AN INVESTIGATION OF STRIATAL ACTIVITY UNDER DELAYED AND EFFORT-BASED LEARNING

by

EKATERINA DOBRYAKOVA

A Dissertation submitted to the

Graduate School-Newark

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Psychology

written under the direction of

Professor Elizabeth Tricomi

and approved by

Dr. Mei-Fang, Cheng

_____

Dr. Mauricio R. Delgado

_____

Dr. Catherine Myers

_____

Dr. Michael Shiflett

_____

Dr. Elizabeth Tricomi

_____

Newark, New Jersey
October, 2012

ABSTRACT OF THE DISSERTATION

An Investigation of Striatal Activity during Delayed and Effort-based Learning

By EKATERINA DOBRYAKOVA

Dissertation Director

Elizabeth Tricomi

Motivation influences human learning and outcome valuation. Depending on the context, one can interpret an outcome in a positive way or not pay attention to the action outcome at all. The striatum is one of the primary structures involved in outcome valuation and learning and of action-outcome contingencies. Striatal activity has been shown to be context-dependent and to reflect individuals' subjective preferences. This dissertation examined striatal activity in the context of delayed and effort-based learning, as well as whether people are willing to overcome effort costs in order to benefit an unfamiliar disadvantaged person. Two functional magnetic resonance imaging experiments were conducted examining striatal activity during performance-related feedback under different time frames (Experiment 1) and following different cognitive effort requirements (Experiment 2). Behavioral Experiment 3 looked at whether individuals are willing to exert cognitive effort during learning to reduce inequity between themselves and a disadvantaged individual. Experiment 1 replicated previous findings of ventral striatal activation to immediate feedback presentation. It was also shown that when feedback is presented after a substantial delay of 25 minutes, processing of feedback switches away from the striatum to posterior parts of the basal ganglia. Experiment 2 revealed that activity of the ventral striatum associated with feedback reflects effort expenditure required to obtain it. Experiment 3 showed that unfair social context can

motivate individuals to exert cognitive effort during learning. This work shows that striatal response to learning outcomes is differentially influenced by delay and effort requirements and that effort costs can motivate learning.

Acknowledgments

First and foremost, I would like to thank my advisor, Elizabeth Tricomi, for guiding me through this difficult but enriching process. She is the best mentor that I had, with patience and kindness that I never knew could exist. I am forever grateful for her contributions to my scientific development. I would also like to thank my committee members, Mauricio Delgado, Mei Cheng, Michael Shiflett, and Catherine Myers, for sharing their interesting ideas, giving practical suggestions regarding my experiments and for providing valuable career guidance. My lab mates, past and present, were also of great help during my graduate career. I specifically would like to thank Samantha DePasque Swanson, Karolina Lempert and Holly Sullivan Toole for numerous moments of brainstorming together and for helping me maintain my sanity during difficult moments. I have been lucky to meet Olga Boukrina who is a close and supportive friend and colleague. I would like to also acknowledge my parents, Irina Dobryakova and Roman Amiragov, for always believing in me and supporting every decision of mine. In conclusion, I thank my husband and my brothers, Andrey Kravtsov, Gerald and Gurgen Amiragov, for being there for me.

Table of Contents

List of Figures

List of Tables

An Investigation of Striatal Activity during Delayed and Effort-based Learning

Chapter One: General introduction

One's level of motivation can greatly impact the learning process and the subjective weight carried by outcomes of one's actions. For a motivated learner, an outcome is a critical component of the learning process since it allows modifying the behavior and achieving a specific goal. For example, a student with a goal of getting a good grade on a test knows from her previous experience that she has to exert effort in order to get a good grade on a test. The student finds out the test result some time later and has to pay attention to the feedback in order to better prepare for the next test or learn that the current strategy leads to success. Learning such action-outcome contingencies requires weighing expected benefits (getting an 'A' vs. getting a 'C') and calculating the costs that an action requires (time and resources pooled while studying). But would the effort spent on studying influence how pleasant the outcome is and would the delay influence how well the information from feedback is learned? The striatum, which is innervated by mid-brain dopamine projection neurons, is one of the critical structures involved in the process of learning from outcomes and in cost-benefit valuation (Assadi, Yucel, & Pantelis, 2009; Botvinick, Huffstetler, & McGuire, 2009; Cocker, Hosking, Benoit, & Winstanley, 2012; Croxson, Walton, O'Reilly, Behrens, & Rushworth, 2009; Floresco, Onge, Ghods-Sharifi, & Winstanley, 2008; Haber & Calzavara, 2009; Klapproth, 2008; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). Therefore, it is important to investigate how striatal activity is affected during learning motivated by different action costs. This dissertation investigates motivational influences

on learning and striatal activity associated with the learning from delayed and effortful outcomes.

Action costs can influence outcome valuation in several ways. For example, costs can come in the form of a delay or increased effort requirements. Therefore, during the course of three experiments, I investigated: 1) striatal activity associated with delay-based learning context, 2) striatal activity associated with effort-based learning context and 3) effort-based learning in a social context. The first two experiments address the functional role of the striatum under delay and effort-based learning conditions. The third experiment was a behavioral experiment examining subject's learning in a social context.

Basal ganglia anatomy

The striatum is part of a small but complex aggregation of nuclei – the basal ganglia (BG). The BG nuclei were previously thought to be responsible primarily for motor function control (Middleton & Strick, 2000). However, today there is plenty of evidence that the BG also contribute to various cognitive and emotional functions and project to cortical areas that share these functional roles with the BG. Therefore, the current proposed functional role of the BG is learning and control of complex behaviors (Graybiel, 2005; Haber, Fudge, & McFarland, 2000; Haber & Knutson, 2010).

The BG consist of the striatum, the globus pallidus, the substantia nigra and the subthalamic nucleus. Evidence from tracing studies shows that the BG are interconnected with functionally specific areas of the prefrontal cortex (PFC) and the thalamus, where specific projections are involved in initiation of movement, emotion and cognition (Haber & Calzavara, 2009).

The striatum can be subdivided onto the dorsal and the ventral striatum. The dorsal striatum includes the body and the dorsal portion of the head of the caudate nucleus and the dorsal putamen, while the ventral striatum consist of the ventral portion of the head of the caudate nucleus, the nucleus accumbens (NAcc) and the anterior putamen (Delgado, 2007). The putamen and dorsal caudate nucleus receive projections from the motor and premotor cortex. However, trace studies show that the caudate nucleus and the putamen are also connected with the dorsolateral prefrontal cortex (DLPFC). This structure plays a role in working-memory processes and is also important in action planning and information integration (Haber, Fudge, & McFarland, 2000). Therefore, the lateral portion of the dorsal striatum is interconnected with motor areas, while the dorsal portion is interconnected with cognitive areas of the cortex. The ventral striatum receives limbic projections and is interconnected with the ventromedial prefrontal cortex (VMPFC) and the anterior cingulate cortex (ACC) (Haber, Fudge, & McFarland, 2000; Haber, Kim, Mailly, & Calzavara, 2006). Even though it is possible to categorize prefrontal projections to the striatum onto those that are responsible for emotion or memory processing, there is also a large degree of convergence of cortical inputs. That is, projections from cognitive regions of the cortex converge with cortical limbic projections in the striatum, making the striatum particularly suitable for representing various aspects of outcome-based learning (Haber & Calzavara, 2009; Haber, Kim, Mailly, & Calzavara, 2006).

Striatal involvement in learning and outcome costs

The evidence from tracing studies cited above are in line with human and animal studies that show striatal involvement, highlighting the role of dopamine in learning of

action-outcome contingencies (Assadi, Yucel, & Pantelis, 2009; O'Doherty et al., 2004; Shohamy, 2011; Tricomi & Fiez, 2012; Williams & Eskandar, 2006). An action-outcome contingency representation is necessary for goal-directed but not habitual action (Tricomi, Balleine, & O'Doherty, 2009). Goal-directed actions can result in immediate effortless outcomes, or might require high amount of effort and/or involve a delay, where after performing an action the outcome is not evident for a specific amount of time. Striatal activity has been shown to be sensitive to such action costs (Croxson, Walton, O'Reilly, Behrens, & Rushworth, 2009; Foerde & Shohamy, 2011).

Delay. The first evidence of timing between the action and the outcome being important in the process of reward-dependent learning comes from original findings on dopaminergic cells' involvement in this process. Apicella, Shultz and others (1992) showed that, during learning the dopamine neurons that project to the striatum increase their firing rate to the presentation of a juice reward and cease their firing rate during outcome omission (Apicella, Scarnati, Ljungberg, & Schultz, 1992; Schultz, 2002; Schultz, Apicella, & Ljungberg, 1993). Hence, it was hypothesized that when there is a temporal gap between the action and the outcome the dopamine signal will not bind the two events together (Cheung & Cardinal, 2005; Maddox, Ashby, & Bohil, 2003). Of note, the differential pattern of firing was not present anymore as the animal learned the action-outcome contingency suggesting that it is the predictive nature of the dopamine signal that aids in learning via updating of subject's expectations about actual outcomes. Neuroimaging studies replicated the pattern of activity observed in animal studies, showing differential activation in the human striatum to presentation of monetary outcomes (gain versus loss) (Delgado, Nystrom, Fissell, Noll, & Fiez, 2000). A similar

pattern of activity was observed during learning through performance-related feedback (correct versus incorrect) that reflected participants' learning success (Tricomi & Fiez, 2008)

Further evidence of the importance of time in dopamine-dependent learning comes from Parkinson's disease (PD) patients. PD patients are impaired at learning through immediate outcomes due to dopamine deficiency in the substantia nigra (SN) that develops during the disease. The SN does not supply enough dopamine into the striatum leading to impairments of PD patients in learning through immediate feedback. However, these patients are able to learn from outcomes if they are presented with a delay of several seconds or when they just observe the stimulus-outcome sequence (Foerde & Shohamy, 2011; Grahn, Parkinson, & Owen, 2008; Shohamy et al., 2004). During these types of learning other non-dopamine-dependent systems are able to compensate the deficiency in the striatal learning system (Voermans et al., 2004). At the same time, it is still unknown how and if the striatum will be engaged when the outcome is substantially delayed, as in the example with a student's test, presented above.

Activity in the striatum has also been shown to track actual magnitude of immediate outcomes (Delgado, Locke, Stenger, & Fiez, 2003; Ino, Nakai, Azuma, Kimura, & Fukuyama, 2010; Pedroni, Koeneke, Velickaite, & Jancke, 2011). That is, striatal blood-oxygen level dependent (BOLD) activity showed a greater increase from baseline in association with the presentation of a large monetary gain and a larger decrease from baseline in association with a large monetary loss. At the same time, smaller monetary gain and loss outcomes resulted in an intermediate BOLD signal

(Delgado, Locke, Stenger, & Fiez, 2003). Therefore, striatal activity tracks objective value of outcomes.

However, the striatum seems to be also sensitive to the subjective outcome value. Studies of choice preference show that human and animal subjects, when given a choice between two outcomes of comparable value, prefer the ones that can be obtained immediately rather than outcomes obtained after a delay (Assadi, Yucel, & Pantelis, 2009; Floresco, Tse, & Ghods-Sharifi, 2008; Frederick, Loewenstein, & O'Donoghue, 2002; Green & Myerson, 2004). Neuroimaging studies, in particular, show that the striatum reflects subjective choice preferences by exhibiting decreased activation in association with delayed outcomes (Ballard & Knutson, 2009; Kable & Glimcher, 2007, 2010; Prevost, Pessiglione, Metereau, Clery-Melin, & Dreher, 2010).

Effort. Studies that looked at effort-based decision-making also suggest that subjects' choice tendencies involving effortful actions are similar to choice tendencies with delayed outcomes. That is, given a choice between two similar outcomes, the outcome that requires less effort is preferred (Botvinick, Huffstetler, & McGuire, 2009; Botvinick & Rosen, 2009; Prevost, Pessiglione, Metereau, Clery-Melin, & Dreher, 2010). Striatal activation decreases in association with outcomes that require high effort and increases in association with outcomes that require low effort expenditure. Therefore, neuroimaging evidence is in line with the notion that effort is disutilitarian; i.e., it carries a negative weight.

However, neuroimaging studies have not looked at how striatal activity tracks cost-benefit valuation in the context of learning. Previous studies focused primarily on how outcomes are valued during choice. For a motivated learner the effort might still be

disutilitarian, and if given a choice, the learner indeed might choose an easier task. But when the choice is not available and one has to master a given task, the subjective reward of mastering a more difficult task might be greater than the subjective reward felt after mastering an easier task.

The importance of the striatum in outcome valuation is also evident from animal studies that show that direct dopamine depletion from the striatum influences subjects' outcome choice that require effort expenditure. Subjects become less motivated to act in a beneficial manner and work to obtain a larger outcome. Administration of dopamine antagonists has similar effects on behavior (Salamone, Correa, Mingote, & Weber, 2003; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006).

Effort in a social context. From an evolutionary and economic perspective, one should act in a self-benefiting fashion (Fehr & Schmidt, 1999; Kogut, 2012). Based on this principle, one should avoid effortful actions and prefer outcomes that are easily obtained and increase one's overall payoffs. Such outcomes should carry greater value than outcomes that reflect someone else's gain. That is, valuable outcomes are those that are relevant to the self and not to another person. For example, a student would be a lot happier to receive an 'A' on a test than if a friend receives the same grade. In addition, a student will not be willing to spend his or her own time studying for a test so that another person can receive a grade for it. Such behavior would be irrational.

However, in a social context people have been shown to behave irrationally. For example, people engage in altruistic punishment, i.e. they sacrifice their own payoffs in order to punish a violator of social norms (Fehr & Fischbacher, 2003; Lee, 2008). This is one of the examples when an action is costly. From an economic perspective, one should

not act in a costly manner and should not care about the actions of a violator if those actions do not affect one's payoffs. However, in a social context, acting according to the social norms and actions that restore fairness seem to dominate the situation and affect the behavior and outcome value. Instead of preferring an outcome that maximizes own gains, people prefer the outcomes that restore fairness and go in accordance with social norms. One idea that has not been explored is whether learning is also affected by a social context.

Learning context and extrinsic versus intrinsic outcomes

As previously mentioned, neuroimaging studies showed that the pattern of striatal activation during the presentation of intrinsic outcomes is similar to that observed during extrinsic outcome presentation (Tricomi, Delgado, McCandliss, McClelland, & Fiez, 2006; Tricomi & Fiez, 2008). Whether there is a cost of delay or effort, subjective valuation of an extrinsic outcome decreases leading to a decrease in striatal activity (Botvinick, Huffstetler, & McGuire, 2009; Kable & Glimcher, 2007). Extrinsic outcomes are primary and secondary rewards, such as food or water and money, respectively (Linke et al., 2010). The motivation with primary extrinsic outcomes is to quench thirst or hunger, which is an internal drive and has to be satisfied for survival. At the same time, the motivation with the secondary extrinsic outcome is to earn more money and to maximize one's benefits.

Intrinsic outcomes are not motivated by tangible objects as extrinsic outcomes are. For example, one might want to get a good grade for a class or gain popularity and get praised for good performance in a sporting competition. In experimental settings intrinsic outcomes are presented in the form of performance-related feedback that reflects

to a participant how successful he or she is on a task. Even though the striatum shows a similar pattern of activity to the presentation of extrinsic and intrinsic rewards, intrinsic rewards are more subjective than extrinsic rewards. Hence, their value is more prone to be modified by the context in which such outcomes are delivered.

It has also been shown that striatal activity is context-dependent (Nieuwenhuis et al., 2005; Tricomi & Fiez, 2012). Activity of the dorsal striatum (i.e. caudate nucleus and putamen) can be greatly influenced by the informational context carried by the outcome (Tricomi & Fiez, 2008, 2012). When performance-related feedback is presented for the same stimulus but carries different information, striatal BOLD reflects the potency of information carried by feedback. Activity of the ventral striatum has also been shown to be sensitive to subjective outcome valuation (Kable & Glimcher, 2007, 2010). Even animal studies present evidence of context dependency of learning mechanisms in the striatum. For example, indirect administration of d-amphetamine, a dopamine antagonist, decreases rats' motivation to wait for a larger food reward in a delay-based task when the rewarded outcome is cued, but it has an opposite effect on animal's behavior when the outcome is not cued (Cardinal, Robbins, & Everitt, 2000).

Given the context of learning and a goal of achieving task success, the delay and effort demands might not influence the value of intrinsic outcomes in a negative way and the outcomes might still be subjectively valuable. That is, in the context of delay-based learning, delaying the outcome might result in learning and activate the striatal system because the outcome would still carry the information about task success and indicate ways for improvement. Similarly, in the context of effort-based learning, the subjective value of an outcome that indicates task success might be subjectively higher if it follows

an effortful task rather than an easy task because outcomes after an effortful action might be more intrinsically motivating to some learners. In an unfair social context, regardless of the disutility of effort, one might still want to perform an effortful action on behalf of another person in order to benefit that person and restore fairness. This dissertation explores how the context of delay and effort influences striatal activity during learning and whether inequity aversion can motivate effortful learning behavior.

Intrinsic feedback is used in all three experiments, the interpretation of which depends on participants' motivation. These studies attempt to further demonstrate how, under specific motivational context, the subjective value of outcomes can change and how the striatal activity tracks outcome value during delay and effort-based learning.

Chapter Two: Experiment 1: Basal Ganglia Engagement during Feedback Processing after a Substantial Delay

Numerous neurophysiological and neuroimaging studies implicate the striatum, the input unit of the basal ganglia, as a structure important for reward processing (Delgado, 2007; Hikosaka, Nakamura, & Nakahara, 2006; McClure, York, & Montague, 2004). The striatum is a major target of midbrain dopamine neurons, which code for prediction errors with a phasic increase in activity in response to unexpected rewards and a decrease below baseline during the omission of expected rewards (Schultz, 2002; Schultz, 2010). The blood oxygen level dependent (BOLD) signal in the striatum also shows greater response to rewards than punishments, including positive versus negative performance-related feedback (Elliott, Frith, & Dolan, 1997; Schultz, Apicella, & Ljungberg, 1993; Seger, 2008; Seger, Peterson, Cincotta, Lopez-Paniagua, & Anderson, 2010; Tricomi, Delgado, McCandliss, McClelland, & Fiez, 2006; Tricomi & Fiez, 2008).

In most of these studies the outcome is presented immediately after the response of the subject or after a short delay of a few seconds (Delgado, 2007; Delgado, Nystrom, Fissell, Noll, & Fiez, 2000; Foerde & Shohamy, 2011; Haber & Knutson, 2010; Kobayashi & Schultz, 2008; Maddox, Ashby, & Bohil, 2003; Schultz, 2002, 2010; Tricomi & Fiez, 2008). This suggests that the striatum is involved in learning when the action and an outcome are in close temporal proximity.

Previous studies emphasize the importance of close temporal proximity of the response and the corresponding feedback during dopamine-dependent learning due to a rapid degradation of the dopamine signal (Foerde & Shohamy, 2011; Maddox, et al., 2003). The dopamine signal is thought to strengthen the link between the original action and its outcome (Maddox, et al., 2003). In addition, if there are many intervening events between the action and the delayed outcome, the dopamine signal might not strengthen the association between the specific action and the appropriate outcome (Cardinal, 2006; Cheung & Cardinal, 2005; Kurniawan, Guitart-Masip, & Dolan, 2011). In daily life, the delay between an action and an outcome can be much longer, and yet humans and animals are still able to learn the association between specific outcomes and the actions that produced them (Cardinal, 2006; Foerde & Shohamy, 2011). For example, in academic testing situations, there is usually a time period between test submission and feedback about test performance. It is still unclear, however, whether learning from substantially delayed feedback engages similar brain structures and leads to similar performance compared with learning from immediate feedback.

In the current study, we investigated whether the neural substrates underlying learning from delayed feedback are similar to those underlying learning from immediate

feedback. Specifically, we wished to examine the effects of presenting subjects with performance-related feedback after a substantial delay of around 25 minutes, with many intervening events, on activity in the striatum and other brain structures. We used as a model an academic testing situation and a feedback-based word association task similar to one used previously to examine striatal activity following immediate feedback (Tricomi & Fiez, 2008). Participants engaged in a study phase performed outside the scanner, followed by fMRI data acquisition as they performed a multiple-choice test of their memory. For some trials, subjects received immediate feedback, but for others, they did not receive feedback until the trials were presented during a second, review phase. As with an academic test, they were shown their previous responses and feedback about whether they were correct. This paradigm allowed us to test several alternative hypotheses. First, delay could have no effect on the pattern of neural activity observed following feedback. In this case, positive feedback should produce stronger activation than negative feedback, regardless of whether the feedback is presented immediately or after a substantial delay. Alternatively, a substantial delay could alter neural processing of feedback. For example, unlike immediate feedback, positive and negative feedback might not produce differential activation in the striatum after a delay. Finally, learning from delayed feedback might be dependent on different neural structures than learning from immediate feedback. For example, recent evidence has suggested that the medial temporal lobe (MTL) may be particularly important for learning from delayed feedback (Foerde & Shohamy, 2011).

Thus, the aim of the current study was to investigate whether subjects would be able to learn from delayed feedback presentation to the same degree as from immediate

feedback, and to map brain regions responsible for the processing of the feedback presentation after a substantial delay.

Methods

Participants

Twenty-four right-handed individuals consented to participate in the experiment for a payment of $50. Four participants were not included in the main analysis due to technical problems and one participant was excluded due to excessive motion. Therefore, data from 19 participants were analyzed (11 females; mean age 23.89 years, SD 3.14). The research was approved by the Institutional Review Boards of Rutgers University and the University of Medicine and Dentistry of New Jersey (UMDNJ).

An additional set of 21 Rutgers University students were recruited for a behavioral version of the experiment after the recruitment for the fMRI experiment was completed in order to gather self-report data. All individuals consented to participate in the experiment for research credit. All of the participants were fluent in English.

Materials

A 3-Tesla Siemens Allegra scanner was used to acquire all fMRI data. Behavioral data acquisition and stimulus presentation was administered using the "E-Prime" software (Schneider, Eschman, & Zuccolotto, 2002).

Procedure

Scan session. A T1-weighted pulse sequence was used to collect structural images in 43 contiguous slices (3x3x3 mm voxels) tilted 30° from the AC-PC line (Deichmann, Gottfried, Hutton, & Turner, 2003). Similarly, 43 functional images were collected using

a single-shot echo EPI sequence amounting to 172 acquisitions (TR= 2500 ms, TE = 25 ms, FOV = 192 mm, flip angle = 80°).

Behavioral paradigm. In this study, participants had to perform a paired-associate word learning task (Tricomi & Fiez, 2008). The words used in the experiment contained 4–8 letters and 1–2 syllables, had Kucera-Francis frequencies of 20–650 words per million, and had high imagibility ratings (score of over 400 according to the MRC database) (Coltheart, 1981). The words were matched for word length and frequency at the trial level. Words presented on the same trial were not semantically related, with a score of less than 0.2 on the Latent Semantic Analysis similarity matrix (Landauer, Foltz, & Laham, 1998) and did not rhyme or begin with the same letter.

At the beginning of the experiment, participants took part in a study phase outside of the scanner, during which they acquired word associations (180 trials). This was done in order for participants to acquire initial learning that would be further augmented via feedback presentation. In addition, it has been shown that the striatum is differentially activated when feedback is informative of one's performance, but not when feedback is only arbitrarily related to one's responses, prior to learning (Tricomi & Fiez, 2008).

The format of experimental trials resembled multiple-choice test questions. That is, on each trial of the study phase, participants were presented with three words, where the top word was the main word with two word options underneath. One of the options was highlighted in green, indicating that this option was the correct match for the main word. Participants were instructed to memorize the main word and the associated highlighted option. Trials were presented in random order for the duration of four seconds, separated by a fixation point lasting for three seconds.

The words learned during the study phase were then randomly assigned to the three feedback presentation conditions (immediate, delayed and no feedback conditions) and were presented during six scanning sessions, each lasting seven minutes (Figure 1a). The conditions were presented randomly in blocks of 10 trials. Each block was separated from the next block by a jittered fixation point (1-5 sec). Each trial lasted approximately 8 seconds and started with a jittered fixation point (1-5 sec) that also contained a label that informed participants about the type of the feedback condition. That is, participants, were informed that they would be presented with 10 trials of each condition in random order and they were reminded about the condition they were doing by the label with the fixation point that preceded each trial. The trial order was fixed, so that the length of time between the presentation of the initial trial and the corresponding delayed feedback event would not vary.

During Scanning Phase 1 (the first three scanning sessions), participants had to select, by pressing a button, one of the options as a match for the main word based on what they remembered from the study phase. Feedback, which reflected whether participants selected a correct match for the main word (green √; red X), was presented for trials in the immediate feedback condition during Scanning Phase 1 (scan sessions one through three). For trials in the other two conditions, participants were not informed whether they selected a correct option, and they were presented with a control screen that showed a black pound sign (#) instead of the feedback. Both feedback and control screens were presented for the duration of one second.

During Scanning Phase 2 (the last three scanning sessions), feedback was presented for trials in the delayed feedback condition that reflected whether participants selected the correct option as a match for the main word during the first scanning phase.

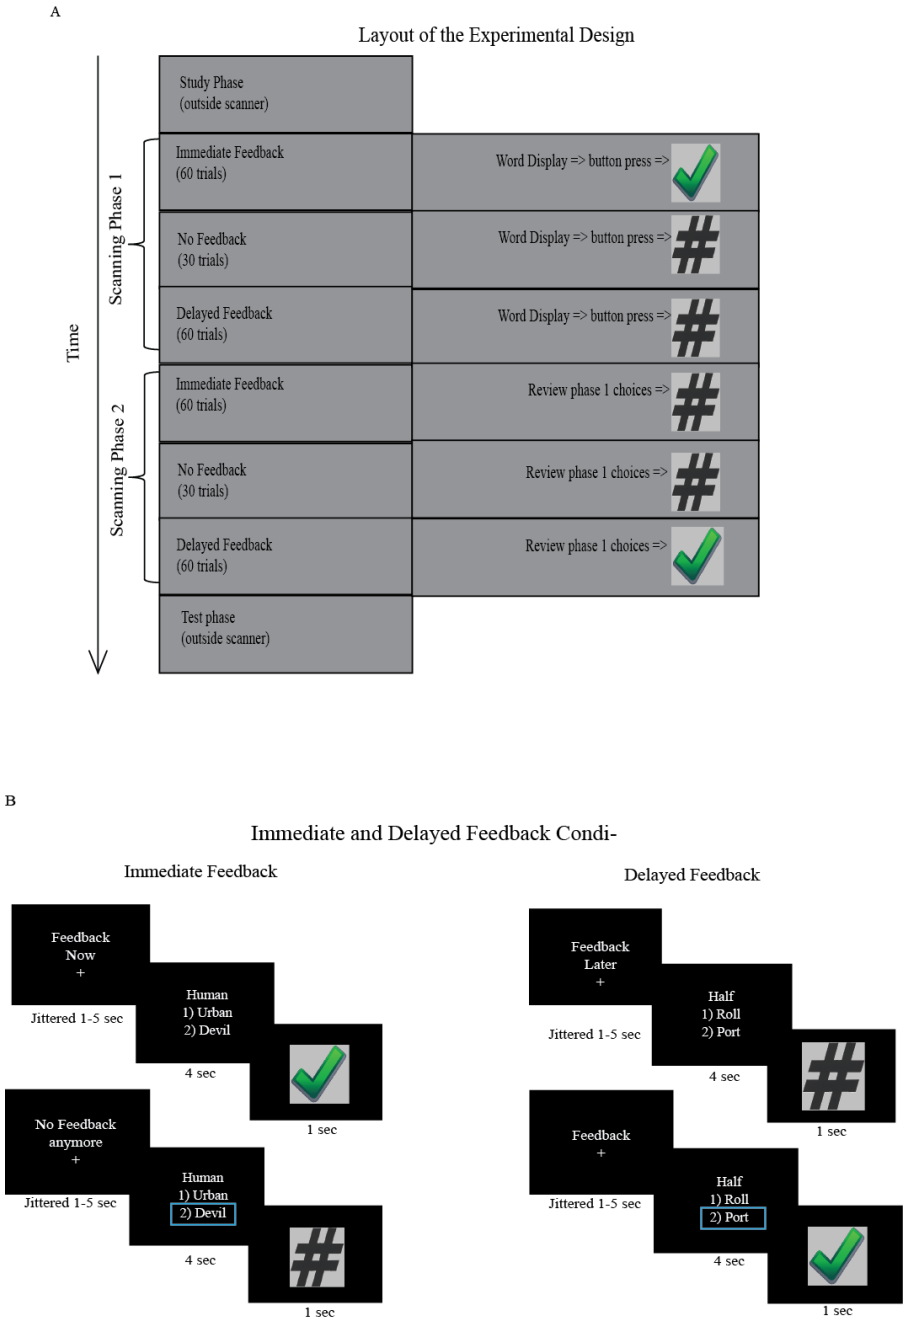

Figure 1. A. Chart of experimental events demonstrating their progression through time. *B.* Depiction of trials for immediate and delayed feedback during Scanning Phase 1 and

2. During Scanning Phase 2, the blue highlight appears for all conditions and indicates the participant's choice made at Scanning Phase 1.

The resulting delay between the action and the outcome was approximately 25 minutes. For trials in the other two conditions participants were presented with a control screen that showed a black pound sign (#) instead of the feedback (green √; red X). Both feedback and control screens were presented for the duration of one second. During Scanning Phase 2, participants were reminded of their choice made during the first three scanning sessions. This was done by presenting a blue highlight around the option that they previously selected, for all conditions (Figure 1b). In order to control for the motor response, when presented with the stimulus, participants were required to press a third button, unrelated to any word option.

During the test phase that occurred outside of the scanner at the end of the experiment, all of the words from all of the feedback conditions were presented in random order. Feedback was not presented at this stage of the experiment in order to test the effect feedback had on memory. Each trial lasted four seconds and was followed by a confidence rating question, where participants were given an unlimited time to indicate how certain they were about their response on the scale from one to seven (1 = complete guess; 7 = completely sure). A fixation point followed and lasted for three seconds. The complete experimental layout is presented in Figure 1a.

Data analysis

Behavioral data. Behavioral analysis was performed on the data from participants whose data were included in the fMRI data analyses. Accuracy data from the scanning phase 1 and accuracy data from the test phase were analyzed with an ANOVA and post-hoc two-sample t-tests.

fMRI data. Preprocessing of the functional data was performed using the Brain Voyager QX software (Version 2.1.2; Brain Innovation, Maastricht, the Netherlands). Preprocessing included three-dimensional correction for motion using six parameters. Images were spatially smoothed (8 mm, FWHM), voxel-wise linearly detrended, and passed through a high-pass temporal filter of frequencies (3 cycles per time course). The resulting data were normalized to the Talairach stereotaxic space (Talairach & Tournoux, 1998).

After image preprocessing, a whole brain analysis was performed on the data. A random-effects general linear model (GLM) analysis was performed on the one-second time period of feedback presentation, where the predictors of interest were immediate feedback (positive and negative), delayed feedback (positive and negative), and no feedback. The no response trials and the six motion parameters were included in the model as regressors of no interest. The GLM analysis resulted in identification of regions of interest (ROIs) thresholded at $p < 0.001$, along with a contiguity threshold of 3 (3x3x3 mm^3) contiguous voxels, determined using the cluster-level statistical threshold estimator in BrainVoyager (Version 2.1; Brain Innovation, Maastricht, The Netherlands). This method corrects for multiple comparisons and produces a cluster level false positive alpha rate of 0.05. Whole brain analyses were aimed at detecting differences associated with immediate and delayed feedback presentation of positive versus negative valence and at detecting differences associated with feedback presentation compared to no feedback. We also performed a whole-brain, voxel-wise ANOVA with delay (immediate versus delayed feedback) and valence (positive versus negative feedback) as within-

subjects factors. For this analysis, the no feedback trials were included in the model as predictors of no interest.

Additional whole brain contrasts were conducted to directly compare immediate and delayed feedback presentation and to compare valence conditions for immediate and delayed feedback, individually. For the delay contrast, the no feedback trials of Scanning Phase 1 versus Scanning Phase 2 were used to control for non-feedback related effects of time, since the immediate feedback trials necessarily occurred during Scanning Phase 1 while the delayed feedback trials occurred during Scanning Phase 2. The valence contrasts were aimed at detecting differences associated with immediate and delayed feedback presentation of positive versus negative valence and at detecting differences associated with feedback presentation compared to no feedback presentation of a corresponding delay. Thus, for these analyses, the predictors of interest were immediate feedback (positive and negative), delayed feedback (positive and negative), and no feedback (Phase 1 and 2). For all analyses, the no response trials and the six motion parameters were included in the model as regressors of no interest. We identified regions of interest (ROIs) thresholded at $p < 0.005$, along with a contiguity threshold of 6 (3x3x3 mm^3) contiguous voxels, determined using the cluster-level statistical threshold estimator in BrainVoyager (Version 2.1; Brain Innovation, Maastricht, The Netherlands). This method corrects for multiple comparisons and produces a cluster level false positive alpha rate of 0.05.

Roadmap of analyses. To prepare a reader, below I include an outline of specific fMRI analyses that were performed on data obtained from this experiment.

The design of the current experiment allowed for conducting an ANOVA analysis that allowed us to identify the regions sensitive to the delay (immediate or delayed feedback presentation) of feedback presentation and to the valence of presented feedback (positive vs. negative).

Based on the ANOVA results, further analyses were conducted where feedback valence was contrasted in order to see whether the striatum is differentially activated during feedback presentation of different valence.

Results

Behavioral results

Accuracy. Figure 2a displays accuracy results for the three feedback conditions of Scanning Phase 1 and of the test phase. We conducted an ANOVA with the factors of feedback type (immediate, delayed and no feedback type) and experimental phase (scanning phase 1 and test phase) in order to see the effect feedback had on learning. The ANOVA revealed a significant interaction of feedback type by experimental phase ($F(2,18) = 7.39$, $p < 0.005$) and main effect of experimental phase ($F(1,18) = 19.97$, $p < 0.0001$).

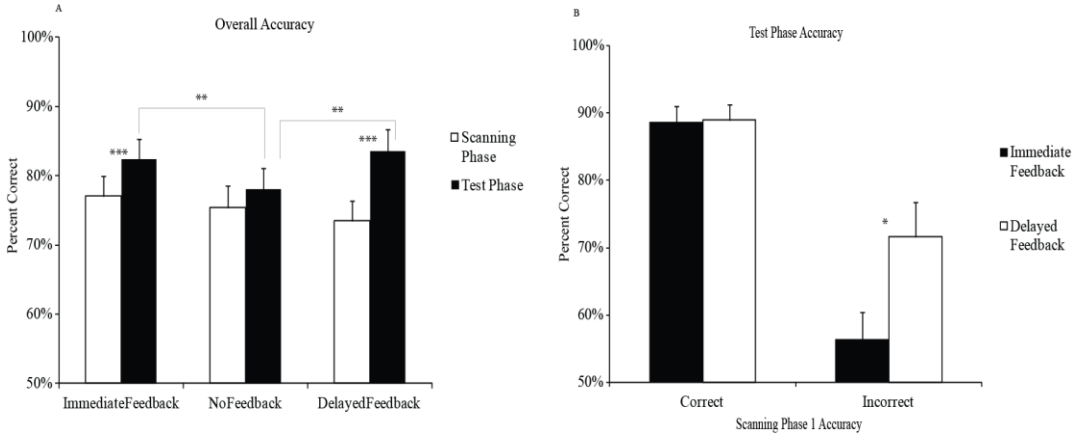

Figure 2 A. Accuracy for the three feedback conditions during the scanning phase and at the test phase. Significant differences between the scanning phase and the test phase

accuracy were detected for immediate and delayed feedback conditions ($p < 0.001$). At the test phase, the immediate feedback condition and delayed feedback conditions differ significantly from the no feedback condition ($p < 0.05$ for both comparisons). *B.* Accuracy at the test phase for positive and negative feedback trials from Scanning Phase 1. Significant difference was observed for negative feedback trials between immediate and delayed feedback, $t(18) = 2.90$, $p < 0.01$. *** indicates p-values less than 0.001; ** indicates p-values less than 0.05; * indicates p-values less than 0.01.

Post-hoc two-tailed paired t-tests showed that participants' accuracy improved significantly in the immediate and delayed feedback conditions. This was indicated by a significant difference between Scanning Phase 1 and the test phase, $t(18) = 4.04$, $p < 0.001$ (immediate feedback condition), $t(18) = 5.39$, $p < 0.001$ (delayed feedback condition). No significant accuracy increase was observed between Scanning Phase 1 and the test phase in the no feedback condition, $t(18) = 1.34$, $p = 0.19$.

In addition, two-tailed paired t-tests revealed significant differences at the test phase between immediate and no feedback conditions and between delayed and no feedback conditions, $t(18) = 2.53$, $p = 0.02$ (immediate vs. no feedback condition), $t(18) = 2.7$, $p = 0.01$ (delayed vs. no feedback condition). At the same time, participants' accuracy in the immediate feedback condition was not significantly different from the accuracy in the delayed feedback condition ($t(18) = 0.64$, $p = 0.53$).

Influence of feedback valence and delay on subsequent performance. Two-tailed paired t-tests were performed on the accuracy data from the positive and negative feedback trials of Scanning Phase 1 in order to see whether there were differences at test phase in learning from immediate and delayed feedback. No significant differences in test phase accuracy were found for positive feedback trials. However, for trials with negative feedback during Scanning Phase 1, at test phase, there was a significant difference between immediate and delayed feedback ($t(18) = 2.90$, $p = 0.009$). That is, significantly

more incorrect trials during the delayed feedback condition were correctly identified at the test phase than the incorrect trials of the immediate feedback condition (Figure 2b). *Questionnaire Results.* An additional set of 21 subjects took part in a behavioral version of the experiment without being scanned. This was done in order to gather more self-report data about subjective value of each type of feedback. Behavioral findings replicated the result of the group of subjects that participated in the fMRI version of the experiment. Two-tailed t-tests revealed significant differences at test phase between immediate and no feedback conditions and between delayed and no feedback conditions, $t(20) = 5.06$, $p < 0.001$ (immediate vs. no feedback condition), $t(20) = 3.38$, $p < 0.001$ (delayed vs. no feedback condition).

Out of the 21 participants, 15 indicated that learning from delayed feedback was more difficult than learning from immediate feedback. In addition, the majority of participants (20 out of 21) indicated that immediate positive feedback presentation felt more rewarding than delayed positive feedback presentation, while the immediate negative feedback presentation (14 out of 21) felt more punishing compared to delayed negative feedback presentation.

fMRI results

ANOVA results We performed a whole-brain, voxel-wise within-subjects ANOVA with delay (Phase 1 vs. Phase 2) and valence (with delay (Phase 1 vs. Phase 2) and valence (positive, negative, and no feedback) as within-subject factors.) as within-subject factors. The resulting clusters of activation with a threshold of $p < 0.005$ and a contiguity threshold of 6 voxels are listed in Table 1. A main effect of delay was found in the lentiform nucleus (the putamen and globus pallidus), bilaterally. A similar region showed

a main effect of valence. An overlap map of the main effect of delay and valence in the lentiform nucleus is presented in Figure 3a. Additionally, a main effect of valence was found more anteriorly, in the caudate nuclei (caudate head), bilaterally. Another more posterior part of the caudate nucleus also showed a main effect of valence (caudate body). Several cortical regions showed an interaction of delay and feedback valence, but no striatal areas were identified.

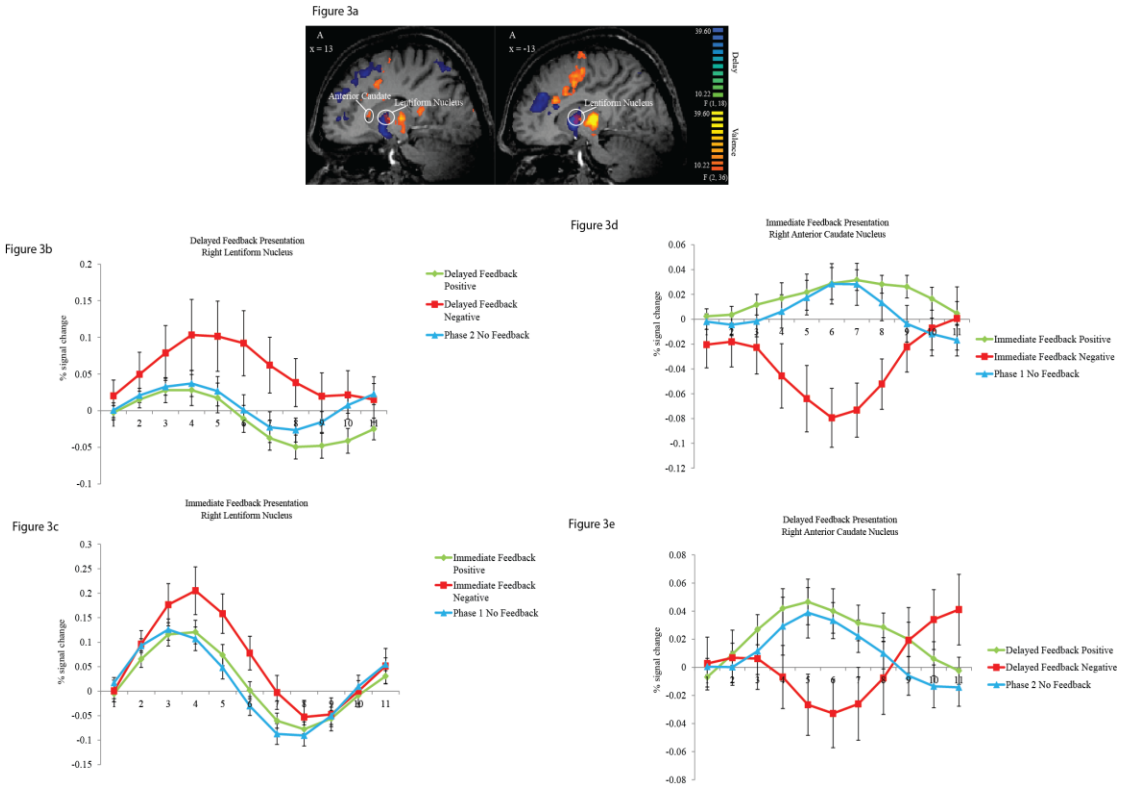

Figure 3. A. An overlap map of regions displaying the main effect of valence (in orange; positive vs. negative vs. no feedback) and the main effect of delay (in blue; Phase 1 vs. Phase 2). Anterior striatum seems to be sensitive to the valence of feedback (positive vs. negative feedback) and while the posterior parts of the basal ganglia seem to be sensitive both to delay (immediate vs. delayed feedback presentation) and valence. *B-C.* Time course of activation to delayed and immediate feedback in the right lentiform nucleus showing main effect of valence and delay. *D-E.* Time course of activation to immediate and delayed feedback in the right anterior caudate nucleus showing main effect of valence.

Effect of delay. To further investigate the brain responses to feedback presented immediately versus after a delay, a whole-brain GLM analysis was conducted to directly

compare immediate and delayed feedback presentation, while controlling for scanning phase. That is, collapsing across valence, immediate feedback presentation was compared to delayed feedback presentation, while including the no feedback presentation trials of the corresponding phase as a control (i.e., (Delayed feedback – Phase 2 no feedback) versus (Immediate feedback – Phase 1 no feedback). This contrast resulted in activity of the lentiform nucleus and the posterior caudate nucleus (Figure 4; Table 2).

Figure 4

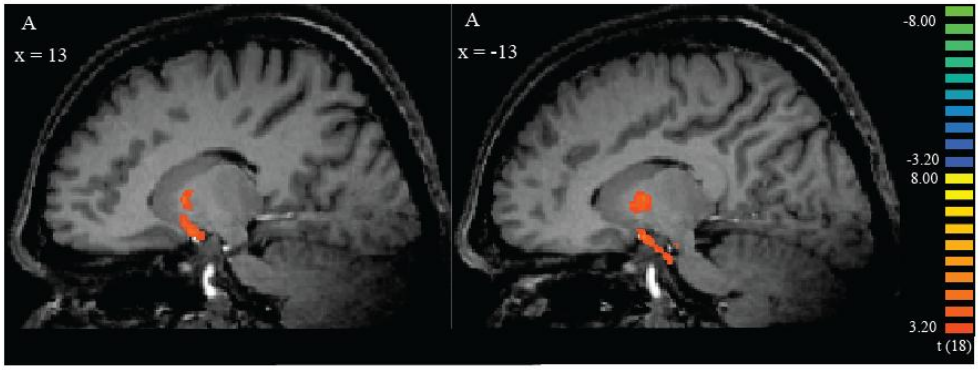

Figure 4. Brain activity associated with feedback presentation after a delay: (Delayed feedback – Phase 2 no feedback) versus (Immediate feedback – Phase 1 no feedback). No feedback presentation trials were included as a control.

Effects of valence for immediate feedback presentation. To gain a better understanding of the effects of valence in our dataset, we also performed contrasts between the different valence conditions for immediate and delayed feedback presentation, individually. During the immediate feedback condition, the contrast between positive feedback presentation versus negative feedback presentation revealed significant differences in activity in the right and left caudate nuclei (Figure 5; Table 3a). The cluster of activity in the right caudate overlapped with the cluster identified by our ANOVA as showing a main effect of valence. No striatal activity was detected for the contrast of immediate positive versus no feedback (Phase 1) (Table 3b). A cluster of activity in the caudate tail

was detected for the contrast of no feedback (Phase 1) versus immediate negative feedback presentation (Table 3c).

Figure 5

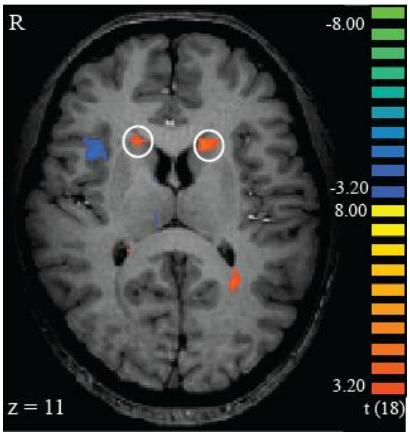

Figure 5. Brain activity associated with immediate positive feedback presentation versus immediate negative feedback presentation.

A cluster of activity was also detected in the right and left dorsal anterior cingulate cortex (dAcc), showing greater activation during the presentation of the negative feedback than positive feedback. These clusters were similar to areas from our ANOVA that showed a main effect of delay (right dACC) and a main effect of valence (bilaterally).

Effects of valence for delayed feedback presentation. Differential activity in the basal ganglia was observed during the contrast of delayed negative feedback versus delayed positive feedback and during the contrast of delayed negative versus no feedback (Phase 2). Specifically, the left lentiform nucleus cluster detected during the contrast of delayed negative versus delayed positive contrast (Figure 6) overlapped with the region detected during the contrast of delayed negative versus no feedback contrast (Table 4a, b, c). At

the same time, the left lentiform nucleus region overlapped with the ANOVA region showing the main effect of valence and delay.

Figure 6

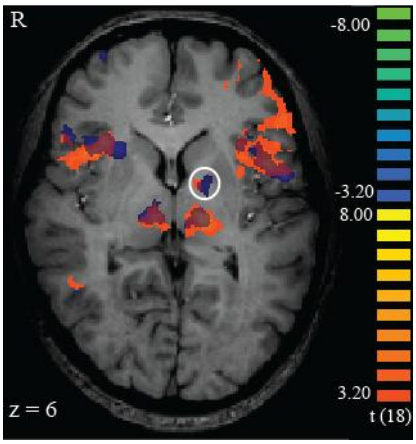

Figure 6. Brain activity associated with delayed negative feedback presentation versus delayed positive feedback presentation.

The anterior insula and ventral anterior cingulate were also activated bilaterally for the presentation of delayed negative feedback versus no feedback and for the presentation for delayed negative feedback versus delayed positive feedback. These regions, as well as the basal ganglia regions that showed sensitivity to delayed feedback versus no feedback, overlapped with the clusters of activity identified by our ANOVA as showing a main effect of valence.

Discussion

Feedback processing after a delay

In this experiment, performance-related feedback was presented either immediately or after a substantial delay of approximately 25 minutes. Similarly to studies with short delay (Foerde & Shohamy, 2011), our behavioral findings revealed equivalent accuracy at the test phase, indicating that subjects were able to learn from both immediate

and delayed feedback. To support learning in the delayed condition, either the striatum must be similarly recruited during both immediate and delayed feedback processing, or delayed feedback processing must be accomplished through a separate neural mechanism. Although basal ganglia activation was observed during the presentation of both types of feedback in our study, separate subregions were identified as playing a role in learning from immediate and delayed feedback.

During immediate feedback processing, we noted the typical increase in caudate activation following positive feedback and a decrease in activity during negative feedback presentation (Delgado, et al., 2000; Tricomi & Fiez, 2008). We did not, however, replicate this effect for the presentation of delayed feedback in the current study. Although this null result should be interpreted cautiously, since the caudate did not show a main effect of delay or interaction of valence and delay, it is in line with recent evidence indicating that the caudate is not critically involved in learning from delayed feedback in the way that it is for immediate feedback (Foerde & Shohamy, 2011). A common explanation for decreased caudate activation in response to immediate negative feedback presentation is that this type of feedback is interpreted as a punishment. It may be that after a substantial delay, negative feedback may no longer elicit as strong an affective response. That is, while the information carried by immediate feedback is processed in terms of its valence, the information carried by delayed feedback may be interpreted more cognitively, with less of the affective component of reward and punishment for one's action. Indeed, according to the questionnaire data, participants reported feeling immediate feedback to be more rewarding and punishing as compared to delayed feedback.

We propose that, after time elapses, negative feedback may be perceived as an opportunity to learn so that one's performance can be subsequently corrected. In line with this interpretation, we found that errors were more likely to be corrected during the test phase if the negative feedback was received after a delay, rather than immediately. Since the test phase occurred closer in time to the presentation of the delayed feedback than the presentation of immediate feedback, it is also possible that the increased improvement in accuracy on the test phase after receiving delayed negative could be due to recency effects. We would expect, however, for recency effects to apply to both positive and negative feedback trials, but this is not the case. Immediate positive feedback and delayed positive feedback did not result in differential performance during the test phase.

We identified a second, more posterior region in the basal ganglia as showing a main effect of delay and valence. A similar region was identified as showing greater activity in response to delayed feedback compared to no feedback and delayed negative versus positive feedback. This is the same region that was also activated during the second order contrast that allowed to compare immediate versus delayed feedback presentation, controlling for the confound of time. This region did not overlap with the more anterior region in the caudate identified in processing immediate feedback valence. This suggests that distinct neural mechanisms may support learning from immediate and delayed feedback. Hence, the anterior part of the basal ganglia might be more valence sensitive, while the posterior part of the basal ganglia might be involved in reinterpretation of negative information after a delay in a more cognitive fashion. Such functional differences among the basal ganglia regions may be explained by the diverse reciprocal projections from the basal ganglia to the prefrontal cortex. Anatomical tracing

studies show major projections from the anterior-ventral striatum to such limbic structures as the amygdala and the orbitofrontal cortex, regions that are involved in processing of affective information (Haber, Fudge, & McFarland, 2000; Haber, Kim, Mailly, & Calzavara, 2006; Haber & Knutson, 2010). Although the posterior-dorsal basal ganglia (caudate body and tail and the putamen and globus pallidus (GP)) is typically thought to be part of the "motor loop" (Middleton & Strick, 2000; Seger, 2008), it also projects to prefrontal cortex structures involved in processing of cognitive information, such as the dorsolateral prefrontal cortex and the dorsal anterior cingulate cortex (dAcc) (Boettiger & D'Esposito, 2005; Haber & Knutson, 2010; Han, Huettel, Raposo, Adcock, & Dobbins, 2010; Longe, Senior, & Rippon, 2009; Mohanty et al., 2007). Indeed, recent studies suggest that in addition to its motor functions, the GP plays an important role in memory processing and learning (Baier, Karnath, & Dieterich, 2010; McNab & Klingberg, 2008).

Although our task involved declarative memory acquisition, we did not find significant feedback-related effects in the MTL. It may be that the MTL was similarly recruited in all conditions of our study because they all relied on declarative memory to a similar degree. Other work, however, has found that the hippocampus is engaged in outcome processing when it is temporally separated from the cue (Foerde & Shohamy, 2011). Even though feedback was presented after a delay in the current study, participants were reminded of the cue before delayed feedback was presented. This feature of the experimental design might be another reason why no hippocampal activity was detected. Tasks that are more specifically aimed at modulating MTL activity may be necessary to identify dopaminergic or reward-related influences on activity in this region

(Foerde & Shohamy, 2011; Sadeh, Shohamy, Levy, Reggev, & Maril, 2011; Shohamy, Myers, Kalanithi, & Gluck, 2008; Wittmann et al., 2005)

Anterior cingulate cortex-striatum network

Anterior cingulate activity was observed during both types of feedback presentation in the current study. Specifically, an overlapping cluster of Acc activity was detected during the presentation of delayed negative feedback when contrasting it to no feedback presentation, and during the presentation of immediate negative feedback when contrasting it to positive feedback and no feedback presentation .

Previous studies also report increased Acc activation in response to feedback, suggesting that the Acc plays an important role in decision making (Rushworth, Behrens, Rudebeck, & Walton, 2007). The Acc is proposed to detect conflict related to improper responding during the task (Daniel & Pollmann, 2010; Holroyd & Coles, 2002; Holroyd et al., 2004; Yeung, Botvinick, & Cohen, 2004). Specifically, the prediction error signal from the midbrain, which reflects how well the expectation of the outcome matches the actual outcome, is used by the Acc (Kennerley, Walton, Behrens, Buckley, & Rushworth, 2006; Milham & Banich, 2005) to signal to other cognitive areas to increase cognitive control and correct future performance (Hong & Hikosaka, 2008).

Activation of the lentiform nucleus in our study was especially pronounced in response to delayed negative feedback. This activity, in conjunction with the observed activation for negative feedback trials in the dAcc, may reflect a role of this network in error correction. Indeed, neurons in the globus pallidus internal capsule, which influence dopamine neurons via their projections to the lateral habenula, are sensitive to prediction errors, and increase their firing rate when a target signals the absence of a reward

(McNab & Klingberg, 2008; Walsh & Phillips, 2010). Even though it is not possible to distinguish the specific part of the globus pallidus activated during our task using fMRI tools, this is one potential mechanism by which the basal ganglia region activated in our study may contribute to learning from feedback after a delay. Further research will be needed, however, to determine why this region might be especially sensitive to negative feedback when it is presented after a substantial delay.

Limitations

Our finding of a greater percentage of errors from the delayed feedback condition that were correctly identified at test phase can be explained by the fact that the test phase followed immediately after the scanning phase 2, during which the feedback for the delayed feedback condition was presented. However, we think that this explanation would be appropriate if there would also have been a difference between the positively identified trials of delayed and immediate feedback condition. This was not observed.

Conclusion

This study suggests that the neural mechanisms involved in feedback processing are affected by the temporal proximity of the feedback to an action. That distinct basal ganglia subregions are involved in immediate and delayed feedback processing suggests that the same type of feedback might be interpreted differently if it is presented after a delay rather than immediately.

Here we replicate previous findings related to the caudate's role in processing immediate feedback and, in addition, our results shed light on a potential role of other basal ganglia nuclei such as the lentiform nucleus in delayed feedback processing. Taken together, our results underscore the importance of the basal ganglia in the performance of

cognitive tasks, and point to a functional heterogeneity within the basal ganglia in supporting learning under different time frames.

Chapter Three: Experiment 2: Effort-based Learning.

Delay between the action and the outcome is just one of the factors that can influence outcome valuation. Effort expenditure is another cost that one has to account for when deciding to perform a goal-directed action.

There are two types of effort that can be distinguished and that can influence outcome valuation: physical effort and cognitive effort. For example, an athlete might expend more effort during training when he wants to win a specific competition. Similarly, a student might study harder if getting a good grade for a class is important. In experimental settings, physical effort tasks require subjects to perform a particular action continuously in order to obtain the desired outcome, such as pressing a lever, climbing a barrier or applying extra force to obtain food or money (Croxson, Walton, O'Reilly, Behrens, & Rushworth, 2009; Floresco, Tse, & Ghods-Sharifi, 2008; Kurniawan et al., 2010; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). Cognitive effort and consequent information processing that it requires can be defined as allocation of increased mental resources for a particular task (Jansma, Ramsey, de Zwart, van Gelderen, & Duyn, 2007). Cognitive effort tasks manipulate attentional demands, working memory demands and require more information integration than less complex tasks (Botvinick & Rosen, 2009; Cocker, Hosking, Benoit, & Winstanley, 2012; Kool, McGuire, Rosen, & Botvinick, 2010; McGuire & Botvinick, 2010; Satterthwaite et al., 2012; Schmidt, Lebreton, Clery-Melin, Daunizeau, & Pessiglione, 2012).

Neural mechanisms postulated to be associated with processing effortful goal-directed behavior and its outcomes seem to be dopamine dependent (Salamone, Correa, Mingote, & Weber, 2003; Treadway et al., 2012). Animal research showed that NAcc dopamine depletion interferes with animals' behavior in such a way that animals lose the motivation to act in an effortful manner. For example, a rat might be presented with two outcome options: 1) a small food reward that is easily accessible (i.e., requires no effort) and 2) a larger food reward that can be obtained if the rat chooses to climb a barrier (i.e., decides to exert extra effort). Intact rats prefer the larger food reward that requires extra effort. However, after dopamine depletion or after dopamine antagonist administration, these rats are not able to overcome behavioral costs such as effort requirements and prefer the option with small amount of food that requires less effort. This behavioral change occurs only during the dopamine depletion from NAcc and thought to be unrelated to motor impairments also associated with dopaminergic projections to the striatum (Assadi, Yucel, & Pantelis, 2009; Ishiwari, Weber, Mingote, Correa, & Salamone, 2004; Phillips, Walton, & Jhou, 2007; Salamone, Correa, Mingote, & Weber, 2003; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006).

Recently, a cognitive effort task for rats was developed. Cocker and others trained rats on a visuospatial discrimination task. The results showed that all subjects preferred the high reward option that was associated with high cognitive effort. Administration of dopamine agonist increased the preference for the high effort and high reward option (Cocker, Hosking, Benoit, & Winstanley, 2012). Therefore, goal-directed behavior that has physical and cognitive effort requirements seem to depend on similar neural mechanisms.

Based on the findings from the animal literature of dopamine involvement in effort-based decision-making, several neuroimaging studies attempted to look at human brain mechanisms involved in processing effort-related information. Neuroimaging findings, paralleling the findings of the animal literature, implicated the ventral striatum as one of the primary regions involved in processing the cost of actions and associated outcomes. However, in contrast to animal studies, increased NAcc activation was observed in association with outcomes that are easily obtained (Botvinick, Huffstetler, & McGuire, 2009; Croxson, Walton, O'Reilly, Behrens, & Rushworth, 2009; Kurniawan et al., 2010). Most recently, only Treadway et al. (2012) presented evidence of ventral striatum activation to low probability rewards delivered after effortful behavior after dopamine agonist administration, utilizing a task that paralleled animal paradigms (Salamone, Correa, Mingote, & Weber, 2003; Treadway et al., 2012). This discrepancy might be explained by paradigm differences between studies. In some studies the reward magnitude was not varied when the effort level was (Botvinick, Huffstetler, & McGuire, 2009), while other experiments presented participants with insignificant reward magnitude difference in conjunction with variable effort demands (Croxson, Walton, O'Reilly, Behrens, & Rushworth, 2009; Kurniawan et al., 2010; Treadway et al., 2012).

One theory is that decreased activation of the human ventral striatum is thought to represent effort-discounting in the human brain. According to this theory, termed effort-discounting, effort carries a negative weight. Effortful actions and associated outcomes will be devalued due to a greater amount of effort required to perform the action and obtain the outcome compared to the effortless action and outcomes associated with it (Botvinick & Rosen, 2009; Kool, McGuire, Rosen, & Botvinick, 2010; McGuire &

Botvinick, 2010). This principle holds for both physical and cognitive effort. That is, one would always perform an action that requires minimal amount of effort exertion and when given two rewards of equal magnitude, the person will prefer the one that requires least effort expenditure (Botvinick & Rosen, 2009; Kool, McGuire, Rosen, & Botvinick, 2010; McGuire & Botvinick).

However, contrast theory makes opposite predictions and postulates that an outcome of a more effortful action would be preferred due to a greater contrast between the aversive action and the rewarding nature of the outcome (Singer, Berry, & Zentall, 2007; Zentall & Singer, 2007). The original study that lent support to the contrast theory account involved pigeons that were trained to associate positive and negative outcomes (such as the presence and absence of food) with either effortful or effortless actions (one peck versus 20 pecks). After training, when pigeons were given a choice between the two cues, the cue associated with effortful action (20 pecks) was chosen more often than the cue associated with the effortless action (one peck). Similar results were reported with humans (Alessandri, Darcheville, Delevoye-Turrell, & Zentall, 2008). Based on contrast theory, one can predict increased activity of the ventral striatum in association with outcomes that follow effortful actions. Therefore, contrast theory suggests that effort associated with a goal-directed action can alter the value of its outcome, but in the opposite direction than the effort-discounting theory predicts.

The studies cited above and the evidence of effort-discounting in the human brain can be grouped into categories of studies that offer *extrinsic outcome* delivered after either low or high effort demands, and tasks that present participants with *a choice* between either low or high effort that leads to a extrinsic outcome. The extrinsic

outcomes presented during the course of some studies differ either insignificantly or do not differ at all, while the level of effort expenditure is varied. It has not been investigated how low and high cognitive effort demands would influence outcome valuation and associated brain activity during learning through performance-related feedback, i.e. through outcomes when one has no choice but to perform an action. That is, would intrinsic outcomes that follow high cognitive effort be subjectively rated as more preferable and valuable and would these subjective outcome preferences result in increased activation in the VS?

In my second study, I looked at the effect of a cognitive effort manipulation on outcome valuation and associated striatal activity during a trial-and-error learning with intrinsic rewards (positive feedback). Participants were presented with trial-and-error learning task, in which they had to learn to associate abstract images with specific responses based on the feedback presented after each trial. Cognitive effort was manipulated by means of presenting either feedback that reflected accuracy for one stimulus or feedback that reflected cumulative accuracy to two stimuli presented at the same time. I expected participants' performance on the effortful condition to be significantly worse than their performance on the low effort condition. In addition, I expected to see an increase in VS activation in association with positive feedback following the high effort condition in conjunction with participants' rating of feedback after the more difficult condition as more preferable.

Methods.

Participants.

Twenty-four individuals participated in the experiment for payment of $50. All participants provided written consent to participate. Data from one of the participants were not included into the main analysis due to diagnosed brain abnormality. Data from one other participant were not collected completely due to the onset of a panic attack. Therefore, data from 22 participants were analyzed (9 females; mean age 23.3 years, SD 5.4). The research was approved by the Institutional Review Boards of Rutgers University.

Materials.

A 3-Tesla Siemens Trio scanner was used to acquire all fMRI data. Behavioral data acquisition and stimulus presentation was administered using the "E-Prime" software (Schneider, Eschman, & Zuccolotto, 2002).

Procedure.

Scan session. A T1-weighted pulse sequence was used to collect structural images in 41 contiguous slices (3x3x3 mm voxels) tilted 30° from the AC-PC line (Deichmann, Gottfried, Hutton, & Turner, 2003). Similarly, 41 functional images were collected using a single-shot echo EPI sequence amounting to 142 acquisitions (TR= 2500 ms, TE = 25 ms, FOV = 192 mm, flip angle = 80°).

Behavioral paradigm. In this experiment, participants had to learn to associate abstract images with one of the four specific buttons on the computer keyboard (1, 2, 3 and 4). Specifically, participants were presented with two learning conditions that represented high and low cognitive effort conditions, and two random feedback conditions that required no cognitive effort but only a motor response. In the 1-step learning condition (low cognitive effort; LE), participants were presented with one

abstract image and had to respond with one of the four buttons, only one of which led to the presentation of the correct feedback (green √). The other three buttons led to the presentation of the incorrect feedback (red X).

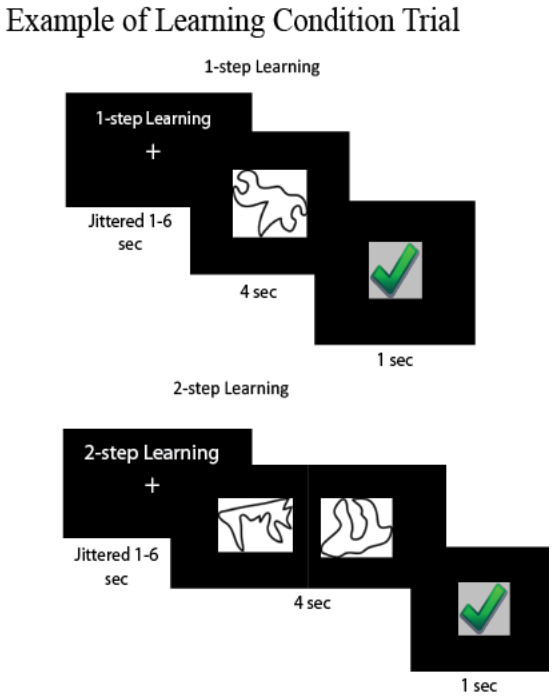

Figure 7. Depiction of trials for 1-step (low effort) and 2-step (high effort) conditions. 1-step random and 2-step random conditions resembled the above set-up; however, random feedback did not reflect performance accuracy.

During the 2-step learning condition (high cognitive effort; HE), participants were presented with two abstract images, side by side (Figure 6). Participants had to respond to both images: buttons 1 or 2 were response options for the first image, while buttons 3 or 4 were response options for the second image. Cumulative feedback was presented after the response to both images was made. That is, correct feedback was presented only when participants responded correctly to both images. At all other times, incorrect feedback was presented. A fixation point of one second separated the two abstract images.

It is important to note here that feedback provided the same amount of information in both conditions. That is, at the beginning of the learning process, when

participants are just guessing, there is a 25% chance of being correct in each trial, since there are four response options per trial in both learning conditions.

The 1-step and 2-step random conditions resembled the conditions described above in all respects. The only difference was that feedback did not reflect participants' accuracy and correct and incorrect feedback presentation was random. Participants were informed of the fact that feedback during random conditions would not reflect their performance and were told that they just had to press the button during these conditions.

Each trial started with a fixation point that contained the label informing participants of the condition they were in. The label and the fixation point lasted 4.5 seconds on average. A fixation point that was presented in between the different conditions also lasted 4.5 seconds (1-8 seconds). All stimuli were presented for four seconds while the feedback screen was presented for one second. There were four trials in each condition presented twice during one run. There were a total of six runs in the experiment.

The experiment ended outside of the scanner and participants were given the questionnaire that inquired: 1) which type of feedback presentation felt more rewarding to participants, 2) during which condition participants felt more engaged in the task, 3) when learning was harder and 4) whether random feedback presentation was rewarding.

Data analysis.

Behavioral data. Behavioral analysis was performed on the data from participants whose data were included in the fMRI data analyses. Accuracy data and response time data from the two learning conditions was analyzed by means of paired t-tests.

fMRI data. Preprocessing of the functional data was performed using the Brain Voyager QX software (Version 2.3.1; Brain Innovation, Maastricht, the Netherlands). Preprocessing included three-dimensional correction for motion using six parameters. Images were spatially smoothed (8 mm, FWHM), voxel-wise linearly detrended, and passed through a high-pass temporal filter of frequencies (3 cycles per time course). The resulting data were normalized to the Talairach stereotaxic space (Talairach & Tournoux, 1998).

After image preprocessing, a whole brain analysis was performed on the data. A random-effects general linear model (GLM) analysis was performed on the one-second time period of feedback presentation. The predictors of interest were: positive 1-step feedback, negative 1-step feedback, positive 2-step feedback, and negative 2-step feedback for the learning conditions; and positive 1-step feedback, negative 1-step feedback, positive 2-step feedback, and negative 2-step feedback for the random conditions. The missed trials and six motion parameters were included in the model as regressors of no interest. The GLM analysis resulted in identification of regions of interest (ROIs) thresholded at $p < 0.001$, along with a contiguity threshold of 3 (3x3x3 mm^3) contiguous voxels, determined using the cluster-level statistical threshold estimator in BrainVoyager (Version 2.3; Brain Innovation, Maastricht, The Netherlands). This method corrects for multiple comparisons and produces a cluster level false positive alpha rate of 0.05. Whole brain contrasts were aimed at detecting differences associated with feedback presentation of positive and negative valence in the 1-step and 2-step learning and random conditions, and at detecting differences associated with feedback

presentation between the 1-step and 2-step learning conditions (i.e. high cognitive effort vs. low cognitive effort condition).

Roadmap of analyses. To prepare the reader, below I include an outline of specific contrasts that are performed for the analysis of data obtained from this experiment.

To obtain an unbiased representation of brain activity that represents all of the conditions independent of a particular contrast, ROI analysis was conducted on the a priori coordinates taken from Tricomi at al. (2010) (-10, 9, 0). An ANOVA was conducted on the beta weights obtained from this ROI.

A whole-brain voxel-wise within-subjects ANOVA was also conducted in order to see whether there are differences between conditions that result in differential brain activity.

To clarify the results from the above analysis, direct contrasts between conditions were performed. Positive and negative feedback presentation was contrasted for each condition in order to see whether the valence of feedback presentation influences the brain activity.

In addition, positive feedback and negative feedback were compared between conditions in order to analyze whether feedback presentation from a specific condition is more potent than feedback presentation of the other condition.

Two second-order contrasts were performed. One of them was looking at positive feedback presentation and was performed in order to compare positive feedback from conditions of different effort levels. The other second-order contrast was aimed at

detecting differences between the conditions associated with overall difficulty. To perform this contrast we collapsed across valence in all conditions.

Results.

Behavioral results.

Accuracy. Figure 8a displays accuracy results for the two learning conditions. A two-tailed t-test was performed to compare the accuracy between the conditions in order to see in which condition feedback resulted in better learning. The t-test revealed a significant difference in accuracy between the two conditions, showing that participants learned significantly better from the 1-step learning condition: $t(21) = 5.37$, $p < 0.001$.

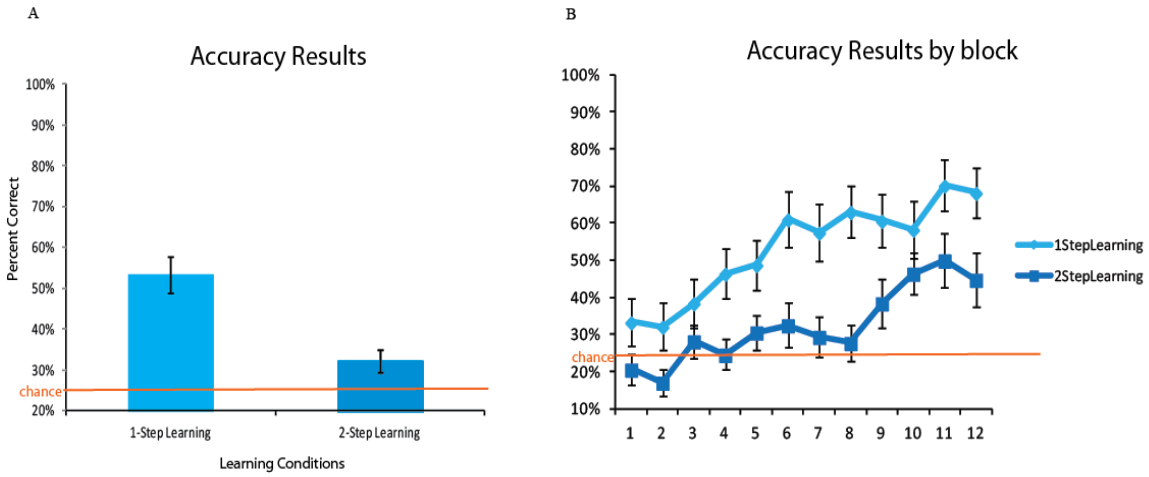

Figure 8. A. Accuracy for the two learning conditions. There is a significant difference in performance between the 1-step (low effort) and 2-step (high effort) conditions. *B.* Learning accuracy by block for the two learning conditions. Chance performance is at 25%.

To see the progress of learning, accuracy was analyzed on a block by block basis. For illustration purposes, accuracy results for the 1-step and the 2-step conditions are plotted by block in Figure 7b. The accuracy for the 1-step learning condition increases faster than accuracy in the 2-step learning condition.

Response time data. It was not possible to perform an ANOVA on the response time (RT) data in this experiment due to a methodological difference between the 2-step learning and the 1-step learning conditions. Specifically, the two conditions were programmed so that the known RT for the 2-step learning condition represented an average of the RTs for the two stimuli presented during this condition. The RT for the 1-step learning condition represented the actual RT to a single stimulus presented during the 1-step condition. Given such set up, the RT of the 2-step learning condition was always greater than the RT of the 1-step learning condition. For that reason, the RT of the 1-step learning condition was compared to the RT of the 1-step random condition and the RT of the 2-step learning condition was compared to the RT of the 2-step random condition. A two-tailed within-subjects t-tests revealed significant differences in RTs between the 1-step learning and random conditions ($t(21) = 4.43$, $p < 0.001$(incorrect 1-step learning vs. random; $t(21) = 2.17$, $p < 0.04$ (correct 1-step learning vs. random)) where incorrect and correct responses for the 1-step learning condition were significantly slower. In addition, RT for correct responses in the learning condition were significantly faster than the incorrect responses, ($t(21) = 2.17$, $p < 0.05$). Similarly, two-tailed within-subjects t-tests revealed significant differences in RTs between the 2-step learning and random conditions ($t(21) = 3.86$, $p < 0.001$(correct 2-step learning vs. random; $t(21) = 4.19$, $p < 0.001$ (incorrect 2-step learning vs. random)) with incorrect and correct responses for the 2-step learning condition being significantly slower. In addition, RT for correct responses in the learning condition were significantly faster than the incorrect responses, ($t(21) = 2.53$, $p < 0.05$).

Questionnaire data. Nine out of 22 participants indicated that they were most engaged in the task during the 1-step learning condition while 8 of the participants were most engaged in the task during the 2-step learning condition. The rest of the participants did not provide a clear response for this question. Similarly, 10 out of 22 participants indicated that feedback in the 2-step learning condition was more rewarding while the 8 indicated that feedback in the 1-step learning condition was more rewarding. Four other participants did not provide a clear response for this question. At the same time, all participants responded that learning during the 2-step learning condition was more difficult. Out of these participants, 5 of the participants who reported preferring feedback after the 2-step learning condition reported to also be more engaged while performing this condition. Similarly, 5 of the participants who reported preferring feedback after the 1-step learning condition reported to also be more engaged while performing this condition. In addition, fourteen participants answered no to the question about whether random feedback was rewarding to them.

fMRI results.

Region of Interest (ROI) analysis. In this study I was specifically interested in the pattern of activity of the VS due to numerous animal and human studies implicating this region in participation of cost-benefit valuation during the decision-making process. Therefore, I conducted an ROI analysis on the VS with data extracted from an 8mm sphere centered around (-10, 9, 0) coordinates (Tricomi, Rangel, Camerer, & O'Doherty, 2010). A within-subject ANOVA with difficulty (2-step vs. 1-step), contingency (learning vs. random) and valence (positive vs. negative) as within-subject factors revealed a significant main effect of valence ($F(1) = 22.39$, $p < 0.0001$). In addition, there was a trend towards

significance for the three-way interaction of difficulty, contingency and valence (F(1) = 3.76, p = 0.07) and the main effect of difficulty (F (1) = 3.46, p = 0.08). The parameter estimates from the VS coordinates are presented in Figure 8.

ANOVA results. A whole-brain, voxel-wise within-subjects ANOVA was performed with difficulty (2-step vs. 1-step), contingency (learning vs. random) and valence (positive vs. negative) as within-subject factors. There were only two regions that showed the three-way interaction were the medial prefrontal cortex (BA 10) and the fusiform gyrus; they are listed in Table 4a ($p < 0.05$ corrected). In addition, an interaction of difficulty by contingency resulted in activation of several prefrontal cortex areas, such as dorsolateral prefrontal cortex (DLPFC; BA 9) and orbitofrontal cortex (BA 11) (Table, 4b; $p < 0.05$ corrected). Interaction results of difficulty by valence and contingency by valence are presented in Tables 4c and 4d ($p < 0.05$ corrected).

Tables 5a through 5c show lists of regions with main effects of difficulty, contingency and valence ($p < 0.05$ corrected). The main effect of contingency was observed in DLPFC (BA 9) and in bilateral dorsal striatum, while the main effect of valence revealed large bilateral clusters of activity in the dorsal and ventral striatum.

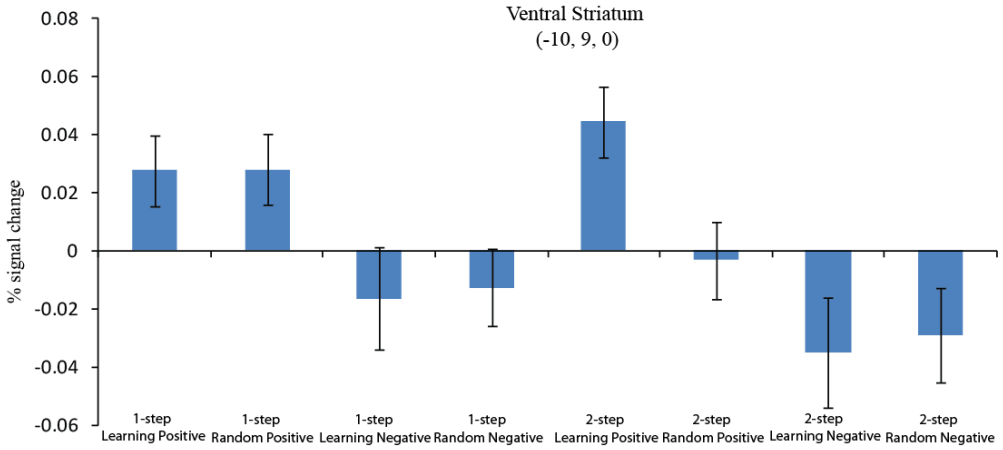

*Figure 9.*Parameter estimates from the left region of interest in the ventral striatum. An ANOVA (difficulty x contingency x valence) revealed a significant main effect of

valence and a trend towards significance for the main effect of difficulty and for the 3-way interaction.

1-step feedback presentation. A whole-brain GLM analysis was conducted to see activations associated with the period of feedback presentation in the 1-step learning and random conditions. Contrasting positive feedback presentation versus negative feedback presentation of the 1-step learning condition revealed stronger activation for positive feedback in the ventral striatum bilaterally, along with activity in the dorsal putamen and other cortical areas (Table 6a, Figure 9, $p < 0.05$ corrected).

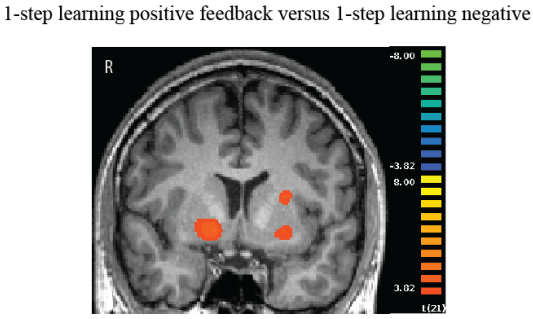

Figure 10. Brain activity associated with the presentation of the 1-step learning positive feedback versus 1-step learning negative feedback presentation.

At the same time, contrasting positive feedback presentation versus negative feedback presentation of the 1-step random condition, revealed stronger activation for positive feedback in the left ventral striatum. Contrasting 1-step learning negative versus 1-step random negative feedback revealed a single cluster of activity in the cerebellum (Table 6b; Figure 10, $p < 0.05$ corrected).

1-step random positive feedback versus 1-step random negative

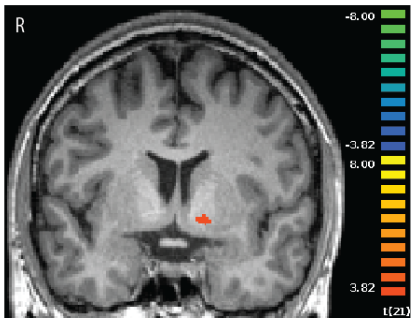

Figure 11. Brain activity associated with the presentation of the 1-step random positive feedback versus 1-step random negative feedback presentation.

Contrasting positive feedback presentation of the 1-step learning condition versus positive feedback presentation of the 1-step random condition only revealed activation in thalamus that was associated with positive feedback of the learning condition (Table 6c; $p < 0.05$ corrected).

2-step feedback presentation. Contrasting positive feedback presentation versus negative feedback presentation of the 2-step learning condition revealed large clusters of activation in the bilateral basal ganglia and frontal areas (Table 7a; $p < 0.05$ corrected). Stronger activation to positive feedback was detected in the bilateral dorsal and ventral striatum in addition to the putamen and globus pallidus (Figure 11).

2-step learning positive feedback versus 2-step learning negative

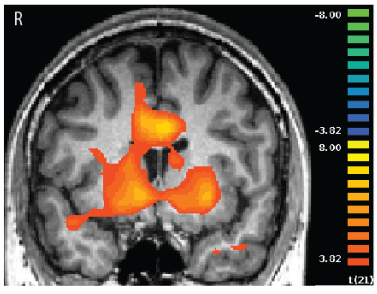

Figure 12. Brain activity associated with the presentation of the 2-step learning positive feedback versus 2-step learning negative feedback presentation.

Additionally, DLPFC (BA 9), MPFC (BA 10) and ACC (BA 32 and 24) showed increased activation to positive feedback presentation than negative feedback presentation (Figure 12a).

The contrast of 2-step random positive feedback presentation versus negative feedback presentation resulted in a small cluster of activity in the right ventral striatum. Several cortical areas were also activated during this contrast (Table 7b; Figure 13; $p < 0.05$ corrected).

Contrasting positive feedback during the learning condition versus positive feedback during the random condition results in bilateral dorsal striatal activation and activation of the right ventral striatum (Figure 14).

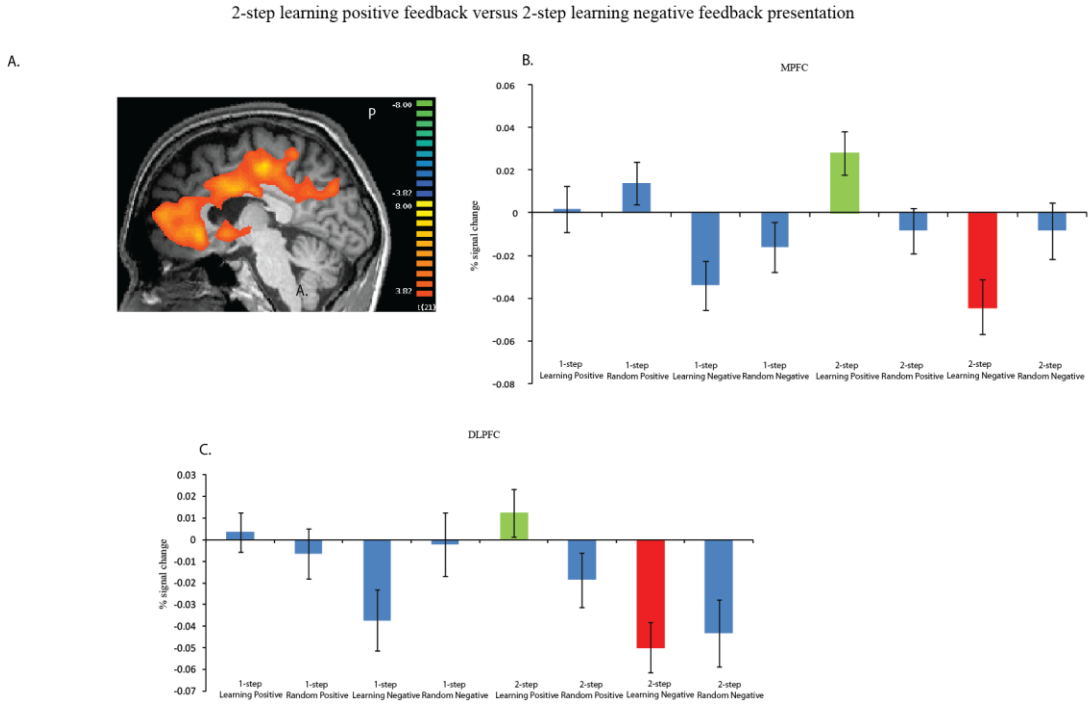

Figure 13. A. Brain activity associated with the presentation of the 2-step learning positive feedback versus 2-step learning negative feedback presentation. *B.* For illustrative purposes, percent signal change in the MPFC associated with differential activation for positive vs. negative feedback presentation during the 2-step learning condition is plotted for all conditions. *C.* Percent signal change in the DLPFC associated

with differential activation for positive vs. negative feedback presentation during the 2-step learning condition is plotted for all conditions.

2-step random positive feedback versus 2-step random negative

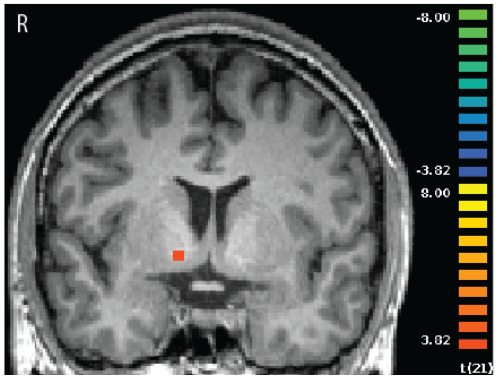

Figure 14. Brain activity associated with the presentation of the 2-step random positive feedback versus 2-step random negative feedback presentation.

2-step learning positive feedback versus 2-step random positive feedback

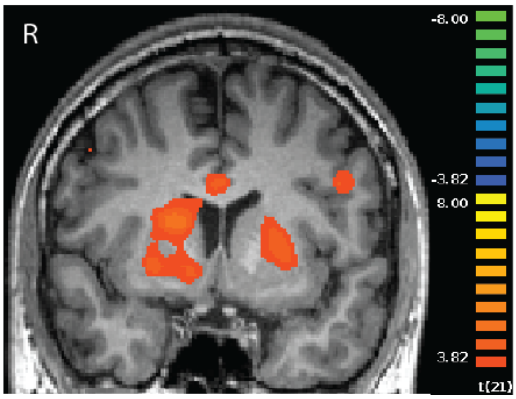

Figure 15. Brain activity associated with the presentation of the 2-step learning positive feedback versus 2-step random positive feedback presentation.

In addition, cortical areas were also activated during this contrast. Specifically, MPFC was activated more for positive feedback presentation during the learning condition. A similar contrast of 2-step learning negative feedback presentation versus 2-step random negative feedback presentation revealed several cortical areas and activations in the cerebellum (Table 7c; $p < 0.05$ corrected).

Second order contrast: HE versus LE. I conducted a contrast of HE versus LE feedback presentation while controlling for the random conditions. Thus this contrast resulted from

contrasting positive feedback presentation of the 2-step learning condition versus positive feedback presentation of the 1-step learning condition (i.e., (2-step learning positive – 2-step random positive) – (1-step learning positive – 1-step random positive)). This contrast revealed activations in the right ventral striatum, ventral ACC (BA 24), DLPFC (BA 9) and bilateral MPFC (BA 10) (Table 8; Figure 15; $p < 0.05$ corrected).

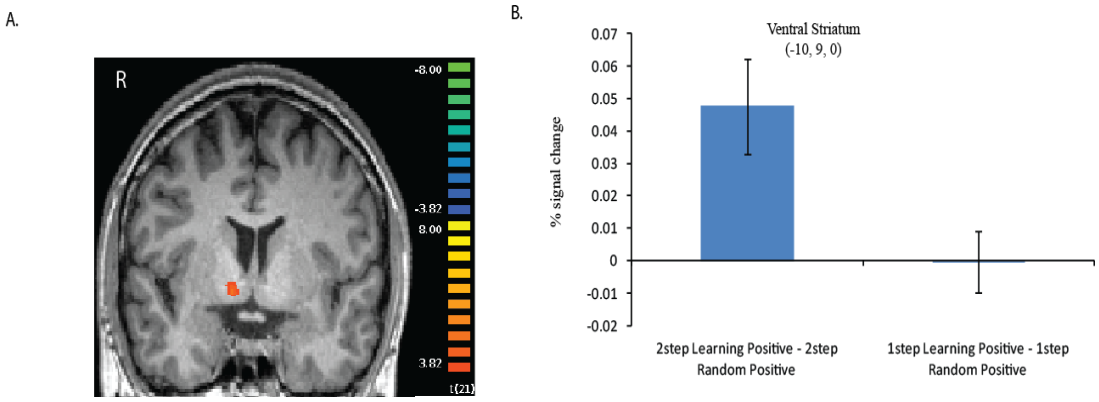

Figure 16. A. Brain activity associated with the presentation of the 2-step learning positive feedback versus 1-step learning positive feedback. *B.* Percent signal change in the ROI (-10, 9, 0) of the VS. Random conditions are included in the contrast as a control.

A similar contrast, between HE and LE conditions, collapsed across valence, was performed in order to see whether DLPFC would be activated in association with the HE condition (i.e, (2-step learning– 2-step random) – (1-step learning– 1-step random)). As expected, DLPFC was more activated for the HE (2-step learning) condition (Table 9; $p < 0.05$ corrected). No striatal activation was detected during this contrast, most likely because this analysis was not intended for looking at positive versus negative feedback presentation. Instead, the DLPFC activity may be due to activation during the cue period. Because there is no jitter between the cue and feedback presentation screens, I cannot objectively segregate BOLD activity related to the cue period only.

Exploratory analysis: 1-step vs. 2-step random positive feedback presentation. An additional contrast was performed in order to see whether the VS is activated for the

comparison of random feedback conditions. A two-tailed within-subjects t-test performed on parameter estimates from the a priori ROI showed a difference between the two conditions at a reduced threshold of $p < 0.01$ ($t(21) = 3.00$). A whole-brain analysis also showed bilateral VS activation for the random positive feedback presentation of 1-step random versus 2-step random conditions a reduced threshold of $p < 0.01$uncorrected.

Discussion

Summary and interpretation

The second dissertation study addressed the question of how cognitive effort manipulation would influence outcome valuation and associated striatal activity during trial-and-error learning with intrinsic rewards (positive feedback). Specifically, cognitive effort was varied by means of presenting performance-related feedback after a response to a single stimulus or after a sequence of two responses to two different stimuli. The results showed that the VS was more active in association with the feedback presentation during the HE condition. This finding supports the neural hypothesis based on the contrast theory, which suggests that feedback presented after high cognitive effort as compared to low cognitive effort might be perceived as more valuable and rewarding. However, there might be other factors driving the striatal activity during feedback presentation of the HE condition that are discussed below along with other results.

Cognitive effort manipulation

As expected, the 2-step learning condition was more cognitively demanding. This was reflected in significantly lower performance on the 2-step learning condition as compared to the 1-step learning condition. Consistent with accuracy data, all of the

participants indicated in a questionnaire that the 2-step learning condition was more difficult to learn than the 1-step learning condition.

Positive versus negative feedback

The ANOVA results showed a robust main effect of valence in the striatum. Simple contrasts revealed that the main effect of valence is driven by significant differences between positive and negative feedback presentation in all four conditions. Therefore, this data replicates previous findings with respect to observing differential striatal activation to positive versus negative outcomes in the striatum (Daniel & Pollmann, 2010; Tricomi, Delgado, & Fiez, 2004; Tricomi, Delgado, McCandliss, McClelland, & Fiez, 2006; Tricomi & Fiez, 2008).

Previous studies showed that the VS reacts to presentation of positive and negative outcomes, even when these outcomes are passively delivered, while the dorsal striatum is sensitive only to outcomes that result from one's goal-directed actions and is not sensitive to non-contingent outcomes (Balleine, Delgado, & Hikosaka, 2007; Bjork & Hommer, 2007; O'Doherty et al., 2004; Tricomi, Delgado, & Fiez, 2004; Zink, Pagnoni, Martin-Skurski, Chappelow, & Berns, 2004). Consistent with previous findings, the present results show activations of the dorsal striatum and the VS in association with positive versus negative feedback presentation during the learning conditions, with especially widespread bilateral activations of the dorsal striatum and the VS during the high effort condition. Clusters of activity in the VS that resulted from positive versus negative feedback contrasts of random conditions were smaller in comparison to clusters activated during the 1-step and 2-step learning conditions but overlapped with activity from the learning conditions.

Having response options also might potentially explain VS activity to non-contingent outcomes observed in the current study. Several studies reported that people prefer to have a choice rather that not to have a choice, and that having a choice might be rewarding and associated with perceived control over the environment (Bown, Read, & Summers, 2003; Leotti & Delgado, 2011; Leotti, Iyengar, & Ochsner, 2010). In this experiment, participants had a choice of four response options though the responses did not actually determine the outcomes. That is, having a choice between four options might have been, to a degree, pleasant, driving activity in the VS.

Learning versus random conditions

The main effect of contingency resulted in activation in the dorsal striatum (caudate nucleus). A simple contrast revealed that the activity in the dorsal striatum was driven by the difference between positive feedback presentation of the 2-step learning and 2-step random conditions. At the same time, large clusters of activity were observed throughout the striatum during this contrast, including the cluster in the left VS. It has been shown that different parts of the striatum are involved in processing outcomes that are passively and actively obtained, with dorsal striatum being activated during instrumentally obtained rewards and the VS being activated in association with outcomes that are passively delivered as well as those that are obtained actively (O'Doherty et al., 2004; Zink, Pagnoni, Martin-Skurski, Chappelow, & Berns, 2004). Indeed, comparing positive feedback presentation of the 2-step learning versus random condition, the activity of the dorsal striatum seems to be affected by action-outcome contingency. However, the current results also suggest that the VS in the current context might be more sensitive to actively obtained rewards that require high cognitive effort.

Additionally, no striatal activation was observed when looking at a similar comparison between the 1-step conditions (learning positive versus random positive). The striatal response to positive feedback presentation after the 1-step learning condition is comparable to the response observed in association with the 1-step random condition. This might be due to the interaction of the 1-step random condition with the 2-step random condition. This result potentially supports the effort discounting theory showing that the striatum in the absence of learning might be activated to a greater extent when physical effort requirements are low and the outcome value is indistinguishable in comparison to the high physical effort condition. Judging by parameter estimates from the VS ROI (Figure 8) of the 1-step random positive and the 2-step random positive feedback presentation, it looks like the positive feedback of the 1-step random condition elicited the activity of the VS to greater extent than the positive feedback presentation during the 2-step random condition, even though the difference between the two conditions is not statistically significant at 0.001. An exploratory analysis did reveal significant VS activation at a lower threshold.

Another explanation of the similar activation patterns associated with 1-step learning and random conditions might be participants' engagement in the task. Individuals are prone to perceiving patterns in their environment, even when it is actually absent. Given that the 1-step learning condition is less cognitively demanding and is easier to learn, it is possible that participants were still engaged in the task during the 1-step random condition, perceiving feedback as being the result of their responses and disregarding experimental instructions (Jessup & O'Doherty, 2011). The 2-step learning condition is more cognitively demanding and, in this case, it is hard to learn the action-

outcome relationship. The cognitive load participants are expecting from the 2-step learning condition might make it easier for them to disengage during the 2-step random condition.

Negative feedback comparisons

Comparing negative feedback presentation between conditions did not reveal any striatal activity. Very few contrasts of negative feedback presentation showed any activation, if at all. Usually, negative feedback presentation is characterized by a decrease in striatal activation and thought to be interpreted as a punishment. This was the case in this experiment, as shown by the positive versus negative feedback presentation comparison. However, the differences between the negative feedback conditions were not significant. If one assumes that negative feedback is in fact interpreted as punishment, negative feedback presentation may have been interpreted as similarly punishing in all conditions, regardless of the cognitive effort required. However, participants in the current experiment were not questioned about subjective valuation of negative feedback.

Role of DLPFC in learning

Dorsolateral prefrontal cortex (Brodmann areas 9 and 46) has been generally implicated in a variety of functions that can be characterized as executive processes. More specifically, DLPFC was implicated in processing working memory, planning, reasoning and information integration over time. The DLPFC location in the human brain is comparable to the brain in monkeys with animal literature ascribing similar functional roles to this region as human literature does (Hoshi, 2006; Krawczyk, 2002). Therefore, there seems to be a consensus that during learning and decision-making, the DLPFC may help support the function of information integration during the need for increased

cognitive demand (Boehler et al., 2011; Boettiger & D'Esposito, 2005; McGuire & Botvinick, 2010; Satterthwaite et al., 2012; Schmidt, Lebreton, Clery-Melin, Daunizeau, & Pessiglione, 2012; Seo, Barraclough, & Lee, 2007; Stoppel et al., 2011; Tanaka et al., 2004; Tanaka et al., 2006).

The current fMRI results seem to be in line with the hypothesized function the DLPFC. In the current experiment, the DLPFC (Brodmann area 9) showed an interaction of difficulty by contingency and was activated during feedback presentation of the HE condition (2-step learning condition) when it was compared to feedback presentation of the LE (1-step learning) condition. This analysis was performed not in order to look at differential activation during feedback presentation of different valence but in order to see whether the more effortful condition would elicit activation in areas associated with task complexity and information integration such as the DLPFC. At the same time, positive feedback presentation of the HE condition might be the primary driving force of DLPFC activity in the current set-up since it allows for the most information integration. The DLPFC was also activated during positive feedback presentation of HE condition when it was compared to positive feedback presentation of the LE condition, with random feedback presentation conditions being controlled for. DLPFC activity did not differentiate between negative feedback presentations. It is positive feedback that provides a great amount of information about correct responses on a specific trial and helps most to eliminate response options for other HE trials. In the current context, the DLPFC might act as a node that collects information from early action outcomes of several trials and integrates this information, passing it to the striatum, so that the correct action can be selected for goal-achievement.

Role of VMPFC in decision-making

Several other areas of the prefrontal cortex also showed consistent activation in the current experiment. Of particular importance to outcome valuation is ventromedial prefrontal cortex (VMPFC). The VMPFC represents a rather large area of the prefrontal cortex and can be further segregated onto the medial prefrontal cortex (MPFC (BA 10)) and orbitofrontal cortex (OFC (BA 11)) (O'Doherty, 2011; Rushworth, Noonan, Boorman, Walton, & Behrens, 2011). Numerous animal and human studies showed that these regions play an important role in goal-directed behavior and processing of affective information (Krawczyk, 2002; Mitchell, 2011). Specifically, there is evidence that this region is involved in calculating action value and subjective value of outcomes. The VMPFC can also provide emotional information about decision options causing a person to favor a specific outcome option. In addition, there is evidence that there is a prediction error signal in VMPFC, similar to the prediction error signal observed in the striatum (Alexander & Brown, 2011; Kable & Glimcher, 2007; Krawczyk, 2002; O'Doherty, 2011; Padoa-Schioppa & Assad, 2006; Rangel & Hare, 2010; Wallis & Miller, 2003).

Consistent and widespread activation of MPFC was observed in association with the HE condition of the current task. This region showed an interaction of difficulty by contingency and valence and a main effect of valence and was consistently activated during positive feedback presentation of the 2-step learning condition when it was compared to negative feedback presentation, random positive feedback presentation and positive feedback of the 1-step learning condition. OFC activity showed the main effect of valence and was activated during the contrast of positive feedback presentation of the 2-step learning and 2-step random conditions, in addition to a similar contrast where the

2-step learning condition was compared to the 1-step learning condition with collapsed valence. It is possible that in the current study, these two prefrontal regions in conjunction with the VS reflect participants' preferences for the outcomes of the HE condition. Specifically, MPFC (BA 10) and OFC (BA 11) might be collecting affective information during learning and reflecting subjective preferences of participants during feedback presentation of the HE condition. This type of feedback might generally be perceived as more pleasant during task performance, even though only half of the participants indicated that they preferred feedback presentation after the HE condition.

The VMPFC and DLPFC share a similar pattern of activity in the current study suggesting that these prefrontal regions might interact during learning and outcome processing that involves high cognitive effort. The DLPFC might use the affective information provided by the VMPFC during outcome presentation as part of its information integration function.

Limitations

Questionnaire results. According to the contrast theory (Zentall & Singer, 2007), outcomes that follow higher effort are preferred due to the aversive nature of the action and due to a contrast between the aversive action and rewarding nature of an outcome. Participants were asked to indicate on the questionnaire the condition during which they were more engaged in the learning process and the type of feedback presentation that they preferred (i.e. feedback presented after the 2-step learning condition or the 1-step learning condition). Questionnaire data indicates that only half of the participants preferred cumulative feedback presented during the 2-step learning condition that was rated as more difficult and cognitively demanding. Unfortunately, because participants

were not questioned about the reason for their preference, it is not clear whether participants preferred 2-step learning outcomes because of an aversion to the higher cognitive effort (and the contrast between this aversion and the positive outcome) or because the HE condition is more challenging. That is, difficulty in learning might differentiate the two learning conditions in terms of how challenging (rather than aversive) they are. Participants might prefer actions that are challenging and would interpret outcomes after such actions as more rewarding. Both cases could lead to the same pattern of striatal activity and outcome valuation.

Prediction error. As stated above, only half of the participants perceive feedback after the 2-step learning conditions as more rewarding while the other half did not and preferred feedback presentation after the easy 1-step learning condition. Therefore, greater activation of the ventral striatum cannot be solely explained by the notion that the outcome that follows the high effort is more valuable.

It has been established that dopamine neurons carry a prediction error signal. That is, the dopamine neurons that project to the striatum increase their firing rate when there is a discrepancy between observed and expected outcomes (i.e. when the real outcome is better than expected). At the same time, these neurons reduce their firing rate when the outcome is worse than expected (Schultz, 2002). A similar pattern has been observed in the hemodynamic response in the striatum (Delgado, Nystrom, Fissell, Noll, & Fiez, 2000). Therefore, given that participants' performance in the 1-step condition is significantly better than participants' performance during the 2-step condition, the prediction error for the 2-step condition would be higher since the outcome remains unpredictable during this condition. This might drive the hemodynamic response and

result in stronger activation associated with the 2-step condition that was observed.

At the same time, the conditions with the highest PE in the current study should be the random conditions, where it was impossible to learn action-outcome contingencies and feedback valence was always unexpected. However, comparing positive and negative feedback presentation in both 2-step and 1-step random conditions revealed only a small activation in the VS. Similarly, comparison of 2-step learning condition to the 2-step random condition resulted in widespread striatal activation. Therefore, without eliminating the possibility that the PE plays a role in the comparison of the 2-step and 1-step learning conditions, intrinsic motivation should also have an effect on VS activity and outcome valuation after the HE condition, contributing to the observed VS activity.

Chapter Four: Experiment 3:

Self versus other outcome relevance – a behavioral study

The two previously described studies show how context and motivation influence reward valuation and striatal activity during learning. The third study attempts to show that outcome valuation can be affected not only by the delay between an action and an outcome or by the effort one has to exert to reach an outcome, but also by social context. Would one make an effort to learn on behalf of another person in order to remove inequity created in a social context between oneself and the disadvantaged individual?

Social context can exert a powerful influence on how outcomes are valued. For example, in competitive situations, people may interpret a competitor's failures or punishments as their own rewards. People often compare themselves to others, and may see others' failings as indicators of their own superiority (de Bruijn, de Lange, von Cramon, & Ullsperger, 2009; Fliessbach et al., 2007; Howard-Jones, Bogacz, Yoo,

Leonards, & Demetriou, 2010; Takahashi et al., 2009). This behavior stems from the notion that people tend to pursue their own interests and care about maximizing their own gains. However, there is a large body of experimental evidence showing that individuals also care about outcomes that do not benefit them personally and that preserve fairness (Fehr & Schmidt, 1999; Gintis & Fehr, 2012; Glockner & Hilbig, 2012; Gummerum, Hanoch, & Keller, 2008; Loewenstein, Bazerman, & Thompson, 1989; Rilling et al., 2002).

Reputation considerations and social norms encourage fairness (Kogut, 2012). Social norms are conventions in a society that usually require acting not in self-interest but for the interest of a particular group (Rossano, 2012). Probably the best examples of individuals following social norms and giving up personal goods for the benefit of others are taxation and charitable donation (Rossano, 2012). That is, besides having a goal of augmenting personal gains and maximizing their own rewards, people tend to care for the welfare of others even when those individuals are genetically unrelated. Individuals behave altruistically in order to benefit another person and punish those who do not follow the common social norms (Fehr & Schmidt, 1999; Lee, 2008; Loewenstein, Bazerman, & Thompson, 1989). Through such actions individuals try to restore fairness in social situations and the outcomes of such actions might be experienced as rewarding (Harbaugh, Mayr, & Burghart, 2007; Kogut, 2012).

The desire to be fair might be one of the factors motivating inequality aversion. Inequality aversion refers to the tendency to oppose inequitable and unfair outcomes. This tendency is often strong enough that one would give some material good to another person in order to reduce the inequality in a situation (Fehr & Fischbacher, 2003; Fehr &

Schmidt, 1999; Lee, 2008; Tricomi, Rangel, Camerer, & O'Doherty, 2010). There are several types of situations that people can perceive as inequitable and unfair. One can feel as if he was treated unfairly and others had put him in a disadvantageous position (disadvantageous inequality) or one can observe another person being treated unfairly by others. Similarly, an individual might be involved in treating someone in an unfair way. In case of advantageous inequality, an individual is endowed with goods and, at the same time, witnesses another person being in a disadvantageous position. People were shown to be sensitive to these types of unfairness. Therefore, individuals seem to be averse to inequality regardless of their role in the social situation. That is, inequality aversion occurs whether one is the agent or the victim of an unfair action or whether one is just an observer (Fetchenhauer & Huang, 2004; Johansson & Svedsater, 2009).

Of interest is the finding that shows that inequality aversion develops with age. Specifically, it has been shown that young children prefer to not share their candies with others, while older children, as they progress in their development, exhibit sharing behavior more often (Fehr, Bernhard, & Rockenbach, 2008). Additionally, children of younger ages know about fairness norms but their behavior is driven by self-interest. They also report greater levels of satisfaction after their non-sharing behavior. Older children do share more; however, their levels of satisfaction from the sharing behavior are not as high as the levels of satisfaction of the oldest children, who exhibited the most sharing behavior and, hence, inequality aversion. It seems that self-benefitting behavior gets overshadowed by the time children get older and incorporate fairness norms (Gummerum, Hanoch, & Keller, 2008; Kogut, 2012).

Hemodynamic activity in the striatum observed when a participant is in a social context or presented with social rewards is comparable to the activity observed during tasks that feature extrinsic and intrinsic rewards (monetary and performance-related feedback) (Izuma, Saito, & Sadato, 2008; Lee, 2008; Rademacher et al., 2010; Takahashi et al., 2009), demonstrating that social context indeed has a powerful influence on outcome valuation and brain activity. Recently, it was also shown that the striatum is one of the structures involved in processing inequality-averse social preferences, as was evidenced in a study by Tricomi et al. (2010). In this study, participants evaluated monetary transfers to either themselves or to the other person after either being endowed with a monetary bonus or not. Stronger activation was observed in the ventral striatum while participants, who received the monetary bonus, rated transfers of payments for the other person who did not receive any monetary bonus. Similarly, stronger activation was observed in the ventral striatum of participants who did not receive any endowment while they rated monetary bonus transfers to themselves. In both cases, payment transfers allowed reduction of the inequality between the two individuals and both individuals rated these transfers as preferable. The activity of the ventral striatum tracked the preference that allowed to reduce the difference in endowment between individuals (Tricomi, Rangel, Camerer, & O'Doherty, 2010).

In a similar vein, if an aversion to inequality and a tendency for fairness is fairly strong in individuals, it might be reflected in their performance in a learning task. For example, would a person exert more effort on a learning task (reflected in a significantly better performance) in order to benefit another individual and that way lessen the inequality gap between oneself and the disadvantaged person? For an inequality-averse

individual, performance-related feedback should be more valuable when it reflects a positive outcome of his or her learning that benefits the disadvantaged individual. My third study attempts to prompt such inequity-averse behavior. In the current study, each participant was placed in an unequal social situation. It is important to note that the unequal situation was set up so that it benefited the participant and was unfair in relation to a confederate. The participant witnessed an unequal distribution of a monetary endowment at the beginning of the study and was given a chance to reduce the created inequality between himself or herself and another person by winning a monetary bonus for the other person in a learning task. Given the tendency of people to be altruistic and reduce the inequality in social situations, participants should try winning money for the disadvantaged person during the task. In such a situation, feedback will not only reflect the monetary gain on behalf of another person, but it will also reflect goal-achievement. In this case, the goal of the participant, who cares for the welfare of an unfamiliar individual, might be to reduce the experimentally created inequality and to restore fairness. However, the goal of the participant who is self-interested and does not care about other's welfare would have a goal of maximizing one's own gain.

As previously mentioned, there were two confederates involved in the experiment. The second confederate received the same amount of money as the participant and performed an identical task. Therefore, in all respects the role of the third person in the experiment resembled the role of the participant. A participant's performance on the learning task for the third person might be motivated by two different factors. First, it is possible that the participant would try to win money for the advantaged confederate due to the feeling of reciprocity, since the participant thought that the second

confederate was doing the same task and receiving the same instructions. At the same time, it is also possible that the participant might not really try to win a monetary bonus for the advantaged confederate because that person was not treated unfairly. That is, the accuracy for the advantaged confederate (the Other High condition) is hypothesized to be equivalent to the participant's performance in the condition during which the participant can win money for himself (the Self condition). Alternatively, participants' accuracy during the Other High condition can be higher than performance on the Self condition due to the feeling of reciprocity. However, I expected to observe better performance when participants are trying to win a monetary bonus for someone who was treated unfairly (Other Low condition) relative to when they were playing for someone in the similar situation as themselves (Other High condition) due to inequality aversion.

Methods.

Participants.

A total of 40 subjects participated in the experiment; however data from only 30 subjects were analyzed (17 females; mean age 19.6 years, SD 2.2) due to some subjects indicating to have not understood the task at the end of the study or forgetting the names of the confederates. Ten females participated in this experiment for the payment of $10. An additional set of 10 females and 20 males participated in this experiment for research credit and were paid $5 as part of the experimental manipulation. All participants provided written consent to participate. Confederates in this experiment were of the same gender as a participant in order to prevent gender interaction effects.

Procedure.

Behavioral paradigm. At the beginning of the experiment, each participant was introduced to two confederates. The participants were told that the experiment in which all three participants were participating was intended for the study of learning strategies. The participant came with the expectation that he or she would be paid for the experiment. However, at the beginning of the experiment, the experimenter informed the participant and the confederates that only two individuals will be paid for performing a challenging computer task, while the third person will fill out a short survey. All three individuals were then asked to draw straws in order to determine which two of them would receive monetary compensation and perform a challenging computer task, and which one of them would receive no monetary compensation and would just complete a survey. The person who drew the shortest straw was assigned to the survey; the procedure of drawing straws was always fixed so that one of the confederates drew the short straw. The participant was led to believe that the order in which the individuals straws was random. After drawing the straws, two $10 (paid version) or $5 (research credit version) bills were distributed to the participant and the confederate who did not draw the short straw. The confederate and the subject were then escorted to separate testing rooms to complete the computer task.

The computer task that the participant completed in the testing room was adapted from Experiment 2. Participants were presented with the 2-step learning condition only, where stimuli presentation occurred sequentially. That is, participants had to respond to the first stimulus before being presented with the second stimulus. Participants were also told that they had the opportunity to earn more bonus money for themselves and for each of the confederates. That is, participants had an opportunity to win the bonus for

themselves (Self condition), for the confederate who was doing the survey and did not get any monetary bonus previously (Other Low condition), and for the confederate who was also doing the computer task and got the monetary bonus (Other High condition).

The experiment ended with a questionnaire that asked: 1) during which condition they felt more engaged in the task, 2) whether they intentionally responded incorrectly, 3) whether they made an effort to get a bonus for the confederate who did not get the bonus previously, and 4) whether they believed the set-up situation. Participants also had to indicate the name of the confederate who did not get any monetary bonus and was doing the survey, and the name of the confederate who got the monetary bonus as the participant. This was a manipulation check to make sure that the participants were paying attention and remembered which person was which when performing the task. Participants were debriefed at the end of the experiment and received an additional payment of $5.

Data analysis.

Accuracy data and response time data were analyzed using a within-subject ANOVA with a factor of beneficiary (Self vs. Other Low vs. Other High). This was done in order to see whether participants learning differed between beneficiary conditions. Post-hoc two-tailed t-tests were also run to determine whether there were specific differences between beneficiaries for whom participants attempted to win more money and whether there were specific response time differences between correct and incorrect trials.

Results.

Accuracy.

A within-subject ANOVA showed that there is a significant difference in participants' accuracy across conditions, F (2, 29) = 5.01, p = 0.03. A post-hoc two-tailed within-subject t-test revealed that participants tried to win money specifically for the confederate who did not receive any monetary bonus rather than for the other confederate who got the same amount of money as the participant. That is, there is a significant difference between Other High and Other Low conditions (t(29) = 2.28, p = 0.03). At the same time, there were no significant differences observed when comparing Self versus Other High conditions and Self versus Other Low conditions (t(29) = 0.39, p = 0.70 and t(29) = 1.33, p = 0.20, respectively) (Figure 16).

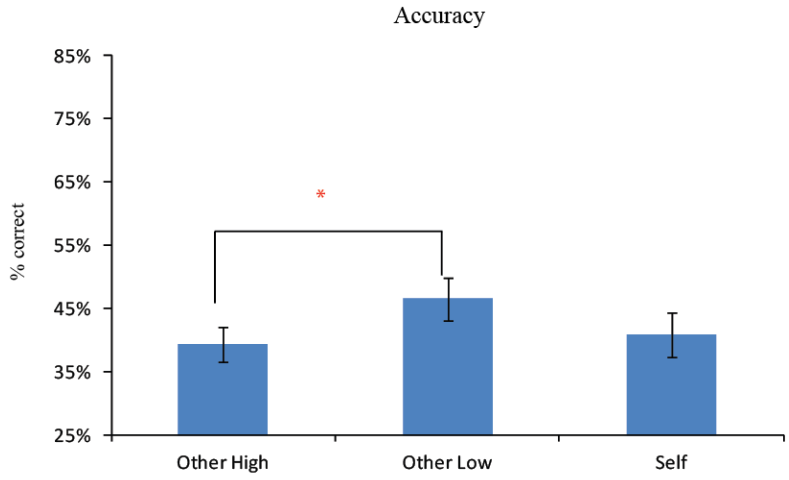

Figure 17. Accuracy results for the three learning conditions. Other High = participants' accuracy while trying to win money for the confederate who got the bonus; Other Low = participants' accuracy while trying to win money for the confederate who did not get the bonus; Self = participants' accuracy while trying to win money for themselves. * denotes significant difference with $p < 0.05$.

Response time data.

A within-subject ANOVA with factors of valence (correct vs. incorrect) and beneficiary (Self vs. Other High vs. Other Low) showed that there is a significant main effect of beneficiary in participants' response time (RT), F (2, 29) = 5.8, p = 0.02. A post-hoc two-tailed within-subject t-test revealed that the main effect of beneficiary was

driven primarily by the significant difference in incorrect RTs between Other Low versus Other High conditions and Self versus Other High conditions (t(29) =2.08, p = 0.05, t(29) = 2.73, p = 0.01, respectively). At the same time, there were no significant differences observed when comparing RTs associated with correct responses (Figure 17).

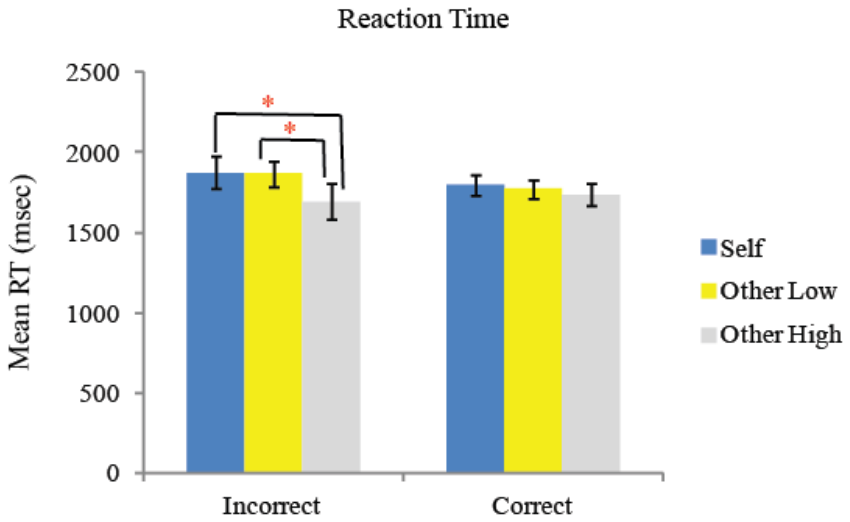

Figure 18. Response time (RT) results for the three learning conditions. Other High = participants' RT to correct and incorrect trials, while trying to win money for the confederate who got the bonus; Other Low = participants' RT to correct and incorrect trials, while trying to win money for the confederate who did not get the bonus; Self = participants' RT to correct and incorrect trials, while trying to win money for themselves. * denotes significant difference with $p < 0.05$.

Questionnaire data.

Based on the questionnaire responses of subjects, 22 out of 30 participants indicated that they tried to win a monetary bonus for the disadvantaged confederate. At the same time, 13 out of 30 participants were more engaged in the learning task during the Self condition, while nine participants indicated that they were more engaged in the learning task during the Other Low condition. The rest of the subjects did not provide clear responses for that question. In addition, none of the participants intentionally responded incorrectly while performing the task.

Reported helpers

An additional analysis was performed on the group of subjects (n = 22) who specifically reported trying to win money for the disadvantaged confederate (Other Low) to see whether they would show a stronger effect between Other High and Other Low conditions, that was observed in the full data set. However, the two-tailed within-subject t-test did not reveal a significant difference between the conditions ($t(21) = 1.85$, $p = 0.08$).

Discussion.

In everyday life, humans are immersed in multiple social situations. Therefore, it is important to study learning and decision-making in a social context. Social context influences how learning and decision-making outcomes are interpreted and evaluated. A variety of studies have investigated human behavior in competitive situations or in situations in which participants behaved in an altruistic manner, putting themselves in disadvantageous position (Bolton & Ockenfels, 2000; de Bruijn, de Lange, von Cramon, & Ullsperger, 2009; Howard-Jones, Bogacz, Yoo, Leonards, & Demetriou, 2010). In my third study, I attempted to show that participants would exert cognitive effort in order to benefit another person. Outcomes that reflect to participants the welfare of a disadvantaged individual and, at the same time, do not benefit them personally were still relevant and resulted in better performance in the feedback-based learning task.

Subjects were immersed in a social interaction with two unfamiliar persons and witnessed an unfair distribution of money. That is, at the beginning of the experiment, the subject received his or her endowment together with one of the confederates, while the second confederate received an endowment of a lesser amount. This manipulation was

performed in an attempt to elicit inequality aversion through an unfair distribution of monetary bonuses, so that the participant would attempt to reduce the inequality by exerting cognitive effort in a learning task and winning money for the disadvantaged individual. As current results showed, participants indeed tried to win money for the confederate who received a lesser amount of money at the beginning of the experiment. This is reflected in a significantly better performance on a learning task when participants' performance in the Other Low condition is compared to participants' performance in the Other High condition. That is, while performing an identical learning task, feedback significantly improved participants' performance in the condition where feedback reflected the welfare of the disadvantaged confederate. In addition, a greater number of participants indicated on the questionnaire that they tried to win money for the disadvantaged confederate.

An alternative interpretation of the current result might be that participants were not trying to win a bonus for the disadvantaged confederate, but rather were not engaged in the task enough while performing the Other High condition. However, questionnaire data indicates that none of the participants intentionally responded incorrectly while 'playing' for the confederate who received the same endowment as they did. This might mean that participants did not dislike or feel competitive towards this confederate.

Participants' accuracy for the Other High condition seems to be similar to their performance for the Self condition. However, there was little evidence of reciprocity, since participants performed slightly worse in the Other High condition than in the Self condition. This is suggestive of the possibility that the role of the confederate who

received the same endowment as the participant was perceived by the participant as very similar to oneself and as someone who was not treated unfairly.

At the same time, some of the participants indicated that they were more engaged in the learning task while they were trying to win money for themselves, even though the accuracy was still better in the Other Low condition than in the Self condition. In addition, there was no significant difference observed between the Self conditions and the Other Low. This might indicate that participants were not completely unselfish and were also trying to maximize their own gains in addition to thinking about the welfare of the disadvantaged confederate and attempting to win money for that confederate.

Another distinguishing component of the current study is that, unlike in previous experiments, where participants had to share or give up some part of their endowment, in the current study the endowment the participants received was fixed. Therefore, in the current study participants had to exert cognitive effort in order to win extra money and make the distribution of money more equal and fair. However, the limitation of this study that weakens the above statement is that participants were never asked about whether they perceived the set-up social situation as fair or unfair.

Furthermore, several studies report increased striatal activation during prosocial behavior, suggesting that outcomes of actions that benefit others might be rewarding (Harbaugh, Mayr, & Burghart, 2007). In addition, the results obtained in the study of Tricomi et al. (2010), showed that striatum was activated to a greater degree when the participant with a monetary endowment was rating monetary transfers to the disadvantaged participant. Therefore, given the current results, one might expect to see an increase in striatal activation when participants respond correctly and are presented with

positive feedback while playing for the disadvantaged confederate (Other Low condition), as participants are trying to reach the goal of eliminating the unfair situation of inequitable distribution of money. This will be an interesting question for future investigations.

To summarize, in the current study participants were presented with an endowment at the beginning of the study and had a chance to act in a self-benefitting manner and in other-benefitting manner during the course of the experiment. That is, participants could win even more money for oneself in the trial-and-error feedback-based learning task. However, participants tried to win money for the unfamiliar individual whom they just met but who did not get any monetary endowment. Even though a large group of participants indicated that they were more engaged in the task while trying to win money for themselves, this unfair social situation and unequal distribution of endowments seemed to motivate learning (maybe even unconsciously) that benefited the disadvantaged individual. Feedback presentation in the current set-up seemed to reflect goal-achievement, i.e. improving one's performance in order to win money for the benefit of another person. That is, the results of this study suggest that performance-related feedback is still relevant and carries value even when it does not reflect personal gain. People are willing to exert cognitive effort on behalf of an unfamiliar individual in order to reduce the inequality.

Chapter Five: General Discussion

Summary

This dissertation work explored striatal functioning during learning from performance-related feedback presented under different time frames and cognitive effort

requirements. It was demonstrated that the striatum is sensitive to the specific conditions of the learning context. The striatum is selectively involved in learning from immediately presented outcomes and is modulated by the degree of cognitive effort required to obtain such outcomes. Additionally, processing of feedback after a substantial delay switches to more posterior parts of the basal ganglia, suggesting that processing of delayed and immediate outcomes might not depend on the same mechanism. Also, learning can be affected by an unfair social context to obtain other-relevant outcomes. It was demonstrated that people are motivated by the unfair social context to exert cognitive effort during learning in order to eliminate the difference in endowment between themselves and an unfamiliar disadvantaged individual. Together, these findings suggest that when given a goal of achieving task success during learning, the subjective outcome value may be affected by delay and effort requirements, modulating striatal BOLD response to such outcomes.

Delay and Effort Costs

Presenting performance-related feedback immediately or after a substantial delay seems to change the affective nature of the feedback. Participants' questionnaire responses were in line with such an interpretation. Immediate positive feedback was rated as more rewarding while negative feedback was rated as more punishing compared to delayed feedback presentation of the corresponding valence. Indeed, during experiment 1, performance-related feedback was characterized by differential activation of the ventral striatum, with a rise and a decrease below baseline shown to be characteristic of rewarding and punishing outcomes (Delgado, Nystrom, Fissell, Noll, & Fiez, 2000), only when presented immediately but not after a delay.

The ventral striatum also showed the characteristic rise and decline below baseline in BOLD response during experiment 2, where each type of feedback was presented immediately but differed in terms of effort requirements. Results of experiment 2 also suggest that cognitive effort required to obtain an outcome greatly influences striatal activity, with outcome-related activity of the striatum scaling up with increased cognitive effort requirements. That is, the striatum exhibited increased activation to feedback that followed higher cognitive effort demands, even though participants' preferences regarding outcomes that follow effortful actions varied.

Even though the two experiments differ in the nature of the learning paradigm, with experiment 1 utilizing a declarative learning task while experiment 2 involved trial-and-error learning, performance-related feedback presentation seems to elicit a similar response in the ventral striatum across both tasks. However, modifying the context of learning so that the feedback is presented after a delay switches the neural mechanisms of learning away from the striatum, engaging regions of the posterior basal ganglia such as the globus pallidus. The different activation patterns suggest that delayed feedback may be processed differently than immediate feedback, possibly in a more cognitive and less affective manner.

Regardless of different neural mechanisms involved in learning from immediate and delayed feedback, the delay between an action and an outcome did not seem to interfere with learning of action-outcome contingencies. This type of feedback presentation resulted in performance similar to that associated with immediate feedback presentation. In addition, feedback after a delay was not rated as rewarding or punishing as the immediate feedback was, also supporting the interpretation that the affective nature

of feedback changes under different time frames. Negative feedback presentation, in particular, might not be perceived as punishing when presented after a substantial delay. Therefore, although a student receiving test results back might be able to incorporate both immediate and delayed feedback to modify his future behavior, a substantial delay between the test and the feedback might permit the student to view the results with a 'cool head' rather than in an emotional way.

Dorsal Striatum

The dorsal striatum was activated in both experiment 1 and experiment 2. The left body of the caudate nucleus was activated to presentation of immediate feedback in experiment 1, while during experiment 2 the dorsal striatum (i.e. the caudate nucleus and the putamen) was activated bilaterally during the 2-step learning condition for the contrast of positive versus negative feedback and for the contrast of positive feedback of the 2-step learning versus random conditions. The left dorsal putamen was activated during the 1-step learning condition for the contrast of positive versus negative feedback.

The dorsal striatum has been shown to be involved in instrumental learning and is considered to play a primary role in implementing instrumental action. In the actor-critic reinforcement learning model, the dorsal striatum is thought to play a role of an 'actor', i.e. the structure that uses the prediction error signal from the ventral striatum to modify the information about action-outcome contingencies (O'Doherty et al., 2004). O'Doherty and others (2004) showed that the dorsal striatum was activated only during instrumental learning task but not during Pavlovian learning during a delivery of liquid rewards, suggesting that the dorsal striatum plays a role in implementing goal-directed actions.

The dorsal striatum acting as an 'actor' during learning also suggests that this region might be engaged when one perceives himself or herself as an agent of an action that causes the outcome. Besides being shown to be activated during instrumental learning tasks, the dorsal stratum has been shown to be engaged when one simply perceives a contingency between one's own action and the outcome (Elliott, Newman, Longe, & Deakin, 2004; Tricomi, Delgado, & Fiez, 2004). For example, Tricomi and others (2004) showed that during random feedback presentation, the dorsal striatum was activated when participants thought that the feedback is dependent on their action, which was not the case when participants were told that the feedback is actually random.

In experiment 2, presentation of random feedback, when participants were aware of feedback being non-contingent on their actions, did not elicit activation of the dorsal striatum but was associated with only ventral striatal activity. Since the random feedback was not contingent upon the subjects' responses, this activation pattern is in line with previous evidence of the ventral striatum activated to non-contingent rewards (O'Doherty et al., 2004) and with the dorsal striatum being activated when one perceives a contingency between an action and an outcome (Tricomi, Delgado, & Fiez, 2004).

Ventral Striatum

The two fMRI experiments of this dissertation showed that the activity of the ventral striatum is affected by the learning context. The ventral portion of the caudate head was activated during immediate feedback presentation of experiment 1, while large bilateral portions of the ventral striatum (including the NAcc and the ventral head of the caudate nucleus) was activated during the learning conditions of experiment 2. This suggests that the activity of the ventral striatum is sensitive to immediate feedback

presentation and reflects cognitive effort demands of the preceding action during presentation of an outcome of that action.

Previous literature suggests that the ventral striatum (NAcc in particular) and dopamine play an important role in delay and effort-based outcome learning (Assadi, Yucel, & Pantelis, 2009; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). That said, it has been suggested that there are separate neural mechanisms for processing of different action costs. As is evident from the current fMRI results, this dissociation of mechanisms might also be present during learning.

Neuroimaging evidence that points to different neural mechanisms involved in delay and effort processing is presented by Prevost and others. In a task designed to compare delay and physical effort discounting, they showed that while behaviorally outcomes were discounted similarly after effort and delay, the ventral striatal activity showed discounting to delayed rewards only. That is, activity of the ventral striatum to the rewarding outcome was shown to be not discounted by effort costs (due to the absence of correlation between subjective reward ratings and activity of the stratum to high reward outcomes) (Prevost, Pessiglione, Metereau, Clery-Melin, & Dreher, 2010). The results of this study go against the statements from previous fMRI experiments that investigated the cost of effortful actions (Botvinick, Huffstetler, & McGuire, 2009; Kurniawan et al., 2010; Treadway et al., 2012), but the results are comparable to those obtained in previous delay-discounting work (Ballard & Knutson, 2009; Kable & Glimcher, 2007). It is hard to reconcile the divergent findings from this particular study, but it does stress the dissociation of mechanisms involved in processing the costs of delay and effort. In the context of the current experiment, given that participants were placed in

the learning environment (unlike in previous experiments) and had a goal to learn the task to the best of their ability, the effort manipulation in experiment 2 was potent enough to influence the subjective valuation of learning outcomes and associated activity of the ventral striatum while the delayed outcome presentation in experiment 1 engaged a different region that receives dopaminergic inputs, namely the globus pallidus.

Globus Pallidus

The globus pallidus (GP) was activated during both fMRI studies of this dissertation. Specifically, the GP showed persistent activation during delayed negative feedback presentation in comparison with no feedback presentation in experiment 1 and in experiment 2 during the learning conditions and the contrast of positive feedback between the 2-step learning and random conditions.

Similar to the dorsal caudate nucleus, the GP is connected to the DLPFC, which has been shown to be involved in information integration and working memory processes (Draganski et al., 2008), while the inputs of the ventral striatum and the amygdala go to the ventral portions of the GP (Haber & Knutson, 2010). Human and animal studies showed this region to be involved in incentive motivation and learning tasks suggesting a shared role of the striatum and GP in outcome valuation that motivates action. For example, Pessiglione and others showed increased activation of the ventral GP to supraliminal and subliminal presentations of monetary rewards of higher magnitude (Pessiglione et al., 2007). At the same time, Tindell and others (2004) showed that ventral GP neurons fire to the first predictor of the conditioned stimulus that signals reward delivery during classical conditioning, but show no firing when the presented stimulus signals the absence of a reward. During the actual reward presentation, the ventral GP

neurons fired before and after learning occurred. These findings suggest that the ventral GP plays a complex role in learning, motivation and reward processing (Tindell, Berridge, & Aldridge, 2004). Similar findings on GP were reported in an fMRI study during a gambling task with increased GP activity during high and increasing monetary gains (Elliott, Friston, & Dolan, 2000).

While showing positive signals to rewarding outcomes, GP has also been shown to reflect negative outcomes, showing positive functional connectivity with the lateral habenula during error detection (Ide & Li, 2011). This is thought to be explained by the lateral habenula projection neurons that can be found in the GP (internal capsule) and that show the reverse firing rate to positive outcomes (i.e.as the lateral habenula neurons ceased firing to the absence of the reward, the projection neurons to the GP increase their firing rate) (Hong & Hikosaka, 2008). It remains to be investigated what attributes of a task activate the GP. This region receives numerous projections that might contribute to rather discordant signals observed during various tasks, making the exact role of the GP in reward-based learning unclear. But given the multitude of inputs and outputs of GP, this region might be a BG node that links reward signals and cognitive and emotional processes (Elliott, Friston, & Dolan, 2000).

Overcoming Effort Costs to Reduce Social Inequity

Effortful actions are costly and individuals prefer to avoid such actions. However, experiment 3 showed that people are willing to exert effort during learning to reduce inequity. Human tendency to be averse to inequality, the tendency to dislike an inequitable distribution of goods (Fehr & Schmidt, 1999), was shown to motivate participants' to exert cognitive effort and learn the task better for the benefit of another

person. This experiment used a novel set up where the social situation was created at the beginning of the experiment with participants witnessing an inequitable distribution of money. During the learning task participants were by themselves, without further interaction with the confederates. That is, the social context had to prime participants for the learning task. Indeed, participants performed significantly better in the learning condition that allowed reducing the gap in endowment created during the set up social situation at the beginning of the experiment. Even though most participants indicated that they were trying to win more money for themselves, the learning accuracy was better for the condition that endowed the disadvantaged confederate with money. This suggests that individuals, some maybe unconsciously, act in other-regarding fashion and outcomes of actions that reflect the benefit of others might still be relevant to them. Therefore, the unique finding of this experiment is that, for the first time, individuals were shown to be willing to exert cognitive effort during learning in order to reach a goal of mastering the task and of winning money for another person to eliminate inequality.

From an evolutionary and economic perspective such behavior might be considered irrational. A person who is trying to maximize own benefits would try to win more money for him or herself during learning. However, it has been shown time and again that humans have other-regarding preferences and care about the wellbeing of people who surround them (Fehr & Schmidt, 1999; Loewenstein, Bazerman, & Thompson, 1989). Other-regarding preferences might be an attribute of human behavior only. Even though non-human primates were shown to have some sense of fairness, at the same time, in situations where they can benefit another individual, they do not exhibit other-regarding preferences (Silk et al., 2005). That being said, inequity aversion is not a

hard-wired human tendency, but a social norm individuals learn to follow. Yong children were shown to dislike behaving according to social norms. But humans become aware of them from a young age and incorporate them into their behavior around age 10 (Kogut, 2012).

Numerous self-report and neuroimaging studies that utilized economic games suggest that altruistic, fair and cooperative behavior is associated with an increase in positive emotions. This in turn, leads to the activation of reward circuitry in the brain, including activation of the striatal regions (Tabibnia & Lieberman, 2007). It remains to be investigated how the striatum will react given the context of this experiment and whether feedback presentation that reflects another person's gain activates the striatum to a greater extent than feedback that reflects personal gain. Given the results from the first two studies, I predict that the subjective value of outcomes would be higher during the Other Low condition than during the Self condition, resulting in greater activation of the striatum. Indeed, some neuroimaging studies showed that the striatum is sensitive to vicarious reward presentation (Mobbs et al., 2009). Specifically, Mobbs and others (2009) reported that observation of another person winning was rated as pleasant and was associated with an increase in the VS activity. Similarly, Cooper and others showed DS involvement when participants were observing another person perform an instrumental task (Cooper, Dunne, Furey, & O'Doherty, 2012).

Conclusion

Human learning mechanisms had to adapt to the complex world that requires calculation of action costs and weighting the benefits of action outcomes. Hence, the learning mechanisms are flexible and sensitive to contextual influences. Previous

investigations of action-outcome learning analyzed neural mechanisms associated with immediate feedback presentation or with subjects' preferences about delayed and effortful outcomes. However, humans often perform complex actions that influence outcomes in the future rather than immediately reflect the consequence of their actions. Additionally, humans are still motivated to perform effortful actions and are able to learn action-outcome contingencies. This dissertation research explored the striatal responses to feedback under conditions of different feedback timing and different effort costs, and demonstrated the influence of social inequity on the willingness to exert cognitive effort in a learning task.

The striatum is an important component of the neural learning mechanism and it is important to investigate the extent of striatal involvement in learning of action-outcome contingencies to further our understanding of the role of the striatum. I showed that the striatum is involved in learning from immediate outcomes and tracks effort requirements of these outcomes. At the same time, the globus pallidus was involved in learning from delayed outcomes and also reflected effort requirements of immediate outcomes. In addition, it was also shown that unfair social context can motivate individuals to exert effort during learning for the benefit of another person. Taken together these experiments suggest that delay and effort costs can modulate outcome value and associated striatal activity in the context of learning of action-outcome contingencies.

References

Alessandri, J., Darcheville, J. C., Delevoye-Turrell, Y., & Zentall, T. R. (2008). Preference for rewards that follow greater effort and greater delay. *Learning & Behavior, 36*(4), 352-358.

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience, 14*(10), 1338-U1163.

Apicella, P., Scarnati, E., Ljungberg, T., & Schultz, W. (1992). NEURONAL-ACTIVITY IN MONKEY STRIATUM RELATED TO THE EXPECTATION OF PREDICTABLE ENVIRONMENTAL EVENTS. *Journal of Neurophysiology, 68*(3), 945-960.

Assadi, S. M., Yucel, M., & Pantelis, C. (2009). Dopamine modulates neural networks involved in effort-based decision-making. *Neuroscience and Biobehavioral Reviews, 33*(3), 383-393.

Ballard, K., & Knutson, B. (2009). Dissociable neural representations of future reward magnitude and delay during temporal discounting. *Neuroimage, 45*(1), 143-150.

Balleine, B. W., Delgado, M. R., & Hikosaka, O. (2007). The role of the dorsal striatum in reward and decision-making. *Journal of Neuroscience, 27*(31), 8161-8165.

Bjork, J. M., & Hommer, D. W. (2007). Anticipating instrumentally obtained and passively-received rewards: A factorial fMRI investigation. *Behavioural Brain Research, 177*(1), 165-170.

Boehler, C. N., Hopf, J. M., Krebs, R. M., Stoppel, C. M., Schoenfeld, M. A., Heinze, H. J., et al. (2011). Task-Load-Dependent Activation of Dopaminergic Midbrain Areas in the Absence of Reward. *Journal of Neuroscience, 31*(13), 4955-4961.

Boettiger, C. A., & D'Esposito, M. (2005). Frontal networks for learning and executing arbitrary stimulus - Response associations. *Journal of Neuroscience, 25*(10), 2723-2732.

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review, 90*(1), 166-193.

Botvinick, M. M., Huffstetler, S., & McGuire, J. T. (2009). Effort discounting in human nucleus accumbens. *Cognitive Affective & Behavioral Neuroscience, 9*(1), 16-27.

Botvinick, M. M., & Rosen, Z. B. (2009). Anticipation of cognitive demand during decision-making. *Psychological Research-Psychologische Forschung, 73*(6), 835-842.

Bown, N. J., Read, D., & Summers, B. (2003). The lure of choice. *Journal of Behavioral Decision Making, 16*(4), 297-308.

Cardinal, R. N., Robbins, T. W., & Everitt, B. J. (2000). The effects of d-amphetamine, chlordiazepoxide, alpha-flupenthixol and behavioural manipulations on choice of signalled and unsignalled delayed reinforcement in rats. *Psychopharmacology, 152*(4), 362-375.

Cheung, T. H. C., & Cardinal, R. N. (2005). Hippocampal lesions facilitate instrumental learning with delayed reinforcement but induce impulsive choice in rats. *Bmc Neuroscience, 6*.

Cocker, P. J., Hosking, J. G., Benoit, J., & Winstanley, C. A. (2012). Sensitivity to Cognitive Effort Mediates Psychostimulant Effects on a Novel Rodent Cost/Benefit Decision-Making Task. *Neuropsychopharmacology*.

Cooper, J. C., Dunne, S., Furey, T., & O'Doherty, J. P. Human Dorsal Striatum Encodes Prediction Errors during Observational Learning of Instrumental Actions. *Journal of Cognitive Neuroscience, 24*(1), 106-118.

Croxson, P. L., Walton, M. E., O'Reilly, J. X., Behrens, T. E. J., & Rushworth, M. F. S. (2009). Effort-Based Cost-Benefit Valuation and the Human Brain. *Journal of Neuroscience, 29*(14), 4531-4541.

Daniel, R., & Pollmann, S. (2010). Comparing the Neural Basis of Monetary Reward and Cognitive Feedback during Information-Integration Category Learning. *Journal of Neuroscience, 30*(1), 47-55.

de Bruijn, E. R. A., de Lange, F. P., von Cramon, D. Y., & Ullsperger, M. (2009). When Errors Are Rewarding. *Journal of Neuroscience, 29*(39), 12183-12186.

Delgado, M. R. (2007). Reward-related responses in the human striatum. In B. W. Balleine, K. Doya, J. Odoherty & M. Sakagami (Eds.), (Vol. 1104, pp. 70-88). Oxford: Blackwell Publishing.

Delgado, M. R., Locke, H., Stenger, V., & Fiez, J. (2003). Dorsal striatum responses to reward and punishment: Effects of valence and magnitude manipulations. *Cognitive, Affective, & Behaviroal Neuroscience, 3*, 27-39.

Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., & Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology, 84*(6), 3072-3077.

Draganski, B., Kherif, F., Kloeppel, S., Cook, P. A., Alexander, D. C., Parker, G. J. M., et al. (2008). Evidence for segregated and integrative connectivity patterns in the human basal ganglia. *Journal of Neuroscience, 28*(28), 7143-7152.

Elliott, R., Friston, K. J., & Dolan, R. J. (2000). Dissociable neural responses in human reward systems. *Journal of Neuroscience, 20*(16), 6159-6165.

Elliott, R., Newman, J. L., Longe, O. A., & Deakin, J. F. W. (2004). Instrumental responding for rewards is associated with enhanced neuronal response in subcortical reward systems. *Neuroimage, 21*(3), 984-990.

Fehr, E., Bernhard, H., & Rockenbach, B. (2008). Egalitarianism in young children. *Nature, 454*(7208), 1079-U1022.

Fehr, E., & Fischbacher, U. (2003). The nature of human altruism. *Nature, 425*(6960), 785.

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics, 114*(3), 817-868.

Fetchenhauer, D., & Huang, X. (2004). Justice sensitivity and distributive decisions in experimental games. *Personality and Individual Differences, 36*(5), 1015-1029.

Fliessbach, K., Weber, B., Trautner, P., Dohmen, T., Sunde, U., Elger, C. E., et al. (2007). Social comparison affects reward-related brain activity in the human ventral striatum. *Science, 318*(5854), 1305-1308.

Floresco, S. B., Onge, J. R. S., Ghods-Sharifi, S., & Winstanley, C. A. (2008). Cortico-limbic-striatal circuits subserving different forms of cost-benefit decision making. *Cognitive Affective & Behavioral Neuroscience, 8*(4), 375-389.

Floresco, S. B., Tse, M. T. L., & Ghods-Sharifi, S. (2008). Dopaminergic and glutamatergic regulation of effort- and delay-based decision making. *Neuropsychopharmacology, 33*(8), 1966-1979.

Foerde, K., & Shohamy, D. (2011). Feedback Timing Modulates Brain Systems for Learning in Humans. *Journal of Neuroscience, 31*(37), 13157-13167.

Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature, 40*(2), 351-401.

Gintis, H., & Fehr, E. (2012). The social structure of cooperation and punishment. *Behavioral and Brain Sciences, 35*(1).

Glockner, A., & Hilbig, B. E. (2012). Risk is relative: Risk aversion yields cooperation rather than defection in cooperation-friendly environments. *Psychonomic Bulletin & Review, 19*(3), 546-553.

Grahn, J. A., Parkinson, J. A., & Owen, A. M. (2008). The cognitive functions of the caudate nucleus. *Progress in Neurobiology, 86*(3), 141-155.

Graybiel, A. M. (2005). The basal ganglia: learning new tricks and loving it. *Current Opinion in Neurobiology, 15*(6), 638-644.

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

Gummerum, M., Hanoch, Y., & Keller, M. (2008). When child development meets economic game theory: An interdisciplinary approach to investigating social development. *Human Development, 51*(4), 235-261.

Haber, S. N., & Calzavara, R. (2009). The cortico-basal ganglia integrative network: The role of the thalamus. *Brain Research Bulletin, 78*(2-3), 69-74.

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.

Harbaugh, W. T., Mayr, U., & Burghart, D. R. (2007). Neural responses to taxation and voluntary giving reveal motives for charitable donations. *Science, 316*(5831), 1622-1625.

Hong, S., & Hikosaka, O. (2008). The Globus Pallidus Sends Reward-Related Signals to the Lateral Habenula. *Neuron, 60*(4), 720-729.

Hoshi, E. (2006). Functional specialization within the dorsolateral prefrontal cortex: A review of anatomical and physiological studies of non-human primates. *Neuroscience Research, 54*(2), 73-84.

Howard-Jones, P. A., Bogacz, R., Yoo, J. H., Leonards, U., & Demetriou, S. (2010). The neural mechanisms of learning from competitors. *Neuroimage, 53*(2), 790-799.

Ide, J. S., & Li, C.-S. R. (2011). Error-related functional connectivity of the habenula in humans. *Frontiers in Human Neuroscience, 5*.

Ino, T., Nakai, R., Azuma, T., Kimura, T., & Fukuyama, H. (2010). Differential activation of the striatum for decision making and outcomes in a monetary task with gain and loss. *Cortex, 46*(1), 2-14.

Ishiwari, K., Weber, S. M., Mingote, S., Correa, M., & Salamone, J. D. (2004). Accumbens dopamine and the regulation of effort in food-seeking behavior:

modulation of work output by different ratio or force requirements. *Behavioural Brain Research, 151*(1-2), 83-91.

Izuma, K., Saito, D. N., & Sadato, N. (2008). Processing of social and monetary rewards in the human striatum. *Neuron, 58*(2), 284-294.

Jansma, J. M., Ramsey, N. F., de Zwart, J. A., van Gelderen, P., & Duyn, J. H. (2007). fMRI study of effort and information processing in a working memory task. *Human Brain Mapping, 28*(5), 431-440.

Johansson, L. O., & Svedsater, H. (2009). Piece of cake? Allocating rewards to third parties when fairness is costly. *Organizational Behavior and Human Decision Processes, 109*(2), 107-119.

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

Klapproth, F. (2008). Time and decision making in humans. *Cognitive Affective & Behavioral Neuroscience, 8*(4), 509-524.

Kogut, T. (2012). Knowing what I should, doing what I want: From selfishness to inequity aversion in young children's sharing behavior. *Journal of Economic Psychology, 33*(1), 226-236.

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision Making and the Avoidance of Cognitive Demand. *Journal of Experimental Psychology-General, 139*(4), 665-682.

Krawczyk, D. C. (2002). Contributions of the prefrontal cortex to the neural basis of human decision making. *Neuroscience and Biobehavioral Reviews, 26*(6), 631-664.

Kurniawan, I. T., Seymour, B., Talmi, D., Yoshida, W., Chater, N., & Dolan, R. J. (2010). Choosing to Make an Effort: The Role of Striatum in Signaling Physical Effort of a Chosen Action. *Journal of Neurophysiology, 104*(1), 313-321.

Lee, D. (2008). Game theory and neural basis of social decision making. *Nature Neuroscience, 11*(4), 404-409.

Leotti, L. A., & Delgado, M. R. (2011). The Inherent Reward of Choice. *Psychological Science, 22*(10), 1310-1318.

Leotti, L. A., Iyengar, S. S., & Ochsner, K. N. (2010). Born to choose: the origins and value of the need for control. *Trends in Cognitive Sciences, 14*(10), 457-463.

Linke, J., Kirsch, P., King, A. V., Gass, A., Hennerici, M. G., Bongers, A., et al. (2010). Motivational orientation modulates the neural response to reward. *Neuroimage, 49*(3), 2618-2625.

Loewenstein, G. F., Bazerman, M. H., & Thompson, L. (1989). SOCIAL UTILITY AND DECISION-MAKING IN INTERPERSONAL CONTEXTS. *Journal of Personality and Social Psychology, 57*(3), 426-441.

Maddox, W. T., Ashby, F. G., & Bohil, C. J. (2003). Delayed feedback effects on rule-based and information-integration category learning. *Journal of Experimental Psychology-Learning Memory and Cognition, 29*(4), 650-662.

McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. *Proceedings of the National Academy of Sciences of the United States of America, 107*(17), 7922-7926.

Middleton, F. A., & Strick, P. L. (2000). Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Research Reviews, 31*(2-3), 236-250.

Mitchell, D. G. V. (2011). The nexus between decision making and emotion regulation: A review of convergent neurocognitive substrates. *Behavioural Brain Research, 217*(1), 215-231.

Mobbs, D., Yu, R. J., Meyer, M., Passamonti, L., Seymour, B., Calder, A. J., et al. (2009). A Key Role for Similarity in Vicarious Reward. *Science, 324*(5929), 900-900.

Nieuwenhuis, S., Heslenfeld, D. J., von Geusau, N. J. A., Mars, R. B., Holroyd, C. B., & Yeung, N. (2005). Activity in human reward-sensitive brain areas is strongly context dependent. *Neuroimage, 25*(4), 1302-1309.

O'Doherty, J. (2011). Contributions of the ventromedial prefrontal cortex to goal-directed action selection. In G. Schoenbaum, J. A. Gottfried, E. A. Murray & S. J. Ramus (Eds.), *Critical Contributions of the Orbitofrontal Cortex to Behavior* (Vol. 1239, pp. 118-129). Oxford: Blackwell Science Publ.

O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science, 304*(5669), 452-454.

Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature, 441*(7090), 223-226.

Pedroni, A., Koeneke, S., Velickaite, A., & Jancke, L. (2011). Differential magnitude coding of gains and omitted rewards in the ventral striatum. *Brain Research, 1411*, 76-86.

Pessiglione, M., Schmidt, L., Draganski, B., Kalisch, R., Lau, H., Dolan, R. J., et al. (2007). How the brain translates money into force: A neuroimaging study of subliminal motivation. *Science, 316*(5826), 904-906.

Phillips, P. E. M., Walton, M. E., & Jhou, T. C. (2007). Calculating utility: preclinical evidence for cost-benefit analysis by mesolimbic dopamine. *Psychopharmacology, 191*(3), 483-495.

Prevost, C., Pessiglione, M., Metereau, E., Clery-Melin, M. L., & Dreher, J. C. (2010). Separate Valuation Subsystems for Delay and Effort Decision Costs. *Journal of Neuroscience, 30*(42), 14080-14090.

Rademacher, L., Krach, S., Kohls, G., Irmak, A., Grunder, G., & Spreckelmeyer, K. N. (2010). Dissociation of neural networks for anticipation and consumption of monetary and social rewards. *Neuroimage, 49*(4), 3276-3285.

Rangel, A., & Hare, T. (2010). Neural computations associated with goal-directed choice. *Current Opinion in Neurobiology, 20*(2), 262-270.

Rilling, J. K., Gutman, D. A., Zeh, T. R., Pagnoni, G., Berns, G. S., & Kilts, C. D. (2002). A neural basis for social cooperation. *Neuron, 35*(2), 395-405.

Rossano, M. J. (2012). The Essential Role of Ritual in the Transmission and Reinforcement of Social Norms. *Psychological Bulletin, 138*(3), 529-549.

Rushworth, M. F. S., Noonan, M. P., Boorman, E. D., Walton, M. E., & Behrens, T. E. (2011). Frontal Cortex and Reward-Guided Learning and Decision-Making. *Neuron, 70*(6), 1054-1069.

Salamone, J. D., Correa, M., Mingote, S., & Weber, S. M. (2003). Nucleus accumbens dopamine and the regulation of effort in food-seeking behavior: Implications for studies of natural motivation, psychiatry, and drug abuse. *Journal of Pharmacology and Experimental Therapeutics, 305*(1), 1-8.

Satterthwaite, T. D., Ruparel, K., Loughead, J., Elliott, M. A., Gerraty, R. T., Calkins, M. E., et al. (2012). Being right is its own reward: Load and performance related ventral striatum activation to correct responses during a working memory task in youth. *Neuroimage, 61*(3), 723-729.

Schmidt, L., Lebreton, M., Clery-Melin, M. L., Daunizeau, J., & Pessiglione, M. (2012). Neural Mechanisms Underlying Motivation of Mental Versus Physical Effort. *Plos Biology, 10*(2).

Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron, 36*(2), 241-263.

Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned-stimuli during successive steps of learning a delayed-response task. . *Journal of Neuroscience, 13*(3), 900-913.

Seo, H., Barraclough, D. J., & Lee, D. (2007). Dynamic signals related to choices and outcomes in the dorslateral prefrontal cortex. *Cerebral Cortex, 17*, I110-I117.

Shohamy, D. (2011). Learning and motivation in the human striatum. *Current Opinion in Neurobiology, 21*(3), 408-414.

Shohamy, D., Myers, C. E., Grossman, S., Sage, J., Gluck, M. A., & Poldrack, R. A. (2004). Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. *Brain, 127*, 851-859.

Silk, J. B., Brosnan, S. F., Vonk, J., Henrich, J., Povinelli, D. J., Richardson, A. S., et al. (2005). Chimpanzees are indifferent to the welfare of unrelated group members. *Nature, 437*(7063), 1357-1359.

Singer, R. A., Berry, L. M., & Zentall, T. R. (2007). Preference for a stimulus that follows a relatively aversive event: Contrast or delay reduction? *Journal of the Experimental Analysis of Behavior, 87*(2), 275-285.

Stoppel, C. M., Boehler, C. N., Strumpf, H., Heinze, H. J., Hopf, J. M., & Schoenfeld, M. A. (2011). Neural processing of reward magnitude under varying attentional demands. *Brain Research, 1383*, 218-229.

Tabibnia, G., & Lieberman, M. D. (2007). Fairness and cooperation are rewarding - Evidence from social cognitive neuroscience. In *Social Cognitive Neuroscience of Organizations* (Vol. 1118, pp. 90-101). Oxford: Blackwell Publishing.

Takahashi, H., Kato, M., Matsuura, M., Mobbs, D., Suhara, T., & Okubo, Y. (2009). When Your Gain Is My Pain and Your Pain Is My Gain: Neural Correlates of Envy and Schadenfreude. *Science, 323*(5916), 937-939.

Tanaka, S. C., Doya, K., Okada, G., Ueda, K., Okamoto, Y., & Yamawaki, S. (2004). Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nature Neuroscience, 7*(8), 887-893.

Tanaka, S. C., Samejima, K., Okada, G., Ueda, K., Okamoto, Y., Yamawaki, S., et al. (2006). Brain mechanism of reward prediction under predictable and unpredictable environmental dynamics. *Neural Networks, 19*(8), 1233-1241.

Tindell, A. J., Berridge, K. C., & Aldridge, J. W. (2004). Ventral pallidal representation of pavlovian cues and reward: Population and rate codes. *Journal of Neuroscience, 24*(5), 1058-1069.

Treadway, M. T., Buckholtz, J. W., Cowan, R. L., Woodward, N. D., Li, R., Ansari, M. S., et al. (2012). Dopaminergic mechanisms of individual differences in human effort-based decision-making. *The Journal of neuroscience : the official journal of the Society for Neuroscience, 32*(18), 6170-6176.

Tricomi, E., Balleine, B. W., & O'Doherty, J. P. (2009). A specific role for posterior dorsolateral striatum in human habit learning. *European Journal of Neuroscience, 29*(11), 2225-2232.

Tricomi, E., Delgado, M. R., & Fiez, J. A. (2004). Modulation of caudate activity by action contingency. *Neuron, 41*(2), 281-292.

Tricomi, E., Delgado, M. R., McCandliss, B. D., McClelland, J. L., & Fiez, J. A. (2006). Performance feedback drives caudate activation in a phonological learning task. *Journal of Cognitive Neuroscience, 18*(6), 1029-1043.

Tricomi, E., & Fiez, J. A. (2008). Feedback signals in the caudate reflect goal achievement on a declarative memory task. *Neuroimage, 41*(3), 1154-1167.

Tricomi, E., & Fiez, J. A. (2012). Information content and reward processing in the human striatum during performance of a declarative memory task. *Cognitive Affective & Behavioral Neuroscience, 12*(2), 361-372.

Tricomi, E., Rangel, A., Camerer, C. F., & O'Doherty, J. P. (2010). Neural evidence for inequality-averse social preferences. *Nature, 463*(7284), 1089-U1109.

Voermans, N. C., Petersson, K. M., Daudey, L., Weber, B., van Spaendonck, K. P., Kremer, H. P. H., et al. (2004). Interaction between the human hippocampus and the caudate nucleus during route recognition. *Neuron, 43*(3), 427-435.

Wallis, J. D., & Miller, E. K. (2003). Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *European Journal of Neuroscience, 18*(7), 2069-2081.

Walton, M. E., Kennerley, S. W., Bannerman, D. M., Phillips, P. E. M., & Rushworth, M. F. S. (2006). Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks, 19*(8), 1302-1314.

Williams, Z. M., & Eskandar, E. N. (2006). Selective enhancement of associative learning by microstimulation of the anterior caudate. *Nature Neuroscience, 9*(4), 562-568.

Zentall, T. R., & Singer, R. A. (2007). Within-trial contrast: Pigeons prefer conditioned reinforcers that follow a relatively more rather than a less aversive event. *Journal of the Experimental Analysis of Behavior, 88*(1), 131-149.

Zink, C. F., Pagnoni, G., Martin-Skurski, M. E., Chappelow, J. C., & Berns, G. S. (2004). Human striatal responses to monetary reward depend on saliency. *Neuron, 42*(3), 509-517.

Table 1

Clusters of activity revealed by ANOVA with delay (immediate vs. delayed feedback) and valence (positive vs. negative feedback vs. no feedback Stage 1 and no feedback Stage 2) as within-subject factors ($p < 0.005$, contiguity threshold of 6 voxels (3x3x3 mm^3)).

Main Effect: Delay

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak F
Precuneus	2866	R	11	-56	60	28.7
Precuneus	511	L	-22	-68	51	22.5
Supplementary Motor Area	714	R	35	7	51	23.3
Precentral Gyrus	1049	R	26	-14	48	24.3
Supplementary Motor Area	956	L	-7	1	48	29.2
Medial Frontal Gyrus/Frontal Eye Fields	2617	R	11	25	48	32.2
Precentral Gyrus	1292	L	-37	13	39	31.6
Precentral Gyrus	761	R	53	-11	33	31.8
Dorsolateral Prefrontal Cortex	4912	R	5	34	30	47.6
Inferior Frontal Gyrus	2610	L	-46	1	24	47.4
Insula	331	R	29	16	9	32.7
Medial Frontal Gyrus	599	R	11	52	6	21.1
Anterior Cingulate	520	L	-19	43	3	27.2
Inferior Frontal Gyrus	2200	R	47	43	3	29.1
Lentiform Nucleus	3225	L	-13	4	3	82.8
Lentiform Nucleus	3677	R	14	4	-9	31.5

Main Effect: Valence

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak F
Supplementary Motor Area/Ventral Anterior Cingulate	26659	R	-7	4	51	33.6
Superior Parietal Lobule	3148	L	-34	-71	45	12.01
Dorsolateral Prefrontal Cortex	371	R	29	25	33	11.4
Caudate Body	455	R	17	-5	24	11.8
White Matter/Posterior Cingulate Cortex	4633	L	-25	-17	24	21.2

Middle Frontal Gyrus	599	R	29	40	21	12.2
Cuneus	347	R	2	-92	15	9.9
Posterior Cingulate Cortex	1042	R	20	-38	12	14.7
Superior Frontal Gyrus	338	L	-37	55	12	9.2
Caudate Head	341	R	17	22	9	11.6
Lentiform Nucleus	302	L	-10	-2	9	9.5
Insula/Infreior Frontal Gyrus/Supplementary Motor Area	14280	R	35	16	6	24.9
Lentiform Nucleus/Caudate	614	R	8	1	6	12.4
Thalamus	3085	L	-13	-17	6	23.4
Insula/Infreior Frontal Gyrus/Supplementary Motor Area/Supramarginal Gyrus	36775	L	-31	19	6	31.1
Thalamus	2767	R	8	-14	3	32.02
Middle Temporal Gyrus	737	L	-58	-38	0	14.5
Superior Frontal Gyrus	1051	R	26	52	-3	13.1
Cerebellum	1243	L	-37	-56	-27	15.7
Cerebellar Tonsil	227	R	38	-41	-33	10.6
<u>Interaction: Delay x Valence</u>						
Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak F
Precentral Gyrus	539	L	-46	-2	51	10.8
Inferior Parietal Lobule	920	L	-52	-47	42	11.2
Supramarginal Gyrus	1149	R	47	-50	33	12.3
Middle Temporal Gyrus	1290	L	-37	-65	30	14.7
Dorsolateral Prefrontal Cortex	473	L	-31	40	27	13.7
Declive of Cerebellum	1568	R	32	-71	-15	22.8
Parahippocampal Gyrus	257	R	14	-14	-27	17.05
Cerebellar Tonsil	362	L	-40	-53	-36	14.4

Table 2

Regions activated for the contrast of delayed feedback presentation versus immediate feedback presentation, with no feedback ($p < 0.005$, contiguity threshold of 6 voxels ($3x3x3\ mm^3$)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Lentiform Nucleus	324	R	11	7	6	5.3
Lentiform Nucleus	743	L	-16	1	6	5.5
Ventral Lentiform Nucleus	401	R	11	4	-6	5.5
Midbrain	607	L	-10	-11	-24	6.1

Table 3a

Regions activated for the contrast of immediate positive feedback presentation versus immediate negative feedback presentation ($p < 0.005$, contiguity threshold of 6 voxels ($3x3x3$ mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Immediate positive > Immediate negative						
Supramarginal Gyrus	505	R	50	-53	30	5.4
Superior Temporal Gyrus	580	L	-43	-59	27	4.9
Caudate Tail	535	R	17	-20	24	4.5
White matter/ Posterior Cingulate Cortex	1765	L	-28	-32	24	6.02
Cuneus	469	L	-19	-86	21	4.5
White matter/Dorsal Posterior Cingulate	539	L	-25	-50	21	4.8
Head of Caudate	459	L	-13	16	12	5.1
Head of Caudate	642	R	14	19	3	5.6
Inferior Temporal Gyrus	372	L	-43	-65	-3	4.8
Immediate negative > Immediate positive						
Postcentral Gyrus	3985	L	-46	-23	45	5.2
Dorsal Anterior Cingulate/Supplementary Motor Area	8433	R	8	22	30	7.3
Insula	1634	R	41	13	15	5.9
Thalamus	633	R	5	-17	6	5.5
Thalamus	519	L	-10	-14	6	4.7
Insula	564	L	-31	19	3	4.4

Table 3b

Regions activated for the contrast of immediate positive and negative feedback presentation versus no feedback presentation ($p < 0.005$, contiguity threshold of 6 voxels ($3x3x3\ mm^3$)).

Region	Cluser size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Immediate positive > No feedback 1						
None						
Immediate negative > No feedback 1						
Insula	1931	R	38	19	3	5.8
Supplementary Motor Area	4091	R	5	4	54	5.7
Thalamus	461	R	8	-14	3	5.1
Insula	932	L	-31	19	3	5
Precentral Gyrus	467	L	-49	1	30	4.9
No feedback 1 > Immediate negative						
Caudate Tail	453	R	20	-23	24	6.5
Somatosensory Cortex	760	L	-1	-26	69	6
White matter/ Posterior Cingulate Cortex	523	L	-25	-17	24	5.9
Cingulate Gyrus	317	R	26	-44	24	5.2
Cuneus	381	R	5	-92	33	4.6

Table 4a

Regions activated for the contrast of delayed negative feedback presentation versus delayed positive feedback presentation ($p < 0.005$, contiguity threshold of 6 voxels ($3x3x3\ mm^3$)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Delayed negative > delayed positive						
Supplementary Motor Area	4814	R	41	-2	45	5.9
Inferior Parietal Lobule	564	R	53	-56	45	4.7
Dorsolateral Prefrontal Cortex	988	R	32	28	33	5.5
Ventral Posterior Cingulate	376	L	-4	-17	30	5.6
Lentiform Nucleus	204	L	-13	-2	9	3.8
Insula	2528	R	32	19	6	5.2
Thalamus	763	R	5	-14	3	5.6
Thalamus	2863	L	-16	-17	3	6.5
Superior Temporal Gyrus	470	R	44	-38	0	5.3
Inferior Frontal Gyrus	10474	L	-31	19	-3	7.2
Middle Temporal Gyrus	1365	L	-67	-29	-6	6.6
Declive of Cerebellum	522	R	26	-77	-21	4.8
Delayed positive > delayed negative						
Posterior Cingulate Cortex	387	L	-31	-44	12	6.03

Table 4b

Regions activated for the contrast of delayed negative feedback presentation versus no feedback presentation ($p < 0.005$, contiguity threshold of 5 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Delayed negative > no feedback						
Supplementary Motor Area/Dorsal Anterior Cingulate Cortex	9991	L	-4	-2	51	6.8
Inferior Parietal Lobule	334	L	-55	-29	45	5.6
Precentral Gyrus	713	R	41	-8	42	5.1
Dorsolateral Prefrontal Cortex	413	R	41	28	36	4.4
Precentral Gyrus	2915	R	32	4	33	5.2
Angular Gyrus	896	L	-28	-62	30	4.2
Dorsolateral Prefrontal Cortex	596	L	-31	43	27	4.3
Inferior Frontal Gyrus	6502	L	-43	7	24	5.3
Lentiform Nucleus	352	L	-16	1	9	4.4
Thalamus	812	R	11	-14	6	5.3
Claustrum	2594	R	29	19	3	5.8
Middle Temporal Gyrus	741	R	44	-47	0	6
Superior Frontal Gyrus	909	R	29	55	-3	4.7
Middle Temporal Gyrus	473	L	-64	-35	-6	5.4
Thalamus	972	L	-13	-11	-9	5.8
Fusiform Gyrus	840	L	-43	-50	-12	5.1
Midbrain	356	R	8	-17	-15	5.2
No feedback > Delayed negative						
Posterior Cingulate Gyrus	2469	L	-22	-47	24	5.9

Table 4c

Regions activated for the contrast of delayed positive feedback presentation versus no feedback presentation ($p < 0.005$, contiguity threshold of 6 voxels ($3x3x3\ mm^3$)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
No feedback > delayed positive						
Cingulate Gyrus	326	L	-16	-38	24	4.2
Cuneus	494	R	8	-95	9	4.6
Delayed positive > no feedback						
None						

Table 5a

Regions activated during the ANOVA showing a difficulty x contingency x valence interaction ($p < 0.005$, contiguity threshold of 6 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	F
Middle Prefrontal Cortex (BA10)	285	R	14	49	12	12.8
Fusiform gyrus	738	R	41	-59	3	14.01

Table 5b

Regions activated during the ANOVA showing a difficulty x contingency interaction ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	F
Premotor cortex	242	L	-25	4	51	16.9
Supramarginal gyrus	326	L	-31	-32	42	23.2
Dorsolateral prefrontal cortex (BA9)	1242	L	-40	22	36	24.6
Anterior Cingulate cortex (BA32)	480	R	8	34	30	25.3
Anteriormedial Prefrontal cortex	467	L	-34	58	-6	22.7
Orbitofrontal Cortex (BA11)	148	L	-19	46	-15	18.3

Table 5c

Regions activated during the ANOVA showing a difficulty x valence interaction ($p < 0.001$, contiguity threshold of 3 voxels ($3x3x3\ mm^3$)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	F
Superior Temporal gyrus	245	R	65	4	-3	22.7

Table 5d

Regions activated during the ANOVA showing a contingency x valence interaction ($p < 0.001$, contiguity threshold of 3 voxels ($3x3x3\ mm^3$)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	F
Dorsal Posterior Cingulate cortex	171	R	5	-23	48	17.9
Dorsal Posterior Cingulate cortex	165	L	-13	-29	45	17.02
Supplementary motor area	403	L	-18	28	36	21.9
Ventral Anterior Cingulate cortex	218	L	-13	7	33	19.6
Ventral Anterior Cingulate cortex	238	R	11	-5	30	26.7
Dorsal Putamen	226	L	-23	7	12	21.5
Posterior Putamen	114	L	-22	-2	-6	20.5

Table 6a

Regions activated during the ANOVA showing a main effect of difficulty ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	F
Superior Parietal Lobule	523	L	-10	-62	54	25.2
Cuneus	280	L	-7	-86	39	19.8
Cingulate cortex	147	L	-13	-41	0	20.3
Cerebellum	110	L	-1	-77	-15	18.2

Table 6b

Regions activated during the ANOVA showing a main effect of contingency ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak F
Supramarginal gyrus	1623	R	44	-38	39	21.9
Frontal eye field region	194	R	47	13	39	18.1
Dorsolateral Anterior Cingulate cortex	3417	R	11	22	36	27. 6
Supplementary motor area/Frontal eye fields/Insula/Caudate nucleus/Putamen/Thalamus	9286	R	26	13	33	33.1
Supramarginal gyrus	289	L	-31	-38	33	25.2
Precuneus	588	L	-22	-59	30	24.1
Dorsolateral Prefrontal cortex (BA9)	554	L	-43	16	30	19.7
Ventral Posterior Cingulate cortex	1406	R	5	-17	27	27.01
Ventral Anterior Cingulate cortex	117	R	2	7	24	21.7
Caudate nucleus/Thalamus	2475	L	-19	-8	18	26.1
Globus Pallidus	122	L	-25	-17	-3	21.3
Inferior Prefrontal gyrus	371	L	-31	22	-3	18.5
Posterior lobe of Cerebellum	6056	L	-37	-65	-27	24. 3
Anterior lobe of Cerebellum	7958	R	38	-41	-30	46.9
Cerebellar tonsil	156	L	-31	-32	-36	25.2

Table 6c

Regions activated during the ANOVA showing a main effect of valence ($p < 0.001$, contiguity threshold of 4 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak F
Paracentral lobule	711	R	5	-29	63	20.1
Supplementary motor area	1045	R	11	-17	48	27.6
Superior Frontal gyrus	2490	L	-22	31	45	24.7
Cuneus	1150	R	14	-71	33	22.01
Premotor cortex	477	L	-34	-8	33	19.8
Angular gyrus	699	R	50	-68	30	19.9
Ventral Anterior Cingulate cortex	7181	R	2	-5	30	30.2
Dorsal Anterior Cingulate cortex	169	R	17	28	27	22.04
Insula	1027	R	38	13	21	22.1
Dorsal Posterior Cingulate cortex	440	L	-31	-68	21	17. 8
Superior Temporal gyrus	624	L	-37	-50	21	21.9
Insula	412	L	-37	25	15	20.7
Dorsolateral Prefrontal cortex (BA46)	844	R	47	28	12	28.9
Posterior Superior Temporal gyrus	637	R	65	-35	9	16.7
Middle Frontal cortex	162	R	38	49	9	17.3
Anterior Superior Temporal gyrus	144	R	62	-23	0	16.3
Superior Temporal gyrus	615	R	62	-14	0	18.5
Ventromedial Prefrontal Cortex (BA 10, 11), rostral Anterior Cingulate Cortex	22763	L	-1	40	-3	27.1
Inferior Prefrontal gyrus	1187	L	-37	37	-6	24.7
Ventral Striatum/Caudate nucleus/Lentiform nucleus	5850	R	14	4	-6	45. 6
Lentiform nucleus/Ventral Striatum/Caudate nucleus	4903	L	19	4	-6	44.5
Middle Temporal gyrus	290	L	-55	-17	-9	18.7
Fusiform gyrus	8181	L	-43	-53	-9	26.2
Hippocampus	1005	L	-31	-14	-9	21.1
Inferior Temporal gyrus	1065	L	-43	-41	-21	30.1
Anterior lobe of Cerebellum	1721	R	23	-29	-24	25.4
Parahippocampal gyrus	342	R	26	-8	-33	23.6

Cerebellar Tonsil	1450	L	-22	-35	-39	34.8
Posterior portion of Posterior Lobe of Cerebellum	671	R	17	-68	-42	20.6
Posterior portion of Posterior Lobe of Cerebellum	4149	L	-16	-71	-42	26.8
Cerebellar Tonsil	350	R	14	-41	-43	23.8

Table 7a

Regions activated for the contrast of 1-step learning positive versus 1-step learning negative feedback presentation ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Positive > Negative						
Premotor cortex	328	L	-7	-23	60	4.7
Superior Frontal gyrus	178	L	-25	25	48	4.3
Superior Temporal gyrus	913	R	62	-50	15	5.3
Putamen	289	L	-22	7	15	4.5
Middle Temporal gyrus	263	R	50	-44	6	4.3
Fusiform Gyrus	704	L	-43	-62	0	4.8
Ventral Striatum	1001	R	14	7	-3	4.8
Ventral Striatum	383	L	-22	7	-3	4.3
Parahippocampal Gyrus/Amygdala	873	L	-34	-17	-9	4.9
Negative > Positive						
None	-	-	-	-	-	-

Table 7b

Regions activated for the contrast of 1-step random positive versus 1-step random negative feedback presentation ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Positive > Negative						
Dorsolateral Prefrontal cortex (BA46)	513	R	47	25	15	6.1
Ventral striatum	103	L	-13	4	-3	4.3
Cerebellum	1574	L	-16	-77	-36	4.98
Negative > Positive						
None	-	-	-	-	-	-

Table 7c

Regions activated for the contrast of 1-step learning positive versus 1-step random positive feedback presentation and 1-step learning negative versus 1-step random negative feedback presentation ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Learning Positive > Random Positive						
Thalamus	550	L	-1	-11	18	4.9
Learning Negative > Random Negative						
Cerebellum	149	L	-37	-62	-43	4.4

Table 8a

Regions activated for the contrast of 2-step learning positive versus 2-step learning negative feedback presentation ($p < 0.001$, contiguity threshold of 4 voxels (3x3x3 mm^3)).

Region	Cluster Size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Positive > Negative						
MFG/ACC/SMA						
Supplementary Motor Area	7345	R	8	-20	48	7.5
Dorsal Posterior Anterior Cingulate cortex	4044	L	-7	-39	39	6.4
Dorsolateral Prefrontal cortex (BA9)	5783	L	-19	40	36	7.7
Ventral Anterior Cingulate cortex	5654	R	2	1	30	7.5
Medial Frontal gyrus (BA10)	4852	R	14	55	15	7.7
Dorsal Anterior Cingulate cortex	5051	L	-10	34	12	7.9
Medial Frontal gyrus (BA 6)	136	R	38	-8	45	4.3
Superior Occipital gyrus	3812	L	-34	-74	27	5.8
Supramarginal gyrus	212	L	-52	-50	18	5.6
Insula	2813	R	28	16	18	6.3
Insula	2812	L	-34	-2	15	5.2
Superior Temporal gyrus	193	L	-43	-32	9	4.4
Fusiform gyrus	294	L	-43	-62	-6	4.3
Middle Temporal gyrus	240	L	-55	-14	-9	4.5
Inferior Frontal gyrus	1248	R	41	31	-18	5.4
Inferior Temporal gyrus	316	L	-40	-5	-21	4.8
Cerebellum	1036	L	-7	-62	-24	4.8
Parahippocampal gyrus	387	L	-28	7	-30	4.5
Cerebellum	406	L	-7	-50	-36	4.7
Basal ganglia						
Dorsal Striatum	1166	L	-7	13	18	4.7
Dorsal Striatum	3335	R	14	10	15	4.9
Lentiform nucleus	3922	l	-25	13	1	5.6
Lentiform nucleus	5177	R	14	-1	-3	6.1
Ventral Striatum	1265	R	8	10	-3	6.2
Ventral Striatum	1161	L	-17	2	-6	6.8
Negative > Positive						

None	-	-	-	-	-	-

Table 8b

Regions activated for the contrast of 2-step random positive versus 2-step random negative feedback presentation ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Positive > Negative						
Superior Frontal gyrus	134	L	-19	46	42	4.2
Posterior Anterior Cingulate cortex	492	R	8	-32	30	4.8
Ventral Posterior Anterior Cingulate cortex	124	L	-1	-23	24	4.2
Superior Temporal gyrus	182	R	59	-41	9	4.3
Ventral Striatum	86	R	11	1	0	4.2
Cerebellum	110	R	23	-32	-27	4.1
Parahippocampal gyrus	148	R	23	-11	-30	6.03
Negative > Positive						
None	-	-	-	-	-	-

Table 8c

Regions activated for the contrast of 2-step learning positive versus 2-step random positive feedback presentation and 2-step learning negative versus 2-step random negative feedback presentation ($p < 0.001$, contiguity threshold of 4 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Learning Positive > Random Positive						
Medial Frontal gyrus (BA8)	388	R	47	13	42	4.5
Inferior Parietal gyrus	133	L	-31	-38	39	5.3
Medial Frontal gyrus (BA8)	570	R	26	16	36	4.5
Angular gyrus	1408	L	-37	-59	30	5.5
Angular gyrus	359	R	35	-62	30	5.1
Medial Frontal gyrus (BA10)	115	R	26	55	24	4.4
Ventral Posterior Anterior Cingulate cortex	1448	R	8	-17	24	7.4
Dorsal Striatum	3279	R	11	16	17	5.9
Thalamus	1999	L	-18	-2	15	5.6
Medial Frontal gyrus (BA10)	495	L	-28	43	12	4.3
Thalamus	2836	R	14	-11	9	5.2
Dorsal Striatum	2615	L	-22	13	6	6.1
Inferior Frontal gyrus	1799	L	-28	28	3	5.4
Putamen/Lentiform	156	L	-28	-14	-3	4.5
Ventral Striatum	1125	R	11	4	-3	4.9
Posterior portion of Posterior Lobe of Cerebellum	157	R	14	-74	-24	4.1
Posterior portion of Posterior Lobe of Cerebellum	1039	L	-19	-65	-27	5.4
Anterior Portion of Posterior Lobe of Cerebellum	489	R	32	-68	-30	4.2
Cerebellar Tonsil	264	L	-1	-50	-33	4.6
Learning Negative > Random Negative						
Inferior Parietal gyrus	257	L	-34	-50	45	4.1
Anterior Cingulate cortex	328	R	8	28	30	4.6

Thalamus	426	R	17	-11	15	4.9
Anterior Insula	148	R	29	22	0	4.4
Superior Frontal Gyrus	121	L	-37	61	-6	4.6
Middle portion of Posterior Lobe of Cerebellum	643	R	5	-62	-12	4.6
Posterior portion of Posterior Lobe of Cerebellum	226	R	29	-65	-21	4.3
Anterior Lobe of Cerebellum	1087	L	-34	-59	-24	4.6
Anterior portion of Anterior Lobe of Cerebellum	421	R	32	-41	-30	4.8
Cerebellar Tonsil	296	L	-25	-65	-33	4.2

Table 9

Regions activated for the contrast of 2-step learning positive versus 1-step learning positive feedback presentation ($p < 0.001$, contiguity threshold of 3 voxels (3x3x3 mm^3)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
2-step learning positive > 1-step learning positive						
Dorsolateral Prefrontal Cortex (BA9)	114	L	-28	22	30	5.6
Insula	122	R	44	-17	21	4.7
Medial Frontal Gyrus	2330	R	14	43	12	7.6
Anterior Cingulate cortex	190	R	2	31	3	4.3
Medial Prefrontal Cortex (BA10)	2532	L	-31	52	0	4.2
Ventral Striatum	289	R	11	-2	-3	5.7
Inferior Frontal gyrus	486	L	-34	16	-3	5.03
Orbitofrontal cortex	731	L	-22	49	-12	4.8
2-step learning negative > 1-step learning negative						
None	-	-	-	-	-	-

Table 10

Regions activated for the contrast of 2-step learning versus 1-step learning feedback presentation ($p < 0.001$, contiguity threshold of 3 voxels ($3x3x3\ mm^3$)).

Region	Cluster size (mm^3)	Hemisphere	Peak X	Peak Y	Peak Z	Peak t
Supplementary Motor area	274	L	-28	4	51	4.3
Inferior Parietal lobe	404	L	-31	-32	42	4.8
Anterior Cingulate cortex	484	R	8	34	30	5.03
Inferior Frontal gyrus	186	L	-49	7	30	4.4
Dorsolateral Prefrontal cortex (BA9)	1840	L	-28	16	27	5. 6
Medial Frontal cortex	431	L	-34	58	-6	4.8
Orbitofrontal cortex (BA11)	145	L	-19	46	-15	4.3

Appendix A

Additional ROI analyses were conducted on data from experiment 1 and 2 in the caudate nucleus and nucleus accumbens (dorsal (DS) and ventral striatum (VS)) a priori coordinates extracted from an 8mm sphere. The a priori coordinates for this analysis were taken from Zink et al. 2003 that reflect the pattern of activity of the dorsal and ventral striatum during the two tasks.

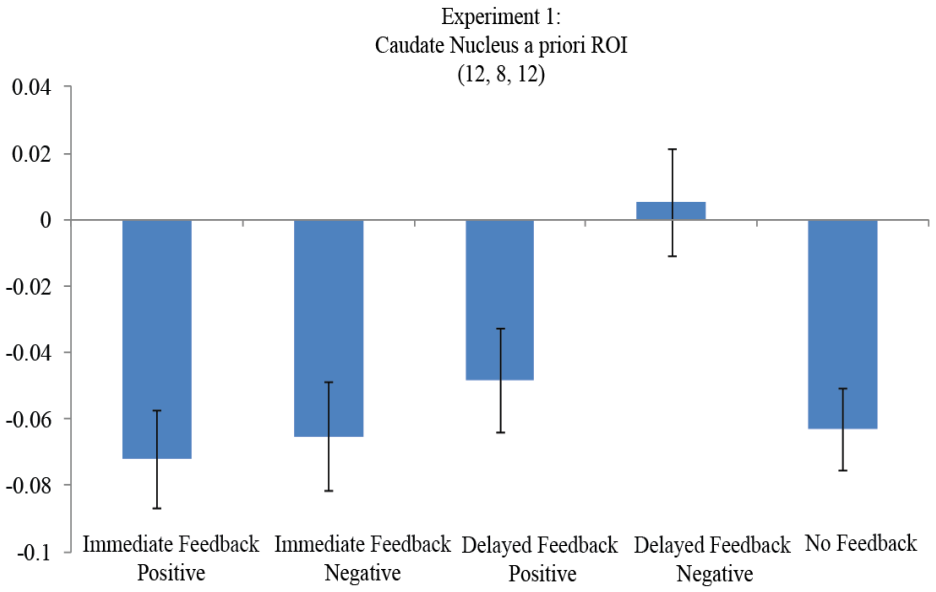

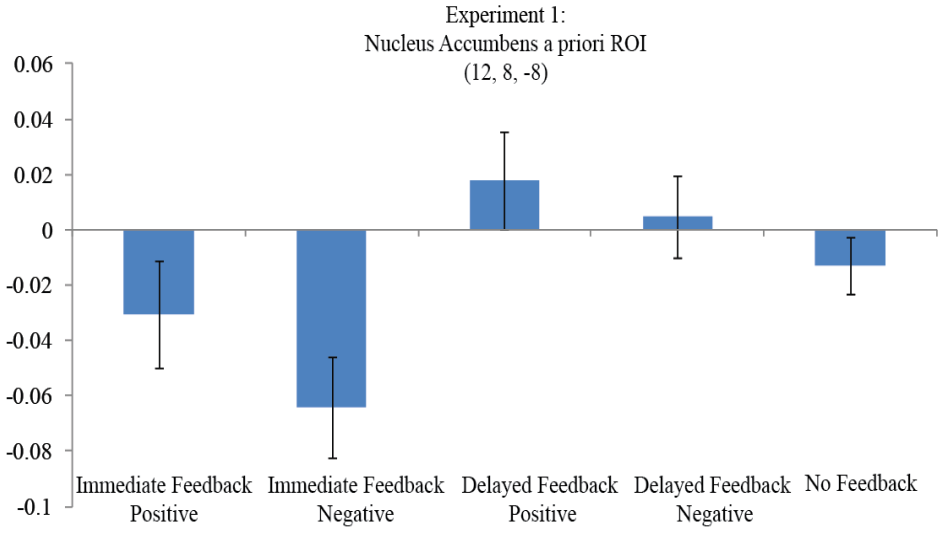

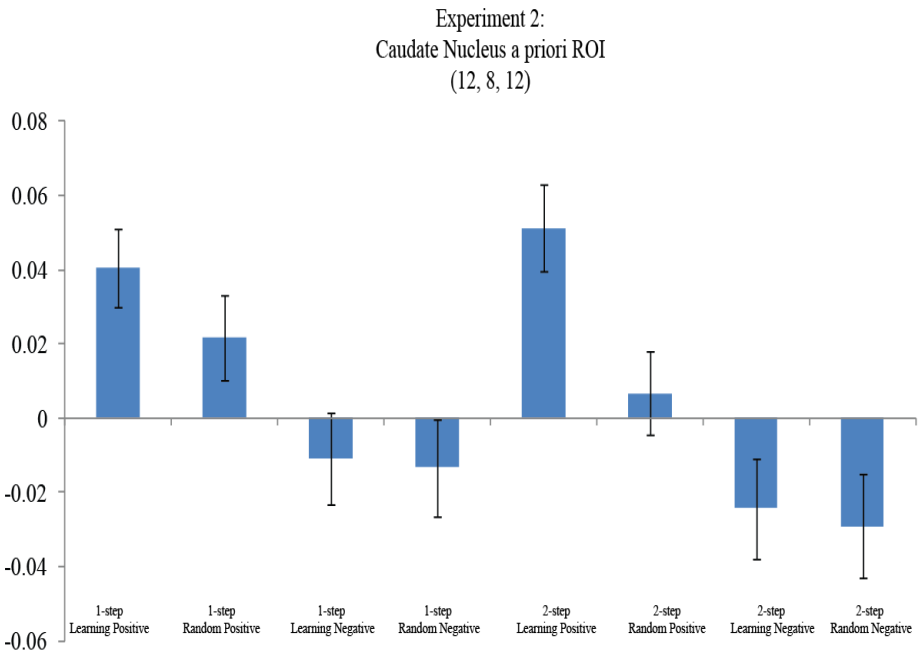

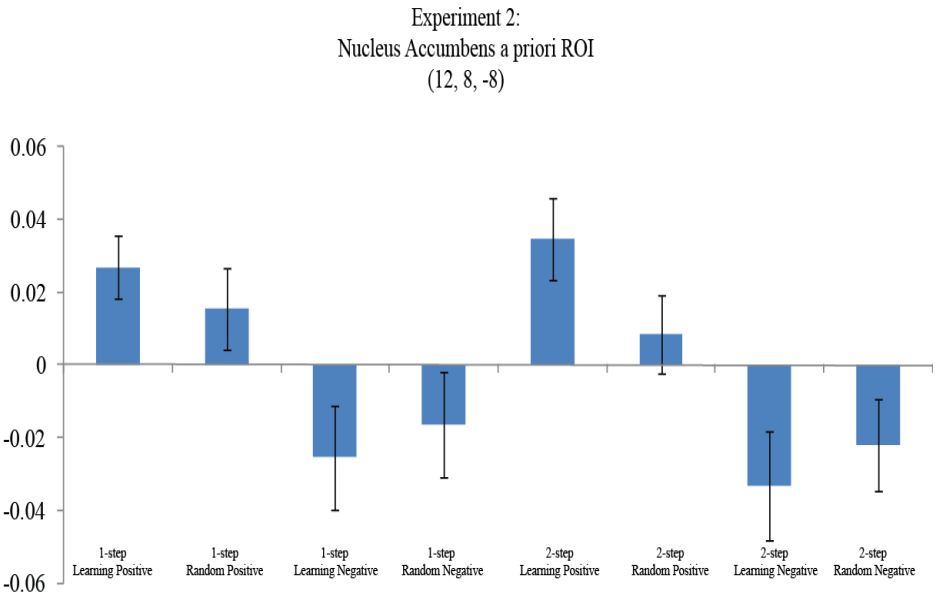

An abundance of neurophysiological and neuroimaging evidence suggests that specific parts of the striatum are involved in processing different types of information. For example, the VS is implicated in processing reward anticipation and the value of outcomes during instrumental, Pavlovian and observational

learning (Cooper, Dunne, Furey, & O'Doherty; Haruno & Kawato, 2006; O'Doherty et al., 2004; Tricomi, Delgado, & Fiez, 2004). This evidence highlights the VS involvement in the limbic corticostriatal loop (Seger, 2008). On the other hand, the DS is involved more in instrumental learning, with a decrease in DS engagement as learning progresses and performance improvement can be seen in participants' behavior (Haruno et al., 2004; Kahnt, Heinzle, Park, & Haynes). This pattern of involvement in learning highlights the DS role in the executive corticostriatal loop dedicated to high-order cognitive processes (Seger, 2008).

In the current work, the VS ROI (nucleus accumbens) shows a trend towards differentiation between positive and negative feedback in both experiments across all of the conditions. That is, the negative feedback presentation seems to always result in less activity in VS than positive feedback presentation, regardless of the cost of feedback (delay versus effort). This is in accord with the role that is attributed to the VS of a region involved in processing of different outcome types in a similar manner. However, delayed outcomes are thought to be subjectively less valuable to individuals, resulting in a weaker response in the VS (Kable & Glimcher, 2007). Such pattern of VS activity can also be observed when looking at the VS activity associated with delayed feedback presentation. There is only slight activation from baseline associated with delayed feedback (Experiment 1). The same idea can be applied to random versus action-contingent feedback presentation (Experiment 2). A less robust activation of the VS is present when feedback is presented during 1-step and 2-step random conditions. This is probably a result of participants valuing the

feedback that is contingent on their action more, and reflects their learning success (1-step and 2-step learning conditions).

The DS ROI (caudate nucleus) in experiment 2 shows a similar pattern of activity as the VS ROI. The DS ROI in experiment 1 shows a more robust activation from baseline with only delayed negative feedback resulting in a very slight increase from baseline. It seems that the DS treats immediate feedback and positive feedback after a delay similarly.

Although the above results from a priori ROI regions do not display the same results across both experiments, the main findings are in line with the current literature and still highlight the multifaceted role of the striatal components in learning of action-outcome contingencies. The DS seems to be more involved in the learning process per se and is more sensitive to contextual influences, while the VS might respond more to the valence of outcomes and is not as responsive to the context of the learning environment.

However, the striatum is not the only brain region involved in learning of action-outcome associations. The main body of the dissertation document mentions prefrontal cortex regions such as the dorsolateral and ventromedial prefrontal cortex (DLPFC and VMPFC, respectively) and the anterior cingulate cortex. These cortical regions together with the striatum represent distinct corticostriatal loops. Even thought this dissertation work primarily focuses on the role and functioning of the striatum in learning though performance-related feedback, it provides a good depiction of the activity of corticostriatal loops.

Several animal and human studies looked at the interplay of striatal and PFC regions. It is not possible to track the flow of information through corticostriatal loops via the analysis conducted in the current work, but current studies can still shed further light on the interplay between prefrontal and striatal components. The striatum was shown to be involved at the beginning of the learning process, while PFC regions take over as the learning progresses (cue-outcome learning) (Kahnt, Heinzle, Park, & Haynes, 2011; Pasupathy & Miller, 2005). At the same time, Ballard and others (2011) showed that the DLPFC provides information to the striatum during motivated learning (action-outcome learning). Together with the DS, the DLPFC comprises the executive corticostriatal loop (Lopez-Paniagua & Seger, 2011; Seger, 2008). In the context of experiment 2, the DLPFC might have gathered information about the action-outcome pairing and influenced the DS in the selection of the correct response during the more difficult and motivating condition (2-step learning).

The VS and the VMPFC comprise the motivational loop and also participate in action selection (Lopez-Paniagua & Seger, 2011; O'Doherty, 2011), but were reported to be involved more in processing of affective information. In the current work, the most prominent activity of the VMPFC was observed during experiment 2. The VMPFC, that was shown to be activated in response to not only positive outcomes but also to negative outcomes, was also shown to reflect the action value of selected actions (Rushworth, Noonan, Boorman, Walton, & Behrens, 2011). In experiment 2 of the current work, the VMPFC showed a robust activity during the positive feedback of the 2-step learning condition. Such

activity might reflect greater valuation of chosen responses that result in positive feedback during the 2-step learning condition.

Another important component of corticostriatal circuitry is the anterior cingulate cortex (ACC). The ACC is a complex region, implicated in error monitoring and in processing emotional and rewarding information (Mohanty et al., 2007; Rushworth et al., 2011). The dorsal parts of the ACC are connected with the DS that is part of the executive loop. The anterior parts of the ACC are part of the motivational loop (Lopez-Paniagua & Seger, 2011; Seger, 2008). Based on the results of the current studies, where ACC-striatum network was consistently activated, the ACC might have been responsible for supplying error information to the striatum and posterior parts of the basal ganglia (as has previously been shown before; e.g. Daniel & Pollman, 2010), but with greater involvement when negative feedback is provided after a delay (Experiment 1).

Appendix B

The 2-step learning of experiment 2 of this dissertation is proposed to require increased cognitive processing in comparison to the 1-step learning condition due to a greater amount of information that has to be processed. Increased cognitive demands, such as working memory (WM) load, planning and cognitive control, are associated with increased activity in regions of the prefrontal cortex (PFC) as observed with functional magnetic resonance imaging (Krawczyk, 2002). In addition, positron emission tomography (PET) studies also point to the involvement of the PFC during increased information processing. Specifically, receptor-specific radioactive tracer absorption increases during tasks that require increased information processing. There is a large concentration of mesolimbic dopamine neurons in the PFC. In particular, there is an abundance of D1 dopamine receptors, the absorption of which can be observed during the PET scans that aim at looking at cognitively demanding tasks (Cools & D'Esposito, 2011; Takahashi, Yamada, & Suhara, 2012).

PFC levels of dopamine neurotransmitter (DA) lie at the core of several neuropsychological disorders such as Parkinson's disease (PD), schizophrenia, attention deficit hyperactivity disorder, and the like. There are several dopamine-based medications that affect dopamine levels in the brain and aim to alleviate certain cognitive deficits of individuals. Therefore, there have been a number of studies performed that looked at the effects of DA medication on such processes as WM and cognitive control in healthy participants and also in populations with PD or traumatic brain injury, where DA deficiency comes as a result of axonal

damage after a brain injury (Cools & D'Esposito, 2011; McDowell, Whyte, & D'Esposito, 1998). However, the results of these studies are mixed, leading to complex conclusions about the effects of the DA medication on the brain. For example, studies with healthy participants who were given bromocriptine, a DA agonist, showed an increase in WM capacity. This effect is also driven by baseline WM abilities. That is, when one has low WM capacity, bromocriptine has a positive effect on WM performance. With high baseline WM capacity, bromocriptine hinders task performance. Baseline WM capacity is dependent on baseline D1 receptor levels, with increased concentration of D1 receptors associated with better WM capacity (Cools & D'Esposito, 2011; Takahashi et al., 2012).

Similarly, patients with frontal lesions after bromocriptine administration exhibit a mix of improvements on executive measures. For example, after 2.5 mg of bromocriptine, improvements were observed in frontal lesion patients on such tasks as the Wisconsin Card Sorting task and the Stroop task but not on the WM task. PD patients have a similar profile to frontal lesion patients on such executive processes tasks (Cools & D'Esposito, 2011; McDowell et al., 1998).

In context of current study, results similar to those obtained on the 1-step learning condition can be expected with healthy participants on bromocriptine. However, it is hard to predict how the results of the 2-step learning condition can be affected. It is possible that it would be easier to make a prediction knowing the baseline working memory capacity of participants. Knowing the baseline working memory capacity would allow seeing whether there is a relationship between

one's performance on the 2-step learning task and working memory capacity. And based on the direction of this relationship, it might be possible to predict effects of bromocriptine. It would be interesting to observe performance of patients with frontal lesions before and after bromocriptine administration to see whether it improves their performance on the 2-step learning condition. Experiment 1 seem to rely on processes other than WM, hence bromocriptine might not have any positive or negative influences on participants' performance.

EKATERINA DOBRYAKOVA

ekaterina@psychology.rutgers.edu

Curriculum Vitae

EDUCATION

2010 – 2012 Rutgers University
PhD, Psychology
Dissertation: *An Investigation of Basal Ganglia Activity during Delayed and Effort-based Learning.* Advisor: Dr. Elizabeth Tricomi

2007 – 2010 Rutgers University
M.A., Psychology
Advisor: Dr. Elizabeth Tricomi

2004 – 2007 Rutgers University
B.A., Psychology, High Honors
Honor's thesis title: *Confusing intention and action: How do we know what we did and did not do?* Thesis committee members: Guenther Knoblich, Natalie Sebanz.

PUBLICATIONS

Dobryakova, E. and Tricomi, E. (submitted). Basal ganglia engagement during feedback processing after a substantial delay.

CONFERENCE PRESENTATIONS

Dobryakova, E. and Tricomi, E. (2012). Subjective cognitive effort modulates activity of the ventral striatum. *Poster to be presented at the 5th Social and Affective Neuroscience Society.* New York, NY.

Dobryakova, E. and Tricomi, E. (2012). Modulation of ventral striatal activity by subjective cognitive effort. *Poster to be presented at the 19th Cognitive Neuroscience Society annual meeting*. Chicago, IL.

Dobryakova, E. and Tricomi, E. (2010). Striatal engagement during feedback processing after a substantial delay. *Poster presented at the Society for Neuroscience annual meeting (2010).* San Diego, CA: Society for Neuroscience.

Dobryakova, E. and Shiffrar, M. (2008). Is that you? Observers' own actions impact their visual perception of human action. *Poster presented at the 16th Annual OPAM, Chicago, IL.*

Dobryakova, E. and Knoblich, G. (2007). Confusing past intentions with past actions. *Poster presented at the 48th Annual Meeting of the Psychonomic Society, Long Beach, CA.*

PROFESSIONAL MEMBERSHIPS

Cognitive Neuroscience Society (2011-present)
Society for Neuroscience (2010-present)
Social and Affective Neuroscience Society (2009-present)