IMAGING THE ROLE OF THE CAUDATE NUCLEUS IN FEEDBACK PROCESSING DURING A DECLARATIVE MEMORY TASK

by

Elizabeth Tricomi

BS, Cornell University, 2000

MS, University of Pittsburgh, 2004

Submitted to the Graduate Faculty of

Arts and Sciences in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2006

UNIVERSITY OF PITTSBURGH

ARTS AND SCIENCES

This dissertation was presented

by

Elizabeth Tricomi

It was defended on

September 6, 2006

and approved by

Charles Perfetti, Professor, Psychology

Erik Reichle, Associate Professor, Psychology

Mark Wheeler, Assistant Professor, Psychology

James McClelland, Professor, Psychology, Stanford University

Dissertation Advisor: Julie Fiez, Professor, Psychology

IMAGING THE ROLE OF THE CAUDATE NUCLEUS IN FEEDBACK PROCESSING DURING A DECLARATIVE MEMORY TASK

Elizabeth Tricomi, M.S.

University of Pittsburgh, 2006

Although the caudate nuclei are involved in processing performance feedback, the specific conditions under which feedback engages the caudate are not well understood. It is unclear whether feedback in a declarative memory task can activate the caudate to the same extent as a non-declarative task. Likewise, it is not known how the type and amount of information carried by feedback affect caudate activation. To answer these questions, I examined brain activation during a feedback-based paired associate word learning task over three rounds of trials with two response options (Experiment 1) and over one round of trials with either two or four response options (Experiment 2). The caudate nuclei were strongly engaged in Experiment 1 only during the second two rounds, when feedback reflected the accuracy of memory. In Experiment 2, differential responses to positive and negative feedback were observed in the caudate nuclei in the 4-choice condition, for which positive feedback provides more information than negative feedback. Responses to positive and negative feedback were not differential in the 2-choice condition and were similar in magnitude to the response elicited by positive feedback in the 4-choice condition. These results indicate that the caudate can be involved in feedback processing during tasks engaging declarative memory, but that activation during such tasks is context dependent. The caudate is most strongly engaged when positive and negative feedback have differential value to the individual. Though not required for learning to occur, caudate activation may nevertheless contribute to feedback-based declarative memory formation.

TABLE OF CONTENTS

PREFACE X

1.0 INTRODUCTION 1

1.1 The Role of the Striatum in Feedback Processing 2

1.2 The Role of the Striatum in Nondeclarative Learning 5

1.3 Striatal Influences on the Declarative Memory System 7

1.4 Modulation of Caudate Activation During Feedback-Based Declarative Learning 11

2.0 EXPERIMENT 1: THE EFFECT OF THE TYPE OF INFORMATION PROVIDED BY FEEDBACK ON CAUDATE ACTIVATION 15

2.1 Methods 17

2.1.1 Participants 17

2.1.2 Materials 18

2.1.3 Procedure 18

2.1.4 Data Analysis 20

2.2 Results 22

2.2.1 Behavioral Results 22

2.2.2 fMRI Results 23

2.3 Discussion 27

3.0 EXPERIMENT 2: THE EFFECT OF THE INFORMATIONAL VALUE OF FEEDBACK ON CAUDATE ACTIVATION 34

3.1 Methods 36

3.1.1 Participants 36

3.1.2 Materials 37

3.1.3 Data Analysis 39

3.2 Results 41
3.2.1 Behavioral Results 41
3.2.2 fMRI Results 43
3.3 Discussion 48
4.0 GENERAL DISCUSSION 54
4.1 The Nature of Feedback-Related Caudate Activation 54
4.1.1 The Rise 55
4.1.2 The "Reward" Response 59
4.1.3 The "Punishment" Response 60
4.2 Frontal Contributions to Feedback Processing 62
4.2.1 Lateral Prefrontal Cortex 63
4.2.2 Anterior Cingulate Cortex 65
4.2.3 Ventromedial Prefrontal Cortex 66
4.3 Learning in a Feedback-Based Declarative memory Task 67
4.3.1 Medial Temporal Lobe Contributions 67
4.3.2 Striatal Contributions 68
4.4 Trial Summary 72
4.5 Conclusions 73
BIBLIOGRAPHY 107

LIST OF TABLES

Table 1: Experiment 1 Accuracy and Reaction Time Across Experimental Rounds 75
Table 2: Experiment 1 Reaction Time for Correct Versus Incorrect Trials 75
Table 3: Experiment 1 Confidence for Correct Versus Incorrect Trials 75
Table 4: Experiment 1 Round X Time ($p<0.0001$) 76
Table 5: Experiment 1 Accuracy X Time ($p<0.0001$) 77
Table 6: Experiment 1 Round X Accuracy X Time ($p<0.0001$) 78
Table 7: Experiment 2 Trial Type X Time ($p<0.001$) 79
Table 8: Experiment 2 Feedback Type X Time ($p<0.001$) 80
Table 9: Experiment 2 Trial Type X Feedback Type X Time ($p<0.001$) 81
Table 10: Experiment 2 4-choice Condition: Feedback Type X Time ($p<0.001$) 82
Table 11: Experiment 2 2-choice Condition: Feedback Type X Time ($p<0.001$) 83
Table 12: Experiment 2 Subsequent Memory Effects 84

LIST OF FIGURES

Figure 1. Experimental Design. 85
Figure 2. Brain Regions Showing a Round by Time Interaction in Experiment 1. 86
Figure 3. Time Course of Activation in the Caudate Region Displaying a Round by Time Effect. 87
Figure 4. Brain Regions Showing an Accuracy by Time Effect in Experiment 1. 88
Figure 5. Time Courses of Activation in the Caudate Region Displaying an Accuracy by Time Effect. 89
Figure 6. Brain Regions Showing a Round by Accuracy by Time Interaction in Experiment 1. 90
Figure 7. Time Courses of Activation in the Caudate Body Region Showing a Round by Accuracy by Time Interaction. 91
Figure 8. Time Courses of Activation in the Left Prefrontal Cortex/Insula Region Showing a Round by Accuracy by Time Interaction. 92
Figure 9. Hippocampal Regions Showing a Main Effect of Time and their Time Courses of Activation. 93
Figure 10. Experiment 2 Behavioral Results from the Immediate Post-test. 94
Figure 11. Experiment 2 Behavioral Results from the One-week Follow-up Post-test. 95
Figure 12. Brain Regions Showing a Trial Type by Time Interaction in Experiment 2. 96
Figure 13. Brain Regions Showing a Feedback Type by Time Interaction in Experiment 2 97
Figure 14. Brain Regions Showing a Trial Type by Feedback Type by Time Interaction in Experiment 2. 98
Figure 15. Time Courses of Activation in the Right Caudate Region Showing a Trial Type by Feedback Type by Time Interaction. 99
Figure 16. Time Courses of Activation in the Left Prefrontal Cortex/Insula Region Showing a Trial Type by Feedback Type by Time Interaction. 100

Figure 17. Brain Regions Showing a Feedback Type by Time Interaction in the 4-choice Condition of Experiment 2..... 101
Figure 18. Brain Regions Showing a Feedback Type by Time Interaction in the 2-choice Condition of Experiment 2..... 102
Figure 19. Time Courses of Activation in the Caudate Regions Showing a Feedback Type by Time Interaction for the 4-choice Condition. 103
Figure 20. Medial Temporal Lobe Regions Showing a Main Effect of Time and their Time Courses of Activation. 104
Figure 21. Brain Regions Showing a Subsequent Memory Effect in Experiment 2. 105
Figure 22. Time Courses of Activation in the Caudate as a Function of Post-test Accuracy.... 106

PREFACE

I owe a debt of gratitude to the many people who have supported me while I have been a graduate student. First and foremost, I'd like to thank my advisor, Julie Fiez, for her guidance these past six years. She has been a supportive mentor, an insightful colleague, and an esteemed role model. I would also like to thank the other members of my doctoral committee, Chuck Perfetti, Erik Reichle, Mark Wheeler, and Jay McClelland, for the training they have given me throughout graduate school, as well as for their advice on the dissertation itself. I have been fortunate to have worked in a lab full of friendly people eager to share their expertise. I am especially grateful to Mauricio Delgado, who pioneered reward research in the Fiez lab, for his active participation in my scientific training. Susan Ravizza and Jason Chein were also extremely helpful as I learned the basics of fMRI research, and Kate Fissell's assistance with the technical aspects of data analysis has been invaluable. I also gratefully acknowledge Corrie Durisko, Mary Conway, and Marge Gibson, for keeping the lab running smoothly, and Sara Guediche and Michelle Moore for making our office a pleasant place to work. Finally, special thanks go to my parents, Al and Bet Tricomi, for instilling in me a love of learning, and to my husband, Mike Shiflett, for everything.

1.0 INTRODUCTION

Feedback is a term borrowed from cybernetics, and is used to describe a process by which the effects of a response control that response (Hill, 1997). In psychological studies of human behavior, the term describes stimuli used to signal performance accuracy, which can then be used to modify future performance (Holroyd & Coles, 2002; Kluger & DeNisi, 1996). Performance feedback not only serves to guide performance, but it also serves an emotive role since subjective feelings are induced by the "reward" or "punishment" of receiving positive or negative feedback (Elliott et al., 1998; Schultz et al., 1998). Recent evidence implicates the striatum, and specifically the caudate nucleus, in feedback processing. However, many questions about the neural basis of feedback-based learning remain unclear. For instance, most work on the neural basis of feedback-based learning has utilized nondeclarative learning tasks, so the role that feedback may play in declarative learning is still unspecified. The goal of the experiments presented here is to use functional magnetic resonance imaging (fMRI) to examine the influence of the caudate, as well as regions more typically engaged in declarative memory, on memory formation in a feedback-based declarative memory task.

1.1 THE ROLE OF THE STRIATUM IN FEEDBACK PROCESSING

Although it has no extrinsic value, the intrinsic value of feedback may be quite potent to motivated learners. If a shared neural system is involved in processing intrinsic and extrinsic rewards and punishments, then feedback should engage neural mechanisms responsible for processing reward-related information. The striatum is the input unit of the basal ganglia and has been implicated by both electrophysiological and imaging studies in the processing of primary and secondary rewards (Berns et al., 2001; Breiter et al., 2001; Delgado et al., 2000; Elliott et al., 2000b; Hikosaka et al., 1989; Knutson et al., 2000; Schultz, 1998). The striatum receives dopaminergic inputs from the substantia nigra (Alexander et al., 1986) and ventral tegmental area (Robbins & Everitt, 1996), as well as afferents from the amygdala, hippocampal formation, and prefrontal cortex (Robbins & Everitt, 1996). It sends information back to the prefrontal cortex by way of the globus pallidus (the output unit of the basal ganglia) and the thalamus, forming a frontal-striatal-thalamic loop (Alexander et al., 1986; Robbins & Everitt, 1996).

Much emphasis has been placed on the importance of the dopaminergic projections from the midbrain to the striatum in the processing of reward-related information. Electrophysiological recording studies of dopamine neurons in monkeys demonstrate that they respond to unpredicted primary rewards such as food and fluid; however, if the reward is preceded by a conditioned, predictive stimulus, the neuronal activity occurs following this learned stimulus and is absent following the actual reward. If a predicted reward does not actually occur, then the dopamine neurons are depressed at the time at which the predicted reward would have occurred. These dopaminergic neurons, however, do not differentiate among different rewards (Schultz et al., 1998). One hypothesis for the mechanism by which feedback may work is that dopamine released in the striatum may act as a reinforcement signal, allowing

behavior to be adaptively modulated to maximize future reward (Barto, 1995; Holroyd & Coles, 2002; Montague & Berns, 2002; O'Doherty et al., 2004).

Striatal neurons are activated during the expectation of predictable task events, with activity beginning after the predictive stimuli and ending immediately after the predicted event. Some of these neurons respond preferentially to certain types of rewards over others (e.g., apple juice vs. water), differentiating them from the nondiscriminating dopaminergic neurons (Schultz et al., 1998). One study involving a visually-guided saccade task in monkeys found that the preferred direction of neurons in the caudate nucleus, a structure in the dorsal striatum, changed as a function of which direction was rewarded, showing that the activity of caudate neurons can be modulated over short time periods through expectation of reward (Kawagoe et al., 1998). In humans, functional magnetic resonance imaging (fMRI) studies have shown that the caudate nucleus is activated by extrinsic monetary rewards (Delgado et al., 2000; Elliott et al., 2000b; Knutson et al., 2000). The caudate differentiates between reward and punishment trials, showing sustained activation following feedback indicating a monetary gain and activity that tends to decrease below baseline following feedback indicating a monetary loss (Delgado et al., 2000).

Neuroimaging studies confirm that the brain processes positive and negative feedback in a manner similar to the way it processes information about extrinsic rewards and punishments, such as money (e.g., Elliott et al., 1997; Tricomi et al., 2006). However, the caudate could be involved in feedback processing without contributing to feedback-based learning. For example, the caudate could simply be responding to the hedonic (pleasurable) aspects of the feedback. This appears, though, not to be the case. The caudate does not exhibit robust activation in response to randomly delivered or cued monetary rewards and punishments. Instead, it shows strong activation only when there is a perceived contingency between a response and the

outcome that follows the response (Tricomi et al., 2004). As feedback-based learning requires such goal-directed action, the caudate could act to facilitate learning by linking actions with their consequences. Recent neuroimaging studies support this idea. For example, activation in the caudate nucleus during a probabilistic learning task was found to occur only when the task was feedback-based (Poldrack et al., 2001). Moreover, performance of a perceptual learning task with feedback produced caudate activation and successful learning, while doing the same task with no feedback had neither of these effects (Tricomi et al., 2006). In addition, caudate activation has been shown to vary as learning occurs; as contingencies between stimuli and correct responses become well-learned, caudate activation decreases (Delgado et al., 2005b; Law et al., 2005). Finally, fMRI studies have shown that reversal learning, which involves learning new simulus-response-outcome contingencies and ignoring old ones, activates the caudate nucleus (O'Doherty et al., 2003; Rogers et al., 2000).

This neuroimaging work is also complemented by neuropsychological work that indicates that patients with Parkinson's disease and Huntington's disease, which disrupt striatal function, are impaired on a variety of feedback-based nondeclarative learning tasks (Monchi et al., 2004; Packard & Knowlton, 2002; Seger, 1994; Shohamy et al., 2004). One recent study showed that Parkinson's patients were specifically impaired on a feedback-based version of a probabilistic category-learning task, but not a non-feedback version of the task (Shohamy et al., 2004). Parkinson's patients are not impaired, however, on typical artificial grammar learning and prototype learning tasks, which involve category learning but do not involve feedback (Reber & Squire, 1999). When artificial grammar learning tasks are modified to be feedback-based, then Parkinson's patients do show deficits (Smith & McDowall, 2006).

1.2 THE ROLE OF THE STRIATUM IN NONDECLARATIVE LEARNING

Thorndike, an early 20th century psychologist, did pioneering research on reinforcement learning, or in other words, on how outcomes of behavior serve to build up stimulus-response connections. His research led him to come up with his influential "law of effect," which states that "what comes after a connection acts upon it to alter its strength" (Thorndike, 1927). According to this law, how the strength of a stimulus-response (S-R) connection was altered depended on whether a "satisfier" or an "annoyer" followed the connection (Thorndike, 1911). Thorndike believed that S-R habits were slowly built up without the encoding of any information about the outcome, i.e., that a "satisfier" simply served to strengthen an S-R bond without any knowledge of a response-outcome contingency being learned (Balleine & Dickinson, 1998; Dickinson & Balleine, 2002).

This idea has had a pervasive influence on scientific thinking about how habit learning (i.e., the gradual building up of S-R associations) occurs (Yin & Knowlton, 2006). Traditionally, the striatum has been thought to carry out the function of habit learning in much the way Thorndike originally proposed, by slowly strengthening S-R associations over many repetitions (e.g., Kantak et al., 2002; Kantak et al., 2001; McDonald & White, 1993). Nevertheless, Thorndike's view of a learning process that is blind to action-outcome contingency turns out to be false. First, this idea doesn't allow for motivational state (e.g., hunger) to modify behavior, but empirical evidence shows that changes in motivational state do affect behavior (Balleine & Dickinson, 1998; Dickinson & Balleine, 2002). Second, the idea also fails because it predicts that as long as an S-R pairing is followed by a reward, the S-R bond should be strengthened, regardless of response-outcome (R-O) contingency. That is, it predicts that if there are also presentations of the reward that don't follow a response (which degrades the contingency

between the response and the outcome) this shouldn't affect future responding. In actuality, doing this depresses instrumental performance (Balleine & Dickinson, 1998). Therefore, performance is sensitive to the information an outcome provides about the consequences of an action. In other words, feedback is used to build up stimulus-response-outcome (S-R-O) associations. So in actuality, the development of a habit consists of two distinct processes: one process, which involves goal-directed behavior (instrumental learning), and one which involves a frequency-based strengthening of a response (habit learning). Early in learning, the former process dominates; if the contingencies between a response and outcome change, the animal will change its behavior (Balleine & Dickinson, 1998). Eventually, however, the response to a stimulus becomes a true habit, and becomes independent of the outcome; the animal will continue to perform the same response, even if the outcome is devalued or omitted (Balleine & Dickinson, 1998; Cardinal et al., 2002; Dayan & Balleine, 2002; Dickinson & Balleine, 2002; Yin et al., 2004).

Recent evidence suggests that distinct subregions of the striatum may be preferentially involved in these two aspects of nondeclarative learning tasks. In rats, the dorsomedial striatum (DMS), also referred to as "central striatum" (Bar-Gad et al., 2003; Haber et al., 2000), is roughly equivalent to the caudate in primates (Yin & Knowlton, 2006). Although its precise role is not yet fully understood, recent work has shown it to be functionally distinct from dorsolateral striatum (DLS), the region roughly equivalent to the putamen in primates (Yin & Knowlton, 2006). Lesion studies provide support for the idea that the DMS is important for goal-directed action, while the DLS seems to subserve habit learning (Devan et al., 1999; Yin et al., 2004). For example, although normal taste aversion learning occurs, DLS-lesioned rats decrease lever pressing after taste aversion, unlike controls or DMS-lesioned rats, who continue to lever press.

This indicates a role of the DLS in maintaining S-R habits even in the face of altered contingencies (Yin et al., 2004). Meanwhile, the DMS has been shown to be especially important for reversal learning tasks. DMS lesions lead to regressive errors; that is, rats go back to performing an old response, even after a new correct response has been made (Ragozzino, 2003).

This evidence suggests that although the striatum is involved in the development and maintenance of habits, which are processes associated with nondeclarative learning, this function is complemented by another role of the striatum. In particular, the DMS (or caudate in primates) seems to be involved in learning S-R-O contingencies, which is the basis of feedback-based learning. Since these two functions appear to be distinct and to be subserved by distinct subregions of the striatum, they may not always go hand in hand. In a situation that involves feedback-based learning, but does not rely exclusively on nondeclarative learning, the caudate could still be involved. For example, through its strengthening of S-R-O associations, the caudate could act to make recall of correct responses less effortful, even when the responses involve semantic knowledge usually thought to be acquired via declarative memory.

1.3 STRIATAL INFLUENCES ON THE DECLARATIVE MEMORY SYSTEM

The striatum is not usually associated with declarative learning; instead, other brain regions, including the hippocampus and adjacent cortex within the medial temporal lobe (MTL), and the prefrontal cortex (PFC), are usually thought of as subserving declarative learning and memory. The often profound deficits of amnesics with MTL damage on declarative memory tasks, along with their relatively preserved ability on nondeclarative memory tasks, underscores

the importance of the MTL for declarative memory in particular (Squire, 2004; Zola & Squire, 2000). Neuroimaging studies have indicated that the MTL is activated during both memory encoding and retrieval, although there is still some debate about whether different subregions are responsible for the two processes (Eldridge et al., 2005; Meltzer & Constable, 2005). Futhermore, activation levels in the hippocampus (Meltzer & Constable, 2005; Reber et al., 2002) and parahippocampus (Brewer et al., 1998; Reber et al., 2002; Wagner et al., 1998) during encoding have been found to correlate with performance on subsequent memory tests, and MTL activation during retrieval has also been found to be correlated with memory accuracy (Law et al., 2005; Meltzer & Constable, 2005; Wittmann et al., 2005). Subsequent memory effects have also been found in the left inferior PFC for memory of words (Wagner et al., 1998), and right PFC for memory of pictures (Brewer et al., 1998). The right PFC also appears to be involved in retrieval effort, with some studies showing increased PFC activation bilaterally during successful retrieval (Nyberg & Cabeza, 2000). The functions of the MTL in memory appear to be independent of conscious effort, while activation in the PFC seems to reflect deliberate effort (Reber et al., 2002).

There are several possibilities of how this system might act during feedback-based declarative learning tasks. The declarative memory system could operate wholly autonomously from the striatal system, in which case both systems might be activated. This activation could have a redundant effect on learning, so that the amount learned would be the same regardless of whether both systems were utilized, or the effects on learning could be additive, in which case more would be learned than if only one of the systems were utilized. Conversely, the two neural systems could interact. Activation of one system could inhibit activation of the other, or

conceivably the two systems could act synergistically, so that activation of one system would potentiate activation in the other.

Each of these possibilities has garnered some support. Squire argues that different memory systems operate "independently and in parallel" (Squire, 2004). In other words, multiple types of memory formation can occur at the same time by different memory systems. One type of memory may seem to dominantly affect behavior, but if the region responsible for the dominant behavior is inactivated, other behavioral strategies will emerge that make it clear that other brain regions have also acquired a memory, in a different form. For example, when a rat uses a declarative strategy to return to a rewarded location rather than to perform the action that was rewarded (e.g., a left turn), this strategy choice becomes reversed when the hippocampus is inactivated with lidocaine injections (Packard & McGaugh, 1996). Likewise, one recent neuropsychology study found that two patients with MTL lesions and profound declarative memory impairments were nevertheless able to slowly acquire a feedback-based discrimination learning task that control subjects learn quickly with a declarative memory strategy, presumably by making use of a striatal learning system (Bayley et al., 2005). These studies leave open the question, however, of whether for humans with no brain damage, the striatum and MTL would be employed simultaneously during tasks such as these.

An alternate view is expressed by Poldrack and Rodriguez (2004), who argue that the striatal and hippocampal systems interact competitively, whereby when one system is active it suppresses activity in the other, with the PFC mediating this interaction. This view is supported by a neuroimaging study, in which an inverse relationship was found between activation in the caudate and medial temporal lobe. In a typical weather-prediction task with feedback (a probabilistic, nondeclarative learning task), initial MTL activation quickly decreased, and the

caudate nucleus became activated. In a paired-associate version of the task, designed to engage declarative memory, activation was greater in the MTL and less in the caudate than for the feedback-based version of the task (Poldrack et al., 2001). However, since there was no condition involving a declarative memory task with feedback, it is impossible to tell whether the inverse relationship between caudate and MTL activation would persist in such a situation.

Finally, some have argued for the third point of view, that is, that the brain's dopamine-mediated reward system, which is usually associated with the striatum's role in learning and memory, facilitates hippocampus-dependent memory formation (Wittmann et al., 2005). Neural circuitry supports the idea that the hippocampus and striatum may interact. Cells in the nucleus accumbens, a structure in the ventral striatum, receive excitatory input from the hippocampus. The nucleus accumbens inhibits the globus pallidus (ventral pallidum), releasing its tonic inhibition of the ventral tegmental area (VTA), a part of the midbrain containing dopaminergic neurons (Lisman & Grace, 2005). Although there are no direct projections between the hippocampus and dorsal striatum, the dorsal striatum is interconnected with the substantia nigra (Haber et al., 2000), a part of the midbrain which, along with the VTA, sends dopaminergic projections to the hippocampus (Lisman & Grace, 2005). Functionally, the MTL and nucleus accumbens have been shown to exhibit greater activation when there is high motivation to learn than when there is low motivation (Adcock et al., 2006). Furthermore, long-term, incidental memory of reward-associated cues has been found to be greater than memory for neutral cues, indicating a facilitatory effect of reward-related activation on long-term memory (Wittmann et al., 2005).

Given the varying viewpoints present in the literature, the issue of the striatum's potential influence on the declarative memory system warrants further investigation. One goal of the

research presented in this thesis was to address this issue by imaging the brain during execution of a feedback-based declarative memory paradigm.

1.4 MODULATION OF CAUDATE ACTIVATION DURING FEEDBACK-BASED DECLARATIVE LEARNING

Much remains to be known about which aspects of feedback processing determine caudate activation. Such activation is not consistent across all tasks involving feedback, but instead seems to be modulated by a number of factors. For example, the caudate nucleus is recruited only when participants perceive that there is a contingency between their action and whether they subsequently receive a reward or punishment (Tricomi et al., 2004). Therefore, feedback might be expected to more strongly activate the caudate when learners feel a strong sense of agency in determining the outcome. If the learner has no prior knowledge of the task, and therefore no basis on which to choose a response, the caudate might be less strongly activated than when the feedback is a true assessment of performance.

Another, potentially related, factor that has been found to modulate caudate activation is the value of feedback to the individual. In the animal literature, this is sometimes termed the "incentive value" of an outcome, and can be thought of as the degree to which it is goal-worthy; this value is not absolute, but may change as a function of motivational state (Balleine & Dickinson, 1998; Dayan & Balleine, 2002; Dickinson & Balleine, 1994). Nothing about feedback displays (e.g., green checkmarks) is intrinsically rewarding. It is the meaning in relation to the task that endows feedback with value. Notably, however, the goals of obtaining positive feedback and improving performance are not identical (Kluger & DeNisi, 1996).

Therefore, the perceived value of feedback should be modulated by whether it provides meaningful information about performance.

The perceived value of an outcome is also affected by the other possible outcomes. Studies of "counterfactual comparisons" indicate that reward-related brain regions respond to an outcome's value relative to the alternative outcomes, rather than to its absolute value (Breiter et al., 2001; Nieuwenhuis et al., 2005; Ursu & Carter, 2005). For example, when the alternative is a monetary gain, winning no money produces a punishment response in reward-sensitive brain regions, including the caudate nucleus, while the same outcome produces a reward response when the alternative is a monetary loss (Nieuwenhuis et al., 2005).

Another factor that could alter the perceived value of feedback is the amount of task-related information provided by feedback. Performance-related information indicates whether the participant's response is accurate or inaccurate, while task-related information indicates *which* response is accurate. Positive feedback indicates that the response chosen is the accurate one, while negative feedback eliminates that response from the possible accurate responses. The use of a declarative memory task allows for the amount of information provided by feedback to be easily operationalized, since the relationships between the stimuli, responses, and outcomes can be clearly delineated and made obvious to the learner. For the purposes of this thesis, I will use as a definition of information the difference between the probability of the learner answering correctly on trial t+1 and trial t, given perfect memory. For example, on trial 1 of a given problem that has two possible response options and no prior knowledge of which is correct, the chance of answering correctly is 0.5. Feedback either confirms the guess or eliminates the response from the set of responses that could be accurate. On trial 2 of the same problem, given accurate memory of the information provided by the feedback on trial 1, the chance of answering

correctly is now 1.0. The probability of answering correctly has increased by 0.5, and this can be taken as a measure of the amount of information provided by the feedback.

Based on this definition, the amount of information provided by positive feedback increases with more response options, or in other words, as the chance of answering correctly decreases. Conversely, negative feedback provides the most information with only two possible options. In both cases, it is when the type of feedback deviates most from expectation that it is most informative. Indeed, the dopaminergic system seems to be sensitive to this measure. Midbrain dopamine neurons, which project to the striatum, have been found to fire in proportion to how expected a reward is. Unexpected rewards produce the most firing, while firing in anticipation of an upcoming reward increases as rewards become more expected (Fiorillo et al., 2003). Further studies have found that reward-related activity in the striatum is greatest when the reward is unpredictable (Berns et al., 2001; McClure et al., 2003). Finally, once a task is well-learned, feedback becomes completely expected, and therefore ceases to provide information. Consistent with the idea that the striatum in sensitive to the informational value of feedback is the finding that activity in the caudate nuclei appears to decrease over the course of learning (Delgado et al., 2005b; Haruno et al., 2004; Law et al., 2005; Pasupathy & Miller, 2005).

In this thesis, I describe two experiments that I performed to examine the role of the caudate in processing performance feedback and how this processing may contribute to the learning process. In Experiment 1, changes in caudate activation were examined over the initial stages of learning. The brains of participants were scanned as they encountered the same trials of an arbitrary paired associate learning task with feedback three times. Experiment 2 examined how the number of response options in a feedback-based declarative memory task affects

activation in the caudate during learning. Participants were scanned while performing the task with trials that had 2 or 4 response options, allowing me to compare the effects of feedback that carries different amounts of information.

2.0 EXPERIMENT 1: THE EFFECT OF THE TYPE OF INFORMATION PROVIDED BY FEEDBACK ON CAUDATE ACTIVATION

The head of the caudate nucleus, which lies within the dorsal striatum, has been implicated in processing reward-related information (Delgado et al., 2000; Elliott et al., 2000b; Knutson et al., 2000), including the processing of performance-related feedback (e.g., (Elliott et al., 1997; Poldrack et al., 2001; Tricomi et al., 2006). However, little is known about the specific role that the caudate is playing in feedback-based learning. The goal of this experiment was to examine the parameters that govern feedback-related caudate activation.

Feedback provides multiple types of information. First, it provides information about what the correct response is. Second, it can serve as an indication of task success. It is unclear whether feedback-related caudate activation discriminates between these two types of information. To examine this issue, this experiment made use of a feedback-based paired associate word-learning task, in which participants encountered the same sets of trials three times. The task was set up in a "multiple choice" format, with a target word and two choices of possible words for the second half of the pair. The participants began with no prior knowledge of the correct pairings, which were arbitrary. On the first trial of this task, they were required to pick an answer and learn based on feedback whether the response guessed was correct. The second time they encountered the same trial, they tried to retrieve the memory of the correct answer and respond accordingly, and feedback indicated whether their answer was correct. Note

that feedback is serving two different functions during these two memory stages. Since the word pairs were arbitrary, performance was also arbitrary during the initial encoding stage. At this point, feedback provides information about what the correct answer is, but does not provide an indication of task success. That is, learners are aware that negative feedback is not a reflection of poor performance, but instead, only of bad luck. In fact, with only two response options, the correct answer can be determined from either negative or positive feedback. Therefore, the two outcomes may not be perceived as having differential value. This is one case in which the goal of mastering a task is not related to obtaining positive feedback. In contrast, feedback during later, "retrieval" trials provides an assessment of memory performance, and consequently of task success. In this instance, positive feedback may be perceived as more rewarding than negative feedback.

The main hypothesis tested in this experiment was that caudate activation would be higher the second time each participant encountered a particular trial than the first. This hypothesis was based on several previous findings related to reward processing in the caudate. First, there is some evidence to suggest that the caudate should be more activated when participants feel a strong sense of agency in determining the outcome. Specifically, in a guessing task with monetary outcomes, the caudate was recruited only when participants believed that there was a contingency between their action and whether they subsequently received a reward or punishment (Tricomi et al., 2004). If a learner has no prior knowledge of the task, and therefore no basis on which to choose a response, the caudate might be less strongly activated than when the feedback is a true assessment of performance.

Another factor that has been found to modulate caudate activation is the value of feedback to the individual. In a guessing task, the caudate responded much more strongly when

feedback indicated monetary gain or loss than when there were no monetary incentives (Delgado et al., 2004). However, in a learning study in which feedback indicated performance accuracy but had no monetary value, robust caudate activation was also observed (Tricomi et al., 2006). In this case, the goal of participants was to learn to perform well on the task, rather than to earn money, so positive feedback was valued highly because it provided information about achievement of this goal. When the information provided by positive and negative feedback is equivalent and unrelated to task success, it may not have much value to the learner, and therefore may not robustly activate the caudate. In contrast, the caudate may activate more strongly on trials involving memory retrieval, since feedback then provides information about the degree to which the learner's goal of mastering the task has been achieved.

Finally, once a task is well-learned, feedback becomes completely expected, and therefore ceases to provide information. At this point, positive feedback may become less rewarding. Several studies have found that reward-related activity in the striatum is greatest when the reward is unpredictable (Berns et al., 2001; McClure et al., 2003). Further studies have indicated that activity in the caudate nuclei appears to decrease over the course of learning (Delgado et al., 2005b; Haruno et al., 2004; Law et al., 2005; Pasupathy & Miller, 2005). These findings led to a second hypothesis that caudate activation might be somewhat attenuated on the third round of trials, as feedback is beginning to become more expected.

2.1 METHODS

2.1.1 Participants

Twenty healthy, right-handed adults were recruited through posted advertisements and

were paid $57 for their participation in the experiment. Of these participants, four were excluded from analysis due to excessive head movement, one was excluded due to a technical problem, and four were excluded due to ceiling performance. Data from the remaining eleven subjects were analyzed (5 women, 6 men; mean age ± SD, 21.9 ± 2.3) All participants signed an informed consent form according to the Institutional Review Board at the University of Pittsburgh.

2.1.2 Materials

A 3 Tesla Siemens head-only scanner and standard radio frequency coil was used for all the MR scanning sessions. Stimulus presentation and behavioral data acquisition was controlled using "E-prime" software (Schneider et al., 2002) and the integrated function imaging system (IFIS) [Pittsburgh, PA].

2.1.3 Procedure

Scan session:

Structural images were collected using a standard T1-weighted pulse sequence, in thirty-eight contiguous slices (3.125 x 3.125 x 3.0 mm voxels) parallel to the AC-PC line. Thiry-eight functional images were collected in the same locations as the structural slices. They were acquired using a one-shot echo-planar imaging (EPI) pulse sequence [TR=2000 ms, TE=25 ms, FOV=20 cm, flip angle = 79°].

Behavioral paradigm:

This experiment involved a paired associate word learning task, and participants were scanned as they performed 3 rounds of trials on this task, that is, the initial encoding round, and

two subsequent rounds, which may include memory retrieval as well as further encoding. Participants received performance-dependent feedback following each trial throughout the scan session.

The scan session consisted of nine six-minute runs, divided conceptually into three rounds of 60 trials. Each trial was 18 s long. On each trial of Round 1 (the first three runs), participants saw a target word and two choices of possible word matches, labeled "a)" and "b)" (Figure 1). The words 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. Additionally, 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 et al., 1998) and did not rhyme or begin with the same letter. Participants were asked to guess which response word went with the target word by pressing one of two buttons on a response glove. Since the words were unrelated, guesses were arbitrary. Participants had 4 s to respond, after which the display was replaced by a feedback display, which was shown for 1 s (Figure 1). On 50% of trials, participants received positive feedback following their response (3 green √s), indicating they guessed correctly, and on the other 50% of trials, participants received negative feedback (3 red Xs), indicating they guessed incorrectly. If the subject made no response, three white hyphens were shown and the trial was excluded from analysis. Participants were asked to use the feedback to try to remember the correct word pairs (the correct pairings remained consistent throughout each scanning session, although which item was the correct one was chosen randomly for each participant). Each trial ended with a 13 s delay period in which participants fixated on a white cross in the center of the screen. Participants were not given any

constraints or advice on how to remember the word pairs.

On Round 2, the same 60 trials were repeated, in random order, with the position of the response options chosen randomly (i.e., which was choice "a" and which was choice "b"). Although the procedure was the same, it should be noted that participants' guesses were no longer arbitrary; instead, participants were asked to pick the correct response based on the feedback they received during Round 1. Feedback was again presented after each trial, indicating whether the participant answered correctly or incorrectly. Finally, on Round 3, the procedure from Round 2 was repeated, with the same sixty trials presented in random order.

Following the scan, participants took a computerized post-test, using the E-prime program, in which the 60 trials were repeated one more time, with no feedback. Following each trial, participants were asked to make a confidence judgment by choosing a number from 1-7 (1 = complete guess, 7 = completely sure). The post-test was self-paced, unlike the trials during the scan. Approximately one week later (7-9 days following the scan session), participants returned to the lab and completed the post-test a second time, with the exception of one participant, who never returned.

2.1.4 Data Analysis

Behavioral Data

Analyses were performed on the behavioral data of the participants who were used in the fMRI data analyses. One-sample t-tests were used on the accuracy data from each round and the post-tests to determine which conditions differed from chance. A repeated measures ANOVA was performed with subject as a random factor and round as a within-subjects factor to determine if accuracy changed over the course of the experiment. A similar ANOVA was

performed to assess whether reaction time changed over the course of the scanning session. Two-tailed t-tests were used to determine whether reaction time and/or confidence differed between correct and incorrect trials. Finally, a linear regression of reaction time on confidence was fit for the post-test data to determine whether reaction time differed as a function of confidence.

FMRI Data

The NeuroImaging Software package (NIS 3.5), developed at the University of Pittsburgh and Princeton University, was used to analyze the fMRI data. Images were reconstructed and corrected for subject motion with Automated Image Registration (AIR 3.08; (Woods et al., 1992). Subjects with head motion that exceeded 4 mm or degrees in any direction were not used in the analysis. The images were detrended to adjust for scanner drift within runs. The structural images of each subject were stripped to remove the skull and co-registered to a common reference brain (chosen from among the participants; (Woods et al., 1993). Functional images were transformed into the same common space, normalized by a mean scaling of each image to match global mean image intensity across subjects, and smoothed using a three-dimensional Gaussian filter (8 mm FWHM) to account for anatomical differences between subjects. This set of data was then analyzed statistically.

A voxel-wise analysis was performed on the fMRI data, through a repeated-measures three-way ANOVA with subject as a random factor and accuracy (correct vs. incorrect), round (1-3), and time period (2 s time periods T1-T9) as within-subject factors. As mentioned above, four participants were excluded from analysis because they showed a ceiling effect; they had fewer than 8 incorrect trials for Round 3. Regions of interest (ROIs) were defined as regions with four or more contiguous voxels showing a significant effect; this contiguity threshold

served as a precaution against type-1 errors (Forman et al., 1995). All ROIs were transformed into Talairach space (Talairach & Tournoux, 1988) using the AFNI program (Cox, 1996).

Further analysis focused on whether levels of activation during Round 1 predicted accuracy on Round 2 (Wagner et al., 1999). To do this, trials were coded according to the accuracy for the identical trial on the following round. A repeated measures ANOVA was then performed with subject as a random factor and current round accuracy (correct vs. incorrect), future accuracy (correct vs. incorrect), and time (2 s time periods T1-T9) as within-subject factors. To increase the power of this analysis, data from participants whose performance was at ceiling on Round 3 were included, as was data from the three participants who were excluded from the main analysis due to excessive movement, since the movement occurred after Round 1; thus, n = 19 for this analysis. A similar analysis was not performed on the data from Rounds 2 and 3, since if performance is accurate on any given trial on Round 2, it is more likely to be accurate subsequently as well, and so effects due to performance and to learning are confounded on these rounds.

2.2 RESULTS

2.2.1 Behavioral Results

Table 1 lists the mean accuracy and reaction times for each round in the scanner and the two post-tests. With the exception of Round 1, which was programmed so that each participant would receive positive feedback 50% of the time, performance was significantly better than chance, indicating that learning occurred. A repeated measures ANOVA indicated that accuracy changed over the course of the experiment ($F(4,39.3) = 56.7$, $p < 0.05$). Reaction time did not

differ significantly over the three rounds in the scanner ($F(2,20) = 1.57$, $p > 0.05$). Reaction times during the scan session were not compared to post-test reaction times, since the post-tests were self-paced.

Two-tailed t-tests showed no significant differences between reaction times for correct and incorrect trials on Rounds 1 and 2, but significant differences on Round 3 and the two post-tests, with faster times for correct trials (Table 2). Confidence was significantly greater for correct trials than incorrect trials on both post-tests (Table 3). Finally, increases in confidence judgments were associated with decreases in reaction time; mixed model linear regression analyses of reaction time on confidence, with subject as a random factor, showed this effect to be significant for both post-tests ($F(6, 641) = 13.8$ and $F(6, 581) = 5.5$, for post-tests 1 and 2, respectively, $p < 0.05$).

2.2.2 fMRI Results

A voxel-wise ANOVA with subject as a random factor and round (1-3), accuracy (correct vs. incorrect), and time period (2 s time periods T1-T9) as within-subject factors was performed. The resulting activation clusters showing a Round by Time effect, at a threshold of $p < 0.0001$ and a contiguity threshold of four voxels (cf. Tricomi et al., 2004), are listed in Table 4 and shown in Figure 2. Although there were subtle differences in activation between Rounds 2 and 3, the most salient differences were between Round 1 activation and the activation during the other two rounds. Some regions showed more activation for Round 1 than Rounds 2 and 3, while other regions showed greater activation for Rounds 2 and 3 than Round 1 (Table 4). A region was identified in the caudate nuclei that showed the latter pattern, with a flat response during Round 1 and an increase in the response during Rounds 2 and 3 (Figure 3). The response

is greatest for Round 2, with a slight attenuation for Round 3. One might expect these differences to be most apparent at the response peak, which occurs at T3, 4-6 s after the trial start. Levels of activation at this time point represent the magnitude of the initial rise in activation that occurs even before the feedback display; due to the sluggishness of the hemodynamic response in the brain, one would not expect outcome-related effects on the BOLD signal to occur until several seconds after the feedback presentation. Pairwise two-tailed t-tests performed on time point T3 revealed significant differences between Rounds 1 and 2 ($t(10) = -4.0$, $p < 0.05$) and between Rounds 1 and 3 ($t(10) = -2.5$, $p < 0.05$), with a trend towards significance between Rounds 2 and 3 ($t(10) = 1.9$, $p < 0.1$).

The activation clusters showing an Accuracy by Time effect at $p < 0.0001$ are listed in Table 5 and shown in Figure 4. Bilateral activation was found in the ventral head of the caudate. The significance threshold was raised until the caudate region separated from contiguous regions, and the time courses of activation, shown separately for each round, are depicted in Figure 5. Although the time course of activation during Round 1 is relatively flat, there is a slight differentiation between trials with positive and negative feedback, with the activation dipping slightly below baseline following negative feedback. For Rounds 2 and 3, after an initial rise, the signal following negative feedback shows a pronounced dip below baseline, whereas the signal following positive feedback does not. In previous work, the BOLD signal has shown significant differentiation between correct and incorrect trials 6-9 s after the feedback display (Tricomi et al., 2006). This corresponds to T6 and T7 in the current experiment. Two-tailed t-tests were performed on these time points for each round. Significant differences were found at T6 on Round 1 ($t(10) = 2.3$, $p < 0.05$); at T6 and T7 on Round 2 ($t(10) = 6.0$ and 3.8, respectively, $p < 0.05$), and at T6 and T7 on Round 3 ($t(10) = 7.9$ and 4.4, respectively, $p < 0.05$).

The activation clusters showing a three-way Round by Accuracy by Time effect at $p < 0.0001$ are listed in Table 6 and shown in Figure 6. In general, the difference between correct and incorrect trials was greatest on Round 3, with some clusters exhibiting more activation for correct trials than incorrect trials and some showing the reverse pattern. The only clusters of activation showing a significant three-way interaction in the caudate were in the caudate body, whereas the clusters showing Round by Time and Accuracy by Time effects were observed in the head of the caudate. The activation pattern in the caudate body clusters shows a distinct pattern of activation. Activation is elicited on all three rounds, with an earlier response on correct trials than on incorrect trials during Round 3, but not Rounds 1 and 2 (Figure 7).

By far the largest cluster showing the three-way interaction was in the left dorsolateral prefrontal cortex (DLPFC) and inferior frontal gyrus (IFG). The time course from this region, plotted by Round, is shown in Figure 8. Although there were separate peaks in the DLPFC and IFG, when the threshold was raised to isolate these subregions, they were found to show similar activation patterns; thus the time course from the larger region has been plotted. While incorrect trials show greater activation than correct trials on all three rounds, the signal most strongly differentiates on Round 3.

Although the hippocampus/parahippocampus did not show any significant interactions, it did show a very strong main effect of time, bilaterally (peak Talairach coordinates: -27, -16, -12 and 23, -12, -15). The threshold was increased to isolate the hippocampal activation clusters from contiguous activated regions. These clusters are shown in Figure 9, along with the time courses of activation. The signal decreases from trial onset, and recovers following the feedback display. This recovery is increasingly slower following negative feedback as round increases.

Previous research has indicated that the parahippocampal/fusiform gyrus and the prefrontal cortex are more active when successful encoding is taking place, as measured by performance on a subsequent memory test (Brewer et al., 1998; Wagner et al., 1998). To see if a similar effect was present in the current task, the data on Round 1 were coded by whether the participant would answer correctly on the corresponding trial on Round 2, and an ANOVA was performed to identify regions showing a Future Accuracy X Time effect. No regions were identified at the significance threshold of $p < 0.0001$, with a cluster threshold of 4 voxels. An exploratory analysis at a threshold of $p < 0.001$, however, did reveal regions similar to those found in previous work that displayed greater activation on trials associated with subsequent correct performance than on trials associated with subsequent incorrect performance. Specifically, such regions included the left fusiform gyrus (peak Talairach coordinates: -50, -43, -9), left posterior lateral inferior frontal gyrus (peak Talairach coordinates: -40, -1, 29), and the left middle frontal gyrus (peak Talairach coordinates: -47, 21, 26), with a location superior to the anterior lateral inferior frontal gyrus area found to show a subsequent memory effect in previous work (cf. Wagner et al., 1998). No regions were identified in the striatum or the hippocampus.

Finally, to test the idea that an increased signal in the caudate on correct trials might act to decrease reaction time on the subsequent trial (by increasing confidence and making the retrieval of the correct response more automatic), the Round 2 trials for which performance was accurate on both Rounds 2 and 3 were coded by the difference in reaction time between the two rounds. A median split was used to divide the trials into those with a large reaction time decrease versus a small reaction time decrease (or an increase in some cases). Although no clusters showing a reaction time difference by time effect were identified in the caudate nuclei at a threshold of $p < 0.0001$, an exploratory analysis with a threshold of $p < 0.001$ did identify a

cluster in the right ventral caudate nucleus (peak Talairach coordinates: -8, 6, 0). The activation peak was higher for trials with a subsequently large reaction time decrease than for a small reaction time decrease. At this threshold, other regions also showed a reaction time difference by time effect, including the midbrain, right cuneus, right postcentral gyrus, and right superior cerebellum.

2.3 DISCUSSION

The main finding of this experiment is that the head of the caudate is more robustly activated when feedback indicates task success, rather than when it is purely informational. Although the same sixty trials were repeated three times, the resulting activation in the caudate, as well as other brain regions, was quite different. This discussion will focus on activation in the caudate and MTL; a discussion of the role of other brain regions involved in this experiment and Experiment 2 is covered in the final chapter of this thesis. On Round 1, when feedback was arbitrarily related to performance, the BOLD signal in the caudate head did not rise above baseline. There was a significant rise in activation on Round 2, once feedback began to reflect task performance. In other words, feedback was then serving as an indication of whether the learner was performing well or poorly. Learners may have felt that if they received negative feedback, it was their "fault" for not remembering the word pair accurately. In contrast, they knew that since there was no way to tell in advance of the experiment what the correct pairs were, that they would receive negative feedback approximately 50% of the time during Round 1, no matter what. The caudate, then, does not seem to process feedback indiscriminately; the

meaning of the feedback in relation to the participant's goals and expectations plays a crucial role in the response of the caudate.

A second finding was that although the caudate response was still significantly greater on Round 3 than on Round 1, the magnitude of the rise was somewhat attenuated from the signal rise on Round 2. This fits within the context of other experiments that have shown diminished activity in the caudate as appropriate responses become well-learned (Delgado et al., 2005b; Haruno et al., 2004; Law et al., 2005; Pasupathy & Miller, 2005). In the present experiment, correct performance was already at 75% for Round 3. One reason for the attenuation effect may be that as correct responses become well-learned, positive feedback becomes more expected. Therefore, it carries less information about performance. In other words, if a learner is already highly confident that a particular response is correct, positive feedback does not tell the learner anything new. If, on the other hand, the learner has only a "hunch" that a given response is correct, positive feedback confirms that hunch, and in doing so provides performance-related information that the learner wouldn't have known otherwise.

Supporting this idea is the fact that reaction time begins to significantly differ between correct and incorrect responses on Round 3. On the post-tests, confidence and reaction time were inversely correlated; participants tended to respond more quickly on trials for which they were most confident. We can assume that this effect was true during the scanning session as well, so although explicit measures of confidence were not obtained during the scanning session, reaction time can be used as a proxy for confidence. Therefore, that reaction times were faster for correct trials than incorrect trials during Round 3 is an indication that participants were more confident of their correct responses. In contrast, on Round 2, no such effect exists, indicating

that although performance was above chance, participants may not yet have been more confident of their correct responses than their incorrect responses.

If activity in the caudate nucleus is contributing to learning, then one might expect that confidence would increase more when the caudate signal is greatest. Indeed, there was some evidence that this was the case. If reaction time tracks inversely with confidence, then large decreases in reaction time should signify increases in confidence, relative to small decreases or increases. The right caudate nucleus did show greater peak activation on correct trials during Round 2 that were associated with a large decrease in reaction time on Round 3. Although this effect was found at an exploratory significance threshold, it raises the intriguing possibility that caudate activation may act to solidify response-outcome knowledge. Over repeated trials, such activation could potentially automatize responding, such that it becomes more habitual.

As in other studies of reward-related processing, the signal in the caudate nucleus differentiates between positive and negative outcomes in this experiment (e.g., Delgado et al., 2000; Tricomi et al., 2006). Specifically, there is a greater signal on trials with positive feedback than negative feedback. There are, however, some subtle, yet interesting, differences between the pattern observed in this experiment compared to prior experiments. Previous experiments have shown a sustained response to outcomes indicating monetary rewards or correct performance, while the response to monetary loss or negative feedback decreases sharply, even dipping below baseline activation in some cases. In this experiment, for Rounds 2 and 3, which show the greater overall response in the caudate, the response to positive feedback returns to baseline relatively quickly following the response peak. It does not dip below baseline, but neither is the response particularly sustained. By time point T5 (4-6 s after the onset of the feedback display) the response has returned to the baseline level, whereas in other work the

"reward" signal remains elevated until about 9 s after the outcome display (cf. Tricomi et al., 2006; Tricomi et al., 2004). The signal following negative feedback, on the other hand, shows an exaggerated dip below baseline on Rounds 2 and 3 (-0.13% decrease, compared to -0.06% in prior work; cf. Tricomi et al., 2004)). This may be attributed to the fact that learning occurred quickly in this task, so that even by Round 2 performance was above chance. In previous work, performance remained at or near chance throughout the scanning session. As discussed above, caudate activation seems to be related to the participant's expectations of the possible outcome. The more a reward is expected, the less "rewarding" it may seem, and the sustained activation that seems to define a "reward" response may be attenuated. Conversely, the more a reward is expected, the more unexpected a punishment response would be, and so the defining dip of a "punishment" response may increase in magnitude.

Though not the focus of this study, the body of the caudate showed an interesting pattern of activation, which was different from that observed in the head of the caudate. Specifically, the caudate body showed a signal that increased from baseline on all rounds, and this signal emerged later in time course than for the caudate head. Differences in the signal for correct and incorrect trials began to emerge as Round increased. Late in the time course, the signal became greater for incorrect trials compared to correct trials, while on Round 3, the signal rose more quickly for correct than incorrect trials. Previous work has also suggested that the activity of the body of the caudate is distinct from that of the caudate head (Seger & Cincotta, 2005). This work showed an increased response in the caudate body over many trials of training on a classification task with feedback, with greater activation on correct than incorrect trials late in training. It is possible that the results of the present study show the beginnings of performance-related differences in the signal produced by the caudate body, which might continue to evolve

with more training. However, further research will be required to pinpoint the functional role of this region relative to the caudate head.

Unlike most studies investigating feedback-based learning, this experiment utilized a task involving declarative memory acquisition. This is not to say that nondeclarative learning was not also taking place; however, the task differed in many respects from tasks designed to primarily engage nondeclarative memory and was instead similar to tasks designed to engage declarative memory. For example, rather than requiring many trials for correct responses to slowly be acquired, as is typical of nondeclarative learning, learning occurred quickly in this experiment, with performance significantly above chance after only one presentation of the 60 trials. Additionally, participants were quite aware of what it was they were learning, whereas for many nondeclarative learning tasks, performance improves without conscious knowledge of what marks the right response. The results of this experiment argue against a competitive systems view of "declarative" and "nondeclarative" brain systems, with activity in the hippocampus and striatum showing an inverse relationship. During each round of this experiment, the hippocampus showed a significant, albeit decreasing, response on each trial. Decreasing responses in the hippocampus are not unusual, and have been attributed to activity during "rest" that is higher than activity during the task itself (Stark & Squire, 2001). Indeed, during this task, the participants may well have been engaging in associative strategies to try to remember the word pairs during the intertrial intervals, causing hippocampal activation, which may have then paused during the beginning of the trial proper as they responded to the new trial and waited for the feedback, resulting in a decreasing response. Meanwhile, after responding very little on Round 1, the caudate was active on Round 2 and 3. If activity in the two brain regions were inversely related, we would expect that either the caudate would not be engaged at

all in this experiment, or that the hippocampus would no longer be engaged once the striatum was recruited. Instead, it appears that both regions can be engaged simultaneously. There is some evidence that the hippocampus, like the caudate, was modulated by feedback valence and round, in that the recovery of the hippocampal response back to baseline is increasingly slower following negative feedback on Rounds 2 and 3. However, this does not necessarily mean that the caudate was causing this modulation, and so the results of this experiment leave open the possibility that the two brain regions act independently.

It is important to note that the caudate does not appear to be processing the information provided by the feedback about which of the response options is correct. If it were, then either it would be active during Round 1, or no learning would take place during Round 1. However, performance is significantly above chance on Round 2, despite the caudate's lack of response on Round 1. Therefore, other brain regions must be able to process the information provided by the feedback about which response is correct to support learning of the appropriate word pairs. Indeed, brain regions that have been found to show subsequent memory effects in other declarative memory tasks, such as the left prefrontal cortex and left fusiform gyrus (Brewer et al., 1998; Wagner et al., 1998), showed a similar effect during Round 1 of this task, whereas the caudate did not. Caudate activation, it appears, is not necessary for feedback-based learning to occur, at least in some situations.

What role, then, does the caudate play in feedback-based learning, if it is not necessary to process the information it provides about which response is the correct one? The results from this experiment suggest that it is processing a different kind of information, that is, information about task performance. The goal of the participants in this experiment is not to obtain positive feedback *per se*, but to learn the correct word pairings. Positive feedback seems to yield a

reward response and negative feedback a punishment response only inasmuch as they are related to whether the goal is being achieved. Feedback in this experiment only carries such goal-related information on Rounds 2 and 3, and likewise, the caudate is only active during these rounds.

Importantly, only two response options were present on each trial of this experiment, which meant that the amount of information provided by feedback on Round 1 was equated. Each type of feedback was equally useful in enabling participants to reach the goal of task mastery. It is possible that if the two types of feedback were differentially effective in facilitating goal achievement, feedback could elicit caudate activation even during an initial round of trials. This possibility is addressed in Experiment 2.

3.0 EXPERIMENT 2: THE EFFECT OF THE INFORMATIONAL VALUE OF FEEDBACK ON CAUDATE ACTIVATION

Previous work indicates that the caudate is involved in processing reward-related information, but that its activity is very context-dependent (e.g., Delgado et al., 2005a; Delgado et al., 2004; Hikosaka et al., 1989; Kawagoe et al., 1998). In some cases, the same outcome can produce differing amounts of caudate activity, depending on how the outcome is interpreted. For example, when monetary wins and losses were thought to be contingent upon a response, the caudate showed much stronger activation than when the outcomes were thought to be independent of a response, even though in both cases the outcomes were actually fixed (Tricomi et al., 2004). In an interactive game, monetary wins and losses that resulted from decisions made by a partner of neutral moral character produced a differential response in the caudate nucleus, whereas identical outcomes thought to be from decisions made by a partner with praiseworthy moral character did not (Delgado et al., 2005a). Another study found that earning no money when the alternative was losing money produced a greater response in the caudate than it did when the alternative was winning money (Nieuwenhuis et al., 2005). Importantly, this study demonstrates that the absolute value of the outcome is less important for driving caudate activation than the outcome's value relative to other possible outcomes.

Experiment 1 indicated that when positive and negative feedback provide equal amounts of information and do not indicate task success, little activation is observed in the caudate nucleus. It did not address, however, whether the caudate would become active if positive and

negative feedback provided differing amounts of information. In this context, positive and negative feedback would have different values relative to one another, which could drive caudate activation. In Experiment 2, the same paired-associate word-learning task that was used in Experiment 1 was used again, except that the amount of information provided by the feedback was manipulated by altering the number of response options.

In this task, participants choose which of several word options correctly forms a pair with the target word. With only two response options, positive and negative feedback provide equal amounts of information about the correct answer; if the feedback is positive, the response chosen is correct, while if the feedback is negative, the response not chosen is correct. However, with more than two response options, positive feedback provides more information about the correct answer than negative feedback, because the correct answer cannot be deduced from negative feedback. Therefore, with more than two response options, positive feedback may be valued more highly than negative feedback, because it better helps the learner reach the goal of learning the correct pairs. Consequently, in this context, positive feedback may serve as a reward, and negative feedback as a punishment, which may produce caudate activation.

In addition, positive feedback would be less expected with more than two response options, and dopamine neurons in the midbrain, which project to the striatum, have been found to fire in proportion to how expected a reward is. Unexpected rewards produce the most firing, while firing in anticipation of an upcoming reward increases as rewards become more expected (Fiorillo et al., 2003). In fact, one can think of information as the deviation of an outcome from expectation; when an outcome is fully predicted, it provides no information, whereas if there is a low probability that any given action is the correct one, then positive feedback provides a large amount of information. If caudate activity were to track with information in this way, and each

response option has an equal probability of being correct, then receipt of positive feedback should produce more caudate activation with more response choices than with fewer. However, according to this account, negative feedback should produce more caudate activation with two choices than with more than two, since it is more expected with more possible options. Such a finding would deviate from the results obtained in Experiment 1, in which purely informational feedback did not produce significant caudate activation for either positive or negative feedback when there were two response options.

To examine how the informational value of feedback affects caudate activation, participants were scanned while performing the feedback-based word-learning task used in Experiment 1 with trials that had two or four response options. A single round of trials were scanned, and the correct responses followed the laws of probability, such that there was a 50% chance of choosing the correct answer with two response options and a 25% chance of choosing the correct answer with four response options. This allowed the amount of information carried by the feedback to be cleanly defined. Post-tests were used to assess memory and confidence immediately following the scan and after one week.

3.1 METHODS

3.1.1 Participants

Nineteen volunteers were recruited through posted advertisements and were paid $60 for their participation in the study. Data from two participants were excluded due to a program error and data from a third participant were excluded due to excessive head motion, leaving 16 participants in the analysis (11 women, 5 men; mean age ± SD, 22.9 ± 1.9). All participants

were healthy, right-handed adults. No one who participated in Experiment 1 was allowed to participate in Experiment 2. All participants signed an informed consent form according to the Institutional Review Board at the University of Pittsburgh.

3.1.2 Materials

As in Experiment 1, a 3 Tesla Siemens head-only scanner and standard radio frequency coil was used for all the MR scanning sessions. Stimulus presentation and behavioral data acquisition was controlled using "E-prime" software (Schneider et al., 2002) and the integrated function imaging system (IFIS) [Pittsburgh, PA].

Procedure:

Scan session:

Structural images were collected using a standard T1-weighted pulse sequence, in 42 contiguous slices (3.125 x 3.125 x 3.0 mm voxels) parallel to the AC-PC line. Thirty-eight functional images were collected in the locations of the middle 38 structural slices. They were acquired using a one-shot echo-planar imaging (EPI) pulse sequence [TR=2000 ms, TE=25 ms, FOV=20 cm, flip angle = 79°].

Behavioral paradigm:

Experiment 2 involved a paired associate word learning task, similar to the one used in Experiment 1, except that varying numbers of response options were presented and only the initial encoding round was scanned. Specifically, there were 60 trials with two options and 120 trials with four options. The experiment was designed such that participants were correct on

50% of trials with two options and 25% of trials with four options. This trial structure ensured that there were at least 30 "correct" trials for each trial type, while maintaining chance levels of performance.

The scan session consisted of 10 6-minute runs. Each trial was 20 s long. The same word lists used in Experiment 1 were used again, along with additional words subject to the same constraints (more words were needed for Experiment 2 than Experiment 1). The words used in the different conditions were matched on all parameters (e.g., word length, number of syllables, etc.). The procedure for each trial was the same as in Experiment 1, except that subjects were given 6 s to respond; this was necessary to allow enough time to read all four response options when present. Trials with two and four response options were randomly intermixed. Again the subjects were told to use the feedback to learn the correct word pairs.

Following the scan, participants took a post-test, using the E-Prime program. Instead of the multiple-choice format, the participants saw the target word and the word they picked on the corresponding trial during the scan, and were asked to respond as to whether the pairing was correct or incorrect. This procedure was used so that for all trial conditions, there were two possible responses allowing performance between conditions to be compared. No feedback was given. Pairs from all 180 trials were presented, in random order. As in Experiment 1, the post-test was self-paced, and participants made confidence judgments following each trial by choosing a number from 1-7 (1 = complete guess, 7 = completely sure). Participants returned to the lab and completed the post-test a second time approximately one week later (4-8 days following the scan session, with 14 of 16 participants returning after 7 days).

3.1.3 Data Analysis

Behavioral Data

Since accuracy during the scan is predetermined, the only useful behavioral data from the scanning session itself is the reaction time data. A two-tailed t-test was performed to see if there were reaction time differences between trial types. The behavioral data for the post-test were analyzed in terms of accuracy, reaction time, and confidence measures. D-prime scores were calculated for each subject for each trial condition. Two-tailed t-tests were used to determine whether performance exceeded chance for each trial type, and to determine whether performance differed between the trial types. Repeated measures ANOVAs were performed to see whether reaction time and/or confidence differed across conditions. Finally, a mixed model linear regression analysis of reaction time on confidence, with subject as a random factor, was performed for each post-test to determine whether differences in confidence judgments were associated with differences in reaction time.

FMRI Data

As in Experiment 1, the NeuroImaging Software package (NIS 3.5) was used to analyze the fMRI data. Images were reconstructed and corrected for subject motion with Automated Image Registration (AIR 3.08; Woods et al., 1992). Runs with head motion that exceeded 4 mm or degrees in any direction were not used in the analysis. The images were detrended to adjust for scanner drift within runs. The structural images of each subject were stripped to remove the skull and co-registered to a common reference brain (chosen from among the participants; (Woods et al., 1993). Functional images were transformed into the same common space, normalized by a mean scaling of each image to match global mean image intensity across

subjects, and smoothed using a three-dimensional Gaussian filter (8 mm FWHM) to account for anatomical differences between subjects. This set of data was then analyzed statistically.

A voxel-wise analysis was performed on the fMRI data, through a repeated-measures three-way ANOVA with subject as a random factor and trial type (number of response options), accuracy (correct vs. incorrect), and time period (2 s time periods T1-T10) as within-subject factors. Regions of interest (ROIs) were defined as regions with four or more contiguous voxels showing a significant effect. Additional follow-up statistical contrasts were performed on the data from each trial type individually.

A subsequent memory analysis was performed to investigate whether levels of activation during the scanning session predict accuracy on the post-test (Wagner et al., 1999). As in previous work (Wagner et al., 1998), trials for which the correct answer was remembered with high confidence on the post-test following the scanning session were compared to trials corresponding to subsequent incorrect responses. Trials were coded as "high confidence" if the confidence score was greater than or equal to 5 on the 7 point scale. A voxel-wise repeated measures ANOVA was performed on the fMRI data with subject as a random factor and post-test accuracy (high confidence correct vs. incorrect), and time (2 s time periods T1-T10) as within-subject factors. Additionally, since there was an *a priori* hypothesis that the caudate might show a subsequent memory effect, the fMRI data in this region were analyzed with respect to the post-test accuracy data.

3.2 RESULTS

3.2.1 Behavioral Results

Scanning session

Behavioral data analyses were performed on the data from the participants that were included in the fMRI analysis. During the scanning session, the average reaction time was 2409 ± 552 ms for the 2-choice condition and 3002 ± 852 ms for the 4-choice condition. The difference between the two conditions was significant ($t(15) = -5.7$, $p < 0.05$, two-tailed), which is to be expected since it takes longer to read through four response options than two.

Immediate Post-test

The behavioral data for the post-tests were analyzed in terms of accuracy, reaction time, and confidence measures. For the immediate post-test, the average d-prime score was 0.81 ± 0.7 for the 2-choice condition and 0.91 ± 0.85 for the 4-choice condition. This performance was better than chance for both conditions ($t(15) = 4.9$, $p < 0.05$ for the 2-choice condition; $t(15) = 4.3$, $p < 0.05$ for the 4-choice condition), indicating that learning occurred. The difference in performance between the conditions was not significant ($t(15 = -0.64$, $p > 0.05$).

The average reaction time on the immediate post-test was 3125 ± 742 ms (mean ± SD). The post-test reaction time data was coded with respect to the feedback the subject received on the corresponding trial during the scan (feedback type), the trial type of the corresponding trial (2 choices vs. 4 choices), and the post-test accuracy (correct vs. incorrect). As shown in Figure 10a, participants responded more quickly on accurate trials and on trials for which they had received positive feedback during the scan. A repeated measures three-way ANOVA indicated

that there were significant main effects of feedback type ($F(1, 18) = 5.5$, $p < 0.05$) and post-test accuracy ($F(1, 15.8) = 7.9$, $p < 0.05$), but no interactions.

The average confidence judgment score (out of 7) on the immediate post-test was 4.5 ± 1.0 (mean ± SD). The factors that acted to decrease reaction time also tended to increase confidence. A repeated measures three-way ANOVA indicated that there was a main effect of feedback type ($F(1, 16.4) = 29.4$, $p < 0.05$), a main effect of post-test accuracy ($F(1, 15.5) = 31.6$, $p < 0.05$), and a trial type X feedback type X post-test accuracy interaction ($F(1, 17.5) = 4.8$, $p < 0.05$). As shown in Figure 10b, confidence was higher for accurate trials and for trials on which the participants had received positive feedback during the scan. For the 4-choice condition, the effect of feedback type was present only for accurate trials. Finally, increases in confidence judgments were associated with decreases in reaction time; a mixed model linear regression analysis of reaction time on confidence, with subject as a random factor, showed this effect to be significant ($F(6, 2790) = 10.0$, $p < 0.05$).

Follow-up Post-test

For the follow-up post-test one week after the scanning session, the average d-prime score was 0.35 ± 0.44 for the 2-choice condition and 0.39 ± 0.49 for the 4-choice condition. This performance was better than chance for both conditions ($t(15) = 3.2$, $p < 0.05$ for the 2-choice condition; $t(15) = 3.1$, $p < 0.05$ for the 4-choice condition). The difference in performance between the conditions was not significant ($t(15) = -0.37$, $p > 0.05$).

The average reaction time on the follow-up post-test was 2495 ± 955 ms (mean ± SD). The effects observed on the immediate post-test are no longer observed on the follow-up post-test. A repeated measures three-way ANOVA indicated that there were no significant main effects, although there were significant interactions of trial type X feedback type ($F(1, 19.5) =$

5.1, $p < 0.05$), post-test accuracy X feedback type ($F(1, 15.7) = 6.1$, $p < 0.05$), and trial type X feedback type X post-test accuracy ($F(1, 15.9) = 5.8$, $p < 0.05$). For the 4-choice condition, reaction time did not differ based on feedback type, but was faster correct than incorrect trials. For the 2-choice condition, reaction time was fastest for incorrect trials for which participants received positive feedback during the scan, slowest for incorrect trials for which they received negative feedback, and in between for correct trials (Figure 11a). Since these effects no longer mirror the results found for the confidence judgments, they are difficult to interpret and may be the result of participants guessing very quickly on some trials and trying harder on others.

The average confidence judgment score on the follow-up post-test was 3.5 ± 1.3 (mean ± SD). A repeated measures three-way ANOVA indicated that there was a main effect of feedback type ($F(1, 15.3) = 8.5$, $p < 0.05$), a main effect of post-test accuracy ($F(1, 15.3) = 14.9$, $p < 0.05$), and a trial type X feedback type interaction ($F(1, 15.7) = 8.6$, $p < 0.05$). This pattern is remarkably similar to that observed for the immediate post-test. Confidence was higher for accurate trials on which the participants had received positive feedback during the scan, and this effect was greater for the 4-choice condition (Figure 11b). Finally, increases in confidence judgments were associated with decreases in reaction time; a mixed model linear regression analysis of reaction time on confidence, with subject as a random factor, showed this effect to be significant ($F(6, 2670) = 15.2$, $p < 0.05$).

3.2.2 fMRI Results

A voxel-wise ANOVA with subject as a random factor and trial type (2-choice vs. 4-choice), feedback (positive vs. negative), and time period (2 s time periods T1-T10) as within-subject factors was performed on the fMRI data. The resulting activation clusters showing a

Trial Type by Time effect, at a threshold of $p < 0.001$ and a contiguity threshold of four voxels (cf. Tricomi et al., 2006), are listed in Table 7 and shown in Figure 12. The most salient aspect of these results is that there is far greater activation in visual cortex for the 4-choice condition compared to the 2-choice condition; this makes sense given that there were two more words to process in the 4-choice condition. The caudate is noticeably absent from the list of regions identified as significant for this contrast.

The regions displaying a significant effect of Feedback Type by Time are listed in Table 8 and shown in Figure 13. Although a number of regions were identified, the montage of the activated regions shows that compared to the Trial Type by Time contrast, the effect of Feedback Type by Time was relatively weak. Again, the caudate was not identified as showing a significant effect for this contrast.

Figure 14 depicts the activation clusters that were identified as showing a three-way Trial Type by Feedback Type by Time interaction at a significance level of $p < 0.001$ and a cluster threshold of four voxels. The corresponding areas are listed in Table 9. The most common pattern of activation among these regions was for them to show greater activation following positive feedback than negative feedback for the 4-choice condition, with this effect either absent or smaller in the 2-choice condition. One such region was identified in the dorsal portion of the head of the right caudate nucleus. The time courses of activation are plotted in Figure 15. Contrary to the findings from Experiment 1, trials in both the 2-choice condition and the 4-choice condition elicited a hemodynamic response in the caudate. Only the 4-choice condition, however, differentiates between positive and negative feedback. Since previous work has found such differentiation 6-9 s following a feedback display (e.g., Tricomi et al., 2006), two-tailed t-tests were performed on the data from the corresponding time points in this experiment, T7 and

T8. No significant differences were found for the 2-choice condition, while a significant difference was found at T8 for the 4-choice condition ($t(15) = 3.3$, $p < 0.05$). Despite this differentiation, the hemodynamic response did not decrease below baseline in the way that is typical of "punishment" responses (e.g., Delgado et al., 2000; Tricomi et al., 2004).

The largest cluster and the one with the highest peak F-value that was identified as showing the three-way interaction was in the left prefrontal cortex (PFC), including activated voxels in both the dorsolateral PFC and the inferior frontal gyrus, as well as the insula. The significance threshold was increased until this region separated into regions of smaller than 100 voxels; the pattern of activation was the same, so the average time course over the entire activated region has been plotted in Figure 16. As in the caudate, there is no differentiation between trials with positive and negative feedback for the 2-choice condition, while there is a greater response to positive feedback than negative feedback for the 4-choice condition. Additionally, the signal following positive feedback in the 4-choice condition exceeded the signal for the 2-choice condition.

To confirm the results of the three-way interaction and to more closely investigate striatal areas that might be showing differential responses to positive and negative feedback in the two conditions, follow-up ANOVAs were performed on the data from the 4-choice condition and the 2-choice condition individually. The resulting activation clusters showing a Feedback by Time effect, at a threshold of $p < 0.001$ and a cluster threshold of 4 voxels, are listed in Tables 10 and 11 and displayed in Figures 17 and 18. A large region in the left DLPFC, IFG, and insula was found to show a significant effect for the 4-choice condition, with a greater response to positive than negative feedback; no such region was identified in the 2-choice condition, which is

consistent with the finding that this region shows a three-way interaction of trial type, feedback type, and time.

Significant clusters were identified in both the left and right caudate nuclei for the 4-choice condition, but not for the 2-choice condition. The location of these clusters was in the ventral head of the caudate, which more closely replicates the location identified in previous work than the more dorsal region identified as showing the three-way interaction in the preceding analysis (cf. Delgado et al., 2000; Tricomi et al., 2006). The time courses, too, show a response profile that is somewhat more similar to that shown in previous work, with a more defined response peak than the more dorsal region (Figure 19). Still, both regions show the same basic pattern of activation; although a response is elicited in both conditions, it only differentiates between positive and negative feedback in the 4-choice condition. Again, the difference in activation between conditions was compared at time points T7 and T8 with two-tailed t-tests; the differentiation was not significant for the 2-choice condition and was significant for the 4-choice condition at T8 for the left caudate ($t(15) = 2.3$, $p < 0.05$) and at T7 and T8 for the right caudate ($t(15) = 2.4$ for both time points, $p < 0.05$). It should be noted, however, that this differentiation appears to begin earlier than this time frame, and is apparent in the time course graphs at T6, which is the time point of peak activation. This, in addition to the fact that the response does not dip below baseline for negative feedback trials, may indicate that the negative feedback response is not analogous to the punishment responses observed in previous work (e.g., Delgado et al., 2000).

As in Experiment 1, the hippocampus/parahippocampal gyrus displayed a strong main effect of time (peak Talairach coordinates: -28, -15, -15 and 24, -18, -12 for left and right sides, respectively). The threshold was increased to isolate the hippocampal activation clusters from

contiguous activated regions, and is shown in Figure 20, along with the time courses of activation. The locations and pattern of activation are consistent with the results from Experiment 1, with the signal decreasing from trial onset, and recovering following the feedback display.

Table 12 lists the regions that showed a post-test accuracy by time effect on the immediate post-test. These regions are also displayed in Figure 21. Most of these areas showed greater activation when the trial would be remembered with high confidence on the post-test following the scanning session than when performance would be inaccurate, although there were two small regions showing the opposite effect. Among the regions displaying a subsequent memory effect were the left and right inferior frontal gyri, left hippocampus/parahippocampus, and left fusiform, which are regions that have previously been associated with subsequent memory (Brewer et al., 1998; Wagner et al., 1999). A similar analysis performed on the data coded by accuracy on the 1-week follow-up post-test revealed five regions showing a post-test accuracy by time effect ($p < 0.001$). These regions were the right parahippocampal gyrus (peak Talairach coordinates: 27, -21, -9), two small regions in the left middle temporal gyrus (-51, -48, 9 and -51, -69, 15), and left inferior parietal lobule (-65, -45, 30). Of these, only the cluster in the right parahippocampal gyrus displayed a clear subsequent memory effect with greater activation for trials that would be remembered with high confidence than subsequently incorrect trials. It is interesting that the region in the right parahippocampal gyrus was not identified as showing a subsequent memory effect for the immediate post-test, but did show such an effect for the follow-up post-test. This indicates that the role of this region in the formation of enduring memories was masked during the first post-test, when a greater proportion of answers were

remembered with high confidence; in other words, the activity in this region was only elevated for trials that would still be remembered with high confidence after a week's delay.

Since there was an *a priori* hypothesis that the caudate might show a subsequent memory effect, the post-test accuracy data (for trials from both the 2-choice and 4-choice conditions) were applied to the caudate activation clusters that displayed a significant feedback type by time effect for the 4-choice condition. For trials on which participants received positive feedback during the scanning session, caudate activation was greater when subsequent performance on the post-test would be accurate than when it would be inaccurate (Figure 22). In the left caudate, this effect was significant at timepoints T6 and T7 (i.e., 4-8 s after the onset of the feedback display; $t(15) = 2.7$ for T6, $t(15) = 2.2$ for T7, $p < 0.05$, two-tailed). The effect was not significant in the right caudate, nor were there any significant differences for trials on which negative feedback was presented during the scan (Figure 22).

3.3 DISCUSSION

In this experiment, the caudate responded to both positive and negative feedback during the first presentation of trials of a multiple-choice task involving learning arbitrary word pairs. This finding stands in contrast to the findings from the first round of trials in Experiment 1, in which the caudate did not show a rise in activation during the same task with only two response options. In fact, the response to positive feedback was actually more sustained than for Round 2 of Experiment 1; the signal does not return back to baseline until the last time point of the trial (12-14 s after the onset of the feedback display), compared to a return to baseline 4-6 s after the feedback presentation in Experiment 1. The presence of this signal in the caudate can be

explained in terms of the difference that feedback is playing in goal-directed behavior in the two experiments. In the first round of Experiment 1, positive and negative feedback helped learners reach their goal of learning the correct answers equally well. In contrast, in this experiment, positive feedback indicated which answer was correct, but in the 4-choice condition, negative feedback did not. Therefore, positive feedback may have been perceived as more rewarding than negative feedback when there were four choices. Potentially because of this, caudate activation was produced, which differentiated between the two feedback types in the 4-choice condition.

Despite this differentiation, it is interesting to note that the pattern of the response to negative feedback is unlike that observed in Experiment 1, and unlike the "punishment" response observed in previous work (e.g., Delgado et al., 2000). Specifically, although the signal following positive feedback is greater than following negative feedback, the negative feedback signal does not dip below baseline. In fact, it remains sustained above baseline, looking more like a "reward" response, than a "punishment" response, albeit a smaller one than for the positive feedback condition. Previous work has found the peak of the punishment response to be at least as great as that for the reward response, with the differentiation occurring later in the time course. In contrast, in this case, the peak response for positive feedback is greater than the peak for negative feedback. It is possible, then, that negative feedback was not interpreted as a punishment in this experiment, but rather as less of a reward compared to positive feedback.

As discussed with respect to Experiment 1, a sense of agency in determining the outcome of one's actions has been linked to an increase in caudate activation (Tricomi et al., 2004). Yet in this experiment, the sense of agency is minimal; as in the first round of Experiment 1, participants know that their responses are arbitrarily related to the feedback they receive, and that there is no way of knowing in advance which response is most likely to be correct. Although

this minimal contingency between action and outcome was sufficient to allow caudate activation in this experiment, the lack of real agency in determining the outcome may help explain why the caudate's response to negative feedback was uncharacteristic of a punishment response. Although negative feedback may have been the less desirable outcome, participants knew it was unavoidable, so it may not have acted as a true punishment in the way that negative feedback did on Rounds 2 and 3 of Experiment 1, when negative feedback was a direct indication of a failure of the learner to remember the correct answer.

The caudate did not differentiate between positive and negative feedback in the 2-choice condition. Each carried the same amount of information, and may have been interpreted as equally rewarding. What is more surprising is that the caudate responded at all in this condition, given the results from Experiment 1, in which the caudate showed no response for the same task with two choices. It seems that by simply intermixing trials with two response options with trials with 4, the interpretation of the feedback was different enough to produce activation in the caudate. Rather than viewing the 2-choice condition and 4-choice condition as fundamentally different, it seems that feedback in these two conditions was interpreted, by the caudate at least, to be part of a single set. Yet positive and negative feedback weren't interpreted categorically across the two conditions; negative feedback produced a greater response in the 2-choice condition than the 4-choice condition. It seems that the experimental paradigm was interpreted as a single task with four possible outcomes. Of these, negative feedback in the 4-choice condition was the only one which did not provide enough information to determine the correct word pair, and the response of the caudate in this situation was distinct from the response to the other three possible outcomes. It would be interesting to see how far the extent of this influence of the 4-choice condition on the caudate's activity in the 2-choice condition would go. If the

trials were blocked, rather than randomly intermixed, would there still be caudate activation for the 2-choice condition? What if all of the 4-choice trials were presented first, followed by the 2-choice trials? These questions remain open, and would further fine-tune our understanding of the influences of context on the caudate's activity.

In the 2-choice condition of this experiment, there is a 50% chance of receiving negative feedback and a 50% chance of receiving positive feedback; thus the two feedback types are equally expected. In the 4-choice condition, there is a 25% chance of receiving positive feedback and a 75% chance of receiving negative feedback. Therefore, overall, positive and negative feedback in the 2-choice condition are each more expected than positive feedback in the 4-choice condition and less expected than negative feedback in the 4-choice condition. Thus, if the caudate were coding specifically for the informational value of the feedback, defined as the degree to which the outcome deviates from expectation, then its response should be greatest in response to positive feedback in the 4-choice condition and least in response to negative feedback in the 4-choice condition, with a response to both negative and positive feedback in the 2-choice condition in between. This is not quite the pattern observed in the caudate. Rather, the response to all conditions other than negative feedback in the 4-choice condition is about the same, with that condition showing a smaller response. Instead of tracking with the informational value of the feedback, defined in this way, the caudate's outcome response seems to be tracking with the probability of knowing the correct answer on the next presentation of the identical trial, given perfect memory. Other brain regions, however, do show a pattern in which the positive feedback response is even greater in the 4-choice condition than the 2-choice condition, and are discussed in the General Discussion chapter of this thesis.

In this experiment, when positive feedback was delivered, the caudate was more active for trials that would subsequently be remembered than for trials that would subsequently be forgotten. This suggests that caudate activation may indeed facilitate feedback-based learning. There were, however, no such differences for trials in which negative feedback was delivered. Interestingly, positive feedback led to higher confidence ratings on the post-tests than negative feedback, which is consistent with the conclusion that the caudate strengthened associations based on positive feedback more than negative feedback. It could be that the caudate only facilitates appetitive learning, or these results could be due to the fact that the response to negative feedback was not truly an "error" response.

Either way, the caudate is not acting in isolation in its involvement in feedback-based learning. Once again, the hippocampus/parahippocampus showed a strong main effect of time, indicating that responses can be elicited in the MTL and striatum simultaneously. As in Experiment 1 and previous work, activation in regions such as the left inferior frontal gyrus, left medial temporal lobe, and left fusiform gyrus were found to predict which trials would subsequently be remembered on the post-test following the scanning session. Interestingly, the right parahippocampal gyrus showed a subsequent memory effect with relation to the post-test occurring one week after the scanning session, but not the post-test occurring immediately after it. This is consistent with the results of Wittman, et al. (2005), who found that reward-associated cues led to activation of the midbrain and striatum and improved memory performance, but this effect was also only apparent after a long delay. The same effect did not hold for "intermediate-term memory," or in other words, of memory that occurs after the period of active maintenance of working memory, but which is still in a labile state (Eichenbaum, 2002). During this time, it is expected that the memory should be dependent on the MTL but not yet consolidated

(McClelland et al., 1995). Wittman, et al. (2005) argue that the reward-related activation helped to determine which memories the MTL would successfully consolidate into long-term memory and which would instead decay. Likewise, even low levels of activation in the MTL might have been enough to support intermediate-term memory in the present experiment, whereas only memories associated with higher BOLD signal during encoding were resistant to decay and were successfully consolidated into long-term memory.

In summary, the results of this experiment build upon the results of Experiment 1 by demonstrating the strong role that context plays in guiding feedback-related caudate activation. By simply including a condition with four response options as well as two, the caudate was engaged in all conditions, whereas this was not the case when only two response options were present in Experiment 1. The critical feature present in Experiment 2 was that positive and negative feedback carried different amounts of information, and thereby were perceived to aid learners in reaching their goal of task mastery to different degrees. In the general discussion that follows, a more thorough discussion of the factors influencing feedback-related caudate activity is presented.

4.0 GENERAL DISCUSSION

In this discussion, I will integrate the results from the two experiments presented in this thesis in order to draw conclusions about the nature of feedback-related caudate activation, as well as the role that other regions may play in feedback-based learning. In addition, I will discuss potential interactions between the reward system and the declarative memory system.

4.1 THE NATURE OF FEEDBACK-RELATED CAUDATE ACTIVATION

The signal produced in the caudate nucleus during feedback-based learning appears to reflect both activity occurring prior to the feedback display (as indicated by the initial rise of the response) and activity that results from processing the feedback itself. This latter component of the signal seems to reflect feedback valence, because depending on the outcome, the signal will follow one of two general trajectories. The "reward" response is characterized by a sustained signal following an outcome of positive value, whereas the "punishment" response is characterized by a sharp decrease in signal following a negative outcome. Factors that affect the initial rise, the "reward" response, and the "punishment" response will each be considered in turn.

4.1.1 The Rise

Unlike the activity of midbrain dopaminergic neurons (Fiorillo et al., 2003), the magnitude of caudate activation prior to outcome delivery does not appear to be proportional to either the probability of an upcoming reward or the uncertainty about the valence of the upcoming outcome. If it were proportional to the probability of an upcoming reward, then the signal in the caudate nuclei would have had to have been greater for the 2-choice condition in Experiment 2 (for which the probability of positive feedback was 50%) than the 4-choice condition (for which the probability of positive feedback was 25%). However, the magnitude of the caudate signal rise did not differ between these two conditions. Additionally, for Experiment 1, the probability of positive feedback is highest during Round 3, when the most learning has occurred, but the magnitude of the signal rise is actually attenuated for this round compared to the signal during Round 2. If the signal rise were proportional to the uncertainty about whether the outcome would be positive or negative, then it would have been greatest during Round 1 of Experiment 1 (in which the probability of positive feedback is 50%, and therefore uncertainty is maximal), yet during this round the signal was actually weakest. Likewise, uncertainty was maximal for the 2-choice condition of Experiment 2, yet the magnitude of the signal rise did not differ between the 2-choice and 4-choice conditions.

The magnitude of the initial rise in caudate signal prior to the presentation of the feedback display instead seems to reflect the degree to which the individual "cares" about the valence of the impending outcome. Two factors seem to be critical in determining this parameter. First, we know from a previous study that rewards and punishments in and of themselves are not enough to trigger caudate activation; instead, they must be perceived as the consequence of an action (Tricomi et al., 2004). In that study, an arbitrary linking of responses

and outcomes was enough for caudate activation to be elicited; it was only when outcomes were perceived to have no relation at all to responses that caudate activity was diminished. This minimal sense of agency in determining the outcome was present even on the first presentation of each trial in the experiments presented here. The fact that caudate activation was elicited in Experiment 2 again shows that only an arbitrary relationship between action and outcome is necessary for caudate activation to be elicited. On the other hand, that a robust signal in the head of the caudate was not present during Round 1 of Experiment 1 shows that this minimal sense of agency is not sufficient to drive feedback-related caudate activation in this region. Still, this does not rule out the possibility than an increased sense of agency on Rounds 2 and 3 resulted in increased caudate activation; it merely suggests that a sense of agency is not the sole factor in modulating the magnitude of the caudate signal rise.

The experiments presented here indicate that a second factor is also critical to the activation of the caudate nuclei: the potential outcomes of the action must differ in their value to the individual. In the case of Round 1 of Experiment 1, the task cannot be learned without paying attention to whether the feedback is positive or negative, and significant learning does occur. In this sense, the learners do care about the feedback. Yet apparently information that can lead to task success is not enough to activate the caudate if it is equivalent on every trial; rather than treating both positive and negative feedback as rewarding, the caudate was not strongly engaged. Both positive and negative feedback provided equal amounts of information about which answer was the correct one, and so both were equally valuable in helping learners achieve their goal of learning the correct answers. Because of this, learners may not have been invested in which type of feedback they received; they may not have felt that their response choice "mattered," and so they may not have cared about the valence of the outcome of their

response. It would be interesting to observe the response in the caudate if trials like these were intermixed with trials in which no feedback was provided. In this case, both positive and negative feedback might act as a "reward," since both provide useful information, while an outcome of no feedback might act as "punishment." In anticipation of these differential outcomes, there might be a rise in caudate activation.

On Rounds 2 and 3 of Experiment 1, positive and negative feedback took on a different meaning from its meaning on Round 1. It provided a new type of information—information about whether learners were achieving their goal of learning the correct word pairs. This information differed in its potential reward value; positive feedback indicated task success, while negative feedback indicated an error in performance. As such, participants were highly invested in the outcome valence, and correspondingly, a signal was produced in the caudate nuclei.

In Experiment 2, positive and negative feedback had different relative values in the 4-choice condition, since positive feedback better helped the learner achieve the goal of learning the correct word pairs. This was enough to produce a rise in caudate signal prior to the outcome display. Interestingly, this effect generalized to the 2-choice condition as well. This leads to the question of whether expectation-based caudate activity changes on a trial-by-trial basis, or whether it always generalizes across trials. There is some evidence that the caudate is able to modulate this activity at the trial level. When trials involving monetary rewards and punishments thought to be contingent upon a button press were randomly intermixed with trials where no contingency was present, only the trials with perceived contingency elicited caudate activation (Tricomi et al., 2004). In that experiment, however, participants may have viewed these conditions as two separate, though interspersed, tasks, whereas trials in the 2-choice and 4-

choice conditions in Experiment 2 may have all been considered to be instances of the same task, leading to a generalization in expectancy-based activation.

Over the course of learning, the signal in the caudate tends to decrease (Delgado et al., 2005b; Haruno et al., 2004; Law et al., 2005; Pasupathy & Miller, 2005). This can be explained in terms of the expected value of the feedback to the learner decreasing. If the learner is already sure of the correct response, then positive feedback will not be useful, since it will not provide any new information. In Experiment 1, this effect is exhibited on Round 3, which shows attenuation in signal rise compared to Round 2, and performance is also higher on Round 3. Furthermore, the initial rise in signal displayed during Round 2 of Experiment 1 appears to be weaker than the signal rise displayed in Experiment 2 (cf. Figures 5 and 19). This may be due to the fact that learners were already performing well above chance on Round 2 of Experiment 1, so positive feedback may have already been more expected than negative feedback. In contrast, overall there was a smaller chance of obtaining positive than negative feedback in Experiment 2.

As mentioned above, this effect stands in contrast to findings from electrophysiological recordings from dopaminergic neurons in the midbrain in monkeys, which indicate that these neurons increase their firing prior to outcome delivery as a function of the probability of reward (Fiorillo et al., 2003). A key difference between the experimental paradigms in these studies and the one used in the experiments presented here is that the cues that predict reward in the monkey studies are conditioned; that is, the animal learns over many presentations that the cue predicts the reward. In contrast, in the experiments presented here, participants know from the outset that they will receive feedback at the end of each trial. It may be that such top-down knowledge of the task structure does not affect midbrain neuron firing in the way that slowly conditioning a cue-outcome expectation does.

Some speculation is warranted on what purpose the initial rise in caudate activation serves. Lauwereyns, et al. (Lauwereyns et al., 2002a; Lauwereyns et al., 2002b) suggest that anticipatory activation in the caudate may exert top-down control on stimulus processing, increasing the activity level in perceptual regions involved in the detecting the reward signal. Similarly, in the experimental paradigm used in the experiments presented here, the pre-outcome signal in the caudate may act to focus the individual's attention on the rewarding or punishing aspects of the upcoming outcome.

4.1.2 The "Reward" Response

Following the presentation of an outcome of positive value, there tends to be a sustained response in the caudate nucleus. This occurs if the outcome is better than the alternatives. However, this response does not occur if there was no rise in activity prior to the outcome display, indicating that it is dependent upon the same factors that produce that initial rise. For example, there is no increase in signal prior to the feedback presentation during Round 1 of Experiment 1, and there is also no increase in signal following positive feedback. Similarly, just as the initial rise exhibited in Round 2 of Experiment 1 was not as great as the rise for Experiment 2, the response to positive feedback is also less sustained than for Experiment 2 (cf. Figures 5 and 19). This suggests that when a positive outcome is expected, it produces not only a small anticipatory response, but also a small outcome response. Conversely, when an action has the potential to produce a positive outcome, but such an outcome is unlikely, there appears to be a larger anticipatory response and also a larger response following the outcome if it does turn out to be positive.

In Experiment 2, the response following positive feedback was greater on trials associated with correct post-test performance than trials associated with incorrect post-test performance. This suggests that this "reward" signal can serve to facilitate learning. This is not to say that it is the only region involved in feedback-based learning, as there were a number of brain regions which showed subsequent memory effects. Nor is the caudate necessary for feedback-based learning to occur; the fact that significant learning occurred during Round 1 of Experiment 1, when there was no activation in the caudate, indicates otherwise. Regions outside of the brain's reward system may perform such functions as using the feedback to determine which response option is correct and forming a memory of the correct answer, while the caudate may interact with these regions, potentially through the attentional system. Perhaps the reward response in the caudate serves as a "priority signal," leading to better processing of reward-related information than neutral information.

4.1.3 The "Punishment" Response

The signal in the caudate nuclei tends to display a sharp decrease from the response peak following outcomes of negative value. The more sharply negative the slope of this decrease, the stronger in magnitude this "punishment" response is. Often, the signal decreases below baseline before returning to its resting state. Unlike the reward signal, the magnitude of the punishment "dip" is not closely linked with the magnitude of the signal prior to feedback delivery. This is seen in Experiment 1, in which there is even a small dip on Round 1, despite a flat response prior to outcome delivery. There was also an especially large dip on Rounds 2 and 3, even though there was a relatively small anticipatory rise. One reason for this pronounced dip may have been that when positive outcomes are expected, negative outcomes are unexpected, and just as

unexpected rewards result in a large reward response, so too may unexpected punishments result in punishment responses of greater magnitude.

A second reason for the pronounced dip may have been the strong sense of agency in determining the outcome during Rounds 2 and 3 of the first experiment. Negative feedback in that situation indicated that participants had made a true error, as opposed to an incorrect uneducated guess. This stands in contrast to the situation in Experiment 2, when for the 4-chocie condition, the caudate signal following negative feedback did not rise to the same level as the reward response, but also showed an uncharacteristically sustained response for a "punishment" response. In fact, this response looks rather similar to the response to positive feedback during Round 2 of Experiment 1 (cf. Figures 5 and 19). Added to this is the fact that the response following negative feedback in the 2-choice condition of Experiment 2 did not differ at all from the response to positive feedback; both responses look like sustained "reward" responses. Therefore, it appears that negative feedback in this experiment was not interpreted as a punishment, but rather as a reward that varied in magnitude depending upon the amount of information it provided (i.e., small for the 4-choice condition and large for the 2-choice condition). Presumably, if the brain were imaged during a second round of trials in Experiment 2, the caudate's response would look more like the responses displayed for Rounds 2 and 3 of Experiment 1. For the two-choice condition especially, negative feedback would likely produce a "punishment" response rather than a "reward" response. For trials in the 4-choice condition, however, negative feedback would not necessarily indicate an error, but might simply mean the learner made a second incorrect guess; therefore, the magnitude of the dip might be expected to be less strong than for the 2-choice condition.

What might be the purpose of the punishment signal? One possibility is that it is a part of an error detection system. Another possibility is that the punishment signal might act to keep the dorsolateral striatum from strengthening a stimulus-response bond. One way it could do this is through its projections back to the midbrain. The caudate projects to cells in the substantia nigra that are intermixed with cells that project to the dorsolateral striatum (Haber et al., 2000). It is possible that the decrease in caudate activity could serve to decrease the firing of dopaminergic cells in the midbrain. This might aid in preventing the reward system from strengthening S-R associations that lead to a negative outcome.

4.2 FRONTAL CONTRIBUTIONS TO FEEDBACK PROCESSING

The pattern of activation shown in the caudate nuclei in the studies presented here suggests that they are subject to top-down control. The activity levels reflect knowledge of such high level notions as what the meaning of the feedback is in relation to the task at hand. As the input unit of the basal ganglia, the striatum receives a wide range of converging input from the cortex (Bar-Gad et al., 2003). Several regions in the frontal cortex have been linked to reward processing and reward-related learning (Elliott et al., 2000a; Elliott et al., 2000b; Fellows & Farah, 2005; Holroyd & Coles, 2002; O'Doherty et al., 2001). The dorsolateral prefrontal cortex (DLPFC) is involved in executive function, such as working memory, and may be involved in reward-based learning tasks to the extent that they require such processes (Boettiger & D'Esposito, 2005; Fellows & Farah, 2005). The ventromedial prefrontal cortex (VMPFC) and orbitofrontal cortex have been shown to be especially important for reversal learning tasks in which stimulus-response-outcome contingencies are changed and participants must adjust their

responses accordingly (Fellows & Farah, 2005; O'Doherty et al., 2003; O'Doherty et al., 2001; Rolls, 1999). Finally, the anterior cingulate cortex (ACC) has been suggested to be involved in error detection or response conflict (Holroyd & Coles, 2002; Holroyd et al., 2004; van Veen et al., 2004). The DLPFC, VMPFC, and ACC all project to the striatum (Haber et al., 2006), and they also receive input from the basal ganglia through basal ganglia-thalamocortical pathways (Middleton & Strick, 2002). Thus, these regions are in a good position to interact with the striatum during feedback-based learning. In this section, the patterns of activation displayed in these frontal regions in the experiments presented here will be discussed.

4.2.1 Lateral Prefrontal Cortex

An expansive region which included voxels in the dorsolateral and inferior frontal portions of the PFC, as well as the adjoining insula, showed an interaction of round, feedback type, and time in Experiment 1. A similar region showed an interaction of trial type, feedback type and time in Experiment 2. At first glance, the pattern of activation in this area across the two experiments appears contradictory. During Rounds 2 and 3 of Experiment 1, the signal is greater following negative feedback than positive feedback, but for the 4-choice condition of Experiment 2, the signal is greater following positive feedback than negative feedback (cf. Figures 7 and 16). How can this apparent discrepancy be explained? This region does not seem to be coding for the valence of the outcome (i.e., whether it is positive or negative), but instead seems to reflect the amount of information provided by the feedback, regardless of its sign.

As discussed previously, as learning progresses, positive feedback carries less and less information as the learner becomes sure of the correct answers, making positive feedback redundant. Reflecting this, the signal following positive feedback decreases with increasing

round in Experiment 1, so that there is only a small response to positive feedback on Round 3 (Fig. 7). Conversely, negative feedback continues to provide information, since it indicates an error, and likewise, the signal following negative feedback remains high for all rounds. For Experiment 2, positive feedback in 4-choice condition provides enough information to determine the correct answer, while negative feedback in the 4-choice condition only provides enough information to eliminate one of the four response options. Again the left PFC reflects this pattern with a smaller response to negative feedback than positive feedback in this condition. For the 2-choice condition, in which both positive and negative feedback provide enough information to determine the correct answer, there is no such differentiation in the signal in the left PFC. Furthermore, since positive feedback is less expected in the 4-choice condition than the 2-choice condition, it can be thought of as more informative in this instance since the outcome deviates more from the expected outcome. Indeed, the signal in the left PFC fits with the explanation that it may reflect this informational value of feedback, since the response to positive feedback is greater in the 4-choice condition than the 2-choice condition.

Such a pattern of activation does not necessarily mean that the specific role of the left PFC is to code for the amount of information provided by feedback. Instead, it could simply reflect the increased processing demands that occur when more information is available. For example, the dorsolateral PFC has been associated with a role in executive control processes, such as working memory, and it may be that with more information, there is an increased load on this system as it updates memory representations.

4.2.2 Anterior Cingulate Cortex

The dorsal anterior cingulate and medial frontal gyrus have been suggested to be involved in detecting errors, either through self-detection or from negative feedback (Holroyd & Coles, 2002; Holroyd et al., 2004). Others have argued that this region is not involved in error detection *per se*, but rather that it increases in activity as conflict between incompatible processing streams increases (Botvinick et al., 2001; van Veen et al., 2004). In Experiment 1, a large activation cluster was identified spanning across voxels in the medial frontal gyrus, superior frontal gyrus, and dorsal anterior cingulate that showed greater activation to negative feedback than positive feedback, and a smaller cluster in the anterior cingulate displayed a round by accuracy by time interaction, with an increased difference in signal between trials with negative and positive feedback as the round increased. This pattern of activation could be consistent with a role for this region in error detection. However, in Experiment 2, a similar activation cluster in the medial frontal gyrus and superior frontal gyrus showed a three-way interaction of trial type, feedback type, and time, with a greater response to positive feedback than negative feedback in the 4-choice condition, but not the 2-choice condition. This finding conflicts with an interpretation of a specific role of this region in error detection. It is more consistent with the idea that this region may become more active with increasing conflict. There is increased conflict when negative feedback indicates a true error, as in Rounds 2 and 3 of Experiment 1, than when it simply indicates an incorrect guess, as in Round 1. Likewise, positive feedback may cause more conflict when it is less expected, as in the 4-choice condition, than when positive and negative feedback are equally expected.

4.2.3 Ventromedial Prefrontal Cortex

The VMPFC is another region that has been implicated in processing negative feedback. Neuropsychology studies indicate that patients with VMPFC damage are impaired at trial-and-error learning, and that this impairment is caused by a specific deficit in the ability to learn from punishment, while the ability to learn from reward remains intact (Fellows & Farah, 2005; Wheeler, 2006). This idea would suggest that the VMPFC should show a difference in the signal following negative and positive feedback in the experimental paradigm utilized in the experiments presented here. However, no such result was found. Although this null finding should be interpreted with caution, due to the fact that this region is particularly prone to susceptibility artifact, a region in the VMPFC was identified as displaying a subsequent memory effect, with greater signal for trials that would be remembered with high confidence than trials for which performance would be incorrect. This does suggest some role of the VMPFC in feedback-based learning. It is consistent with the recent finding that in addition to reflecting knowledge of higher order task structure, the VMPFC displays an "update signal" that reflects the change in the probability that a certain choice will be correct (Hampton et al., 2006). It may be that the VMPFC plays a role in learning from both positive and negative feedback, but that other regions can compensate and allow learning from positive reinforcement to proceed if the VMPFC is damaged, whereas there may not be similar compensatory mechanisms for learning from negative feedback.

4.3 LEARNING IN A FEEDBACK-BASED DECLARATIVE MEMORY TASK

4.3.1 Medial Temporal Lobe Contributions

Neuropsychological work demonstrates the crucial importance of the MTL in forming arbitrary associations, like those that were learned in the experiments presented here (Squire, 2004; Zola & Squire, 2000). Yet, this region showed a decreasing response profile in both Experiment 1 and Experiment 2. Previous work has also found that MTL activation tends to be below baseline when that baseline is unconstrained (Stark & Squire, 2001). These decreases can be avoided if a baseline task is utilized that disrupts MTL recruitment (Stark & Squire, 2001). However, this would seem to change the nature of the task, eliminating not only task-unrelated thought that might be associated with activity in the MTL, but also reducing the amount of time when the MTL can be recruited to aid in memory formation. Despite the negative response profile, there is evidence that the MTL is playing a role in memory formation in the experiments presented here. In addition to showing a main effect of time, the MTL displayed subsequent memory effects. In particular, in Experiment 2, a region in the left hippocampus and parahippocampal gyrus showed greater activation for trials on the immediate post-test for which participants responded correctly with high confidence compared to incorrect trials. Additionally, a similar effect was found for the delayed post-test in the right parahippocampal gyrus, and the location of the region showing this effect was quite close to the region identified as showing a main effect of time (peak coordinates: 27, -21, -9 and 24, -28, -12, respectively). Therefore, although the MTL seems to display a high level of tonic activity, which then tends to decrease when attention is focused on a learning a simple arbitrary association, this relatively low level of activity appears to still support memory formation. It should be noted that arbitrary word pairs

are much less rich in detail than the more complex associations that we make every day, and indeed, the hippocampus has been noted to be especially important for memory of rich configural associations, such as those required in spatial navigation (Maguire et al., 1997). The lack of detail of the word pairs may cause them to activate the MTL at a relatively low level, but this low level of activation could still be critical to memory formation and cause amnesics with MTL damage to have great difficulty in making arbitrary word pair associations.

Law, et al (2005) performed a study which also involved an arbitrary multiple choice task with feedback. Participants learned to associate kaleidoscope images with one of four spatial locations. The effect of positive versus negative feedback on brain activation was not examined, but instead the authors focus on the effects of memory strength on brain activation. Participants were trained on sets of trials with 4-12 new associations to be learned, and each set was repeated until it was well learned before moving on to a new set. MTL activation was found to increase with memory strength, while activation in other regions, including the DLPFC and right caudate nucleus, were found to decrease once an association was well-learned. Early indications of the latter effect were observed in Experiment 1 of the present work, but the MTL did not show increases in activity over the three rounds of trials, indicating that more repetitions may be necessary for such an effect to come about.

4.3.2 Striatal Contributions

Brain regions involved in the formation of declarative memories, including the hippocampus/parahippocampus, left fusiform gyrus, left inferior frontal gyrus were recruited in these experiments. This was true irrespective of activation of the striatum; similar regions showed a subsequent memory effect for both Round 1 of Experiment 1, in which no caudate

activation was present and for Experiment 2, in which the caudate was recruited. Therefore, it appears that the engagement of declarative memory does not preclude the recruitment of the striatal reward system, nor does engagement of the striatum prevent the engagement of the declarative memory system.

A more difficult question is whether the engagement of the striatum contributes to learning in a declarative memory task. Previous studies have suggested that the striatum does contribute to learning in situations which do not robustly engage the declarative memory system (Knowlton et al., 1996; Poldrack et al., 2001; Poldrack et al., 1999; Tricomi et al., 2006). Therefore, the striatal system does not seem to be simply registering the presence of response-contingent rewards or punishments, but seems to be using that reward-related information to form associations.

One test of this would be to determine whether it would be possible to eventually master the task utilized in the experiments presented here without relying on the declarative memory system. In all probability, the sheer number of trials would be prohibitive. One neuropsychology study did find that two patients with MTL lesions were able to slowly master an 8-pair object discrimination task, in which they learned on the basis of feedback which object from each pair was the correct choice. However, it took over a thousand trials over many weeks of training for task mastery to occur (Bayley et al., 2005). Scaling up this type of task to mastering 60-180 distinct trials, rather than 8, would be impractical, if not impossible. Nevertheless, the study indicates that a task that normally relies on declarative memory can still be learned without a properly functioning declarative memory system, so presumably a scaled-down version of the feedback-based paired associate word-learning task used in the experiments presented here would also be able to be acquired in a similar manner.

In the absence of a declarative memory system, the striatum is capable of the slow formation of stimulus-response-outcome contingencies. Does this preclude the striatum from acting to strengthen memory representations over much shorter time periods when the declarative memory system is intact? The fact that the caudate showed a subsequent memory effect in Experiment 2 for the immediate post-test (i.e., after only a single presentation of each trial) suggests otherwise. There are several possible mechanisms by which caudate activity could facilitate this type of learning, which are not necessarily mutually exclusive. One possibility is that activation in the caudate on correct trials, by strengthening S-R-O associations, could automatize the retrieval of correct responses. Even though the experiments presented here involved few repetitions of the same trials, there is some evidence that such an effect was already occurring. In Experiment 1, reaction time differences emerged between correct and incorrect responses on Round 3, which is the first round where effects of caudate activation should be observed, since no caudate activation was present during Round 1. Likewise, increased signal in the caudate on correct trials during Round 2 was associated with decreased reaction times for the corresponding trials on Round 3. Consistent with these findings, positive feedback in Experiment 2 led to decreased reaction times and increased confidence on the post-test, relative to negative feedback.

A second way that the striatum could influence memory acquisition is by interacting with the hippocampus in a way that facilitates memory for reward-relevant associations. This idea is supported by neural circuitry. Dopaminergic neurons in both the VTA and the substantia nigra project to both the hippocampus and the striatum (Lisman & Grace, 2005). Thus, information coded in the activity of dopaminergic neurons regarding reward-related information should be received by the declarative memory system as well as the striatum, and it is possible that such

dopaminergic input could facilitate memory. Indeed, long-term potentiation in the hippocampus in enhanced by activation of dopamine receptors, and administration of dopamine has been shown to improve memory performance in both rats and humans (Lisman & Grace, 2005). Therefore, the increased activation following correct feedback that was observed in the caudate for trials that would subsequently be remembered could be a reflection of increased firing of dopaminergic neurons on those trials, which would also have led to increased release of dopamine in the hippocampus, presumably enhancing declarative memory. Alternatively, it is possible that caudate activity actually affects the firing of dopaminergic neurons, since it projects back to the substantia nigra (Haber et al., 2000). For example, one possibility is that the caudate could integrate information about the presence of rewarding stimuli with contextual cues to either amplify or attenuate the activity in dopaminergic neurons, which then could modulate LTP in the hippocampus.

Interestingly, no subsequent memory effect was found in the caudate for negative feedback, which may indicate that the caudate is more effective at strengthening associations than forming negative associations. This fits with the finding that overcoming interference from negative responses seems to be dependent on an intact MTL, whereas memory is much less impaired if amnesics never generate an incorrect response. There may be a natural tendency for the cortex to reinforce incorrect associations via Hebbian learning when an incorrect response is made, which the MTL can counteract (McClelland et al., 1999). It is possible that over many repetitions with negative feedback, the striatum could overcome this tendency as well, but that the MTL can do so much more efficiently.

Performance-related feedback is a very common instructional tool, used in declarative and nondeclarative learning tasks alike (Goodman, 1998; Kluger & DeNisi, 1996; Schmidt &

Bjork, 1992). Because of the engagement of the reward system, feedback-based learning seems to be in some way fundamentally different from rote memorization, even if the content of what is learned is similar. What is it about the engagement of the reward system that might cause feedback to be so commonly used? One possibility is that compared to passively memorizing, feedback-based learning may increase motivation by causing learners to feel a stronger sense of agency in the learning process. There is an affective component to the "reward" or "punishment" of positive and negative feedback, which may make feedback-based learning seem almost like playing a game. If learners are invested in the outcome of their responses, they may attend more to the task than if the task involved no feedback. If this is the case, there might be practical implications of research on the neural basis of feedback-based learning. Feedback might be expected to have higher impact on motivation in situations where it elicits a response in the striatum, such as when the task includes instances with more than two response options. However, such conjectures are speculative and would require further investigation.

4.4 TRIAL SUMMARY

The presentation of the word display marks the beginning of each trial. This display acts as a cue that a response is required, and also acts as a temporal predictor that feedback will be presented shortly. At this point, if there is uncertainty about the value that feedback will provide to the learner, the caudate signal begins to rise. The learner chooses a response and waits for the feedback. During this interval, the rise in caudate activation may act to ready the attentional system to focus on the rewarding or punishing aspects of the upcoming outcome. If the outcome is better than expected, feedback provides information and the caudate signal remains relatively

elevated. This "reward" response appears to strengthen the particular stimulus-response-outcome associations involved on that trial; the greater the caudate signal is, the more solidified that knowledge of the correct word pair becomes, as evidenced by decreased reaction time or increased chance of an accurate response on the next trial involving the same stimuli. Eventually knowledge of the right answer becomes automatic. There becomes less uncertainty about the value of feedback, since positive feedback becomes expected. Therefore, the caudate tends to activate less strongly. The outcome matches expectation, providing no new information, and the caudate's response likewise remains relatively flat. If the outcome is worse than expected, feedback provides information by indicating an error, and the response decreases rapidly from its peak. This "punishment" signal may act to help prevent the strengthening of the incorrect S-R-O association; however there is little evidence to suggest that this signal facilitates performance on the next trial involving the same stimuli. Finally, the intertrial interval allows the caudate signal to return to baseline. During this time, however, other brain regions, such as the MTL, may still be engaged as learners invoke associative strategies to aid them in remembering the correct answer.

4.5 CONCLUSIONS

The work presented here has demonstrated that the caudate nuclei are involved in feedback processing when positive and negative feedback have differential value to the individual. Specifically, the value of feedback appears to be related to the degree to which it provides goal-related information, whether this be task-related information that can aid performance or performance-related information that indicates whether one's goal is being met.

This is not to say that recruitment of the caudate is necessary for feedback-based learning to occur. In the experimental paradigm used in this dissertation, by holding in mind the memory of the response and then incorporating the information provided by the feedback about whether that response was correct, learning could take place in much the same way as it would in a typical paired associate memory task. Nevertheless, due to the ways in which feedback provides information about the consequences of our responses, the caudate processes the reward value of feedback and may serve to reinforce memories or to prevent such reinforcement. Through its examination of the factors that influence the engagement of the reward system during feedback-based learning, this dissertation furthers our understanding of the neural mechanisms underlying learning and memory.

Table 1: Experiment 1 Accuracy and Reaction Time Across Experimental Rounds

	Accuracy (Mean ± SD)	t-value (compared to chance performance, *indicates p< 0.05, two-tailed)	Reaction time (ms; Mean ± SD)
Round 1	50% ± 0	N/A	2012 ± 330
Round 2	61.97% ± 7.74	t(10) = 5.13*	2133 ± 220
Round 3	75.15% ± 7.13	t(10) = 11.70*	2039 ± 214
immediate post-test	83.03% ± 8.06	t(10) = 13.60*	2854 ± 595
1 week post-test	71.67% ± 11.60	t(9) = 5.91*	3639 ± 752

Note: The 1 week post-test results do not include data from one participant, who failed to return for the post-test.

Table 2: Experiment 1 Reaction Time for Correct Versus Incorrect Trials

	Correct Trials (ms; Mean ± SD)	Incorrect Trials (ms; Mean ± SD)	t-value (correct vs. incorrect trials, *indicates p< 0.05, two-tailed)
Round 1	1997 ± 365	2027 ± 306	t(10) = -0.74
Round 2	2152 ± 259	2128 ± 255	t(10) = 0.30
Round 3	1986 ± 229	2203 ± 189	t(10) = -5.06*
immediate post-test	2803 ± 595	3234 ± 602	t(10) = -3.32*
1 week post-test	3501 ± 688	4024 ± 963	t(9) = -3.21*

Note: The 1 week post-test results do not include data from one participant, who failed to return for the post-test.

Table 3: Experiment 1 Confidence for Correct Versus Incorrect Trials

	Correct Trials (Mean ± SD)	Incorrect Trials (Mean ± SD)	t-value (correct vs. incorrect trials, *indicates p< 0.05, two-tailed)
immediate post-test	6.2 ± 0.2	4.5 ± 0.7	t(10) = 7.6*
1 week post-test	5.2 ± 0.5	4.4 ± 0.6	t(9) = 4.8*

Note: The 1 week post-test results do not include data from one participant, who failed to return for the post-test.

Table 4: Experiment 1 Round X Time (p<0.0001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
Round 1 > Rounds 2 & 3					
	Medial frontal gyrus (L)	6	61	-11, 1, 58	6.2
	Superior frontal gyrus (L) (decreasing activation)	8, 9	44	-21, 48, 32	5.3
	Middle frontal gyrus (R)	8, 9	15	32, 24, 35	4.7
	Inferior frontal gyrus (L)	44, 45	48	-47, 24, 9	5.2
	Inferior parietal lobule (R) (decreasing activation)	40	13	57, -39, 26	3.8
	Superior temporal gyrus (L)	22	5	-56, -30, 6	3.7
	Middle temporal gyrus (L)	37	9	-50, -56, 6	3.8
	Fusiform gyrus (L)	37	8	-40, -40, -9	3.7
Rounds 2 & 3 > Round 1					
	Medial frontal gyrus (L)	6	32	-1, 18, 44	4.8
	Dorsolateral prefrontal cortex (L)	8, 9	66	-44 18, 32	5.8
	Dorsolateral prefrontal cortex (R)	9	26	48, 21, 29	5.9
	Middle frontal gyrus (L)	10	13	-38, 46, 15	4.2
	Lateral orbitofrontal cortex (L)	11	10	-21, 42, -9	4.3
	Inferior parietal lobule, angular gyrus (L)	7, 39	157	-40, -58, 38	7.2
	Cuneus, precuneus, posterior cingulate cortex (L)	18, 19, 31	279	-11, -71, 23	9.9
	Precuneus (R) (decreasing activation)	31	4	11, -65, 26	3.5
	Lingual gyrus (R)	18	7	2, -80, -9	3.6
	Middle occipital gyrus (L)	18	10	-27, -83, -9	3.6
	Insula (R)	13	24	26, 24, 3	5.5
	Putamen (R)		58	26, -5, 15	5.6
	Caudate nuclei (bilateral)		32	8, 8, 6	3.8
	Superior medial cerebellum (L)		9	-11, -78, -21	3.6

Table 5: Experiment 1 Accuracy X Time (p<0.0001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
correct > incorrect					
	Postcentral gyrus (R) (decreasing activation)	3	11	30,-36, 44	5.41
	Anterior cingulate (R)	32	6	18, 27, 9	5.48
	Inferior parietal lobule (L)	40	7	-28, -40, 47	5.35
	Inferior parietal lobule (R) (decreasing activation)	40	5	29 -49, 47	4.94
	Cingulate (R) (decreasing activation)	31	5	11, -27, 44	6.07
	Posterior cingulate/retrosplenial cortex (R)	29	5	8, -36, 9	5.16
	Parahippocampal gyrus, posterior cingulate (L)	19	91	-21, -52, 12	7.19
	Superior temporal gyrus (R) (decreasing activation)	22	39	51, -5, -3	6.56
	Cuneus, middle occipital gyrus, lingual gyrus, parahippocampal gyrus, fusiform gyrus (R)	19	356	23, -59, -6	10.94
	Cuneus, middle occipital gyrus, inferior occipital gyrus (L)	19	77	-31, -75, -6	7.06
	Cuneus (L) (decreasing activation)	18	4	-18, -87, 23	4.99
	Cuneus (R)	18	6	8, -77, 15	5.17
incorrect > correct					
	Medial frontal gyrus/skull		9	4, -14, 73	6.02
	Medial frontal gyrus, superior frontal gyrus, dorsal anterior cingulate (midline)	6, 8, 32	275	-2, 0, 53	9.67
	Superior frontal gyrus (L)	9	19	-24, 39, 29	6.21
	Middle frontal gyrus, inferior frontal gyrus (L)	9, 45, 46	109	9, 45, 46	7.83
	Inferior frontal gyrus (L)	26	6	-24, 27,-9	6.11
	Precentral gyrus (R) (decreasing activation)	4	5	54, -8, 38	5.41
	Postcentral gyrus (L)	1, 2, 3	70	-43, -22, 53	7.52
	Superior parietal lobule, inferior parietal lobule, angular gyrus (L)	7, 39, 40	130	-37, -56, 38	9.59
	Inferior parietal lobule (R)	40	16	45, -58, 47	7.78
	Insula, inferior parietal lobule (R) (decreasing activation)	13	10	42, -27, 23	5.89
	Insula (L)	13	10	-31, 21, 9	5.41
	Middle temporal gyrus, inferior temporal gyrus, fusiform gyrus, anterior cerebellum (L)	21, 37	168	-40, -53, -24	9.3
	Middle temporal gyrus (R)	21	5	60, -33, -9	4.93
	Medial inferior cerebellum (bilateral)		104	14, -75, -24	8.99
multiple patterns					
	Posterior cingulate, caudate nuclei, thalamus, putamen		998	-8, 8, 0	15.45
correct > incorrect	caudate nuclei (bilateral)			-8, 8, 0	15.5
incorrect > correct	posterior cingulate	23		2, -36, 20	7.1
	putamen (R)			23, -5, 0	8.2
	putamen (L)			-27, -5, 0	11.0
	thalamus, caudate body (L)			-14, -13, 18	12.5
	thalamus, caudate body (R)			14, -11, 15	10.8

Note: When clusters contained more than 100 voxels and had multiple peaks, the threshold was raised until they split into clusters of fewer than 100 voxels, and the activation patterns in these smaller clusters were examined. In one case, the smaller regions showed different patterns of activation, and so the region is listed under "multiple patterns," with the patterns of the smaller regions listed beneath.

Table 6: Experiment 1 Round X Accuracy X Time (p<0.0001)

	Region of Activation	BA	Size (# voxels)	Peak Talairach Coordinates (x, y, z)	Maximum F Value
correct > incorrect difference greatest on Round 3					
	medial frontal gyrus (R)	6	24	11, -5, 53	4.3
	Precentral gyrus (R)	6	5	38, -5, 38	3.77
	Superior parietal lobule (R) (decreasing activation)	7	6	24, -49, 58	3.49
	Precuneus (R) (decreasing activation)	7	5	17, -45, 41	3.37
	Posterior cingulate (R) (decreasing activation)	29	9	20, -47, 9	3.96
incorrect > correct difference greatest on Round 3					
	superior frontal gyrus (L)	8	42	-5, 36, 50	4.42
	superior frontal gyrus (L)	8	7	-8, 52, 44	3.72
	Middle frontal gyrus (R)	6	6	32, 9, 56	3.4
	Middle frontal gyrus (L)	6	10	-39, 5, 47	3.51
	Middle frontal gyrus (L)	10	10	-37, 58, 12	4.43
	medial frontal gyrus (midline) (decreasing activation)	9	17	2, 42, 20	4.13
	dorsolateral prefrontal cortex, inferior frontal gyrus/ventrolateral prefrontal cortex, insula (L)	9, 13, 46, 47	201	-37, 17, 6	5.03
	anterior cingulate (midline)	32	25	-2, 24, 38	4.33
	posterior cingulate (midline)	31	44	-2, -39, 29	5.66
	cingulate gyrus (L)		7	-18, 9, 23	3.58
	Posterior cingulate (L) (decreasing activation)	31	5	-20, -61, 15	3.45
	precuneus (R)	7	22	-5, -55, 67	4.12
	superior parietal lobule, inferior parietal lobule, angular gyrus, middle temporal gyrus (L)	7, 39, 40	178	-47, -61, 44	4.65
	middle temporal gyrus, inferiror temporal gyrus, fusiform gyrus (L)	20, 21, 37	53	-56, -24, -6	4.21
	caudate body (R)		35	17, -8, 23	4.55
	caudate body (L)		11	-18, -16, 23	3.78
	Medial inferior cerebellum (L)		38	-14, -87, -24	4.62
	Lateral superior cerebellum (R)		9	42, -72, -24	4.27
	Medial inferior cerebellum (R)		7	14, -74, -24	3.48

Table 7: Experiment 2 Trial Type X Time (p<0.001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
2 choices > 4 choices					
	superior frontal gyrus, middle frontal gyrus, inferior frontal gyrus (R), medial frontal gyrus, anterior cingulate (bilateral)	6, 8, 9, 10, 32, 44	1228	6, 29, 48	9.67
	medial frontal gyrus (L) (decreasing activation)	6	6	-18, 10, 51	3.88
	middle frontal gyrus (L) (decreasing activation)	8	11	-31, 13, 42	3.87
	precentral gyrus (L) (decreasing activation)		11	-31, -, 27	3.75
	paracentral lobule (midline) (decreasing activation)	5	37	3, -27, 54	4.92
	posterior cingulate (midline)	23, 31	60	3, -38, 33	4.95
	inferior parietal lobule, supramarginal gyrus, angular gyrus, middle temporal gyrus (R)	39, 40	645	51, -48, 39	11.58
	precuneus (R) (decreasing activation)	7	10	6, -59, 33	3.85
	insula (L)	13	5	-45, -2, 18	3.86
	insula (R)		5	24, 21, -6	3.9
	superior temporal gyrus, middle temporal gyrus (R) (decreasing activation)	21, 22	310	48, -24, -6	7.83
	middle temporal gyrus (L) (decreasing activation)	21	82	-48, -21, -6	7.18
	middle temporal gyrus (L) (decreasing activation)	21	5	-48, -5, -15	3.79
	inferior temporal gyrus (R) (decreasing activation)	20	4	44, -8, -18	3.91
	caudate body/ventricle		33	6, 7, 21	5.88
4 choices > 2 choices					
	medial frontal gyrus, anterior cingulate (midline)	6, 24	184	-4, 1, 45	7.95
	middle frontal gyrus (R)	6	9	24, -9, 45	4.06
	precentral gyrus (R)	6	25	34, -15, 57	6.41
	precentral gyrus (L)	6	11	-58, 3, 30	4.42
	precuneus, cuneus, middle occipital gyrus, lingual gyrus, inferior occipital gyrus, fusiform gyrus, sup. cerebellum (bilateral)	7, 17, 18, 19	3212	3, -77, -3	33.97
	thalamus (L)		20	-14, -21, 6	4.55
multiple patterns					
	precentral gyrus, postcentral gyrus, inferior parietal lobule, supramarginal gyrus, angular gyrus, middle temporal gyrus (L)	2, 3, 4, 6, 39, 40	679	-38, -31, 48	9.87
2 choices > 4 choices	inferior parietal lobe, supramarginal gyrus	40		-58, -48, 36	7.18
4 choices > 2 choices	precentral gyrus (L)	4		-31, -21, 60	7.24
	postcentral gyrus (L)	2, 3		-38, -31, 48	9.87

Note: When clusters contained more than 100 voxels and had multiple peaks, the threshold was raised until they split into clusters of fewer than 100 voxels, and the activation patterns in these smaller clusters were examined. In one case, the smaller regions showed different patterns of activation, and so the region is listed under "multiple patterns," with the patterns of the smaller regions listed beneath.

Table 8: Experiment 2 Feedback Type X Time (p<0.001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
positive feedback > negative feedback					
	precentral gyrus, middle frontal gyrus (L)	6, 8, 9	38	-44, -3, 45	4.31
	middle frontal gyrus/inferior frontal gyrus (L)	45, 46	41	-48, 20, 24	5.18
	inferior frontal gyrus (L)	46	5	-44, 36, 6	3.7
	anterior cingulate (R) (decreasing activation)	24	5	14, 23, -3	3.86
	insula (L)	13	8	-34, 20, 6	4
	insula (R)	13	9	27, 25, 0	4
	inferior temporal gyrus, middle occipital gyrus (L)	19, 37	35	-48, -67, 0	5.45
	middle temporal gyrus (L)	21	13	-58, -34, -3	4.11
	middle occipital gyrus (R)	19	58	34, -69, -6	5.56
	thalamus (bilateral); midbrain (L)		45	-8, -14, 3	6.11
negative feedback > positive feedback					
	superior frontal gyrus/middle frontal gyrus	8, 9, 46	65	30, 38, 33	5.36
	superior frontal gyrus (L) (decreasing activation)	10	4	-10, 71, 21	3.75
	middle frontal gyrus (R) (decreasing activation)	6	18	24, 1, 51	4.65
	middle frontal gyrus (L) (decreasing activation)	8, 9	5	-27, 23, 39	3.46
	middle frontal gyrus (L) (decreasing activation)	9	4	-27, 36, 27	3.62
	precentral gyrus (L)	4	36	-31, -15, 45	6.58
	cingulate gyrus (R) (decreasing activation)	31	23	10, -38, 39	6.53
	cingulate gyrus (L)	24	8	-4, 1, 39	3.88
	inferior parietal lobule (L)	40	5	-58, -27, 30	3.63
	precuneus (R) (decreasing activation)	31	10	10, -68, 24	4.2
	lingual gyrus (R)	18	5	9, -83, 0	3.87
	inferior occipital gyrus (L)	18	9	-38, -83, -3	3.77
	inferior occipital gyrus (R)	17	8	24, -101, -9	4.31
	Lateral superior cerebellum (R)		24	27, -70, -22	5.19

Table 9: Experiment 2 Trial Type X Feedback Type X Time (p<0.001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
positive feedback > negative feedback for 4-choice condition only					
	superior frontal gyrus (L)	6	11	-28, -3, 63	4.46
	superior frontal gyrus, medial frontal gyrus (bilateral)	6, 8	156	-4, 26, 48	6.76
	superior frontal gyrus (L)	10	7	-21, 55, 21	3.66
	dorsolateral prefrontal cortex, inferior frontal gyrus/ventrolateral prefrontal cortex, insula (L)	9, 13, 44, 45, 46, 47	542	-51, 17, 3	10.22
	inferior frontal gyrus (ventrolateral PFC) (R)	45, 47	49	48, 20, 9	6.05
	inferior frontal gyrus (R)	47	15	37, 32, -12	4.14
	anterior cingulate (midline) (decreasing activation)	32	17	3, 43, 0	4.6
	inferior parietal lobule (L)	40	73	-45, -59, 42	5.83
	insula (R)	13	10	27, 20, 3	4.17
	middle temporal gyrus (L)	21	38	-62, -45, 0	5.04
	dorsal head of the caudate nucleus (R)		10	10, 7, 15	4.19
negative feedback > positive feedback for 4-choice condition only					
	anterior cingulate (R) (decreasing activation)	32	26	10, 42, 15	4.59
	insula (L) (decreasing activation)	13	4	-35, -18, 24	3.8
	precuneus (R) (decreasing activation)	31	13	3, -66, 21	3.99
	precuneus (L) (decreasing activation)	31	9	-21, -59, 18	4.1
	parahippocampal gyrus/ventricle (decreasing activation)	19	4	-31, -52, 0	4.31
	fusiform gyrus (R)	19	28	24, -80, -12	3.71
positive feedback > negative feedback for 2-choice condition only					
	parahippocampal gyrus (R) (decreasing activation)	36	22	27, -21, -25	4.87
negative feedback > positive feedback for 2-choice condition only					
	superior frontal gyrus, middle frontal gyrus (decreasing activation)	6, 8	14	30, 26, 54	3.96
	middle frontal gyrus (R) (decreasing activation)	6	15	24, 7, 45	4.46
	superior temporal gyrus (R) (decreasing activation)	38	4	52, 17,-25	3.75
	Inferior cerebellum (R)		12	17, -69, -25	3.72

Table 10: Experiment 2 4-choice Condition: Feedback Type X Time (p<0.001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
positive feedback > negative feedback					
	superior frontal gyrus (midline)	8	73	-4, 29, 48	5.58
	dorsolateral prefrontal cortex, ventrolateral prefrontal cortex, insula (L)	9, 13, 44, 45, 46, 47	838	-51, 19, 27	11.38
	middle frontal gyrus (R)	10	7	36, 46, 0	3.64
	inferior frontal gyrus/VLPFC, insula (R)	13, 45, 47	168	48, 20, 15	8.18
	anterior cingulate (bilateral) (decreasing activation)	32	109	-4, 36, 18	4.93
	anterior cingulate (L)		4	-20, 26, 15	3.65
	inferior parietal lobule (L)	40	71	-48, -56, 42	6.11
	middle temporal gyrus (L)	21, 22	147	-58, -38, -3	8.09
	caudate nucleus (L)		50	-11, 10, 6	4.93
	caudate nucleus (R)		22	6, 10, 3	4.56
	thalamus/midbrain (bilateral)		26	-4, -15, 3	4.69
	cuneus/lingual gyrus (R)	18	15	26, -96, 0	4.33
	Inferior cerebellum (R)		63	20, -73, -25	5.67
negative feedback > positive feedback					
	medial frontal gyrus (R) (decreasing activation)	6	9	20, 1, 51	3.76
	precentral gyrus (L) (decreasing activation)	6	7	-37, -15, 36	4.29
	cingulate gyrus (R) (decreasing activation)	31	8	7, -38, 39	4.62
	inferior parietal lobule (L)	40	11	-48, -32, 27	4.34
	insula (L) (decreasing activation)	13	4	-38, -15, 21	3.9
	precuneus (bilateral) (decreasing activation)	31	135	10, -66, 21	6.75
	parahippocampal gyrus (R) (decreasing activation)	30	14	17, -38, -3	5.71
	parahippocampal gyrus (L) (decreasing activation)	19	8	-31, -42, -3	4.38
	fusiform gyrus (L)	18	8	-31, -97, -15	3.87

Table 11: Experiment 2 2-choice Condition: Feedback Type X Time (p<0.001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
positive feedback > negative feedback					
	superior frontal gyrus (R) (decreasing activation)	8	8	24, 35, 51	3.93
	anterior cingulate (R)		4	14, 26, -9	3.78
	middle occipital gyrus (R)	19	23	37, -65, -6	4.49
	parahippocampal gyrus (R) (decreasing activation)	36	5	28, -15, -22	3.62
negative feedback > positive feedback					
	superior frontal gyrus (R) (decreasing activation)	8	10	21, 16, 45	4.43
	middle frontal gyrus (L)	10	5	-38, 36, 24	4.04
	precentral gyrus (L)	6	7	-31, -11, 51	3.86
	anterior cingulate (L)	32	29	-4, 10, 42	4.9
	cingulate gyrus (R) (decreasing activation)	31	6	13, -35, 39	3.7

Table 12: Experiment 2 Subsequent Memory Effects:
Post-test accuracy (high confidence correct trials vs. incorrect trials) X Time (p<0.001)

	Region of Activation	*BA*	*Size (# voxels)*	*Peak Talairach Coordinates (x, y, z)*	*Maximum F Value*
subsequent high confidence correct responses > incorrect responses					
	postcentral gyrus (L)	3	6	-48, -18, 57	3.61
	middle frontal gyrus (R)	9	32	41, 4, 36	5.15
	inferior frontal gyrus (L)	10, 47	137	-51, 22, 3	6.68
	anterior inferior frontal gyrus (R)	47	5	50, 43, 0	4.07
	ventromedial prefrontal cortex (L)	10	59	-8, 55, -3	5.4
	posterior middle temporal gyrus (L) (decreasing activation)	39	5	-47, -69, 21	3.83
	middle temporal gyrus (L) (decreasing activation)	21	7	-61, -47, 9	3.79
	middle temporal gyrus (L) (decreasing activation)	21	48	-65,-25, -3	4.61
	middle temporal gyrus (R) (decreasing activation)	21	6	54, -39, -3	4
	lingual gyrus (L) (decreasing activation)	18	13	-8, -59, 3	4.08
	lingual gyrus (R)	18	5	13,-73, 0	3.68
	lingual gyrus (L)	18	11	-14, -97, -3	4.15
	inferior occipital gyrus (R)	18	13	41, -80, -6	3.88
	middle occipital gyrus, fusiform gyrus (L)	18	45	-45, -77, -9	4.29
	anterior caudate/ventricle (L) (decreasing activation)		8	-13, 26, 9	4.14
	thalamus (L)		4	-4, -8, 12	3.64
	Hippocampus/parahippocampus (L)		11	-31, -15, -6	4.56
subsequent incorrect responses > high confidence correct responses					
	posterior inferior frontal gyrus (R)	44	4	48, 4, 18	4.67
	precuneus (R) (decreasing activation)	31	7	6, -46, 36	4.21

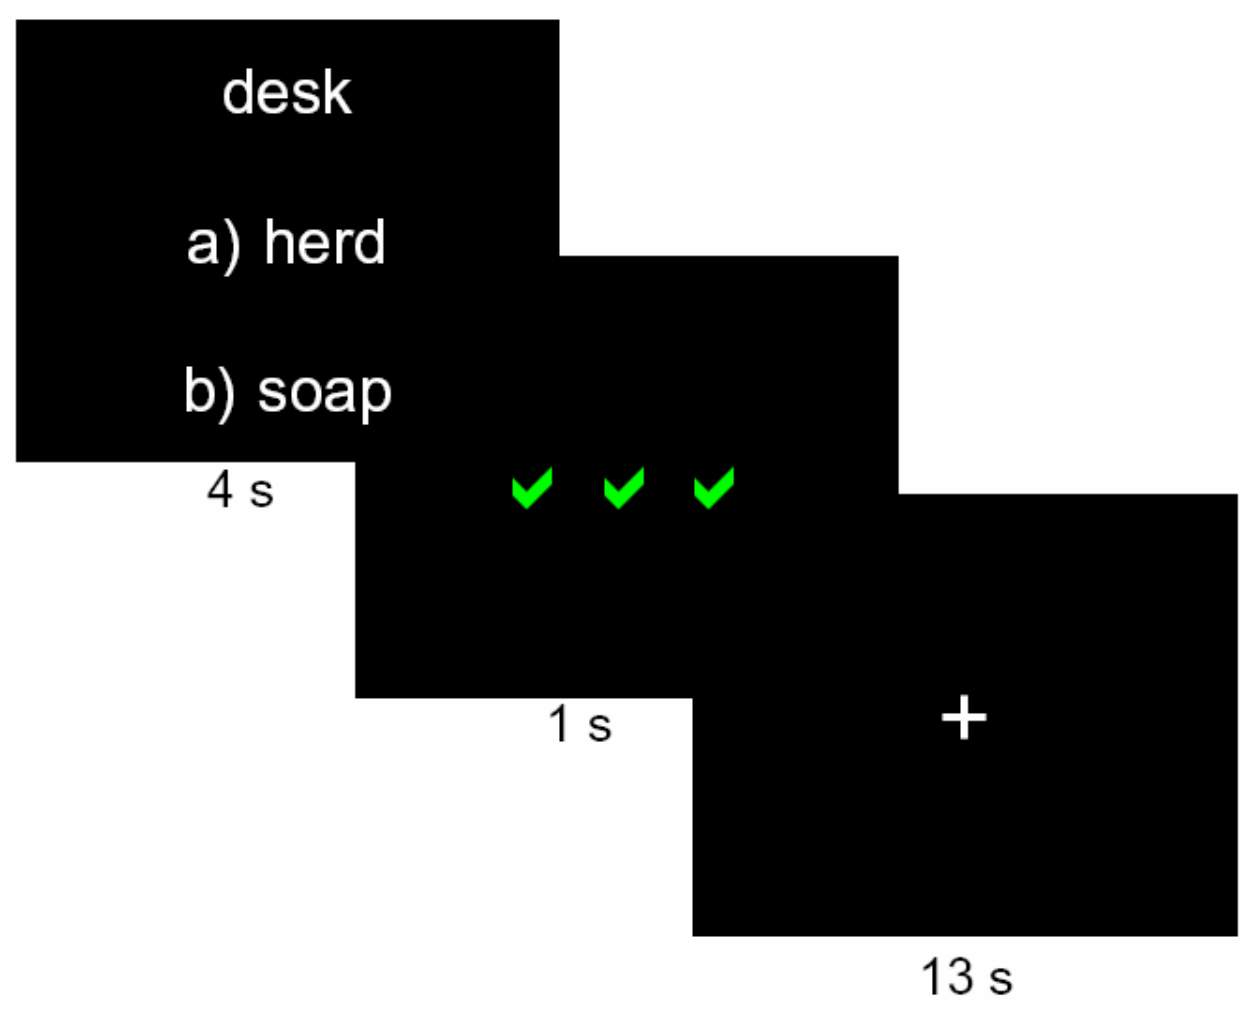

Figure 1. Experimental Design.

Each trial, a target word was presented, along with options for possible word matches, labeled as in a multiple-choice test. After a 4 s response period, the display was replaced with a 1 s feedback display of 3 green √s, indicating a correct response, 3 red Xs, indicating an incorrect response, or 3 white hypens, indicating that no response was made. After a 13 s delay, with a screen showing a fixation cross, the next trial began.

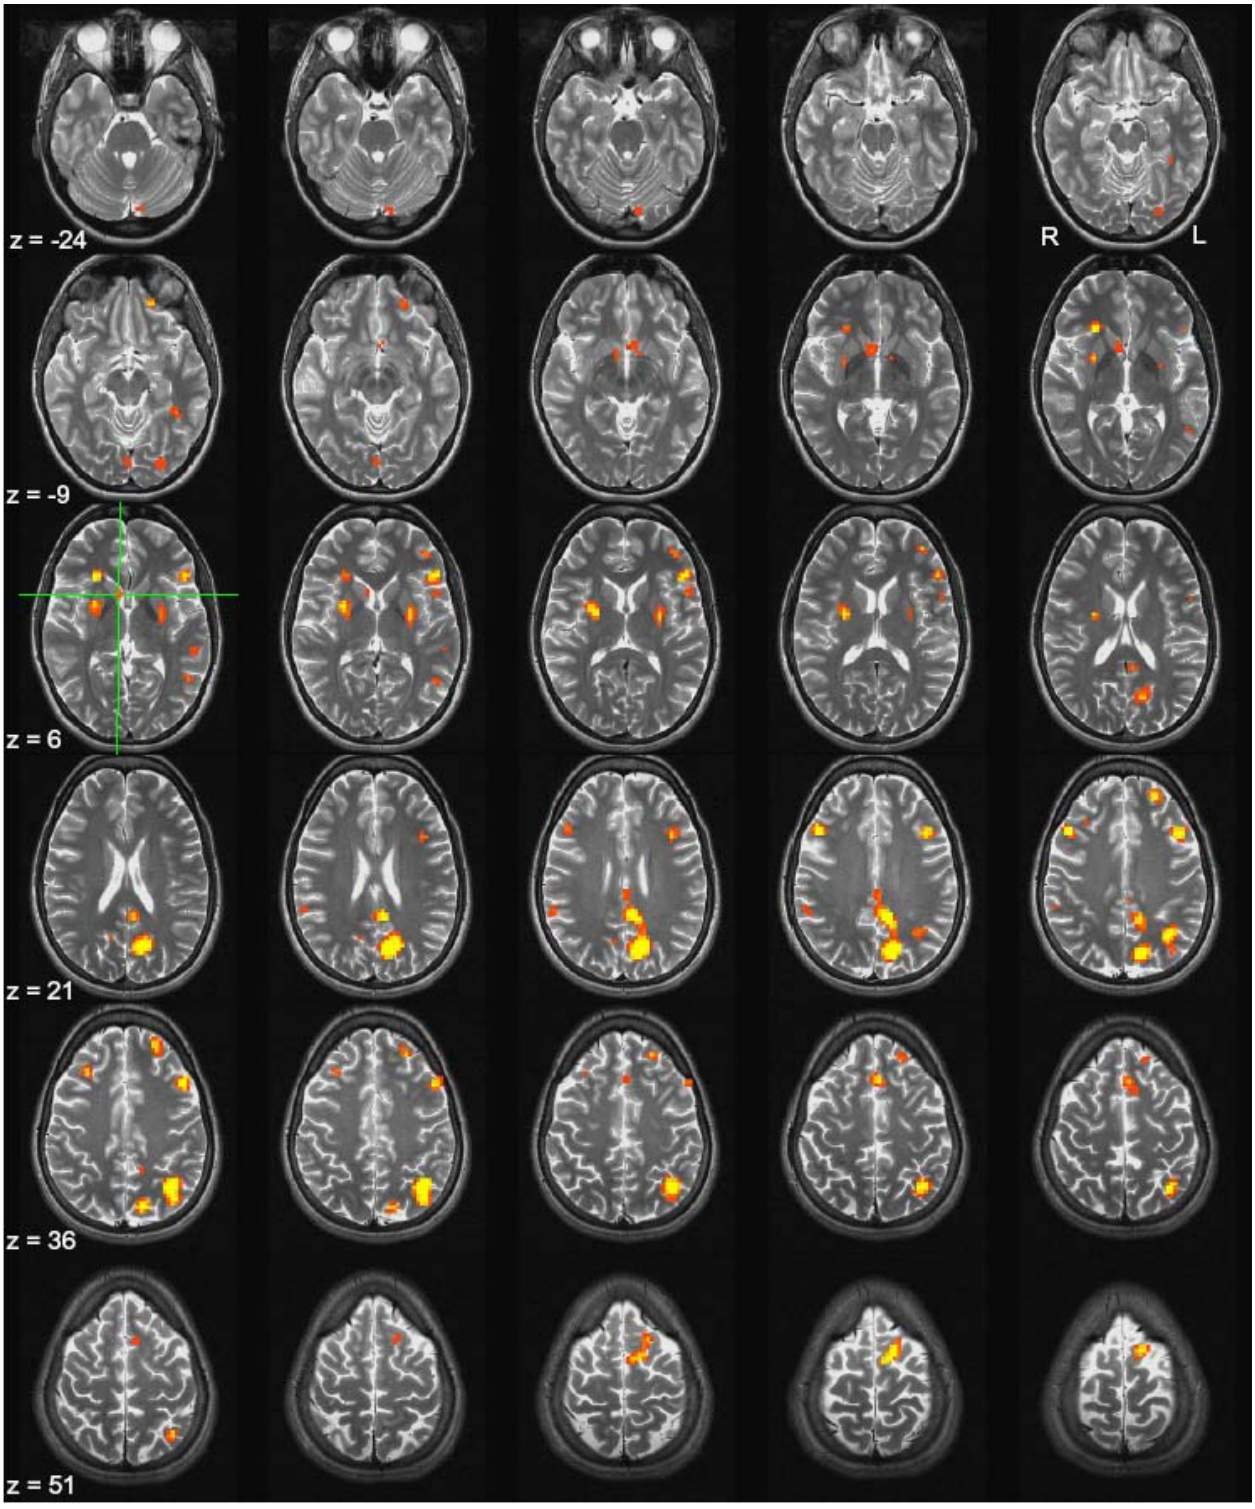

Figure 2. Brain Regions Showing a Round by Time Interaction in Experiment 1.

Voxel clusters displaying a Round by Time interaction in Experiment 1 are shown ($p < 0.0001$; contiguity threshold of 4 voxels). The green crosshair marks the caudate voxel with the peak F-value. Images are left-right reversed.

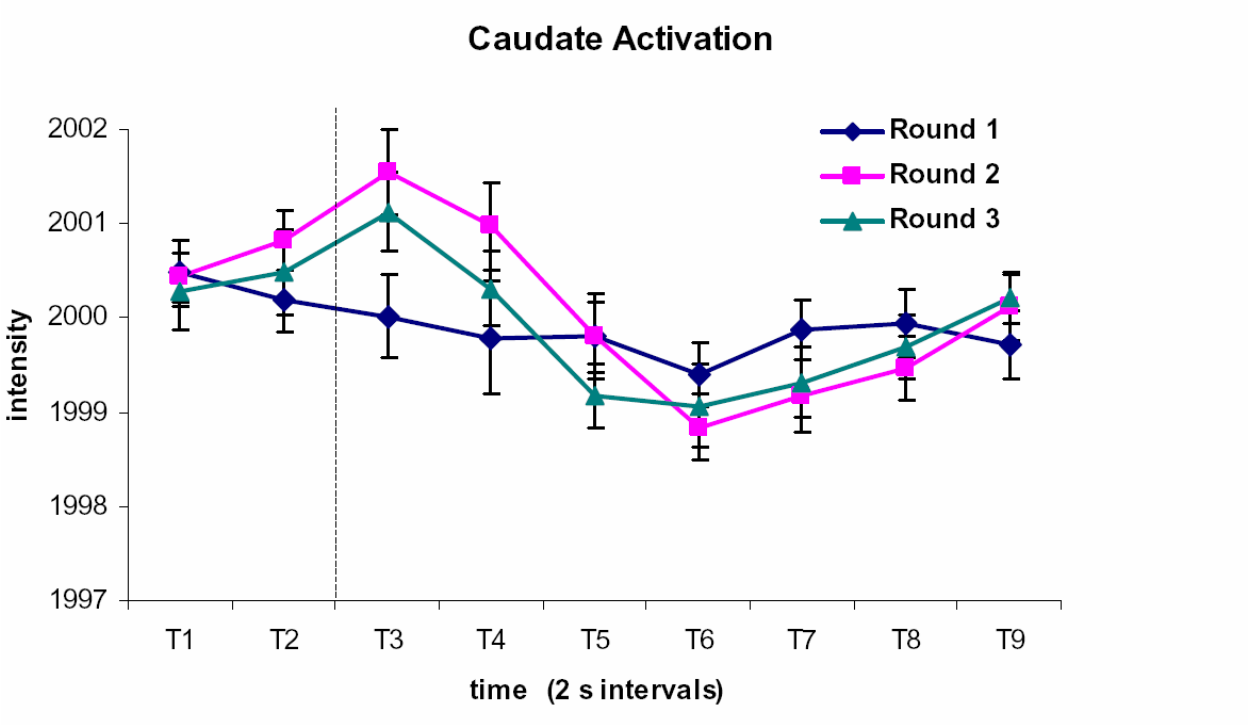

Figure 3. Time Course of Activation in the Caudate Region Displaying a Round by Time Effect.

The caudate nuclei displayed a relatively flat pattern of activation during Round 1 trials, while there is greater activation during Rounds 2 and 3. The activation is somewhat attenuated on Round 3 relative to Round 2. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display.

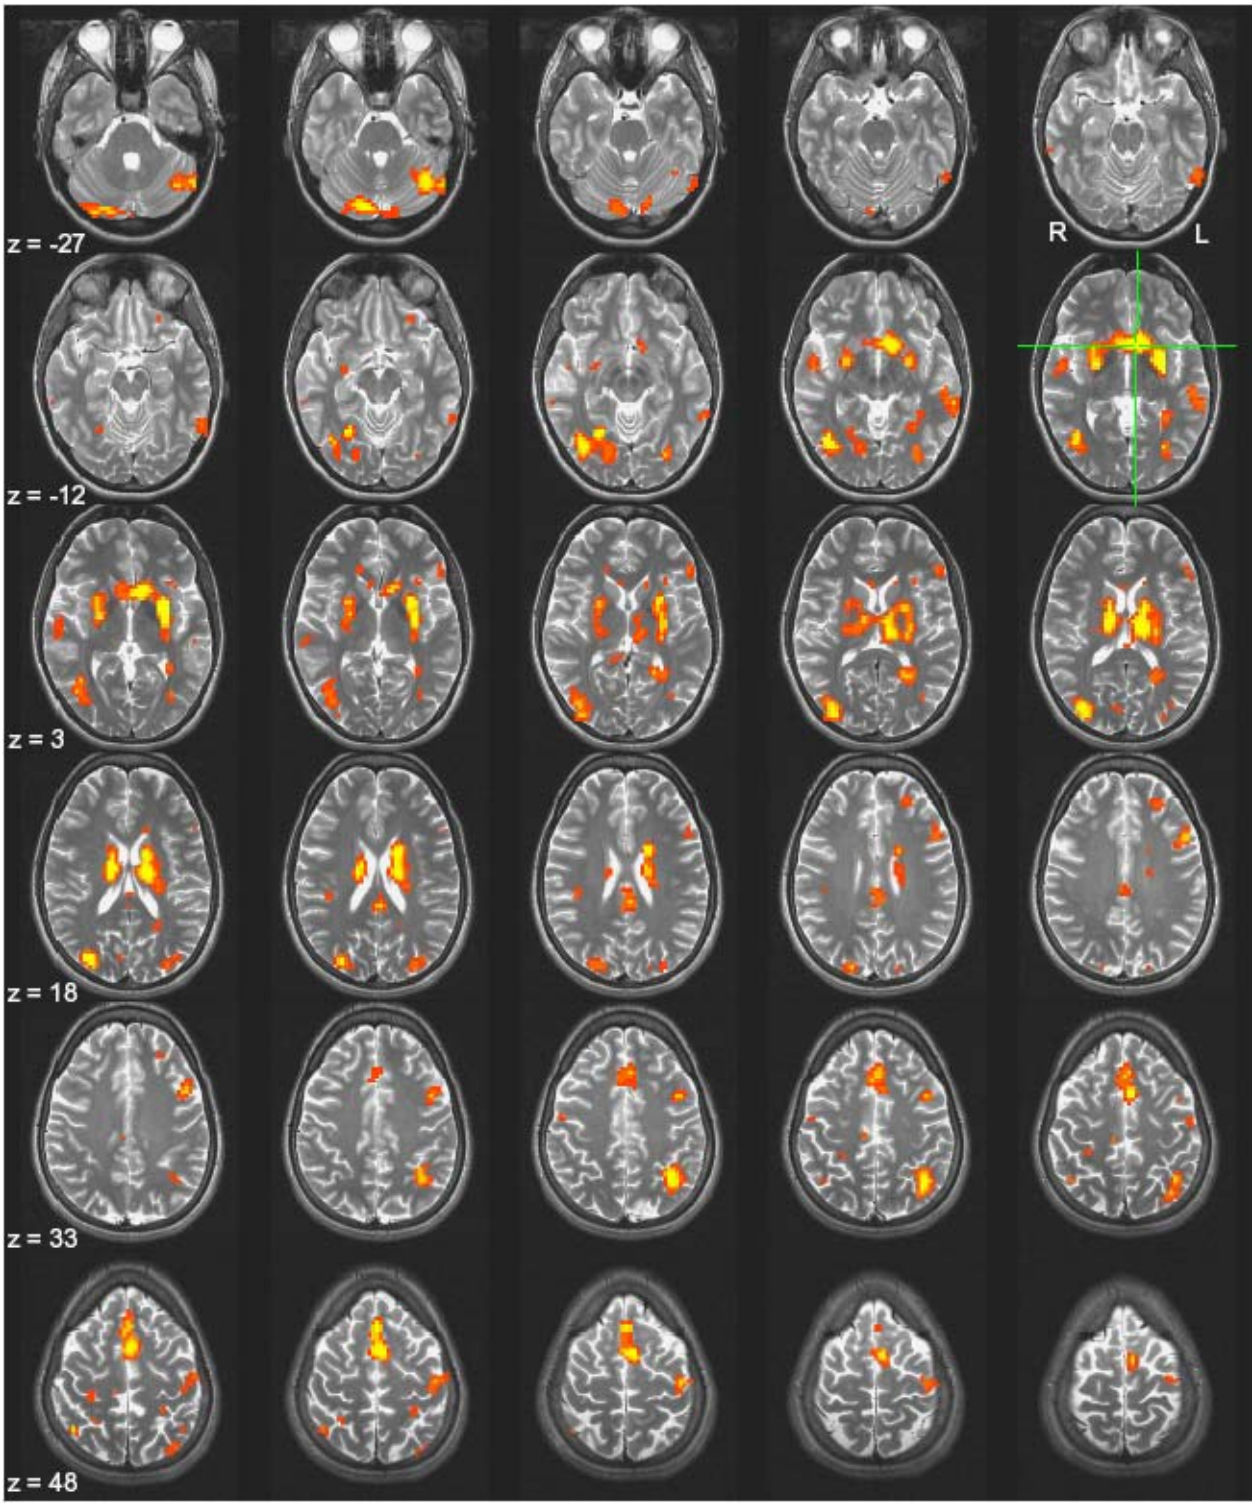

Figure 4. Brain Regions Showing an Accuracy by Time Effect in Experiment 1.

Voxel clusters displaying an Accuracy by Time interaction in Experiment 1 are shown ($p < 0.0001$; contiguity threshold of 4 voxels). The green crosshair marks the caudate voxel with the peak F-value. Images are left-right reversed.

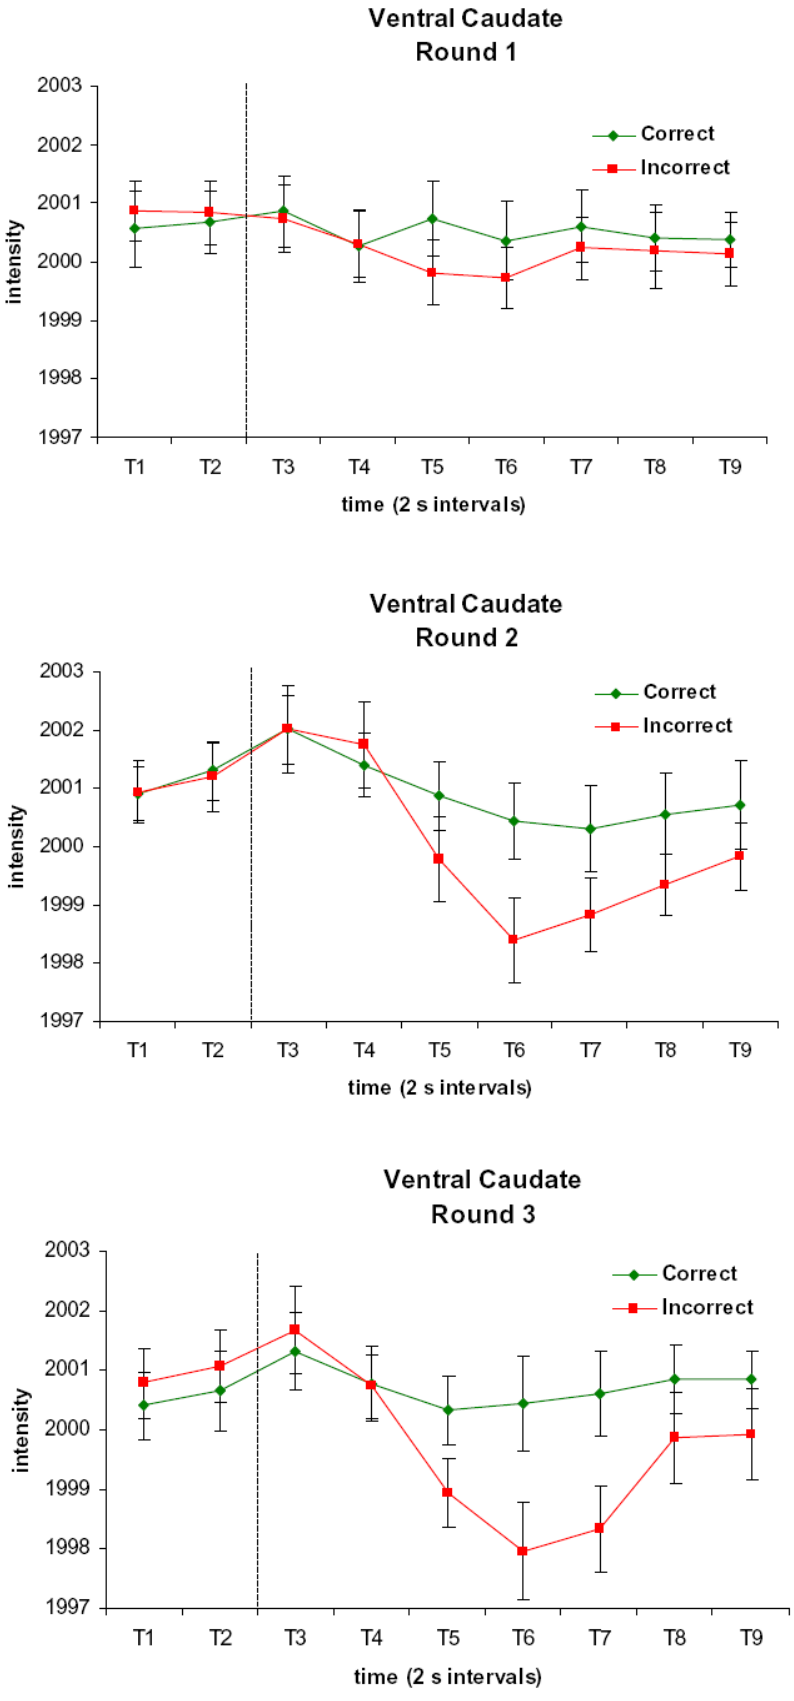

Figure 5. Time Courses of Activation in the Caudate Region Displaying an Accuracy by Time Effect.

There is a slight differentiation between trials with positive and negative feedback during Round 1, whereas for Rounds 2 and 3, the signal following negative feedback shows a pronounced dip below baseline, unlike the signal following positive feedback. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display.

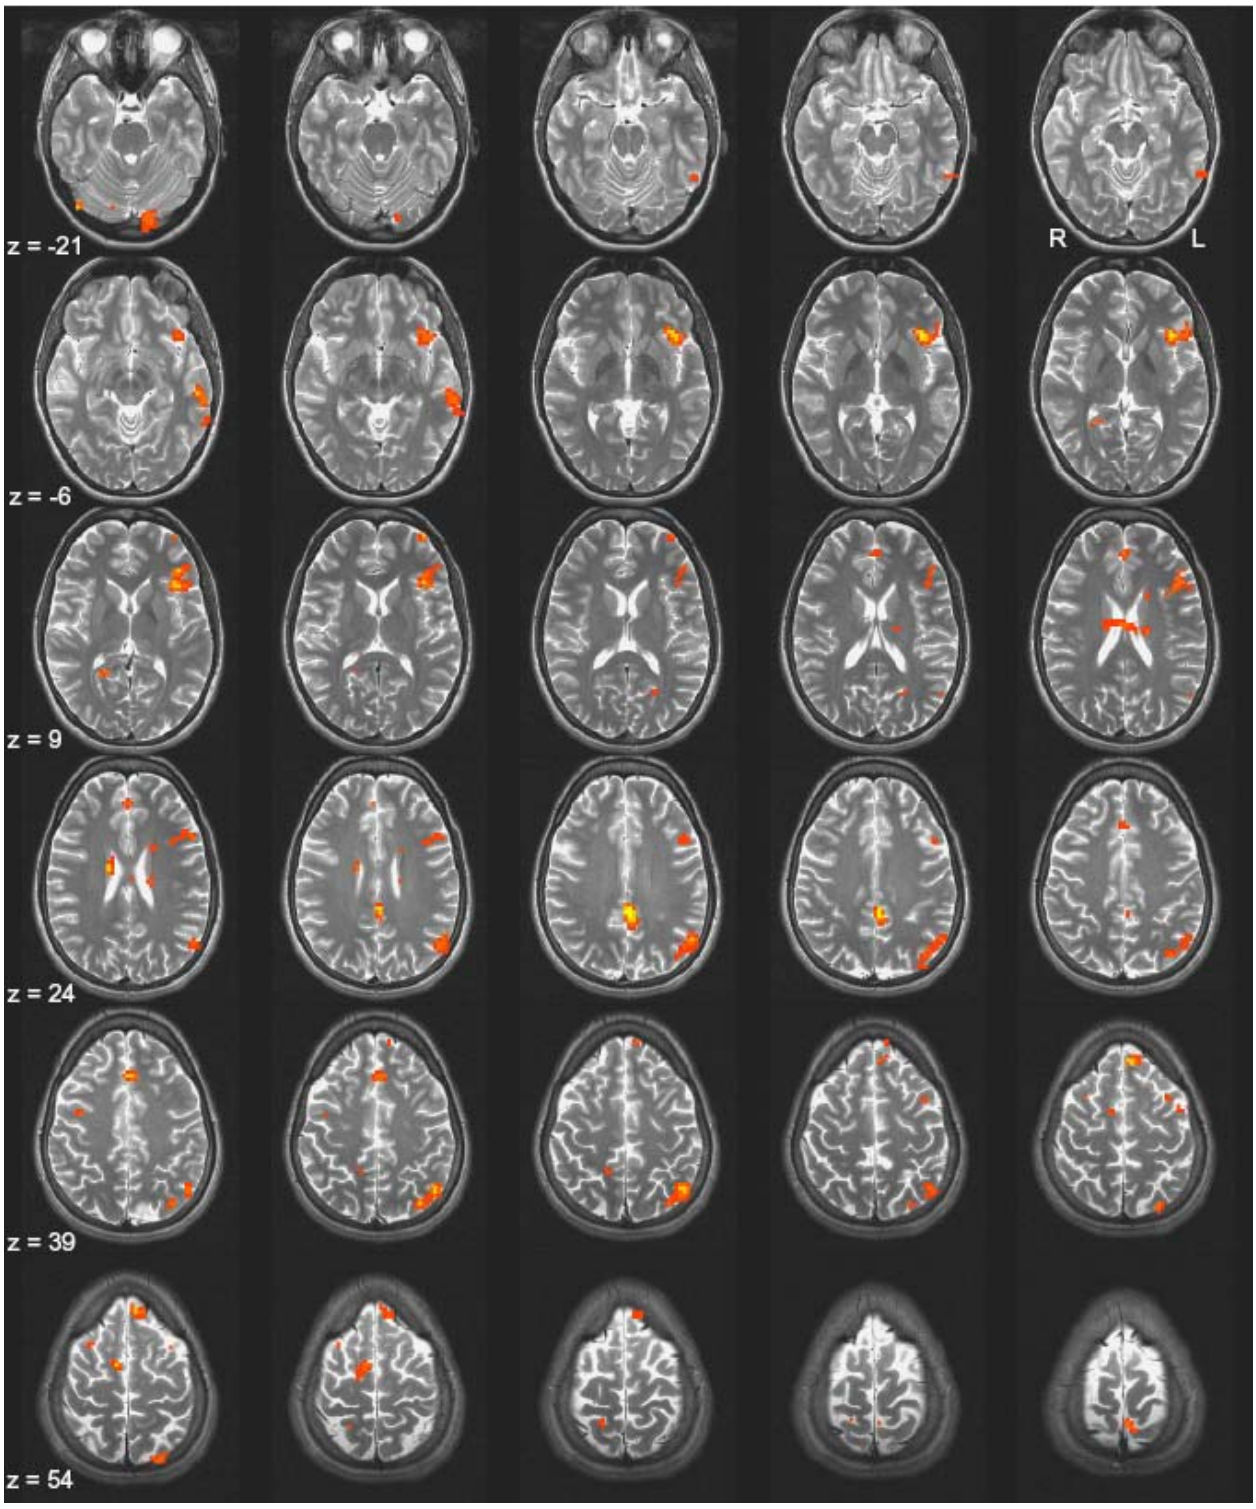

Figure 6. Brain Regions Showing a Round by Accuracy by Time Interaction in Experiment 1.

Voxel clusters displaying a Round by Accuracy by Time interaction in Experiment 1 are shown ($p < 0.0001$; contiguity threshold of 4 voxels). Images are left-right reversed.

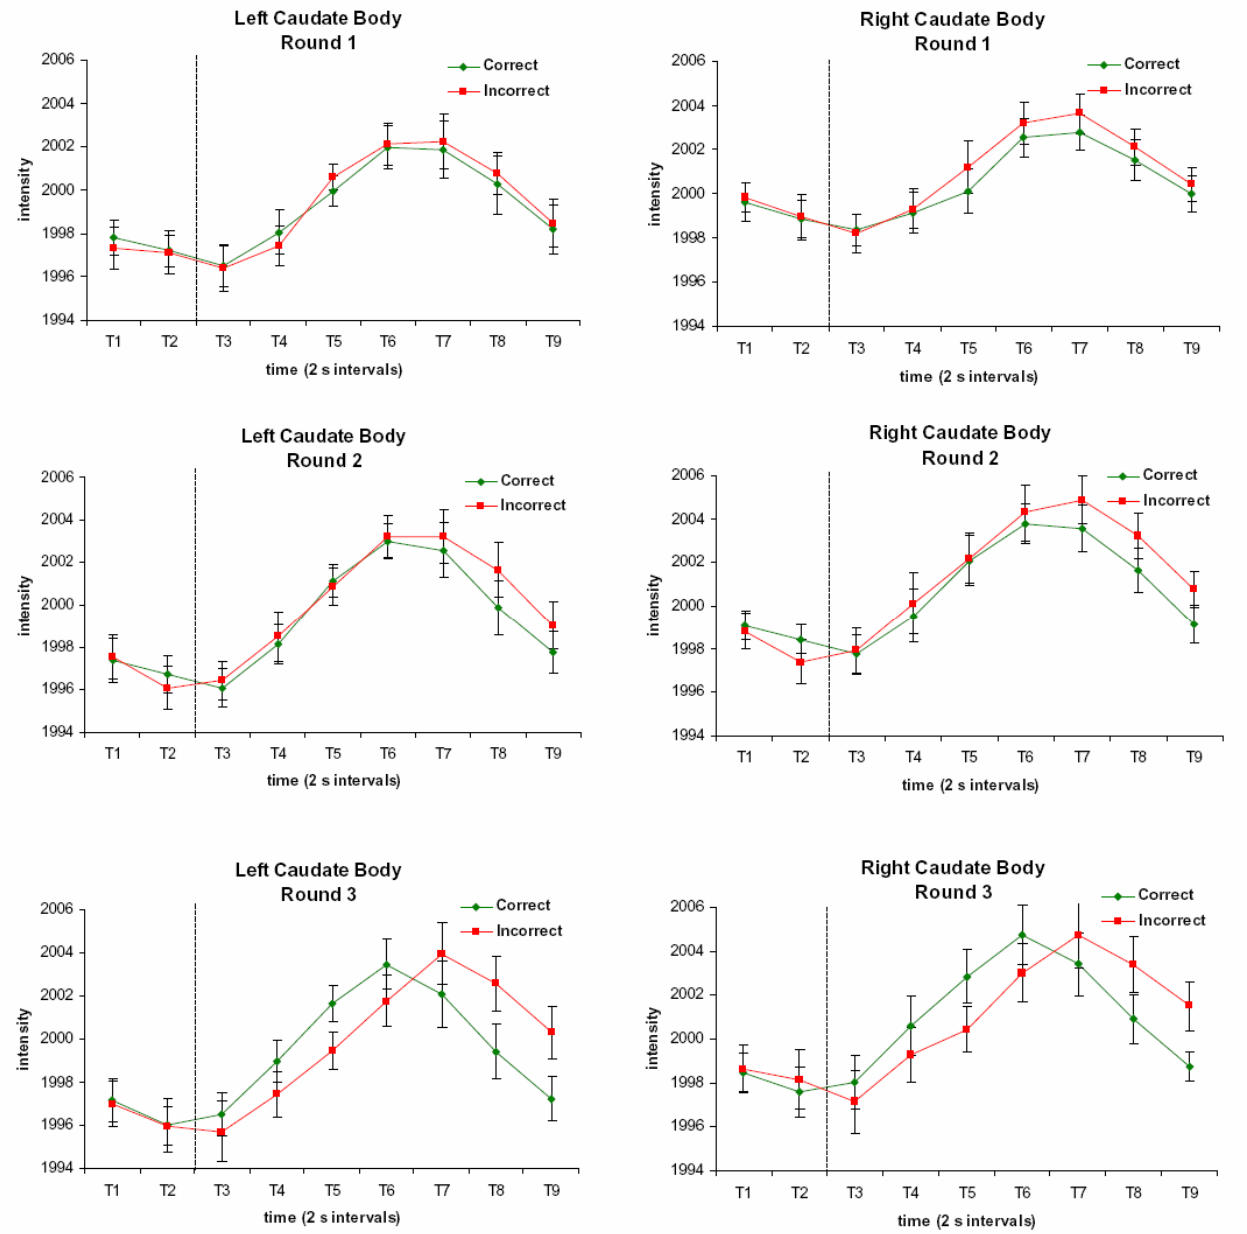

Figure 7. Time Courses of Activation in the Caudate Body Region Showing a Round by Accuracy by Time Interaction.

Activation is elicited on all three rounds in the caudate body, with an earlier response on correct trials than on incorrect trials during Round 3, but not Rounds 1 and 2. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display.

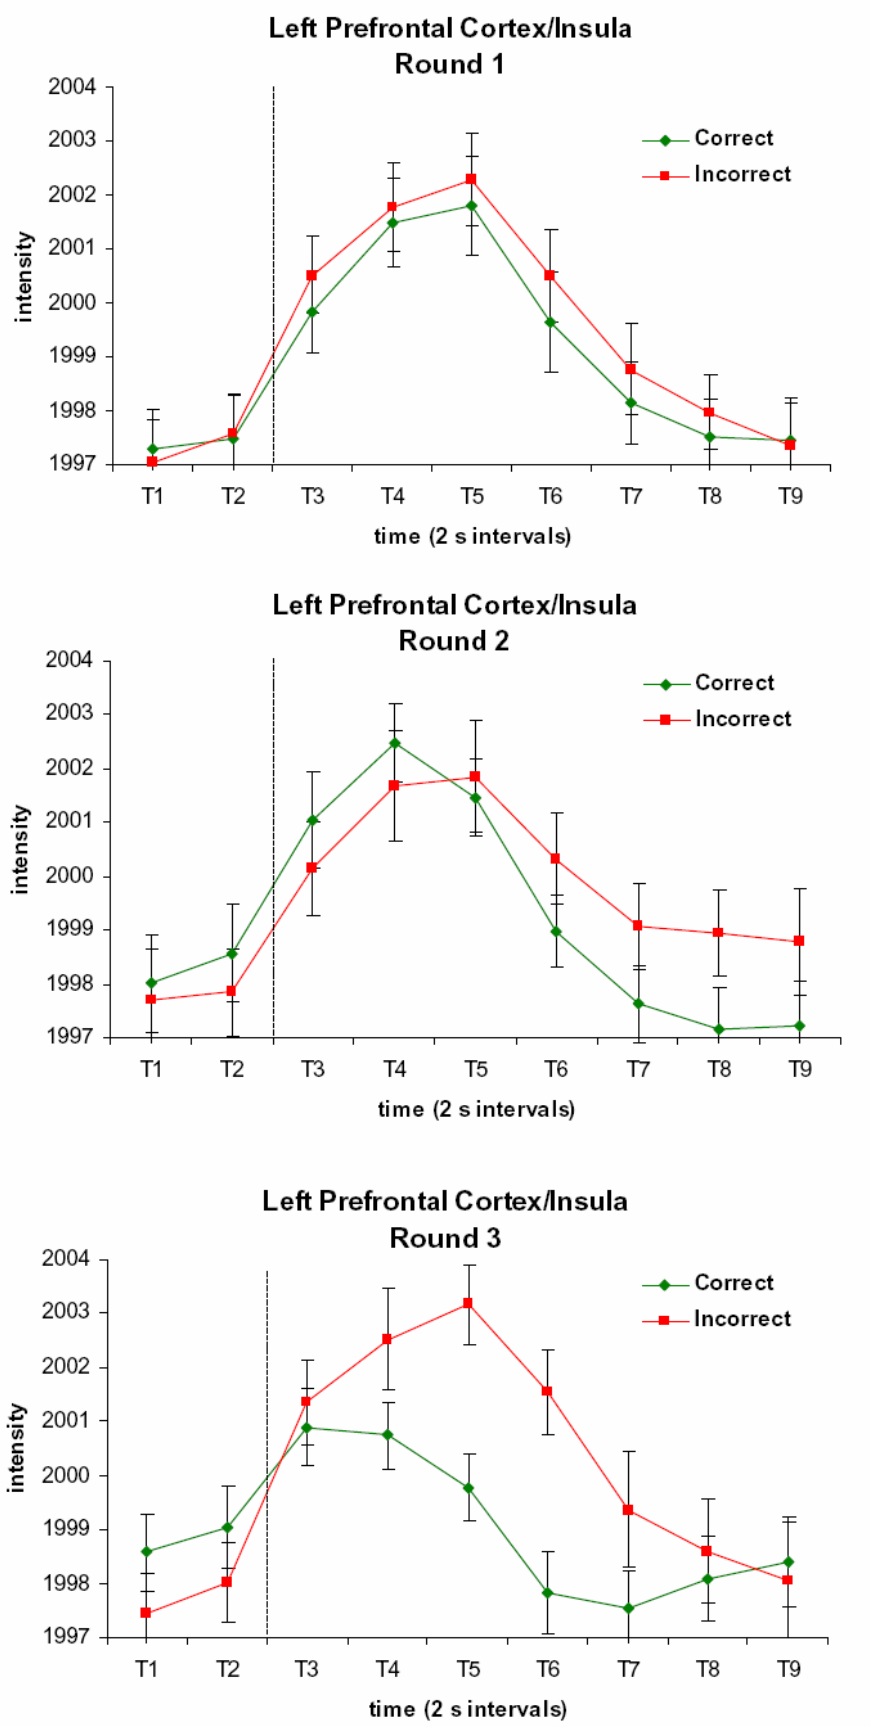

Figure 8. Time Courses of Activation in the Left Prefrontal Cortex/Insula Region Showing a Round by Accuracy by Time Interaction.

The signal is greater following incorrect than correct trials, and this difference is greatest on Round 3. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display.

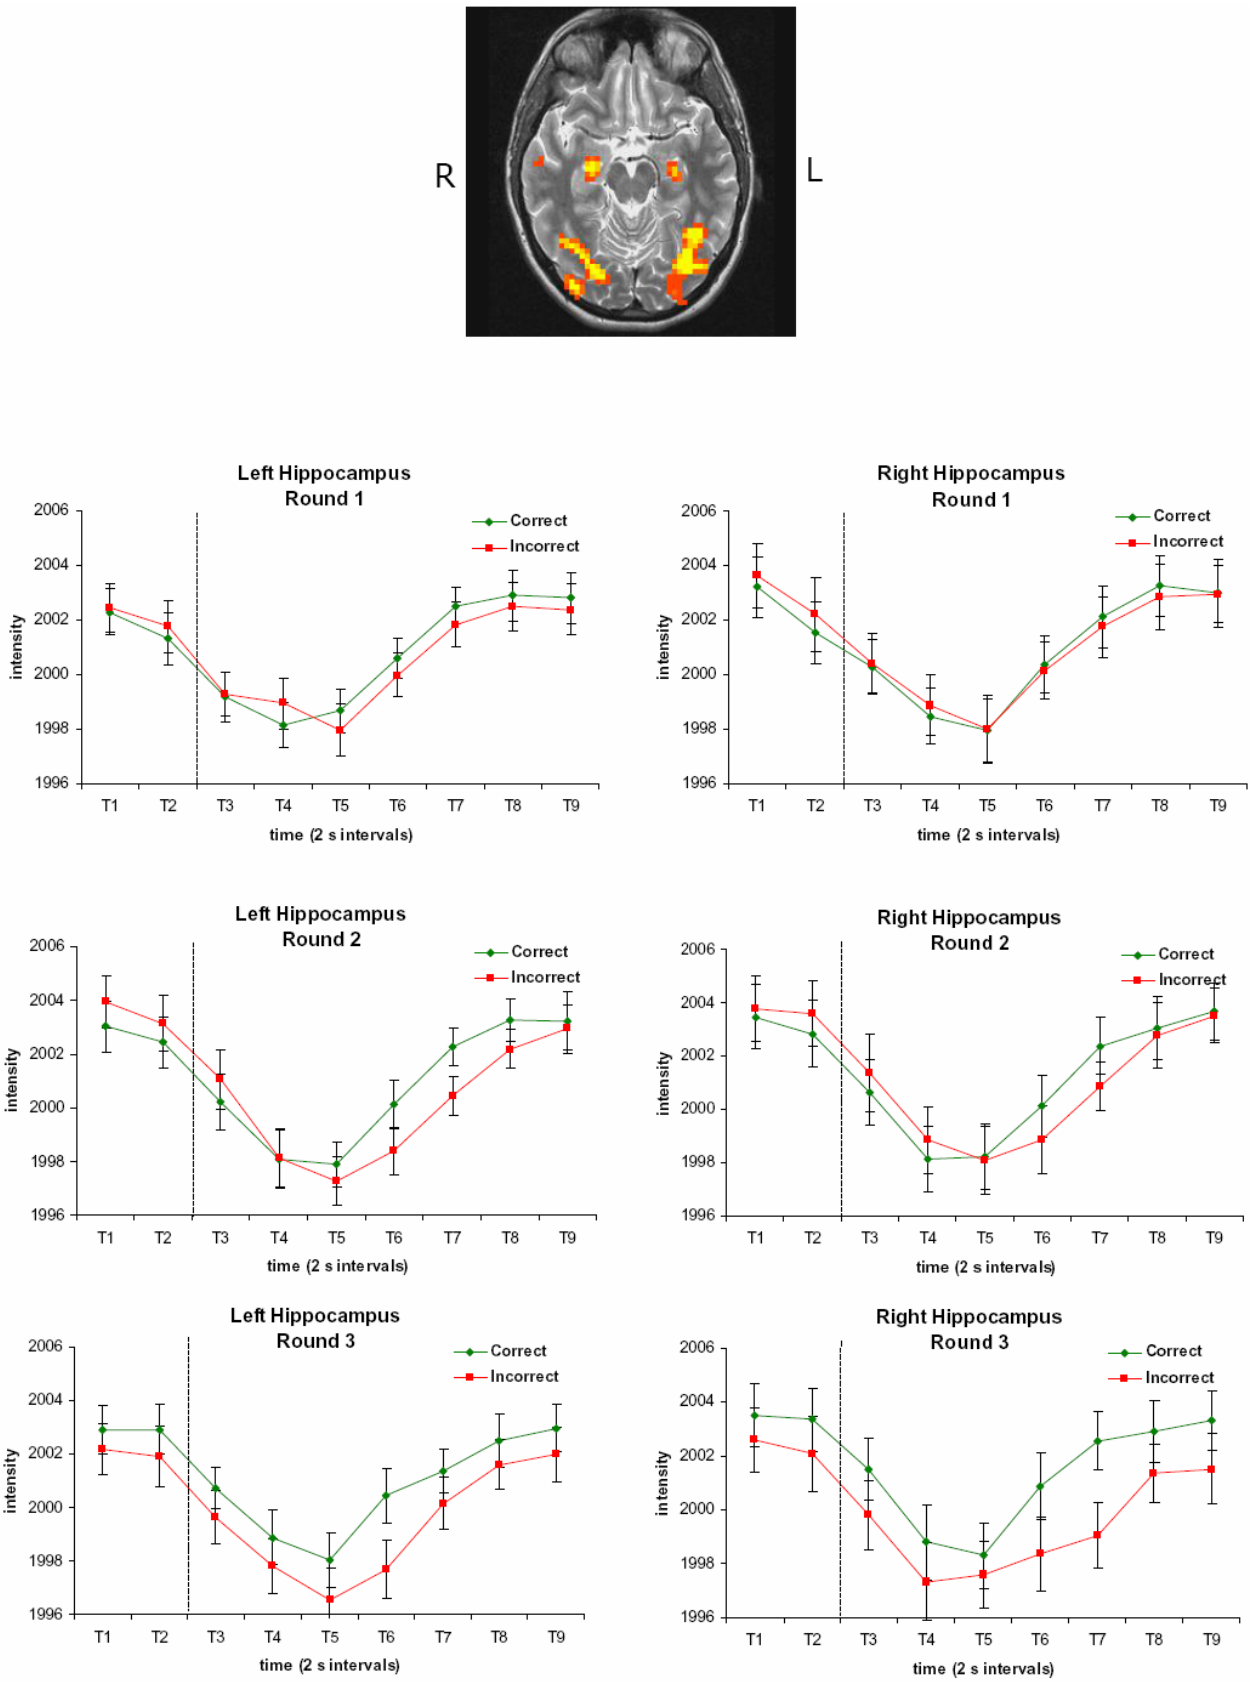

Figure 9. Hippocampal Regions Showing a Main Effect of Time and their Time Courses of Activation. Hippocampal activation clusters displaying a main effect of time are shown, along with the time courses of activation in these regions. The signal decreases from trial onset, and recovers following the feedback display. This recovery is increasingly slower following negative feedback as round increases. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display. The brain image is left-right reversed.

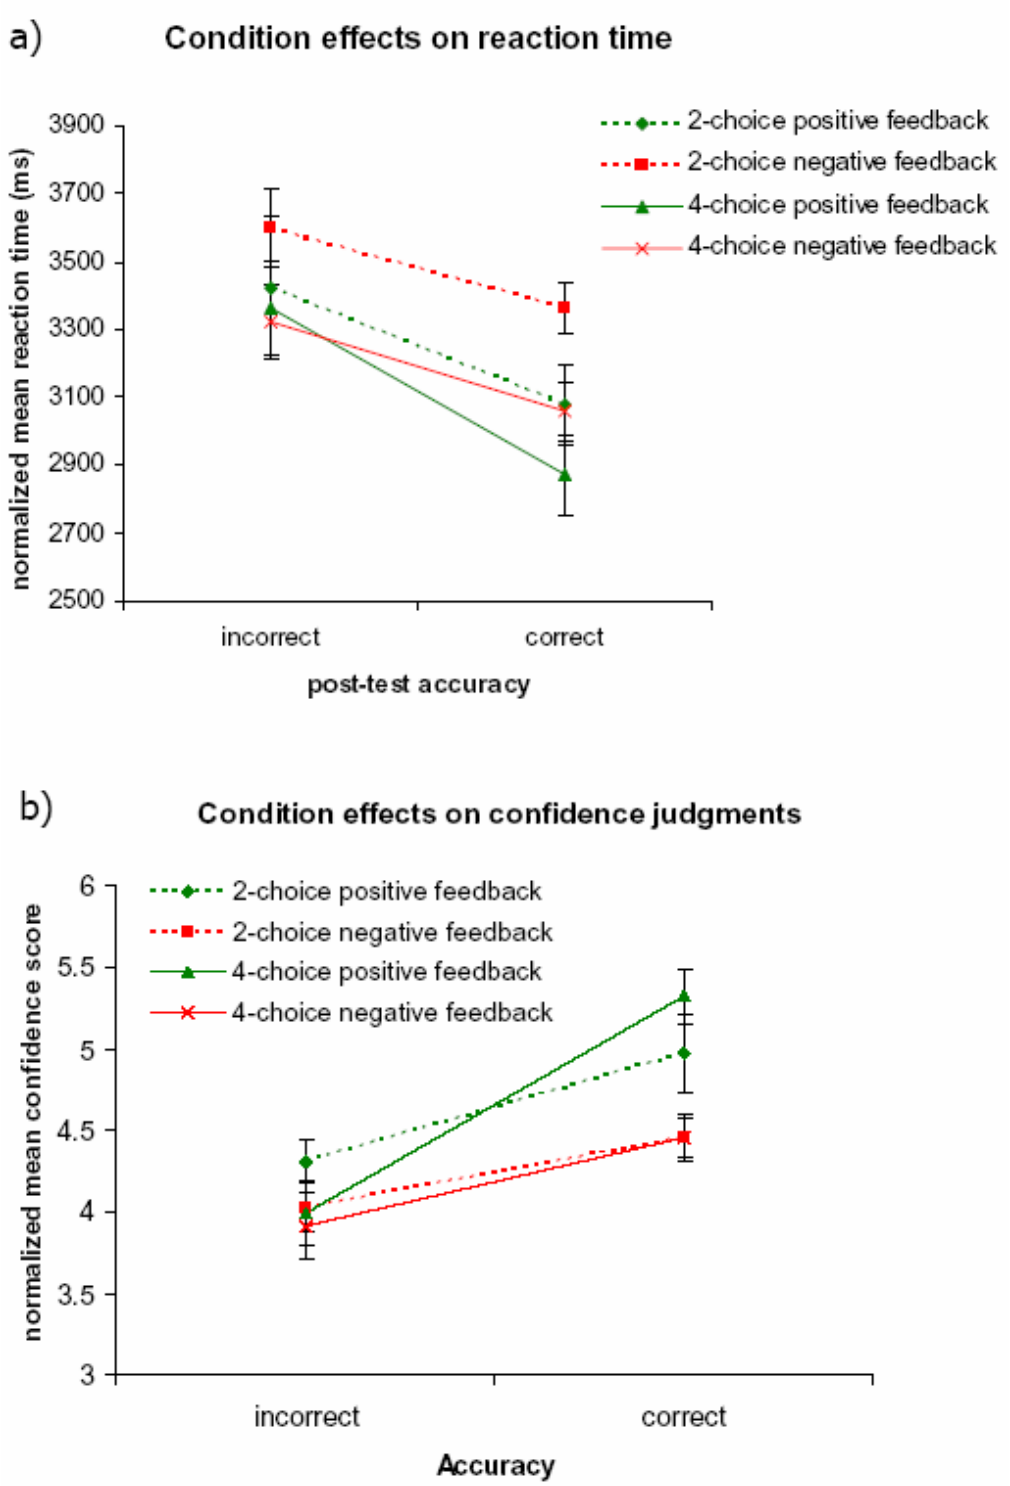

Figure 10. Experiment 2 Behavioral Results from the Immediate Post-test
a) Mean reaction times (normalized to the subject mean) are graphed as a function of post-test accuracy, as well as the trial type and feedback type of the corresponding trial from the scanning session. Participants responded more quickly on accurate trials for which they had received positive feedback during the scan. b) Mean confidence scores (normalized to the subject mean) were higher for accurate trials than inaccurate trials and higher for trials for which participants received positive feedback rather than negative feedback during the scan. For accurate trials, the latter effect was more pronounced for the 4-choice condition than the 2-choice condition.

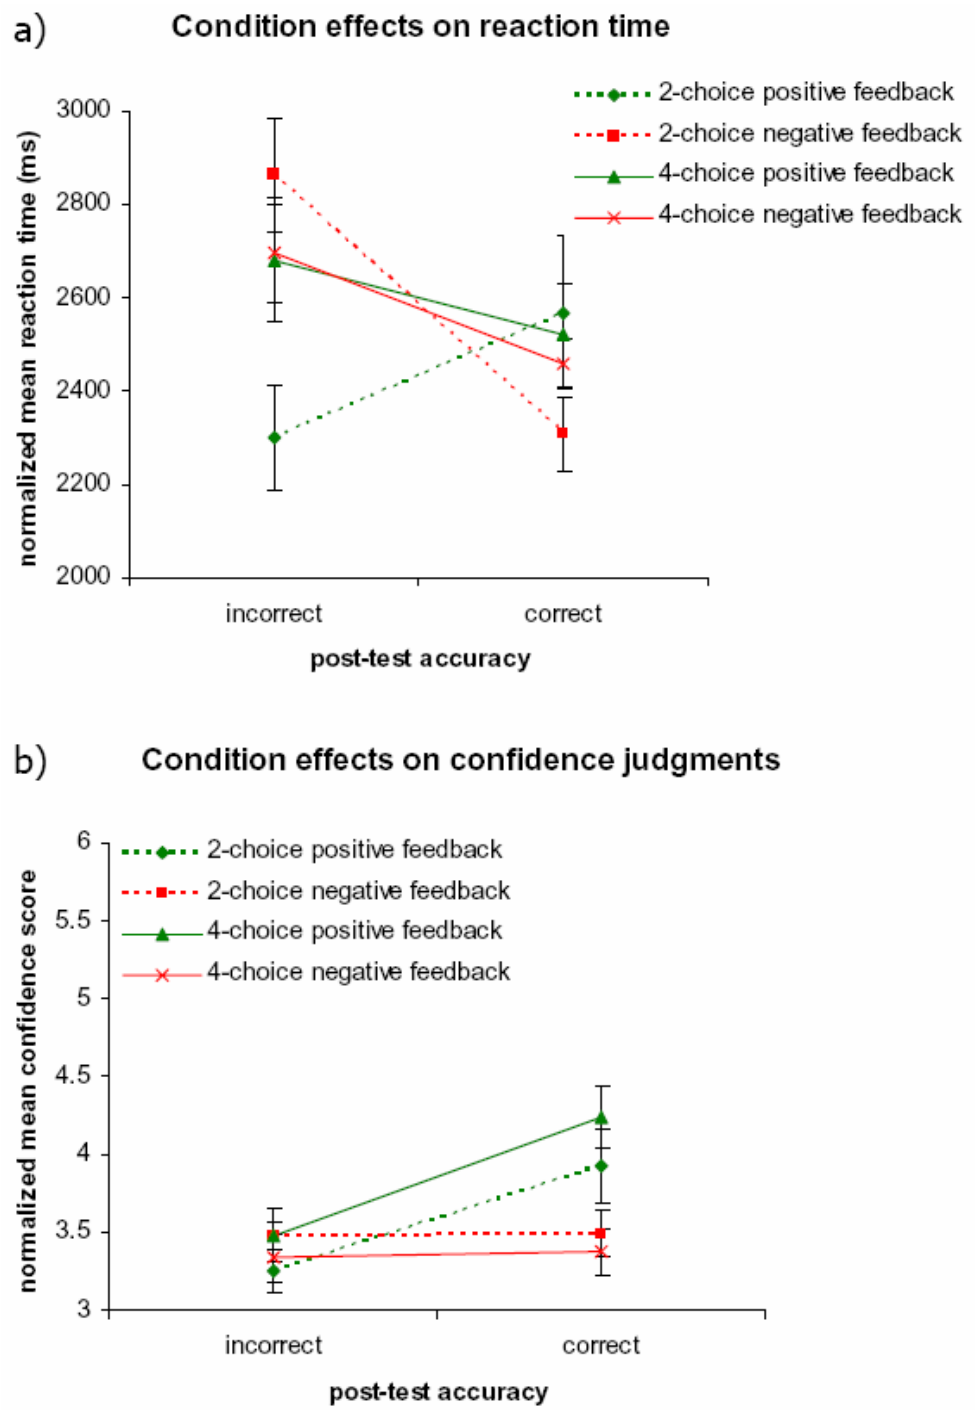

Figure 11. Experiment 2 Behavioral Results from the One-week Follow-up Post-test.
a) Mean reaction times (normalized to the subject mean) are graphed as a function of post-test accuracy, as well as the trial type and feedback type of the corresponding trial from the scanning session. For the 4-choice condition, reaction time did not differ based on feedback type, but was faster for correct than incorrect trials. For the 2-choice condition, reaction time was fastest for incorrect trials for which they received positive feedback during the scan, slowest for incorrect trials for which they received negative feedback, and in between for correct trials. b) Although mean confidence scores (normalized to the subject mean) were generally lower for the follow-up post-test than the immediate post-test, the pattern of results is the same. That is, confidence was higher for accurate trials than inaccurate trials and higher for trials for which participants received positive feedback rather than negative feedback during the scan. For accurate trials, the latter effect was more pronounced for the 4-choice condition than the 2-choice condition.

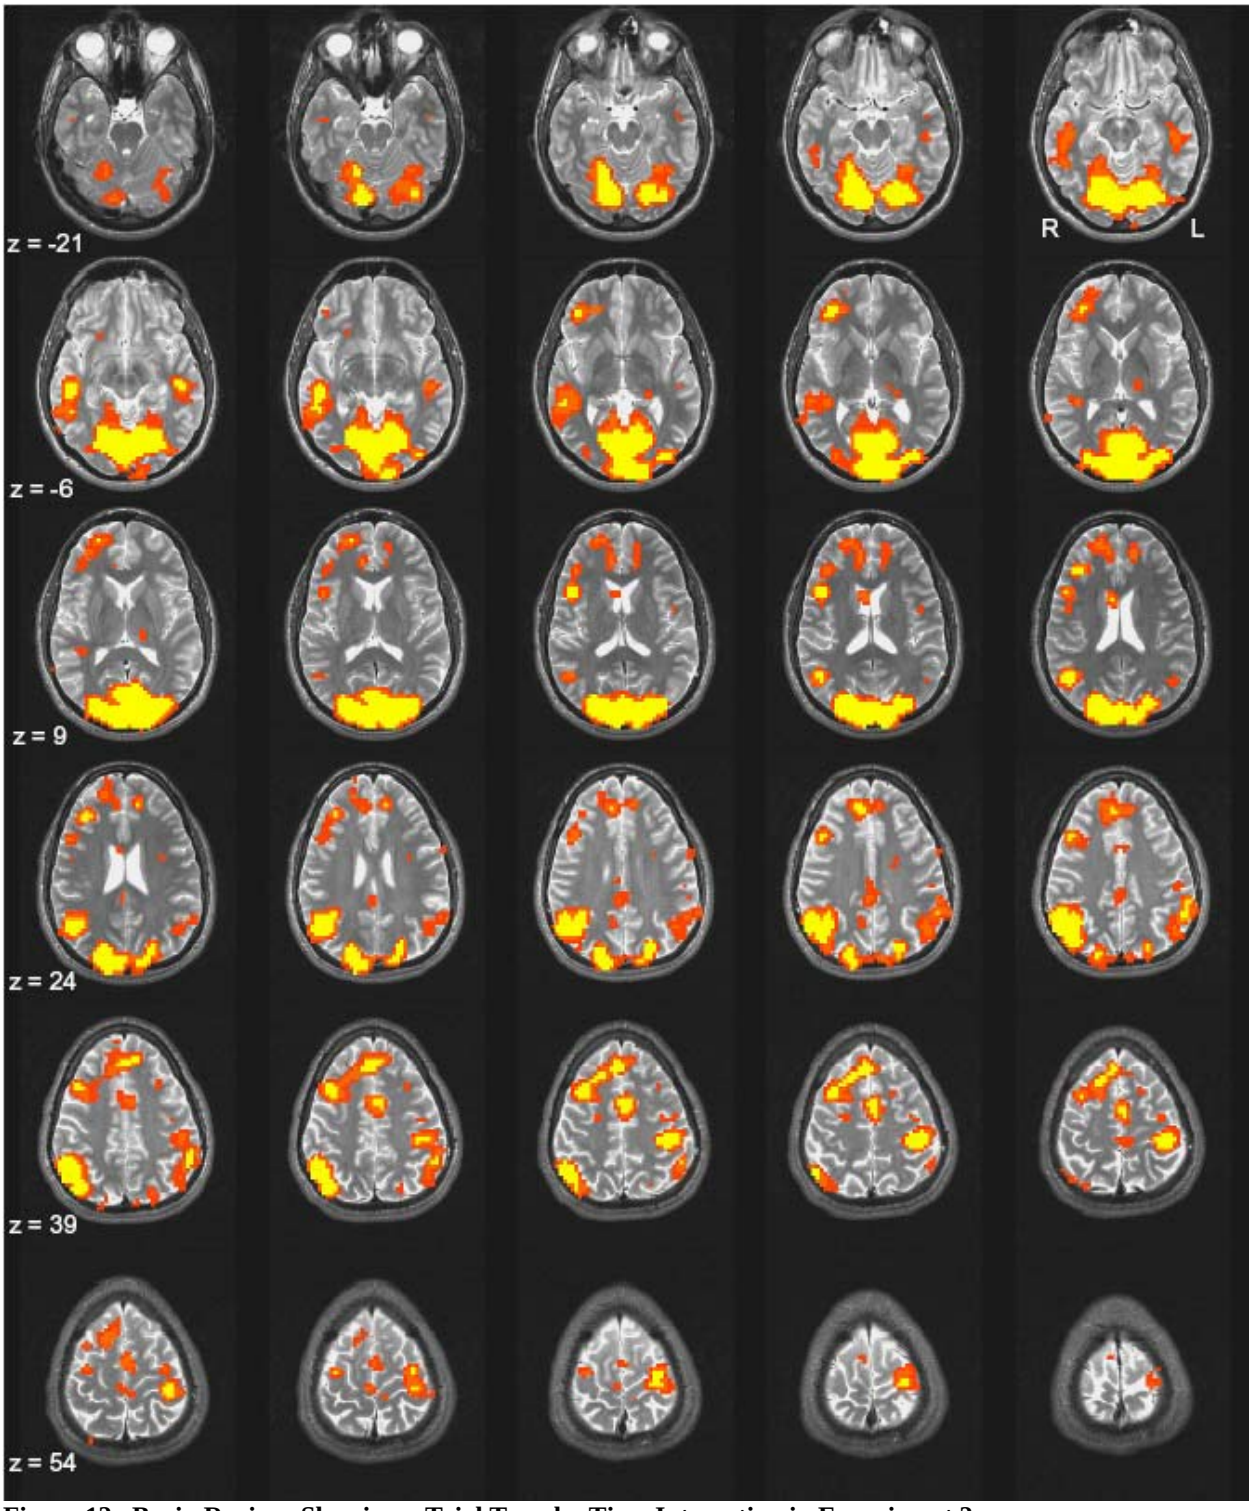

Figure 12. Brain Regions Showing a Trial Type by Time Interaction in Experiment 2.

Voxel clusters displaying a Trial Type by Time interaction in Experiment 2 are shown ($p < 0.001$; contiguity threshold of 4 voxels). Images are left-right reversed.

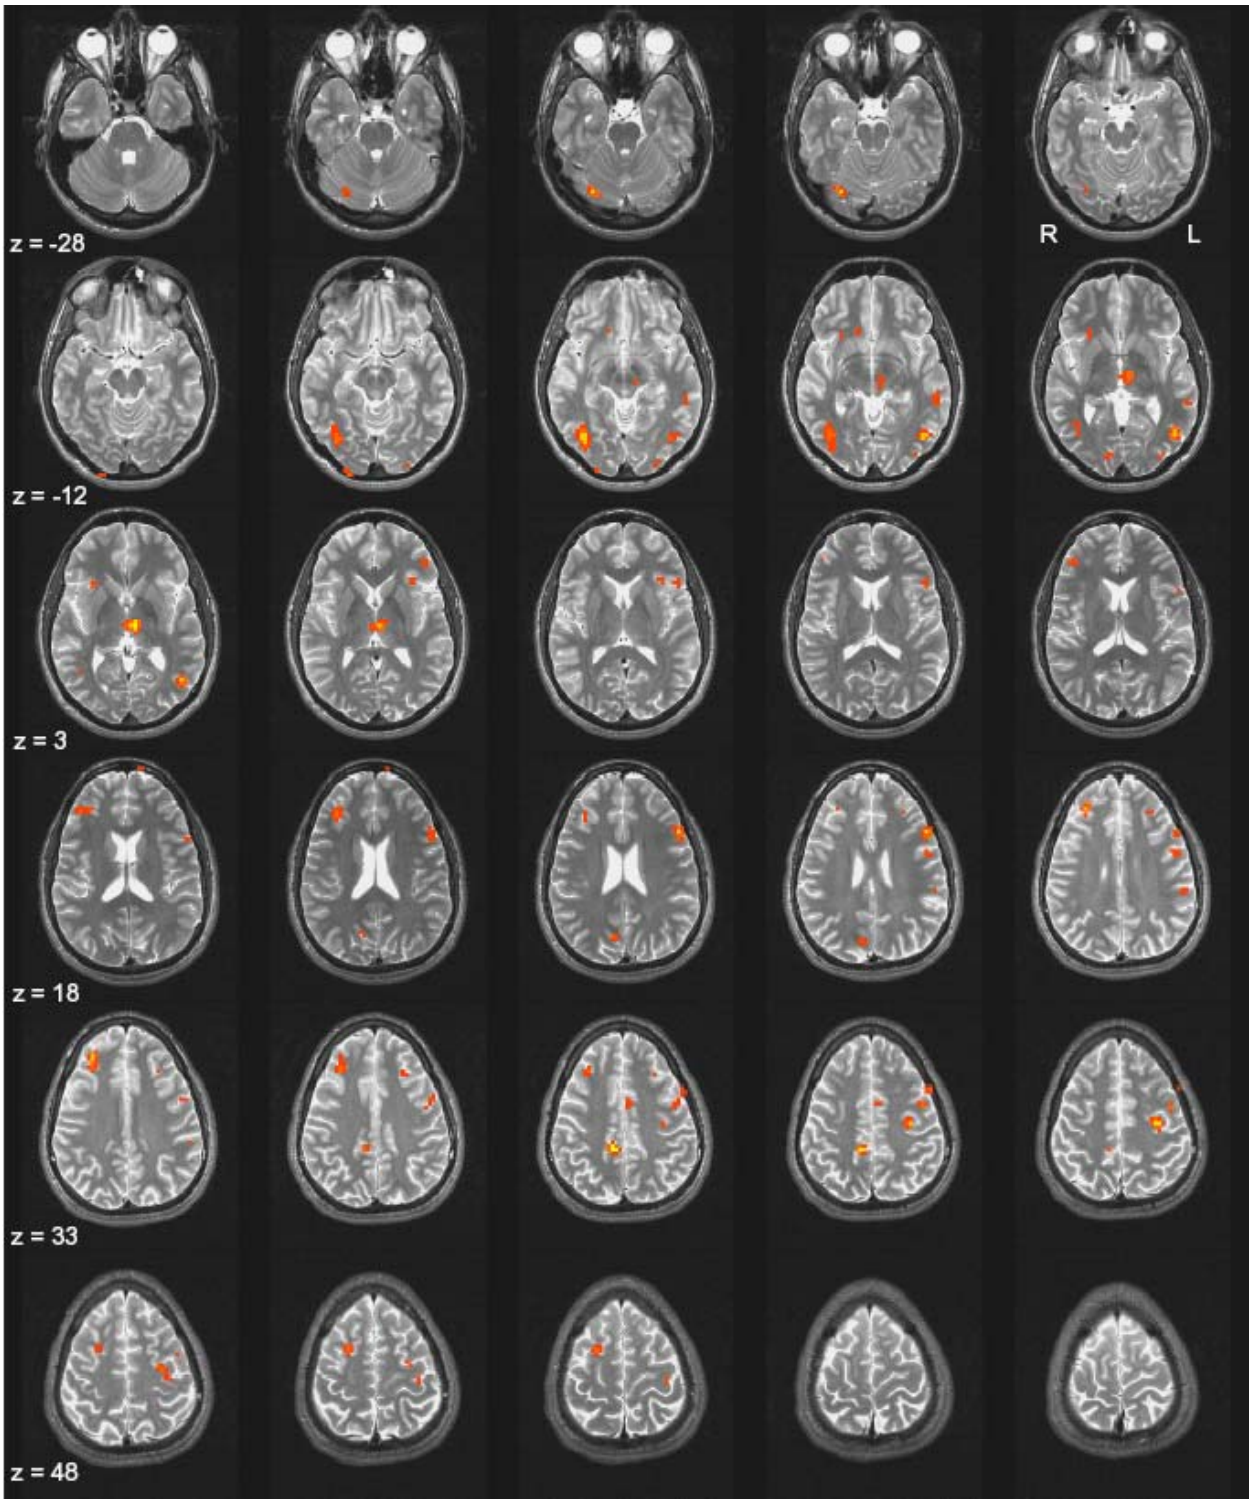

Figure 13. Brain Regions Showing a Feedback Type by Time Interaction in Experiment 2

Voxel clusters displaying a Feedback Type by Time interaction in Experiment 2 are shown ($p < 0.001$; contiguity threshold of 4 voxels). Images are left-right reversed.

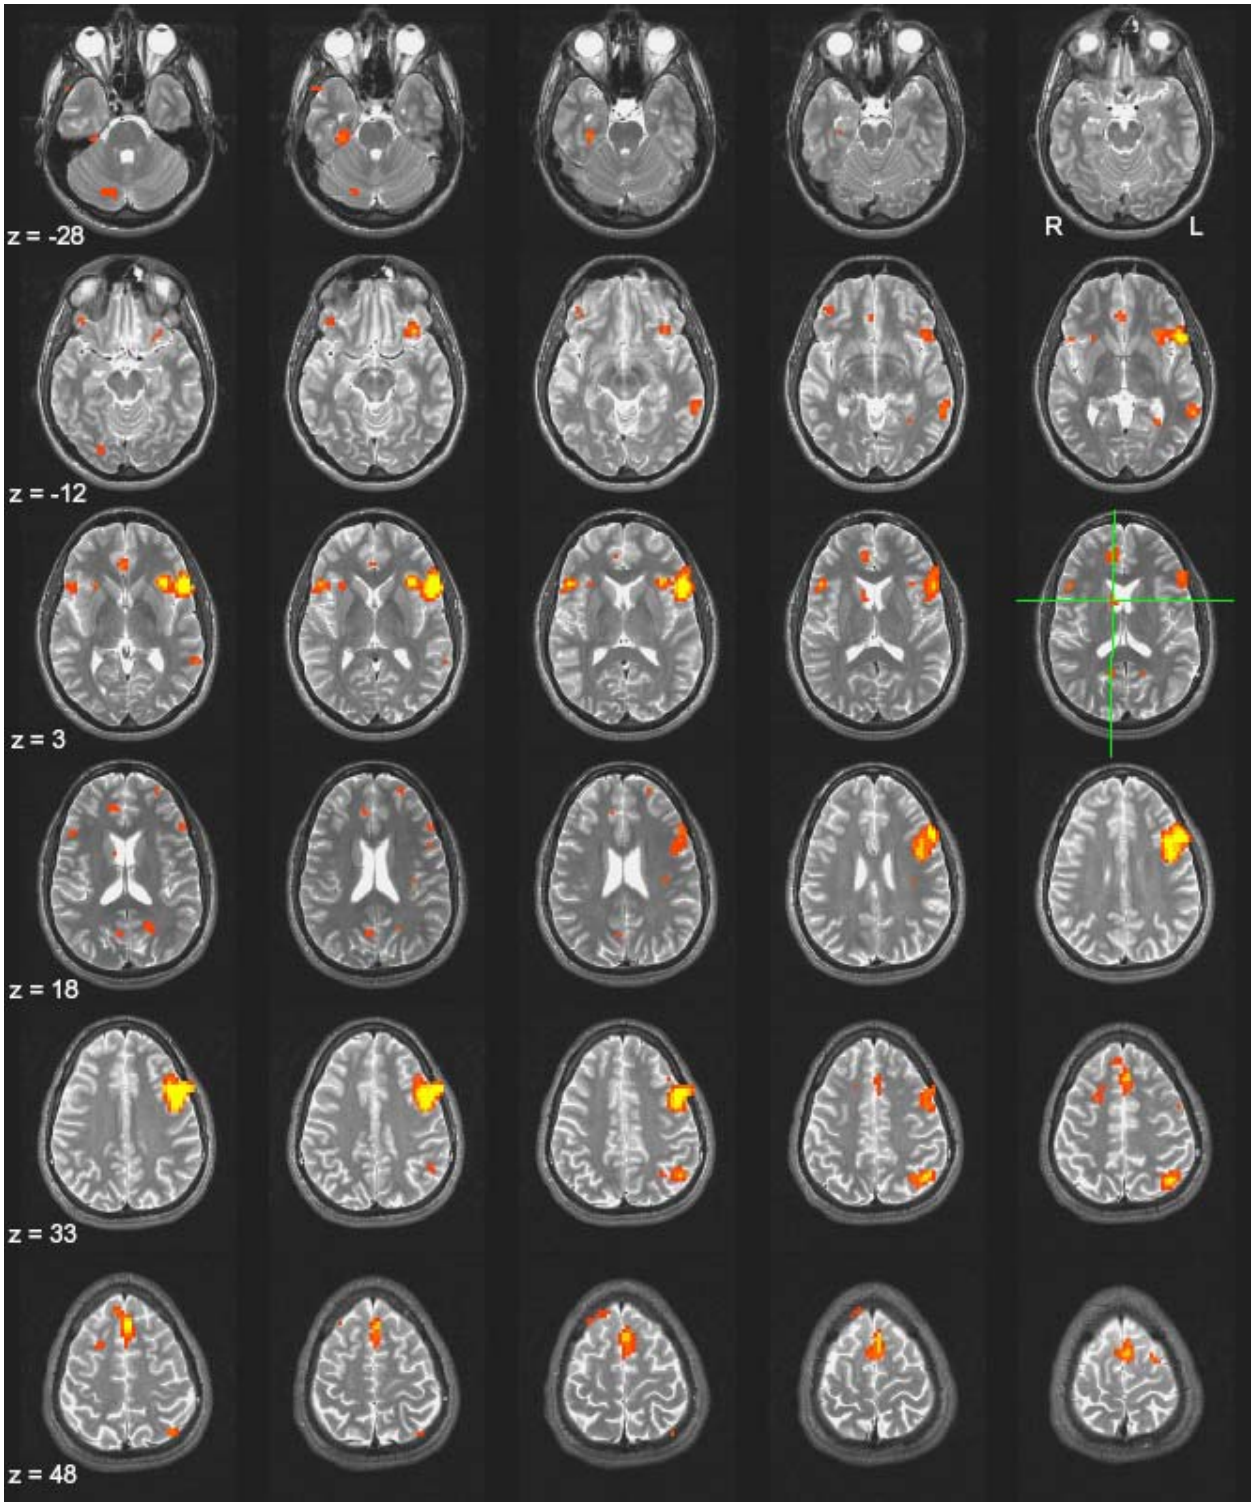

Figure 14. Brain Regions Showing a Trial Type by Feedback Type by Time Interaction in Experiment 2.

Voxel clusters displaying a Trial Type by Feedback Type by Time interaction in Experiment 2 are shown ($p < 0.001$; contiguity threshold of 4 voxels). The green crosshair marks the caudate voxel with the peak F-value. Images are left-right reversed.

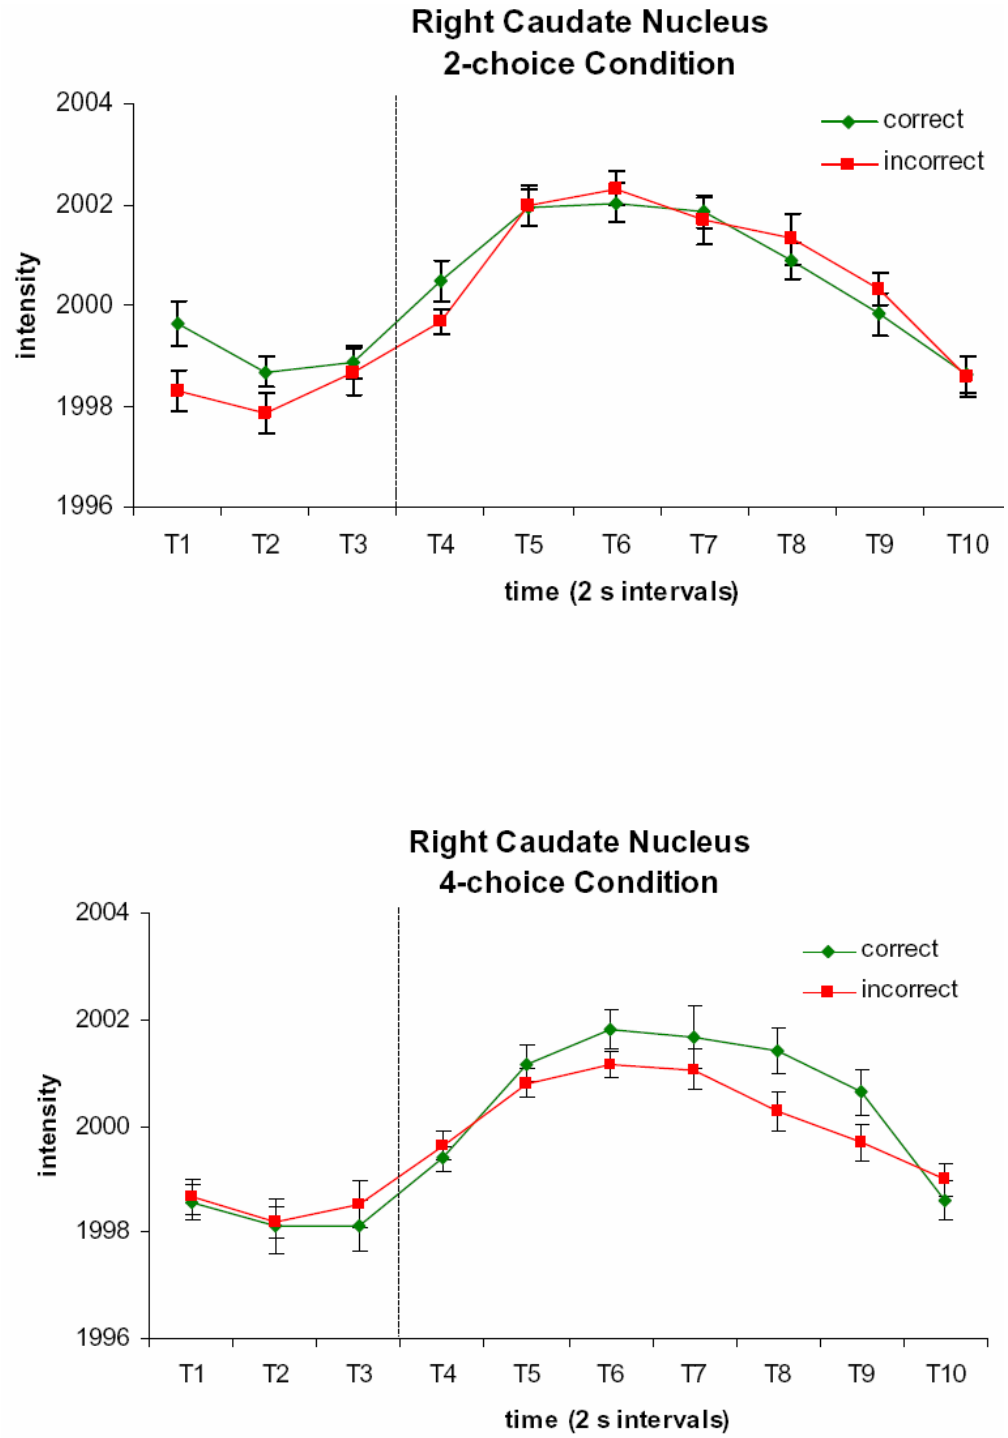

Figure 15. Time Courses of Activation in the Right Caudate Region Showing a Trial Type by Feedback Type by Time Interaction.

The signal is greater following positive feedback than negative feedback in the 4-choice condition but not the 2-choice condition. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display.

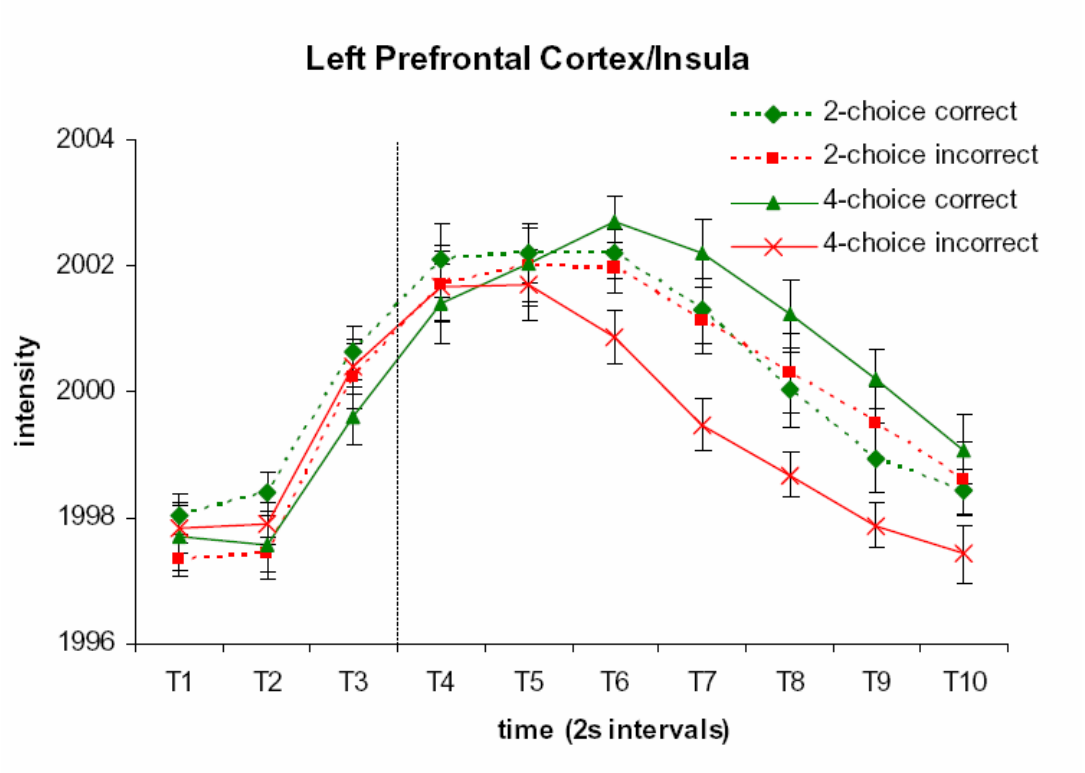

Figure 16. Time Courses of Activation in the Left Prefrontal Cortex/Insula Region Showing a Trial Type by Feedback Type by Time Interaction.

The signal is greatest following positive feedback in the 4-choice condition and least strong following negative feedback in the 4-choice condition, with the signal following both positive and negative feedback displaying an intermediate level of activation. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display.

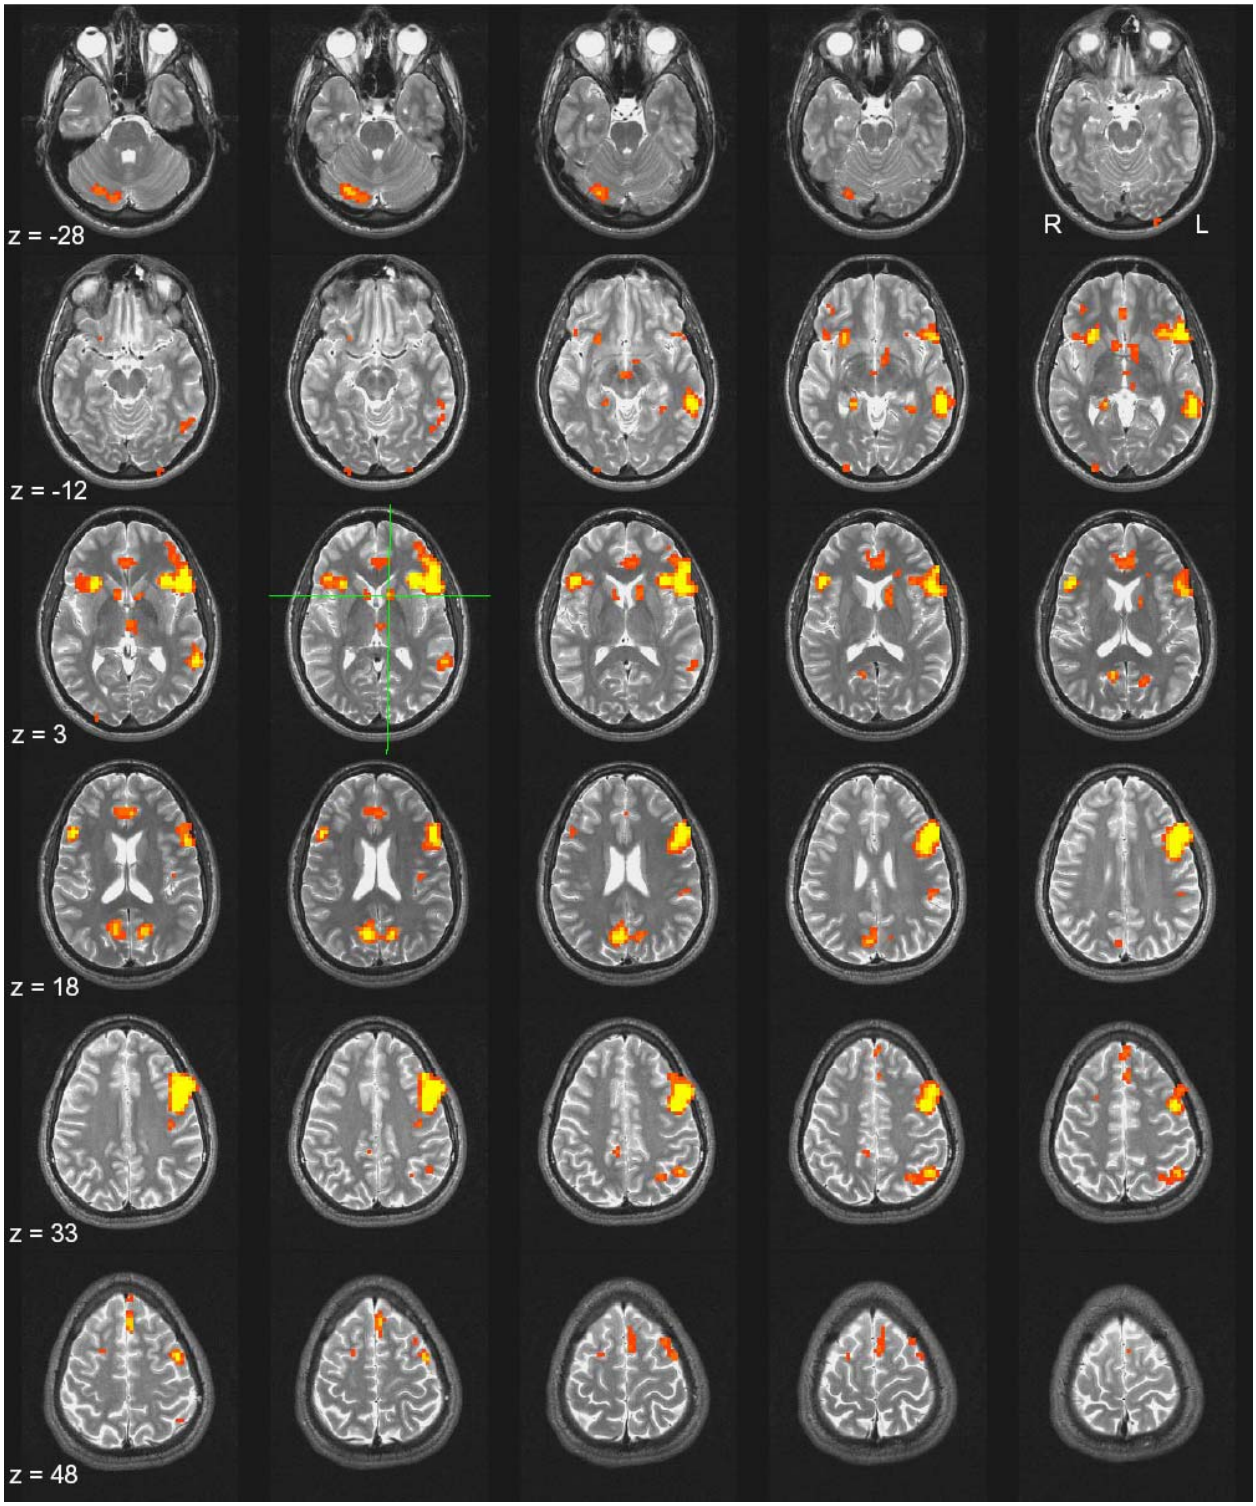

Figure 17. Brain Regions Showing a Feedback Type by Time Interaction in the 4-choice Condition of Experiment 2.

Voxel clusters displaying a Feedback Type by Time interaction in the 4-choice condition of Experiment 2 are shown ($p < 0.001$; contiguity threshold of 4 voxels). The green crosshair marks the caudate voxel with the peak F-value. Images are left-right reversed.

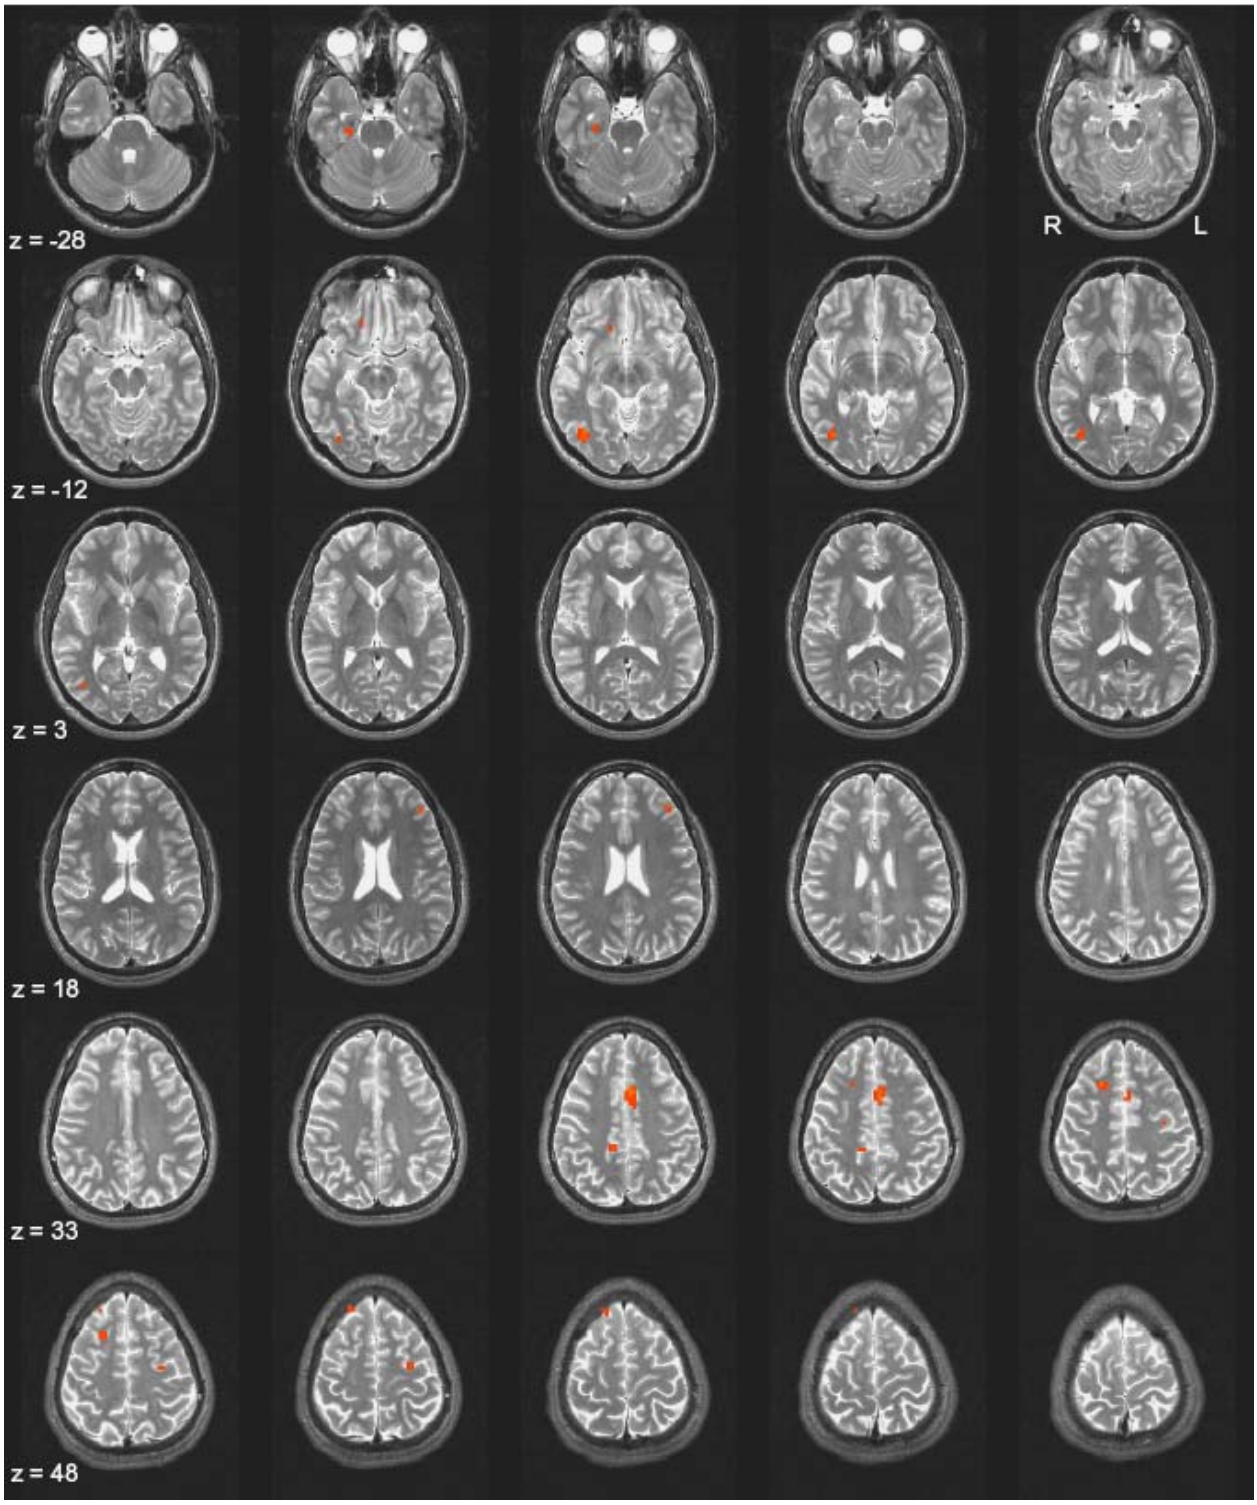

Figure 18. Brain Regions Showing a Feedback Type by Time Interaction in the 2-choice Condition of Experiment 2.

Voxel clusters displaying a Feedback Type by Time interaction in the 2-choice condition of Experiment 2 are shown ($p < 0.001$; contiguity threshold of 4 voxels). Images are left-right reversed.

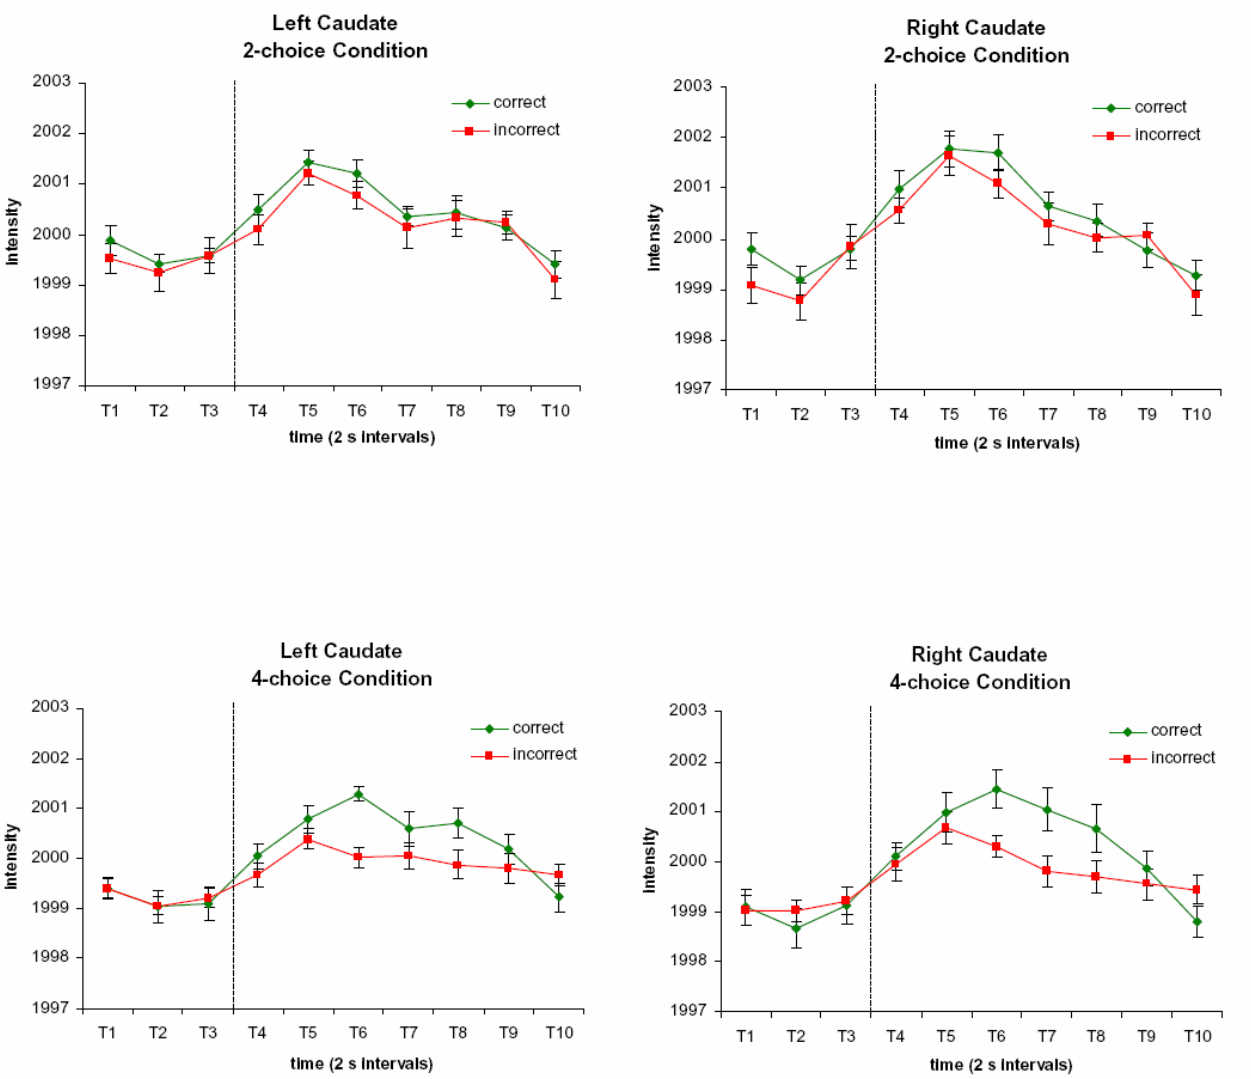

Figure 19. Time Courses of Activation in the Caudate Regions Showing a Feedback Type by Time Interaction for the 4-choice Condition.

The signal is greater following positive feedback than negative feedback in the 4-choice condition but not the 2-choice condition. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display.

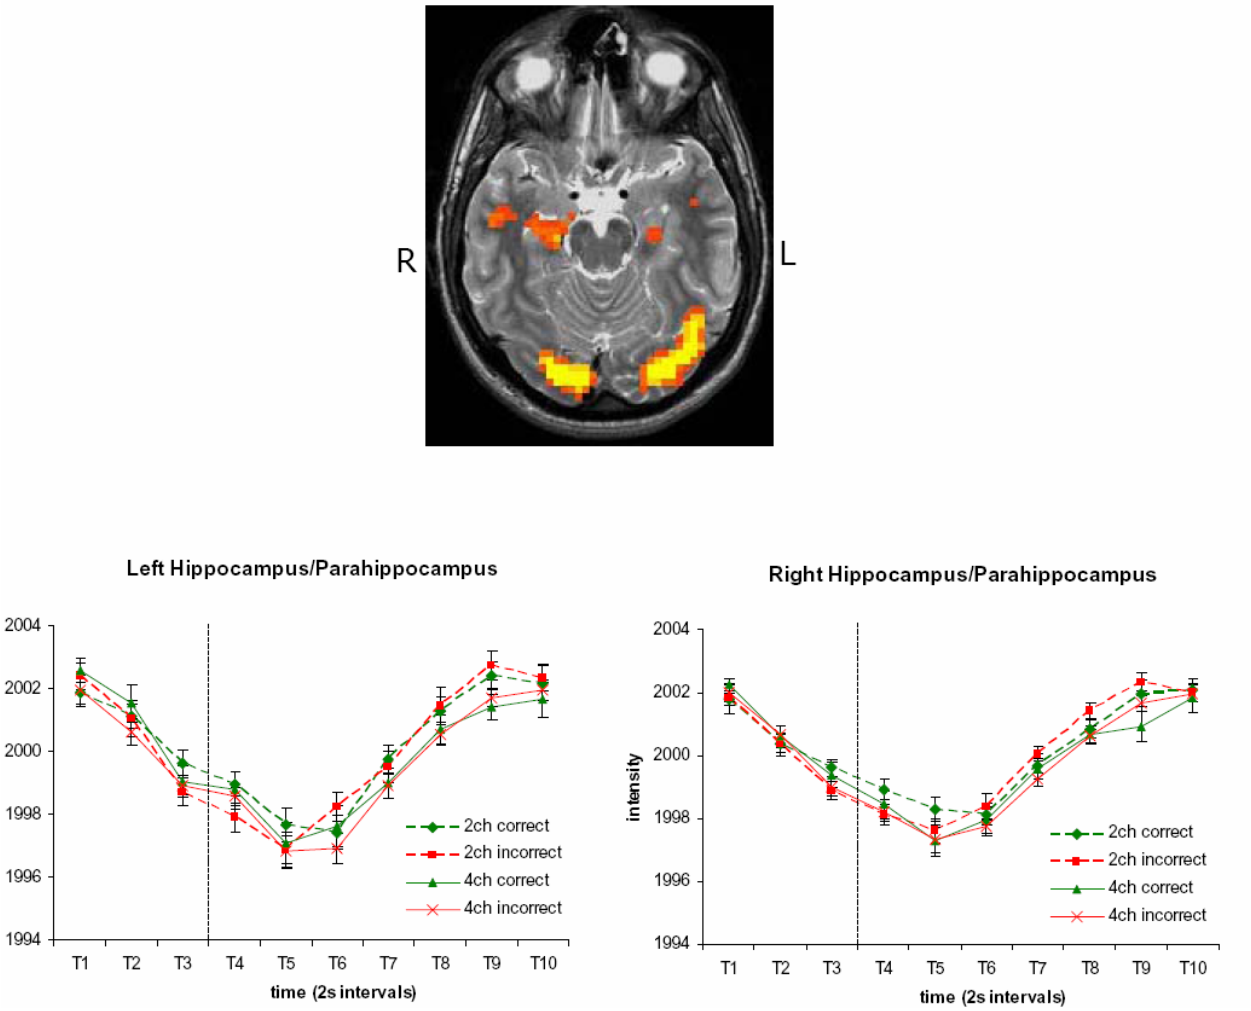

Figure 20. Medial Temporal Lobe Regions Showing a Main Effect of Time and their Time Courses of Activation.

Activation clusters in the hippocampus/parahippocampus displaying a main effect of time are shown, along with the time courses of activation in these regions. As in Experiment 1, the signal decreases from trial onset, and recovers following the feedback display. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display. The brain image is left-right reversed.

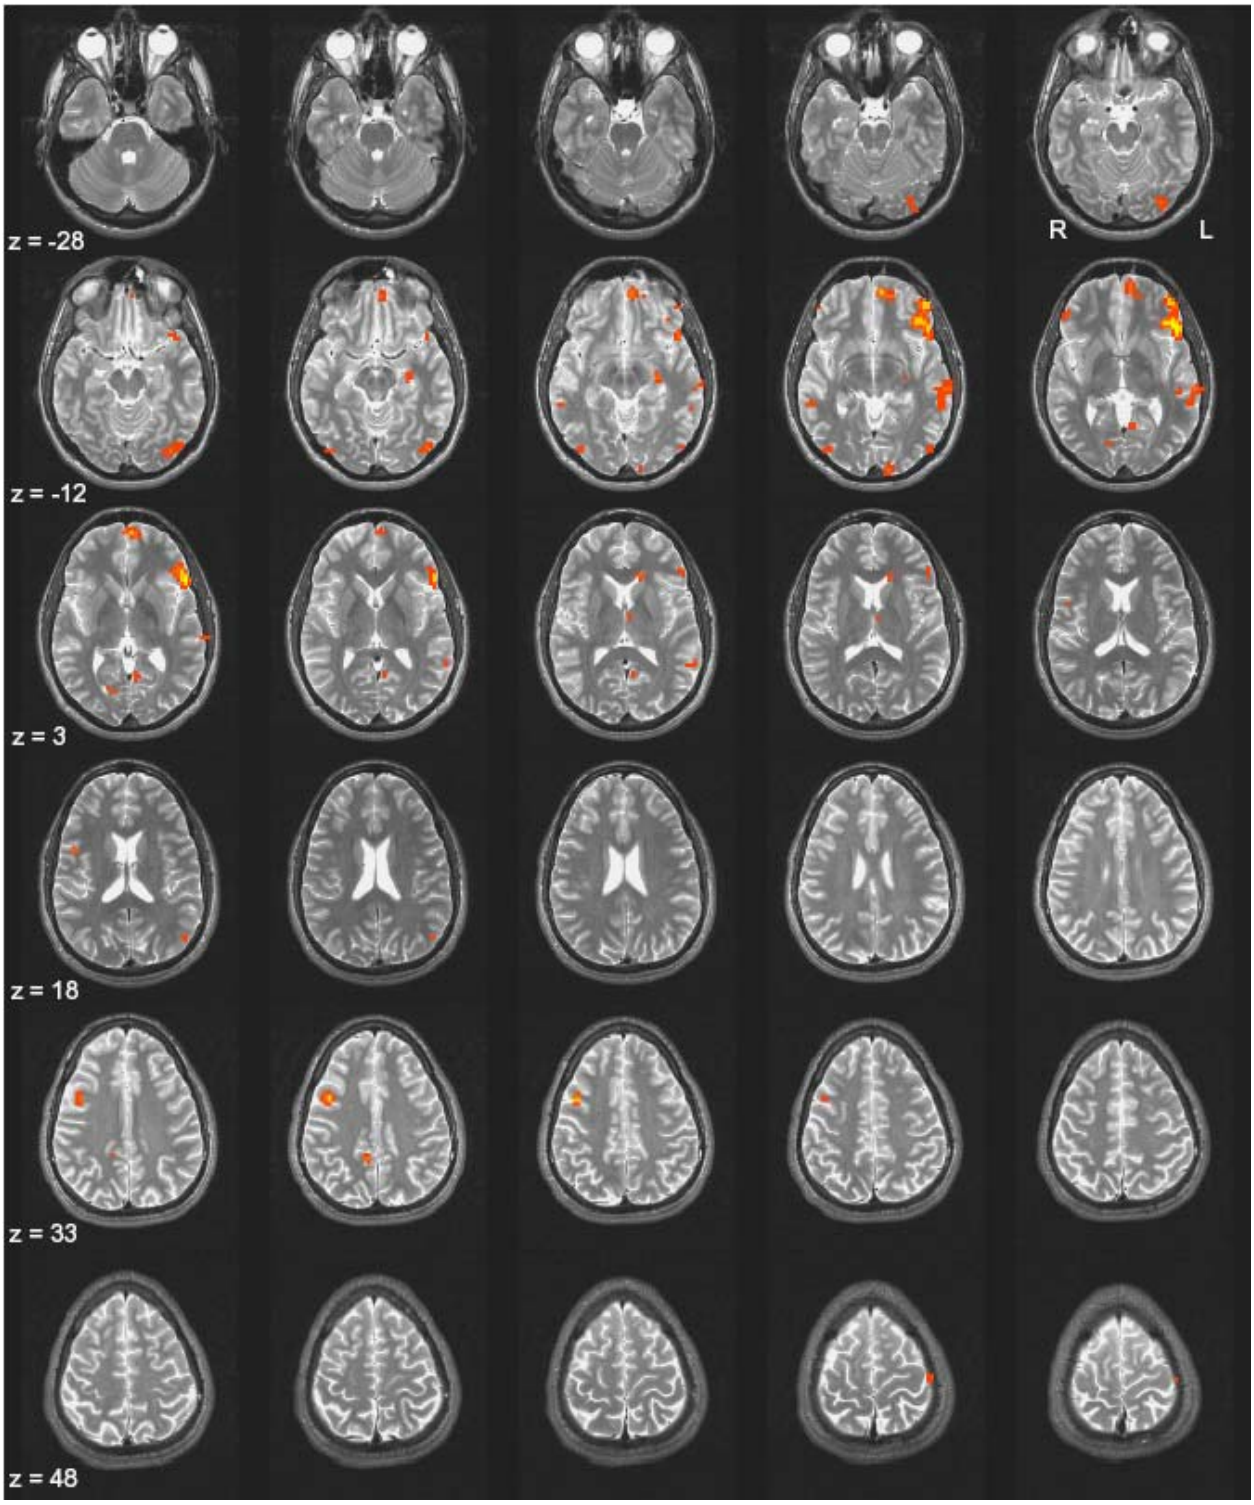

Figure 21. Brain Regions Showing a Subsequent Memory Effect in Experiment 2.

Voxel clusters displaying differential activation on trials remembered with high confidence versus incorrect trials in Experiment 2 are shown ($p < 0.001$; contiguity threshold of 4 voxels). Images are left-right reversed.

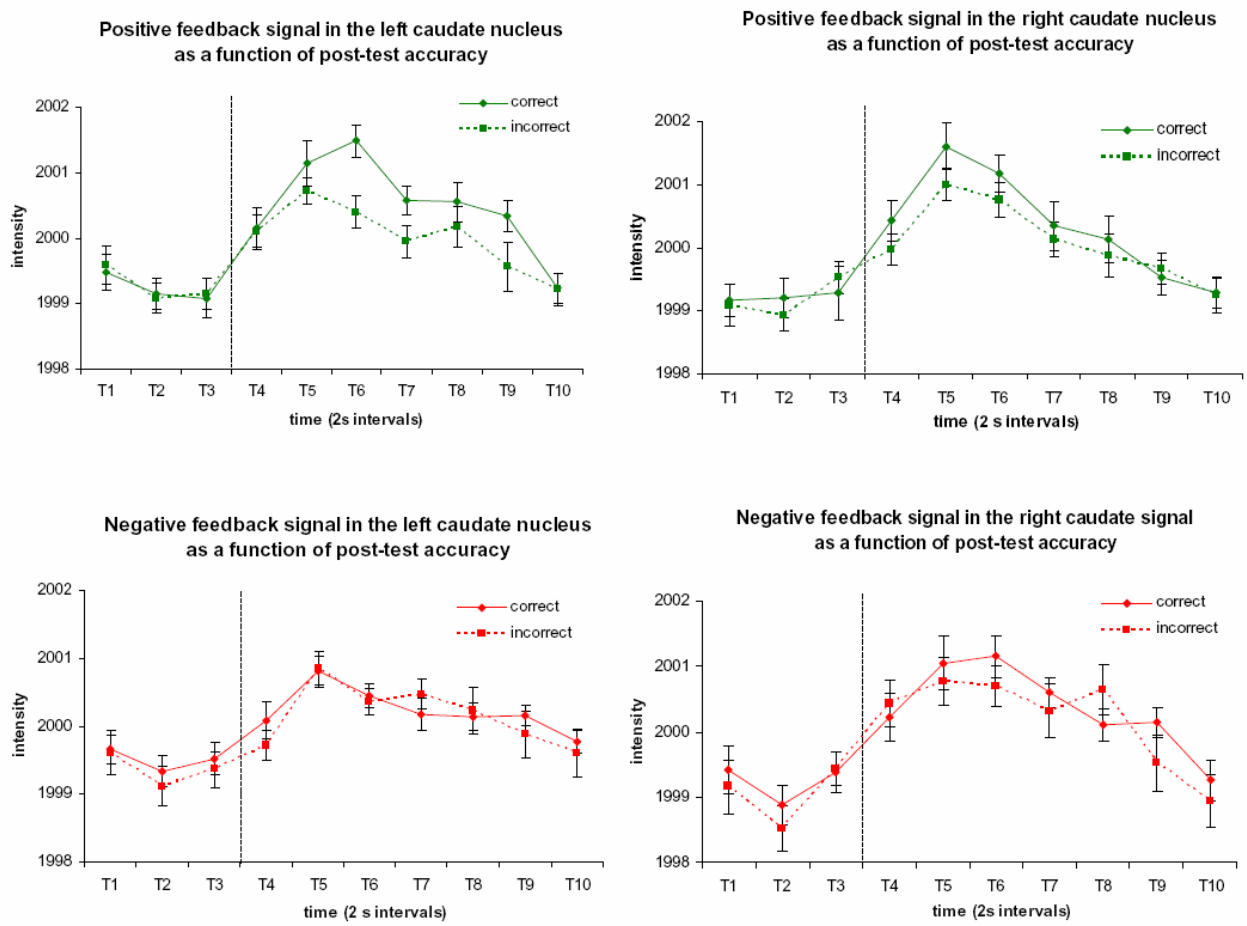

Figure 22. Time Courses of Activation in the Caudate as a Function of Post-test Accuracy.

The signal in the caudate regions identified as showing a Feedback by Time Effect in Experiment 2 is plotted separately for trials associated with correct and incorrect performance on the immediate post-test. In the left caudate, the signal following positive feedback is greater for trials that will later be remembered than trials that will be forgotten. This effect is not significant in the right caudate nucleus. No differences are observed in the signal following negative feedback. Each time period (T1, T2, etc.) respresents a 2 s image acquisition. The dotted line indicates the onset of the feedback display. The brain image is left-right reversed.

BIBLIOGRAPHY

Adcock, R. A., Thangavel, A., Whitfield-Gabrieli, S., Knutson, B., & Gabrieli, J. D. E. (2006). Reward-motivated learning: Mesolimbic activation precedes memory formation. *Neuron, 50*(3), 507.

Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci, 9*, 357-381.

Balleine, B. W., & Dickinson, A. (1998). Goal-directed instrumental action: Contingency and incentive learning and their cortical substrates. *Neuropharmacology, 37*(4-5), 407-419.

Bar-Gad, I., Morris, G., & Bergman, H. (2003). Information processing, dimensionality reduction and reinforcement learning in the basal ganglia. *Prog Neurobiol, 71*, 439-473.

Barto, A. G. (1995). Adaptive critics and the basal ganglia. In J. C. Houk, J. L. Davis & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 215-232). Cambridge, Massachusetts: MIT Press.

Bayley, P. J., Franscino, J. C., & Squire, L. R. (2005). Robust habit learning in the absence of awareness and independent of the medial temporal lobe. *Nature, 436*, 550-553.

Berns, G. S., McClure, S. M., Pagnoni, G., & Montague, P. R. (2001). Predictability modulates human brain response to reward. *J Neurosci, 21*(8), 2793-2798.

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

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychol Rev, 108*, 624-652.

Breiter, H. C., Aharon, I., Kahneman, D., Dale, A., & Shizgal, P. (2001). Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron, 30*, 619-639.

Brewer, J. B., Zhao, Z., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. E. (1998). Making memories: Brain activity that predicts how well visual experience will be remembered. *Science, 281*, 1185-1187.

Cardinal, R. N., Parkinson, J. A., Hall, J., & Everitt, B. J. (2002). Emotion and motivation: The role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci Biobehav Rev, 26*(3), 321-352.

Coltheart, M. (1981). The MRC psycholinguistic database. *Q J Exp Psychol A, 33*, 497-505.

Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res, 29*(3), 162-173.

Dayan, P., & Balleine, B. W. (2002). Reward, motivation, and reinforcement learning. *Neuron, 36*, 285-298.

Delgado, M. R., Frank, R. H., & Phelps, E. A. (2005a). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nat Neurosci, 8*(11), 1611-1618.

Delgado, M. R., Miller, M. M., Inati, S., & Phelps, E. A. (2005b). An fMRI study of reward-related probability learning. *NeuroImage, 24*(3), 862-873.

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. *J Neurophysiol, 84*(6), 3072-3077.

Delgado, M. R., Stenger, V. A., & Fiez, J. A. (2004). Motivation-dependent responses in the human caudate nucleus. *Cereb Cortex, 14*, 1022-1030.

Devan, B. D., McDonald, R. J., & White, N. M. (1999). Effects of medial and lateral caudate-putamen lesions on place- and cue- guided behaviors in the water maze: Relation ot thigmotaxis. *Behav Brain Res, 100*, 5-14.

Dickinson, A., & Balleine, B. (1994). Motivational control of goal-directed action. *Anim Learn Behav, 22*, 1-18.

Dickinson, A., & Balleine, B. (2002). The role of learning in the operation of motivational systems. In C. R. Gallistel (Ed.), *Stevens' handbook of experimental psychology: Learning, motivation, and emotion* (Vol. 3, pp. 497-534). New York: Wiley and Sons.

Eichenbaum, H. (2002). *The cognitive neuroscience of memory*. New York: Oxford University Press.

Eldridge, L. L., Engel, S. A., Zeineh, M. M., Bookheimer, S. Y., & Knowlton, B. J. (2005). A dissociation of encoding and retrieval processes in the human hippocampus. *J Neurosci, 25*(13), 3280-3286.

Elliott, R., Dolan, R. J., & Frith, C. D. (2000a). Dissociable functions in the medial and lateral orbitofrontal cortex: Evidence from human neuroimaging studies. *Cereb Cortex, 10*, 308-317.

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

Elliott, R., Frith, C. D., & Dolan, R. J. (1997). Differential neural response to positive and negative feedback in planning and guessing tasks. *Neuropsychologia, 35*(10), 1395-1404.

Elliott, R., Sahakian, B. J., Michael, A., Paykel, E. S., & Dolan, R. J. (1998). Abnormal neural response to feedback on planning and guessing tasks in patients with unipolar depression. *Psychol Med, 28*(3), 559-571.

Fellows, L. K., & Farah, M. J. (2005). Different underlying impairments in decision-making following ventromedial and dorsolateral frontal lobe damage in humans. *Cereb Cortex, 15*, 58-63.

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

Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., & Noll, D. C. (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): Use of a cluster-size threshold. *Magn Reson Med, 33*(5), 636-647.

Goodman, J. S. (1998). The interactive effects of task and external feedback on practice performance and learning. *Organiz Behav Hum Decis Process, 76*(3), 223-252.

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. *J Neurosci, 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. *J Neurosci, 26*(32), 8368-8376.

Hampton, A. N., Bossaerts, P., & O'Doherty, J. P. (2006). The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *J Neurosci, 26*(32), 8360-8367.

Haruno, M., Kuroda, T., Doya, K., Toyama, K., Kimura, M., Samejima, K., et al. (2004). A neural correlate of reward-based behavioral learning in caudate nucleus: A functional magnetic resonance imaging study of a stochastic decision task. *J Neurosci, 24*(7), 1660-1665.

Hikosaka, O., Sakamoto, M., & Usui, S. (1989). Functional properties of monkey caudate neurons. III. Activities related to expectation of target and reward. *J Neurophysiol, 61*, 814-832.

Hill, W. F. (1997). *Learning* (6th ed.). New York: Addison-Wesley Educational Publishers Inc.

Holroyd, C., & Coles, M. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev, 109*(4), 679-709.

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R. B., Coles, M. G. H., et al. (2004). Dorsal anterior cingulate cortex shows fmri response to internal and external error signals. *Nat Neurosci, 7*(5), 497-498.

Kantak, K. M., Black, Y., Valencia, E., Green-Jordan, K., & Eichenbaum, H. B. (2002). Stimulus-response functions of the lateral dorsal striatum and regulation of behavior studied in a cocaine maintenance/cue reinstatement model in rats. *Psychopharmacology, 161*, 278-287.

Kantak, K. M., Green-Jordan, K., Valencia, E., Kremin, T., & Eichenbaum, H. B. (2001). Cognitive task performance following lidocaine-induced inactivation of different sites within the basolateral amygdala and dorsal striatum. *Behav Neurosci, 115*, 589-601.

Kawagoe, R., Takikawa, Y., & Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nat Neurosci, 1*(5), 411-416.

Kluger, A. N., & DeNisi, A. (1996). The effects of feedback intervention on performance: A historical review, a meta-analysis and a preliminary feedback intervention theory. *Psychol Bull, 119*, 254-284.

Knowlton, B. J., Mangles, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science, 273*, 1399-1402.

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

Landauer, T. K., Foltz, P. W., & Laham, D. (1998). Introduction to latent semantic analysis. *Discourse Processes, 25*, 259-284.

Lauwereyns, J., Takikawa, Y., Kawagoe, R., Kobayashi, S., Koizumi, M., Coe, B., et al. (2002a). Feature-based anticipation of cues that predict reward in monkey caudate nucleus. *Neuron, 33*(3), 463-473.

Lauwereyns, J., Watanabe, K., Coe, B., & Hikosaka, O. (2002b). A neural correlate of response bias in monkey caudate nucleus. *Nature, 418*, 413-417.

Law, J. R., Flanery, M. A., Wirth, S., Yanike, M., Smith, A. C., Frank, L. M., et al. (2005). Functional magnetic resonance imaging activity during the gradual acquisition and expression of paired-associate memory. *J Neurosci, 25*(24), 5720-5729.

Lisman, J. E., & Grace, A. A. (2005). The hippocampal-VTA loop: Controlling the entry of information into long-term memory. *Neuron, 46*(5), 703-713.

Maguire, E. A., Frackowiak, R. S., & Frith, C. D. (1997). Recalling routes around London: Activation of the right hippocampus in taxi drivers. *J Neurosci, 17*(18), 7103-7110.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev, 102*, 419-457.

McClelland, J. L., Thomas, A. G., McCandliss, B. D., & Fiez, J. A. (1999). Understanding failures of learning: Hebbian learning, competition for representational space, and some preliminary experimental data. In J. Reggia, E. Ruppin & D. Glanzman (Eds.), *Progress in brain research* (Vol. 121, pp. 75-80). Amsterdam: Elsevier.

McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron, 38*, 339-346.

McDonald, R. J., & White, N. M. (1993). A triple dissociation of memory systems: Hippocampus, amygdala, and dorsal striatum. *Behav Neurosci, 107*, 3-22.

Meltzer, J. A., & Constable, R. T. (2005). Activation of human hippocampal formation reflects success in both encoding and cued recall of paired associates. *NeuroImage, 24*(2), 384-397.

Middleton, F. A., & Strick, P. L. (2002). Basal-ganglia 'projections' to the prefrontal cortex of the primate. *Cereb Cortex, 12*(9), 926-935.

Monchi, O., Petrides, M., Doyon, J., Postuma, R. B., Worsley, K., & Dagher, A. (2004). Neural bases of set-shifting deficits in Parkinson's disease. *J Neurosci, 24*, 702-710.

Montague, P. R., & Berns, G. S. (2002). Neural economics and the biological substrates of valuation. *Neuron, 36*, 265-284.

Nieuwenhuis, S., Heslenfeld, D. J., Alting von Geusau, N. J., 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.

Nyberg, L., & Cabeza, R. (2000). Brain imaging of memory. In E. Tulving & F. I. M. Craik (Eds.), *The oxford handbook of memory*. Oxford: Oxford University Press.

O'Doherty, J., Critchley, H., Deichmann, R., & Dolan, R. J. (2003). Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *J Neurosci, 23*(21), 7931-7939.

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*, 452-454.

O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nat Neurosci, 4*(1), 95-102.

Packard, M. G., & Knowlton, B. J. (2002). Learning and memory functions of the basal ganglia. *Annu Rev Neurosci, 25*, 563-593.

Packard, M. G., & McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol Learn Mem, 65*(1), 65-72.

Pasupathy, A., & Miller, E. K. (2005). Different time course of learning-related activity in the prefrontal cortex and striatum. *Nature, 433*, 873-876.

Poldrack, R. A., Clark, J., Pare-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., et al. (2001). Interactive memory systems in the human brain. *Nature, 414*(6863), 546-550.

Poldrack, R. A., Prabhakaran, V., Seger, C. A., & Gabrieli, J. D. E. (1999). Striatal activation during acquisition of a cognitive skill. *Neuropsychology, 13*(4), 564-574.

Poldrack, R. A., & Rodriguez, P. (2004). How do memory systems interact? Evidence from human classification learning. *Neurobiol Learn Mem, 82*(3), 324-332.

Ragozzino, M. E. (2003). Acetylcholine actions in the dorsomedial striatum support the flexible shifting of response patterns. *Neurobiology of Learning and Memory, 80*(3), 257-267.

Reber, P. J., Siwiec, R. M., Gitelman, D. R., Parrish, T. B., Mesulam, M.-M., & Paller, K. A. (2002). Neural correlates of successful encoding identified using functional magnetic resonance imaging. *J Neurosci, 22*(21), 9541-9548.

Reber, P. J., & Squire, L. R. (1999). Intact learning of artificial grammars and intact category learning by patients with Parkinson's disease. *Behav Neurosci, 113*(2), 235-242.

Robbins, T. W., & Everitt, B. J. (1996). Neurobehavioural mechanisms of reward and motivation. *Curr Opin Neurobiol, 6*, 228-236.

Rogers, R. D., Andrews, T. C., Grasby, P. M., Brooks, D. J., & Robbins, T. W. (2000). Contrasting cortical and subcortical activations produced by attentional-set shifting and reversal learning in humans. *J Cogn Neurosci, 12*, 142-162.

Rolls, E. T. (1999). The functions of the orbitofrontal cortex. *Neurocase, 5*, 301-312.

Schmidt, R. A., & Bjork, R. A. (1992). New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychol Sci, 3*, 207-217.

Schneider, W., Eschman, A., & Zuccolotto, A. (2002). *E-prime user's guide*. Pittsburgh: Psychological Software Tools Inc.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *J Neurophysiol, 80*, 1-27.

Schultz, W., Tremblay, L., & Hollerman, J. R. (1998). Reward prediction in primate basal ganglia and frontal cortex. *Neuropharmacology, 37*, 421-429.

Seger, C. A. (1994). Implicit learning. *Psychol Bull, 115*, 163-196.

Seger, C. A., & Cincotta, C. M. (2005). The roles of the caudate nucleus in human classification learning. *J Neurosci, 25*(11), 2941-2951.

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*(4), 851-859.

Smith, J. G., & McDowall, J. (2006). When artificial grammar acquisition in Parkinson's disease is impaired: The case of learning via trial-by-trial feedback. *Brain Res, 1067*(1), 216.

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiol Learn Mem, 82*(3), 171-177.

Stark, C. E. L., & Squire, L. R. (2001). When zero is not zero: The problem of ambiguous baseline conditions in fMRI. *Proc Natl Acad Sci U S A*, 12760-12766.

Talairach, J., & Tournoux, P. (1988). *Co-planar stereotaxic atlas of the human brain: An approach to medical cerebral imaging*. Stuttgart; New York: G. Thieme; Thieme Medical Publishers.

Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. New York: Macmillan.

Thorndike, E. L. (1927). The law of effect. *Am J Psychol, 39*, 212-222.

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. *J Cogn Neurosci, 18*(6), 1029-1043.

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

Ursu, S., & Carter, C. S. (2005). Outcome representations, counterfactual comparisons and the human orbitofrontal cortex: Implications for neuroimaging studies of decision-making. *Brain Res Cogn Brain Res, 23*(1), 51-60.

van Veen, V., Holroyd, C. B., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2004). Errors without conflict: Implications for performance monitoring theories of anterior cingulate cortex. *Brain Cogn, 56*(2), 267-276.

Wagner, A. D., Koutstaal, W., & Schacter, D. L. (1999). When encoding yields remembering: Insights from event-related neuroimaging. *Phil Trans R Soc Lond B, 354*, 1307-1324.

Wagner, A. D., Schacter, D. L., Rotte, M., Koutstaal, W., Maril, A., Dale, A. M., et al. (1998). Building memories: Remembering and forgetting of verbal experiences as predicted by brain activity. *Science, 21*, 1188-1191.

Wheeler, E. Z. (2006). *Examing theories of ventromedial prefrontal cortex function.* Unpublished Doctoral dissertation, University of Pittsburgh, Pittsburgh.

Wittmann, B. C., Schott, B. H., Guderian, S., Frey, J. U., Heinze, H.-J., & Duzel, E. (2005). Reward-related fMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron, 45*(3), 459-467.

Woods, R. P., Cherry, S. R., & Mazziotta, J. C. (1992). Rapid automated algorithm for aligning and reslicing PET images. *J Comput Assist Tomogr, 16*(4), 620-633.

Woods, R. P., Mazziotta, J. C., & Cherry, S. R. (1993). MRI-PET registration with automated algorithm. *J Comput Assist Tomogr, 17*(4), 536-546.

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nat Rev Neurosci, 7*, 464-476.

Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2004). Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *Eur J Neurosci, 19*, 181-189.

Zola, S. M., & Squire, L. R. (2000). The medial temporal lobe and the hippocampus. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford handbook of memory* (pp. 485-500). Oxford: Oxford University Press.