Applying electrophysiological methods to investigate the brain mechanisms involved in the processing of rewards, punishments, and performance feedback

By

Travis Edward Baker

B.A. Malaspina University-College, 2004

A Thesis Submitted in Partial Fulfillment of
the Requirements for the Degree of

Master of Science

In the Department of Psychology

University of Victoria

Applying electrophysiological methods to investigate the brain mechanisms involved in the processing of rewards, punishments, and performance feedback

By

Travis Edward Baker
B.A. Malaspina University-College, 2004

Supervisory Committee

Clay B. Holroyd, Ph. D. (Department of Psychology, University of Victoria)
Supervisor

Ronald W. Skelton, Ph. D. (Department of Psychology, University of Victoria)
Co-Supervisor or Departmental Member

Tony W. Robertson, Ph. D. (Department of Psychology, Malaspina University-College/University of Victoria
Departmental Member

Supervisory Committee

Clay B. Holroyd, Ph.D
Supervisor

Ronald W. Skelton, Ph.D
Co-Supervisor or Departmental Member

Tony W. Robertson, Ph.D
Departmental Member

ABSTRACT

To elucidate the neural generator of the feedback error-related negativity (fERN), which is an event-related brain potential (ERP) component elecited following negative feedback, this study utilized a converging method approach by applying electrophysiological methods to replicate a reward task modeled after a functional magnetic resonance imaging, single-cell recording, and primate studies that activated the anterior cingulate cortex (ACC) with negative feedback. This study involved two ERP experiments. The first involved a replication of the reward task and the second experiment involved a modification of the reward task that separated the valence and instructional aspects of the feedback stimulus. Experiment 1 demonstrated that methodological issues can arise when using ERP methods to replicate a paradigm used in other neuroimaging techniques. Experiment 2 was conducted to control for these methodological issues. For experiment 2, the reward condition following a switch movement elicited a large N2 component, which consequently overlapped with the fERN. I concluded that if participants are incorrectly expecting negative feedback after switch trials, the 'unexpected' positive feedback following the switch trails violated their expectation, thereby eliciting a large N2, and not a fERN. This hypothesis will be tested in several follow up experiments by modifying the paradigm in Experiment 2 so that all the possible combinations of condition mappings can be applied.

Table of Contents Page Number

I. Title Page i

II. Supervisor Committee ii

III. Abstract iii

IV. Table of Contents iv

V. List of Figures v

VI. Acknowledgments vi

VII Dedication vii

VI. Introduction 1
- a) The Reinforcement Leaning Theory of the ERN 4

VIII. Experiment 1 15
- a) *Methods*
  - Participants 16
  - Task 16
  - Electrophysiological Recordings 19
  - Data Analysis 19
- b) *Results* 21
- c) *Discussion* 22

IX. Experiment 2 24
- a) *Methods*
  - Participants 26
  - Task 27
  - Electrophysiological Recordings 28
  - Data Analysis 28
- b) *Results* 31
- c) *Discussion* 35

X. General Discussion 35

XI. Reference Material
- a) Bibliography 42
- b) Appendix A 50
- c) Appendix B 53
- d) Appendix C 56
- c) University of Victoria partial copyright license 58

List of figures Page Number

Figure 1. a) Typical example of ERPs associated with negative and positive feedback b) Typical example of ERPs associated with error and correct trials 3

Figure 3. The sequence of events that comprise one trial for experiment 1 18

Figure 4. Electrode montage used for electrophysiological recording. 20

Figure 5. (Top)Averaged ERP waveforms recorded at channel FCz for Stay-Reward, Switch-Reward, Switch-No Reward conditions. (Bottom) Averaged difference waves associated with channels FCz and PO4. 22

Figure 6. The sequence of events that comprise one trial for experiment 2. 29

Figure 7. Stimulus probablility for experiment 2. 29

Figure 8. a) Averaged ERP waveforms recorded at midline channels (Fz, Cz, Pz) for Stay-Reward, and Switch-Reward conditions. b) Scalp voltage maps for the Switch-Reward condition at 308 msec following feedback onset. c) Scalp voltage maps for the Stay-Reward condition at 308 msec following feedback onset. 32

Figure 9. Averaged ERP waveforms recorded at channel FCz for Switch-Reward, and Switch-No Reward conditions. 33

Figure 10. a) ERPs for the Switch-reward, switch no-reward, and associated difference wave at FCz.b) Scalp voltage maps for the difference wave at 232 msec following feedback onset.
c) Scalp voltage maps for the difference wave at 300 msec following feedback onset. 34

Figure 11. Proposed follow-up experiments. 38

Figure 13. Anterior cingulate cortex anatomy 54

Figure 12 a) A pyramidal cell during neurotransmission.
b) A stimulated region of cortex. 52

Acknowledgments

I like to start by expressing my deepest gratitude to my committee for all their well received input throughout my thesis work. I like to begin by acknowledging Tony Robertson, who introduced me to this field and provided me with invaluable knowledge throughout the years of my academic life. I like to express thanks to Ron Skelton for sharpening my empirical eyes, which has permitted me to observe experiments with scientific scrutiny. And to Clay Holroyd, who has allowed me to gently stand on his shoulders throughout the years and thereby allowing me to feel like a young scientist. Finally, I can't express in words alone of all the support my family has given me throughout the years to get me to this point. An accomplishment shared by all.

To Veronica and Elmer Edwards

Everybody makes mistakes, but we try to learn from our mistakes. Such learning depends crucially on the ability to discriminate between positive and negative outcomes following behaviour. How we learn in this respect has been of great interest to theorists who study reinforcement learning (RL). The subject of RL is based on Thorndike's Law of Effect, which holds that if an action is followed by a reward (positive feedback), then that action will likely be performed again, whereas if the action is followed by a punishment (negative feedback), then that action will not likely be performed again (for review see Cantania,1999). Here, the terms "reward" and "punishment" are used as proxies for any outcomes that are respectively consistent or inconsistent with an organism's task goals, such as positive and negative feedback stimuli. Recently, advances in the field of cognitive neuroscience have provided researchers with a window onto the neural and cognitive mechanisms that underlie this reinforcement learning principle. It is not surprising, therefore, that there is evidence from a variety of sources that indicates that the brain responds differentially to positive and negative feedback.

In particular, recent studies have utilized event-related potentials (ERPs) to examine how an evaluative system that processes positive and negative feedback is implemented in the brain (for review of ERPs, please see Appendix A). Specifically, in studies where human participants receive feedback in trial and error learning tasks, analysis of the ERPs following the feedback stimulus revealed that, following negative feedback, the potential became more negative. This ERP component became known as the feedback-ERN (fERN), which is characterized by a negative deflection at fronto-central recording sites that peaks approximately 250 ms following negative feedback presentation (see figure 1a; Holroyd & Coles, 2002; Holroyd et al., 2005; Miltner, Braun,

& Coles, 1997; Nieuwenhuis, et al., 2004). The fERN is commonly quantified by either identifying the maximal negative amplitude of the ERP waveform within a particular time window (e.g extending from 150-300 ms) following the presentation of the feedback stimulus or by subtracting the response to positive feedback from the response to negative feedback, otherwise known as a difference wave analysis. The evaluative system that produces the fERN appears to classify outcomes into binary categories - as events that either do, or do not, indicate that a task goal has been achieved. In fact, a recent study found that the fERN can reflect either utilitarian information (monetary gains vs. losses) or performance information (correct vs. incorrect choices), depending on which aspect of the feedback is made salient to the subjects (Nieuwenhuis et al., 2004).

Since a seminal report on the fERN by Miltner et al. in 1997, several empirical studies have elucidated the cognitive basis and functional significance of the fERN. In one experiment, Hajcak et al. (2006) examined whether or not the fERN is sensitive to the magnitude of feedback. In this study, two experiments were conducted where participants performed simple gambling tasks in which they could receive a range of potential outcomes on each trial. In the first experiment, participants could choose one of four doors on each trial, and received one of four possible outcomes that varied in magnitude and valence (i.e. gained 25¢, gained 5¢ or lost 5¢, lost25¢). In the second experiment, participants could chose between one of five doors on each trial, and received one of five possible outcomes indicating whether they gained 25¢, gained 5¢, broke even, lost 5¢, or lost 25¢. Unlike feedback indicating monetary gain, feedback indicating non-rewards was associated with a fERN in both experiments; however, the amplitude of the fERN did not demonstrate sensitivity to the amount of losses in either experiment. These data suggest that the fERN reflects the early appraisal of feedback based on a binary classification of good versus bad outcomes (see also Holroyd, Larsen, & Cohen, 2004a; Holroyd, Hajcak & Larsen, 2006).

When the source of the fERN was estimated using equivalent dipole analysis procedures, a generator in or near the anterior cingulate cortex (ACC) was suggested (Miltner et al., 1997). In this respect, the fERN closely resembled another event-related brain potential, the response error-related negativity (rERN). The response-ERN (rERN) is a negative deflection in the ERP that occurs approximately 100 ms following errors in speeded-response time tasks (Gehring, 1992; Falkenstein, Hohnsbein, Hoormann, & Blanke,1991; Holroyd & Coles, 2002; Holroyd, Yeung, Coles & Cohen, 2005). The rERN is not observed on correct trials (Figure. 1b). Importantly, the rERN and the fERN appear to be different manifestations of the activity of a single evaluative system when it determines that ongoing events are worse than anticipated (for more detail see Holroyd & Coles, 2002). This resemblance led to the proposal that the negativities following

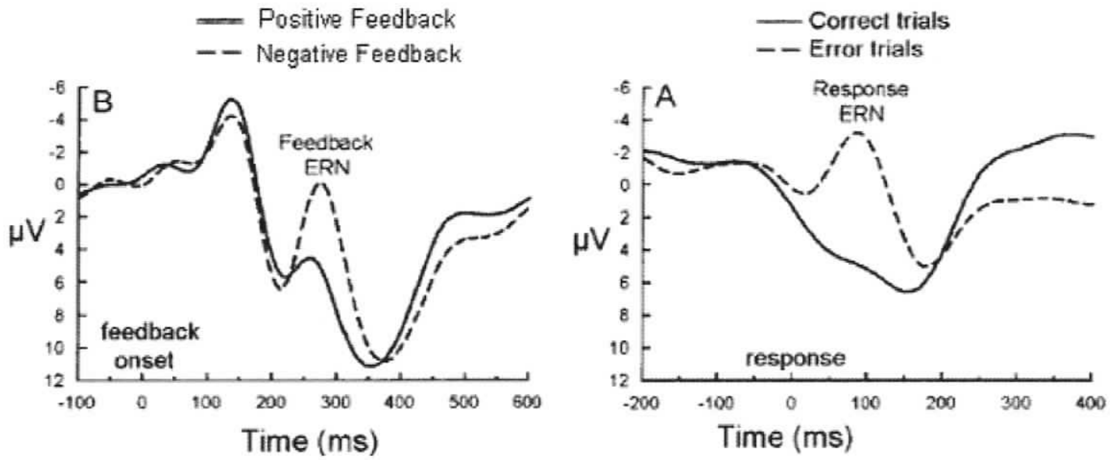

Figure 1. a) Typical example of ERPs associated with negative and positive feedback. fERN is the negative deflection occurring around 250 ms following feedback onset b) Typical example of ERPs associated with error and correct trials. Response-ERN is the negative deflection occurring around 100 ms following error trials (Reproduced from Holroyd et al., 2004b). Note that negative is plotted up by convention.

response errors (the rERN) and negative feedback (the fERN) were associated with the same neural and cognitive error-detection process.

The Reinforcement Learning Theory of the ERN

A recent theory has extended the hypothesis of Miltner et al. (1997) that the fERN is sensitive to an evaluation process that monitors the consequence of action, and that it is generated in the ACC. The Reinforcement Learning Theory of the ERN (RL-ERN theory), proposed by Holroyd and Coles (2002), suggests that the fERN is associated with the impact of a dopamine (DA) signal carried by the midbrain dopamine system (i.e. ventral tegmental area and substantia nigra pars compacta) onto the motor areas of the ACC for the adaptive modification of behaviour (Holroyd & Coles, 2002). The RL-ERN theory is based on the findings of Schultz and colleagues, who showed that DA neurons respond to changes in the prediction of the "goodness" of ongoing events (Schultz, 1998, Schultz et al., 1995). They demonstrated that a positive DA signal (a phasic increase in DA activity seen as a burst of action potentials) is elicited when an event is better than expected (for example, when a monkey receives an unexpected reward), and a negative DA signal (a phasic decrease of DA activity seen as a transient cessation from baseline firing rate) results when an event is worse than expected (for example, when a monkey does not receive an expected reward; Schultz et al., 1995; Schultz, 1998; Previc, 1999, Schultz et al., 2000, Tobler et al., 2005, Samejima et al., 2005). These DA signals are conveyed to various brain structures (i.e. basal ganglia, prefrontal cortex and the ACC) involved in motor control, motivation and goal-directed behavior and it is widely believed that these signals are used by these neural areas for the purpose of reinforcement learning (Schultz, 1998; Previc, 1999, Schultz et al., 2000, Tobler et al., 2005, Samejima et al., 2005). The RL-ERN theory holds specifically that ACC uses these dopamine signals to improve performance on the task at hand according to principles of RL (see Appendix B for anatomical description of ACC).

Furthermore, there is evidence that tonic DA activity exerts an inhibitory affect on pyramidal cells in prefrontal cortex by modulating synaptic input into the cells (Yang, & Seamans, 1996). According to the RL-ERN theory, it is proposed that the fERN is elicited when a negative DA signal carried by the DA system disinhibits the apical dendrites of Layer V pyramidal cells along the ventral bank of the cingulate sulcus of the ACC, which are continuously receiving excitatory input from the various neural structures that project to the ACC (Holroyd & Coles, 2002). Therefore, according to the theory, a reduction of dopaminergic input to the apical dendrites of Layer V pyramidal cells disinhibits the pyramidal cells there. This effectively serves to decrease the gain of the pyramidal cells, enabling a wider range of external input to stimulate the ACC neurons and allowing more excitatory postsynaptic potentials (EPSPs) to summate along the ventral bank of the cinguluate sulcus, which in turn is suggested to generate the fERN (for more detail, please see Holroyd & Coles, 2002). Conversely, it has also been suggested that a positive deflection in the ERP may occur following positive feedback when a positive DA signal carried by the DA system inhibits the apical dendrites of the Layer V pyramidal cells in the same region (Holroyd & Coles, 2002; Holroyd, 2004; Pakzad-Vaezi et al. 2006). According to this position, an increase in DA elicited by the receipt of positive feedback increases the gain and restricts the response of the pyramidal cells to only the strongest inputs, leading to a reduction in the amplitude of the fERN.

Additionally, the phasic increases and decreases in the activity of the DA system appear to induce, respectively , long-term potentiation and long-term depression at synapses linking the presynaptic excitatory inputs to the postsynaptic ACC pyramidal cells (Yang, & Seamans, 1996; Seamans, & Yang, 2004), consistent with the neural implementation of a reinforcement learning rule (Wickens et al., 2003).

Because the RL-ERN theory holds that the amplitude of the fERN is driven by DA, and because DA is sensitive to unexpected positive and negative outcomes, then it follows that the fERN should also be sensitive to unexpected positive and negative outcomes. Several experiments have confirmed this. For example, Holroyd et al. (2003) showed that the amplitude of fERN is influenced by the expected outcome of the feedback. In this study, participants were engaged in a pseudo-reinforcement learning task in which they were required to pick one of four balloons that were displayed on a computer screen. They were told that only one of the balloons contained money, and that they should earn as much money as possible throughout the experiment. Following each response, a feedback stimulus was presented, signifying that the balloon either contained 5 cents US (positive feedback) or was empty (negative feedback). Unknown to participants, the type of feedback was delivered at random (75%, 25%): in the reward condition, the participants were exposed mostly to positive feedback, and in a non-reward condition, the participants were exposed mostly to negative feedback. According to the RL-ERN theory, negative feedback in the non-reward condition would be associated with small fERNs (because presentation of negative feedback would confirm the expectation of the evaluating system), whereas the negative feedback in the reward condition would be associated with large fERNs (because presentation of negative feedback would violate the expectations of the evaluating system). The results were consistent with this theoretical prediction: the fERN was larger in the condition in which rewards were expected than in the condition in which rewards were unexpected. These results have recently been replicated and extended using a time estimation task (Holroyd & Krigolson, submitted).

A complicating factor in understanding the fERN is that it shares a similar latency and scalp distribution with another ERP component, the N2. The N2 is also commonly elicited by unexpected events or infrequent events, specifically in psychological tasks with a simple 'oddball' structure. In the oddball task, participants are presented with frequent and infrequent stimuli and are required to count the occurrences of the infrequent target stimulus (Luck, 2005). Following the occurrence of the target stimulus, the N2 is elicited (Ruggs & Coles, 1995). This component is a large negative deflection with a latency of between 200-260 ms that is distributed across the front of the scalp (see figure 2). The hallmark of the N2 component is that its amplitude increases as the eliciting stimulus occurs less frequently, and the N2 has been linked to the cognitive processes of stimulus identification and distinction (Patel & Assam, 2005). In addition, the P3 is an ERP component that is elicited immediately following the N2 (300 ms or later), and that is distributed over posterior areas of the scalp (see figure 2). Like the N2, it is also sensitive to stimulus probability (Dochin & Coles, 1988, Ruggs & Coles, 1995). For example, several investigators have noted that the P3 amplitude gets larger as

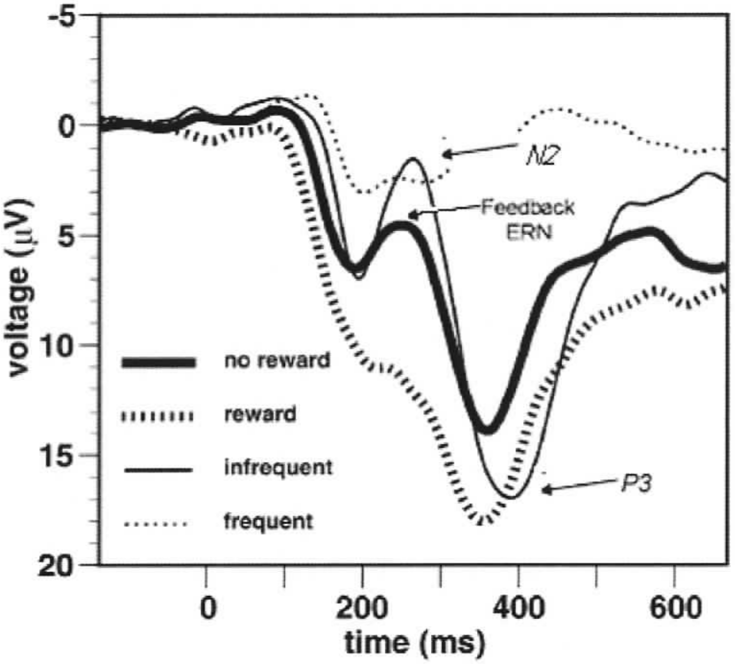

Figure 2. ERPs elicited by imperative stimuli in an oddball task and by feedback stimuli in a reinforcement learning task. "Infrequent, frequent": ERPs elicited by the target (10% of trials) and non-target (90% of trials) stimuli, respectively, in an oddball task. "reward, no reward": ERPs elicited by feedback stimuli indicating rewarded trials and non-rewarded trials, respectively, in a reinforcement learning task. Stimulus onset occurs at 0 ms. The oddball N2 on infrequent trials and the feedback ERN on no reward trials can be seen peaking at about 250 ms. The data from both tasks were recorded in the same group of subjects at channel FCz (reproduced from Holroyd, 2004).

stimulus probability gets smaller; the rarer the event, the larger the P3 (Dien, Spencer, & Donchin, 1999, Nieuwenhuis et al., 2004; and Nieuwenhuis, Aston-Jones, & Cohen, 2005). The N2 and P3 have been referred to as the N2/P3 complex because they tend to occur together, and because they are associated with cognitive processes involved in directing attention to infrequent events.

The N2 has the same temporal and spatial distribution of the fERN and is elicited by unexpected or infrequent stimuli, but the relationship between the fERN and N2 remains an unresolved issue (Holroyd, 2004). Because of the similarities in their scalp distribution and latency, the question arises as to whether or not the two components are in fact produced by the same underlying cognitive process, rather than to distinct cognitive processes as is commonly believed (Holroyd, 2004). On the one hand, the two components have been functionally dissociated, suggesting that they indeed reflect different phenomena. Specifically, the amplitude of the fERN is larger for unfavorable events than for favorable events, even when the frequencies of both events are the same (Miltner et al., 1997; Holroyd & Coles, 2002; Holroyd), but the N2 is larger for infrequent events than for frequent events, in conditions (such as the oddball paradigm) in which rewards and punishments play no obvious roll (Towey et al., 1980). These findings suggest the existence of two systems, one involved in reward or performance processing and a second sensitive to stimulus frequency. Despite their functional dissociation, and the remarkable similarity between the morphologies of the N2 and the fERN, it is has been suggested that the fERN and N2 may reflect the operation of a common underlying mechanism that is sensitive to both stimulus probability and valence (Holroyd, 2004). Recently, to account for this issue Holroyd and colleagues have proposed a modification of the RL-ERN theory which holds that 1) unexpected events elicit a large negativity, which is seen as the N2 in oddball tasks, 2) if the unexpected event is an negative outcome, then the amplitude of this negativity may be increased, however the degree of increase appears to be small; 3) if the unexpected event is an positive outcome, then the amplitude of this negativity is reduced, such that infrequent positive feedback stimuli do not produce a large negativity. At the neural level, infrequent events elicit ACC activity by excitatory input coming from other cortical areas; but infrequent positive outcomes also elicit an increase in DA that inhibits this input and reduces that excitation.

To investigate this issue, Pakzad-Vaezi et al. (2006) conducted an experiment in which they compared ERPs elicited by infrequent positive and negative feedback in a time estimation task under 3 probability conditions (Control 50/50; Easy-25% error/75% correct; and Hard-75%error/25%correct), with the ERP elicited by infrequent events in an oddball task under the same probability conditions. For the time estimation task, participants pressed a mouse button when they believed that 1 sec had elapsed following presentation of an auditory cue. Immediately following, they were presented with feedback indicting whether or not they got it right or wrong. The ERPs associated with the feedback where then compared to the ERPs associated with frequent and infrequent stimuli in an oddball task which the same participants completed either before or after the time estimation task. The results revealed that all of the infrequent events (e.g., infrequent oddball, infrequent negative feedback, and infrequent positive feedback) elicit an N2. In addition, the authors found that positive feedback elicited a positivity with a frontal-central scalp distribution (the "correct-related positivity"), the amplitude of which increased in proportion to the unexpectedness of the feedback, as indicated by a difference wave analysis. These results suggest that the N2 is generated by unexpected events on every trial, but that unexpected positive events produce a frontal-central positivity that cancels out the N2. Thus, the negativity elicited by negative feedback maybe simply a common N2, and the reduction of N2 amplitude associated with positive feedback results from the superposition of a frontal-central positivity.

An important assumption underlying the RL-ERN theory is that the fERN is generated in the ACC. The results of several experiments converge in suggesting that the ACC may indeed be the source of the fERN (e.g. Miltner et al., 1997; Holroyd et al., 2004b). However, other source localization studies have provided conflicting evidence (e.g. Van Veen et al. 2004, Nieuwenhuis, et al., 2005). Thus it remains uncertain whether the fERN is in fact generated in the ACC.

This controversy over the source of the fERN results in part from the inverse problem, which is that any ERP scalp distribution can be accounted for by an infinite number of possible intracranial source configurations. For this reason it is impossible to determine the exact location and orientation of the neural sources responsible for producing an observed voltage distribution on the basis of voltage distribution alone. Nevertheless, it is possible to constrain the set of possible source configurations using data collected with other experimental methodologies. In this way, the strength of one experimental method can compensate for the weakness of another method (and vice-versa). This is known as the "converging method" approach". As explained by Luck:

> "…a commonly cited principle of scientific inference is the idea of converging evidence, which was first developed in the context of perception research but is now widely used in cognitive neuroscience. The basic idea is that many interesting questions about the mind cannot be answered by means of any single method, but a clear answer can be obtained when many methods with different strengths and weaknesses converge on the same conclusion. This is a common use of source localization models." (Luck, 2005, pp. 296)

In the case of the fERN, the solution to the source localization algorithm can be constrained by empirical results associated with other experimental methods. Specifically, to provide converging evidence that the ACC is the source of the fERN, it must be demonstrated that: 1) the ACC is differentially activated by positive and negative feedback, as indicated by the results of an experiment using the other method; 2) the feedback stimuli in the same task used with the other experimental method elicit the fERN; and 3) the source of the fERN in this task can be localized to the area of the ACC indicated by the other experimental method. Together, these results would show that positive and negative feedback stimuli both differentially activate the ACC and differentially modulate the amplitude of the fERN, and further that the fERN can be accounted for by a source located in the ACC. Such mutually consistent evidence would strongly suggest that the fERN is in fact produced in the ACC, as assumed by the RL-ERN theory.

In the present study, I draw on the results of three previous experiments - single-cell recording in monkeys, event-related fMRI in humans, and single-cell recording in humans – that have systematically applied a specific reward based decision-making paradigm to demonstrate that ACC is involved in utilizing rewards and punishments for modifying behaviour (point #1, above; Bush et al, 2002; Williams et al, 2004; Shima & Tanji, 1998). Following the converging method approach, I intend to show that the feedback stimuli in this task also elicit the fERN (point #2, above), and further, that the fERN in this task can be accounted for with a source localized to the ACC (point #3, above). Note that although these points have been addressed in previous studies, to my knowledge no previous study has successfully localized the fERN to the ACC using the same experimental task across the different methodologies. If successful, the use of a single task across experiments would strongly suggest that fERN is in fact elicited in the ACC.

The experimental paradigm I adopt here originated when Shima and Tanji (1998) used a reward based decision-making task in a primate single-cell study to investigate reward sensitive neurons in the ACC (cingulate motor areas as described by authors). In this study, activity in ACC cells was recorded while the monkeys either pushed or turned a handle to receive feedback (apple juice). All trials began in response to a visual trigger signal after a 2-7 second waiting period. In the majority of the trials (Stay-Reward), the monkeys received positive feedback (apple juice) and were trained to stay with the same movement on the following trial. The amount of reward remained constant for twelve trials until the reward was reduced (less apple juice). For the reduced reward trials (Switch-no reward), the monkeys were trained to switch to the other movement on the following trial and the reward reverted to the full amount and a new series of Stay-reward trials began. To control for the possibility that ACC activity may be related nonspecifically to the visual trigger signal, rather than specifically to the reward information, the monkeys were trained to alternate their movement to receive positive feedback when they heard an auditory tone during the waiting period (Switch-Reward).

The researchers found that different cell populations in ACC responded to both motor response and reward processing. Specifically, four types of cells (81 of 221 cells recorded) in the ACC exhibited changes in activity during the time between the onset of the reduced reward and the initiation of a new selected movement. The same cells were not activated in both the Stay-reward and/or Switch-reward trials. From these observations, the authors suggested that these specific cells within the ACC use feedback-related information for behavioural modification. To further support these findings, the authors demonstrated that by injecting the forelimb part of the ACC with muscimol, which deactivated this area, the monkey failed to select a correct movement when the reward was reduced, but had no problems when the alternation was cued with the auditory tone. In addition, these effects were not observed when muscimol was injected into the hind limb area of the ACC. In sum, these findings indicate that specific cells within the ACC respond vigorously to movement selection based on reduction in reward or negative feedback. Thus, the ACC appears to use error information for behavioural adaptation.

Following the success of this paradigm in demonstrating that cells within the ACC employ feedback information to drive behaviour, Bush and colleagues (2002) utilized this task in a human event related fMRI study. The study was modeled after Shima and Tanji's paradigm with some appropriate adaptations for human subjects. Participants were given a button box with two buttons and instructed to use their right index and middle fingers to press buttons 1 and 2. Immediately following their button press, they were visually presented with a stimulus that provided both movement and feedback information. On 80% of the trials (Stay-reward), they were presented with positive feedback ($$$$$) indicating that they received 15 cents for that button press and to continue with the same button press on the following trials. On 10% of the trials (Switch-no reward), the feedback stimulus ($$$) indicated that they would receive 9 cents and that they should press the alternate button on the following trial. On the remaining 10% of the trials (Switch-reward), the feedback stimulus indicated that they would receive the full reward and they should press the alternate button on the following trial.

As predicted, Bush et al. (2002) found significant activation in the ACC after the onset of Switch-no reward trials. More importantly, activity during Switch-no reward trials was higher than during Switch-reward trials, indicating that the greatest ACC activity was elicited by negative feedback associated with a change in behaviour in the context of reduced reward, not by a change of behaviour in and of itself. These data provide additional support for the role of the ACC in reward processing.

In a third study, Williams et al. (2004), recorded from single neurons in the ACC of human participants who were scheduled to undergo surgical cingulotomy to treat major depression, obsessive-compulsive disorder and/or bipolar-affective disorder. The study was modeled after the task design used in the Bush et al. (2002) fMRI study. As predicted, neuronal activity in the ACC increased when participants altered movement direction, and the greatest activity occurred when this change in direction followed negative feedback. In addition, after ACC ablation, participant's error rates increased significantly for the Switch-no reward trials compared to the Switch-reward trials. This suggests that the ACC plays an important role in signaling alternative actions, especially in the context of negative feedback. These findings agree with the previous experiments, suggesting that the ACC in humans plays an important role in reward processing.

Experiment 1

Together, these previous studies have strongly implicated the ACC in feedback processing. Further, as described above, the fERN is thought to be produced in the ACC, although this has yet to be demonstrated. To provide evidence in support of this position, I recorded ERPs from participants engaged in the reward-based decision-making task used in these three studies (i.e., human fMRI, human single-cell recording, and non-human primate; Bush et al., 2002, Shima & Tanji, 1998, Williams et al., 2004). I predicted that the fERN would be associated with the same feedback conditions that modulated ACC activity in the previous studies. Taken together, this result would indicate that the fERN is elicited under conditions that are known to activate the ACC, and suggest that the fERN is in fact generated in the ACC.

Methods

Participants

Twenty undergraduate students (12 male and 8 female, aged 18-25) recruited from the University of Victoria were included in experiment 1. All participants had normal or corrected-to-normal vision and none had a history of head injury. Each received course credit as well as a monetary bonus associated with the experimental task. The amount of money depended on the probability of the reward, as described below. All participants gave informed consent and the study was approved by the local research ethics committee and was conducted in accordance with the ethical standards prescribed in the 1964 Declaration of Helsinki.

Task

Participants were seated comfortably in an electromagnetically shielded, soundproof booth and used a standard SRX Button Box. The stimuli were viewed from a distance of about 70 cm on a 17-inch, 1024×768 computer monitor, and was controlled using E-Prime experiment control software (Psychological Software Tools, Pittsburgh, PA). Subtended viewing angles were 3.3° horizontally and 5° vertically. Participants were asked to rest their forearms on the flat desktop and position their hand and forearm so that the fingertips of the index fingers would rest comfortably on the button box placed in front of them. Participants received both written and verbal instructions, which explained the procedure and stressed the importance of correct posture, and minimizing head movement and eye blinks. They were instructed to use their left index finger to press button 1 and their right index finger to press button 2.

The task used in this experiment was a basic replication of the reward based decision-making task previously used in the experiments discussed (Bush et al., 2002; Williams et al., 2004; Shima & Tanji, 1998). Trials began with the display of the fixation cross '+' and participants were to make their button choice following its appearance. On the first trial, they were asked to guess a response. Immediately afterwards, one of the feedback stimuli was presented at the centre of the screen and remained visible for 500 ms. The feedback stimuli were displayed at random and were as follows; on 80% of the trials, *Stay-reward condition*, participants were presented with a feedback stimulus ($$$) indicating that they received a large reward (3 cents) and that they should press the same button on the following trial; on 10% of the trials, *Switch-reward condition*, participants were presented with a feedback stimulus (<$$$>) indicating that they received a large reward (3 cents) but should press the other button on the following trial; and on 10% of the trials, *Switch-no reward condition*, the participants were presented with a feedback stimulus (<$>) indicating that they received a small reward and should press the other button on the following trial. The feedback stimulus was then followed by a blank screen delay for 1000 ms. The sequence of events that comprise one trial is depicted in Figure 3.

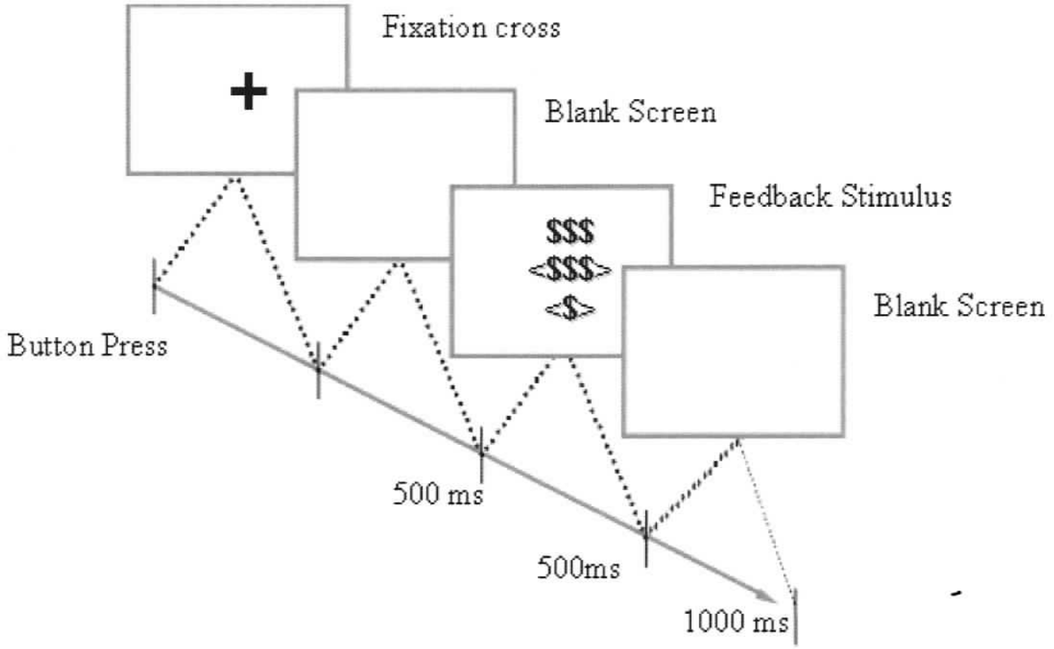

Figure 3. The sequence of events that comprise one trial.

In keeping with the Bush et al. (2001) protocol, participants were told that there is no absolute right answer (i.e., 1 or 2) for any trial, but that they should continue with the same strategy when given the stay cue and change to the other button when given the switch cue. In the occurrence of an error, participants were presented with a Switch-No reward cue. The task consisted of a practice block followed by 4 blocks of 100 trials each, with self-timed breaks between blocks. Following each block, the participant's total earnings were displayed on the computer screen along with a written reminder of the meaning of the instruction and feedback cues. At the end of the experiment, the participant's performance bonus was awarded and the participants were asked to fill out a short questionnaire for the purpose of assessing motivation, reward expectancies, and comprehension of the instructions (See Appendix A). Following the completion of the questionnaire participants were then fully debriefed about the overall purpose of the experiment.

Electrophysiological Recordings

The electroencephalogram (EEG) was recorded using 63 electrode montage in accordance with the extended international 10-20 system (Jasper 1958), as shown in Figure 4. Signals were acquired using sintered Ag/AgCl ring electrodes mounted in a nylon electrode cap (Falk Minow Services, Herrsching) and amplified by low-noise electrode differential amplifier with a frequency response of DC 0.017 – 67.5 Hz (90dB octave roll off). The amplified signals were digitized at a rate of 250 samples per second. Digitized signals were recorded to disk using Brain Vision Recorder software (Brain Products GmbH, Munich). Inter-electrode impedances were maintained below 15 KΩ. Two electrodes were also placed on the left and right mastoid (M1 and M2). During recording, all activity was referenced to an overall average. The electroculogram (EOG) was recorded for the purpose of artifact correction; horizontal EOG was recorded from the external canthi of both eyes, and vertical EOG was recorded from the sub orbit of the right eye and electrode channel Fp2.

Data Analysis

Post processing and data visualization were performed using Brain Vision Analyzer software (Brain Products GmbH, Munich). The digitized signals were filtered using a 4-th order digital Butterworth filter with a pass band of .10 – 20 Hz. For all experimental stimuli, a 1200 ms epoch of data extending from 200 ms prior to 1000 ms following stimulus onset (fixation cross, instruction cue, and feedback cue) was extracted from the continuous data file for analysis. Ocular artifacts were corrected using an eye movement algorithm described by Gratton et al. (1983). The EEG data were re-referenced to linked mastoids electrodes, M1 and M2. The data were baseline corrected

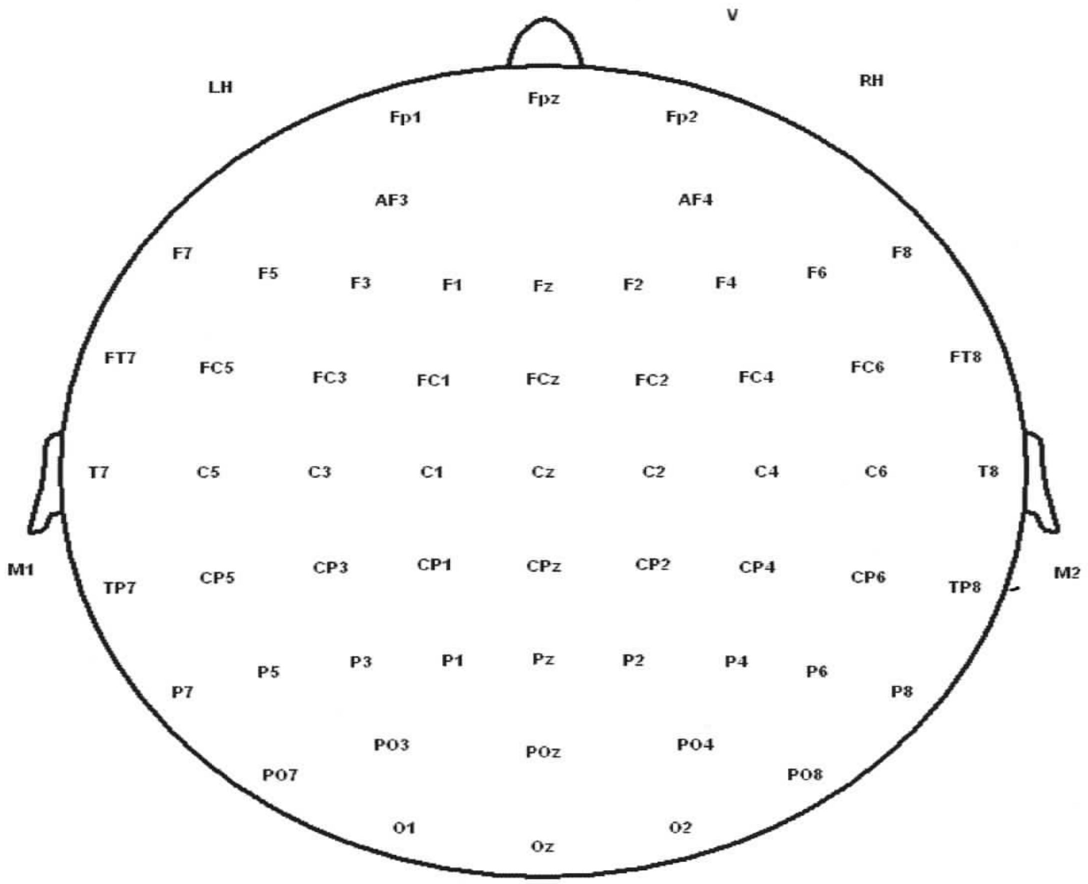

Figure 4. Electrode Montage (63 in total)

by subtracting from each sample the mean voltage associated with that electrode during the 200 ms interval preceding stimulus onset. Muscular and other artifacts were removed using a ±100 μV level threshold and a ±50 μV step threshold as rejection criteria. The EEG data were then re-segmented by condition. For the feedback cues (Stay-Reward, Switch-Reward, and Switch-No Reward), epochs of 800 ms were segmented, extending from 200 ms prior to 600 ms following the onset of the feedback cue. ERPs were then created for each electrode and participant by averaging the single-trial EEG according to feedback type.

fERN Analysis. The fERN was measured at channel FCz, where it reaches maximum amplitude (Holroyd et al., 2004b; Nieuwenhuis et al. 2004; Miltner et al., 1997). fERN amplitude was quantified for each electrode and participant by measuring the peak amplitude of the difference wave constructed by subtracting the Switch-Reward ERPs (10% probability) from the corresponding Switch-No reward ERPs (10% probability). The peak amplitude of this difference wave was obtained by detecting its maximum deflection within a 600 ms window following the onset of the feedback stimulus. The peak positive/negative amplitude of this maximum deflection was statistically tested against zero using a one sample t-test and an alpha of 0.05.

The difference-wave approach is used to separate ERP components of interest from other overlapping ERP components (Luck, 2005). Here, I used the difference-wave to isolate the fERN from other ERP components that may overlap with it, such as the N2 and P3. Otherwise, it would be difficult to determine if any differences observed between the reward and no reward trials were due to a difference in the amplitude of the fERN or due to differences in some other ERP component. The difference-wave approach provides a relatively pure measure of the brain's differential activity to positive feedback vs. negative feedback.

Results

Figure 5 shows the ERPs elicited by the Stay-Reward, Switch-Reward, and Switch-No Reward conditions recorded at channel FCz. The fERN analysis on the difference wave at the electrode site FCz revealed that the maximal negative deflection was significantly different than 0; t(19) = -11.33, p<0.001, (*M* = -5.17 uV, *SD* = 2.04), but had a peak latency inconsistent with previous observations (*M* = 330 ms, *SD* = 130 ms). In addition, I found that the largest amplitude deflection was maximal over most posterior areas of the scalp, specifically PO4, t(19) = -12.53, p<0.001, (*M* = -5.65 uV, *SD*

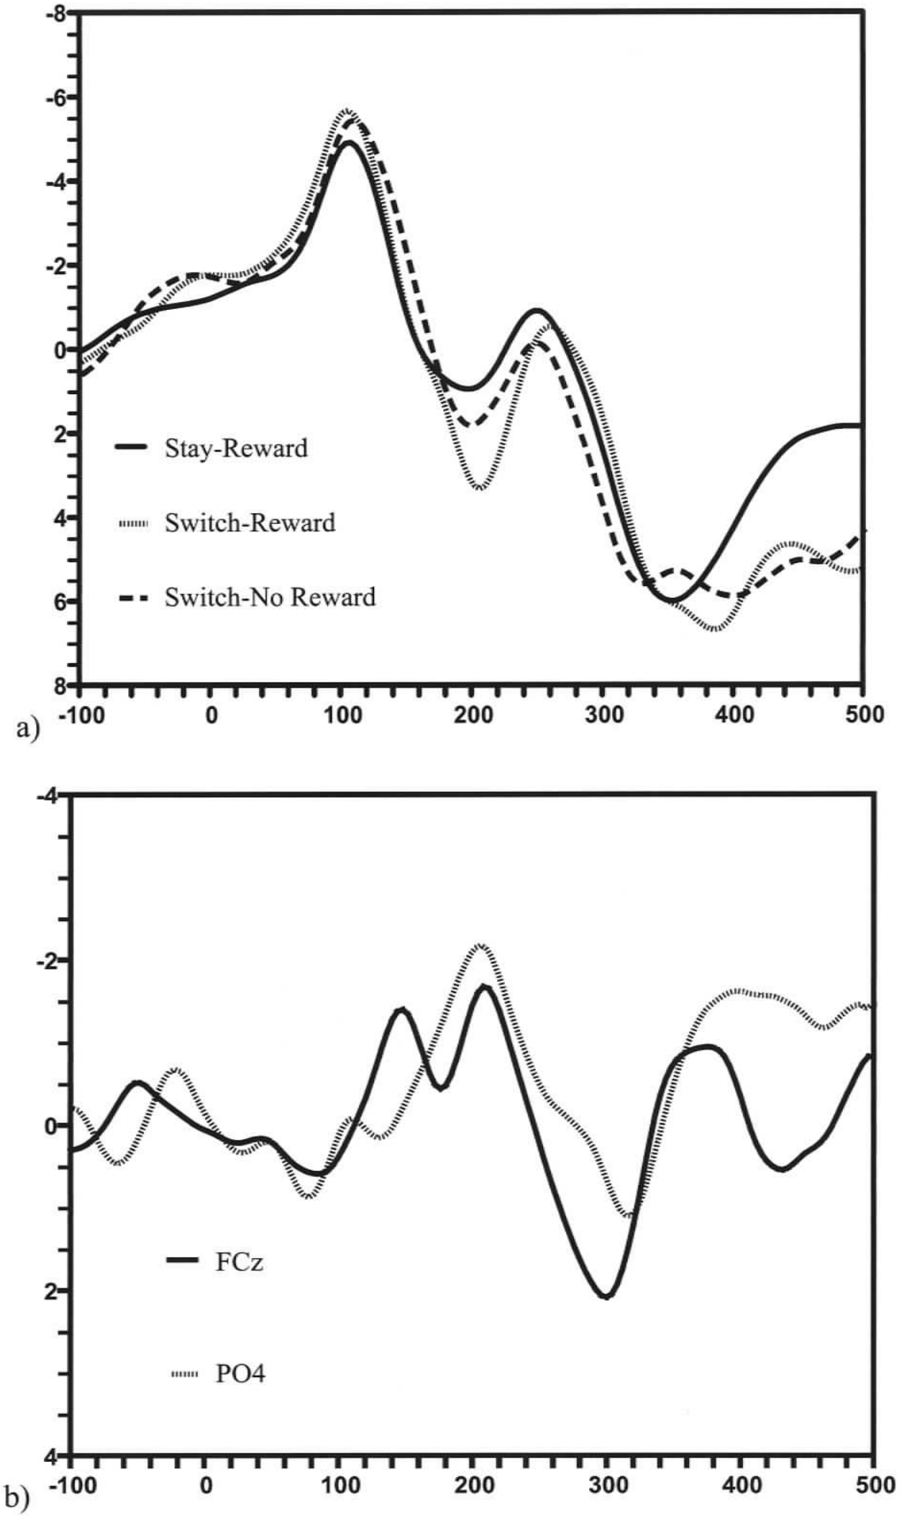

Figure 5. a) Averaged ERP waveforms recorded at channel FCz for Stay-Reward, Switch-Reward, Switch-No Reward conditions. b) Averaged difference waves associated with channels FCz and PO4

= 2.01), with a peak latency also outside the temporal range of the fERN (*M* = 366 ms, *SD* = 110). The scalp distribution and timing of this result is inconsistent with previous observations of the fERN.

Due to the high variability of the fERN peak latency across participants, I decided to take a further step in our analysis by constraining the temporal window of the maximum peak analysis to 200 to 320 ms. This analysis demonstrated that the maximal negative deflection was greatest at FCz and was significantly different than 0; t(19) = -8.51, p<0.001, (*M* = -3.60 uV, *SD* = 1.89), with a peak latency consistent with previous obseravations (*M* = 249 ms, *SD* = 44 ms). However, I also found that the posterior channels, specifically PO4, also displayed comparable amplitudes to FCz in the same time range (*M* = 245 ms, *SD* = 44), and this was significant when tested against 0; t(19) = -12.53, p<0.001, (*M* = -3.57 uV, *SD* = 2.03).

Discussion

The results demonstrate morphological differences between the ERPs elicited by Switch-Reward feedback and the ERPs elicited by Switch-No Reward feedback. However, the topography and latency of this difference is inconsistent with the predictions of the RL-ERN theory. It is important to note that this task involves an important methodological difference from the tasks used to study the fERN. Tasks commonly used to elicit the fERN utilize feedback stimuli that only convey information concerning the outcome of a behavioural response, typically a symbol (+, O) indicating monetary rewards and punishments. Such feedback information is said to be "scalar" in that it provides information along only one dimension (good vs. bad), as opposed to "vector" feedback that provides information along multiple dimensions (e.g., good vs.

bad and what the correct response should be). The dopamine system is believed to carry only scalar signals, not vector signals (Montague, Hyman, & Cohen, 2004), so it is unlikely that feedback with complex information can impact the fERN.

Consistent with this position are the results from a study by Mars et al. (2004), which suggest that if the information conveyed by the feedback stimulus is complex, then this may draw attention from the simple good-or-bad characteristic of the feedback. That study revealed that as the amount of information increased in the feedback, the amplitude of the fERN was reduced. In this study, the feedback provided vector information indicating not only how much reward the subject received, but also which direction to move on the following trial. Thus, it is possible that the complexity of the feedback information reduced the amplitude of the fERN. The main objective of Experiment 2 was to investigate this issue directly.

Experiment 2

I examined the issue of whether the complexity of the feedback information reduced the amplitude of the fERN by conducting a follow-up ERP experiment to experiment 1, using a modified version of the reward based decision-making. The main modification was to separate the two characteristics of the stimulus (instruction and feedback) from each other. To ensure that this new version of the task was consistent with the stimulus probabilities of the previous experiments, the instruction cue was displayed before the feedback stimulus (please see methods section for further description). Presumably, the impact of the valence dimension of the feedback is increased when the feedback only contains information related to reward. Furthermore, on the basis of the previous studies (which found ACC activity associated with negative feedback) and in conjunction with the RL-ERN theory, I predicted that a fERN would be elicited by negative feedback following the presentation of the switch cue, but not by positive feedback following the presentation of the switch cue. Taken together with the findings of Bush et al., (2002), Williams et al., (2004), and Shima & Tanji (1998), positive results would provide supporting evidence that; 1) the absence of the fERN in experiment 1 was due to the complexity of the feedback stimuli, 2) the fERN is elicited under conditions that are known to activate the ACC, and are consistent with the RL-ERN theory, and 3) suggest that the fERN is in fact generated in the ACC.

It is important to note that this modified version of the reward-based-decision making task raises a few methodological concerns. The first issue is whether or not the participants would disassociate the two reward conditions. These two conditions are both associated with reward, but follow different instruction cues. If subject were to combine these two outcomes into a single category (i.e., reward), then the subjective probability of a reward would be 90%, rather than 10% on switch-reward trials and 80% on stay trials. This result would make it difficult to disassociate the fERN from other overlapping ERP components such as the N2. Although the N2 has the same temporal and spatial distribution of the fERN and is elicited by infrequent stimuli, it is commonly believed to be generated by a different cognitive process. Nevertheless, the relationship between the fERN and N2 is an unresolved issue (Holroyd, 2004). Because feedback stimuli provide both frequency and valence information, it can be expected to produce both an N2 and an ERN, therefore complicating the interpretation of the results. To address this issue, the approach taken was to examine whether the ERP waveforms elicited by positive feedback in both reward conditions were different. It was inferred that if these two ERP waveforms are different in any respect, then the participants are differentially processing the two reward outcomes. This issue was addressed by examining the P3. Since the characteristics of the P3 can be manipulated by the probability of the stimulus presentation, it would be used to confirm that subjects are in fact categorizing the Stay-reward and Switch-reward outcomes separately, rather than as a single category. I expected that the (infrequent) Switch-reward ERPs would yield significant differences from the corresponding (frequent) Stay-reward ERPs. Specifically, the amplitude of the P3 would be larger for the infrequent Switch-reward condition and would be maximal over posterior channels, specifically Pz. If so, this would provide evidence that the participants were processing the two reward conditions differently and rule out the possibility that participants were combining both reward outcomes into a single subjective category.

A second methodological issue that needs to be addressed is whether the participants were ignoring the feedback altogether and instead focusing their attention only on the instruction cue. If attention is drawn away from the feedback stimulus, then the amplitude of the N2 would be reduced or absent. To address this issue, I compared N2 amplitude across the ERPs associated with Switch-reward and Switch-no reward (see methods for more detail). It can be inferred that if the amplitude of the N2 differs across the two conditions, then the participants were paying attention to the feedback stimuli.

Methods

Participants

Twelve undergraduate students, (7 male and 5 female, aged 18-26), were recruited from the University of Victoria via SONA (UVIC Psychology Research Participation

System). Although participation for this study is strictly voluntary, each participant received extra-credit in their first year psychology course as well as a monetary bonus associated with their performance of the experimental task. Participants that reported a history of a brain injury were excluded from the study.

All participants were asked to provide informed consent as approved by the Office of the Vice-President of Human Research, University of Victoria. This study was conducted in accordance with the ethical standards prescribed in the 1964 Declaration of Helsinki.

Task

In line with experiment 1, a modified version of the reward based decision-making task was used for experiment 2. The sequence of events that comprised one trial is depicted in Figure 6. Each trial began with the display of a fixation cross in the form of a '+' sign, which remained at the centre of the screen until the participant made their button choice. On the first trial they were asked to guess a response. The button choice was followed by a blank screen delay for 500 ms. Immediately afterwards, one of two instruction cues, presented in yellow, was displayed in the centre of the screen for a duration of 500 ms. The instruction cues consisted of a Stay and Switch Cue. For the Stay cue (| |), which appeared in 80% of the trials, participants were instructed to stay with the same button choice on the subsequent trial. For the Switch Cue (< >), which appeared in 20% of the trials, participants were instructed to switch their button choice on the subsequent trial.

Following the presentation of the instruction cue and after a short delay (500 ms), one of two feedback cues, presented in green, was displayed in the centre of the screen as follows: 1) Reward Cue '\$', indicating that participants would receive 3 cents for that trial; 2) No Reward Cue '0', indicating that participants would receive no reward for that trial. To control for probability effects, as well as keeping in line with the stimulus probabilities set out in the previous research designs, the feedback probabilities were set accordingly. Following the stay cue, only the reward cue was displayed (Stay-Reward), constituting 80% of the total presentation of the feedback cue (i.e. 100% of stay cues were followed by reward cues). Following the switch cue, one of two feedback cues was displayed: 1) The reward cue was displayed on 50% of those trials, constituting 10% of the total presentation of the feedback cue (Switch-Reward). 2) The no reward cue was displayed on the remaining trials, constituting 10% of the total presentation of the feedback cue (Switch-No Reward). See Figure 7. The feedback cue was then followed by a blank screen delay for 1000 ms.

Electrophysiological recordings

Electrophysiological recordings methods were identical to Experiment 1.

Data Analysis

fERN Analysis. The method of fERN analysis was identical to Experiment 1 except the peak amplitude of the difference waves was obtained by detecting its maximum deflection within a temporal window extending from 200 to 320 ms following the onset of the feedback stimulus, a typical window size for studying the fERN. This difference wave approach provides a relatively pure measure of the brain's sensitivity to positive and negative feedback.

P3 Analysis. To resolve the methodological issue of whether the participants were processing the two reward conditions differently, a peak amplitude measure was

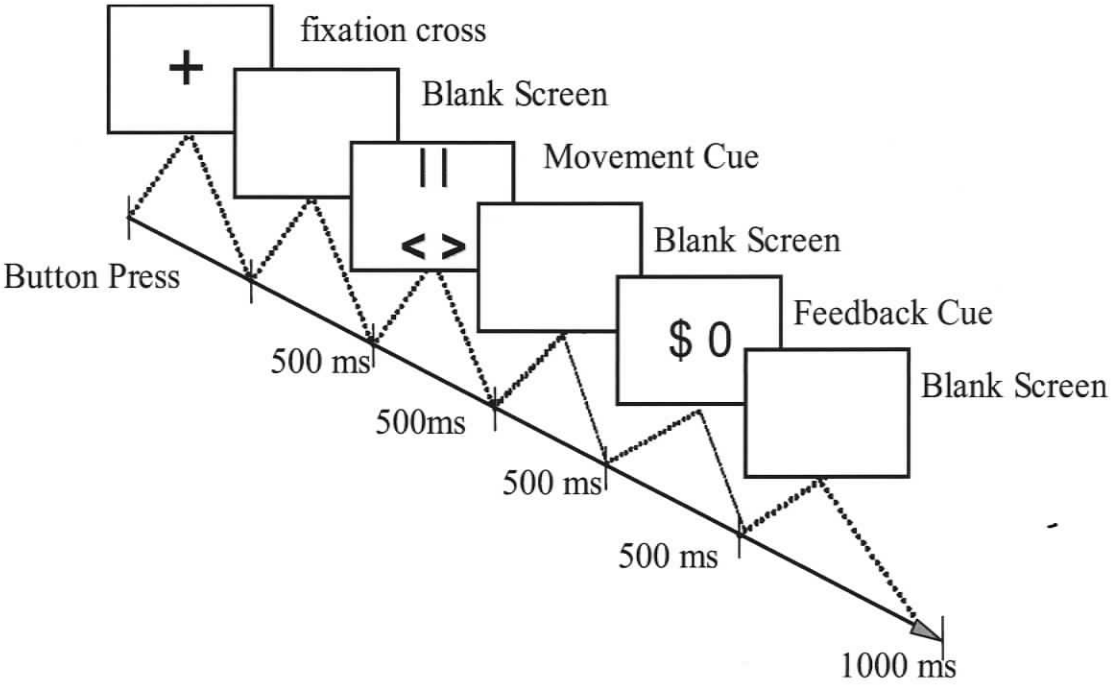

Figure 6. The sequence of events that comprise one trial.

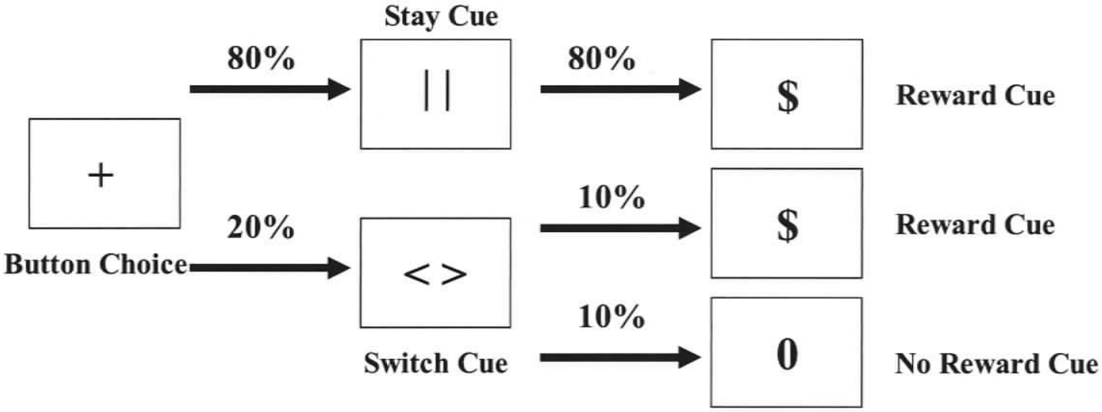

Figure 7. Stimulus Probablility.

conducted on the P3 for the reward conditions. Scalp sites over the parietal, central, and frontal locations (Pz, Cz, Fz), were selected to assess regional differences in P3 amplitude (Dien, Spencer, & Donchin, 1999; Nieuwenhuis, Aston-Jones, & Cohen, 2005). Peak amplitudes were measured by identifying the maximal positive amplitude of the ERP within a window extending from 300 to 600 ms following the presentation of the feedback stimulus. The latency of this peak amplitude was defined as the peak latency measure. This algorithm was applied to the Stay-Reward and Switch-Reward ERPs associated with the 3 midline electrode sites (Pz, Cz, Fz), and the results were submitted to an analysis of variance (ANOVA) with repeated measures and to paired t-tests. The Greenhouse-Geisser correction for repeated measures was applied where appropriate to correct for non-sphericity.

N2 Analysis. To resolve the issue of whether the participants ignored the feedback stimuli, a peak amplitude measure was conducted on the N2 for the Switch-Reward and Switch-No Reward conditions. Scalp sites over the FCz were selected to assess local differences in N2 amplitude. Particularly, the negativity associated with the N200 was measured by identifying the maximal negative amplitude of the ERP within a window extending from 150-300 ms following the presentation of the feedback stimulus. The latency of this peak amplitude was selected as the peak latency measure. This algorithm was applied to the Switch-Reward and Switch-No Reward and electrode sites and values obtained were submitted to paired t-tests.

Results

P3

Figure 8 presents stimulus-locked grand averages for reward conditions across the midline (Fz, Cz, and Pz). To confirm that subjects are in fact categorizing the Stay-reward and Switch-reward outcomes separately, I conducted a 3 (electrode site) X 2 (reward) repeated measures ANOVA on the amplitude of the P3. Consistent with our predictions, this analysis revealed a significant main effect of electrode site ($F(2,22) =$

12.17, $p < 0.0001$). Also, there was a significant main effect of reward probability on the P300 amplitude, ($F(1,11) = 48.51$, $p < 0.0001$), but no interaction between electrode site and reward probability ($p > 0.21$). Post-hoc tests indicated that the amplitude of the P3 was significantly larger at electrode site Pz than at Cz ($p < 0.05$), as well as at Fz ($p < 0.001$). Additionally, post-hoc tests revealed that the amplitude of the P3 was significantly larger for the Switch-Reward condition than for the Stay-Reward Condition ($p < 0.0001$). Taken together, these results suggest that subjects were processing the Stay-reward and Switch-reward outcomes separately.

N2

Figure 9 presents stimulus-locked grand averages for the Switch-Reward and Switch-No Reward conditions at FCz. To confirm that subjects were attending to the feedback stimulus, I conducted a paired sample T-test on the amplitude of the N2. Consistent with my predictions, this analysis revealed that the peak amplitude of the N2 was significantly different between the Switch-Reward ($M = 1.03$, $SD = 3.9$) and to the Switch-No Reward Condition ($M = 4.2$, $SD = 2.20$), $t(12) = 4.1$, $p < 0.05$. These results suggest that subjects were attending to the feedback stimulus and not ignoring it.

fERN

Figure 9 presents stimulus-locked grand averages recorded at FCz for Switch-Reward, and Switch-No Reward conditions. A robust frontal negative deflection during the window of analysis was apparent across both conditions. This negative deflection was larger for the Switch-Reward condition than for the Switch-No Reward condition. This observation was unexpected, as I predicted the opposite. Figure 9 illustrates the Switch-No Reward and Switch-Reward difference wave associated with channel FCz. The

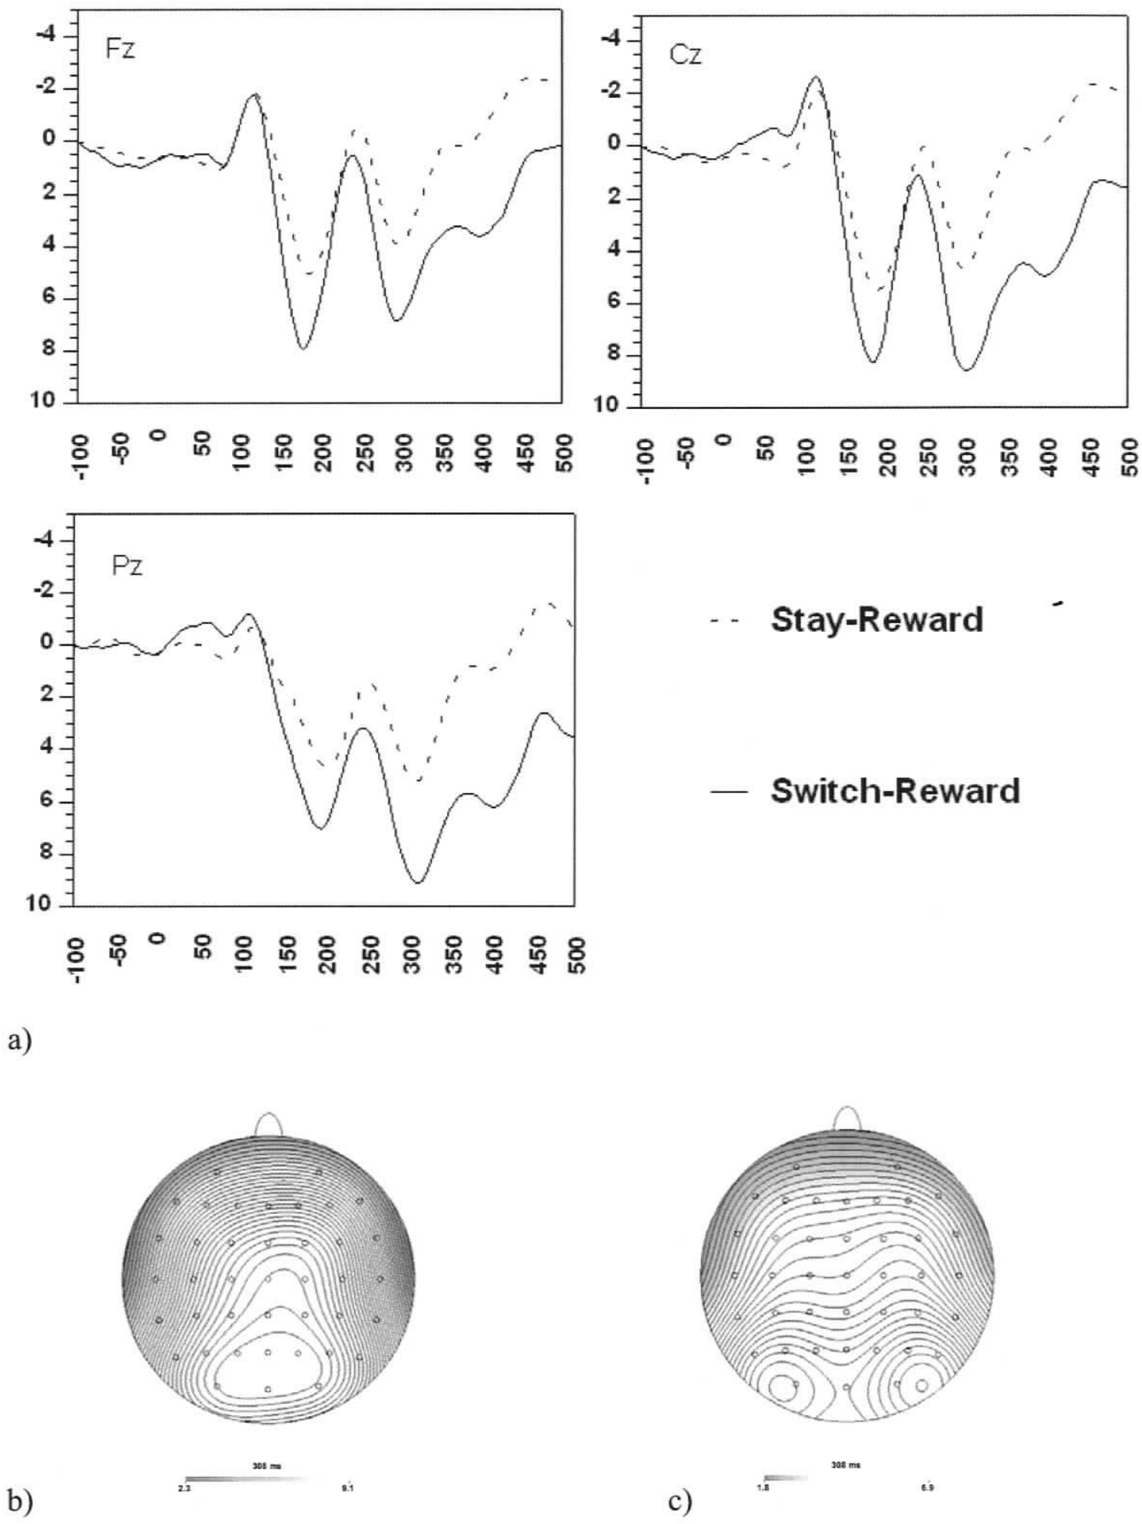

Figure 8. a) Averaged ERP waveforms recorded at midline channels (Fz, Cz, Pz) for Stay-Reward, and Switch-Reward conditions. b) Scalp voltage maps for the Switch-Reward condition at 308 ms following feedback onset. c) Scalp voltage maps for the Stay-Reward condition at 308 ms following feedback onset.

frontal maximum positive deflection peaking around 230 ms reflects this unexpected outcome. The peak analysis of this difference wave indicated that this positive deflection was maximal at FCz (4.9 uV, ± 1.04), peaked within the time window of the fERN (230 ms), and was significantly different from zero, t(11) = 4.54, $p < 0.001$. In addition, a negative peak was also observed in the difference wave. This negative deflection was found to be maximal at FCz (*M* = -3.08 uV, ± .89), and was significantly different from zero, t(11) = -3.44, $p < 0.005$. Interestingly, the peak analysis of this negative deflection peaked within the time window of the fERN (300 ms).The lower panel of figure 9 shows voltage maps for the maximal positive deflection around 230 ms (b), and the maximal negative deflection around 300 ms (c).

These results were confirmed by a spatial-temporal principle component analysis (PCA) with varimax rotation. The first component extracted by the PCA, which accounted for 43% of the overall variance, exhibited maximal loadings over frontal regions of the scalp. Further, three primary epochs associated with this component were identified; one from 372ms to 400 ms preceding feedback onset (accounting for 30.3% of the total variance), one from 216 ms to 244 ms preceding feedback onset (accounting for 18.8% of the total variance), and one from 276ms to 300 ms (accounting for 17.6% of the total variance). Importantly, the latter two epochs coincide within the time window of the fERN (Hajcak et al., 2006; Holroyd et al., 2004b; Nieuwenhuis et al., 2004; Holroyd & Coles, 2002; Miltner et al., 1997).

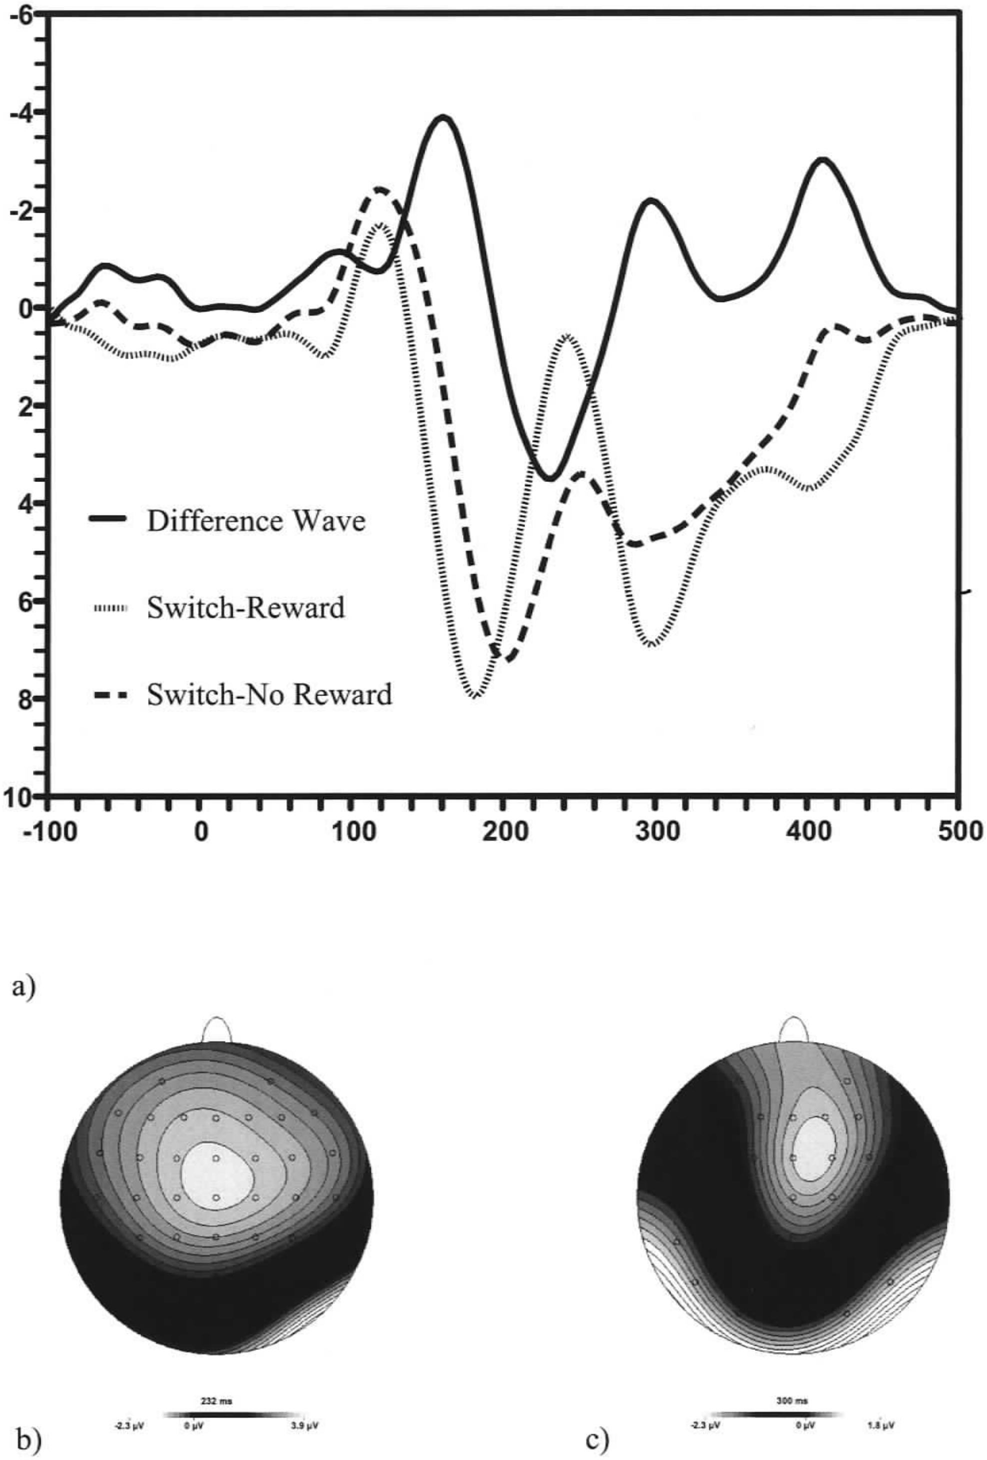

Figure 9. a) ERPs for the Switch-reward, switch no-reward, and associated difference wave at FCz (Switch-No Reward - Switch-Reward). b) Scalp voltage maps for the maximal positive deflection in the difference wave at 232 ms following feedback onset. c) Scalp voltage maps for the maximal negative deflection in the difference wave at 300 ms following feedback onset.

Discussion

Experiment 2 was conducted, in part, to rule out the possibility that the complexity of the feedback information reduced the amplitude of the fERN in experiment 1 (e.g., reward vs. no-reward and what the correct response should be). It was expected that the negative feedback in and of itself would elicit a fERN. Contrary to the prediction, it was found that the Switch-Reward condition elicited a larger negativity than the Switch-No reward condition, thereby producing a positive rather than negative deflection in the difference wave during the time period associated with the fERN. The difference wave analysis confirmed that this positive deflection was statistically significant, was maximal at front-central electrodes sites, and had a peak latency of 230 ms. This result was very surprising, given that the RL-ERN theory predicts that the fERN should be elicited by negative feedback when ongoing events are worse than expected (Holroyd and Coles, 2002). Although the results of experiment 2 are inconsistent with the fundamental prediction of the RL-ERN theory, recent empirical investigation of the N2 component may provide a better insight into these findings (Please see general discussion).

General Discussion

To determine the source of the fERN, this study followed a converging method approach by applying electrophysiological methods to a reward task that had been used in three previous studies that showed that positive and negative feedback differentially activate the ACC: a primate single-cell recording study, a human event-related fMRI study, and human single cell recording study. Experiment 1 presented a straight forward replication of the reward task used in the previous neuroscience studies. Although the results demonstrated morphological differences between the ERPs elicited by positive feedback and negative feedback, the topography and latency of this difference were inconsistent with the predictions of the RL-ERN theory. I interpreted these findings by inferring that the complexity of the feedback information may have reduced the amplitude of the fERN. To investigate this possibility, I conducted a second experiment that separated the instruction cue from the feedback cue. The Switch-Reward condition elicited a larger negativity than the Switch-No reward condition, a result that was inconsistent with the RL-ERN theory. Indeed, inspection of figure 9 illustrates a robust frontal negative deflection elicited by the onset of the feedback stimulus that was largest for the Switch-Reward condition. This finding may perhaps be understood in terms of the N2 component.

As described in the introduction (see page 6), the N2 is commonly elicited by an unexpected or infrequent stimulus in an oddball paradigm. If this large negative deflection elicited by the Switch-Reward condition is a N2 component, it can be argued that this may result from unexpected presentation of the reward stimulus following the switch cue. In experiment 2, two instruction cues preceded the feedback cues. The instruction cues consisted of a Stay and Switch Cue. For the Stay cue (| |), which appeared in 80% of the trials, participants were instructed to stay with the same button choice on the subsequent trial. Following this cue, the participants received positive feedback indicating that they received 3 cents for that trial. The instruction stay cue could therefore be interpreted as an early indicator of a reward, or favorable outcome. In this respect, the evaluating system would come to expect (or predict) that the stay cue would always indicate positive feedback. For the Switch Cue (< >), which appeared in 20% of the trials, participants were instructed to switch their button choice on the subsequent trial. Following this cue, participants received either negative feedback (no reward) or positive feedback. Since the positive feedback was only presented following the stay trials, it can be speculated that because participants were biased to expect positive feedback following stay trials, they incorrectly expected negative feedback following switch trials. If this were the case, then positive feedback following switch trials would come as a surprise. As described above, it is known that unexpected events elicit a large N2-P3 complex, for example by infrequent targets in the oddball paradigm. In experiment 2, the P3 was larger following switch-reward feedback than following switch-no reward feedback, consistent with hypothesis that the former feedback was more unexpected. I suggest that because subjects may have been biased to expect negative feedback following the switch cue, positive feedback was unexpected and therefore elicited a large N2, which has a latency and scalp distribution comparable to the fERN (Holroyd, 2004).

I propose to test this hypothesis with the following three experiments. I will apply the same paradigm used in experiment two, but modify the paradigm so that all the possible combinations of condition mappings can be applied (see figure 10). In experiment 3, trials following the switch cue will be followed by positive feedback, and trials following the stay cue will either be followed by positive or negative feedback (50% probability). In Experiment 4, the Stay cue, which will appear in 80% of the trials, will be followed by negative feedback. For the Switch Cue, which will appear in 20% of the trials, participants will either receive negative feedback, or positive feedback. Experiment 5 will follow the same feedback probabilities as Experiment 4 but the instruction cues will be reversed. If the participants are utilizing the instruction cue to develop their expectations of the subsequent feedback cue, and when these exceptions are

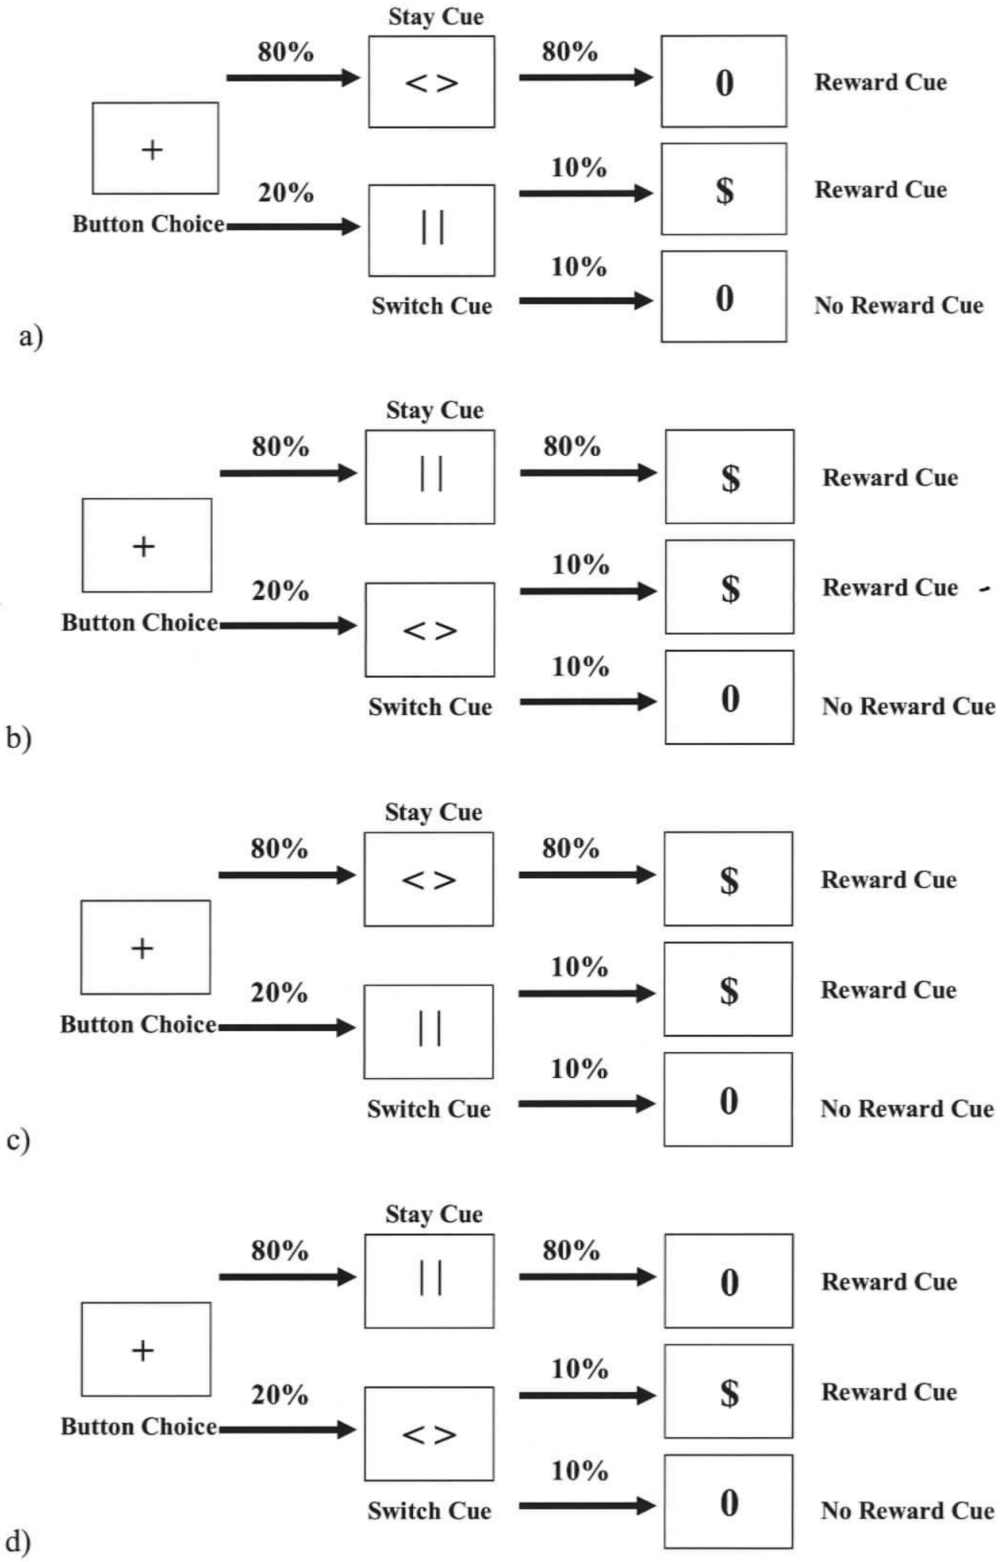

Figure11. a) Experiment 2. b) Experiment 3. c) Experiment 4. d) Experiment 5

violated, (i.e. Experiment 3: Stay-Reward, Experiment 4: Switch-No Reward, and Experiment 5: Stay-No Reward), a large N2 component will result. This would reproduce the findings in Experiment 2. Positive results will therefore replicate common findings in several electrophysiological investigations that a N2 is elicited by unexpected outcomes or expectancy violations.

Thus, I suggest that the unexpected stimuli in experimental 2 may have elicited a massive N2 that dominated the brain's response to positive feedback following the switch cue. It is also very interesting that our difference wave analysis revealed a small, but significant, negative deflection at 308 ms that was maximal over frontal central electrode sites, specifically FCz. This negative deflection is consistent with time window and scalp distribution of the fERN. Given that the previous studies on which our present experiment is modeled after did elicit ACC activity to negative feedback, our results can possibly be interpreted as follows: 1) the maximal positive deflection in the difference wave at front central brain regions at 230 ms may result from a large N2 elicited by unexpected rewards; 2) the maximal negative deflection in the difference wave at front central brain regions at 300 ms may reflect a delayed fERN. As shown in figure 11, this negativity in the difference wave occurs during the onset of the P3 wave. The Switch No-Reward P3 is slightly more negative than the Switch-Reward P3, and therefore possibly reflects component overlap between the P3 and the fERN. It can be inferred that the negativity elicited by a delayed fERN may reduce the positive deflection of the P3 associated with the Switch-no reward trial, since both conditions have the same stimulus probability and therefore should not differ in P3 amplitude. To examine this more thoroughly, I will monitor this negative deflection, or reduced P3, in the future experiments.

Many of the empirical studies that have led to our understanding the neural basis and functional significance of the fERN have been inspired by the RL-ERN theory. However, an important unproven assumption of the theory is that the ACC generates the fERN. In the past, studies have used the converging method approach to address this issue (Van Veen et al. 2004, Nieuwenhuis, et al., 2005). Although these studies used tasks – in contrast to the present study - elicited an fERN, they failed to find evidence of differential ACC activation by positive and negative feedback when using the same task in human fMRI experiments. In the present study, I found the converse: building on previous studies that showed that positive and negative feedback differentially activate the ACC, I was unable to demonstrate that the feedback in this same task also elicits the fERN. These results are consistent with the results of the previous converging method studies in that all of the results argue against the ACC as the source of the fERN.

A second concern is why the two experiments in this study did not elicit a typical fERN, even apart from the question of source localization. This result has revealed an apparent inconsistency with the RL-ERN theory. As noted above, this appears to be due to the effect of a large N2 with an apparently delayed fERN. If valid, the account of the systems that gives rise to the N2 and fERN may provide insight into how large-scale neural networks for processing unexpected vs. expected events and negative vs. positive feedback are mediated in biological systems and the role they play in executive control. In either case, it still remains unclear whether the systems that generates the N2 and fERN constitute two competing, parallel systems, or whether they work in concert, each contributing specialized functions. Thus, how these compete or cooperate in reinforcement learning tasks (with both frequency and feedback variables) constitutes an important and unresolved psychological problem. Future research is needed to resolve this issue.

References

Allman, J.M., Hakeem, A., Erwin, J.M., Nimchinsky, E., & Hof, P. (2001) Anterior cingulate cortex: The evolution of an interface between emotion and cognition. *Annals of the New York Academy of Sciences*, 935, 107-17. Review

Allman, J.M., Hakeem, A., & Watson, K. (2002). Two phylogenetic specializations in the human brain. *The Neuroscientist 8*, 335-46.

Baker, T. E., & Holroyd, C. B. (2006). Which way do I go? Neural activation in response to feedback processing and decision making in a virtual T-Maze. *Psychophysiology*, *43*, S22.

Bates, J. F., & Goldman-Rakic, P. S. (1993). Prefrontal connections of medial motor areas in the rhesus monkey. *Journal of Comparative Neurology, 336,* 211–228.

Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulated cortex. *Nature, 402,* 179-182.

Bush, G., Vogt, B. A., Holmes, J., Dale, A. M., Greve, D., Jenike, M. A., & Rosen, B. R. (2002). Dorsal anterior cingulate cortex: A role in reward-based decision making. *Proc. Natl. Acad. Science USA 99,* 523-528.

Bush, G., Luu, P., & Posner, M. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences, 4,* 215-223.

Cantania, C. A. (1993). Thorndike's legacy: Learning, selection, and the Law *of* Effect. *Journal of Experimental Analysis of Behaviour, 72,* 425-454.

Dien, J., Spencer, K. M., & Donchin, E. (1999). Parsing the late positive complex: Mental chronometry and the ERP components that inhabit the neighborhood of the P300. *Psychophysiology, 41,* 665-678.

Dien, J., Spencer, K. M., & Donchin, E. (2003). Localization of the event-related potential novelty response as defined by principle components analysis. *Cognitive Brain Research, 17,* 637-650.

Dochin, E., & Coles, M. G. N. (1988). Is the P300 component a manifestation of context updating? *Behavior Brain Science, 11,* 357–74.

Dum, R. P., & Strick, P. L. (1993). Cingulate motor areas. In B. A. Vogt, & M. Gabriel (Eds.), Neurobiology of cingulate cortex and limbic thalamus*: A comprehensive handbook* (pp. 415-441). Boston: Birkhauser.

FitzGerald, M. T., Folan-Curran., J. (2002) Clinical Neuroanatomy and related neuroscience. *Fourth Edition*. New York: W.B. Sanders.

Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1991). Effects of crossmodal divided attention on late ERP components. II. Error processing in choice reaction tasks. *Electroencephalography and clinical Neurophysiology, 78*, 447-455.

Gehring, W., Gratton, G., Coles, M., & Donchin, E. (1992). Probability effects on stimulus evaluation and response processes. *Journal of Experimental Psychology, 18,* 198-216.

Gaspar, P., Berger, B., Febvret, A., Vigny, A., & Henry, J. P. (1989). Catecholamine innervation of the human cerebral cortex as revealed by comparative immunohistochemistry of tyrosine hydroxylase and dopamine-beta-hydroxylase. *Journal of Comparative Neurology, 279,* 249–271.

Hajcak, G., Moser, J. S., Holroyd, C. B., & Simons, R. F. (2006). The feedback-related negativity reflects the binbary evaluation of good versus bad outcomes. *Biological Psychology, 71,* 148-154.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review, 109,* 679-709.

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., & Cohen, J. D. (2003). Errors in reward prediction are reflected in the event-related brain potential. *Cognitive Neuroscience and Neuropsychology Neuroreport, 14,* 2481-2484.

Holroyd, C. B. (2004). A note on the N200 and the feedback ERN. In M. Ullsperger & M. Falkenstein (Eds.), Errors, Conflicts, and the Brain. Current Opinions on Performance Monitoring, (pp. 211-218). Leipzig: MPI of Cognitive Neuroscience.

Holroyd, C. B., Larsen, J. T., & Cohen, J. D. (2004a). Context dependence of the event-related brain potential associated with reward and punishment. *Psychophysiology, 41,* 245-253.

Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R., Coles, M. G. H., & Cohen, J. D. (2004b). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature Neuroscience, 7,* 497-498.

Holroyd, C. B., Hajcak, G., & Larsen, J. T. (in press). The good, the bad and the neutral: Electrophysiological responses to feedback stimuli. *Brain Research.*

Holroyd, C. B., Yeung, N., Coles, M. G. H., & Cohen, J. D. (2005). A mechanism for error detection in speeded response time tasks. *Journal of Experimental Psychology: General, 134,* 163-191.

Jasper, H. H. (1958). The ten twenty electrode system of the international federation. *Electroencephalography and Clinical Neurophysiology, 10,* 371–375.

Luck, S. J. (2005). An Introduction to the Event-Related Potential Technique. Cambridge, MA: MIT Press.

Macdonald, A., Cohen, J., Stenger, A., & Carter, C. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cortex in cogntive control. *Science, 288,* 1835-1839.

Mars, R. B., De Bruijn, E. R., Hulstijn, W., Miltner, W. H. R., & Coles, M.G. H. (2004). What if I told you: 'You were wrong'? Brain potentials and behavioral adjustments elicited by performance feedback in a time-estimation task. In: M. Ullsperger, M. Falkenstein (Eds). Errors, Conflicts, and the Brain. Current Opinions on Performance Monitoring, (pp. 129-134). Leipzig: MPI of Cognitive Neuroscience.

Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a generic neural system for error detection. *Journal of Cognitive Neuroscience, 9,* 788-798.

Montague, P. R., Hyman, S. E., & Cohen, J. D. (2004). Computational roles for dopamine in behavioral control. *Nature, 431,* 760-766.

Morecraft, R. J., & Van Hoesen, G. W. (1998). Convergence of limbic input to the cingulate motor cortex in the rhesus monkey. *Brain Research Bulletin, 45,* 209–232.

Nieuwenhuis, S., Yeung, N., Holroyd, C. B., Schurger, A., & Cohen, J. D. (2004) Sensitivity of electrophysiological activity from medial frontal cortex to utilitarian and performance feedback. *Cerebral Cortex 14*, 741--747.

Nieuwenhuis, S., Holroyd, C. B., Mol, N., & Coles, M. G. H. (2004). Reinforcement-related brain potentials from medial frontal cortex: Origins and functional significance. *Neuroscience and Biobehavioural Reviews, 28,* 441-448.

Nieuwenhuis, S., Slagter, H. A., von Geusau, N. J. A., Heslenfeld, D., & Holroyd, C. B. (2005). Knowing good from bad: Differential activation of human cortical areas by positive and negative outcomes. *European Journal of Neuroscience, 21,* 3161-3168.

Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus-norepinephrine system. *Psychological Bulletin, 131,* 510-532.

Pakzad-Vaezi, K., Krigolson, O. E., & Holroyd, C. B. (2006). Underlying unity: The relationship between the N200 and the feedback ERN. *Psychophysiology, 43,* S75.

Previc, F. H. (1999). Dopamine and the origins of human intelligence. *Brain and Cognition, 41*, 299-350.

Nimchinsky, E. A., Gilissen, E., Allman, J. M., Perl, D. P., Erwin, J. M., & P. R. Hof. (1999). A neuronal morphologic type unique to humans and great apes. *Proceedings of the National Academy of Sciences*. 96, 5268-5273.

Pandya, D. N., Van Hoesen, G. W., & Mesulam, M. M. (1981). Efferent connections of the cingulate gyrus in the monkey. *Experimental Brain Research, 42,* 319–330.

Posner, M. I., & DiGirolamo, G. J. (1998). Executive attention: Conflict, target detection, and cognitive control. In R. Parasuraman (Ed.), *The attentive brain* (pp. 401–423). Cambridge, MA: MIT Press.

Samejima, K., Ueda, Y., Doya, K., & Kimura, M. (2005). Representation of Action-Specific Reward Values in the Striatum. *Science, 25,* 1337 – 1340.

Schultz, W., Romo, R., Ljungberg, T., Mirenowicz, J., Hollerman, J. R., & Dickinson, A. (1995). Reward-related signals carried by dopamine neurons. In J. Houk, J. Davis, & D. Beiser (Eds.), *Models of information processing in the basal ganglia.* Cambridge, MA: MIT Press.

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

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

Schultz, W., Tremblay, L., & Hollerman, J. R. (2000). Reward processing in primate orbitofrontal cortex and basal ganglia. *Cerebral Cortex, 10,* 272-283.

Seamans, J. K., & Yang, C. R. (2004). The principal features and mechanisms of dopamine modulation in the prefrontal cortex. *Progress in Neurobiology, 74,* 1–57.

Shima, K., & Tanji, J. (1998). Role for Cingulate Motor Area Cells in Voluntary Movement Selection Based on Reward. *Science 282*, 1335-1338.

Suri, R., & Schultz, W. (2001) Temporal Difference Model Reproduces Anticipatory Neural Activity. *Neural Computation, 13,* 841-862.

Reynolds, J. N. J., Hyland, B. I., & Wickens, J. R. (2001). A cellular mechanism of reward-related learning. *Nature, 413,* 67-70.

Rugg, M., & Coles, M. (1995). Electrophysiology of mind: Event-related brain potentials and cognition. Oxford, NY: Oxford University Press

Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive Coding of Reward Value by Dopamine Neurons. *Science, 307,* 1642-1645.

Ullsperger, M., & Von Cramon, D. Y. (2003). Error processing using external feedback: Specific roles of the habenular complex, the reward system, and the cingulate motor area revealed by fMRI. *Journal of Neuroscience, 23,* 4308-4314.

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 and Cognition, 56,* 267-276.

Vogt, B. A. (1993). Structural organization of cingulate cortex: Areas, neurons, and somatodendritic transmitter receptors. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 19–70). Boston: Birkhauser.

Wickens, J. R., Reynolds, J. N., & Hyland, B. I. (2003). Neural mechanisms of reward-related motor learning. *Current Opinion in Neurobiology, 13,* 685-690.

Williams, Z. M., Bush, G., Rauch, S. L., Cosgrove, G. R., & Eskandar, E. N. (2004). Human anterior cingulated neurons and the integration of monetary reward with motor responses. *Nature Neuroscience Advance Online Publication, 21,* 1-6.

Yang, C. R., & Seamans, J. K. (1996). Dopamine D1 receptor actions in layers V–VI rat prefrontal cortex neurons in vitro: Modulation of dendritic–somatic signal integration. *Journal of Neuroscience, 16,* 1922–1935.

Yeung, N., Botvinick, M., Cohen, J. (2004). The Neural Basis of Error Detection: Conflict Monitoring and the Error-Related Negativity. *Psychological Review, 111*, 931-959.

Appendix A

Event-Related Brain Potential

The electroencephalogram (EEG) is a recording of the electrical activity of the cerebral cortex with electrodes placed on the scalp (Luck, 2005). The raw EEG represents the sum of extracellular current flow associated with dendritic excitatory and inhibitory postsynaptic potentials of hundreds of thousands of cortical connections. Specifically, consider the electrical activity of a single pyramidal cell activated by an afferent pathway (See figure 12a). Here, pre-synaptic neurons send excitatory signals which depolarize the apical dendrites of the post-synaptic neuron and produce excitatory post-synaptic potentials. This inward current at the synapse, which is referred to as the "sink", flows down the apical dendrite and towards sites distant from the synapse, especially the cell body, where it moves outward across the cell body. These sites are referred to as the "source" and this outward flow of positive charge leaves the cell. At this instant, the combination of current source and sink describes a dipole, with a relatively negative charge at the distal part of the apical dendrite and a positive charge closer to the cell body. These dipoles are thought to be physically located in cortical tissue, within cortical columns of neurons that have specialized functions and that act together. Thus, an electrode placed near the end of the apical dendrite detects a negative potential or negative deflection in the EEG (active site) (note that this activity is measured with respect to a distant reference electrode; see methods). Although it is impossible to record the individual post-synaptic potentials at a single neuron, two cortical properties permit us to record the brain's electrical potential. First, pyramidal neurons tend to be organized in parallel, having the same relative orientation and, therefore, dipole polarity. Second, many neurons are synchronously activated. The summation of the dipoles, created at each of thousands of neurons, creates an electrical potential detectable at the scalp (See figure 12b). And by time-locking the beginning of the EEG epochs to a series of events (e.g., stimulus presentations or motor responses) and then averaging together these epochs to reduce the signal to noise ratio, we can observe the many voltage changes that are specifically related to the brain's response to that stimulus or response. It is these voltage changes that constitute the ERP.

The waveform features (i.e. peaks or troughs) are characterized by their amplitude (in uV), latency (in ms), polarity (+, -), and scalp distribution (electrode sites). Amplitude can be measured by numerous methods, but is commonly measured in relation to some other feature of the waveform or in relation to a baseline. Latency is most commonly measured in terms of the temporal relationship between the feature of the waveform following the event of interest (for more detail, see Rugg & Coles, 1995; Luck, 2005). Thus, ERP components are defined by their amplitude, polarity, latency, and scalp distribution which, by systematic experimental manipulation, can be shown to be functionally related to an underlying cognitive process. This technique provides a measure of cortical activity with high temporal resolution, but with limited spatial resolution. For more extensive descriptions of ERPs and ERP components, see reviews by Coles and Rugg (1995), Hillyard and Picton (1987), Naatanen and Picton (1986), and Luck (2005).

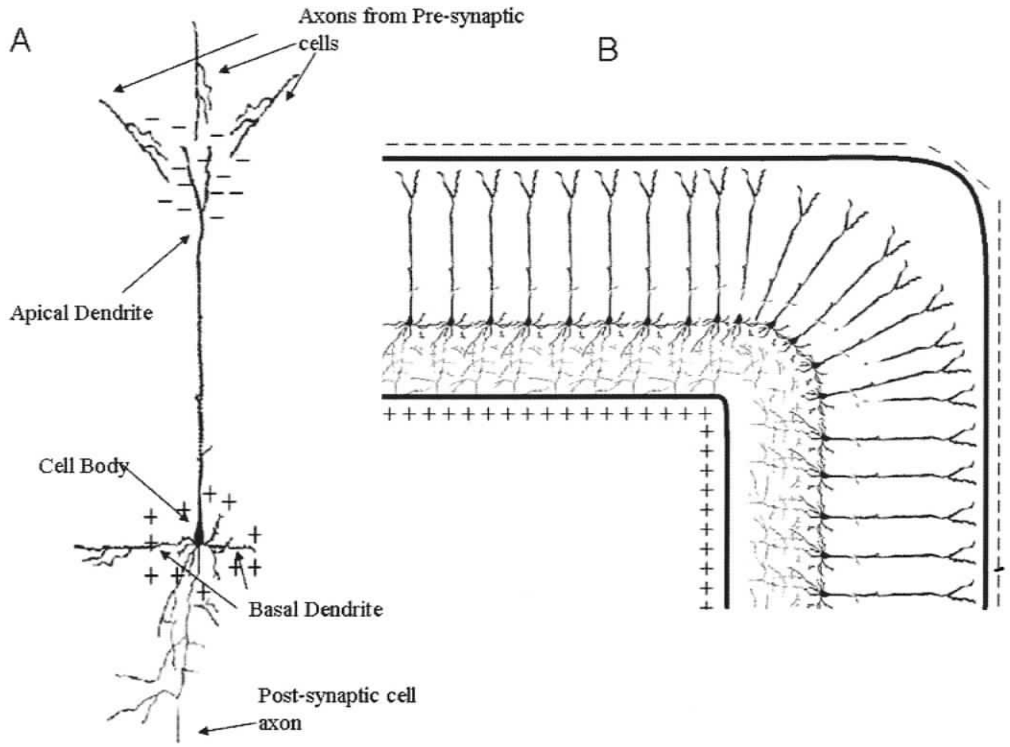

Figure 11. A) A pyramidal cell during neurotransmission. Pre-synaptic neurons send excitatory signals that depolarize the apical dendrites of the post-synaptic neurons, giving rise to post-synaptic potentials. This creates a current sink at the distal part of the dendrite and a current source closer to the cell body, yielding a small dipole. B) A stimulated region of cortex. The summation of the dipoles created at each of thousands of neurons creates an electrical potential detectable at the scalp.

Appendix B

The Anterior Cingulate Cortex

The ACC is the frontal part of the cingulate cortex and includes Broadmann's areas 24 (ventral ACC) and 32 (dorsal ACC). The ACC forms a collar around the corpus callosum, and its sulcus is a prominent feature of the medial wall, composing about half of its surface (see figure 13). Characterized by a primitive cytoarchitecture, the ACC is distinguished from much of the neocortex by a reduced or absent layer 4 and a well-developed layer 5, which is a characteristic of motor areas. Stimulation to this area evokes movement and this area is somatotopically mapped (Allman et al., 2001). Layer 5 of the ACC notably houses two classes of neurons, pyramidal neurons and spindle neurons. As shown in Figure 12, the pyramidal neurons have triangular cell bodies with extensive dendritic arborizations. The dendrites consist of a single apical dendrite extending towards the pial surface of the cortex, and multiple basal dendrites extending from the cell body. From the cell body, a single axon projects towards several areas concerned with directing motor behavior, such as the basal ganglia, supplementary motor area, primary motor area, or spinal cord (Dum & Strick, 1993; Van Hoesen et al., 1993). Spindle neurons, by contrast, are characterized by a large spindle shaped soma, a large apical dendrite extending toward the pial surface of the cortex and a single large basal dendrite extending toward the underlying white matter (Allman et al., 2001; Allman et al., 2002). Although spindle neurons have only recently been discovered and their function is poorly understood, it is apparent that they are only found in the ACC of humans and other hominids (Nimchinsky et al., 1999).

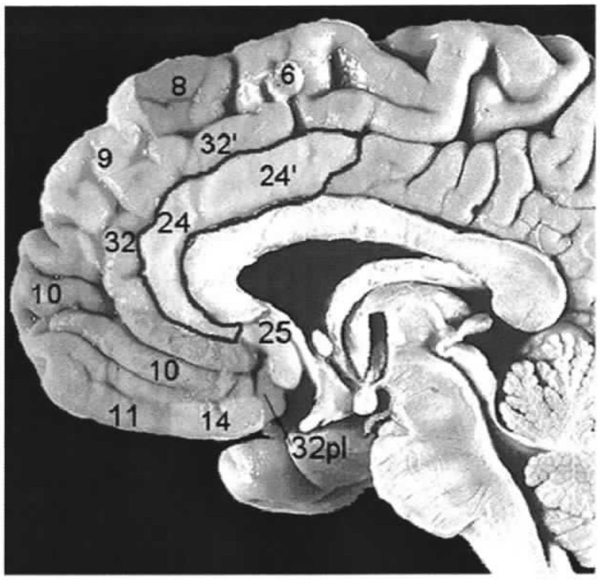

Figure 12. Anterior cingulate cortex, Area 24 and 32 (Bush et al. 2000).

The ACC belongs to the rostral limbic system and receives widespread afferent projections from the limbic lobe, which includes the amygdala, ventral striatum, orbitalfrontal cortex, and the anterior insular cortex (Fitzgerald & Folan-Curran, 2002). For that reason, it is proposed that the ACC affords a critical pathway devoted to the regulation of emotional and motivational factors that influence motor activity (Morecraft & Van Hoesen, 1998; Pandya, Van Hoesen, & Mesulam, 1981; Bush et al., 2000). Moreover, the ventral bank of the ACC is richly interconnected with the dorsolateral prefrontal cortex and tends to be concurrently activated when there are increased demands both for strategic processes and for evaluating the output of the monitoring system (Macdonald et al. 2000). In such cases, activity of the dorsolateral prefrontal cortex is thought to be responsible for maintaining and representing attentional demands of a task at hand, and the generation of contextually appropriate behaviours in the absence of external stimulation, while the ACC is involved in an evaluative process, such as monitoring the occurrence of errors (Yeung, Botvinick, & Cohen, 2004). As a result, it is suggested that the ACC motor areas provide an important route for the dorsolateral prefrontal cortex to influence motor output (Bates & Goldman-Rakic, 1993; Morecraft & Van Hoesen, 1991). Furthermore, converging evidence from event-related brain-potential studies (ERP), functional magnetic resonance imaging (fMRI) studies, human single-cell recording studies, and primate single-cell has suggested that the caudal region of the ACC contributes to high-level cognitive control of motor behavior, especially by using rewards to guide action selection (Holroyd & Coles, 2002; Holroyd et al., 2005; Miltner, Braun, & Coles, 1997; Nieuwenhuis, Holroyd, Mol, & Coles, 2004; Bush et al., 2002;Williams et al., 2004; Shima & Tanji, 1998). In all these cases, it has been suggested that functions such as reward processing, error detection, anticipation of tasks, motivation, and modulation of emotional responses are attributed to the ACC (Bush et al., 2000; Nieuwenhuis et al., 2001; Posner & DiGirolamo, 1998).

Appendix C

Participant Questionnaire

You may decline to respond to any or all questions on this questionnaire.
Answer the first set of questions using the five point scale below the question.

1. How fatigued were you before you started the experiment? Please rate on a scale of 1 to 5, with 1 meaning "not at all" and 5 meaning "a lot".

   Please circle:

   1 2 3 4 5

2. How bored did you get during the experiment? Please rate on a scale of 1 to 5, with 1 meaning "not at all" and 5 meaning "very".

   Please circle:

   1 2 3 4 5

3. How hard was it to make a decision regarding making the correct button press? Please rate on a scale of 1 to 5, with 1 meaning "not at all" and 5 meaning "a lot".

   Please circle:

   1 2 3 4 5

4. Did earning money mean much to you? Please rate on a scale of 1 to 5, with 1 meaning "not at all" and 5 meaning "a lot".

   Please circle:

   1 2 3 4 5

5. Did not earning money mean much to you? Please rate on a scale of 1 to 5, with 1 meaning "not at all" and 5 meaning "a lot".

   Please circle:

   1 2 3 4 5

6. How well do you think you did on the choice task? Please rate on a scale of 1 to 5, with 1 meaning "Poor" and 5 meaning "Great".

   Please circle:

   1 2 3 4 5

7. Did you ever feel that the feedback stimuli was incorrect? (yes / no)?

8. Did you develop any strategies about how to respond? (yes / no)?

   If yes, please explain:

9. Did you think you were going to be paid at the end of the experiment?

   (yes / no)?

10. If you could change anything in this experiment, what would it be.

11. Can you please describe what the following symbols indicated.

    1. +
    2. ||
    3. <>
    4. $
    5. 0

10. In your own words, can you please describe the choice task you just performed?

Thank you for your participation.