal lobe (Petrides & Pandya, 2002). The SLF can be parsed into 3 subdivisions, SLF I, II, and III, which differ on where the tracts originate. Note that the classification of these regions is based on cytoarchitectonic designations of Petrides and colleagues in monkeys, rather than those of Brodmann in humans (Pandya & Seltzer, 1982; Petrides & Pandya, 1994). Through two of these tracts, the mid-DLPFC is interconnected to two areas of the PPC, PG and PF, which both lie in monkey IPL.

These large interconnections may help explain why the PPC and DLPFC have involvement in similar functions. Both regions are involved in working memory functions and cognitive control functions (i.e. PPC in top-down attention, and DLPFC in cognitive control more generally). Furthermore, both the PPC and LPFC have neurons responsive to quantities (Nieder & Miller, 2004). Importantly, the co-involvement of both regions in these functions may help explain why both the PPC and DLPFC show similar activation patterns on fMRI studies of intertemporal choice. The next section of this chapter will offer some potential explanations of what functions these two regions of the brain may be engaged in during intertemporal choice and choice involving risk.

Potential roles of the DLPFC and PPC in Intertemporal Choice and Choice involving Risk

Valuation

In choice situations, it is reasonable to assume that individuals tend to choose options that they subjectively value more highly than other options. As previously discussed, a network of brain regions including the striatum, MPFC, and posterior cingulate are all associated with subjective valuation of rewards on intertemporal and risky choice tasks (Kable & Glimcher, 2007, 2010; Levy et al., 2010; Peters & Buchel, 2009, 2010). Less is known about how the brain codes for the subjective value of negative incentives on choice tasks, but some evidence implicates the MPFC and posterior cingulate in such a role (FitzGerald et al., 2009).

Prior fMRI research indicates that activation in two of the regions responsive to subjective valuation (i.e. MPFC and striatum) is associated with both choice of small immediate over larger delayed rewards (McClure et al., 2007; Wittmann et al., 2010) and with choice of potentially larger risky over smaller less risky rewards (Kuhnen & Knutson, 2005; Matthews et al., 2004). An important question remains. Why is there greater activation in regions that respond to subjective value only for certain types of choices? If people choose what they subjectively value highly, then one would expect that brain activation that positively scales with subjective value would be greater for the chosen option, regardless of its risk level or time of receipt.

One possible answer to this puzzle is the proposal of McClure and colleagues (2004) that there are two reward evaluation systems in the brain - both an "impulsive system" containing the ventral striatum, MPFC, and posterior cingulate (i.e. Beta regions)

that places a large value on rewards that are available immediately and a more patient system containing both the DLPFC and PPC (i.e. Delta regions) that discounts the value of delayed rewards relative to immediate rewards less steeply. Similarly, there might be dual evaluation systems for incentives in other types of choices, such as for losses in intertemporal choices and for gains in choices involving risk. For all of these types of choices, the DLPFC and PPC could both be involved in one of these systems and could place a different value on incentives as a function of their attributes (e.g. delay or risk) than do other regions.

If the DLPFC and PPC respond to subjective value, then why haven't past intertemporal and risky choice studies found consistent patterns of activation related to subjective value in these regions? The lack of findings may be a result of the models used to capture brain activity. If activity in these regions scales to subjective value differently than does activity in other regions, one might not expect to find it with the same statistical models.

There are several reasons one could predict a role for the DLPFC in the evaluation of options on intertemporal and risky choice tasks. First, activation in the DLPFC is associated with the subjective value of items during basic choice (Hare et al., 2009) and during bidding tasks (Plassmann et al., 2007; Plassmann et al., 2010). Second, disruption of the DLPFC with rTMS affects valuation of items (Camus et al., 2009). Third, activation in the DLPFC scales with reward probability (Tobler et al., 2008). In contrast, there is not much reason to predict a role for the PPC in valuation processes during choice tasks, since it does not typically respond to predicted value or its components.

Manipulation of Values within working memory

In order to choose the most highly valued option on a risky choice or intertemporal choice task, a person may need to hold monetary values online in working memory, manipulate them to account for time or risk, and then compare them. Functions of brain regions directly involved in subjective valuation may automatically take into account the time and risk of options to determine value, but working memory could be used to transform the monetary value of options in a more conscious manner, such as by transforming values by an inflation factor or by computing expected values. It would be expected that such value manipulation processes would rely on the DLPFC and PPC since both of these regions are engaged during working memory manipulation tasks (Champod & Petrides, 2007, 2010; Wager et al., 2004).

Evidence suggests, however, that manipulation of values within working memory during intertemporal choice may not be dependent on the DLPFC. Disruption of the left DLPFC with low frequency rTMS leads to greater choice of smaller more immediate rewards on trials pitting an immediate against a delayed option, but not on trials pitting two delayed options against each other (Figner et al., 2010). This pattern of findings suggests that disruption of the DLPFC affects choices without impairing the ability to manipulate values in working memory, because trials with two delayed options may be expected to lead to a larger working memory load than trials with a delayed option and an immediate option. On trials with two delayed options, the monetary values of two options may both need to be manipulated to determine the present value of the options, whereas on trials with an immediate and a delayed option, only the monetary value of the delayed option should need to be manipulated to obtain the present value of the options.

Cognitive Control

The DLPFC and PPC may be involved in cognitive control operations during choice tasks, specifically on trials involving seductive or aversive immediate options. In intertemporal choice tasks, when there are options for immediate monetary gain or loss, the control processes of the DLPFC and PPC may help one select a larger delayed gain over a smaller immediate gain or help one select a smaller immediate loss over a larger delayed loss. In risky choice tasks for monetary gain, cognitive control processes of the DLPFC and PPC may help one select a smaller certain gain over a potentially larger risky gain. In both intertemporal and risky choice tasks involving positive incentives, cognitive control could be used to help individuals overcome prepotent urges to choose rewarding options that might have better immediate values than alternatives. Similarly, in tasks involving negative incentives, cognitive control could be used to help individuals overcome prepotent urges to avoid choosing aversive options that might have worse immediate values than alternatives.

Cognitive control could be used on choice tasks to bias behavior towards specific goals and rules. People may approach intertemporal choice tasks with the goal of obtaining the largest long-term gain, and may have specific rules that help them compare the values of different options. In intertemporal choice tasks, these goals and rules may be easy to follow except in situations where an option is immediate, because such options may lead to large activations in areas coding for subjective value that in turn influence choice. In these situations, cognitive control may be needed to help one overcome a prepotent urge to select the option that offers the best immediate value, and to choose in line with one's goals and rules. Similarly, in risky choice tasks, cognitive control might

be used to help individuals follow goals to avoid choosing options with excessive risk, or might be used to help individuals use rules to compare the different values (e.g. computations of expected value).

One might expect that the DLPFC would be particularly important in cognitive control on choice tasks, given the findings implicating the DLPFC (MacDonald et al., 2000), and the PFC more generally (Miller & Cohen, 2001), in cognitive control functions. However, the PPC could be involved as well. The PPC might be expected to be involved in the cognitive control function of top-down attention, since this is a key function of the region (Corbetta & Shulman, 2002). Top-down attention might be needed to focus on specific attributes of choice options (e.g. magnitude or probability) that are determined by goals, and might be especially needed when distracting options that offer immediate incentives are present.

A cognitive control account of the DLPFC and PPC could offer an alternative to the suggestion that choice always follows directly from subjective valuation, and could explain why activation in these regions scales with choice patterns but does not appear to scale with subjective value on the same types of tasks. This is because cognitive control might be able to affect choices without affecting subjective valuation.

Numerical Comparisons

Given its role in numerical comparison processes, the left PPC may be involved in comparing the numerical values of the two options and their components (i.e. magnitude, time, and probability of outcomes) on economic choice tasks (Andres et al., 2005; Cappelletti et al., 2007). This proposed function overlaps with that of working memory,

since individuals may hold numerical values online and manipulate them in working memory when they compare them.

Ultimately which specific numerical values are compared will depend on how an individual decides to manipulate the components of the choice options. An individual may not only compare the presented monetary, time, and probability values, but may compare transformations of them. Since both the DLPFC and PPC are implicated in manipulation of items in working memory (but principally the PPC), such numerical value manipulation processes could also rely on the DLPFC (Champod & Petrides, 2007, 2010; Wager et al., 2004).

Research Plan

Experiment 1 was designed to see if transiently disrupting either the DLPFC or PPC with low frequency rTMS would bias intertemporal choices involving monetary gains and intertemporal choices involving monetary losses towards options with a better immediate value but a worse long-term value. No prior studies have used neuromodulation to investigate the role of the PPC in intertemporal choices for gains or the role of either the DLPFC or PPC in intertemporal choices for losses. It was predicted that disruption of either the DLPFC or PPC would lead to greater selection of small immediate monetary rewards over larger delayed monetary rewards than would sham stimulation to the same regions. For choices involving losses, it was predicted that disruption of the DLPFC or PPC would lead to greater selection of larger delayed losses over smaller immediate losses than would sham stimulation. This was predicted for

choices involving losses because larger delayed losses have a better immediate value than smaller immediate losses. A secondary aim of this experiment was to see if greater choice of options with a better immediate value was associated with deficits in the ability to compare numerical values, since numerical comparison deficits could underlie difficulties in choosing options with the best long-term value. This was tested by administering a numerical comparison task following disruption of the DLPFC or PPC. A further question that was investigated was whether right vs. left hemispheric disruption of the DLPFC and PPC differentially affected intertemporal choice.

Experiment 2 was designed to see if transiently disrupting the DLPFC with low frequency rTMS would bias choices towards greater risk seeking. Subjects completed a risky choice task that involved choices between a certain option and a risky option that offered a potentially larger probabilistic reward. It was predicted that disruption of the DLPFC would lead to greater choice of risky options than would sham stimulation. Choices were limited to the gain domain (i.e. choices could lead to reward or no reward, but never a loss) to see if disruption of activity in the DLPFC would lead to greater risk seeking in choices that did not involve the possibility of loss, which has not previously been investigated. Additionally, the probability of the risky option was varied across trials to see if the effects of disrupting the DLPFC on risk preferences would vary with the risk level. As in Experiment 1, a further question was whether right vs. left hemispheric disruption of the DLPFC differentially affected choice.

Utilizing rTMS offered several advantages over other approaches such as fMRI. While fMRI can help establish that a particular brain region is involved in a behavior, it cannot be used to tell whether a brain region plays an essential role in that behavior.

rTMS can be used to infer causality, since if a behavior is abnormal when a particular brain region is disrupted, one can infer a role for the region in that particular behavior. rTMS also has advantages over studies of patients with brain lesions in that one can target the same region with rTMS across subjects, whereas lesions are rarely confined to the same specific region across individuals and following brain injury, reorganization of structure-function relationships occurs. In the context of Experiments 1 and 2, using rTMS allowed us to examine if disruption of activity in specific brain regions affected choice in healthy individuals.

It is currently unclear whether the tendencies to discount the value of different types of monetary incentives are all related to a common personality trait. Delay discounting of monetary gains is positively correlated with both delay discounting of losses and with probability discounting of gains (Mitchell, 1999; Murphy et al., 2001; Myerson et al., 2003; Richards et al., 1999), suggesting that these different types of discounting may have similar relationships with some trait. A number of studies have investigated the relationship between one trait, impulsivity, and delay discounting of monetary gains, and have found a positive relationship (Kirby et al., 1999; Madden et al., 1997; Mitchell et al., 2005; Petry, 2001, 2002; Reynolds, Richards et al., 2006). However, there has been little research on the relationship between impulsivity and other types of monetary discounting. Experiment 3 was designed to investigate the relationships between different facets of impulsivity and the tendencies to discount different types of monetary incentives (i.e. delayed gains, delayed losses, risky gains). A secondary aim was to examine relationships among different types of discounting in the

same subjects more directly, by examining whether the tendencies to discount different types of monetary incentives were correlated.

CHAPTER II

EXPERIMENT 1 – CHANGES IN INTERTEMPORAL CHOICE FOLLOWING DISRUPTION OF THE DLPFC AND THE PPC

Introduction

A wealth of intertemporal choice research shows that individuals discount the value of delayed monetary rewards relative to immediate monetary rewards (Frederick et al., 2004). All else being equal, individuals tend to prefer immediate over delayed rewards, and will often choose smaller immediate options over larger delayed alternatives. Individuals tend to discount the value of delayed monetary losses as well (Estle et al., 2006), often preferring larger delayed losses over smaller immediate ones. Nevertheless, people are also sometimes able to choose the option with the best long-term value.

Recently, it has been shown that functions performed by the DLPFC help individuals make intertemporal choices for gains with the best long-term value (Figner et al., 2010). Following disruption of the left DLPFC with low frequency rTMS, individuals are more likely to choose smaller immediate over larger delayed monetary rewards than they are following sham stimulation, especially on difficult trials in which the two options are close in value. Importantly, only intertemporal choices that contain an immediate option are affected, suggesting that DLPFC processes are only needed to make the better long-term choice when an enticing immediate option is present.

Previous research reveals that another brain region, the PPC, has similar patterns of brain activity to the DLPFC during intertemporal choice (McClure et al., 2007; McClure et al., 2004). Activation in both regions is greater on more difficult than on easy intertemporal choice trials (McClure et al., 2004). This suggests a role for both the DLPFC and PPC in the choice process itself, since brain areas that are involved in task relevant computations generally show increased activity as difficulty increases (Assmus, Marshall, Noth, Zilles, & Fink, 2005; Braver et al., 1997; Gould, Brown, Owen, ffytche, & Howard, 2003). Activity in both regions may help individuals make better long-term choices through cognitive control, as cognitive control functions have been linked to both the DLPFC and PPC (Cole & Schneider, 2007; Corbetta & Shulman, 2002; Mevorach et al., 2010; Miller & Cohen, 2001; Wager et al., 2004). Alternatively activity in both regions could help individuals make better long-term choices through valuation processes that ascribe greater value to delayed options than do the functions of other brain regions. Consistent with the proposal that PPC activity helps individuals make choices for the better long-term option are findings showing greater activation in the PPC when individuals choose larger delayed rewards over smaller immediate alternatives (Wittmann et al., 2007). Also consistent with this are findings that there is more activation in a set of brain regions that includes the DLPFC and PPC than in a set of other regions sensitive to reward immediacy (i.e. ventral striatum, MPFC, Posterior Cingulate) when individuals choose larger delayed rewards (McClure et al., 2004).

However, PPC activation during intertemporal choice tasks may occur for a different reason – because intertemporal choices require numerical comparisons. In order to choose the reward with the better long-term value, an individual must quantitatively

assess the value of each option, both the monetary value of the reward and the time of receiving it. Both human fMRI and lesion studies indicate that one region of the PPC, the intraparietal sulcus (IPS), is involved in numerical manipulations (Dehaene et al., 2003). Furthermore, disruption of the left IPS using TMS or low frequency rTMS impairs performance on numerical comparison tasks as demonstrated by increased reaction times (Andres et al., 2005; Cappelletti et al., 2007). Therefore, increased activation seen in the PPC during intertemporal choice tasks could reflect engagement of numerical comparison functions rather than engagement of valuation or cognitive control functions.

Another question of interest is whether the DLPFC and PPC have a similar role in intertemporal choices involving monetary losses and those involving monetary gains. The neural circuitry that helps individuals decide what rewards to approach may be different from the circuitry that helps them decide what aversive outcomes to avoid, an idea that is consistent with different theories of motivation (Gray, 1981; Panksepp et al., 2002). If so, then the DLPFC and PPC could both play a function in intertemporal choices for gains, but not in intertemporal choices for losses. fMRI research, however, reveals that the PPC and DLPFC are similarly activated during intertemporal choices involving only monetary gains and those involving only monetary losses (Xu et al., 2009). These findings suggest the alternative possibility that the functions of the PPC and DLPFC cut across motivational lines and may perform the same functions in both intertemporal choices involving gains and those involving losses. If so, then in both types of choices, the DLPFC and PPC could help individuals make choices for options with the better long-term value when another option has a better immediate value. This would be supported if disruption of the same regions of the DLPFC or PPC with low frequency

rTMS increased choice of both monetary gain and loss options with better immediate values. In choices involving losses, such an effect would be seen as a greater preference for larger delayed over smaller immediate losses.

To examine whether the DLPFC and PPC are both essential for making choices for options with a better long-term but worse immediate value than alternatives, we utilized low frequency rTMS to transiently suppress functioning in the DLPFC and PPC while individuals performed an intertemporal choice task. We predicted that disruption of both regions would lead to greater selection of options with a better immediate value, and that we would see this effect both with choices involving monetary gains and with choices involving monetary losses. That is, individuals who received real rTMS were predicted to choose more immediate gains and fewer immediate losses than were individuals who had received sham rTMS. Additionally, we administered a numerical comparison task to see whether greater selection of options with a better immediate value following disruption of brain regions would be associated with deficits in the ability to compare two numerical values.

Methods

Participants

64 right-handed individuals (43.75% female) between the ages of 18 and 30 (*M* Age = 21.06, SD = 2.79) from Vanderbilt University and the Nashville community participated in this study. All of these participants reported having no history of neurological or psychiatric problems, no females were currently pregnant, and no

participants had previously received TMS. 95 additional subjects were consented but excluded or withdrew. Reasons subjects were excluded included baseline time preferences revealed in Session 1 that did not allow us to make task stimuli (see criteria below), inconsistent preferences as revealed in Session 1, no loss discounting as revealed on the loss discounting prescreen (for approximately 20 subjects this was given following consent), having previously received TMS, risk factors that could increase the chances of having negative effects from TMS (e.g. neurological conditions), excessive movement during TMS, and experimenter error. Additionally, two subjects were excluded because they exhibited choice preferences that were extremely unstable across sessions on the intertemporal choice task (mean choice of delayed gains or losses changed by more than 75% across sessions 2 and 3). Both of these subjects were in the sham stimulation group. No other participants in the sham or real stimulation groups exhibited changes in preference across sessions of this magnitude. All participants completed written informed consent approved by the Vanderbilt IRB.

Session One – Indifference Point Procedure

We screened subjects prior to consent with a questionnaire to ensure that we only recruited participants who discounted delayed monetary losses. The screen included 4 questions, each which required participants to write the magnitude of a hypothetical monetary loss to be incurred in 4 weeks that they valued equally to an immediate hypothetical monetary loss of a specific magnitude. Screening was essential for the study of monetary losses since prior research shows that some individuals actually prefer to

suffer a monetary loss or experience other aversive events sooner rather than later (Frederick et al., 2004).

Following consent, we determined subject indifference points between immediate and delayed monetary rewards. To do so, subjects completed a task on a computer in which we determined 4 immediate equivalents (i.e. dollar amount of an immediate monetary reward that a person values equally to a delayed monetary reward) for 4 different delayed monetary gains: $2.50, $5.00, $7.50, and $10.00 available at 4 weeks in the future. The order of determining immediate equivalents was from low to high magnitudes. Subjects were told that they would receive the amount of money for one random choice at the time associated with that choice. For timing of one trial of this task and presentation of items, see Figure 4A.

To determine each immediate equivalent, participants were presented with an initial trial in which they made a choice between a delayed monetary gain and an immediate gain of half the value of the delayed option. If participants chose the delayed option, the value of the immediate option increased by half, and if they chose the immediate option it decreased by half. On the next trial, the immediate value changed in a similar way as on the previous trial but only by one quarter of the original value. Over six trials, the immediate value increased or decreased by progressively smaller amounts (i.e. by 1/(2^x) where x was trial number) depending on participant responses so that the subjective value of the immediate amount would iteratively approach that of the delayed amount. After the sixth trial, a final catch trial was presented in which the immediate value was higher than the just-calculated immediate equivalent. This provided a check to ensure that subjects were answering according to their preferences (i.e. were answering

consistently). If they did not choose the immediate value on the catch trial, the immediate equivalent for that specific delayed magnitude reward was determined again (i.e. the six trial procedure and check was repeated). After answering consistently, or after completing the indifference point procedure three times, participants then performed the indifference point procedure for the next delay-magnitude pair. If participants answered inconsistently for three indifference point procedures in a row for a given delay-magnitude pair they were excluded.

Figure 4: Trial Structure of the Intertemporal Choice task. A. Intertemporal Choice*. Following 500 ms of fixation, two different money/time pairs were presented on a computer screen (side randomized). The pairs always consisted of a smaller amount of money to be gained or lost immediately and a larger amount of money to be gained or lost, respectively, four weeks in the future. These time points were chosen because prior research has shown disrupted intertemporal choice with these time points following rTMS (Figner et al, 2011). The words "For Sure" indicated that subjects were making choices for options that were certain. Subjects decided which money/time pair they would prefer by pressing the "z" or "m" key and had up to 6500 ms to respond. Immediately after responding, the triangle under the subject's selected option turned red and the other triangle disappeared to indicate the subject's choice; this was displayed for 250 ms. If subjects did not respond, the next part of the trial began after 6500 ms. B. Numerical Comparison. The same two monetary values that had been shown for the intertemporal portion remained on the screen. Additionally, the words "Larger" or "Smaller" were displayed above both monetary values, indicating whether participants were instructed to select the larger or smaller monetary value. Participants had up to 4000 ms to respond, and immediately following response the selected option was displayed (as was done for the intertemporal choice portion) for 250 ms. However, if subjects did not respond within 4000 ms, the next trial began. (* The timing and display of items in the Indifference Point Task were identical to those of the Intertemporal Choice part of the Intertemporal Choice Task (i.e. A. above) except there was no time limit for response).

Following the indifference point procedure for delayed monetary gains, participants performed the same indifference point procedure for delayed monetary losses. For the loss procedure, the same magnitudes of delayed incentives were used as were used for the gains task, except now subjects chose which of the two values they preferred to lose. They were told that they would have to pay the amount of money for one random choice at the time associated with that choice.

Subjects whose immediate equivalents indicated they did not discount the value of the future gains and losses were excluded from further sessions. Additionally, subjects were excluded if any immediate equivalent was smaller than 10 cents or if the difference between the immediate equivalent and associated delayed value was less than 10 cents. This was done to ensure that there would be an adequate range of values above and below each subject's immediate equivalents to create choice stimuli for further sessions.

At the end of session one, one random trial of intertemporal choice for gain and one random trial of intertemporal choice for loss were selected for payout. Subjects then either received or paid (depending on trial type) the amount of money at the specified time associated with their choice on that trial.

Sessions Two and Three

Participants that met eligibility requirements after session one were equally divided into 4 groups of 16 subjects. Group one received rTMS to the left DLPFC in one session and to the left PPC in a separate session (with counterbalanced ordering). Group two received rTMS to the right DLPFC in one session and to the right PPC in a separate session (with counterbalanced ordering). Group three received sham rTMS to the same

areas as group one (half DLPFC first/half PPC first), while group four received sham rTMS stimulation to the same areas as group two (half DLPFC first/half PPC first). Each of the four groups had the same gender distribution (9 Males and 7 females). Except for the different brain regions stimulated across sessions, procedures in sessions 2 and 3 were identical.

Prior to receiving rTMS, participants performed shortened practice versions of the tasks, in order to make sure they understood the tasks. Subjects were told that one gain choice and one loss choice from all the task trials performed after stimulation would be randomly selected for payment at the end of each session.

After completing the practice tasks, participants received real or sham rTMS stimulation for 30 minutes to either the DLPFC or PPC (details outlined under TMS Methods section).

Intertemporal Choice Task

Immediately following completion of stimulation, participants completed a 72 trial intertemporal choice task on a computer. On each trial, participants completed an intertemporal choice for monetary gain or loss (randomly mixed) followed by a numerical comparison (See Figure 4). Half of the choices involved losses and half involved gains.

The values of the delayed options were the same as in session 1 ($2.50, $5.00, $7.50, and $10.00). The magnitude of each immediate option was calculated as a specific percentage difference from that subject's session one immediate equivalent for the delayed option. This allowed us to vary the subjective value of each immediate option

(as revealed in session one) relative to the associated delayed option which let us predict which option an individual should choose on each trial (See Table 1).

Table 1: Immediate Values for Different Trials in the Intertemporal Choice Task.

Immediate Relative Value	Total	Description
1. Immediate Subjective Value Lower than Delayed Subjective Value*	32	Magnitude of immediate option is below immediate equivalent (i.e. immediate value at indifference point). There were 4 different percent distances below the immediate equivalent (-10%, -20%, -35%, and -50%). Subjects who receive sham stimulation should choose the delayed gain if it is a choice for gains, and the immediate loss if is a choice for losses.
2. Immediate Subjective Value Equal to Delayed Subjective Value	8	Magnitude of immediate option is equal to immediate equivalent (i.e. percent distance equal 0%). Neither immediate nor delayed choice is predicted for sham group.
3. Immediate Subjective Value Higher than Delayed Subjective Value	32	Magnitude of immediate option is above immediate equivalent. There were 4 different percent distances above the immediate equivalent (10%, 20%, 35%, and 50%). Subjects who receive sham stimulation should choose the immediate gain if it is a choice for gains, and the delayed loss if is a choice for losses.

Each delayed gain and each delayed loss magnitude was shown nine times. Each delayed option was paired once with an immediate option located at each of the 9 percent distances from the immediate equivalent of the delayed value. Values at negative distances were taken as a function of the percentage difference between the immediate equivalent and zero. A -35% distance indicated that the immediate value was 35% less than the immediate equivalent. Thus if the immediate equivalent of a $2.50 delayed option was $2.00, the immediate value was $1.30 (i.e. 2.00 - .35*2.00 = 1.30). Values at positive distances were calculated as a function of the percentage difference between the immediate equivalent and the delayed value. A 35% distance indicated that the immediate value was greater than the immediate equivalent by 35% of the distance between the immediate equivalent and the delayed value. Thus, if the immediate equivalent of a $2.50 delayed option was $2.00, the immediate value was $2.18 (i.e. 2.00 + (.35 * (2.50 -2.00)) = 2.18). *Lower relative subjective value means less positive value (for monetary gains) and less negative value (for monetary losses) than the alternative. Conversely higher relative subjective value means more positive value (for gains) and more negative value (for losses) than the alternative.

Following each intertemporal choice, individuals were asked to make a numerical comparison of the two monetary values that were available on the previous intertemporal choice trial (See Figure 4). This was performed to examine whether subjects had deficits in the ability to compare two monetary values. On half of the trials (determined randomly), participants indicated which of the two values was larger and on the other half of trials indicated which of the two values was smaller.

Indifference Point Task

After completing all of the aforementioned trials, subjects were given the identical indifference point procedure as in session one. However, this time the iterative procedure was only done to calculate the immediate equivalent for delayed values twice, first for a $10.00 gain available at 4 weeks and then for a $10.00 loss to be paid in 4 weeks. Additionally, no consistency checks were performed in order to limit the time to complete the task. Collecting subject indifference points post-stimulation provided an added check of whether real rTMS to the DLPFC or PPC increased impulsive choices.

At the end of sessions two and three, one random trial of intertemporal choice for gain and one random trial of intertemporal choice for loss were selected for payout. Subjects then either received or paid (depending on trial type) the amount of money at the specified time associated with their choice on that trial.

TMS Methods

Low frequency (1 Hz) rTMS was delivered with a MagStim TMS double 70 mm (Figure 8) coil (Magstim, Wales, UK) at 54% power. The rTMS parameters used were within currently recommended guidelines (Rossi, Hallett, Rossini, & Pascual-Leone,

2009) and stimulation with these parameters leads to suppression of excitability in the targeted region for a period of time following stimulation (Robertson, Theoret, & Pascual-Leone, 2003). We note that 1 Hz rTMS applied to the DLPFC at 54% power has previously led to altered intertemporal choices following stimulation (Figner et al., 2010). Sham stimulation was delivered with a MagStim placebo coil, which produced clicks that resembled the sound of rTMS, but without a magnetic pulse. Given the potential for subjects to identify whether they received sham or real stimulation, it was prudent for subjects to receive only one of these types of stimulation. Subjects were blind to the type of stimulation (e.g. sham vs. real) they received.

Positioning of the TMS coil was accomplished by using the 10-20 EEG System, which has previously been used to deliver TMS to identified brain regions with reasonable structural accuracy (Herwig, Satrapi, & Schonfeldt-Lecuona, 2003). This method was used to ensure proper positioning of the coil over the DLPFC (BA 46/9) and PPC (centered on the intraparietal sulcus). Specifically, to localize the PPC, the center of the coil was held tangentially to the participant's head with the handle pointing rostrally and placed over P3/P4 which has been shown to lie over this area (Andres et al., 2005; Herwig et al., 2003). To localize the DLPFC, the center of the coil was held tangentially to the participant's head with the handle pointing caudally, and placed one cm antero-lateral to F3/F4 which has been suggested to provide better coverage over BA 9/46 than do the points F3/F4 (Herwig et al., 2003). These specific points for stimulation were marked on a lycra swimcap that subjects wore; the position of the swimcap was placed in a consistent position on the head across subjects, by using the nasion, inion, and preauricular points as physical landmarks for placement. During stimulation, all

participants wore earplugs as protection against the noise of the rTMS pulse. Participants maintained their head position during rTMS administration using a chin/head rest and were visually monitored to ensure that no movement had occurred. In cases where a participant moved their head, the coil was immediately repositioned over the target.

All individuals received either real or sham rTMS to the DLPFC or PPC in session two for 30 minutes and the same type of stimulation to the other brain region (PPC or DLPFC) on the ipsilateral side for 30 minutes in session three. The order of stimulating these two regions was counterbalanced across subjects in each group. We used an "offline" rTMS paradigm; subjects completed tasks after stimulation was completed. Since impairments in behavior following low frequency rTMS have been shown to last for half the time of the previous stimulation (Mottaghy, Gangitano et al., 2002), the tasks following rTMS were limited to the first 15 minutes post-stimulation.

Statistical Analysis Methods

Using PASW Statistics 18 (SPSS Inc., Chicago, IL), data were analyzed with Generalized Estimating equations (GEE) which allow one to model effects while accounting for correlations within observations of individual subjects (Liang & Zeger, 1986). In order to test the effects of real stimulation relative to sham, separate GEE models were created for responses following stimulation to each region (i.e. left DLPFC, left PPC, right DLPFC, right PPC). To examine within subject differences across stimulation regions in subjects who received real rTMS, separate GEE models were created for data for subjects who received stimulation to each hemisphere (i.e. left, right). All models contained an intercept term.

We used GEE models with a logit link function and binomial distribution to predict choice of an immediate option (0 = Chose Delayed, 1 = Chose Immediate) at the trial level on the Intertemporal Choice Task. Separate models were created to predict choice on gain and loss trials Independent variables in all models included Stimulation Type (0 = Sham rTMS, 1 = Real rTMS), Immediate Relative Value (i.e. percent distance of the immediate option from the immediate equivalent of the delayed option, which ranged from -50% to 50%) and Session Number (0 = Session 2, 1 = Session 3). Additionally, we included the interaction of Immediate Relative Value by Stimulation Type if it was significant ($p < .05$). To examine within subject differences across region of stimulation in the real rTMS groups, we created similar GEE models, except the within subject variable Region (0 = DLPFC, 1 = PPC) was used instead of the between subject variable Stimulation Type; the interaction of Region by Immediate Relative Value was only included if significant.

All other GEE models we created had an identity link function and normal distribution. To predict reaction time on trials of the Intertemporal Choice Task (measured in milliseconds (ms)), we created separate GEE models for gain and loss trials. Between subject models (each limited to one region) and within subjects models included the same independent variables as the choice models, plus up to two additional variables: Immediate Relative Value ^{2} (i.e. Immediate Relative Value squared) and the Interaction between this variable and Stimulation Type (or Region, depending on analysis). Interactions were only included in the final models if significant.

We next used GEE models to predict reaction time in ms on the numerical comparison portion of the Intertemporal Choice Task; models contained data from all

trials pooled together (i.e. both gain and loss) and were limited to correct trials. Models contained the same independent variables as the models predictive of choice, and an interaction term was only included if significant. We did not create models to examine numerical comparison accuracy, as both real and sham rTMS groups performed at near perfect accuracy on the task (Mean accuracy for real rTMS group= .976, SD = .037; Mean accuracy for sham rTMS group= .976, SD = .023).

To predict effects of real relative to sham stimulation on immediate equivalents measured on the Indifference Point Task, we utilized the Generalized Linear model function in PASW Statistics 18 (SPSS Inc., Chicago, IL) with an identity link function and normal distribution. Each model examined the effects of stimulation at one of the four regions for either gain or loss values and contained the following predictors: Stimulation Type, Session Number, and Session 1 Immediate Equivalent. This last variable was included as a covariate because immediate values at the indifference point (i.e. immediate equivalents) following stimulation should be related to subject's initial pre-rTMS (i.e. Session 1) values, and thus significant effects of Stimulation Type might not emerge without controlling for this variable. To predict immediate equivalents across region of stimulation for subjects who received real rTMS, we used GEE models that included Region (-1 = neither (i.e. for Session 1), 0 = DLPFC, 1 = PPC) and Session Number (-1 = Session 1, 0 = Session 2, 1 = Session 3) as predictors. Separate models were run for gain and loss immediate equivalents.

For all GEE models, we used unstructured working correlation matrices. However, because the majority of analyses performed on numerical comparison reaction time failed to converge, we also constructed GEE models for this dependent variable with

exchangeable working correlation matrices. In order to ease comparison of models for each dependent variable, all presented results for a specific dependent variable were based on models with a similar working correlation matrix structure. Thus, while all reported results for numerical comparison reaction time were based on exchangeable matrices, all others were based on unstructured matrices.

RESULTS

Session One Indifference Points

Using independent sample t-tests, we first compared mean immediate equivalents measured in session 1 on the Indifference Point Task for individuals in each real rTMS group (i.e. left or right hemisphere) with individuals in the group that would receive sham rTMS to regions of the same hemisphere. This was done to ensure that session 1 choice preferences across stimulation groups were not different. Sixteen independent sample t-tests were performed in all: eight for subjects in the groups that would receive stimulation to left hemisphere regions and eight for subjects in the groups that would receive stimulation to right hemisphere regions. Each t-test compared mean immediate equivalents across stimulation groups for one of the eight types of delayed incentives (i.e. 4 magnitudes of delayed gain and 4 magnitudes of delayed loss). There were no significant differences in mean session one immediate equivalents between the left hemisphere real and sham rTMS groups, and no significant differences between the right hemisphere real and sham rTMS groups. Importantly, these results reveal that prior to receiving rTMS (i.e. in session 1), mean levels of delay discounting were similar across

real and matched sham rTMS groups. For mean session 1 immediate equivalents, see Table 2.

Table 2: Mean Session One Immediate Equivalents in the Indifference Point Task.

	Left Real		Left Sham		Right Real		Right Sham	
	Mean	(SD)	Mean	(SD)	Mean	(SD)	Mean	(SD)
Gain $2.50	1.226	(.529)	1.146	(.456)	1.424	(.598)	1.510	(.619)
Gain $5.00	2.964	(1.096)	2.804	(1.153)	3.391	(.828)	3.234	(1.157)
Gain $7.50	4.608	(1.767)	4.576	(1.526)	4.701	(1.683)	4.748	(1.836)
Gain $10.00	6.836	(2.055)	6.194	(2.266)	6.523	(2.274)	6.681	(2.090)
Loss $2.50	1.509	(.510)	1.570	(.588)	1.456	(.633)	1.389	(.555)
Loss $5.00	3.361	(.988)	3.168	(1.109)	3.014	(1.163)	2.673	(1.306)
Loss $7.50	5.119	(1.523)	4.640	(1.699)	4.843	(1.769)	4.916	(1.587)
Loss $10.00	7.103	(2.130)	6.710	(2.223)	6.621	(2.282)	6.164	(2.307)

Mean session 1 immediate equivalents for specific magnitudes of delayed gains and losses (listed in left column) available in 4 weeks. Averages are constructed separately for subjects in each of the four rTMS groups (listed in column titles).

Intertemporal Choices

For descriptive statistics of overall mean proportion of choices for immediate gains and losses following each type of stimulation to each side of each brain region, see Table 3.

Table 3: Mean Choice of Immediate Options in the Intertemporal Choice Task.

	Left DLPFC		Left PPC		Right DLPFC		Right PPC	
	Mean	(SD)	Mean	(SD)	Mean	(SD)	Mean	(SD)
Real: Gain	.389	(.213)	.389	(.168)	.511	(.190)	.474	(.205)
Sham: Gain	.396	(.197)	.413	(.156)	.440	(.247)	.379	(.200)
Real: Loss	.755	(.178)	.728	(.172)	.653	(.190)	.644	(.194)
Sham: Loss	.749	(.183)	.711	(.197)	.718	(.176)	.749	(.245)

Mean proportion of choices for the immediate option across subjects computed for each type of stimulation (i.e. Real vs. Sham) to each side and region (listed in column titles). Statistics are collapsed across all gain trials and all loss trials separately.

Model parameters for all GEE models that predict choice of an immediate monetary option are presented in Table 4. All between subject models contained Stimulation Type, Immediate Relative Value, Session Number, and an intercept as predictors. We also included the interaction of Immediate Relative Value by Stimulation Type in all initial models, to assess whether effects of real rTMS on choice depended on the relative value of the immediate reward. However, if this interaction term was not significant ($p >= .05$), we did not include it in the final model.

Disruption of the right DLPFC led to greater choice of options with a better immediate but worse long-term value. Importantly, this was seen in both the gain and loss domains. Individuals who received real rTMS to the right DLPFC chose more immediate monetary gains (Odds Ratio(OR) = 1.645, $p <.01$) than did individuals who received sham stimulation to the same region (see Figure 5A). They also chose less immediate monetary losses (OR = .653, $p < .05$) than did individuals who received sham stimulation to the same region (see Figure 5B). These effects were not dependent on the relative value of the immediate option, as similar differences in choice behavior emerged across the range of gain and loss values. In contrast to the effects seen following disruption of the right DLPFC, disruption of the left DLPFC did not affect choices.

Disruption of the right PPC also led to greater choice of options with a better immediate, but worse long-term value, and occurred in both the gain and loss domains. In contrast to the effects seen following disruption of the right DLPFC, the effects seen following disruption of the right PPC were dependent on the relative subjective value of the immediate option. There was an interaction of Immediate Relative Value by Stimulation Type on choice of both immediate monetary gains ($p <.01$) and immediate

monetary losses ($p < .05$). At low relative subjective values of the immediate option, individuals who received real stimulation chose similarly to those who received sham stimulation. However, as the subjective value of the immediate option increased relative to that of the delayed option, individuals became increasingly more likely to choose immediate monetary gains (see Figure 5C) and increasingly less likely to choose immediate monetary losses than individuals who received sham stimulation to the same region (see Figure 5D). Across both gain and loss choices, those who received real rTMS became increasingly more likely than subjects who received sham rTMS to choose the option with the better immediate, but worse long-term value, as the relative value of the immediate option increased.

Disruption of the left PPC also affected choice behavior, but only for gain trials. Real rTMS administered to the left PPC led to fewer choices for immediate monetary gains than did sham stimulation ($OR = .751$, $p < .05$) (see Figure 5E). The effect of left PPC stimulation contrasted with that of right PPC stimulation for gain trials in two ways. First, disruption of the left PPC led to decreased rather than increased choice of monetary gains with the better immediate value. Second, the effect occurred across the range of gain trials rather than interacting with Immediate Relative Value.

Table 4: Between and Within Subject GEE Models for Intertemporal Choice.

	GAIN Model			LOSS Model		
	B	(SE)	OR	B	(SE)	OR
Between: Left DLPFC						
Intercept	-.563***	(.155)	.569	.848***	(.121)	2.335
StimType	.138	(.149)	1.148	.162	(.161)	1.176
Imm. Rel. Value	.038***	(.003)	1.039	-.027***	(.002)	.973
Session	.194	(.153)	1.215	.588***	(.164)	1.800
Between: Left PPC						
Intercept	-.379***	(.102)	.685	1.264***	(.168)	3.538
StimType	-.286*	(.115)	.751	.065	(.172)	1.067
Imm. Rel. Value	.037***	(.003)	1.038	-.032***	(.002)	.968
Session	-.040	(.115)	.961	-.405*	(.170)	.667
Between: Right DLPFC						
Intercept	-.599***	(.116)	.549	.943***	(.109)	2.568
StimType	.498**	(.158)	1.645	-.426*	(.167)	.653
Imm. Rel. Value	.039***	(.002)	1.039	-.031***	(.002)	.969
Session	.421**	(.159)	1.523	.210	(.164)	1.233
Between: Right PPC						
Intercept	-.177	(.158)	.838	1.307***	(.180)	3.696
StimType	.442**	(.168)	1.555	-.719***	(.213)	.487
Imm. Rel. Value	.036***	(.003)	1.036	-.022***	(.004)	.978
Session	-.757***	(.171)	.469	.185	(.200)	1.203
StimType X Imm. Rel. Value	.011**	(.004)	1.011	-.011*	(.005)	.989
Real Within: Left						
Intercept	-.646***	(.081)	.524	1.090***	(.142)	2.974
Region	.051	(.073)	1.053	-.067	(.064)	.936
Imm. Rel. Value	.036***	(.003)	1.037	-.030***	(.003)	.970
Session	.070	(.077)	1.072	.178**	(.065	1.195
Real Within: Right						
Intercept	.148	(.077)	1.160	.404***	(.099)	1.498
Region	-.208***	(.050)	.812	.052	(.067)	1.054
Imm. Rel. Value	.043***	(.003)	1.044	-.032***	(.002)	.969
Session	-.205***	(.049)	.815	.249***	(.070)	1.283

Models predict choice of immediate option on each gain trial (in columns listed GAIN Model) or on each loss trial (in columns listed LOSS Model). In left column model predictors are listed below each model (name of model listed in **Bold**). Models listed "Between": contain data from subjects who received real rTMS and from subjects who received sham rTMS to region listed after word "Between". Models listed "Real Within" contain data from subjects who received real rTMS to regions of the hemisphere listed after words "Real Within". OR = Odds Ratio (e^B). SE = Standard Error of B. Intercept = Model Intercept. Imm. Rel. Value = Immediate Relative Value (percent distance of the immediate option from immediate equivalent of the delayed option). Region = Stimulation Region (0 = DLPFC, 1 = PPC). StimType=Stimulation Type (0 = Sham rTMS, 1 = Real rTMS). Session = Session Number (0 = Session 2, 1 = Session 3). X = Predictor is interaction of two terms. Models only contained interaction if significant at $p < .05$.* $p < .05$; ** $p < .01$; *** $p < .001$.

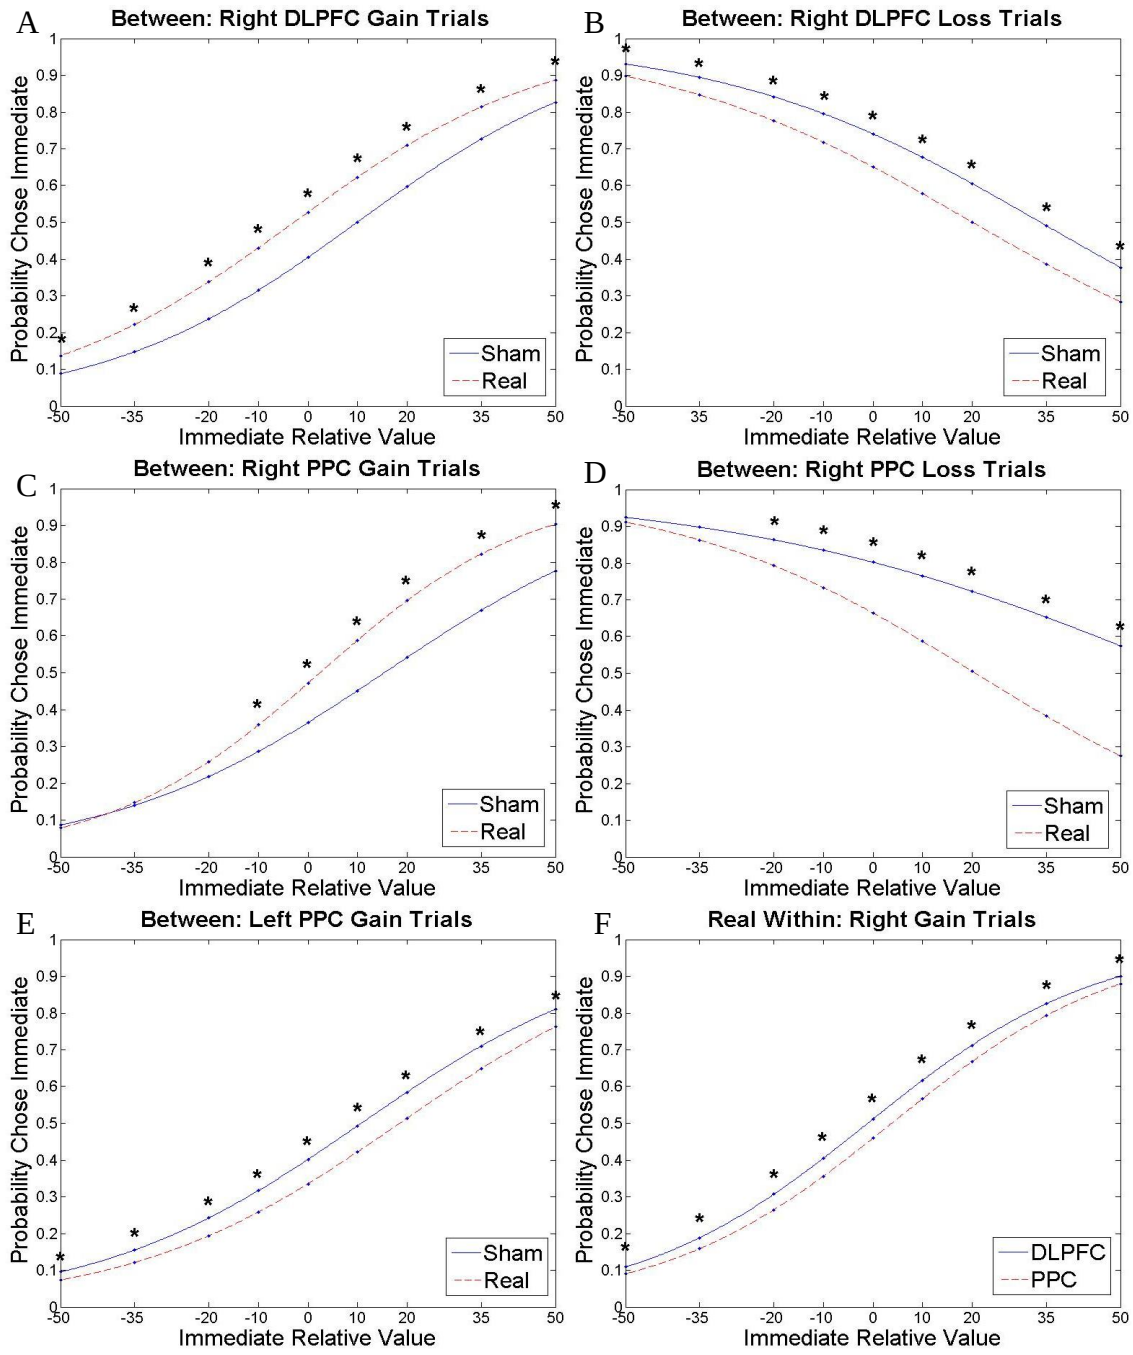

Figure 5: Intertemporal Choice Graphs. Probability of choosing the immediate option as a function of the relative value of the immediate option. Numbers on x-axis are the percent distance between the immediate option and the immediate equivalent of the delayed option (i.e. x-axis represents Immediate Relative Value). Graphs show choice at average session (Session = .5). A-E. Compares choice on gain or loss trials (listed in title) between real vs. sham stimulation to one region (listed in title). F. Compares choice on gain trials across stimulation region for those who received real rTMS to right hemisphere regions. * Paired comparison significant at $p < .05$.

To determine whether there was a stronger effect of rTMS in the right PPC or right DLPFC, we examined whether there were significant effects of stimulation region within subjects who received real rTMS (see Table 4). Individuals chose fewer immediate monetary gains following disruption of the right PPC than of the right DLPFC (OR = .812, $p < .001$) (see Figure 5F). However, the magnitude of this effect was small, indicating that choice patterns were not very divergent following disruption of either region. Because of the small size of the effect, it is not discussed further. In contrast to the effect of region of disruption in the right hemisphere on choice for gains, there was no effect of region of disruption in the right hemisphere on choice for losses. Additionally, region of disruption in the left hemisphere did not affect choice of either gains or losses. Additionally, there were no significant interactions of Immediate Relative Value by Region on choice in any within subject models.

Both between and within subject analyses reveal that participants were sensitive to the relative subjective values of the two options as predicted. In all monetary gain models, there was a positive effect of Immediate Relative Value such that subjects were more likely to chose the immediate gain as the relative value of the immediate option increased (all $p < .001$). In contrast, in all monetary loss models, there was a negative effect of Immediate Relative Value, such that subjects were less likely to choose the immediate loss as the relative value of the immediate loss increased (i.e. as the magnitude of the immediate loss increased) (all $p < .001$). This reveals that while disruption of several regions with rTMS affected choice, disruption did not make participants insensitive to the relative values of the two options.

We found some effects of Session Number on choice of monetary gains and losses in between subject models, but these effects did not exhibit consistent patterns across models. In between subject models, there was a positive effect of Session Number on choice of immediate gains following stimulation to the right DLPFC ($p < .01$), but a negative effect of Session Number following stimulation to the right PPC ($p < .001$). In the domain of losses, there was a positive effect of Session Number on choice of immediate losses following stimulation to the left DLPFC ($p < .001$), but a negative effect of Session Number following stimulation to the left PPC ($p < .05$). It is difficult to interpret how Session Number affected choice across all subjects, because there was no consistent effect of Session Number in either the gain or loss domain.

There were, however, more interpretable effects of Session Number in within subject models. Amongst individuals who received real rTMS, Session Number was consistently positively associated with greater choice of the option with the better long-term value. Individuals were less likely to choose immediate gains following disruption of right hemisphere regions ($p < .001$) and more likely to choose immediate losses following disruption of regions of either hemisphere (both $p < .01$) in session 3 than in session 2. This suggests that individuals who received real rTMS improved in their ability to choose the better long-term option over time; they were more likely to choose the larger delayed gain and the smaller immediate loss in session 3 than in session 2.

As a supplemental analysis, we investigated whether subjects who received real rTMS to right hemisphere regions were more likely to choose the option with the better immediate value than were subjects who received real rTMS to the same regions in the left hemisphere. To assess this, we constructed between subject GEE models similar to

those produced across real and sham groups, except the predictor Stimulation Type was replaced by Hemisphere (0 = Left, 1 = Right). Each model only included data from participants who received real rTMS to a brain region (PPC or DLPFC). Consistent with the between subject results, disruption of either right DLPFC or right PPC was associated with a greater tendency to choose the option with the better immediate value. Individuals who received real rTMS to the right DLPFC were more likely to choose immediate gains (OR = 2.137, $p < .001$) and less likely to choose immediate losses (OR = .461, $p < .001$) than were individuals who received real rTMS to the left DLPFC. Similarly, those who received real rTMS to the right PPC were more likely to choose immediate gains (OR = 2.064, $p < .001$) and less likely to choose immediate losses (OR = .636, $p < .01$) than were those who received real rTMS to the left PPC. There was no interaction of Immediate Relative Value by Hemisphere in any of the models. Thus, unlike in the between stimulation type analyses, in the between hemisphere analyses, the effects of disruption of the right PPC on choice were not dependent on the relative value of the immediate option. The greater tendency to choose the better immediate, but worse long-term gains and losses seen following disruption of right than of left hemisphere regions supports the results of our analyses investigating differences in choice following real vs. sham rTMS to the right DLPFC and right PPC. Together, both analyses indicate that disruption of the right DLPFC and the right PPC increased choice of gain and loss options with better immediate, but worse long-term values.

We performed another supplemental analysis to see if changes in choice patterns seen following disruption of the DLPFC or PPC were dependent on how close the subjective values of the two options were to each other. This was done because prior

research shows that disruption of the DLPFC with rTMS affects choice more on difficult intertemporal choices (Figner et al., 2010), and it might be more difficult to compare the values of the two options and make a choice as the subjective values approach each other. To assess this, we reran all eight of our between subject GEE models that compared choice between individuals in the real and sham rTMS groups, but added two terms to the models: a quadratic effect of Immediate Relative Value and an interaction of this term with Stimulation Type. The quadratic effect, Immediate Relative Value Squared (i.e. Immediate Relative Value^2), was used because the value of this term decreased as the subjective values of the two options approached each other. These analyses revealed a significant interaction between Immediate Relative Value^2 and Stimulation Type on loss trials following stimulation to the left DLPFC (p=.003). Despite this interaction, there were no significant differences between groups on choice. Because of the lack of significant differences, this finding is not further discussed. In contrast to the interaction seen following stimulation to the left DLPFC, there were not any significant interactions between Immediate Relative Value^2 and Stimulation Type following stimulation to other regions. Importantly, these results indicate that the effects on choice following disruption of the right DLPFC and the right PPC were not dependent on how close the relative subjective values of the two options were to each other.

Intertemporal Choice Reaction Time

We next tested whether reaction time was different across groups using GEE models (see Table 5) to see whether changes in choice patterns following disruption of brain regions were associated with changes in the time it took to choose between the

options. We assumed that as comparisons and choices became more difficult, reaction time would increase. Choices could become more difficult as the relative subjective values of the two options approach each other, and also as the objective monetary values of the two options approach each other. To examine both of these effects, all reaction time models included both linear and quadratic effects of Immediate Relative Value. As the linear term, Immediate Relative Value, increased, the monetary value of the immediate option approached that of the delayed option. In contrast, as the quadratic effect, Immediate Relative Value squared (i.e. Immediate Relative Value^2), decreased, the subjective values of the two options approached each other. As the relationship between Immediate Relative Value and choice could differ across groups, we also probed for interactions between Stimulation Type and the linear and quadratic effects of Immediate Relative Value. These interactions were only included in final models if significant.

Real rTMS over the right DLPFC or right PPC slowed reaction times relative to sham stimulation on loss trials (both $p < .05$). These slowed reaction times could indicate that it was harder following real stimulation to choose loss options with the better long-term value, since real stimulation to these regions decreased the tendency to choose these options. However, one should not make a strong claim based on the patterns of these effects. This is because changes in choice and reaction time interacted differently with the relative values of the two options. There were significant negative effects of Immediate Relative Value^2 on reaction time following real rTMS to the right PPC and to the right DLPFC but not following sham rTMS to these regions, as indicated by the significant interactions of Stimulation Type by Immediate Relative Value^2 (both $p <.05$)

and the nonsignificant main effects of Immediate Relative Value^2. This reveals that reaction time slowed as the subjective values of the two options approached each other for the real rTMS group, but not for the sham group. Additionally, it suggests that real rTMS made it more difficult for individuals to choose amongst losses as the subjective values of the two monetary loss options approached each other (see Figure 6A & 6B). If this reaction time effect were linked to choice, one would expect that individuals who received real rTMS would make fewer choices for options with the better long-term outcome as the subjective values of the options approached each other as well. However, as discussed earlier this was not the case. There were no significant interactions between Immediate Relative Value^2 and Stimulation Type on choice following stimulation to right hemisphere regions.

We did not find evidence for a relationship between changes in reaction time and changes in choice patterns on gain trials following disruption of the right DLPFC or the right PPC. There was an interaction of Immediate Relative Value by Stimulation Type on intertemporal choice reaction time for gain trials following stimulation to the right PPC ($p < .05$). While individuals who received sham rTMS to this region responded more slowly as Immediate Relative Value increased, individuals who received real stimulation did not. Despite this difference, reaction time following disruption of right PPC was not significantly slower than reaction time following sham stimulation (see Figure 6C), even though these subjects were more likely to choose immediate gains than those in the sham group. Additionally, there was no effect of Stimulation Type to the right DLPFC on reaction time for gain trials, even though choice was affected. The lack of significant differences in reaction time on gain trials across groups that had significant

differences in choice patterns on these trials indicates that changes in the ability to make good long-term choices on gain trials were not associated with changes in reaction time, following disruption of right hemisphere regions.

In contrast, there did appear to be a relationship between changes in reaction time and choice on gain trials following disruption of the left PPC. Relative to individuals who received sham stimulation, those who received real rTMS to the left PPC both had decreased reaction time on gain trials ($p < .001$; see Figure 6D) and made fewer choices for immediate gains. Neither the effects on reaction time nor the effects on choice following disruption of this region were dependent on the relative value of the immediate option, showing a similarity between the two effects. Individuals responded more quickly and were more likely to choose the better long-term option following disruption of the left PPC.

We found some between subject effects on reaction time in the left DLPFC as well. However, these effects appeared to be modest and were not linked to any changes in choice patterns, since disruption of the left DLPFC did not affect choice. Following stimulation delivered to the left DLPFC, we saw a significant interaction between Immediate Relative Value^2 and Stimulation Type ($p < .05$), such that reaction time on loss trials for the sham group slowed more as the two subjective values approached each other than it did for the real rTMS group (see Figure 6E). However, as no reaction time differences between groups reached statistical significance, this effect is not discussed further.

Table 5: Between Subject GEE Models for Intertemporal Choice Reaction Time.

	GAIN Model		LOSS Model	
	B	(SE)	B	(SE)
Between: Left DLPFC				
Intercept	3136.488***	(75.437)	3396.698***	(94.894)
StimType	-15.041	(89.065)	-91.084	(126.894)
Imm. Rel. Value	3.435***	(.522)	4.736***	(.706)
Imm. Rel. Value^2	-.070***	(.014)	-.153***	(.031)
Session	-414.228***	(89.483)	-375.119***	(100.187)
StimType X Imm. Rel. Value	---†		-1.172	(1.250)
StimType X Imm. Rel. Value^2	---		.086*	(.040)
Between: Left PPC				
Intercept	3375.831***	(77.763)	3468.686***	(80.793)
StimType	-348.695***	(72.604)	-6.946	(84.845)
Imm. Rel. Value	.282	(.829)	5.102***	(.524)
Imm. Rel. Value^2	-.113***	(.015)	-.036*	(.015)
Session	-479.025***	(72.556)	-550.298***	(90.422)
StimType X Imm. Rel. Value	2.451*	(1.086)	---	
Between: Right DLPFC				
Intercept	3330.106***	(57.810)	3387.930***	(66.438)
StimType	46.987	(76.358)	232.525*	(109.378)
Imm. Rel. Value	2.952***	(.771)	6.213***	(.881)
Imm. Rel. Value^2	-.110***	(.019)	.017	(.018)
Session	-452.088***	(82.957)	-449.830***	(95.091)
StimType X Imm. Rel. Value	---		-1.633	(1.116)
StimType X Imm. Rel. Value^2	---		-.075*	(.035)
Between: Right PPC				
Intercept	3599.815***	(84.868)	3433.702***	(69.620)
StimType	-32.280	(81.097)	452.172***	(91.683)
Imm. Rel. Value	2.526**	(.884)	8.206***	(.883)
Imm. Rel. Value^2	-.117***	(.015)	.001	(.025)
Session	-573.447***	(82.021)	-673.276***	(78.212)
StimType X Imm. Rel. Value	-3.057*	(1.289)	-2.267	(1.214)
StimType X Imm. Rel. Value^2	---		-.128***	(.033)

Models predict reaction time (ms) on gain or loss trials. Imm. Rel. Value^2 = Immediate Relative Value squared. ---†= Interaction not included in model (models only contained interaction if it was significant at $p < .05$). Other abbreviations same as in Table 4. * $p < .05$; ** $p < .01$; *** $p < .001$.

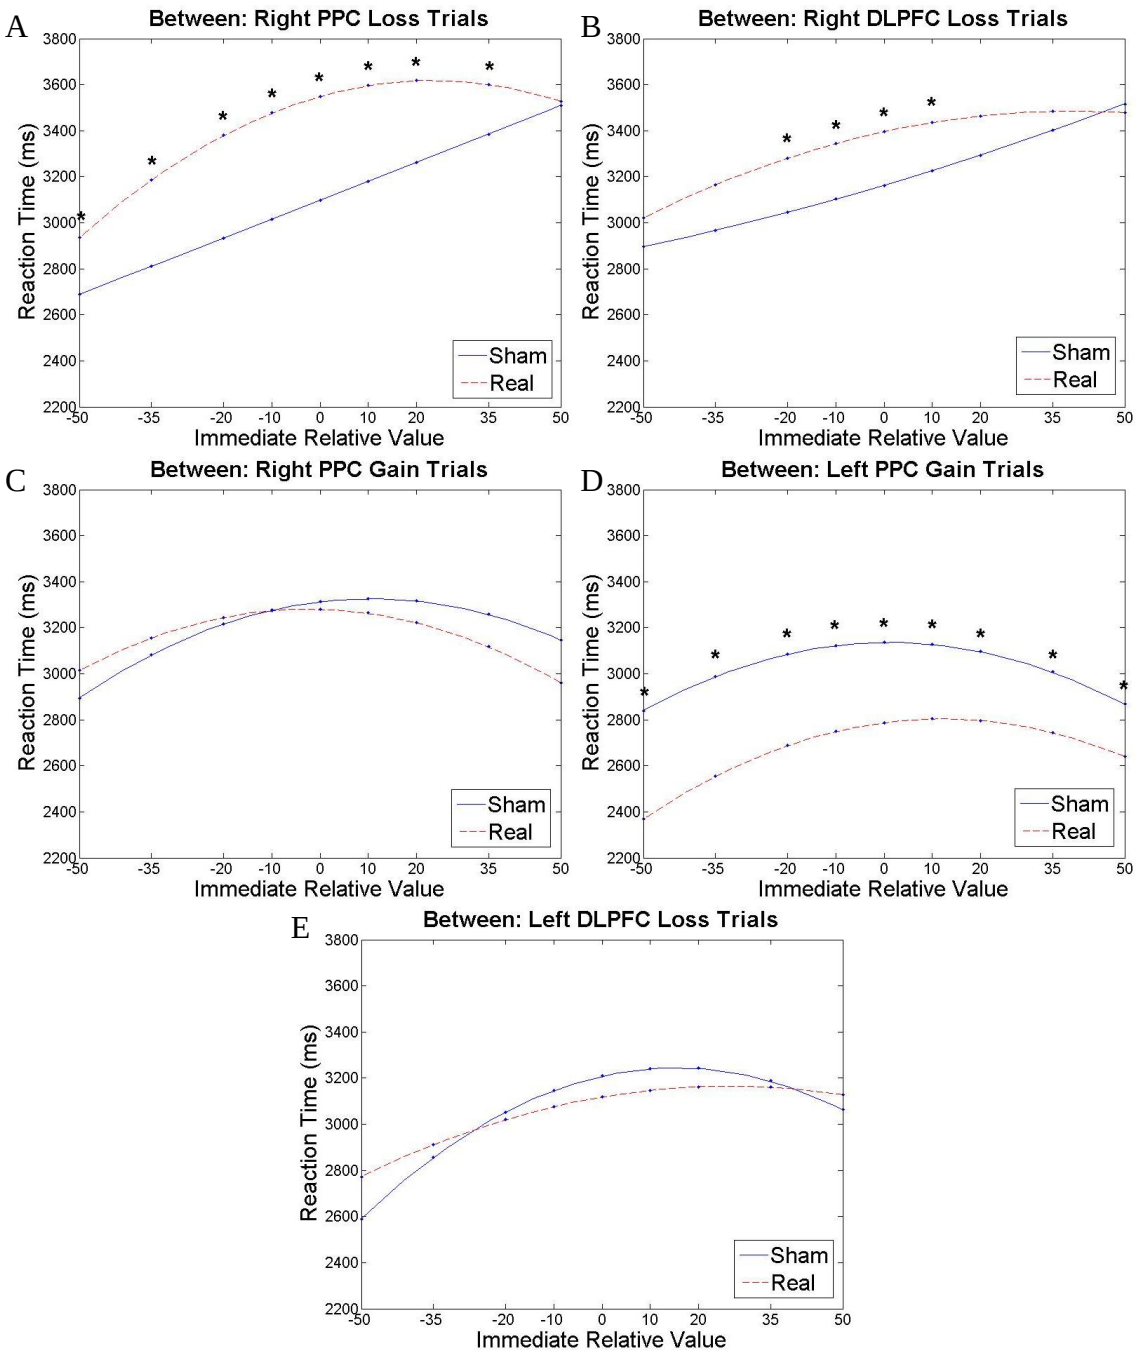

Figure 6: Between Subject Intertemporal Choice Reaction Time Graphs. Choice reaction time (in ms) as a function of the relative value of the immediate option. Numbers on x-axis are the percent distance between the immediate option and the immediate equivalent of the delayed option. Graphs show reaction time at average session (Session = .5). A-E. Compares reaction time on gain or loss trials (listed in title) between real vs. sham stimulation to one region (listed in title). * Paired comparison significant at $p < .05$.

The results of our between subject reaction time analyses suggest that the relationship between choice patterns and reaction time is dependent on the region of stimulation and the type of choice (i.e. gain or loss). Some effects on choice were associated with changes in reaction time, particularly on gain trials following rTMS to the left PPC, while others were clearly dissociable (i.e. effects following rTMS to right hemisphere regions on gain trials). There was some evidence that slowed reaction time was associated with greater choice of better immediate but worse long-term loss options, and that speeded reaction time was associated with greater choice of better long-term but worse immediate gain options. However, with the exception of the effects following rTMS to the left PPC, the relationships between choice and reaction time appeared to be weak.

For subjects who received real rTMS, there were some differences in reaction time that depended on region of stimulation, but only following left hemisphere stimulation (see Table 6). Following real rTMS to the left DLPFC and the left PPC, there were effects of Region on loss trials that were dependent on both Immediate Relative Value and Immediate Relative Value^2 (both $p < .001$; see Figure 7B), and effects of Region on gain trials that were dependent on Immediate Relative Value^2 ($p < .001$; see Figure 7A). To investigate whether the divergent effects of Immediate Relative Value^2 across regions were coupled with changes in choice patterns, we reran our left hemisphere within subject GEE gain and loss intertemporal choice models for subjects in the real rTMS group, but this time included additional terms for Immediate Relative Value^2 and the interaction of this term with Region. Both terms were nonsignificant, revealing that reaction time effects were not coupled with effects on choice. Since

analyses revealed there were no effects of Region on intertemporal choice following real rTMS to the left hemisphere, these reaction time effects are not discussed further.

Table 6: Within Subject GEE Models for Intertemporal Choice Reaction Time.

	GAIN Model		LOSS Model	
	B	(SE)	B	(SE)
Real Within: Left				
Intercept	3213.539***	(53.855)	3339.426***	(75.206)
Region	75.152	(46.643)	8.890	(49.856)
Imm. Rel. Value	3.159***	(.631)	3.394***	(.866)
Imm. Rel. Value^2	-.065**	(.020)	-.113***	(.021)
Session	-568.657***	(31.514)	-398.777***	(40.386)
Region X Imm. Rel. Value	.508	(.629)	3.305***	(.951)
Region X Imm. Rel. Value^2	-.119***	(.036)	.084***	(.022)
Real Within: Right				
Intercept	3374.700***	(57.030)	3624.722***	(71.978)
Region	61.083	(37.630)	31.714	(44.156)
Imm. Rel. Value	1.267*	(.609)	5.409***	(.558)
Imm. Rel. Value^2	-.102***	(.017)	-.070***	(.018)
Session	-289.310***	(38.950)	-339.540***	(43.682)

Models predict reaction time (ms) for subjects who received real rTMS on gain or loss trials. Imm. Rel. Value^2 = Immediate Relative Value squared. Models only contained interaction if it was significant at $p < .05$. Other abbreviations same as in Tables 4. * $p < .05$; ** $p < .01$; *** $p < .001$.

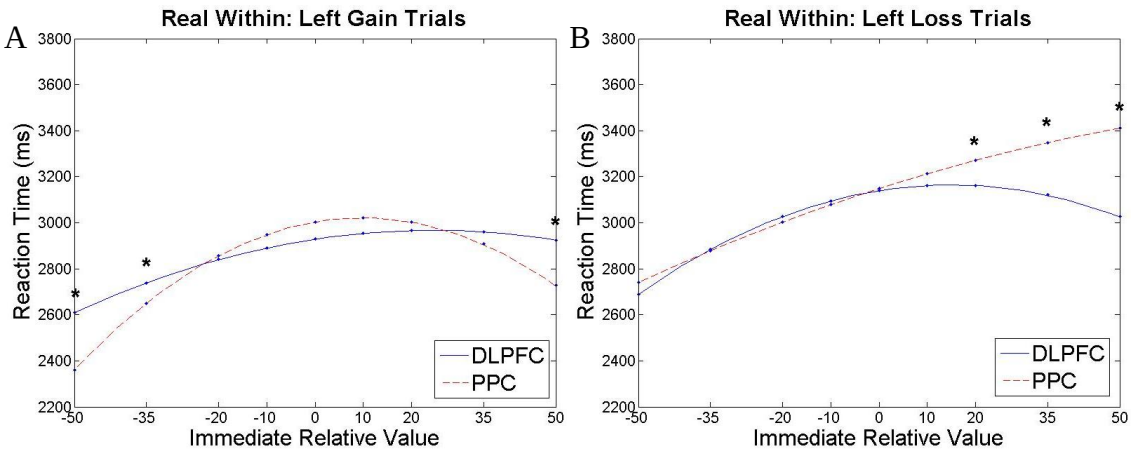

Figure 7: Within Subject Intertemporal Choice Reaction Time Graphs. Intertemporal choice reaction time (in ms) as a function of the relative value of the immediate option. Numbers on x-axis are the percent distance between the immediate option and the immediate equivalent of the delayed option. Graphs show reaction time at average session (Session = .5). A-B. Compares reaction time on gain or loss trials (listed in title) across region of stimulation for those who received real rTMS to left hemisphere regions. * Paired comparison significant at $p < .05$.

All intertemporal choice reaction time models revealed an effect of Session Number, such that subjects responded more quickly in session 3 than in session 2 (all $p < .001$); this improvement in reaction time was likely due to familiarity with the task. Additionally, in all models there was a significant positive linear effect of Immediate Relative Value ($p < .05$) and/or a significant negative quadratic effect of Immediate Relative Value (i.e. of Immediate Relative Value^2) ($p < .05$), as predicted. Reaction time thus increased as the monetary values of the two options approached each other (i.e. as Immediate Relative Value increased) and/or increased as the subjective value of the two options approached each other (i.e. as Immediate Relative Value^2 decreased). These reaction time effects suggest that trials did become more difficult both as the monetary values and as the subjective values of the two options approached each other.

Numerical Comparison Reaction Time

We next used GEE models to examine the effects of Stimulation Type, Immediate Relative Value, and Session Number on numerical comparison reaction time for correct trials to see if an increased tendency to choose options with the better immediate value on the Intertemporal Choice Task was associated with an impaired ability to perform numerical comparisons (see Table 7). We included Immediate Relative Value (i.e. of the option that had been immediate on the preceding intertemporal choice trial) as an independent variable, because numerical comparisons should become more difficult and slower as the monetary values of the two numbers approach each other, consistent with past research (Moyer & Landauer, 1967). As with other analyses, we only included an interaction term containing Immediate Relative Value if it was significant.

There were no significant effects of Stimulation Type on numerical comparison reaction time following stimulation to the right DLPFC or right PPC. There were also no significant interactions between Stimulation Type and Immediate Relative Value following stimulation to these regions. Thus, there was no evidence that a greater tendency to choose the option with the better immediate value, following disruption of right hemisphere regions, was due to a disrupted ability to perform numerical comparisons.

Table 7: Between and Within Subject GEE Models for Numerical Comparison Reaction Time.

	B	(SE)
Between: Left DLPFC		
Intercept	1229.931***	(60.732)
StimType	-10.794	(53.764)
Imm. Rel. Value	1.791***	(.331)
Session	-146.037**	(53.735)
Between: Left PPC		
Intercept	1206.382***	(45.622)
StimType	-6.429	(57.072)
Imm. Rel. Value	1.567***	(.442)
Session	-133.138*	(57.058)
StimType X Imm. Rel. Value	-1.051*	(.512)
Between: Right DLPFC		
Intercept	1157.663***	(55.137)
StimType	-17.158	(84.379)
Imm. Rel. Value	1.946***	(.352)
Session	89.258	(84.527)
Between: Right PPC		
Intercept	1317.983***	(60.156)
StimType	42.651	(70.440)
Imm. Rel. Value	.900**	(.279)
Session	-286.806***	(70.442)
Real Within: Left		
Intercept	1199.186***	(38.567)
Region	-12.359	(23.388)
Imm. Rel. Value	1.929***	(.401)
Session	-106.250***	(23.403)
Region X Imm. Rel. Value	-1.412**	(.451)
Real Within: Right		
Intercept	1234.196***	(66.061)
Region	33.197	(27.166)
Imm. Rel. Value	2.109***	(.377)
Session	-100.784***	(27.106)
Region X Imm. Rel. Value	-1.291*	(.531)

Models predict reaction time (ms) on accurate trials for subjects who received real or sham rTMS to listed region in bold. Models only contained interaction if it was significant at $p < .05$. Imm. Rel. Value is the Immediate Relative Value of the option that was immediate on the preceding choice trial. Other abbreviations same as in Table 4. * $p < .05$; ** $p < .01$; *** $p < .001$.

Disruption of only one region, the left PPC, affected numerical comparison reaction time differently across real and sham stimulation. Following rTMS delivered to this region, there was an interaction of Immediate Relative Value by Stimulation Type on reaction time ($p < .05$). Although reaction time increased as Immediate Relative Value increased following both real and sham rTMS, this value dependent increase was greater for sham than real stimulation. Despite this difference, however, there were no significant differences in reaction time across groups (see Figure 8A), revealing that this effect was modest.

We next looked at the effects of stimulation region on reaction time within subjects who received real rTMS (see Table 7). We found significant interactions of Immediate Relative Value by Region on numerical comparison reaction time among individuals who received real rTMS to the left hemisphere and among those who received it to the right hemisphere (both $p < .05$). Although reaction time increased as Immediate Relative Value increased following disruption of all regions, this value dependent increase was greater for DLPFC regions than for PPC regions (see Figure 8B & 8C). This suggests that while it became more difficult to compare the values of the numbers as they approached each other following disruption of either the DLPFC or PPC to either hemisphere, the level of difficulty increased more as the two values approached each other following disruption of the DLPFC.

All models revealed that subjects were sensitive to the relative numerical values of the two options. In every GEE model, there was a positive effect of Immediate Relative Value, such that as the value of the smaller option increased and the monetary values of the two options approached each other, reaction time increased (all $p < .01$).

This was expected as numerical comparisons should become more difficult as the distance between the two numbers decreases. Additionally, almost all models revealed an effect of Session Number. All significant effects of Session Number were in the negative direction, such that participants responded more quickly in session 3 than in session 2 (all $p < .05$), suggesting that performance improved with familiarity.

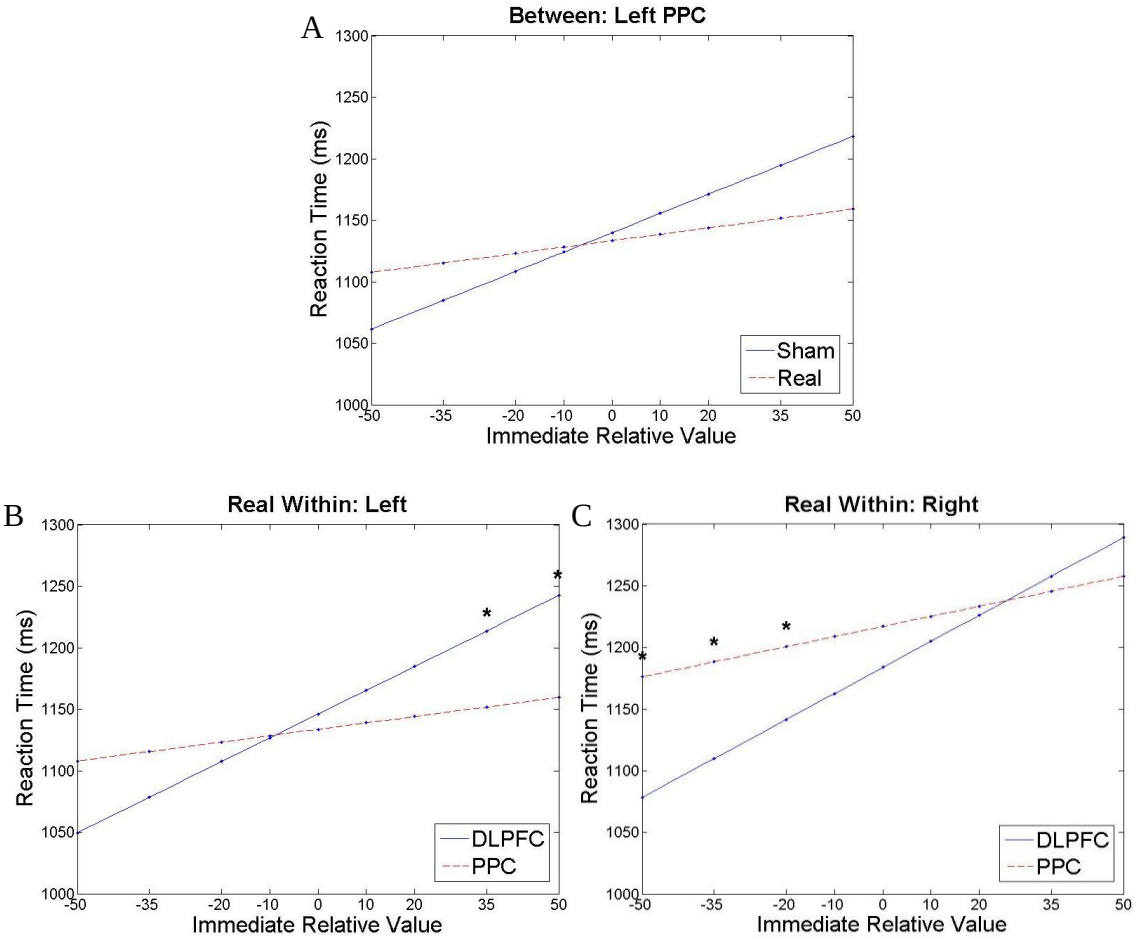

Figure 8: Numerical Comparison Reaction Time Graphs. Numerical comparison reaction time (in ms) as a function of the relative value of the immediate option on the preceding choice trial. Numbers on x-axis are the percent distance between the immediate option and the immediate equivalent of the delayed option. Graphs show reaction time at average session (Session = .5). A. Compares reaction time between real vs. sham stimulation to the left PPC. B-C Compares reaction time across region of stimulation for those who received real rTMS to regions in one hemisphere (listed in title). * Paired comparison significant at $p < .05$.

As a supplementary analysis, we reran our numerical comparison reaction time models by only modeling reaction time on correct trials in which the reaction time was >= 100 ms and that were within two standard deviations of each subject's mean reaction time. This was done to remove trials in which subjects may have hastily responded and trials that were outliers. It also allowed us to better compare our results with those of prior studies that have investigated the effects of TMS or rTMS on numerical comparison abilities, since it is typical to remove outliers before analysis (Andres et al., 2005; Cappelletti et al., 2007; Sandrini et al., 2004). This supplementary analysis led to similar results as the primary reaction time analysis, except there was now no longer any significant interaction of Immediate Relative Value by Stimulation Type following rTMS to the left PPC. Importantly, both the supplementary and primary analyses show that increased choice of better immediate, but worse long-term options, following disruption of the right DLPFC and the right PPC, was not associated with an impairment in the ability to perform numerical comparisons.

Indifference Points

As a second approach to examining changes in intertemporal choice, we examined whether subject's immediate equivalents for $10 delayed gains and losses were affected by disruption of the DLPFC and PPC. In order to examine the difference in immediate equivalents following real and sham rTMS to each hemisphere, we used the Generalized Linear Model, which included Stimulation Type, Session Number, and Session One Immediate Equivalent as predictors. We also created within subject GEE

models for subjects who received real rTMS to more thoroughly investigate how real rTMS affected immediate equivalents.

Contrary to the findings of the Intertemporal Choice Task, there were no effects of real rTMS on immediate equivalents on the indifference point procedure for monetary gains or losses. There was both no effect of Stimulation Type in between subject models and no effect of Region in within subject models for individuals who received real rTMS. Thus, greater choice of options with the better immediate value on the Intertemporal Choice Task was not associated with altered levels of delay discounting on the Indifference Point Task.

The lack of findings on the Indifference Point Task may have been partially due to decay of rTMS effects since this task was completed after the Intertemporal Choice Task. However, this is unlikely, since all subjects completed the Indifference Point Task before the effects of rTMS should have completely decayed (i.e. all subjects completed the task within 15 minutes after the termination of stimulation). Maximum completion time for the Indifference Point Task was 12.110 minutes (min) post stimulation (Mean = 8.242 min, SD = 1.102 min), and maximum completion time for the preceding Intertemporal Choice Task was 10.275 minutes post stimulation (Mean = 7.413 min, SD = .965 min).

The lack of findings could also have been due to differences in the components of the intertemporal choice and indifference point tasks. The smaller number of trials on the Indifference Point Task may have made choices insensitive to neuromodulation relative to choices on the Intertemporal Choice Task, which provided both a greater number of trials and a greater range of immediate and delayed values.

Discussion

We found that disruption of either the right DLPFC or the right PPC with rTMS led to greater choice of options with a better immediate, but worse long-term value, supporting the hypothesis that the functions of both of these regions help individuals make choices that optimize long-term value relative to immediate value. Increased choice of options with a better immediate value following rTMS was found in both the gain and loss domains. Individuals were more likely to choose small immediate monetary gains and less likely to choose small immediate monetary losses than were sham subjects and subjects who received real rTMS to contralateral brain regions

Although disruption of the right DLPFC or right PPC led to increased choice of options with a better immediate value, the effects seen following disruption of each brain region were different. Disruption of the right DLPFC led to increased choice of options with a better immediate value (i.e. immediate gains and delayed losses) generally. In contrast, the increased tendency to make choices for options with the better immediate value following disruption of the PPC was dependent on the value of the immediate option. Following disruption of the right PPC, individuals became increasingly more likely than sham subjects to choose options with a better immediate, but worse long-term value as the relative value of the immediate option increased (i.e. as immediate gains and losses increased). This suggests that both regions are involved in different functions during intertemporal choice.

Past neuroeconomic studies have outlined three major models to account for intertemporal choice behavior: single-valuation (Kable & Glimcher, 2007, 2009), dual-

valuation (McClure et al., 2004), and self-control (Figner et al., 2010). The single and dual valuation accounts both posit that choice directly follows from valuation, but differ in the number of valuation processes that influence choice. The single-valuation model holds that a number of regions of the brain, including the MPFC and ventral striatum, respond to the value of both immediate and delayed rewards, but respond less for rewards as they become more delayed. According to this model, disrupting the DLPFC or PPC should not lead to increased choice for either type of option because it would not change valuation of the choice options. In contrast, the dual-valuation model holds that there are two valuation processes, an impatient system, that values immediacy, and a more patient system, that includes the LPFC and PPC, which discounts delayed rewards much less steeply than the impatient system. Disruption of a region involved in the patient system should lead to decreased valuation of delayed rewards and increased choice of options with the better immediate value. Self-control accounts on the other hand hold that activity in the DLPFC, and possibly other brain regions, can override or modulate valuation-related activity in other regions so that individuals can choose the better long-term option (Figner et al., 2010; Hare et al., 2009). The self-control processes in the intertemporal choice theory of Figner and colleagues (2010) are characterized as classic cognitive control mechanisms that are deliberate and that have the ability to override more prepotent responses (i.e. prepotent responses to choose the option with the best immediate value), in favor of goals (Miller & Cohen, 2001).

The effects on choice seen following disruption of the right DLPFC are most consistent with the dual-valuation account, suggesting that functions of the DLPFC place more weight on the long-term value of options than do other regions. Disruption of this

region led to a shift in preference for options with a better immediate value, and this shift did not depend on the relative values of the two options. This is the effect one would expect if one disrupted activity in a brain region that does not discount the value of delayed rewards, relative to immediate rewards, as much as other regions do. Specifically, if there are parallel patient and impatient valuation systems, and the patient system is disrupted, the impatient system's valuation will dominate the decision process. In the dual-valuation model of McClure et al. (2004), impatient systems (i.e. ventral striatum, MPFC) would become dominant, once regions of the patient system (i.e. DLPFC and PPC) are disrupted. Our findings are consistent with the model of McClure and colleagues, and suggest that the DLPFC is part of a "patient valuation" system.

The authors of a past rTMS study of intertemporal choice have argued against a valuation role for the DLPFC based on their finding that disruption of the DLPFC affected choices but not valuation of the options (Figner et al., 2010). However, their lack of an effect could have been due to limited power to detect differences on the valuation task, as there were only 12 trials. This is quite possible, as the findings of several other studies do suggest that the DLPFC has an important role in valuation. Activation in the DLPFC has been shown to scale with the subjective value of items on both bidding and choice tasks (Hare et al., 2009; Plassmann et al., 2007; Plassmann et al., 2010). Furthermore, disruption of the DLPFC with low frequency rTMS has led to altered valuation of items on a bidding task (Camus et al., 2009). In contrast to the findings of Figner et al., the findings of these other studies are supportive of our suggestion that the DLPFC is involved in valuation of items during intertemporal choice.

Our effects of disruption of the right DLPFC do not support a self-control account for this region in intertemporal choice. A self-control account could predict that a region involved in self-control would show more involvement as the motivational drive to choose an immediate gain (or to avoid choosing an immediate loss) becomes more prepotent, under the assumption that it should be harder to override this prepotent urge as it increases. Since this prepotent urge should increase as the value of the immediate option increases, this idea holds that disruption of a region involved in self-control would lead to greater choice of the best immediate option as the value of the immediate option increases and becomes closer to the monetary value of the larger delayed option. Alternatively, a self-control account could predict that a region involved in self-control would show more involvement as intertemporal choices become more difficult, that is as the values of the two options approach each other, under the assumption that it should be harder to override a prepotent urge to choose an immediate gain or to avoid choosing an immediate loss when the two options are close in subjective value. This idea holds that disruption of a region involved in self-control would lead to greater choice of options with the better immediate value as the subjective values of the two options approached each other. However, disruption of the right DLPFC did not lead to either of the changes in choice that might be expected to occur after disruption of a self-control process. As the value of the immediate option increased, individuals who received real rTMS did not become increasingly more likely than those who received sham stimulation to choose the option with the better immediate value, as demonstrated by the non-significant interactions between Stimulation Type and Immediate Relative Value. Similarly, supplementary analyses revealed that individuals who received real rTMS to the right

DLPFC did not become increasingly more likely than those who received sham stimulation to choose the better immediate option as the subjective values of the two options approached each other; there were non-significant interactions between Stimulation Type and the quadratic effect of Immediate Relative Value (i.e. Immediate Relative Value^2). Rather, the effect of rTMS appeared consistent across the range of values studied. Thus, the right DLPFC data appear most consistent with a dual-valuation account.

The effects seen after disruption of the right PPC, however, are more consistent with a self-control account than a dual-valuation account. Disruption of the right PPC led to changes in choice that were dependent on the relative values of the two options. Importantly, when the value of the immediate option was low, disruption of the right PPC did not lead to increased choice of the option with the better immediate value. As the value of the immediate option increased, however, individuals who received real rTMS became increasingly more likely than those who received sham stimulation to choose the option with the better immediate value. In other words, following disruption of the right PPC, individuals became increasingly less likely than those in the sham group to choose the option with the better long-term, but worse immediate value as the value of the immediate option increased. This suggests that the right PPC helps individuals override or inhibit prepotent drives to select options with better immediate value so that options with better long-term value can be selected. Supplementary analyses revealed that individuals who received real rTMS to the right PPC did not become increasingly more likely than those who received sham stimulation to choose the better immediate option as the subjective values of the two options approached each other, as demonstrated by the

non-significant interactions between Stimulation Type and the quadratic effect of Immediate Relative Value. Together, this suggests that cognitive control functions of the right PPC play an increasingly important role in helping individuals choose the best long-term option as the motivational drive to choose an immediate gain (or to avoid choosing an immediate loss) increases, rather than as the subjective values of the two options approach each other. Thus, the data suggest that the right PPC helps individuals choose in line with their goals when there is a conflict between those goals and motivational drives.

Although we have suggested that PPC self-control functions help a person choose the option with a better long-term value as the value of an immediate option increases, a person might not want to engage these self-control processes to choose the best long-term option when the value of the immediate option is sufficiently high. It is unlikely, however, that the values of our immediate choice options were high enough to discourage a person from wanting to use self-control. Across subjects, the average percent difference between immediate and delayed monetary values when the Immediate Relative Value was at its highest were within a range that others have shown is not associated with an exclusive bias for choosing either better immediate or better long-term intertemporal choice options (i.e. percent difference in monetary value between 4% and 34%) (McClure et al., 2004).

This is the first study to show that disruption of the DLPFC or the PPC has similar effects on both intertemporal choices involving gains and losses. We found that disruption of either of these regions in the right hemisphere led to greater choice of options that were associated with a better immediate value but a worse long-term value

than alternatives. Disruption of the right DLPFC led to a general shift in preference (towards immediate gains and away from immediate losses) overall. In contrast, disruption of the right PPC led to changes in preference that depended on the value of the immediate alternative, such that there was a greater tendency to choose immediate gains and to not choose immediate losses as the values of the immediate options increased (i.e. as gains and losses became more positive and more negative, respectively). This has important implications given that past theories of motivation have postulated different neural circuits for motivations to approach rewards and motivations to avoid aversive outcomes (Gray, 1981; Panksepp et al., 2002). Contrary to the predictions of these theories, our results reveal that disruption of the same brain regions leads to a greater tendency to both approach immediate rewards and avoid immediate losses. Our findings show that the right DLPFC and PPC are involved in optimization of long-term vs. immediate value in general during choice, rather than being specific to one motivational system.

An important question is how the right DLPFC and right PPC interact with other brain areas involved in the decision making process. For instance, the DLPFC and PPC could exert an influence on choice by modulating activity in other regions responsive to value. Evidence supports this conjecture for the DLPFC. When individuals make choices for food based on attributes associated with better long-term value (i.e. health) rather than better immediate value (taste), activation in the DLPFC is functionally connected with activation in the VMPFC, an area with activation related to subjective valuation of food items on the task (Hare et al., 2009). This suggests that the DLPFC lends more weight to the value of delayed options in the decision process by modulating

value signals in other brain regions. The PPC might also modulate value signals on choice tasks, but this has not previously been investigated. Another possibility is that activity in the DLPFC and PPC more directly affect choice, by overriding activity in other regions that strongly favors immediate value. There is also a third possibility, that the DLPFC and PPC could affect choice by affecting processes prior to valuation. The PPC could engage controlled attentional processes that lead to greater focus on delayed alternatives, or on alternatives with the highest monetary value, and such a process might be engaged more as the smaller immediate options become more distracting (i.e. larger in value). Previous research has linked the PPC to controlled attention, making this idea plausible (Corbetta et al., 2000; Corbetta & Shulman, 2002; Friedrich et al., 1998). Future fMRI work that examines the time courses of activation in different brain regions during different components of the intertemporal choice process (i.e. as a choice unfolds and is made) and that looks at functional connectivity between regions will help differentiate between these alternative accounts.

While we have postulated a different role for the right PPC and right DLPFC based on our results, one finding suggests an alternative possibility that the right PPC may be involved in a similar valuation process as the right DLPFC during intertemporal choice. A close look at the between subject loss graphs in Figure 5B and 5D reveals that the interaction involving losses following stimulation to the right PPC actually may have been due to differences in the sham group rather than in the real rTMS group, as demonstrated by the shallower slope of the loss line following right DLPFC sham stimulation than following right PPC sham stimulation. We note that if the right PPC between subject loss model is rerun without an interaction term, the main effect of

Stimulation Type is significant ($p < .05$), and in a similar direction as the main effect of Stimulation Type in the right DLPFC loss model. In contrast, the interaction following stimulation to the right PPC on gain trials does not appear to have been driven by changes in the sham group. This suggests that self-control rather than valuation was affected following disruption of the right PPC on gain trials. Given the parallel interactions between Stimulation Type and Immediate Relative Value seen following stimulation to the right PPC on both gain and loss trials, we believe the data more strongly support a role for the right PPC in self-control than in valuation on loss trials as well. However, it is unclear why individuals performed differently on loss trials following sham stimulation to the right DLPFC and right PPC. It does suggest that preferences on choice tasks may be somewhat unstable over time. This is a possibility given that individuals received stimulation to each region on different days.

Although disruption of right hemisphere regions led to greater choice of options with a better immediate value, it did not make subjects insensitive to the relative subjective values of the two options. Both subjects who received real rTMS and those who received sham rTMS to regions of the right hemisphere were more likely to choose the immediate gain and less likely to choose the immediate loss as the relative subjective value of the immediate option increased. Interestingly, it was the sham subjects rather than the real rTMS subjects who had choices that were more inconsistent with their prior preferences. When the relative value of the immediate option was equal to that of the delayed option, sham subjects were more likely to make a choice for the option with the better long-term value than for the option with the better immediate value, even though they were predicted to choose each option with equal frequency. In contrast, at the same

Immediate Relative Value, subjects who received real rTMS to right hemisphere regions had choice patterns closer to chance level.

A greater tendency to make choices for the option with the better long-term value following sham stimulation may have occurred because of the different structures of our intertemporal choice task and our indifference point task. The numerical comparison portion of our intertemporal choice task may have led individuals to focus more on the magnitude of the two monetary options rather than the available time of the two options on the Intertemporal Choice Task, since the numerical comparison task required individuals to compare the two monetary values but not the times of payment. If subjects indeed focused more on magnitude, they may have exhibited less delay discounting than they did on the Indifference Point Task, which did not have a numerical comparison component. This in turn would have led to more choices for the option with the better long-term value than were predicted in the sham group, as we observed. Prior research has shown that how delays are framed (i.e. date vs. delay) affects delay discounting (Read, Frederick, Orsel, & Rahman, 2005); our results raise the possibility that whether an individual attends to reward magnitudes vs. delays affects delay discounting as well.

It is important to note that we found a lack of coupling between changes in choice and the ability to compare numerical values following disruption of right hemisphere regions. Although previous research has shown slowed reaction time on numerical comparison tasks following disruption of the left PPC (Andres et al., 2005; Cappelletti et al., 2007)**,** we did not observe altered numerical comparison reaction time following disruption of either side of the PPC. Our lack of differences may have been due to an inability to localize the specific region of the IPS important for numerical comparisons,

or may have been due to differences in our numerical comparison task from those used in other studies. Unlike typical numerical comparison tasks, subjects in our study did not all receive the same stimuli. This was done to ensure that individuals compared numbers in our task that had been in the previous intertemporal choice, but since subjects had already seen the numbers before the beginning of the numerical comparison, our task may have been easier than typical tasks. Nevertheless, we found no support that altered intertemporal choice following disruption of regions of the DLPFC and PPC was due to an impaired ability to compare the values of the two alternatives, suggesting that neither of these brain regions may be necessary for numerical comparisons in intertemporal choice. However, we note that this does not rule out a role for these regions in more complex value comparisons. The numerical comparisons that accompanied our numerical comparison task may have been less demanding than the comparisons that are required to make an intertemporal choice. While the numerical comparison task required comparison of two values (i.e. two magnitudes), the intertemporal choice task required comparison of 4 (i.e. two magnitudes and two time points).

In contrast to our finding of increased choice for the better immediate option following disruption of the right PPC in both the gain and loss domains, we found that disruption of the left PPC led to decreased choice of the better immediate option, but only involving gains. Also unlike the effects seen following disruption of the right PPC, disruption of the left PPC led to a general shift in the choice function, rather than changes that interacted with the relative values of the two options. The increased tendency to choose delayed monetary gains following disruption of the left PPC was coupled with decreased reaction time, suggesting that when the left PPC was disrupted, individuals

were able to more efficiently choose delayed gains. One possible explanation for the different pattern of results following disruption of the left vs. the right PPC is that the functions of the left PPC may normally inhibit self-control functions of the contralateral right PPC during intertemporal choice. If so, then disrupting the left PPC could lead to greater activity in the right PPC and increase the likelihood and speed of choosing better long-term options. Why individuals should be more likely to choose larger delayed monetary rewards after disruption of the left PPC, but have unaltered choices involving losses is unclear. The inability to find an effect in the loss domain may have been because the effects of disrupting the left PPC on choice were modest. The magnitude of the effect on choice following disruption of the left PPC was smaller than those seen following disruption of right hemisphere frontal and parietal regions. Because of the small magnitude of this effect, it is not discussed further.

Our reaction time measures revealed that individuals were sensitive to both the relative subjective values and the objective monetary values of the two options during intertemporal choice. Reaction time slowed as both the subjective values and the monetary values approached each other, revealing that intertemporal choices are not only more difficult when the subjective values of two options are close, but also when the monetary values are close. We found slower choice reaction time following disruption of right hemisphere regions in the loss, but not the gain domain. This effect was greatest when the subjective values of the two options were close. Given that this finding only emerged in the loss domain and emerged as the choices became more difficult suggests that intertemporal choices involving losses are more difficult than those involving gains. While the choice itself may be more difficult, we did not find evidence that the ability to

choose the better long-term option was harder when choices involved losses. In fact, it appeared to be easier for individuals to make such choices for losses than for gains. As demonstrated by the right hemisphere between subject choice graphs in figure 5(A-D), when the relative values of the two options were equal, both sham and real rTMS groups were more likely to make choices for the better long-term outcome in the loss domain (i.e. choose immediate loss) than they were in the gain domain (i.e. choose delayed gain). Past evidence shows that individuals are less likely to discount the value of delayed losses than gains (Baker et al., 2003; Benzion et al., 1989; Estle et al., 2006; Frederick et al., 2004; Murphy et al., 2001). Our findings extend this by suggesting that even when the subjective values of delayed and immediate options are equal, it may be easier to choose the option with the better long-term value in the loss domain than in the gain domain.

Reaction time would be expected to be slower when individuals use self-control processes during choice than when they do not, since it should take more time to make decisions when using deliberative cognitive control processes than when solely using more automatic valuation processes. This is because that performance on tasks is thought to take longer when it relies on controlled rather than on automatic processes (Cohen, Dunbar, & McClelland, 1990). If the right PPC is involved in self-control processes during intertemporal choice, then one might predict that disruption of it would affect both choice and reaction time more as the relative value of an immediate option increases. The reason for this is that self-control processes that are used to help one choose the best long-term option should be engaged more as the relative value of an immediate option increases. If these self-control processes are taken offline following real rTMS to the

right PPC, then individuals who receive real rTMS would be expected to have increasingly faster reaction times than sham subjects as the relative value of the immediate option increases. We, however, did not find that reaction times were faster following real than following sham rTMS to the right PPC. Rather, reaction time increased following real rTMS, but only on loss trials. These reaction time findings thus are not consistent with our suggestion that the right PPC is involved in self-control during choice. However, increases in reaction time following real rTMS may have occurred because other processes in addition to self-control processes were disrupted. Our intertemporal choice task required subjects to switch between performing different sub-tasks: choice amongst gains, choice amongst losses, and numerical comparisons. Disruption of the right PPC may have led to increased reaction time because it disrupted the ability to switch between different tasks, a proposal that is plausible given prior findings linking the right PPC to task-switching (Brass & von Cramon, 2004). Self-control processes may have been disrupted following real rTMS as well, but reaction time may have been more affected by disruption of task-switching processes than by disruption of self-control processes.

While previous findings show that disruption of the left DLPFC leads to increased choice for small immediate over larger delayed rewards (Figner et al., 2010), we found that disruption of the right, but not the left DLPFC, had this effect. Although our lateralized effect is inconsistent with the rTMS study by Figner and colleagues (2010), it is consistent with prior fMRI work. Intertemporal choice studies often show task-related activation in the right DLPFC, or nearby lateral PFC regions, as well as or instead of activations in the left hemisphere (McClure et al., 2007; McClure et al., 2004; Xu et al.,

2009). PPC activations are also often seen in the right hemisphere or bilaterally on intertemporal choice tasks (McClure et al., 2007; McClure et al., 2004; Xu et al., 2009). Taken together, this reveals the possibility that neural activity in both the right and left DLPFC (and PPC) may increase the chances of choosing the best long-term option during intertemporal choice. However, the relative importance of each hemisphere may depend on the characteristics of the intertemporal choice task being performed. Our task was most likely more difficult than that employed by Figner and colleagues. While their subjects only performed intertemporal choices for monetary gains, our subjects performed randomly mixed trials of choices involving monetary gains and losses. Furthermore, on each trial of our task there were two sub-tasks, an intertemporal choice and a numerical comparison, further increasing difficulty since subjects had to switch from doing one task to another. This raises the possibility that the right DLPFC is more involved in making intertemporal choices in difficult situations. This possibility may be further supported by our inability to find differences on our indifference point task following disruption of the right DLPFC or right PPC, which was probably easier than the Intertemporal Choice Task, since it did not mix gain and loss trials and did not contain a numerical comparison task. Given that our results are based on a larger number of trials than those employed by Figner and colleagues, and the fact that we saw similar effects in both the gain and loss domain, our results provide strong evidence that the right DLPFC is more crucial than the left in the ability to choose the better long-term, but worse immediate, incentive.

In conclusion, we found that disruption of either the right DLPFC or the right PPC with low frequency rTMS led to greater choice for monetary options with a better

immediate, but worse long-term value in both the gain and loss domain. This reveals that activity in both regions helps individuals make choices that optimize long-term positive and negative value in relation to immediate value. However, our findings also suggest that the specific roles of the DLPFC and PPC are different in intertemporal choice. Effects seen following disruption of the right DLPFC suggest it is involved in valuation of items, and that valuation processes of the DLPFC put less weight on the value of immediate options than do valuation functions in other brain regions. In contrast, effects seen following disruption of the right PPC suggest that is involved in cognitive control functions that help individuals make better long-term choices, rather than valuation functions, during intertemporal choice. Our findings help clarify the roles of the DLPFC and PPC during intertemporal choice, and indicate that the functions performed by these regions do not depend on whether choices involve positive or negative incentives.

CHAPTER III

EXPERIMENT 2 – CHANGES IN RISKY CHOICE FOLLOWING DISRUPTION OF THE DLPFC

Introduction

Many economic choices that people make are between options with varying levels of risk. For instance, when deciding whether to buy a stock or a bond or when deciding whether to buy a lottery ticket or save one's money, a person is confronted with options that have different levels of risk. Such situations involve deciding whether or not to select an option that could lead to a better outcome, but that also has a higher probability of leading to a worse outcome than the alternative option. Although others have defined risk level in different ways (Rothschild & Stiglitz, 1970; Weber & Johnson, 2009), here we define it as a direct function of the probability of outcomes on a choice task. Risk level can be said to increase as the probability that an option leads to a rewarding outcome decreases, or conversely, risk level can be said to increase as the probability that an option leads to a negative, or non-rewarding, outcome increases. Even though there are limitations to the choice scenarios this definition can be applied to, we think that this definition of risk is useful because it is able to directly capture the effects of different probabilities of gain or loss on choice.

Several pieces of evidence suggest that in choices involving risk, the right DLPFC has an important role in helping individuals avoid taking risks. In an important study, Knoch and colleagues (2006) demonstrated that disruption of the right DLPFC with low

frequency rTMS leads to greater selection of a riskier, but potentially larger reward option over a less risky reward option on "Roger's Risk Task" (Rogers et al., 1999), in which individuals make choices between two options that differ in terms of probability and magnitude of reward and loss. Evidence that increasing activity in the right relative to the left DLPFC with tDCS leads to fewer choices for the riskier option on Roger's Risk Task is consistent with the rTMS findings (Fecteau et al., 2007).

Several other findings suggest that the left DLPFC also has an important role in choice involving risk. Matthews and colleagues found that fMRI BOLD signal in the left DLPFC, rather than the right DLPFC, predicted choice of a smaller certain over a larger risky reward (Matthews et al., 2004). Additionally, both increasing activity in the left relative to the right DLPFC, or increasing activity in the right relative to the left DLPFC with tDCS has led to greater choices for the riskier option on Roger's Risk Task in older adults (Boggio et al., 2010). The divergent results across studies, however, leave unclear whether the role of the left DLPFC is different from that of the right DLPFC in choices involving risk.

An important question that has not previously been adequately addressed is whether disruption of the DLPFC only leads to a greater preference for risk in situations that have specific levels of risk. Disruption of the DLPFC could affect risk preferences in general, shifting preferences towards greater preference for choosing risky options. Alternatively, shifts in preference could interact with the level of risk. An increased preference for risk might occur at different levels of risk but be greater at higher than at lower levels of risk. This could occur if the processes in the DLPFC that increase risk-aversion are engaged for different levels of risk but become engaged to a greater extent as

the risk level increases. While some evidence suggests that neuromodulation of the DLPFC can affect the relationship between the level of risk and choice (Fecteau et al., 2007), how the risk level of various options affects risk preference has not been reported.

Moreover, it is unclear whether disrupting activity in the DLPFC would lead to greater risk taking on tasks that do not involve the possibility of a loss, as prior studies have utilized Roger's Risk Task which offers the possibility of a potential loss in addition to the possibility of a potential gain on every trial. This is an important question given that individuals' propensities to take risks are different when losses and gains are involved (Kahneman & Tversky, 1979; Tversky & Kahneman, 1981, 1992). It may be that different brain regions bias behavior towards risk-aversion when risky options can lead to losses than when they cannot. If so, the DLPFC might only bias behavior towards risk-aversion when there is a possibility of a loss. Alternatively, the DLPFC may bias behavior towards risk-aversion when there is a possibility of a loss and when there is not. This possibility would be supported if disrupting activity in the DLPFC was shown to lead to greater risk taking in choices that do not involve the possibility of a loss.

To examine whether activity in the DLPFC underlies risk-aversion for choices involving monetary gains, but not losses, with differing levels of risk, we utilized low frequency rTMS to transiently suppress functioning in the left or right DLPFC while individuals performed a risky choice task that involved choices between a certain option and a risky option that offered a larger reward with one of two probabilities. By only varying the probability of one option, our design allowed us to more precisely examine how the probability of a risky option affects choices than have prior neuromodulation studies. We varied the relative subjective values of the options on the choice task

according to subject's pre-stimulation preferences and predicted that individuals would become more risk seeking following disruption of either side of the DLPFC. There were two competing hypotheses regarding how increases in risk seeking would interact with risk level. If functions of the DLPFC lead to risk-aversion in general, then increased risk seeking following disruption of this region should be similar under different levels of risk. Alternatively, if the functions of the DLPFC are more likely to lead to risk aversion under higher levels of risk, then there should be a greater increase in risk seeking at higher than at lower levels of risk, following disruption of this region.

Methods

Participants

30 right-handed individuals (46.67% female) between the ages of 18 and 29 (*M* Age = 21.17, SD = 3.10) from Vanderbilt University and the Nashville community participated in this study. All of these participants reported having no history of neurological or psychiatric problems, and no females were currently pregnant. 55 additional subjects were consented but excluded or withdrew. Reasons subjects were excluded included baseline risk preferences as revealed in Session 1 that did not allow us to make task stimuli (see criteria below), having previously received TMS, risk factors that could increase the chances of having negative effects from TMS (e.g. neurological conditions), and experimenter error. All participants completed written informed consent approved by the Vanderbilt IRB.

Session One – Indifference Point Procedure

We screened subjects prior to consent with a questionnaire to ensure that we only recruited participants who were risk averse or risk neutral. There were 6 questions, each which required participants to write the magnitude of a hypothetical monetary gain available with 100% probability that they valued equally to a risky hypothetical monetary gain available with a less than certain probability. On half of the questions the risky option was available with 70% probability and on the other half it was available with 90% probability. Any subjects who were risk seeking (i.e. if any of the certain values were greater than the expected value of the matched risky gain) were excluded. This was done to increase power because if the DLPFC modulates risk aversion and prevents risk-seeking behavior, disruption of this region may not have effects or may have smaller effects on individuals who are already risk-seeking.

Following consent, we determined subject indifference points between certain and risky monetary rewards. The risky rewards offered a reward with a probability of 70% or 90%, and offered no reward with the opposing probability (i.e. 100 - probability). Subjects completed an indifference point task on a computer in which we determined 8 certainty equivalents (i.e. dollar amount of a certain monetary reward that the person values equally to a risky monetary reward) for 8 different risky monetary gains: 4 different magnitudes (i.e. $2.50, $5.00, $7.50, and $10.00) available with one of two probabilities (i.e. 70% or 90%). The order of determining certainty equivalents was from low to high magnitudes, and the lower probability risky option was always determined first (i.e. order was $2.50 with 70%, $2.50 with 90%.... $10 with 90%). Subjects were told that one of their choices would randomly be selected for payment. If the selected

choice was certain, they would definitely receive the money, but if it was risky, they would have a chance to receive it (associated with the probability of the option). For timing of one trial of this task and presentation of items, see Figure 9.

To determine each certainty equivalent, participants were presented with an initial trial in which they made a choice between a risky monetary gain available with 70% or 90% probability and a certain option of half the maximum value of the risky option. If participants chose the risky option, the value of the certain option increased by half, and if they chose the certain option it decreased by half. On the next trial, the certain value changed in a similar way as in the previous trial but only by one quarter of the original value. Over six trials, the certain value increased or decreased by progressively smaller amounts (i.e. by 1/(2^x) where x was trial number) depending on participant responses so that the subjective value of the certain amount would iteratively approach that of the risky amount. After the sixth trial, a final catch trial was presented in which the certain value was higher than the just-calculated certainty equivalent. This provided a check to ensure that subjects were answering according to their preferences (i.e. were answering consistently). If they did not choose the certain value on the catch trial, the certainty equivalent for that specific risky magnitude reward was determined again (i.e. the six trial procedure and check was repeated). After answering consistently, or after completing the indifference point procedure three times, participants then performed the indifference point procedure for the next probability-magnitude pair. If participants answered inconsistently for three indifference point procedures in a row for a given probability-magnitude pair, they were excluded.

Figure 9: Trial Structure of the Risky Choice task*. A. Decision Phase. Following 500 ms of fixation, two different money/probability pairs were presented on a computer screen (side randomized). The pairs always consisted of a smaller amount of money to be gained with 100% probability (written as "For Sure") and a larger amount of money to be gained with either 70% or 90% probability. The words "Today" indicated that subjects were making choices for options to be paid that day. Subjects decided which money/probability pair they would prefer by pressing the "z" or "m" key and had up to 6500 ms to respond. B. Post Choice Phase. Immediately after responding, the triangle under the subject's choice turned red and the other triangle disappeared to indicate subjects' choice; this was displayed for 250 ms. If subjects did not respond, this post choice phase was skipped, and the next trial began after 6500 ms. (* The timing and display of items in the Indifference Point Task were identical to those of items in the Risky Choice Task except there was no time limit for response).

Subjects whose certainty equivalents indicated they were risk-seeking were excluded from further sessions (i.e. they were excluded if any certainty equivalent was greater than the expected value of the paired risky option). Additionally, subjects were excluded if any certainty equivalent was smaller than 10 cents. This was done to ensure that there would be an adequate range of values below each subject's certainty equivalents to create choice stimuli for further sessions.

At the end of session one, one random trial was selected for payout. Subjects then received the amount of money associated with their choice on that trial. However, if the chosen option was for a risky selection, subjects rolled a 10 sided die to determine whether they would be paid. If they rolled a number less than or equal to the tens digit of

the probability listed on that trial (i.e. less than 7 for 70%), then they were paid the chosen amount.

Sessions Two and Three

Participants that met eligibility requirements after session one were equally divided into 2 groups of 15 subjects. Group one received rTMS to the left DLPFC in one session and to the right DLPFC in a separate session (with counterbalanced ordering). Group two received sham rTMS to the same areas as group one (half left DLPFC first/half right DLPFC first). Each of the two groups had the same gender distribution (8 Males and 7 females). Except for the different brain regions stimulated across sessions, procedures in sessions 2 and 3 were identical.

Prior to receiving rTMS, participants performed shortened practice versions of the tasks, in order to make sure they understood the tasks. Subjects were told that one choice from the tasks performed after stimulation would be randomly selected for payment at the end of each session, and that they would receive that payment with the associated probability.

After completing the practice tasks, participants received real or sham rTMS stimulation for 30 minutes to either the left DLPFC or right DLPFC (details outlined under TMS Methods section).

Risky Choice Task

Immediately following completion of stimulation, participants completed a risky choice task on a computer. On each trial, participants made a choice between a certain monetary gain and a risky monetary gain of potentially larger magnitude (See Figure 9).

The values of the risky options were the same as in session 1 ($2.50, $5.00, $7.50, and $10.00). Half of the risky options were available with 70% probability and half were available with 90% probability. The magnitude of each certain option was calculated as a specific percentage difference from that subject's session one certainty equivalent for the risky option on that trial. This allowed us to vary the subjective value of each certain option (as revealed in session one) relative to the associated risky option, which let us predict which option an individual should choose on each trial (See Table 8). There were 72 of these trials on the Risky Choice Task.

There were also 40 additional trials on the Risky Choice Task with different stimulus properties than those just discussed. On 8 of the additional trials, the risky option had an equal expected value to that of the certain option, and on the other 32 additional trials, the risky option had a lower expected value than did the certain option. On trials with two options of equal expected value, the mean proportion of choices for the risky option were .283 (SD = .260) for the real rTMS group and .130 (SD = .193) for the sham rTMS group. On trials that contained a risky option that had a lower expected value than the certain option, the mean proportion of choices for the risky option were .138 (SD = .153) for the real rTMS group and .075 (SD = .122) for the sham rTMS group. These additional trials are not discussed further, as they were not created to test our primary hypotheses.

Indifference Point Task

After completing all of the aforementioned trials, subjects were given the identical indifference point procedure as in session one. However, this time the iterative procedure was only done to calculate the certainty equivalent for risky values twice, first for a

$10.00 gain available with 70% probability and then for a $10.00 gain available with 90% probability. Additionally, no consistency checks were performed in order to limit the time to complete the task. Collecting subject indifference points post-stimulation provided an added check of whether real rTMS to the DLPFC increased risk seeking.

Table 8: Certain Values for Different Trials in the Risky Choice Task.

Certain Relative Value	Total	Description
1. Certain Subjective Value Lower than Risky Subjective Value	32	Magnitude of certain option is below certainty equivalent (i.e. certain value at indifference point). There were 4 different percent distances below the certainty equivalent (-10%, -20%, -35%, and -50%). Subjects who receive sham stimulation should choose the risky option.
2. Certain Subjective Value Equal to Risky Subjective Value	8	Magnitude of certain option is equal to certainty equivalent (i.e. percent distance equal 0%). Neither certain nor risky choice is predicted for the sham group.
3. Certain Subjective Value Higher than Risky Subjective Value	32	Magnitude of certain option is above certainty equivalent. There were 4 different percent distances above the certainty equivalent (10%, 20%, 35%, and 50%). Subjects who receive sham stimulation should choose the certain option.

Each 70% and each 90% risky option magnitude was shown nine times. Each risky option was paired once with a certain option located at each of the 9 percent distances from the certainty equivalent of the risky value. Values at negative distances were taken as a function of the percentage difference between the certainty equivalent and zero. A -35% distance indicated that the certain value was 35% less than the certainty equivalent. Thus if the certainty equivalent of a $2.50 risky option available with 70% probability was $1.50, the certain value was $.98 (i.e. 1.50 - .35*1.50 = .98). Values at positive distances were calculated as a function of the percentage difference between the certainty equivalent and the risky value. A 35% distance indicated that the certain value was greater than the certainty equivalent of the risky option by 35% of the distance between the certainty equivalent and the risky value. Thus if the certainty equivalent of a $2.50 risky option available with 70% probability was $1.50, the certain value was $1.85 (i.e. 1.50 + (.35 * (2.50 -1.50)) = 1.85).

At the end of sessions two and three, one random trial was selected for payout. Subjects then received the amount of money associated with their choice on that trial.

However, if the chosen option was for a risky selection, subjects rolled a 10 sided die to determine whether they would be paid. If they rolled a number less than or equal to the tens digit of the probability listed on that trial (i.e. less than 7 for 70%), then they were paid the chosen amount.

TMS Methods

Low frequency (1 Hz) rTMS was delivered with a MagStim TMS double 70 mm (Figure 8) coil (Magstim, Wales, UK) at 54% power for 30 minutes. The rTMS parameters used were within currently recommended guidelines (Rossi et al., 2009) and stimulation with these parameters leads to suppression of excitability in the targeted region for a period of time following stimulation (Robertson et al., 2003). We note that 1 Hz rTMS applied to the DLPFC at 54% power has previously led to altered decision making following stimulation (Figner et al., 2010). Sham stimulation was delivered with a MagStim placebo coil, which produced clicks that resembled the sound of rTMS, but without a magnetic pulse. Given the potential for subjects to identify whether they received sham or real stimulation, it was prudent for subjects to receive only one of these types of stimulation. Subjects were blind to the type of stimulation (e.g. sham vs. real) they received.

Positioning of the TMS coil was accomplished by using the 10-20 EEG System, which has previously been used to deliver TMS to identified brain regions with reasonable structural accuracy (Herwig et al., 2003). This method was used to ensure proper positioning of the coil over the DLPFC (BA 46/9). Specifically, to localize the DLPFC, the center of the coil was held tangentially to the participant's head with the

handle pointing caudally, and placed one cm antero-lateral to F3/F4 which has been suggested to provide better coverage over BA 9/46 than do the points F3/F4 (Herwig et al., 2003). These specific points for stimulation were marked on a lycra swimcap that subjects wore; the position of the swimcap was placed in a consistent position on the head across subjects, by using the nasion, inion, and preauricular points as physical landmarks for placement. During stimulation, all participants wore earplugs as protection against the noise of the rTMS clicking. Participants maintained their head position during rTMS administration using a chin/head rest and were visually monitored to ensure that no movement had occurred. In cases where a participant moved their head, the coil was immediately repositioned over the target.

All individuals received either real or sham rTMS to the left or right DLPFC in session two for 30 minutes and the same type of stimulation to the DLPFC on the contralateral side for 30 minutes in session three. The order of stimulating the two sides of the DLPFC was counterbalanced across subjects in each group. We used an "offline" rTMS paradigm; subjects completed tasks after stimulation was completed. Since impairments in behavior following low frequency rTMS have been shown to last for half the time of the previous stimulation (Mottaghy, Gangitano et al., 2002), the tasks following rTMS were limited to the first 15 minutes post-stimulation.

Statistical Analysis Methods

Using PASW Statistics 18 (SPSS Inc., Chicago, IL), data were analyzed with Generalized Estimating equations (GEE) which allow one to model effects while accounting for correlations within observations of individual subjects (Liang & Zeger,

1986). All models contained an intercept term and an unstructured working correlation matrix.

In order to examine whether subjects chose differently under conditions of differing risk, separate within-subject GEE models with a logit link function and binomial distribution were created for subjects who received real stimulation and for subjects who received sham stimulation. These models were created to predict choice of a risky option (0 = Chose Certain, 1 = Chose Risky) at the trial level on the Risky Choice Task. Independent variables in both models were Probability of the risky option (0 = 70%, 1 = 90%), Certain Relative Value (i.e. percent distance of the certain option from the certainty equivalent of the risky option, which ranged from -50% to 50%), and Session Number (0 = Session 2, 1 = Session 3).

We also used GEE models with a logit link function and binomial distribution to see how choice of risky options that had a specific probability of reward were affected by stimulation type and stimulation side. Separate models were created to predict choice of the risky option on 70% risky probability trials and 90% risky probability trials. In order to test the effects of real stimulation relative to sham, separate GEE models were created for responses following stimulation to each hemisphere (i.e. left DLPFC, right DLPFC). Independent variables in all models included Stimulation Type (0 = Sham rTMS, 1 = Real rTMS), Certain Relative Value, and Session Number. Additionally, we included the interaction of Certain Relative Value by Stimulation Type if it was significant ($p < .05$). To examine within subject differences in the real rTMS group, we created similar GEE models, except the within subject variable Side (0 = Left, 1 = Right) was used instead of

the between subject variable Stimulation Type; the interaction of Side by Certain Relative Value was only included if significant.

To predict reaction time on trials of the Risky Choice Task (measured in milliseconds (ms)), we also created separate GEE models for 70% risky probability trials and 90% risky probability trials. These models contained an identity link function and normal distribution. Between subject models (each limited to one hemisphere) and within subjects models included the same independent variables as the choice models, plus one additional variable: Certain Relative Value^2 (i.e. Certain Relative Value squared). No interactions were included in the models.

To predict effects of real relative to sham stimulation on certainty equivalents measured in the indifference point procedure, we utilized the Generalized Linear Model function in PASW Statistics 18 (SPSS Inc., Chicago, IL) with an identity link function and normal distribution. Each model examined the effects of stimulation in one hemisphere for either 70% risky probability or 90% risky probability certainty equivalents and contained the following predictors: Stimulation Type, Session Number, and Session 1 Certainty Equivalent. This last variable was included as a covariate because certain values at the indifference point (i.e. certainty equivalents) following stimulation should be related to subject's initial pre-rTMS (i.e. Session 1) values, and thus significant effects of Stimulation Type might not emerge without controlling for this variable. To predict certainty equivalents across region for subjects who received real rTMS, we used GEE models with an identity link function and normal distribution that included Side (-1 = neither (i.e. for Session 1), 0 = Left DLPFC, 1 = Right DLPFC) and Session Number (-1 = Session 1, 0 = Session 2, 1 = Session 3) as predictors. Separate

models were created for 70% risky probability and 90% risky probability certainty equivalents.

Results

Session One Indifference Points

Using independent sample t-tests, we compared mean certainty equivalents measured in session 1 on the Indifference Point Task for individuals in the real rTMS group with individuals in the sham rTMS group. This was done to ensure that session 1 risk preferences across stimulation groups were not different. Eight t-tests were performed in all; each t-test compared mean certainty equivalents across stimulation groups for one of the eight types of risky gains (i.e. 4 magnitudes of risky gains available with 70% probability and 4 magnitudes of risky gains available with 90% probability). There were no significant differences in mean session 1 certainty equivalents between real and sham rTMS groups. Importantly, these results reveal that prior to receiving rTMS (i.e. in session 1), mean levels of probability discounting were similar across real and matched sham rTMS groups. For mean session 1 certainty equivalents, see Table 9.

Effects of Risk Level on Choice

For descriptive statistics of overall mean proportion of choices for risky options following each type of stimulation to each side of the DLPFC, see Table 10.

Table 9: Mean Session One Certainty Equivalents in the Indifference Point Task.

	Real rTMS		Sham rTMS	
	Mean	(SD)	Mean	(SD)
$2.50 with 70%	1.275	(.494)	1.375	(.394)
$5.00 with 70%	2.597	(.904)	2.715	(.992)
$7.50 with 70%	3.892	(1.352)	3.692	(1.447)
$10.00 with 70%	4.963	(1.792)	4.987	(2.032)
$2.50 with 90%	1.907	(.477)	1.818	(.451)
$5.00 with 90%	3.511	(1.190)	3.610	(.974)
$7.50 with 90%	5.139	(1.906)	5.224	(1.513)
$10.00 with 90%	6.686	(2.429)	7.350	(1.842)

Mean session 1 certainty equivalents for specific magnitudes of risky gains available with 70% or 90% probability (listed in left column). Averages are constructed separately for subjects in the real and sham rTMS groups (listed in column titles).

Table 10: Mean Choice of Risky Options in the Risky Choice Task.

	Left DLPFC				Right DLPFC			
	70% Probability		90% Probability		70% Probability		90% Probability	
	Mean	(SD)	Mean	(SD)	Mean	(SD)	Mean	(SD)
Real	.665	(.284)	.625	(.245)	.601	(.265)	.566	(.259)
Sham	.381	(.222)	.563	(.217)	.368	(.254)	.504	(.226)

Mean proportion of choices for the risky option across subjects computed for each type of stimulation (i.e. Real vs. Sham) to each side of the DLPFC (side listed in column title). Statistics are collapsed separately for trials with a different level of risk (i.e. risky option available with 70% Probability or 90% Probability).

We first tested whether the tendency to choose the risky option was dependent upon the risk level (i.e. probability) of this option. As this relationship could differ following disruption of the DLPFC, separate within subject models were created for subjects who received real and sham rTMS. Each model predicted choice of a risky reward as a function of Probability of the risky option, Certain Relative Value, and Session Number. For a list of model parameters, see Table 11.

Table 11: GEE Models for Risky Choice as a Function of the Probability of the Risky Option.

	Real Within			Sham Within		
	B	(SE)	OR	B	(SE)	OR
Intercept	1.079***	(.130)	2.941	-.922***	(.074)	.398
Cert. Rel. Value	-.028***	(.004)	.973	-.036***	(.002)	.965
Session	.156*	(.067)	1.168	.194***	(.039)	1.214
Prob.	-.109	(.093)	.897	.801***	(.077)	2.228

Models predict choice of risky option on each trial of the Risky Choice Task. In the left column are model predictors. Models listed "Real Within" and those listed "Sham Within" contain data from subjects who received real or sham stimulation, respectively. OR = Odds Ratio (e^B). SE = Standard Error of B. Intercept = Model Intercept. Prob = Probability that risky option leads to reward (0 = 70%, 1 = 90%). Cert. Rel. Value = Certain Relative Value (percent distance of the certain option from certainty equivalent of the risky option). Session = Session Number (0 = Session 2, 1 = Session 3). * $p < .05$; ** $p < .01$; *** $p < .001$.

While the likelihood of choosing the risky option differed as a function of the risk level for the sham group, it did not for the real rTMS group. Sham subjects were less likely to choose the risky option on trials when its probability was 70% than when it was 90% ($p < .001$; see Figure 10B). In contrast, subjects who received real rTMS chose similarly under different levels of risk (see Figure 10A).

As expected, individuals in both groups were less likely to choose the risky option as the relative subjective value of the certain option increased. In both GEE models, there was a negative effect of Certain Relative Value, such that both sham and real rTMS subjects were less likely to choose the risky option as the relative value of the certain option increased (both $p < .001$) (See Figure 10A & 10B). All subsequent GEE models predicting choices on the Risky Choice Task showed these same effects (see Table 12). This demonstrates that disruption did not make participants completely insensitive to the relative subjective values of the two options.

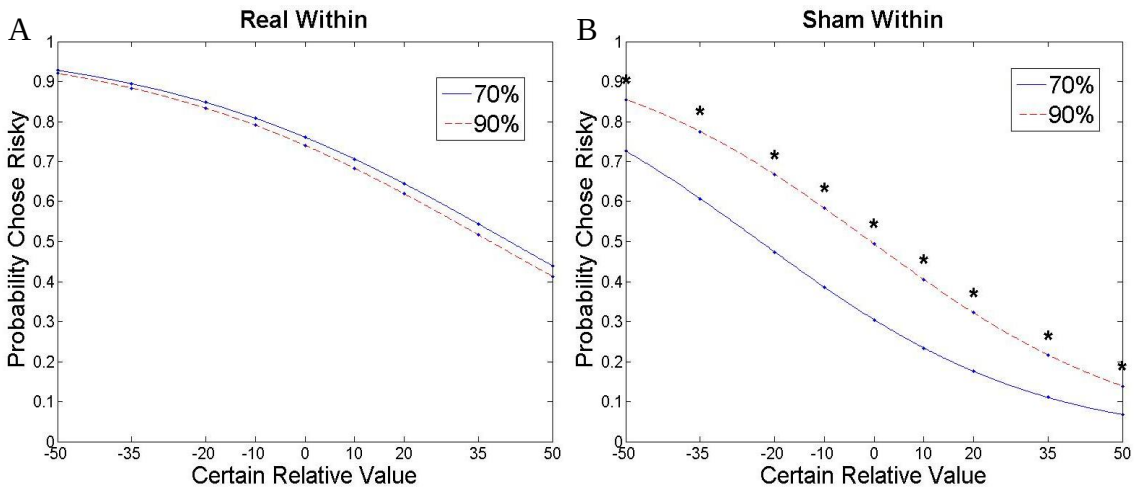

Figure 10: Graphs for Risky Choice as a Function of the Probability of the Risky Option. Probability of choosing the risky option as a function of both the relative value of the certain option and of the probability that the risky option led to reward (i.e. 70% or 90%). Numbers on x-axis are the percent distance between the certain option and the certainty equivalent of the risky option (i.e. x-axis represents Certain Relative Value). A. Compares choice for real rTMS group. B. Compares choice for sham group. Graphs show choice at average session (Session = .5). * Paired comparison significant at $p < .05$.

Additionally, there were positive effects of Session Number within both groups (both $p < .05$). This indicates that choice of the risky option was more likely in session 3 than in session 2 for both groups.

Because there was a significant effect of risk level on choice in the sham group, all subsequent analyses limited themselves to choice on trials with a specific level of risk. Thus behavior on 70% risky probability trials was analyzed separately from behavior on 90% risky probability trials.

Risky Choice

We created separate between subject models to see how risky choice differed as a function of Stimulation Type, Certain Relative Value, and Session Number. Each model

focused on one stimulation site (i.e. left or right DLPFC), and one probability level of the risky option. For all model parameters, see Table 12.

Table 12: Between and Within Subject GEE Models for Risky Choice for each Probability of the Risky Option.

	70% Probability Trials			90 % Probability Trials		
	B	(SE)	OR	B	(SE)	OR
Between: Left DLPFC						
Intercept	-.596***	(.168)	.551	.025	(.128)	1.025
StimType	1.316***	(.210)	3.729	.243	(.164)	1.275
Cert. Rel. Value	-.032***	(.002)	.969	-.029***	(.002)	.971
Session	.046	(.206)	1.047	.500**	(.162)	1.648
StimType X Cert. Rel. Value	.013***	(.004)	1.013	---†		
Between: Right DLPFC						
Intercept	-.761***	(.178)	.467	.246	(.154)	1.279
StimType	.900***	(.181)	2.459	.233	(.166)	1.262
Cert. Rel. Value	-.023***	(.002)	.977	-.024***	(.002)	.977
Session	.401*	(.180)	1.493	-.408*	(.165)	.665
Real Within						
Intercept	.501***	(.123)	1.650	.367***	(.095)	1.443
Side	-.288***	(.042)	.750	-.138***	(.040)	.871
Cert. Rel. Value	-.020***	(.002)	.981	-.023***	(.002)	.978
Session	.246***	(.042)	1.278	.152***	(.039)	1.164

Models predict choice of risky option on each trial of Risky Choice Task. In left column model predictors are listed below each model (listed in **Bold**). Models listed "Between": contain data from subjects who received real rTMS and from subjects who received sham rTMS to the side of DLPFC listed after word "Between". Models listed "Real Within" contain data from subjects who received real rTMS. OR = Odds Ratio (e^B). SE = Standard Error of B. Intercept = Model Intercept. StimType=Stimulation Type (0 = Sham rTMS, 1 = Real rTMS). Session = Session Number (0 = Session 2, 1 = Session 3). Cert. Rel. Value = Certain Relative Value (percent distance of the certain option from certainty equivalent of the risky option). Side = Stimulation Side (0 = Left, 1 = Right). X = Predictor is interaction of two terms. ---†= Interaction not included in model (models only contained interaction if it was significant at $p < .05$). * $p < .05$; ** $p < .01$; *** $p < .001$.

Disruption of either side of the DLPFC increased choice of risky options on the higher risk trials of the task. Choice of the risky option was greater for individuals who received real rTMS to the right hemisphere than for individuals who received sham

stimulation to the same hemisphere when the probability that the risky option led to reward was 70% (Odds Ratio (OR) = 2.459; $p < .001$) (see Figure 11C). Similarly, on trials with this level of risk, choice of the risky option was greater for individuals who received real rTMS to the left hemisphere than for matched sham subjects. Unlike the effect seen in the right hemisphere, the effect of stimulation in the left hemisphere was dependent on the relative value of the certain option. Although the real rTMS group was more likely to choose the risky option at all relative values of the certain option following left hemisphere stimulation, as the relative value of the certain option increased, the increased tendency for the real rTMS group to choose the risky option also increased (see Figure 11A). When the relative value of the certain option was at its highest (i.e. Certain Relative Value was 50), the effect of stimulation on risky choice in the left hemisphere was particularly large (OR = 7.143, $p <.001$).

Unlike the effects of stimulation seen on trials when the probability of the risky option was 70%, when the probability of the risky option was 90%, there were no differences in risky choice between individuals who received real rTMS and those who received sham rTMS (see Figure 11B & 11D). Thus the effects of disrupting the DLPFC on choice were not present on the lower risk trials of the task.

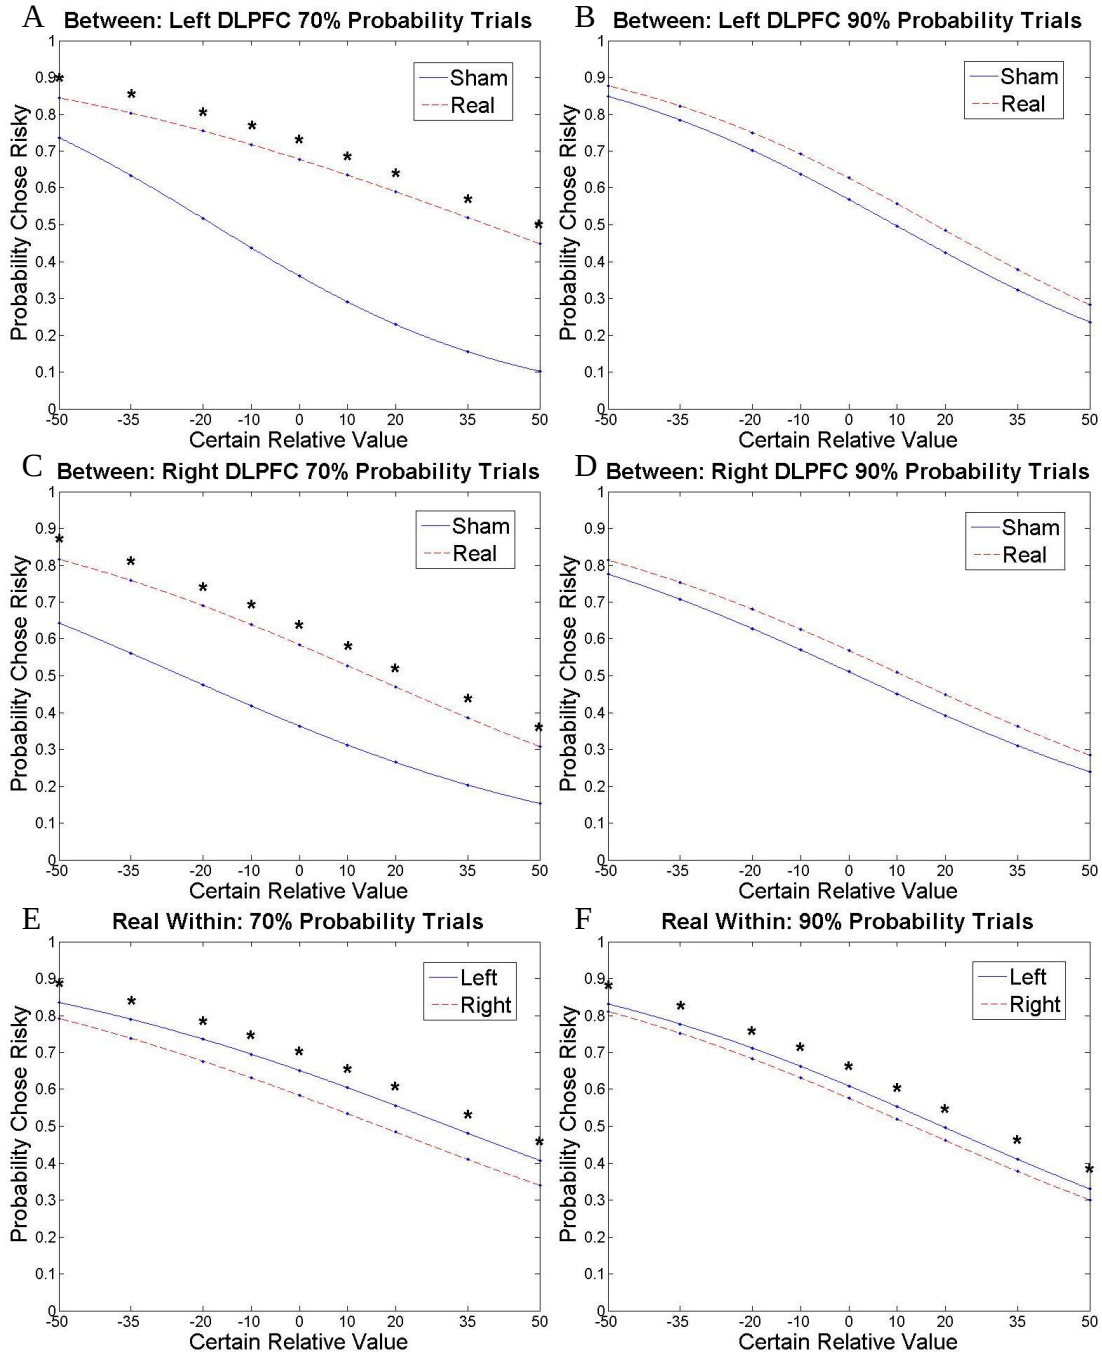

Figure 11: Risky Choice Graphs for each Probability of the Risky Option. Probability of choosing the risky option as a function of the relative value of the certain option. Numbers on the x-axis are the percent distance between the certain option and the certainty equivalent of the risky option (i.e. x-axis represents Certain Relative Value). A-D: Compares choice on trials with a 70% or 90% risky option probability (listed in title) between real and sham stimulation to one side of DLPFC (listed in title). E-F. Compares choice on trials with a 70% or 90% risky option probability (listed in title) across side of stimulation for those who received real rTMS. Graphs show choice at average session (Session = .5). * Paired comparison significant at $p < .05$.

We also investigated whether there was greater choice of risky options following real rTMS delivered to one side of the DLPFC than to the other. There was a lesser tendency to choose the riskier option following disruption of the right than of the left DLPFC both on trials where the probability of the risky option was 70% ($p < .001$; OR = .750) and on trials where the probability of the risky option was 90% ($p < .001$; OR = .871) (See Table 12 and Figure 11E & 11F). These effects were not dependent on the relative value of the certain option, indicating that disruption of the left DLPFC led to a greater overall increase in risky choice than did disruption of the right DLPFC.

We found different effects of Session Number on choice in left hemisphere and right hemisphere between subject models Following stimulation to the left DLPFC, there was a positive effect of Session Number on choice of the risky option when the risky probability was 90% ($p <.01$), but no effect when it was 70%. In contrast, following stimulation to the right DLPFC, there was a negative effect of Session Number on choice of the risky option when the risky probability was 90%, but a positive effect when it was 70% (both $p < .05$). It is difficult to interpret how Session Number affected choice across all subjects, because there was no consistent effect of Session Number on either choice of 90% or of 70% risky rewards.

There were, however, more interpretable effects of Session Number in within subject models. Amongst individuals who received real rTMS, Session Number was consistently positively associated with greater choice of the risky option. In both within subject models for higher risk and lower risk trials, there was a positive effect of session (all $p <.001$), such that subjects who received real rTMS were more likely to choose the

riskier option in session 3 than in session 2. This reveals that individuals who received real rTMS exhibited greater risk seeking over time.

We performed a supplemental analysis to see if the increased tendency to choose risky options on the higher risk trials following disruption of the DLPFC was dependent on the difficulty of the choices. Choices should become more difficult as the subjective values of the two items approach each other, and disruption of the DLPFC could make it especially hard to avoid choosing risky options when the two items are close in subjective value. To assess this, we reran our two between subject GEE models that compared choice between individuals in the real and sham rTMS groups for choices with risky options available with 70% probability, but added two terms to the models: a quadratic effect of Certain Relative Value and an interaction of this term with Stimulation Type. The quadratic effect, Certain Relative Value Squared (i.e. Certain Relative Value^2), was used because the value of this term decreased as the subjective values of the two options approached each other. We note that trials could also have become more difficult as the expected values of the two options approached each other, but given that expected values of the certain options differed across subjects depending on their preferences, we did not include a variable based upon how close the expected values were to each other. We do note that many subjects were close to risk-neutral in session 1, and because of this, the expected values of the two options were often closest in value when the two options were equally subjectively valued. Thus adding a term based on expected value would have exhibited substantial collinearity with the Certain Relative Value^2 predictor. This supplementary analysis revealed that there were no significant interactions between Certain Relative Value^2 and Stimulation Type. This suggests that the increased risk-

seeking on trials with a larger amount of risk following disruption of the DLPFC was not dependent on the difficulty of the choices.

Choice Reaction Time

We next tested whether reaction time was different across real rTMS and sham groups with GEE models to see whether changes in risky choice following disruption were associated with specific changes in reaction time. We assumed that as comparisons and choices became more difficult, reaction time would increase. We thus included a quadratic effect of Certain Relative Value, Certain Relative Value Squared (i.e. Certain Relative Value^2), as a predictor in all models, because trials should become more difficult as the subjective value of the two options approach each other (i.e. as Certain Relative Value^2 decreases). We also included the linear effect of Certain Relative Value in all models, since this would ease interpretation of the quadratic effect. For model parameters, see Table 13.

Disruption of either hemisphere led to faster responding on the Risky Choice Task. Individuals who received real rTMS to the left DLPFC responded more quickly than individuals who received sham stimulation to the same region both on trials where the probability of the risky option was 70% and on trials where it was 90% (both $p < .001$). Individuals who received real rTMS to the right DLPFC also responded more quickly than individuals who received sham stimulation to the same region on trials where the probability of the risky option was 90% ($p < .01$). Additionally, there was a trend for individuals who received real rTMS to the right DLPFC to respond more quickly than the matched sham group on trials with a risky option probability of 70% (p

< .10). Faster reaction times on the choice task following disruption may have been associated with less deliberation before responding. They do not appear, however, to have been associated with a greater tendency to make risky choices. This is because reaction time differences between stimulation groups emerged on all trials, while differences in the tendency to make risky choices were only present on trials where the probability of the risky option was 70%.

In all models but one, there was a significant negative effect of Certain Relative Value^2 as expected (all p <.001). Thus, as the subjective values of the two options approached each other, reaction time slowed, supporting the conjecture that difficulty increased as well. It is unclear, however, why there was not a significant effect of Certain Relative Value^2 on 90% risky reward trials following left hemisphere stimulation.

There was a negative effect of Certain Relative Value on trials with a 70% probability following left DLPFC stimulation ($p < .05$), indicating that reaction time decreased as the value of the certain option increased. However, since no linear trends of Certain Relative Value were significant in other models and there was no hypothesized relationship between this term and reaction time, this predictor is not further discussed.

We found significant effects of Session Number on choice reaction time in the left hemisphere between subject models, but not in the right hemisphere between subject models. Individuals who received stimulation to the left DLPFC responded more quickly in Session 3 than in Session 2 both on trials with 70% and those with 90% risky option probabilities (both $p < .001$). These faster reaction times may have reflected increased familiarity with the task. It is surprising, however, that there were no effects of

Session Number on reaction time following stimulation to the right hemisphere, as the level of familiarity of the task should have been the same regardless of the side of stimulation. One possibility is that the different effects across models were driven by different effects of Session Number on reaction time following real rTMS to the left and right DLPFC, which could have occurred if disruption of the right, but not the left DLPFC, interfered with learning. A divergent effect following rTMS to each side of the DLPFC, however, would not have been captured by our between subject models that included data from both sham and real rTMS subjects, because we did not include a predictor in our models for the interaction of Session Number with Stimulation Type.

Table 13: GEE Models for Risky Choice Reaction Time.

	70% Probability Trials		90 % Probability Trials	
	B	(SE)	B	(SE)
Between: Left DLPFC				
Intercept	3263.657***	(69.485)	3128.853***	(92.375)
StimType	-388.575***	(82.483)	-348.732***	(90.935)
Cert. Rel. Value	-1.195*	(.489)	.660	(.577)
Cert. Rel. Value^2	-.093***	(.016)	-.015	(.018)
Session	-1026.217***	(82.376)	-978.375***	(92.057)
Between: Right DLPFC				
Intercept	2546.824***	(118.588)	2582.488***	(88.957)
StimType	-235.720+	(122.148)	-312.514**	(101.369)
Cert. Rel. Value	.452	(.572)	-.341	(.541)
Cert. Rel. Value^2	-.062***	(.015)	-.071***	(.013)
Session	2.591	(121.157)	.797	(105.103)

Models predict reaction time (ms) in Risky Choice Task. Cert. Rel. Value^2 = Certain Relative Value squared. Other abbreviations same as in Table 12. + p <.10, * p < .05; ** p <.01; *** p <.001.

To see if the divergent effects of Session Number were driven by changes following real rTMS, we reran all four of our between subject GEE models for reaction time but included the interaction of Session Number with Stimulation Type. In no cases

was this interaction significant, revealing that changes in reaction time across sessions were not driven by changes following real rTMS. This suggests that the effects of Session Number on reaction time in the left hemisphere between subject models were not due to familiarity with the task, as familiarity should have been the same following sham stimulation to either hemisphere.

Indifference Points

As a second approach to examining changes in risky choice, we examined whether subject's certainty equivalents for $10 risky gains available with 70% and 90% probability were affected by disruption of the left or right DLPFC. In order to examine the difference in certainty equivalents following real and sham rTMS to each hemisphere, we used the Generalized Linear Model which included Stimulation Type, Session Number, and Session 1 Certainty Equivalent as predictors (See Table 14). In all between subject models there was a significant effect of Session 1 Certainty Equivalent (all $p < .001$), which was expected, since certainty equivalents across time should be related.

Consistent with the results on the Risky Choice Task, real rTMS led to a greater preference for risk than did sham stimulation when the probability of the risky reward was 70%. Following disruption of either the left or the right DLPFC, individuals had higher certainty equivalents for the risky reward available with 70% probability than did matched sham subjects (both $p < .05$), revealing that they became more risk-seeking. However, they did not have higher certainty equivalents than the matched sham group for the risky reward available with 90% probability (although there was a trend for them to have higher certainty equivalents following disruption of the left DLPFC). This is

consistent with our findings on the Risky Choice Task; disruption of either the left or the right DLPFC led to a greater tendency to choose the risky option than did sham stimulation on trials of the Risky Choice Task where the risky option led to reward with 70% probability, but not when it led to reward with 90% probability.

Table 14: Between Subject GENLIN Models for Certainty Equivalents.

	70% Probability Trials		90 % Probability Trials	
	B	(SE)	B	(SE)
Between: Left DLPFC				
Intercept	.951	(.963)	3.465***	(.848)
StimType	1.427*	(.592)	.796+	(.441)
Session	-.360	(.595)	.510	(.437)
Session1CE	.775***	(.160)	.522***	(.105)
Between: Right DLPFC				
Intercept	.728	(1.071)	2.154+	(1.101)
StimType	1.349*	(.677)	.300	(.560)
Session	.448	(.681)	-.111	(.555)
Session1CE	.781***	(.184)	.751***	(.133)

Models Predict Certainty Equivalents for Risky Rewards on Indifference Point Task. Session1CE = Session 1 Certainty Equivalent. Other abbreviations same as in Table 12. * p < .05; ** p <.01; *** p <.001.

We also created within subject GEE models for subjects who received real rTMS to more thoroughly investigate how real rTMS affected certainty equivalents (see Table 15). Consistent with the between subject findings, disruption of either side of the DLPFC increased preference for risk when the probability of the risky reward was 70%. Certainty equivalents for the risky reward available with 70% probability increased following disruption of the left or the right DLPFC (both $p < .05$). There was also some evidence for increased preference for risk when the probability of the risky reward was 90%. Certainty equivalents for the risky reward available with 90% probability were significantly higher than Session 1 values following disruption of the left DLPFC (p

<.001), and there was a trend for them to be higher following disruption of the right DLPFC ($p < .01$).

As a supplementary analysis, we also created within subject GEE models for subjects who received sham rTMS to assess the stability of certainty equivalents across sessions (see Table 15). For subjects who received sham stimulation, there were no significant effects of Side of stimulation or of Session Number, indicating that certainty equivalents for individuals who did not receive real rTMS were stable over time.

Table 15: Within Subject GEE Models for Certainty Equivalents.

	70% Probability Trials		90 % Probability Trials	
	B	(SE)	B	(SE)
Real Within				
Intercept	4.963***	(.447)	6.686***	(.606)
Right Side	1.046*	(.489)	.812+	(.468)
Left Side	.847*	(.400)	1.174**	(.435)
Session	.143	(.172)	.014	(.337)
Sham Within				
Intercept	4.987***	(.507)	7.350***	(.459)
Right Side	.180	(.452)	-.297	(.395)
Left Side	-.398	(.513)	-.288	(.567)
Session	-.083	(.249)	.345	(.323)

Models Predict Certainty Equivalents for Risky Rewards on Indifference Point Task. Right Side = rTMS delivered to right DLPFC (predicts stimulation to right relative to session one baseline value). Left Side = rTMS delivered to Left DLPFC (predicts stimulation to left relative to session one baseline value). Session = Session Number (-1 = Session 1, 0 = Session 2, 1 = Session 3). Real within = within subject models for subjects who received real rTMS. Sham within = within subject models for subjects who received sham rTMS. Other abbreviations same as in Table 12. + p <.10, * p < .05; ** p <.01; *** p <.001.

Discussion

We found that disruption of either the left or right DLPFC with rTMS increased preferences for risk on two monetary reward tasks. On the Risky Choice Task, individuals who received real rTMS had a greater tendency to choose the risky option than did individuals who received sham rTMS when the risky option had a 70% chance of leading to reward, but not when it had a 90% chance. Similarly, certainty equivalents for the 70% risky option as measured on the Indifference Point Task were significantly higher for the real rTMS group than for the sham rTMS group, while certainty equivalents for the 90% risky option were not significantly different across groups. Importantly, while subjects who received sham rTMS exhibited risk preferences on the Risky Choice Task that were modified by the probability level of the risky option, those who received real rTMS did not.

Our data suggest that the DLPFC is involved in devaluation of rewards that are risky, and that the amount of devaluation is dependent upon the level of risk. When functioning of the DLPFC was disrupted, individuals responded similarly on trials with different levels of risk. In contrast, individuals who received sham stimulation were more likely to choose the risky option on trials with a lower level of risk than on trials with a higher level of risk, suggesting that functions of the DLPFC devalued risky options with a higher level of risk. This is further supported by the between subject findings on the Risky Choice Task. Individuals who received real rTMS chose more risky options on trials with a higher level of risk than did sham subjects, but did not choose more risky options on trials with a lower level of risk, suggesting that valuation of

only the more risky items was greater following disruption. Together, this suggests that DLPFC functions alter the subjective value of options available with risk (i.e. with less than certain probability), and that this effect is dependent on the risk level.

An important question is whether the DLPFC devalues all risky rewards, but devalues them more as the risk level increases, or only devalues risky rewards that have a certain level of risk. Our data provide some support for the first possibility. Although there were no significant effects of Stimulation Type on choice on the lower risk trials of the Risky Choice Task, there were some indications that risk preference did increase at this lower level of risk following real rTMS. Certainty equivalents for the 90% risky reward on the Indifference Point task were higher following disruption of the left DLPFC, and marginally higher following disruption of the right DLPFC, than they were for the same individuals prior to stimulation. Additionally, there was a trend for certainty equivalents for the 90% risky reward to be higher following real rTMS than following sham stimulation to the left DLPFC. Together with the findings on the Risky Choice Task, this suggests that the DLPFC devalues risky rewards of different risk levels, but devalues them more as the risk level increases.

Although we have suggested that the DLPFC devalues all risky rewards, with the level of devaluation increasing as the level of risk increases, the DLPFC may only devalue risky rewards of a certain level. It might even overvalue risky rewards of some risk levels. There is some evidence that activation in the DLPFC responds to risk in a nonlinear manner (Tobler et al., 2008). Tobler and colleagues found that activation in the DLPFC associated with the predicted probability of reward was fit by an inverse S-function of probability, which overweighted low probabilities and underweighted high

probabilities relative to a linear function. In contrast to the response seen in the DLPFC, activation in the striatum scaled with risk in a linear manner. This different pattern raises the possibility that compared to other regions of the brain, the DLPFC devalues risky rewards available with high probabilities but overvalues risky rewards available with low probabilities. Would such a pattern fit with our findings? It could if risky rewards available with a sufficiently high probability were not devalued, or devalued less than those available with a slightly lower probability. The findings of Tobler et al. are not inconsistent with this possibility since the highest non-certain reward probability in their study was closer to that of our higher risk level (i.e. 75% in their study) than that of our lower risk level. If the DLPFC does overvalue rewards available with low probabilities (i.e. rewards with a high level of risk), then disruption of the DLPFC might actually lead to greater risk aversion at high levels of risk. To answer this question, future rTMS studies should parametrically vary the level of risk over a wide range of levels.

For intertemporal choices, it has been suggested that the brain contains multiple valuation systems, which are differentially sensitive to immediate rewards (McClure et al., 2004). Our results suggest that the brain has multiple valuation systems for choice involving risk as well, and that the DLPFC is part of one of these valuation systems. Past neuroimaging studies have shown that activation in several regions of the brain, including the striatum and MPFC, scales with the subjective value of risky rewards on choice tasks (Levy et al., 2010; Peters & Buchel, 2009). Given our results, activation in the DLPFC would be expected to scale with subjective value as well. However, this activation might not be well captured by the same models that have been used to capture activation associated with the subjective value of risky rewards in other brain regions. This is

because the DLPFC appears to ascribe a different weight to the level of risk of rewards than do other regions involved in valuation.

How does the DLPFC interact with other brain areas involved in the decision making process? One possibility is that risk related responses in the DLPFC could bias individuals towards greater risk aversion by modulating activity associated with subjective value in other brain regions. There is some evidence that activity in the DLPFC can modulate activity related to valuation on choice tasks that do not involve risk. When individuals make choices between two rewarding items, activation in the DLPFC is functionally connected with activation in the VMPFC, an area with activation that responds to the subjective value of the items on the task (Hare et al., 2009). However, it has not previously been investigated whether risk related responses in the DLPFC are functionally connected to subjective value related responses in the VMPFC or other regions on risky choice tasks.

Previously, it has been suggested that disruption of the DLPFC leads to increased choice of risky options because it disrupts self-control processes that are used to help one choose the safer option (Knoch & Fehr, 2007; Knoch et al., 2006). Self-control accounts of choice hold that the DLPFC, or other frontal regions, perform a cognitive control function during choice by helping one override, modulate, or inhibit prepotent urges that favor selection of superficially seductive options (Figner et al., 2010; Hare et al., 2009; Knoch et al., 2006). In choices involving risks, prepotent urges could favor selection of options that might lead to the highest possible reward, even if they are risky. If so, then self-control might be needed to choose a safer option.

Our data, however, do not support the suggestion that the DLPFC is involved in a self-control function that is used to prevent risk seeking. It could be argued that self-control would be increasingly needed to choose the smaller certain option as the value of the certain option decreases with respect to that of the risky option. This is because as the value of the certain option decreases, the motivational drive to choose the risky option should become more prepotent, and it should become harder to override this urge. If so, then disruption of a region involved in self-control would be expected to lead to greater choice of the risky option as the value of the certain option decreases (i.e. as it becomes more prepotent). Alternatively, a self-control account could predict that a region involved in self-control would show more involvement as the choice becomes more difficult, that is as the values of the two options approach each other, under the assumption that it should be harder to override a prepotent urge to choose the riskier, but potentially more rewarding option when the two options are close in subjective value. This idea holds that disruption of a region involved in self-control would lead to greater choice of risky options as the subjective values of the two options approached each other. Disruption of the DLPFC, however, did not lead to either of the changes in choice that might be expected to occur after disruption of a self-control process that helps prevent risk seeking. Individuals who received real rTMS did not become increasingly more likely than individuals in the sham group to choose the risky option as the relative value of the certain option decreased. They also did not become increasingly more likely than individuals in the sham group to choose the risky option as the relative subjective values of the two options approached each other. Thus, the pattern of our results suggest that the

increased risk seeking that occurred following disruption of the DLPFC was not due to disrupted self-control processes.

One difficulty in the neuromodulation literature is imprecision in the use of the terminology for risk. Risk can be defined in various ways, such as the possibility of a loss, the magnitude of a possible loss, the probability of a non-positive outcome, or as a property of the spread of possible outcomes (e.g. variance or mean-preserving spread). These terms have been conflated in prior neuromodulation studies, making it unclear which of these variables is affected by disruption of DLPFC functioning (Boggio et al., 2010; Fecteau et al., 2007; Knoch et al., 2006). Our conclusions are limited as well, in that the item with the greater spread of possible outcomes in our task was also always the least probable. This, however, does not lessen the importance of our findings, as our study was designed to observe how subjective valuation of risky options changed rather than tease apart different risk constructs. By only varying the probability of one option, we were able to see how choices between certain and risky rewards that were matched for subjective value were affected by disruption of the DLPFC.

An important finding of our study is that we were able to show that risk taking increased following disruption of the DLPFC on a task that offered no possibility of loss. This reveals that the DLPFC has an important role in helping people avoid choosing risky options, even when options cannot lead to losses. Given our findings that disruption of the DLPFC increased risk seeking for choices limited to the gain domain, an important future question is whether it would increase risk seeking for choices limited to the loss domain as well. Compared to other regions involved in valuation, the DLPFC might devalue risky options in both domains, or alternatively might only devalue risky options

that can lead to monetary gain. The divergent risk preferences that people exhibit in the gain and loss domains suggest that how the brain responds to the level of risk may differ when risks involve gains and when risks involve losses (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). This raises the possibility that the DLPFC might evaluate the desirability of risk differently in both domains.

Our results reveal that disruption of either the left or the right DLPFC increases preference for risks, suggesting that the functions of both sides of the DLPFC bias individuals towards risk aversion. Although one past study has found a similar role for both the left and right DLPFC in risk preferences (Boggio et al., 2010), the findings of a larger number of studies have suggested that the roles of the two sides of the DLPFC are different in risk preferences, and that only the functions of the right DLPFC bias individuals towards risk aversion (Fecteau et al., 2007; Gianotti et al., 2009; Knoch et al., 2006). While the right DLPFC may have a more central role in risk-aversion, one possibility is that under conditions of increased task difficulty, the left DLPFC may also come online to help prevent risky choices. It may even have a larger role in risk aversion than the right DLPFC in more difficult conditions, as demonstrated by our within subject findings that individuals chose more risky options on the Risky Choice Task following disruption of the left than of the right DLPFC. Our task was most likely more difficult than the risk task (i.e. Roger's Risk Task) employed by Knoch and others (Boggio et al., 2010; Fecteau et al., 2007; Knoch et al., 2006; Rogers et al., 1999). Options on the trials of Roger's Risk Task that were analyzed in past neuromodulation studies had large differences in expected value. While the expected value of the less risky option was always positive, that of the riskier option was always negative. Because of these large

differences in expected value, the less risky options may have clearly dominated the riskier ones, making it relatively easy to make a risk-averse choice. In contrast, on our risky choice task, stimulus pairs were based on participants' pre-stimulation subjective valuations. We presented two options that were often very close in previously established relative subjective values, and thus it may have been harder to compare the values of the two options and avoid choosing the risky option on our task than on Rogers' Risk Task. The hypothesis that both hemispheres of the DLPFC are necessary for risk-aversion under more difficult situations may also help explain why Boggio and colleagues (2010), who focused on more elderly subjects, observed a similar role for both the right and left DLPFC on Roger's Risk Task after neuromodulation. This task may have been more difficult for the elderly individuals who participated in that study, than for the younger individuals who participated in the study by Knoch et al. (2006).

One unanticipated finding was that sham subjects made more risk-averse choices on the high risk trials of the Risky Choice Task than was predicted by previously established subjective values. A close look at figure 10 reveals that on trials with a lower amount of risk (i.e. risky probability of 90%), sham subjects chose the risky option about half the time when the two options had equal relative subjective values, as predicted. However, on trials with a higher amount of risk (i.e. risky probability of 70%), they chose the risky option approximately 30% of the time when the two options had equal relative subjective values, revealing that they were more risk-averse than predicted. An important question is why sham stimulation led to greater risk-aversion than predicted by session one preferences, but only on trials that had a higher level of risk. We believe this was likely due to the differing characteristics of our task used to elicit preferences (i.e. the

Indifference Point Task) and the Risky Choice Task. This conjecture is supported by supplementary analyses showing that the sham group did not subjectively value risky rewards differently on the Indifference Point Task across sessions.

There are several possible reasons why sham subjects were more risk averse than predicted on the trials of the Risky Choice Task that had a higher level of risk. One possible reason preferences changed is that cross session behavior differed across tasks due to a framing effect. While the Indifference Point Task presented all trials with one level of risk before presenting trials with the other level of risk, the Risky Choice Task randomly mixed trials with different levels of risk. Individuals may have been more risk averse when confronted with a risky option that led to reward with a 70% probability on the Risky Choice Task because it seemed relatively riskier when it was mixed with less risky trials than when it was not mixed with them. Another possible reason why subjects exhibited more risk aversion on the Risky Choice Task than on the Indifference Point Task is because the former task had a time limit while the latter did not. Prior research has shown that individuals are more risk averse when they have less time to respond (Ben Zur & Breznitz, 1981). However, this does not explain why sham subjects only exhibited more risk aversion on the trials with a higher level of risk, unless the effect of response time only occurs at certain levels of risk. Future research investigating this can help resolve this issue. A third possibility is that preferences are less stable when there is a higher level of risk than a lower level of risk. This possibility does not seem very convincing, however, because preferences on the Indifference Point Task for sham subjects did not change across sessions. Certainty equivalents for sham subjects were no

less stable for the risky reward available with 70% probability than for the one available with 90% probability.

Although the preference shifts seen for the sham group on the Risky Choice Task make it difficult to determine how much preferences for risk shifted on the task following real rTMS, data from the Indifference Point Task provide direct evidence that disruption of the DLPFC increased risk seeking above baseline levels. Between subject analyses on this task showed that individuals had a greater preference for risk, as demonstrated by higher certainty equivalents for the risky reward available with 70% probability following real than following sham rTMS. This directly indicates that individuals became more risk seeking following real rTMS, because those in the sham group did not have different certainty equivalents across sessions. Within subject analyses for those who received real rTMS provide even more direct evidence that subjects' preferences for risk were greater following disruption of the DLPFC than they were prior to disruption. Certainty equivalents for the risky reward available with 70% probability were greater following disruption of the left or the right DLPFC than they were for the same subjects in session one (i.e. disruption led to greater preference for risk). There was also some evidence that disruption led to a greater preference for risk for the risky reward available with 90% probability. Certainty equivalents were significantly higher following disruption of the left DLPFC and there was a trend for them to be higher following disruption of the right DLPFC than they were for the same subjects in session one.

We found some evidence that following disruption of the left, but not the right DLPFC, subjects were less sensitive to the relative subjective values of the two items. The slope of the choice curve that predicted whether subjects chose the risky option on

the higher risk trials of the Risky Choice Task was shallower for those who had real rTMS to the left DLPFC than for those who had sham rTMS to this region, indicating that disruption led to choice patterns that were less dependent on the relative subjective values of the options. Other results, however, indicate that even following disruption of the left DLPFC, subjects were not completely insensitive to the relative subjective values of the two options. Following disruption of either side of the DLPFC, individuals were less likely to choose the risky reward with a 70% probability and less likely to choose the risky reward with a 90% probability as the relative value of the certain option increased.

Subjects who received real rTMS had faster reaction times on the Risky Choice Task than did subjects who received sham stimulation. Disruption of either the left or the right DLPFC may have been associated with decreased reaction time because subjects deliberated less prior to making a choice. One could imagine that less deliberation might be associated with greater choice of potentially larger risky options if individuals spend less time focusing on the relative risk levels of the two options before making a decision. Unlike changes in choice on the Risky Choice Task following real rTMS, however, speeded reaction time occurred following real rTMS both for trials with a 70% risky reward and for trials with a 90% risky reward. This divergent pattern of results across analyses reveals that differences in the amount of time spent deliberating over the options cannot fully account for the differences in choice between groups.

In conclusion, our findings extend those of prior neuromodulation studies by showing that the precise changes in risk taking that occur following disruption of the DLPFC depend on the level of risk. We found that disruption of activity in either side of the DLPFC with low frequency rTMS led to significantly greater choice of risky options

than did sham stimulation on trials with a higher level of risk, but not on trials with a lower level of risk, on a task that offered no possibility of monetary loss. This difference was largely due to a relative insensitivity to different levels of risk following disruption. Importantly, our results suggest that both sides of the DLPFC devalue some risky rewards more than do other regions of the brain involved in valuation. Whether the DLPFC solely devalues risky rewards more than other regions, or devalues risky rewards of some levels of risk and values risky rewards of other levels of risk more than other regions is an important question for future research.

CHAPTER IV

EXPERIMENT 3 – INDIVIDUAL DIFFERENCES IN DELAY AND PROBABILITY DISCOUNTING AND IN IMPULSIVITY

Introduction

The subjective values of delayed monetary gains and losses are discounted with respect to those of immediate incentives. Intertemporal choice research has shown that the rate of delay discounting for both delayed gains and losses can be modeled by a hyperbolic or quasi-hyperbolic curve (Estle et al., 2006; Kirby & Marakovic, 1995; Madden et al., 2003; Murphy et al., 2001; Myerson & Green, 1995; Rachlin et al., 1991). Hyperbolic delay discounting is expressed with the following equation (Mazur, 1987):

$$V = A/(1+KD). \qquad \text{(Equation 1)}$$

In this equation, the discounted value (V) of a delayed monetary incentive (i.e. gain or loss) is a function of the objective amount of the incentive (A), the delay to receiving (or paying) the incentive (D), and the individual's discount rate (K). The amount of discounting across individuals is variable, and is captured by the discount rate parameter. Higher rates of delay discounting indicate that the individual places greater weight on the monetary value of immediate relative to delayed incentives. An individual's discount rate, K, may differ for monetary gains of different magnitudes, and for gains and losses. In fact, individuals tend to discount delayed monetary gains more than delayed monetary losses (Baker et al., 2003; Benzion et al., 1989; Estle et al., 2006; Frederick et al., 2004; Murphy et al., 2001).

Empirical data reveal that individuals discount the value of risky monetary rewards (i.e. available with less than 100% probability) relative to certain rewards in a similar manner to how they discount delayed rewards relative to immediate ones. When probability is converted to odds against reward, the subjective value of a risky reward decreases as odds against reward increases, and this decrement in value can be modeled by a hyperbolic or quasi-hyperbolic curve (Myerson et al., 2003; Rachlin et al., 1991). Hyperbolic probability discounting is expressed with the following equation (Rachlin et al., 1991):

$$V = A/(1+HO). \qquad \text{(Equation 2)}$$

The discounted value of a risky monetary gain (V) is a function of the objective amount of the gain (A), the probability discount rate (H), and the odds against receiving the gain (O, where O = (1/probability of reward) – 1). As with delay discounting, individuals differ in their rates of probability discounting (H), and higher rates of discounting reveal that a person places greater weight on the monetary value of certain relative to risky incentives.

Given the individual differences seen in both forms of discounting, it is worthwhile to search for personality traits that may be related to discounting of different types of incentives. Since delay discounting of monetary gains involves a preference for immediate over delayed rewards, a logical place to look for relationships are self-report measures of impulsivity. Impulsive behavior can be defined as the tendencies to give into urges and to respond quickly without planning (Buss & Plomin, 1975). Both of these traits could be related to delay discounting, since a person may choose immediate rewards over larger delayed alternatives either because they cannot override urges to

choose immediate options or because they respond without thinking of the long-term consequences of their choices. Impulsivity could similarly be related to delay discounting of monetary losses, since individuals may prefer larger delayed losses over smaller immediate ones either because of a lack of planning or because of a failure to override urges that encourage avoidance of immediate losses.

Although the tendency to discount risky incentives relative to certain incentives of the same magnitude is not commonly thought of as impulsive behavior, the positive correlations that have been observed between discounting of delayed and risky rewards suggests that a similar process may underlie both types of discounting (Mitchell, 1999; Myerson et al., 2003; Richards et al., 1999). Other evidence suggesting a common process are fMRI findings that the subjective value of both delayed and risky rewards during choice tasks are positively associated with brain activation in similar brain regions, including the MPFC and striatum (Kable & Glimcher, 2007; Peters & Buchel, 2009, 2010). These similarities raise the possibility that both probability discounting and delay discounting have a similar relationship to an impulsivity trait. Alternatively, these similarities may reflect something other than impulsivity. In fact, both types of discounting might actually be expected to have divergent relationships with an impulsivity trait. This is because in the domain of risk, an individual who is highly impulsive might have a tendency to take risks. Such an individual could exhibit low levels of probability discounting If so, then high levels of an impulsivity trait could be related to both low levels of probability discounting and high levels of delay discounting.

To date, a large number of studies have examined the relationship between delay discounting of monetary gains and impulsivity. Many of these studies have found

positive correlations between measures of impulsivity, such as the Eysenck Impulsivity Scale (Eysenck, Pearson, Easting, & Allsop, 1985), Eysenck Personality Questionnaire (Eysenck & Eysenck, 1978), or Barrett Impulsiveness Scale (BIS: Patton, Stanford, & Barratt, 1995), and levels of delay discounting (Kirby et al., 1999; Madden et al., 1997; Mitchell et al., 2005; Petry, 2001, 2002; Reynolds, Richards et al., 2006). In contrast, a study that has investigated the relationship between probability discounting of monetary gains and impulsivity found a negative correlation (Mitchell, 1999). The findings from these studies suggest that delay discounting and probability discounting of gains may have divergent relationships with impulsivity. However, support for this hypothesis is limited because of the pattern of findings in the Mitchell (1999) study. First, the study used more than one impulsivity measure, and the two types of discounting (probability and delay) were not significantly correlated with the same measure. Second, correlations between delay discounting and impulsivity were negative, contrary to the findings of other studies. Furthermore, in one of the largest studies that investigated the relationship between impulsivity and delay discounting, no significant correlations were found (Reynolds, Ortengren, Richards, & de Wit, 2006).

One possible reason for the inconsistencies across studies is that many of the studies have used relatively small sample sizes on the order of 30 to 50 people. This is a hindrance because studies with small samples are more likely to have results that are driven by outliers, and thus may find relationships that are very different from those in the population. Another possible reason is that impulsivity is not a unitary construct, and different measures of impulsivity, such as the BIS and Eysenck Impulsivity Scale, may not in fact be measuring the same underlying trait.

In order to better understand the relationships between impulsivity and different types of discounting, it would be helpful to observe how different facets of impulsivity that are common across measures of impulsivity are related to delay and probability discounting. One existing impulsivity scale that fits this description is the UPPS Impulsive Behavior Scale (Whiteside & Lynam, 2001). Items on the UPPS were selected by performing an exploratory factor analysis on a number of existing personality scales, including impulsivity measures that have shown significant relationships with monetary discounting (e.g. BIS, Eysenck Impulsivity Scale). Items that had high loadings on specific factors were used to construct subscales for different facets of impulsivity. The UPPS Impulsive Behavior Scale measures four facets of impulsivity: Negative Urgency, Lack of Premeditation, Lack of Perseverance, and Sensation Seeking. A revised version of the scale (UPPS-P) adds a fifth facet: Positive Urgency. This fifth facet captures the tendency to engage in impulsive actions when in a positive mood (Cyders & Smith, 2007). Prior research reveals that the level of delay discounting of monetary gains is positively correlated with impulsivity on the Lack of Premeditation facet of the UPPS (Lynam & Miller, 2004), but as no studies have utilized the UPPS-P, it is unclear if delay discounting is related to Positive Urgency. Additionally, no studies have looked at how levels of probability discounting of gains are related to scores on any of the UPPS-P facets. This is a critical comparison which will help uncover whether similar impulsivity traits have divergent relationships with probability and delay discounting.

As prior studies have also not investigated the relationships between delay discounting of losses and facets of the UPPS-P, another important question is whether delay discounting of gains and losses are both related to a similar facet of impulsivity on

this measure. The positive correlations that have been observed between both types of discounting (Murphy et al., 2001) suggest that a similar process underlies delay discounting of gains and losses. If so, then it might be expected that delay discounting of losses would be positively associated with impulsivity, since delay discounting of gains often shows this relationship with impulsivity. However, there has been very little research on the relationship between delay discounting of losses and impulsivity. Ostaszewski & Karzel (2005) divided subjects up into three groups depending on their levels of impulsivity on the Eysenck Personality Questionnaire (Eysenck, Eysenck, & Barrett, 1985), and found that those who had the lowest impulsivity scores discounted the value of delayed losses more than those who were the most impulsive. A limitation of this study is that they did not look at the relationship for the full range of individuals in their sample. To date, no single study has investigated the relationships between impulsivity and both delay discounting of gains and of losses. One might expect that the Lack of Premeditation facet on the UPPS-P would be positively correlated with levels of both types of delay discounting. Individuals who discount the value of both delayed gains and losses more than others might do so because they spend less time thinking about the long-term value of options before making choices.

The current study was designed to investigate the relationships between different facets of impulsivity on the UPPS-P and different types of discounting (delay discounting of gains, delay discounting of losses, and probability discounting of gains) in the same subjects. It was predicted that the Lack of Premeditation facet of the UPPS-P would be positively correlated with levels of delay discounting of gains and losses, but negatively correlated with levels of probability discounting of gains. As lack of consensus seen in

the literature may be due to previous studies utilizing small sample sizes, we utilized a larger sample than is typical, with only one previous study using a sample of similar size (Lynam & Miller, 2004). This gave us sufficient power to uncover relationships between impulsivity and discounting. It also gave us sufficient power for our secondary aim: to look at relationships between all three types of discounting. Understanding the relationships between the tendencies to discount incentives in different domains is important, because it can help reveal whether shared mechanisms underlie different types of discounting.

Methods

Participants

225 right-handed individuals (45.78% female) between the ages of 18 and 30 (*M* Age = 21.13, SD = 2.87) from Vanderbilt University and the Nashville community participated in this study. All of these participants reported having no history of neurological or psychiatric problems, and no females reported being currently pregnant. Specifically, these individuals were consented to participate in either Experiment 1 or Experiment 2 of this dissertation. For numbers consented in each study please see Chapters II and III. Note that some individuals were consented in both Experiments 1 and 2 because their session 1 data made them ineligible for one study but not ineligible for the other (these subjects only completed session 1 once). The study outlined in this chapter is based on session 1 data from all subjects consented in Experiment 1 or 2. All of the procedures in session 1, but the prescreening procedures, were identical in

Experiments 1 and 2. Additional procedures not mentioned in the preceding two chapters are mentioned in this chapter, as they were not pertinent to the hypotheses of the other experiments but are pertinent to those of the experiment outlined here.

26 additional subjects were consented but excluded or withdrew. Reasons subjects were excluded included no loss discounting as revealed on the loss discounting prescreen (for approximately 20 subjects this was given following consent), having previously received Transcranial Magnetic Stimulation (TMS), risk factors that could increase the chances of having negative effects from TMS (e.g. neurological conditions), and experimenter error. All participants completed written informed consent approved by the Vanderbilt IRB.

Indifference Point Procedure

Subjects received one of two questionnaires prior to consent depending on whether they were initially consented in Experiment 1 or Experiment 2 (144 subjects were consented in Experiment 1 and 81 subjects were consented in Experiment 2). The questionnaire for Experiment 1 assessed delay discounting of monetary losses, while the questionnaire for Experiment 2 assessed risk preferences for monetary gains. For details on these prescreens see methods of experiments in Chapters II and III. Subjects consented in Experiment 1 were excluded if the delay discounting prescreen indicated they did not discount the value of delayed monetary losses, while those consented in Experiment 2 were excluded if the risk prescreen indicated they were risk seeking.

Following consent, we determined subject indifference points between immediate and delayed monetary rewards. To do so, subjects completed a task on a computer in

which we determined 12 immediate equivalents (i.e. dollar amount of an immediate monetary reward that person values equally to a delayed monetary reward)) for 4 different delayed monetary gains ($2.50, $5.00, $7.50, and $10.00) available at three different time points in the future (2 weeks, 4 weeks, 8 weeks). The order of determining immediate equivalents was from low to high magnitudes, and from nearer to farther time points in the future (i.e. order was $2.00 available in 2 weeks, $2 available in 4 weeks, $2 available in 8 weeks, … $10 available in 8 weeks). Subjects were told that they would receive the amount of money for one random choice at the time associated with the choice. For timing of one trial of this task and presentation of items, see Figure 12.

Figure 12: Trial Structure of the Indifference Point Task*. A. Decision Phase. Following 500 ms of fixation, two different options were presented on a computer screen (side randomized). Each option had an associated probability (i.e. 70%, 90% or 100%, the latter labeled "For Sure"), valence (i.e. Gain or Loss), amount, and time. The figure above depicts a trial of the Indifference Point task for risky gains. For this task, only the probability and amounts varied across trials. The word "Today" indicated that subjects were making choices for options to be paid that day. Subjects decided which option they would prefer by pressing the "z" or "m" key and there was no time limit to respond. B. Post Choice Phase. Immediately after responding, the triangle under the subject's choice turned red and the other triangle disappeared to indicate the subject's choice; this was displayed for 250 ms. For the example given above, the value of the certain option would be increased on the following trial, since the risky option was just chosen. (*Example shows the indifference point task for risky rewards. Trial structure of the indifference point task for delayed incentives was the same, except that time, rather than probability, was varied across trials.)

To determine each immediate equivalent, participants were presented with an initial trial in which they made a choice between a delayed monetary gain and an immediate gain of half the value of the delayed option. If participants chose the delayed option, the value of the immediate option increased by half, and if they chose the immediate option it decreased by half. On the next trial, the immediate value changed in a similar way as in the previous trial but only by one quarter of the original value. Over six trials, the immediate value increased or decreased by progressively smaller amounts (i.e. by 1/(2^x) where x was trial number) depending on participant responses so that the subjective value of the immediate amount would iteratively approach that of the delayed amount. After the sixth trial, a final catch trial was presented in which the immediate value was higher than the just-calculated immediate equivalent. This provided a check to ensure that subjects were answering according to their preferences (i.e. were answering consistently). If they did not choose the immediate value on the catch trial, the immediate equivalent for that specific delayed magnitude reward was determined again (i.e. the six trial procedure and check was repeated). After answering consistently, or after completing the indifference point procedure three times, participants then performed the indifference point procedure for the next delay-magnitude pair.

Following the indifference point procedure for delayed monetary gains, participants performed the same indifference point procedure for delayed monetary losses. For the delayed loss procedure, the same times and magnitudes of delayed incentives were used as were used for the delayed gains task, except now subjects chose which of the two values they preferred to lose. They were told that they would have to pay the amount of money for one random choice at the time associated with the choice.

Subjects next performed another indifference point procedure, in which they made choices between certain and risky monetary rewards. The risky rewards offered a reward with a probability of 50%, 70%, or 90%, and offered no reward with the opposing probability (i.e. 100 - probability). We determined certainty equivalents (i.e. dollar amount of a certain monetary reward that the person values equally to a risky monetary reward) for 12 different risky monetary gains: 4 different magnitudes (i.e. same as those used for delayed rewards) available with one of three probabilities (i.e. 50%, 70%, or 90%). The order of determining certainty equivalents was from low to high magnitudes, and from lower to higher probability (i.e. order was $2.50 with 50%, $2.50 with 70%.... $10 with 90%). Subjects were told that one of their choices would randomly be selected for payment. If the selected choice was certain, they would definitely receive the money, but if it was risky, they would have a chance to receive it (associated with the probability of the option). On the first trial of this task, individuals were presented with a choice between a risky option and a certain option of half the maximum value of the risky option. Other details of this task were the same as for the indifference point procedure for delayed incentives, except here the magnitude of the certain option was varied rather than that of the immediate option. For all trials of the indifference point procedures, the side of the screen that each option appeared on was determined randomly.

At the end of the session, three random trials were selected for payment: one trial from the indifference point procedure for delayed gains, one trial from the indifference point procedure for delayed losses, and one trial from the indifference point procedure for risky gains. Subjects then either received or paid (depending on trial type) the amount of money at the specified time associated with their choice on that trial. However, if the

chosen option for the selected trial from the procedure for risky gains was a risky option, subjects rolled a 10 sided die to determine whether they would be paid. If they rolled a number less than or equal to the tens digit of the probability listed on that trial (i.e. less than 7 for 70%), then they were paid the chosen amount.

Impulsivity Questionnaire

Immediately after completing the indifference point procedure (and before being paid), individuals completed the UPPS-P Impulsive Behavior scale (Lynam, Smith, Whiteside, & Cyders, 2006). This 59-item scale assesses five facets of impulsivity: Negative Urgency (NU), Lack of Perseverance (LPers), lack of Premeditation (LPrem), Sensation Seeking (SS), and Positive Urgency (PU). Higher scores on each facet indicate more impulsivity within that domain. On the UPPS-P, one rates on a 4-point scale how much they agree or disagree with a number of statements related to their own behavior. We computed scores for each facet by taking the mean score of all items within a facet (note that some items were reverse-coded).

We had complete UPPS-P data for 88.4% of the subjects. There were two reasons for missing data. We had no UPPS-P data for one subject due to experimenter error; the subject was not given the questionnaire. All other subjects with missing data were given the questionnaire and completed some of the items, but did not answer all of them. If subjects had completed at least 70% of the items on a facet, we computed that facet score by replacing missing values with the average of the other scores for that subject on that facet. This allowed us to compute all facet scores for all subjects for whom we had data, with one exception. We did not compute one facet score (i.e. PU) for one subject who

had completed less than 70% of the items for that facet, as recommended in the UPPS-P scoring instructions.

Statistical Analysis Methods

We analyzed data with PASW Statistics 18 (SPSS Inc., Chicago, IL). Delay discount rates (k) were computed by fitting hyperbolic discount functions (Equation 1) to subjects' immediate equivalents across the range of delays (measured in days) using nonlinear regression. Probability discount rates (h) were computed similarly by fitting hyperbolic discount functions (Equation 2) to subject's certainty equivalents; in order to compute these discount rates, the probability that the risky option led to reward was converted to the odds against receiving a reward. We note that hyperbolic discount functions could only be fit to a subject's data if their immediate or certainty equivalents for a specific magnitude incentive varied as the delay or probability of the incentive changed. For each subject, we computed eight time discount rates, one for each gain and loss magnitude, and four probability discount rates, one for each gain magnitude. Higher k or h values indicate a greater degree of discounting. We assessed the model fit of each nonlinear regression by computing Shrunken R^2 (i.e. $R^2 = 1 - (\text{Sum of Squared Error/Corrected Total Sum of Squares})$). We also computed three average discount rates, by averaging the four magnitude-specific discount rates for each of the three types of discounting (delayed gain, delayed loss, and risky gain). Average model fit scores for the three types of discounting were computed by averaging the four respective shrunken R^2 values.

We also assessed discounting using a model-free measure of discounting, area under the curve (AUC) (Myerson, Green, & Warusawitharana, 2001). For each subject, we calculated the area under each of their 12 indifference curves (i.e. four delayed gain magnitude curves, four delayed loss magnitude curves, four risky gain magnitude curves). To do so, we first normalized the subjective value of delayed or risky monetary gain or loss (i.e. from immediate or certainty equivalents at indifference points) and normalized the delay to the incentive or odds against reward, such that the undiscounted value of the incentive was equal to one, and the maximum delay or maximum odds against reward was equal to one. By normalizing the curves, it allows the area under the discounting curve to vary from 0 at the low end (complete discounting) to 1 at the high end (no discounting), and makes the measure of AUC comparable across magnitudes. Thus, lower AUC values indicate a greater degree of discounting. For further details on how to compute AUC, see Myerson et al. (2001). We also computed three average AUC values, by averaging the four magnitude-specific AUC rates for each type of discounting (delayed gain, delayed loss, and risky gain). All subsequent analyses were computed with AUC values rather than k or h values, because data were available for more subjects, and because the distribution of AUC values was much closer to a normal distribution (there was much less skew). However, AUC values were still significantly non-normal (Shapiro-Wilk $p < .05$).

Relationships between discounting and impulsivity, and relationships between different types of discounting were assessed with Spearman correlations (2-tailed) between average AUC values and each facet of impulsivity. We assessed differences

between delayed gain and delayed loss discounting by computing Wilcoxon signed-rank tests (2-tailed) for averaged AUC values.

As a supplemental analysis, we examined the relationship between magnitude of incentives and discounting by using Generalized Estimating Equations (GEE), which model effects while accounting for correlations within observations of individual subjects (Liang & Zeger, 1986). Using PASW Statistics 18 (SPSS Inc., Chicago, IL), we created GEE models with an unstructured correlation matrix and a normal distribution to predict averaged AUC values as a function of magnitude. A separate GEE model was created for each of the three types of incentives.

Data involving delayed gains for a subject were not analyzed if we were missing any of that subject's delayed gain immediate equivalents. Similarly, data involving delayed losses and data involving risky gains were not analyzed if we did not have any of a subject's delayed loss immediate equivalents or risky gain certainty equivalents, respectively. This only occurred if subjects answered inconsistently 3 times in a row on the indifference point procedure for at least one item on that task. In all, 13 subjects did not answer consistently on the delayed gain task, 9 subjects did not answer consistently on the delayed loss task, and 6 subjects did not answer consistently on the risky gain task.

Results

Levels of Discounting

For each subject, we calculated hyperbolic discount rates for four magnitudes of delayed gain, four magnitudes of delayed loss, and four magnitudes of risky gain. In a

large number of cases, however, we were not able to calculate hyperbolic discount rates, because immediate or certainty equivalents for incentives of a specific magnitude were identical at all delays or probabilities of the incentive. Three average discount rates (i.e. for delayed gain, delayed loss, and risky gain) were created by averaging the four magnitude specific discount rates for that type of discounting, and were only created for subjects who had calculable discount rates for all four magnitudes of that type. For median values of discounting, number of subjects with calculable discount rates, and indices of fit for hyperbolic discount rates, see Table 16. Because all discount rates were largely positively skewed, the median and the interquartile range (range) were used to describe the distribution of discount rates rather than the median and standard deviation. The median discount rates averaged across magnitudes for delayed gains, delayed losses, and risky gains were .030 (range = .013 - .056), .026 (range = .008 - .059), and 1.220 (range = .903 - 1.792), respectively. For a visual representation of median averaged hyperbolic discount rates and of median subjective valuation of the delayed or risky incentives (i.e. immediate or certainty equivalents) averaged across magnitudes, see Figure 13.

While we were able to fit hyperbolic discount functions to most individuals' risky gain data (N= 211), we were not able to fit hyperbolic discount functions to a large number of individuals' delayed gain or delayed loss data. Of the subjects who answered consistently on the task, we were only able to calculate hyperbolic discount rates for all magnitudes of delayed gains for 73.6% of subjects (N = 156) and for all magnitudes of delayed losses for 51.9% of subjects (N= 112). This indicates that many subjects had immediate equivalents for delayed incentives that did not vary as a function of delay.

Furthermore, across subjects whose data was fit by hyperbolic delay discount functions, the fit was often quite poor. Median R^2 averaged across magnitudes for delayed gains and losses were .347 and -.024, respectively (See Table 16). These poor model fits reveal that a number of subjects did not exhibit hyperbolic delay discounting. This was especially common when delayed options were losses, consistent with past research showing that standard hyperbolic discounting equations provide poor fits to some individuals' intertemporal choices involving losses (Murphy et al., 2001). In contrast to the poor model fits to our delayed gain and loss data, risky gain data were well fit by hyperbolic discount functions (Median R^2 averaged across magnitudes = .843).

Table 16: Median Hyperbolic Discount Rates and Hyperbolic model fit.

	Magnitude of Delayed or Risky Incentive				
	$2.50	$5.00	$7.50	$10.00	Average
Delayed Gain					
Median k	.032	.023	.019	.015	.030
(range for k)	(.011 - .060)	(.008 - .049)	(.008 - .052)	(.007 - .052)	(.013 - .056)
Median R^2	.704	.687	.593	.783	.347
(range for R^2)	(.063-.817)	(-.428 - .850)	(-1.383 - .884)	(-.574 - .922)	(-2.084 - .764)
N	187	179	173	173	156
Delayed Loss					
Median k	.022	.018	.018	.014	.026
(range for k)	(.007 - .043)	(.006 - .053)	(.006 - .053)	(.005 - .045)	(.008 - .059)
Median R^2	.470	.459	.310	.591	-.024
(range for R^2)	(-.970 - .780)	(-1.387 - .788)	(-3.216 - .801)	(-1.614 - .849)	(-4.213 - .620)
N	138	139	136	139	112
Risky Gain					
Median k	1.112	1.181	1.129	1.050	1.220
(range for k)	(.836 - 1.821)	(.933 - 1.983)	(.831 - 1.692)	(.811 - 1.722)	(.903 - 1.792)
Median R^2	.864	.848	.884	.894	.843
(range for R^2)	(.646 - .943)	(.572 - .963)	(.624 - .957)	(.583 - .980)	(.242 - .919)
N	217	215	214	216	211

Range = Interquartile range. R^2 = Shrunken R^2 index of model fit (note this could be less than 0). N = Number of subjects whose preferences were modeled by corresponding hyperbolic functions.

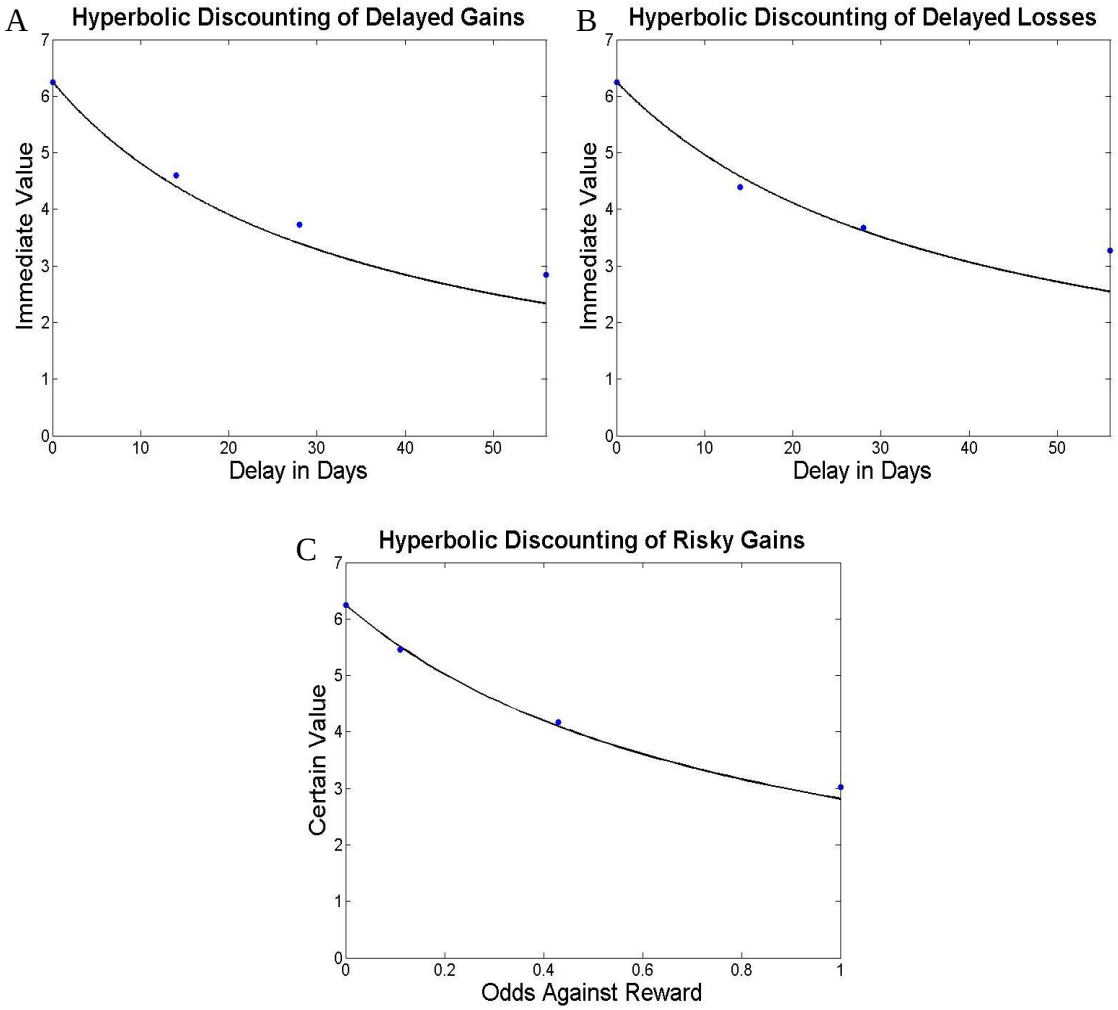

Figure 13: Median Hyperbolic Discounting of Incentives. Curves show the immediate value of incentives of average magnitude as a function of delay (A-B) or show the certain value of rewards of average magnitude as a function of odds against reward (C). Hyperbolic curves in figure created from median values of the discount rate averaged across magnitudes. Dots represent median values of the immediate equivalent averaged across magnitudes for delayed gains (A) or for delayed losses (B), or median values of the certainty equivalent averaged across magnitudes for risky gains (C).

Because data for a large number of subjects were not adequately fit using hyperbolic equations, we also assessed discounting by calculating AUC (i.e. the normalized area under each individual's indifference curve). Calculating AUC allowed us to calculate values of discounting for all subjects, since it did not require any assumptions of how individuals discounted delayed or risky incentives. We calculated separate AUC values for each magnitude of delayed gain, delayed loss, and risky gain, and calculated 3 average discount rates by averaging across magnitudes (i.e. one for each type of discounting).

In contrast to hyperbolic discount rates for which greater discounting is associated with larger discount rate values (i.e. k or h), greater discounting as assessed using AUC is associated with lower AUC values. While no discounting is associated with an AUC value of 1, complete discounting (i.e. immediate or certainty equivalents equal to 0) is associated with an AUC value of 0. However, in the current study subjects could not have these extreme values. Due to the design of our indifference point task, certainty and immediate equivalents always had to be less than the risky or delayed incentive they were associated with by one to seven cents and always greater than zero by one to seven cents (exact amounts depended on the magnitude of the risky or delayed incentive). Because of this, individuals who exhibited no discounting would have had AUC values of approximately .99 and those who exhibited complete discounting would have had AUC values of approximately .13 for delayed incentives and AUC values of approximately .06 for risky gains. For the median AUC values and the number of subjects used to calculate each median, see Table 17. The median averaged AUC values for delayed gains, delayed losses, and risky gains were .690 (range = .529 - .866), .827 (range = .592 - .980), and

.670 (range = .585 - .724), respectively. For a visual representation of the distribution of AUC values, see Figure 14.

Since we had AUC values for many more subjects than we had hyperbolic discount rates for, all further analyses were computed with AUC values. This allowed us to look for relationships across types of discounting and impulsivity that might not have been captured if we had excluded subjects whose choices were not fit by a hyperbolic discount function.

Table 17: Median AUC Values.

	Magnitude of Delayed or Risky Incentive				
	$2.50	$5.00	$7.50	$10.00	Average
Delayed Gain					
Median AUC	.595	.695	.725	.759	.690
(k range)	(.469 - .820)	(.520 - .871)	(.506 - .896)	(.536 - .900)	(.529 - .866)
N	212	212	212	212	212
Delayed Loss					
Median AUC	.829	.835	.831	.835	.827
(k range)	(.593 - .997)	(.567 - .995)	(.596 - .993)	(.589 - .994)	(.592 - .980)
N	216	216	216	216	216
Risky Gain					
Median AUC	.689	.670	.676	.690	.670
(k range)	(.577 - .747)	(.568 - .730)	(.594 - .734)	(.587 - .735)	(.585 - .724)
N	219	219	219	219	219

Range = Interquartile range. N = Number of subjects.

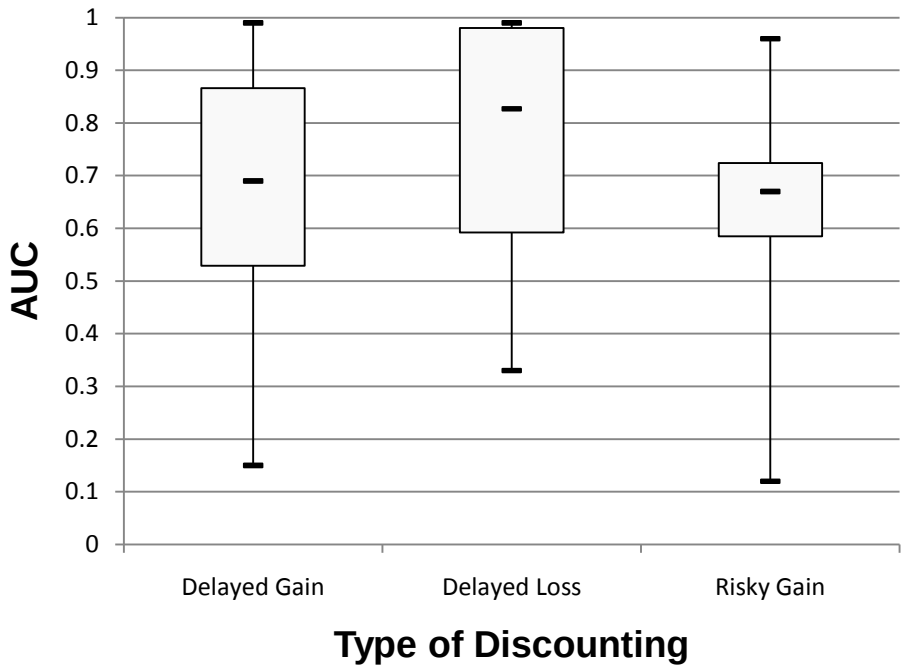

Figure 14: Boxplots of AUC values averaged across magnitudes. Edges of box indicate interquartile range of averaged AUC values. Whiskers indicate minimum and maximum of averaged AUC values. Horizontal line in middle of each box is median of averaged AUC. AUC of 1 represents no discounting, while AUC of 0 represents complete discounting (i.e. delayed option's immediate equivalent or risky option's certainty equivalent is 0).

Relationships between Discounting and Impulsivity

To examine relationships between each type of discounting and impulsivity, we computed Spearman correlations between AUC values averaged across magnitude of each type of discounting and the five facets of impulsivity as measured on the UPPS-P. Spearman correlations were used, because averaged AUC values were non-normal (all Shapiro-Wilk $p < .05$). Averaged AUC values for delayed losses were negatively correlated with the Positive Urgency factor of the UPPS-P ($r = -.150$, $p = .029$), indicating that individuals with higher levels of Positive Urgency discounted the value of delayed losses more. However, we note that this correlation was no longer significant

after correcting for multiple comparisons. No other correlations were significant. For a full list of correlations, see Table 18.

Table 18: Correlations between Averaged AUC values and facets of the UPPS-P.

Impulsivity Facet	Mean	SD	Future Gain AUC	Future Loss AUC	Risky Gain AUC
NU	2.130	.523	-.056	-.073	.016
LPrem	1.984	.486	-.027	.017	-.012
LPers	1.810	.455	-.027	-.048	.058
SS	3.022	.580	.035	-.014	.046
PU	1.770	.573	-.123	-.150*	.115

N for all correlations between all facets (except PU) and Future Gain AUC, Future Loss AUC, and Risky Gain, are 211, 215, and 218, respectively (while N is one less for all correlations with PU). * = $p < .05$. For meanings of impulsivity abbreviations, see Methods. N for descriptive statistics of all Impulsivity facets is 224, except PU (N = 223).

Since more than 10% of subjects were missing complete data on the UPPS-P, it was possible that the observed relationships with Positive Urgency on the UPPS-P were dependent upon the values we imputed for missing responses on the questionnaire. To investigate whether this was the case, we redid all correlation analyses on the subset of participants who had complete UPPS-P data. No relationships were significant. However, consistent with the analyses performed over the entire sample, there was a trend for averaged AUC values for delayed losses to be negatively correlated with Positive Urgency ($r = -.131$, $p = .071$).

Relationships between Types of Discounting

We looked at relationships between each type of discounting, by computing Spearman correlations between the three averaged AUC values. This revealed that all

types of discounting were significantly positively related (all $p < .01$), suggesting some shared trait or underlying mechanism associated with discounting across domains. The relationship between delayed gain and loss discounting was particularly high ($r = .620$, $p < .001$), while those between each type of delay discounting and probability discounting of gains were lower (with delayed gain: $r = .298$, $p < .001$, with delayed loss: $r = .212$, $p = .002$). The greater correlation between different types of delay discounting than between delay and probability discounting suggests that the underlying mechanisms, or traits, of different types of delay discounting are partially distinct from those that underlie probability discounting.

Wilcoxon Signed Rank tests comparing averaged levels of discounting revealed that individuals discounted delayed gains more than delayed losses ($Z = -6.561$, $p < .001$). We did not assess differences between probability and delay discounting, because AUC values across tasks were not comparable, given that the values of the different independent variables used to normalize the data (i.e. of odds against receiving rewards and the delays to incentives) were different.

In order to examine whether there were relationships between the magnitude of incentives and levels of discounting, we used GEE models to predict AUC values as a function of the magnitude of the incentive (i.e. $2.50, $5.00, $7.50, or $10.00). This was done separately for each type of discounting. Analyses revealed that greater magnitudes of incentives predicted greater averaged AUC values for delayed gains (i.e. greater magnitudes predicted less delay discounting) ($B = .007$, $p < .001$). In contrast, the magnitude of incentives did not predict levels of discounting for delayed losses or risky gains ($p = .455$ and $p = .474$, respectively).

Discussion

We did not find evidence supporting our hypothesis that scores on the Lack of Premeditation facet of the UPPS-P would show divergent relationships between delay and probability discounting. No correlations between impulsivity measures on this facet and levels of discounting were significant. Only one impulsivity facet, Positive Urgency, was significantly correlated with any measure of discounting; higher scores on this facet were positively associated with greater delay discounting of losses, but were not associated with other types of discounting. Although we did not find that a single facet of impulsivity was significantly correlated with different types of discounting, we found that all three types of discounting were significantly positively correlated with each other, with correlations between delayed gains and losses being particularly strong. This reveals that a single facet of impulsivity on the UPPS-P is unable to account for the relationships between different types of discounting.

We found that greater levels of delay discounting of losses as measured by AUC scores were associated with higher scores on the Positive Urgency factor of the UPPS-P. In contrast, levels of probability discounting and delay discounting of gains were not significantly related to scores on this facet. At first glance, it may seem strange that Positive Urgency only showed a significant relationship with discounting in the loss domain. The facet of Positive Urgency assesses how much someone endorses taking impulsive actions when they are in positive mood states (Cyders & Smith, 2007). Given the relationship between this impulsivity facet and positive mood, one might expect this impulsivity facet to be particularly related to choices in the gain domain. However, a

close look at the items that make up this scale reveals that a number of them assess the tendency to engage in actions while in a positive mood that may have bad consequences. One reason people may engage in actions that could have bad consequences is because they discount the value of delayed losses and other delayed aversive outcomes. This potential relationship could explain why Positive Urgency was only significantly correlated with levels of discounting in the loss domain. This facet may tap into discounting in the gain domain as well, but our nonsignificant relationships with other forms of discounting suggest that if it does so, then it is only to a minor extent.

In contrast to our predictions, we did not find that any types of discounting were related to the Lack of Premeditation Facet on the UPPS-P. One prior study with a similar sample size to ours found that higher scores on this facet were associated with greater levels of delay discounting (Lynam & Miller, 2004). Given that this facet of impulsivity reflects "the tendency to think and reflect on the consequences of an act before engaging in that act" (Whiteside & Lynam, 2001, p. 685), the results of Lynam and Miller (2004) suggest that individuals may discount delayed monetary gains because they choose immediate options without considering the benefits of delayed options. While there may be a relationship between Lack of Premeditation and levels of delay discounting, our inability to find relationships with this scale indicate that if present they are very small.

Importantly, our findings reveal that behavior on self-report scales of impulsivity has little relationship with behavior on tasks that measure levels of monetary discounting. Neither delay discounting of gains nor probability discounting of gains was significantly correlated with any facet of impulsivity on the UPPS-P. Although delay discounting of losses was correlated with one facet, Positive Urgency, the size of the correlation

coefficient was small, revealing that there was little shared variance between scores on the two measures. This suggests that self-report measures of impulsivity and monetary choice tasks measure different underlying traits, and that both the tendency to choose a better immediate, but worse long-term monetary option and the tendency to choose a potentially larger, but riskier monetary option have little or no relationship with psychometric measures of impulsivity. Investigators should be careful when using the term impulsivity or impulsive choice to describe behavior on intertemporal or risky choice tasks because it may lead others to attribute relationships between psychometric measures of impulsivity and behavior that do not exist.

Alternatively, our inability to find relationships between most forms of discounting and impulsivity could have been due to characteristics of the UPPS-P, rather than a limitation of self-report impulsivity measures in general. Self-report measures of impulsivity might be able to tap into monetary discounting behavior if they assessed impulsivity in the economic domain. Impulsivity could be domain specific, and a trait that leads individuals to discount the value of delayed or risky economic incentives might only be tapped by a measure that contains items involving economic behavior. This is a possible explanation for our findings, as no items on the UPPS-P assess behavior in the economic domain. Researchers should consider examining whether levels of monetary discounting are related to behavior on the DOSPERT (Weber, Blais, & Betz, 2002), which assesses risk taking across multiple domains, and should consider devising new self-report measures to tap into facets of impulsivity, such as lack of premeditation or positive urgency, with items in the economic domain.

We replicated past findings of positive correlations between delay discounting of gains and both delay discounting of losses (Murphy et al., 2001) and probability discounting of gains (Mitchell, 1999; Myerson et al., 2003; Richards et al., 1999). In our study, the strength of the positive correlations between these forms of discounting was also similar to those seen in past research. Myerson and colleagues (2003) observed positive correlations of .17 and of .26 between delay discounting of gains and probability discounting of gains in two samples of over 100 individuals; these correlations are of similar magnitude to our correlation of .298. Although other studies have found higher positive correlations between these two types of discounting (r between .40 and .75), these numbers are less likely to reflect the relationship between the two types of discounting in the general population due to the small sample sizes of the studies (40 and 24, respectively) (Mitchell, 1999; Richards et al., 1999). The strength of our positive correlation between discounting of delayed gains and losses (i.e. r = .620) was remarkably similar to that seen in a previous study (i.e. r = .570) (Murphy et al., 2001). Together with past findings, our results suggest that there is a greater positive correlation between delay discounting of gains and delay discounting of losses than there is between delay discounting of gains and probability discounting of gains. Importantly, we know of no other studies that have looked at both of these relationships in the same individuals.

We also extend findings of past studies by showing that delay discounting of losses is positively correlated with probability discounting of gains. The magnitude of the correlation between these two types of discounting was similar to the magnitude of the correlation between delay discounting of gains and probability discounting of gains. Thus, we found that delay discounting of gains and delay discounting of losses were

more strongly correlated with each other than either of them were with probability discounting of gains. This suggests that there may be greater similarities between the traits and neural processes that lead to delay discounting of different valenced incentives, than there are between those that lead to discounting of delayed and risky incentives. However, there are also likely to be similar processes that lead to discounting of delayed gains, delayed losses, and risky gains. This is because all forms of discounting were positively correlated with each other in our sample. Our data indicate, however, that measures of impulsivity do not account for the shared processes underlying different types of discounting, as no impulsivity facets were significantly correlated with multiple types of discounting.

We replicated prior findings by showing that individuals discounted the value of delayed gains more than they discounted the value of delayed losses (Baker et al., 2003; Benzion et al., 1989; Estle et al., 2006; Frederick et al., 2004; Murphy et al., 2001). Additionally, we replicated the results of prior studies by showing that the level of delay discounting decreased as the magnitude of the delayed gain increased (Benzion et al., 1989; Estle et al., 2006; Green et al., 1997; Green et al., 1999; Kirby, 1997; Kirby et al., 1999; Thaler, 1991). In the current study, there were, however, no effects of magnitude on discounting of either delayed losses or risky gains. The relationship with delay discounting of losses is in line with that of other studies that have shown that the magnitude of an outcome has less of an effect on levels of delay discounting for losses than for gains (Baker et al., 2003; Estle et al., 2006). In contrast to our findings, however, prior studies have shown that the level of probability discounting increases as the magnitude of the gain increases (Christensen et al., 1998; Estle et al., 2006; Myerson

et al., 2003). One potential reason for our inability to find magnitude effects in the domain of risk is that we used monetary gains of small magnitude, in contrast to the large magnitude gains used in other studies (Estle et al., 2006; Myerson et al., 2003). All of our monetary amounts were relatively small, and were never greater than $10. We used such small amounts because it allowed us to have subjects make choices for real incentives. However, it may have prevented us from seeing magnitude effects on levels of probability discounting. Prior research has shown that the effect of magnitude on the level of probability discounting is smaller when different risky reward magnitudes are both relatively small (i.e. $1 and $10) than when one of them is large (i.e. $10 and $10,000) (Christensen et al., 1998).

The intertemporal choice data of many of our subjects was not fit by hyperbolic discount functions. For approximately 50% of subjects, we were not able to model their preferences for at least one magnitude of delayed monetary loss with a hyperbolic discount function. For preferences involving delayed gains, this figure was approximately 25%. Some evidence suggests that for some individuals, standard hyperbolic delay discount functions provide poor fits to choice data, but more so in the loss than in the gain domain (Murphy et al., 2001). The reason that many preferences in our study could not be modeled with hyperbolic discount functions was because subjects often responded identically regardless of delay (i.e. immediate equivalents were the same across delays). That is, many subjects either did not discount delayed incentives at all, or discounted all delayed incentives by the same amount regardless of the delay. For subjects who did have intertemporal choice data that could be fit by hyperbolic equations, the fit was often extremely poor, indicating that many subjects' intertemporal choice

preferences were not explained by hyperbolic discounting. Model fit was worse for losses, which had a median fit index of less than zero. In contrast to the poor fits of hyperbolic delay discount functions, for over 95% of subjects, all of their risky gain data could be modeled by hyperbolic discount functions, and hyperbolic fits were relatively good.

While at first glance our data suggest the possibility that the subjective value of delayed incentives are modeled more poorly with hyperbolic discount functions than are the subjective value of risky incentives, a closer look at the details of our task reveals that the poor fits may have been due to the stimuli used to model discounting. The amount of data points used to fit hyperbolic functions in our study was small. While each curve was only constructed from preferences at three points (i.e. for three delays or probabilities), it is common to use five or more points (Kirby & Marakovic, 1995; Madden et al., 1997; Mitchell, 1999; Murphy et al., 2001; Myerson & Green, 1995; Petry, 2001, 2002). This may have explained the poor fits, but still does not explain why many subjects responded similarly across time points. The lack of variability in immediate equivalents across delays may have been partially due to our small range of delays. We did not use any long delays in our study, as delays ranged from 2 weeks to only 8 weeks. In contrast, in other studies it is typical to have some delays as long as 6 months, with some studies even including some longer delays on the order of years (Madden et al., 1997; Murphy et al., 2001; Myerson & Green, 1995; Petry, 2001, 2002). Perhaps if we had used longer delays, individuals would have shown more variability in their preferences with respect to time. Although many individuals' immediate equivalents did not vary as a function of time, most individuals' certainty equivalents did vary as a function of probability. The

greater variability seen in certainty equivalents may have been because we used stimuli whose attributes varied over a larger range. On the indifference point task for risky rewards, the range of probabilities was not small, and ranged from 50% to 90%.

Levels of delay discounting may have been fit better by a more complex equation, such as a quasi-hyperbolic function that contained two free parameters (Myerson et al., 2003). However, a more complex function still would not have fit the data of the large number of subjects who exhibited no variability in their responses as a function of time. Because of this, we believed the best way to measure levels of discounting was to calculate the area under the subject's discounting curves (AUC). This model-free approach allowed us to examine levels of discounting in all subjects, because it could be computed even for those subjects who had identical immediate equivalents for delayed incentives at all levels of delay. It also revealed that subjects tended to discount the values of delayed gains, delayed losses, and risky gains, as expected.

We examined relationships between impulsivity and three different types of monetary discounting (i.e. delayed gain, delayed loss, risky gain) with a larger sample than has been used in most prior studies. Although levels of discounting were positively correlated across all three types of discounting measured, no facet of impulsivity had a similar relationship with each type of discounting. One facet of impulsivity, Positive Urgency, was positively correlated with greater delay discounting of losses, but the magnitude of this correlation was low. In contrast, no facets of impulsivity were significantly correlated with delay discounting of gains or probability discounting of gains. The positive correlations seen across different types of discounting suggest that a common trait underlies the tendencies to discount different types of incentives, but our

results reveal that it is not related to impulsivity constructs as measured by the UPPS-P Impulsive Behavior scale. Many of the relationships seen between behavioral measures of monetary discounting and self-report measures of impulsivity in past research may have been due to the small sample sizes of the studies.

Chapter V

GENERAL DISCUSSION

In the remaining part of this dissertation a number of broad questions related to intertemporal and risky choice are addressed. The questions build off of the findings of the experiments that were performed. Specific attention is drawn to the roles played by the DLPFC and PPC in economic choices. Some limitations of the experiments are also described, including the limitations of the subject population enlisted and of the inferences that can be drawn from rTMS. This chapter concludes with an overall summary of the key findings of this dissertation.

Questions Concerning Choice

What are the roles of the DLPFC and PPC in intertemporal and risky choice?

The findings from Experiments 1 and 2 suggest that the DLPFC is involved in valuation of delayed and risky incentives rather than in cognitive control functions during economic choice tasks. Disruption of the DLPFC with rTMS led to general shifts in choice preferences across the entire range of values on the Intertemporal Choice Task (Experiment 1) and across the entire range of values on the more risky trials of the Risky Choice Task (Experiment 2), suggesting that valuation of incentives changed in both tasks. Compared to individuals who received sham stimulation, individuals who received real rTMS were more likely to choose better immediate, but worse long-term incentives

(Experiment 1) and were more likely to choose risky incentives (Experiment 2). If cognitive control functions of the DLPFC helped individuals choose better long-term, but worse immediate options and less risky options, then one would expect that such control would be increasingly needed to select these options as the motivational drive to select other options became more prepotent (i.e. as the values of immediate or risky options increased with respect to those of delayed or certain options, respectively). Disruption of the DLPFC, however, did not lead to greater changes in choice as the relative value of the immediate or risky options increased, suggesting that the DLPFC is not involved in cognitive control functions that help individuals make better long-term intertemporal choices and less risky choices.

The shifts in choice seen in Experiments 1 and 2 following real rTMS suggest that the DLPFC evaluates the desirability of incentives on risky and intertemporal choice tasks differently than do other brain regions. The results, however, do not support the conjecture that DLPFC functions are necessary to evaluate incentives and make economic decisions. This is because following disruption of the DLPFC in both Experiments 1 and 2, individuals were still sensitive to differences in the previously established subjective values of the two options. Like individuals who received sham stimulation, individuals who received real rTMS were more likely to make choices for options as their relative values increased, indicating that they could evaluate the desirability of the two options. This pattern of results suggests that the DLPFC is only one brain region that evaluates incentives, and that when valuation processes of the DLPFC are disrupted, choice follows directly from valuations performed by other brain regions.

The results of Experiment 1 suggest that the PPC is involved in cognitive control functions, rather than valuation, during economic choice tasks. In contrast to the effects seen following disruption of the DLPFC, disruption of the right PPC did not lead to general shifts in preference relative to sham stimulation. When the value of the immediate option was low, disruption of the right PPC did not affect choice. However, as the relative value of the immediate option increased, individuals who received real rTMS became increasingly more likely than individuals who received sham rTMS to choose options with a better immediate, but worse long-term value. This evidence supports a cognitive control account of the PPC because cognitive control should be increasingly needed to choose the better long-term option as the value of the immediate option increases. Although the data of Experiment 1 do not indicate which particular cognitive control function the PPC could be involved in during economic choice, one might expect that it would be involved in the control function of top-down attention, since this is a key function of the region (Corbetta & Shulman, 2002).

What do the studies performed tell us about how economic choices are made?

There are two key processes that typically come into play while making an economic choice. First, an individual evaluates the desirability of each option. Once choice options are ascribed values, a person compares the value of each option so that the more highly valued option can be selected. The results of Experiments 1 and 2 provide important information about the first of these processes: valuation.

Importantly, the results suggest that a multiple valuation model of choice is more tenable than a single-valuation model. The key difference between these models is in the

number of different valuation systems that are involved in making a choice. A single-valuation model holds that one region or a set of regions in the brain computes subjective value of options in a similar way and that a choice is made as a result of subjective value coding in this system (Glimcher, 2009; Kable & Glimcher, 2009). In contrast, a multiple valuation model of choice holds that different regions of the brain compute different subjective values for different options, by responding differently as a function of choice parameters, and that these divergent valuation signals are summated in some way to arrive at a decision (Glimcher, 2009; McClure et al., 2004).

For intertemporal choices, an influential multiple valuation model has held that a "patient system" which includes the DLPFC discounts the value of delayed options relative to that of immediate options less steeply than does a set of other brain regions (McClure et al., 2004). The data of Experiment 1 support this idea and those of Experiment 2 build upon it by suggesting that there are multiple valuation processes that ascribe value to risky choice options as well. The results of these two experiments suggest that there are at least two valuation systems that respond differently to choice options as a function of risk level, and at least two valuation systems that respond differently to choice options as a function of delay. There may also be divergent systems that respond to only the subjective values of choice options involving losses or gains. However, in the DLPFC, one system appears to evaluate the desirability of both positive and negative incentives during intertemporal choice.

An important question is how information from different valuation systems influences choice. Valuation signals from each valuation system might be integrated with each other to arrive at a decision, or conversely, each system might act in parallel

and compete with the other in determining a decision. These possibilities are not mutually exclusive, as some signals from each system might be integrated, while others might compete. Additionally, in some regions of the brain, both of these processes could occur.

In Chapters II and III, I suggested that valuation signals from the DLPFC might be integrated with valuation signals in the MPFC, based on evidence showing functional connectivity between the DLPFC and this region during choice (Hare et al., 2009). While there may be integration of value signals here, other findings suggest that the striatum might have a more central role in the integration of values and ultimately determining what choices an individual makes. Like the MPFC, the striatum is responsive to the subjective value of options on intertemporal and risky choice tasks (Kable & Glimcher, 2007, 2010; Levy et al., 2010; Peters & Buchel, 2009, 2010). However, findings suggest that the striatum is more sensitive than the MPFC to many other value related signals. As discussed in Chapter I, the region that is most consistently activated in humans to predictions of value and its components is the striatum, not the MPFC. The striatum is responsive to predictions of positive and negative value, and to parameters of value, including the expected value, magnitude, and probability of rewards. The greater tendency to respond to value related signals in the striatum make it a more plausible candidate than the MPFC for integration of values.

The connectivity of the striatum with other regions makes it a good candidate for both integration of value and selection of options. Other regions involved in valuation, including the MPFC (both dorsal and ventral), OFC, DLPFC, and the dopaminergic midbrain all project to the striatum (Haber & Knutson, 2010). Although neurons from

the DLPFC mostly project to more dorsal regions of the striatum than neurons from regions of the MPFC, some neurons from both the DLPFC and MPFC project to common regions of the striatum (Haber, Kim, Mailly, & Calzavara, 2006). Some neurons in these striatal regions could integrate valuation signals from different valuation systems. Alternatively, competitive interactions between signals from different valuation systems could occur in these common regions. For instance, some neurons in these regions might only be influenced by signals from one valuation system at any time; which valuation system influences them could be determined by the relative strength of signals from different valuation systems.

Other value signals from different regions of the striatum might be integrated with each other or might interact in more indirect ways, through loops of connections between the striatum and midbrain dopamine neurons. Through these connections, striatal neurons that are located more ventrally are able to influence those that are located more dorsally (Haber, Fudge, & McFarland, 2000). These connections also give a possible mechanism linking valuation with choice. Value signals from more ventral striatal regions may be able to bias behavior towards actions to choose specific items by indirectly influencing activity in the dorsal striatum, a region important for motor behavior (Alexander, Delong, & Strick, 1986; Haber, 2003). All of these properties make the striatum a good candidate for where values are compared and decisions are made.

If future research confirms that there are multiple valuation systems that each differentially influence what choices are made, it will be important to examine what factors might change the relative weight of information from each valuation system in the choice process. Cognitive control processes might have such a role, and the results of

Experiment 1 suggest that the PPC might be a key region in cognitive control processes during choice. Given the many reciprocal connections between the PPC and DLPFC, control functions of the PPC would be in a good position to alter valuation functions of the DLPFC.

Could the PPC be involved in valuation rather than cognitive control during choice?

The significant interactions seen between Stimulation Type and Immediate Relative Value on choice following rTMS to the right PPC in Experiment 1 are supportive of the proposal that this region is involved in cognitive control functions that help individuals choose in line with goals to select options with the best long-term value. Following disruption of the right PPC, subjects were increasingly more likely than sham subjects to choose the better immediate option as the relative value of the immediate option increased.

However, there is another way to interpret these interactions that does not rely on the assumption that individuals have specific goals to choose options with the best long-term value during intertemporal choice tasks. The right PPC could be involved in a valuation process during intertemporal choice that places greater relative weight on the value of delayed vs. immediate incentives than do other regions responsive to value, but only under some circumstances. The right PPC does not appear to place more weight in general on the value of delayed incentives than do other regions, because if it did one would have expected to see a general shift in the choice curves following disruption of this region, similar to what was seen following disruption of the right DLPFC in Experiment 1. However, the right PPC could evaluate delayed vs. immediate incentives

differently than do other regions as a function of the magnitudes of the incentives. Other regions responsive to value might be more responsive to changes in the magnitude of immediate incentives than to changes in the magnitude of delayed incentives; the value attributed to incentives might increase at a greater rate as the magnitude of immediate incentives increases than as the magnitude of delayed incentives increases. In contrast, the right PPC might be more similarly responsive to changes in the magnitude of both immediate and delayed incentives. If so, then valuation processes of the right PPC would be expected to diverge more from those of other regions as the relative magnitude of an immediate incentive increases with respect to that of a delayed incentive; as the relative magnitude of an immediate incentive increases, the right PPC would be expected to place greater relative weight on the value of delayed vs. immediate incentives than do other regions of the brain. This suggestion would be consistent with the effects on choice seen following disruption of the right PPC. Individuals who had real rTMS to this region became increasingly more likely than sham subjects to choose the better immediate option as the relative value of the immediate option increased.

Thus, the data of Experiment 1 are consistent with a suggestion that both the right DLPFC and right PPC are involved in more "patient" valuation processes that place greater relative weight on the value of delayed vs. immediate incentives than do other regions of the brain. However, the different effects seen following disruption of each region suggest that if both regions are involved in valuation processes, then they do not evaluate options similarly. Unlike the purported valuation processes of the right DLPFC, those of the right PPC might only ascribe more relative weight to delayed vs. immediate incentives than do other regions when the magnitude of an immediate incentive on an

intertemporal choice task reaches a certain level with respect to the magnitude of a delayed incentive. Despite this conjecture, I believe it is more likely that the right PPC is involved in a cognitive control function than a valuation function during intertemporal choice. This is because that activation in the PPC does not typically scale with the subjective value of items on choice tasks or bidding tasks, or to the components of value prediction (see Chapter I).

Does the PPC perform a cognitive control function in Risky Choice?

An important question is whether the PPC has a role in cognitive control in risky choice. In Chapter I of this dissertation, I postulated that cognitive control might help an individual choose a certain reward instead of a potentially larger, but potentially worse risky reward, just as it might help an individual choose a better long-term, but worse immediate intertemporal choice option. This assumes that in risky choice tasks cognitive control would be needed to override a prepotent urge to choose a potentially larger risky reward and in intertemporal choice tasks cognitive control would be needed to override a prepotent urge to choose the best immediate reward.

There are, however, some reasons to suggest that prepotent urges might generally lead to selection of certain options, rather than of risky options, in risky choice tasks. Individuals tend to devalue both risky and delayed rewards with respect to certain immediate ones (Frederick et al., 2004; Myerson et al., 2003; Rachlin et al., 1991). Additionally, BOLD signal tends to decrease in many brain regions, as does firing of midbrain dopamine neurons, as both delay to reward increases and probability of receiving a reward decreases (Fiorillo et al., 2008; Kable & Glimcher, 2007, 2010; Peters

& Buchel, 2009; Tobler et al., 2005). These similarities between processing immediacy and certainty suggest that motivational drives to choose certain and immediate options might often be prepotent with respect to those to choose risky and delayed ones, respectively. If so, then cognitive control may help one override a prepotent urge to choose a better immediate option in order to select a better long-term option on an intertemporal choice task, yet help one override a prepotent urge to choose a less risky option in order to select a riskier option with a higher expected value on a risky choice task.

However, this is only a conjecture. In order to test whether the PPC has such a role in cognitive control during risky choice, one could administer rTMS to the PPC and have individuals perform a risky choice task afterwards. If it was found that disruption of the PPC led to changes in choice that were dependent on the relative values of the options in a manner that was similar to what was found following disruption of the right PPC in Experiment 1 (i.e. if disruption of the PPC led to greater choice of certain options, and this effect increased as the certain value increased), then it would suggest that the PPC has a role in cognitive control during both risky and intertemporal choice. If it did not lead to any changes, then it might suggest that individuals did not attempt to use cognitive control on the risky choice task. Alternatively, it could indicate that there is something specific about intertemporal choices that engage control, or other, functions of the PPC.

Does a single valuation process in the DLPFC both devalue risky rewards more than other brain regions and value delayed incentives more than other brain regions?

It is not intuitive that a single valuation process of the DLPFC both devalues risky incentives more than valuation processes of other regions, yet values delayed incentives more than valuation processes of other regions. Delays are somewhat risky, and thus one might expect a process that devalues risky incentives more than the processes of other regions to devalue delayed incentives more than the processes of other regions as well. Yet, the perception of risk of delayed options is likely to be highly dependent upon the context of the decision. The amount of risk that an individual believes is associated with a delayed reward may be so low that risk valuation processes of the DLPFC do not evaluate the risk of the reward differently than do other risk valuation processes in the brain. This could be the case, given that the results of Experiment 2 suggest that the DLPFC might not evaluate low levels of risk differently than do other brain regions.

One way to investigate whether the same brain processes in the DLPFC are involved in evaluating both the risk level of incentives and the delay to receiving them is to see whether the same neurons in the DLPFC are responsive to both the levels of risk and delay of incentives. Importantly, one would want to search for responses to different levels of risk and delay that were different in the DLPFC than in other regions responsive to these parameters, since the results of Experiments 1 and 2 suggest that valuation processes of the DLPFC evaluate these parameters differently than do those of other regions. One could search for neurons in the DLPFC that have a larger decrement in response as the risk level of rewards increases (i.e. greater devaluation of risky rewards) and a smaller decrement in response as delay to rewards increases (i.e. greater valuation of delayed rewards) than neurons in other regions of the brain responsive to valuation.

Are the functions of the DLPFC and PPC similar in risky choices involving losses and in risky choices involving gains?

An important finding of Experiment 1 was that disruption of either the right DLPFC or the right PPC affected choices similarly in both choices involving gains and in choices involving losses. In both circumstances, disruption led to a greater preference for options with a better immediate, but worse long-term value, highlighting a role for these regions in optimization of long-term vs. immediate value. An important question is whether disruption of these regions with rTMS would lead to parallel effects for both gains and losses in choices involving risk as well.

Because gains become more risky as they become less probable (i.e. less likely to lead to a positive outcome), but losses become more risky as they become more probable, using the term "level of risk" can lead to confusion across the gain and loss domains. To simplify discussion, here I will use the term "level of probability" rather than level of risk. Importantly, the construct "level of probability" leads to unique hypotheses of the function of the DLPFC across the gain and loss domains.

One might expect a valuation function, such as has been proposed for the DLPFC, that encourages risk aversion for gains by lowering the positive value of rewards as they become less probable (i.e. devalues them), to encourage risk seeking for losses by lowering the negative value of losses as they become less probable (i.e. they would become less negative). This idea assumes that the DLPFC gives options less overall value (i.e. they are valued as less positive or less negative) than do other regions involved in valuation as the probability level decreases. The appeal of this model is that it suggests that the DLPFC has a role in leading to the pattern of risk preferences that people typically have in the gain and loss domains. Individuals tend to be risk averse with gains,

but risk seeking with losses (Kahneman & Tversky, 1979) (although as noted below, this pattern is different when options lead to losses or gains with low probability). If the DLPFC ascribes losses with less overall negative value as they become less probable than do other brain regions, then disruption of the DLPFC would be expected to cause greater risk aversion, rather than greater risk seeking, in choices involving only losses.

How disruption of the DLPFC with rTMS affects risky choices in just the loss domain is a critical question that should be investigated in future studies. As with gains, specific changes in risk taking may be dependent upon the probability level of risky options. Probability level should be modified parametrically to see how disruption of the DLPFC affects risk preference for losses over a large range of probabilities. Crucially, this needs to be done for risky choices involving only gains as well, since it is not known how disruption of the DLPFC affects risk preferences for gains as the probability of the risky option is varied over a large range. Effects on risk preference might be expected to scale in an opposite manner with the probability of gains and with the probability of losses. Such a proposal is not unlikely given that individuals tend to be risk averse with high probability gains and risk seeking with low probability gains, yet tend to show opposite risk preferences with high and low probability losses (Tversky & Kahneman, 1992); functions of the DLPFC could underlie these different patterns of risk preference.

What do correlations among different types of choice preferences suggest about brain function?

The positive correlations seen in Experiment 3 among the tendencies to discount the value of delayed gains, of delayed losses, and of risky gains suggest that some shared neural process or trait may lead to discounting these three types of incentives. An

intriguing possibility is that individual differences in functions of regions of the brain such as the striatum and MPFC, which have activation which positively scales with the subjective value of risky and delayed rewards, might be related to individual differences in all three types of discounting (Kable & Glimcher, 2007, 2010; Levy et al., 2010; Peters & Buchel, 2009, 2010). Individuals who have lower levels of probability and delay discounting could have value related responses in these regions that diminish less as probability decreases and delay increases, than do individuals with higher levels of discounting. One limitation, however, to this proposal is that it is unclear how the subjective value of risky and delayed losses would scale with brain activity on choice tasks, as no prior studies have investigated this.

It is unlikely, however, that processes of the DLPFC underlie the common tendencies to discount delayed gains, delayed losses, and risky gains. If they did, one would expect that all three of these types of discounting would change in the same direction following disruption of the DLPFC (i.e. all increase or decrease). Instead, in Experiments 1 and 2 there were divergent effects on discounting of incentives in the delayed and risky domains following real rTMS: discounting of delayed incentives increased, while discounting of risky incentives decreased. This makes it unlikely that the neural substrates leading to the covariation of delay and probability discounting are related to DLPFC processes, and instead, suggests that DLPFC processes may actually limit the degree of correlation between delay and probability discounting.

What are limitations of the inferences that can be drawn from rTMS?

If functioning is altered following rTMS, then one can infer that the brain structures affected by rTMS had a role in that function. However, one cannot necessarily infer that the targeted brain region is essential for that function. This is because rTMS not only affects the brain region over which it is applied, but also has indirect effects on brain regions connected to the targeted region (Eisenegger, Treyer, Fehr, & Knoch, 2008). I have claimed roles for the DLPFC and PPC in economic choices based on the results of Experiments 1 and 2 since rTMS targeting these regions altered choice behavior. However, the changes in choice patterns seen following rTMS may have been due to disruption of functions in regions connected to the DLPFC and PPC, rather than (or in addition to) disruption of functions in the DLPFC and PPC.

Previous findings show that changes in regional cerebral blood flow following low frequency rTMS administered to the DLPFC occur in both the DLPFC and ipsilateral VLPFC (Eisenegger et al., 2008). This raises the possibility that changes in choice patterns seen in Experiments 1 and 2 were due to disruption of functions of the VLPFC, rather than the DLPFC. In intertemporal choice studies, it has been shown that the VLPFC exhibits similar activation patterns to the DLPFC, and has been claimed to be part of the "patient system" that helps individuals choose better long-term incentives (McClure et al., 2007; McClure et al., 2004). Thus there is some reason to believe that disruption of this region might have led to the changes in choice patterns seen in Experiment 1. In the future, it will be important to try to tease apart the roles of the DLPFC and VLPFC in both intertemporal and risky choice.

Even though the data from Experiments 1 and 2 are not supportive of a cognitive control account of DLPFC functioning during choice, this does not mean that the DLPFC is not involved in such a function during intertemporal and risky choice tasks. TMS does not disrupt all functioning in a region equally, as the effects of TMS depend on the orientation of neurons in the target region with respect to the orientation of the coil (Wagner, Rushmore, Eden, & Valero-Cabre, 2009). Thus even though functions in a region are disrupted following TMS, not all functions will be disrupted equally, and some might not be disrupted at all. It is thus possible that some of the key functions of the DLPFC and PPC that are involved in making economic choices were not disrupted by rTMS. This limitation reveals the importance of using a variety of different neuroscientific approaches when studying the relationships between the brain and behavior.

Do the inclusion criteria of the experiments limit the generalizability of the findings?

All three of the experiments limited inclusion to individuals with specific risk or time preferences. Subjects were excluded from Experiment 1 if they did not discount delayed losses because inclusion of these subjects would not have allowed us to always have individuals make a choice between a small immediate loss and a larger delayed loss. Having this reason for exclusion was essential in testing a self-control account of intertemporal choice, since individuals should not be motivated to use cognitive control processes to choose an immediate loss that is larger than a delayed loss. Subjects who were in Experiment 2 were excluded if they were risk seeking. This was done to increase statistical power, under the assumption that if the DLPFC modulates risk aversion and

prevents risk-seeking behavior, disruption of this region might not have effects or might have smaller effects on individuals who were already risk seeking. Some subjects in Experiment 3 were excluded for each of the reasons subjects were excluded in the other experiments, since individuals in Experiment 3 were consented to participate in Experiments 1 or 2. Because of the inclusion criteria, the findings from all three experiments might not generalize to individuals with different types of preferences.

Effects on choice following disruption of cognitive control functions might be expected to differ across individuals with different preferences, because cognitive control might be used to accomplish different choice goals across individuals. Similarly, effects on choice following disruption of valuation functions might differ across individuals, because valuation functions might normally respond differently to different choice parameters (i.e. risk or time) in individuals with different preferences.

The results of one fMRI study suggest that valuation processes in the PFC respond differently to risk level in those with different risk preferences. Tobler and colleagues (2009) observed that the relationship between activation in the right lateral frontal pole and risk level (as defined as spread of reward outcomes) depended on risk preferences. Individuals who were risk seeking had greater activation in this region for a high risk than a low risk option, while those who were risk averse had greater activation for the low risk option. Even though they found activation in a different region of PFC than I targeted with rTMS in Experiment 2 of this dissertation, it does raise the possibility that the purported valuation functions that were affected with rTMS respond differently to the level of risk in individuals with different risk preferences.

How do people solve intertemporal choice or risky choice problems?

One difficulty in determining the neural mechanisms of intertemporal and risky choices is that there are many different ways a person could make economic choices. Different people may make choices in different ways, or the same person might make choices in different ways at different times. To illustrate, people may make choices based on subjective values that are determined automatically. Alternatively, they could utilize cognitive operations to manipulate values in working memory and make a choice based on manipulated values. The cognitive operations performed could differ depending on an individual's goals and task rules. For example, in risky choices some individuals may compute expected values, while in intertemporal choices some people may multiply future values by a discount rate. People could also hold online different decision rules, such as criteria that if one option dominates another by a certain level of a parameter, they will choose it. How much attention is given to various attributes could differ as well, and this could be a function of goals or task rules. An additional factor in intertemporal choice is that some individuals may use prospection to imagine the desirability of future outcomes.

All of these different factors reveal the complexity of determining the brain regions involved in making intertemporal and risky choices. If different individuals are approaching these choices in different ways, then the brain functions used to make choices will differ as well. Future intertemporal and risky choice studies would benefit by asking individuals to report how they are making their decisions. This could reveal patterns between levels of delay and probability discounting and the ways individuals make choices. From a clinical point of view, it could help identify methods to guide

individuals in decision making. One limitation to this approach, however, is that individuals may not be able to accurately report how they are making decisions on economic choice tasks, as they might not remember or be aware of how they are making decisions.

In future neuroimaging studies, investigators could ask individuals to perform choices in different ways, to see how the neural mechanisms of choice differ with respect to the different cognitive operations used. A good example of such an approach has recently shown that the brain processes which respond to the subjective value of intertemporal choice options differ when prospection is used at the time of decision making (Peters & Buchel, 2010). When performing such studies, investigators should keep in mind that there may be multiple valuation systems in the brain that respond differently to choice parameters and that the relative levels of activity in different brain regions responsive to valuation may vary as a function of the operations used.

Brain processes involved in cognitive control would be predicted to have different effects on choice depending on the task rules or goals that an individual has while performing a choice task, since individuals can use cognitive control to make choices in line with their intentions. The different effects on choice might arise due to different effects of cognitive control functions on valuation systems. More specifically, valuation signals in multiple regions of the brain may be modulated differently by control functions as a result of a person's goals. Future behavioral and neuroeconomic studies could benefit by trying to uncover what cognitive control is being used for. This could be roughly assessed by asking individuals to report what their goals or strategies are during choice tasks.

Under what circumstances are the left vs. the right DLPFC important in choice?

Disruption of either side of the DLPFC affected choices involving risk, yet only disruption of the right DLPFC affected intertemporal choices. Other low frequency rTMS studies of risky or intertemporal choices have found different patterns of effects across side of the DLPFC (Figner et al., 2010; Knoch et al., 2006). An important question is what accounted for these different effects of side of stimulation across studies. The relative involvement of each side of the DLPFC in valuation processes may depend on the components of the tasks used, such as task difficulty. Alternatively it might depend on individual differences in lateralization of brain functions.

Although it has been proposed that the two sides of the frontal lobe have different roles in approach and avoidance behavior (Davidson, 1992), the findings of Experiment 1 suggest that the relative involvement of the left vs. right DLPFC in economic choice is not dependent on whether choices involve approach vs. avoidance. Disruption of the right DLPFC led to similar changes for both intertemporal choices involving gains and losses, while disruption of the left DLPFC did not significantly affect either type of intertemporal choice. This is not the pattern one would expect to see if one side of the DLPFC had a larger role in approach and the other side had a larger role in avoidance. This is because choices involving gains would be expected to engage systems involved in approach behavior more than do choices involving losses, while choices involving losses would be expected to engage systems involved in avoidance behavior more than do choices involving gains.

What are the implications regarding the relationship between self-report impulsivity measures and choice preferences?

Although a number of scientists have labeled intertemporal choices for options with a better immediate, but worse long-term value than alternatives as impulsive choices (Cardinal, 2006; Deluty, 1978; Figner et al., 2010; Logue, 1988; Madden et al., 1997; Pine et al., 2009), the results of Experiment 3 reveal that the relationships between the tendency to make "impulsive choices" and other constructs of impulsivity are at most very modest. Levels of impulsivity derived from self-report scales do not adequately tap into either levels of delay or probability discounting. Because of this, scientists should be careful in using the term "impulsive choice" to describe behavior on choice tasks, as it might lead others to assume that there are relationships between choice behavior and measurements of impulsivity on self-report scales that do not exist.

The lack of strong relationships between impulsivity and choice preferences suggest that choice preference tasks measure an important component of behavior that is not captured by self-report impulsivity scales. Assessing levels of delay and probability discounting in the laboratory might actually provide a better indicator of the likelihood of engaging in unhealthy choices in everyday life than do self-report scales. Many problem behaviors such as eating unhealthy foods, using substances, or deciding not to receive an immunization involve tradeoffs between immediate vs. long-term value. Other problem behaviors such as engaging in unprotected sex or gambling involve assessments of risk. Levels of delay and probability discounting might tap into the tendencies to engage in such unhealthy behaviors better than do other types of measures, because the decision processes engaged during choice tasks may more closely approximate those engaged in outside the laboratory that lead to unhealthy behaviors.

Concluding Remarks

The experiments of this dissertation have expanded our knowledge of the neural mechanisms of intertemporal choice and choice involving risk. They have also provided important information on the relationships between different choice preferences, and between choice preferences and impulsivity. In Experiment 1, it was shown that disruption of the right DLPFC led to increased choice of options for monetary gain and for monetary loss that had a better immediate, but worse long-term value than alternatives. Disruption of the right PPC was also shown to lead to increased choice for these options, but this effect was dependent on the relative value of the two options, suggesting a role for the PPC in cognitive control. In Experiment 2, disruption of either side of the DLPFC led to increased choice of options for potentially larger risky monetary gains, an effect that was dependent on the probability level of the risky option. The pattern of results following disruption of the DLPFC in both Experiments 1 and 2 suggests that the DLPFC has an important role in valuation during both risky and intertemporal choices. In Experiment 3, there were positive correlations between levels of three types of monetary discounting – i.e. delayed gains, delayed losses, and risky gains. However, there were no strong relationships between any forms of monetary discounting and impulsivity. Together this suggests that a shared trait may be related to preferences across choice domains, and that this trait has little relationship with impulsivity as measured on self-report scales.

Understanding the neural mechanisms of economic choices will help reveal how consumers make economic decisions. But more importantly, it will help reveal how

humans and other animals make all kinds of choices, because the processes that make up an economic choice, valuation and value comparison, are important components for any choice between potentially rewarding or potentially aversive options.

REFERENCES

Alexander, G. E., Delong, M. R., & Strick, P. L. (1986). Parallel Organization of Functionally Segregated Circuits Linking Basal Ganglia and Cortex. *Annual Review of Neuroscience, 9*, 357-381.

Allais, M. (1953). Le Comportement de l'Homme Rationnel devant le Risque, Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica, 21*, 503-546.

Andres, M., Seron, X., & Olivier, E. (2005). Hemispheric lateralization of number comparison. *Cognitive Brain Research, 25*(1), 283-290.

Asaad, W. F., Rainer, G., & Miller, E. K. (2000). Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology, 84*(1), 451-459.

Assmus, A., Marshall, J. C., Noth, J., Zilles, K., & Fink, G. R. (2005). Difficulty of perceptual spatiotemporal integration modulates the neural activity of left inferior parietal cortex. *Neuroscience, 132*(4), 923-927.

Bach, D. R., Seymour, B., & Dolan, R. J. (2009). Neural Activity Associated with the Passive Prediction of Ambiguity and Risk for Aversive Events. *Journal of Neuroscience, 29*(6), 1648-1656.

Badre, D., & D'Esposito, M. (2007). Functional magnetic resonance imaging evidence for a hierarchical organization of the prefrontal cortex. *Journal of Cognitive Neuroscience, 19*(12), 2082-2099.

Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nature Reviews Neuroscience, 10*(9), 659-669.

Baker, F., Johnson, M. W., & Bickel, W. K. (2003). Delay discounting in current and never-before cigarette smokers: Similarities and differences across commodity, sign, and magnitude. *Journal of Abnormal Psychology, 112*(3), 382-392.

Baldo, J. V., & Dronkers, N. F. (2006). The role of inferior parietal and inferior frontal cortex in working memory. *Neuropsychology, 20*(5), 529-538.

Ben Zur, H., & Breznitz, S. J. (1981). The effect of time pressure on risky choice behavior. *Acta Psychologica, 47*, 89-104.

Benzion, U., Rapoport, A., & Yagil, J. (1989). Discount Rates Inferred From Decisions: An Experimental Study. *Management Science, 35*(3), 270-284.

Bernoulli, D. (1738; 1954). Exposition of a new theory on the measurement of risk. *Econometrica, 22*, 23-36.

Berns, G. S., Laibson, D., & Loewenstein, G. (2007). Intertemporal choice - toward an integrative framework. *Trends in Cognitive Sciences, 11*(11), 482-488.

Bjork, J. M., Knutson, B., Fong, G. W., Caggiano, D. M., Bennett, S. M., & Hommer, D. W. (2004). Incentive-elicited brain activation in adolescents: Similarities and differences from young adults. *Journal of Neuroscience, 24*(8), 1793-1802.

Bodner, M., Kroger, J., & Fuster, J. M. (1996). Auditory memory cells in dorsolateral prefrontal cortex. *Neuroreport, 7*(12), 1905-1908.

Boggio, P. S., Campanha, C., Valasek, C. A., Fecteau, S., Pascual-Leone, A., & Fregni, F. (2010). Modulation of decision-making in a gambling task in older adults with transcranial direct current stimulation. *European Journal of Neuroscience, 31*(3), 593-597.

Brass, M., & von Cramon, D. Y. (2004). Decomposing components of task preparation with functional magnetic resonance imaging. *Journal of Cognitive Neuroscience, 16*(4), 609-620.

Braver, T. S., Cohen, J. D., Nystrom, L. E., Jonides, J., Smith, E. E., & Noll, D. C. (1997). A parametric study of prefrontal cortex involvement in human working memory. *Neuroimage, 5*(1), 49-62.

Brodmann, K. (1909). *Vergleichende Lokalisationslehre der Grosshirnrinde in ihren Prinzipien dargestellt auf Grund des Zellenbaues.* Leipzig: Barth.

Bunge, S. A. (2004). How we use rules to select actions: a review of evidence from cognitive neuroscience. *Cognitive Affective & Behavioral Neuroscience, 4*, 564-579.

Buss, A. H., & Plomin, R. (1975). *A temperament theory of personality development*. New York: Wiley.

Butters, N., & Pandya, D. (1969). Retention of Delayed-Alternation - Effect of Selective Lesions of Sulcus Principalis. *Science, 165*(3899), 1271-&.

Butters, N., Pandya, D., Sanders, K., & Dye, P. (1971). Behavioral Deficits in Monkeys after Selective Lesions within Middlethird of Sulcus Principalis. *Journal of Comparative and Physiological Psychology, 76*(1), 8-14.

Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of Cognitive Neuroscience, 12*(1), 1-47.

Caminiti, R., Chafee, M. V., Battaglia-Mayer, A., Averbeck, B. B., Crowe, D. A., & Georgopoulos, A. P. (2010). Understanding the parietal lobe syndrome from a neurophysiological and evolutionary perspective. *European Journal of Neuroscience, 31*(12), 2320-2340.

Camus, M., Halelamien, N., Plassmann, H., Shimojo, S., O'Doherty, J., Camerer, C., et al. (2009). Repetitive transcranial magnetic stimulation over the right dorsolateral prefrontal cortex decreases valuations during food choices. *European Journal of Neuroscience, 30*(10), 1980-1988.

Cappelletti, M., Barth, H., Fregni, F., Spelke, E. S., & Pascual-Leone, A. (2007). rTMS over the intraparietal sulcus disrupts numerosity processing. *Experimental Brain Research, 179*(4), 631-642.

Cardinal, R. N. (2006). Neural systems implicated in delayed and probabilistic reinforcement. *Neural Networks, 19*(8), 1277-1301.

Cardinal, R. N., & Howes, N. J. (2005). Effects of lesions of the nucleus accumbens core on choice between small certain rewards and large uncertain rewards in rats. *BMC Neuroscience, 6*, 37.

Caspers, S., Geyer, S., Schleicher, A., Mohlberg, H., Amunts, K., & Zilles, K. (2006). The human inferior parietal cortex: Cytoarchitectonic parcellation and interindividual variability. *Neuroimage, 33*(2), 430-448.

Champod, A. S., & Petrides, M. (2007). Dissociable roles of the posterior parietal and the prefrontal cortex in manipulation and monitoring processes. *Proceedings of the National Academy of Sciences of the United States of America, 104*(37), 14837-14842.

Champod, A. S., & Petrides, M. (2010). Dissociation within the Frontoparietal Network in Verbal Working Memory: A Parametric Functional Magnetic Resonance Imaging Study. *Journal of Neuroscience, 30*(10), 3849-3856.

Chao, L. L., & Knight, R. T. (1997). Prefrontal deficits in attention and inhibitory control with aging. *Cerebral Cortex, 7*(1), 63-69.

Christensen, J., Parker, S., Silberberg, A., & Hursh, S. (1998). Trade-offs in choice between risk and delay depend on monetary amounts. *Journal of the Experimental Analysis of Behavior, 69*(2), 123-139.

Christopoulos, G. I., Tobler, P. N., Bossaerts, P., Dolan, R. J., & Schultz, W. (2009). Neural Correlates of Value, Risk, and Risk Aversion Contributing to Decision Making under Risk. *Journal of Neuroscience, 29*(40), 12574-12583.

Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the Control of Automatic Processes - a Parallel Distributed-Processing Account of the Stroop Effect. *Psychological Review, 97*(3), 332-361.

Cole, M. W., & Schneider, W. (2007). The cognitive control network: Integrated cortical regions with dissociable functions. *Neuroimage, 37*(1), 343-360.

Corbetta, M., Kincade, J. M., Ollinger, J. M., McAvoy, M. P., & Shulman, G. L. (2000). Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nature Neuroscience, 3*(3), 292-297.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience, 3*(3), 201-215.

Culham, J. C., & Kanwisher, N. G. (2001). Neuroimaging of cognitive functions in human parietal cortex. *Current Opinion in Neurobiology, 11*(2), 157-163.

Cyders, M. A., & Smith, G. T. (2007). Mood-based rash action and its components: Positive and negative urgency. *Personality and Individual Differences, 43*(4), 839-850.

Davidson, R. J. (1992). Emotion and Affective Style - Hemispheric Substrates. *Psychological Science, 3*(1), 39-43.

Dehaene, S., Piazza, M., Pinel, P., & Cohen, L. (2003). Three parietal circuits for number processing. *Cognitive Neuropsychology, 20*(3-6), 487-506.

Deluty, M. Z. (1978). Self-Control and Impulsiveness Involving Aversive Events. *Journal of Experimental Psychology-Animal Behavior Processes, 4*(3), 250-266.

Doyon, J., Owen, A. M., Petrides, M., Sziklas, V., & Evans, A. C. (1996). Functional anatomy of visuomotor skill learning in human subjects examined with positron emission tomography. *European Journal of Neuroscience, 8*(4), 637-648.

Eisenegger, C., Treyer, V., Fehr, E., & Knoch, D. (2008). Time-course of "off-line" prefrontal rTMS effects - a PET study. *Neuroimage, 42*(1), 379-384.

Ernst, M., Nelson, E. E., McClure, E. B., Monk, C. S., Munson, S., Eshel, N., et al. (2004). Choice selection and reward anticipation: an fMRI study. *Neuropsychologia, 42*(12), 1585-1597.

Estle, S. J., Green, L., Myerson, J., & Holt, D. D. (2006). Differential effects of amount on temporal and probability discounting of gains and losses. *Memory & Cognition, 34*(4), 914-928.

Eysenck, S. B. G., & Eysenck, H. J. (1978). Impulsiveness and venturesomeness: their position in a dimensional system of personality description. *Psychological Reports, 43*, 1247-1255.

Eysenck, S. B. G., Eysenck, H. J., & Barrett, P. (1985). A revised version of the Psychoticism scale. *Personality and Individual Differences, 6*, 21-29.

Eysenck, S. B. G., Pearson, P. R., Easting, G., & Allsop, J. F. (1985). Age norms for impulsiveness, venturesomeness, and empathy in adults. *Personality and Individual Differences, 6*, 613-619.

Fecteau, S., Knoch, D., Fregni, F., Sultani, N., Boggio, P., & Pascual-Leone, A. (2007). Diminishing risk-taking Behavior by modulating activity in the prefrontal cortex: A direct current stimulation study. *Journal of Neuroscience, 27*(46), 12500-12505.

Fellows, L. K., & Farah, M. J. (2005). Dissociable elements of human foresight: a role for the ventromedial frontal lobes in framing the future, but not in discounting future rewards. *Neuropsychologia, 43*(8), 1214-1221.

Figner, B., Knoch, D., Johnson, E. J., Krosch, A. R., Lisanby, S. H., Fehr, E., et al. (2010). Lateral prefrontal cortex and self-control in intertemporal choice. *Nature Neuroscience, 13*(5), 538-539.

Fiorillo, C. D., Newsome, W. T., & Schultz, W. (2008). The temporal precision of reward prediction in dopamine neurons. *Nature Neuroscience, 11*(8), 966-973.

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science, 299*(5614), 1898-1902.

FitzGerald, T. H. B., Seymour, B., & Dolan, R. J. (2009). The Role of Human Orbitofrontal Cortex in Value Comparison for Incommensurable Objects. *Journal of Neuroscience, 29*(26), 8388-8395.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2004). Time Discounting and Time Preference: A Critical Review. In C. F. Camerer, G. Loewenstein & M. Rabin (Eds.), *Advances in Behavioral Economics* (pp. 162-222). Princeton: Princeton University Press.

Friedrich, F. J., Egly, R., Rafal, R. D., & Beck, D. (1998). Spatial attention deficits in humans: A comparison of superior parietal and temporal-parietal junction lesions. *Neuropsychology, 12*(2), 193-207.

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic Coding of Visual Space in the Monkey's Dorsolateral Prefrontal Cortex. *Journal of Neurophysiology, 61*(2), 331-349.

Fuster, J. M. (1997). *The prefrontal cortex : anatomy, physiology, and neuropsychology of the frontal lobe* (3rd ed.). Philadelphia: Lippincott-Raven.

Fuster, J. M. (2001). The prefrontal cortex - An update: time is of the essence. *Neuron, 30*(2), 319-333.

Fuster, J. M. (2002). Physiology of Executive Functions: The Perception-Action Cycle. In D. T. Stuss & R. T. Knight (Eds.), *Principles of Frontal Lobe Function* (pp. 96-108). New York: Oxford University Press.

Fuster, J. M., & Alexander, G. E. (1971). Neuron Activity Related to Short-Term Memory. *Science, 173*(3997), 652-654.

Gianotti, L. R., Knoch, D., Faber, P. L., Lehmann, D., Pascual-Marqui, R. D., Diezi, C., et al. (2009). Tonic Activity Level in the Right Prefrontal Cortex Predicts Individuals' Risk Taking. *Psychological Science, 20*(1), 33-38.

Glimcher, P. W. (2003). *Decisions, uncertainty, and the brain : the science of neuroeconomics*. Cambridge, Mass.: MIT Press.

Glimcher, P. W. (2008). Understanding risk: A guide for the perplexed. *Cognitive Affective & Behavioral Neuroscience, 8*(4), 348-354.

Glimcher, P. W. (2009). Choice: Towards a Standard Back-pocket Model. In P. W. Glimcher, C. F. Camerer, E. Fehr & R. A. Poldrack (Eds.), *Neuroeconomics Decision Making and the Brain* (pp. 503-521). Amsterdam: Elsevier.

Gobel, S., Walsh, V., & Rushworth, M. F. S. (2001). The mental number line and the human angular gyrus. *Neuroimage, 14*(6), 1278-1289.

Goldenberg, G., & Karnath, H. O. (2006). The neural basis of imitation is body part specific. *Journal of Neuroscience, 26*(23), 6282-6287.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In V. B. Mountcastle, F. Plum & S. R. Geiger (Eds.), *Handbook of Physiology* (Vol. 5, pp. 373-417). Bethesda: American Physiological Society.

Goldman-Rakic, P. S. (1998). The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. In A. C. Roberts, T. W. Robbins & L. Weiskrantz (Eds.), *The Prefrontal Cortex: Executive and Cogntive Functions* (pp. 87-101). New York: Oxford Unviersity Press.

Goodale, M. A., & Milner, A. D. (1992). Separate Visual Pathways for Perception and Action. *Trends in Neurosciences, 15*(1), 20-25.

Gould, R. L., Brown, R. G., Owen, A. M., ffytche, D. H., & Howard, R. J. (2003). fMRI BOLD response to increasing task difficulty during successful paired associates learning. *Neuroimage, 20*(2), 1006-1019.

Gray, J. A. (1981). A critique of Eysenck's theory of personality. In H. J. Eysenck (Ed.), *A model for personality* (pp. 246-276). Berlin: Springer-Verlag.

Green, L., & Myerson, J. (2004). A discounting framework for choice with delayed and probabilistic rewards. *Psychological Bulletin, 130*(5), 769-792.

Green, L., Myerson, J., & McFadden, E. (1997). Rate of temporal discounting decreases with amount of reward. *Memory & Cognition, 25*(5), 715-723.

Green, L., Myerson, J., & Ostaszewski, P. (1999). Amount of reward has opposite effects on the discounting of delayed and probabilistic outcomes. *Journal of Experimental Psychology-Learning Memory and Cognition, 25*(2), 418-427.

Haber, S. N. (2003). The primate basal ganglia: parallel and integrative networks. *Journal of Chemical Neuroanatomy, 26*(4), 317-330.

Haber, S. N., Fudge, J. L., & McFarland, N. R. (2000). Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience, 20*(6), 2369-2382.

Haber, S. N., Kim, K. S., Mailly, P., & Calzavara, R. (2006). Reward-related cortical inputs define a large striatal region in primates that interface with associative cortical connections, providing a substrate for incentive-based learning. *Journal of Neuroscience, 26*(32), 8368-8376.

Haber, S. N., & Knutson, B. (2010). The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology, 35*(1), 4-26.

Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System. *Science, 324*(5927), 646-648.

Herwig, U., Satrapi, P., & Schonfeldt-Lecuona, C. (2003). Using the International 10-20 EEG System for Positioning of Transcranial Magnetic Stimulation. *Brain Topography, 16*(2), 95-99.

Hollerman, J. R., & Schultz, W. (1998). Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience, 1*(4), 304-309.

Hoshi, E., Shima, K., & Tanji, J. (1998). Task-dependent selectivity of movement-related neuronal activity in the primate prefrontal cortex. *Journal of Neurophysiology, 80*(6), 3392-3397.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science, 310*(5754), 1680-1683.

Hyvärinen, J. (1982). *The parietal cortex of monkey and man*. Berlin: Springer-Verlag.

Juckel, G., Schlagenhauf, F., Koslowski, M., Wustenberg, T., Villringer, A., Knutson, B., et al. (2006). Dysfunction of ventral striatal reward prediction in schizophrenia. *Neuroimage, 29*(2), 409-416.

Kable, J. W., & Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience, 10*(12), 1625-1633.

Kable, J. W., & Glimcher, P. W. (2009). The Neurobiology of Decision: Consensus and Controversy. *Neuron, 63*(6), 733-745.

Kable, J. W., & Glimcher, P. W. (2010). An "As Soon As Possible" Effect in Human Intertemporal Decision Making: Behavioral Evidence and Neural Mechanisms. *Journal of Neurophysiology, 103*(5), 2513-2531.

Kahneman, D., & Tversky, A. (1979). Prospect Theory - Analysis of Decision under Risk. *Econometrica, 47*(2), 263-291.

Kane, M. J., & Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: An individual-differences perspective. *Psychonomic Bulletin & Review, 9*(4), 637-671.

Keren, G., & Roelofsma, P. (1995). Immediacy and Certainty in Intertemporal Choice. *Organizational Behavior and Human Decision Processes, 63*(3), 287-297.

Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology-General, 126*(1), 54-70.

Kirby, K. N., & Marakovic, N. N. (1995). Modeling Myopic Decisions - Evidence for Hyperbolic Delay-Discounting within-Subjects and Amounts. *Organizational Behavior and Human Decision Processes, 64*(1), 22-30.

Kirby, K. N., Petry, N. M., & Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology-General, 128*(1), 78-87.

Knight, R. T., & Stuss, D. T. (2002). Prefrontal Cortex: The Present and the Future. In D. T. Stuss & R. T. Knight (Eds.), *Principles of Frontal Lobe Function* (pp. 573-597). New York: Oxford University Press.

Knoch, D., & Fehr, E. (2007). Resisting the power of temptations - The right prefrontal cortex and self-control. *Annals of the New York Academy of Sciences, 1104*, 123-134.

Knoch, D., Gianotti, L. R. R., Pascual-Leone, A., Treyer, V., Regard, M., Hohmann, M., et al. (2006). Disruption of right prefrontal cortex by low-frequency repetitive transcranial magnetic stimulation induces risk-taking behavior. *Journal of Neuroscience, 26*(24), 6469-6472.

Knudsen, E. I. (2007). Fundamental components of attention. *Annual Review of Neuroscience, 30*, 57-78.

Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *Journal of Neuroscience, 21*(16), RC159.

Knutson, B., Bhanji, J. P., Cooney, R. E., Atlas, L. Y., & Gotlib, I. H. (2008). Neural responses to monetary incentives in major depression. *Biological Psychiatry, 63*(7), 686-692.

Knutson, B., Bjork, J. M., Fong, G. W., Hommer, D., Mattay, V. S., & Weinberger, D. R. (2004). Amphetamine modulates human incentive processing. *Neuron, 43*(2), 261-269.

Knutson, B., Fong, G. W., Adams, C. M., Varner, J. L., & Hommer, D. (2001). Dissociation of reward anticipation and outcome with event-related fMRI. *Neuroreport, 12*(17), 3683-3687.

Knutson, B., Fong, G. W., Bennett, S. M., Adams, C. M., & Hommer, D. (2003). A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: characterization with rapid event-related fMRI. *Neuroimage, 18*(2), 263-272.

Knutson, B., & Gibbs, S. E. B. (2007). Linking nucleus accumbens dopamine and blood oxygenation. *Psychopharmacology, 191*(3), 813-822.

Knutson, B., & Greer, S. M. (2008). Anticipatory affect: neural correlates and consequences for choice. *Philosophical Transactions of the Royal Society B-Biological Sciences, 363*(1511), 3771-3786.

Knutson, B., Westdorp, A., Kaiser, E., & Hommer, D. (2000). FMRI visualization of brain activity during a monetary incentive delay task. *Neuroimage, 12*(1), 20-27.

Kobayashi, S., & Schultz, W. (2008). Influence of Reward Delays on Responses of Dopamine Neurons. *The Journal of Neuroscience, 28*(31), 7837-7846.

Koechlin, E., Ody, C., & Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science, 302*(5648), 1181-1185.

Koechlin, E., & Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences, 11*(6), 229-235.

Koenigs, M., Barbey, A. K., Postle, B. R., & Grafman, J. (2009). Superior Parietal Cortex Is Critical for the Manipulation of Information in Working Memory. *Journal of Neuroscience, 29*(47), 14980-14986.

Koopmans, T. C. (1960). Stationary Ordinal Utility and Impatience. *Econometrica, 28*(2), 287-309.

Kuhnen, C. M., & Knutson, B. (2005). The neural basis of financial risk taking. *Neuron, 47*(5), 763-770.

LaBar, K. S., Gitelman, D. R., Parrish, T. B., & Mesulam, M. M. (1999). Neuroanatomic Overlap of Working Memory and Spatial Attention Networks: A Functional MRI Comparison within Subjects. *NeuroImage, 10*, 695-704.

Laird, A. R., Fox, P. M., Price, C. J., Glahn, D. C., Uecker, A. M., Lancaster, J. L., et al. (2005). ALE meta-analysis: Controlling the false discovery rate and performing statistical contrasts. *Human Brain Mapping, 25*(1), 155-164.

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural Representation of Subjective Value Under Risk and Ambiguity. *Journal of Neurophysiology, 103*(2), 1036-1047.

Levy, R., & Goldman-Rakic, P. S. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental Brain Research, 133*(1), 23-32.

Liang, K. Y., & Zeger, S. L. (1986). Longitudinal Data-Analysis Using Generalized Linear-Models. *Biometrika, 73*(1), 13-22.

Loewenstein, G., & Prelec, D. (1992). Anomalies in Intertemporal Choice - Evidence and an Interpretation. *Quarterly Journal of Economics, 107*(2), 573-597.

Logue, A. W. (1988). Research on Self-Control - an Integrating Framework. *Behavioral and Brain Sciences, 11*(4), 665-678.

Lynam, D. R., & Miller, J. D. (2004). Personality pathways to impulsive behavior and their relations to deviance: Results from three samples. *Journal of Quantitative Criminology, 20*(4), 319-341.

Lynam, D. R., Smith, G. T., Whiteside, S. P., & Cyders, M. A. (2006). *The UPPS-P: Assessing five personality pathways to impulsive behavior* (Technical Report). West Lafayette: Purdue University.

MacDonald, A. W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science, 288*(5472), 1835-1838.

Madden, G. J., Begotka, A. M., Raiff, B. R., & Kastern, L. L. (2003). Delay discounting of real and hypothetical rewards. *Experimental and Clinical Psychopharmacology, 11*(2), 139-145.

Madden, G. J., Petry, N. M., Badger, G. J., & Bickel, W. K. (1997). Impulsive and self-control choices in opioid-dependent patients and non-drug-using control participants: Drug and monetary rewards. *Experimental and Clinical Psychopharmacology, 5*(3), 256-262.

Malmo, R. B. (1942). Interference Factors in Delayed Response in Monkeys after Removal of Frontal Lobes. *Journal of Neurophysiology, 5*, 295-308.

Matsumoto, M., & Hikosaka, O. (2009a). Representation of negative motivational value in the primate lateral habenula. *Nature Neuroscience, 12*(1), 77-84.

Matsumoto, M., & Hikosaka, O. (2009b). Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature, 459*(7248), 837-841.

Matthews, S. C., Simmons, A. N., Lane, S. D., & Paulus, M. P. (2004). Selective activation of the nucleus accumbens during risk-taking decision making. *Neuroreport, 15*(13), 2123-2127.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 5. The effect of delay and of intervening events on reinforcement value* (pp. 55-73). Hillsdale, NJ: Erlbaum.

McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience, 27*(21), 5796-5804.

McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science, 306*(5695), 503-507.

Mevorach, C., Hodsoll, J., Allen, H., Shalev, L., & Humphreys, G. (2010). Ignoring the Elephant in the Room: A Neural Circuit to Downregulate Salience. *Journal of Neuroscience, 30*(17), 6072-6079.

Mevorach, C., Humphreys, G. W., & Shalev, L. (2006). Opposite biases in salience-based selection for the left and right posterior parietal cortex. *Nature Neuroscience, 9*(6), 740-742.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24*, 167-202.

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience, 16*(16), 5154-5167.

Milner, B. (1963). Effects of Different Brain Lesions on Card Sorting. *Archives Neurology, 9*(1), 90-100.

Mishkin, M. (1964). Perseveration of central sets after frontal lesions in monkeys. In J. M. Warren & K. Abert (Eds.), *The Frontal Granular Cortex and Behavior* (pp. 219-241). New York: McGraw-Hill.

Mitchell, J. M., Fields, H. L., D'Esposito, M., & Boettiger, C. A. (2005). Impulsive responding in alcoholics. *Alcoholism-Clinical and Experimental Research, 29*(12), 2158-2169.

Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology, 146*(4), 455-464.

Mort, D. J., Malhotra, P., Mannan, S. K., Rorden, C., Pambakian, A., Kennard, C., et al. (2003). The anatomy of visual neglect. *Brain, 126*, 1986-1997.

Mottaghy, F. M., Doring, T., Muller-Gartner, H. W., Topper, R., & Krause, B. J. (2002). Bilateral parieto-frontal network for verbal working memory: an interference approach using repetitive transcranial magnetic stimulation (rTMS). *European Journal of Neuroscience, 16*(8), 1627-1632.

Mottaghy, F. M., Gangitano, M., Sparing, R., Krause, B. J., & Pascual-Leone, A. (2002). Segregation of areas related to visual working memory in the prefrontal cortex revealed by rTMS. *Cerebral Cortex, 12*(4), 369-375.

Moyer, R. S., & Landauer, T. K. (1967). Time Required for Judgements of Numerical Inequality. *Nature, 215*(5109), 1519-1520.

Murphy, J. G., Vuchinich, R. E., & Simpson, C. A. (2001). Delayed reward and cost discounting. *Psychological Record, 51*(4), 571-588.

Myerson, J., & Green, L. (1995). Discounting of Delayed Rewards - Models of Individual Choice. *Journal of the Experimental Analysis of Behavior, 64*(3), 263-276.

Myerson, J., Green, L., Hanson, J. S., Holt, D. D., & Estle, S. J. (2003). Discounting delayed and probabilistic rewards: Processes and traits. *Journal of Economic Psychology, 24*(5), 619-635.

Myerson, J., Green, L., & Warusawitharana, M. (2001). Area under the curve as a measure of discounting. *Journal of Experimental Analysis of Behavior, 76*(2), 235-243.

Nieder, A., & Miller, E. K. (2004). A parieto-frontal network for visual numerical information in the monkey. *Proceedings of the National Academy of Sciences of the United States of America, 101*(19), 7457-7462.

O'Reilly, R. C. (2010). The What and How of prefrontal cortical organization. *Trends in Neurosciences, 33*(8), 355-361.

Ostaszewski, P., & Karzel, K. (2005). Temperament and discounting of delayed and probabilistic losses. *Polish Psychological Bulletin, 36*(2), 117-126.

Pandya, D. N., & Seltzer, B. (1982). Intrinsic Connections and Architectonics of Posterior Parietal Cortex in the Rhesus-Monkey. *Journal of Comparative Neurology, 204*(2), 196-210.

Panksepp, J., Knutson, B., & Burgdorf, J. (2002). The role of brain emotional systems in addictions: a neuro-evolutionary perspective and new 'self-report' animal model. *Addiction, 97*(4), 459-469.

Patton, J. H., Stanford, M. S., & Barratt, E. S. (1995). Factor structure of the Barratt Impulsiveness Scale. *Journal of Clinical Psychology, 51*(6), 768-774.

Perenin, M. T., & Vighetto, A. (1988). Optic Ataxia - a Specific Disruption in Visuomotor Mechanisms .1. Different Aspects of the Deficit in Reaching for Objects. *Brain, 111*, 643-674.

Perret, E. (1974). Left Frontal Lobe of Man and Suppression of Habitual Responses in Verbal Categorical Behavior. *Neuropsychologia, 12*(3), 323-330.

Peters, J., & Buchel, C. (2009). Overlapping and Distinct Neural Systems Code for Subjective Value during Intertemporal and Risky Decision Making. *Journal of Neuroscience, 29*(50), 15727-15734.

Peters, J., & Buchel, C. (2010). Episodic Future Thinking Reduces Reward Delay Discounting through an Enhancement of Prefrontal-Mediotemporal Interactions. *Neuron, 66*(1), 138-148.

Petrides, M. (1998a). Specialized systems for the processing of mnemonic information within the primate frontal cortex. In A. C. Roberts, T. W. Robbins & L. Weiskrantz (Eds.), *The Prefrontal Cortex: Executive and Cognitive Functions* (pp. 103-115). New York: Oxford University Press.

Petrides, M. (1998b). Specialized systems for the processing of mnemonic information within the primate frontal cortex. In A. C. Roberts, T. W. Robbins & L. Weiskrantz (Eds.), (pp. 103-116). Oxford: Oxford University Press.

Petrides, M. (2005). Lateral prefrontal cortex: architectonic and functional organization. *Philosophical Transactions of the Royal Society B-Biological Sciences, 360*(1456), 781-795.

Petrides, M., Alivisatos, B., & Evans, A. C. (1995). Functional Activation of the Human Ventrolateral Frontal Cortex during Mnemonic Retrieval of Verbal Information. *Proceedings of the National Academy of Sciences of the United States of America, 92*(13), 5803-5807.

Petrides, M., & Pandya, D. (2002). Association Pathways of the Prefrontal Cortex and Functional Observations. In D. T. Stuss & R. T. Knight (Eds.), *Principles of Frontal Lobe Function* (pp. 31-50). New York: Oxford University Press.

Petrides, M., & Pandya, D. N. (1994). Comparative architectonic analysis of the human and the macaque frontal cortex. In F. Boller & J. Grafman (Eds.), *Handbook of neuropsychology* (Vol. 9, pp. 17-58). Amsterdam: Elsevier.

Petry, N. M. (2001). Delay discounting of money and alcohol in actively using alcoholics, currently abstinent alcoholics, and controls. *Psychopharmacology, 154*(3), 243-250.

Petry, N. M. (2002). Discounting of delayed rewards in substance abusers: relationship to antisocial personality disorder. *Psychopharmacology, 162*(4), 425-432.

Phelps, E. S., & Pollak, R. (1968). On Second-Best National Saving and Game-Equilibrium Growth. *Review of Economic Studies, 35*, 185-199.

Piazza, M., Izard, V., Pinel, P., Le Bihan, D., & Dehaene, S. (2004). Tuning curves for approximate numerosity in the human intraparietal sulcus. *Neuron, 44*(3), 547-555.

Pindyck, R. S., & Rubinfeld, D. L. (2006). *Microeconomics* (6th ed.). New Delhi: Prentice-Hall of India Private Ltd.

Pine, A., Seymour, B., Roiser, J. P., Bossaerts, P., Friston, K. J., Curran, H. V., et al. (2009). Encoding of Marginal Utility across Time in the Human Brain. *Journal of Neuroscience, 29*(30), 9575-9581.

Plassmann, H., O'Doherty, J., & Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *Journal of Neuroscience, 27*(37), 9984-9988.

Plassmann, H., O'Doherty, J. P., & Rangel, A. (2010). Appetitive and Aversive Goal Values Are Encoded in the Medial Orbitofrontal Cortex at the Time of Decision Making. *Journal of Neuroscience, 30*(32), 10799-10808.

Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective Probability and Delay. *Journal of the Experimental Analysis of Behavior, 55*(2), 233-244.

Rao, S. C., Rainer, G., & Miller, E. K. (1997). Integration of what and where in the primate prefrontal cortex. *Science, 276*(5313), 821-824.

Raz, A., & Buhle, J. (2006). Typologies of attentional networks. *Nature Reviews Neuroscience, 7*(5), 367-379.

Read, D., Frederick, S., Orsel, B., & Rahman, J. (2005). Four score and seven years from now: The date/delay effect in temporal discounting. *Management Science, 51*(9), 1326-1335.

Reynolds, B., Ortengren, A., Richards, J. B., & de Wit, H. (2006). Dimensions of impulsive behavior: Personality and behavioral measures. *Personality and Individual Differences, 40*(2), 305-315.

Reynolds, B., Richards, J. B., & de Wit, H. (2006). Acute-alcohol effects on the Experiential Discounting Task (EDT) and a question-based measure of delay discounting. *Pharmacology Biochemistry and Behavior, 83*(2), 194-202.

Richards, J. B., Zhang, L., Mitchell, S. H., & de Wit, H. (1999). Delay or probability discounting in a model of impulsive behavior: Effect of alcohol. *Journal of the Experimental Analysis of Behavior, 71*(2), 121-143.

Rizzolatti, G., & Matelli, M. (2003). Two different streams form the dorsal visual system: anatomy and functions. *Experimental Brain Research, 153*(2), 146-157.

Robertson, E. M., Theoret, H., & Pascual-Leone, A. (2003). Studies in cognition: The problems solved and created by transcranial magnetic stimulation. *Journal of Cognitive Neuroscience, 15*(7), 948-960.

Rogers, R. D., Owen, A. M., Middleton, H. C., Williams, E. J., Pickard, J. D., Sahakian, B. J., et al. (1999). Choosing between small, likely rewards and large, unlikely rewards activates inferior and orbital prefrontal cortex. *Journal of Neuroscience, 19*(20), 9029-9038.

Rossi, S., Hallett, M., Rossini, P. M., & Pascual-Leone, A. (2009). Safety, ethical considerations, and application guidelines for the use of transcranial magnetic stimulation in clinical practice and research. *Clinical Neurophysiology, 120*, 2008-2039.

Rothschild, M., & Stiglitz, J. (1970). Increasing risk: a definition. *Journal of Economic Theory, 2*, 225-243.

Russell, J. A. (1980). A Circumplex Model of Affect. *Journal of Personality and Social Psychology, 39*(6), 1161-1178.

Samanez-Larkin, G. R., Gibbs, S. E. B., Khanna, K., Nielsen, L., Carstensen, L. L., & Knutson, B. (2007). Anticipation of monetary gain but not loss in healthy older adults. *Nature Neuroscience, 10*(6), 787-791.

Samuelson, P. (1937). A Note on Measurement of Utility. *Review of Economic Studies, 4*, 155-161.

Sandrini, M., Rossini, P. M., & Miniussi, C. (2004). The differential involvement of inferior parietal lobule in number comparison: a rTMS study. *Neuropsychologia, 42*(14), 1902-1909.

Schultz, W. (2006). Behavioral theories and the neurophysiology of reward. *Annual Review of Psychology, 57*, 87-115.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science, 275*(5306), 1593-1599.

Takayama, Y., Sugishita, M., Akiguchi, I., & Kimura, J. (1994). Isolated Acalculia Due to Left Parietal Lesion. *Archives of Neurology, 51*(3), 286-291.

Thaler, R. (1991). Some Empirical Evidence on Dynamic Inconsistency. In R. Thaler (Ed.), *Quasi-rational economics* (pp. 127-133). New York, N.Y.: Russell Sage Foundation.

Tobler, P. N., Christopoulos, G. I., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2008). Neuronal Distortions of Reward Probability without Choice. *Journal of Neuroscience, 28*(45), 11703-11711.

Tobler, P. N., Christopoulos, G. I., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2009). Risk-dependent reward value signal in human prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America, 106*(17), 7185-7190.

Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science, 307*(5715), 1642-1645.

Tobler, P. N., O'Doherty, J. P., Dolan, R. J., & Schultz, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *Journal of Neurophysiology, 97*(2), 1621-1632.

Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature, 428*(6984), 751-754.

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science, 315*(5811), 515-518.

Tversky, A., & Kahneman, D. (1981). The Framing of Decisions and the Psychology of Choice. *Science, 211*(4481), 453-458.

Tversky, A., & Kahneman, D. (1992). Advances in Prospect-Theory - Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty, 5*(4), 297-323.

Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale & R. J. W. Mansfield (Eds.), *Analysis of Visual Behavior* (pp. 549-586). Cambridge: MIT Press.

Vallar, G., & Perani, D. (1987). The anatomy of spatial neglect in humans. In M. Jeannerod (Ed.), *Neurophysiological and neuropsychological aspects of spatial neglect* (pp. 235-258). Amsterdam: North-Holland Elsevier Science.

Vendrell, P., Junque, C., Pujol, J., Jurado, M. A., Molet, J., & Grafman, J. (1995). The Role of Prefrontal Regions in the Stroop Task. *Neuropsychologia, 33*(3), 341-352.

Von Bonin, G., & Bailey, P. (1947). *The Neocortex of Macaca Mulatta*. Urbana: University of Illinois Press.

Von Economo, C. (1929). *The Cytoarchitectonics of the Human Cerebral Cortex*. London: Oxford University Press.

Wager, T. D., Jonides, J., & Reading, S. (2004). Neuroimaging studies of shifting attention: a meta-analysis. *Neuroimage, 22*(4), 1679-1693.

Wager, T. D., & Smith, E. E. (2003). Neuroimaging studies of working memory. *Cognitive, Affective, & Behavioral Neuroscience, 3*(4), 255-274.

Wagner, T., Rushmore, J., Eden, U., & Valero-Cabre, A. (2009). Biophysical foundations underlying TMS: Setting the stage for an effective use of neurostimulation in the cognitive neurosciences. *Cortex, 45*(9), 1025-1034.

Weber, B. J., & Huettel, S. A. (2008). The neural substrates of probabilistic and intertemporal decision making. *Brain Research, 1234*, 104-115.

Weber, E. U., Blais, A. R., & Betz, N. E. (2002). A domain-specific risk-attitude scale: Measuring risk perceptions and risk behaviors. *Journal of Behavioral Decision Making, 15*(4), 263-290.

Weber, E. U., & Johnson, E. J. (2009). Decisions Under Uncertainty: Psychological, Economic, and Neuroeconomic Explanations of Risk Preference. In P. W. Glimcher, C. F. Camerer, E. Fehr & R. A. Poldrack (Eds.), *Neuroeconomics Decision Making and the Brain* (pp. 127-144). Amsterdam: Elsevier.

White, I. M., & Wise, S. P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. *Experimental Brain Research, 126*(3), 315-335.

Whiteside, S. P., & Lynam, D. R. (2001). The Five Factor Model and impulsivity: using a structural model of personality to understand impulsivity. *Personality and Individual Differences, 30*(4), 669-689.

Wittmann, M., Leland, D. S., & Paulus, M. P. (2007). Time and decision making: differential contribution of the posterior insular cortex and the striatum during a delay discounting task. *Experimental Brain Research, 179*(4), 643-653.

Wittmann, M., Lovero, K. L., Lane, S. D., & Paulus, M. P. (2010). Now or Later? Striatum and Insula Activation to Immediate Versus Delayed Rewards. *Journal of Neuroscience, Psychology, and Economics, 3*(1), 15-26.

Wrase, J., Kahnt, T., Schiagenhauf, F., Beck, A., Cohen, M. X., Knutson, B., et al. (2007). Different neural systems adjust motor behavior in response to reward and punishment. *Neuroimage, 36*(4), 1253-1262.

Xu, L. J., Liang, Z. Y., Wang, K., Li, S., & Jiang, T. Z. (2009). Neural mechanism of intertemporal choice: From discounting future gains to future losses. *Brain Research, 1261*, 65-74.