


2014

Effects of expectancy and agency on the feedback-related negativity

Stephen J. Anderson
Iowa State University

Follow this and additional works at: <http://lib.dr.iastate.edu/etd>

 Part of the [Cognitive Psychology Commons](#), and the [Neuroscience and Neurobiology Commons](#)

Recommended Citation

Anderson, Stephen J., "Effects of expectancy and agency on the feedback-related negativity" (2014). *Graduate Theses and Dissertations*. 14098.
<http://lib.dr.iastate.edu/etd/14098>

This Thesis is brought to you for free and open access by the Graduate College at Iowa State University Digital Repository. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Effects of expectancy and agency on the feedback-related negativity

by

Stephen J. Anderson

A thesis submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Major: Psychology

Program of Study Committee:
Robert West, Major Professor
Alison Morris
Kristi Costabile

Iowa State University

Ames, Iowa

2014

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENTS	iii
ABSTRACT	iv
CHAPTER 1 BACKGROUND	1
Introduction	1
FRN: History and Theory	3
Previous Research from Our Lab	11
Theoretical Justification	14
CHAPTER 2 EXPERIMENT 1	15
Hypothesis	15
Introduction	15
Method	16
Results	20
Discussion	24
CHAPTER 3 EXPERIMENT 2	26
Hypothesis	26
Introduction	26
Method	29
Results	32
Discussion	34
CHAPTER 4 GENERAL DISCUSSION	36
REFERENCES	40

ACKNOWLEDGEMENTS

This thesis benefitted greatly from the contributions of several people. I would like to thank Kira Bailey and Paul Kieffaber for their creative input on the research that inspired this study. I especially thank Robert West for his time and resources that were vital to the completion of this project, and for his helpful comments and support throughout the development of this thesis.

ABSTRACT

There is an evolving body of research using Event-Related Potentials (ERPs) to study electrophysiological activity in the brain associated with feedback processing in gambling tasks. In particular, investigators have examined the roles of outcome valence (Gehring & Willoughby, 2002) and outcome probability (Bellebaum & Daum, 2008) on an ERP component known as the Feedback-Related Negativity (FRN), whose amplitude has been shown to distinguish positive and negative outcomes. The current study examined the possibility that FRN amplitude might also be modulated by a person's subjective expectation of a particular outcome (expectancy), as well as whether or not an outcome was the result of their own actions or those of another (agency). Consistent with past research, the results of Experiment 1 of this study show an increase in the amplitude of the FRN when participants are given feedback indicating a loss during a decision-making task. This increased negativity in response to negative feedback was greater for trials in which participants were led to expect a win, though only for trials in which the value of the outcome was relatively smaller. In Experiment 2, the FRN was also found to be greater in amplitude on trials where losses resulted from choices made by participants as opposed to the choices of a computerized opponent. The findings of the current study suggest a somewhat nuanced account of the FRN in which amplitude is sensitive to individual expectations about winning and losing as well as the individual's perceived role in wins and losses.

CHAPTER 1

BACKGROUND

Introduction

The purpose of this study was to examine the effects of expectation and agency on neural activity associated with feedback processing. As psychology is commonly defined as the study of human behavior, a major goal within our academic discipline is to understand the process of decision making. Specifically, we want to be able to predict how an individual will behave in circumstances where multiple courses of action are available, given that her selection among them is underdetermined. Integral to this goal is understanding feedback processing, i.e., task-relevant neural activity that occurs following the presentation of stimuli containing information about the outcomes of behaviors that may then be used to inform similar decisions in the future. Identifying specific patterns of neural activity present during feedback processing may help researchers understand how adaptive processing occurs in the normal healthy brain, as well as help clinicians understand impaired decision processes in individuals who have suffered damage to the relevant areas. Further, identifying real-time neural activity relevant to feedback processing and observing how this activity is modulated by decision outcomes may further elucidate the process of human learning by showing specific features of feedback to which our brains are sensitive.

A growing body of research, including studies using functional neuroimaging (Carter, Botvinick, & Cohen, 1999), lesion (Janer & Pardo, 1991), and event-related potential (ERP) (Gehring & Willoughby, 2002; Brown & Braver, 2005) methodologies, has identified the anterior

cingulate cortex (ACC) as a region of the brain commonly recruited during feedback processing. In particular, the feedback-related negativity (FRN) has been identified in electrophysiological studies as an event-related potential (ERP) component that is sensitive to several aspects of decision outcomes, and there is converging evidence that the source of this component lies in the ACC (Gehring & Willoughby, 2002; Walsh & Anderson, 2012).

A recent study by our lab investigating the neural correlates of feedback processing in a virtual blackjack game produced the surprising finding that the amplitude of the FRN varied significantly between two different types of negative outcomes: busts (i.e., the dealer wins because the player has drawn cards whose total value exceeds 21 points) and losses (i.e., the dealer wins because the dealer has a superior hand and neither party has exceeded 21 points) (West, Bailey, Tiernan, Boonsuk, & Gilbert, 2012; West, Bailey, Anderson, & Kieffaber, 2014). This effect was not anticipated based upon previous studies of the FRN (see Gehring & Willoughby, 2002; Miltner, Braun, & Coles, 1997) or models of ACC functioning (see Alexander & Brown, 2011). Here I explore two explanations to account for the observed differences between busts and losses. First, I hypothesized that the FRN was indexing differences in players' expectations of the likelihood of these two outcomes. In this case, while both busts and losses might be similar outcomes in terms of their negative consequences, the FRN for busts might be larger because busts were a less expected outcome. Second, I hypothesized that the relatively larger amplitude of the FRN following a bust was the result of players' increased subjective responsibility for the outcome. In this case, a loss might be characterized as being beaten by a dealer with a superior hand, whereas a bust might be characterized as having beaten oneself.

To test these two hypotheses, a simple binary selection task was created in which subjects were asked to choose between two boxes appearing on a display in an attempt to find a winning gold token. By manipulating participants' expectation of trial outcomes and whether or not those outcomes were the result of their choice or that of an opponent, this study was intended to shed light on how the FRN responds to conditions of varied expectancy and agency. If it can be shown that the FRN is sensitive to either of these variables, it may indicate that one or both factors influence the early stages of feedback processing.

FRN: History and Theory

Studying the neural correlates of feedback processing using an ERP framework generally involves identifying waveforms with parameters (i.e., latency and amplitude) that vary in response to different characteristics of an outcome. The FRN is an ERP component studied extensively in conjunction with feedback processing related to volitional behavior, and has been associated with a neural error-detection system (Miltner et al., 1997; Walsh & Anderson, 2012). The FRN represents a transient negativity at frontocentral electrodes that peaks between 250-350 ms following feedback (Gehring & Willoughby, 2002; Walsh & Anderson, 2012). There is a general consensus that the FRN is sensitive to the valence of decision outcomes, with amplitude increasing in response to negative feedback (Miltner et al., 1997; Holroyd & Coles, 2002; Gentsch, Ullsperger, & Ullsperger, 2009). However, the response properties of the FRN are not a straightforward matter.

The FRN was first described by Gehring and Willoughby (2002). The authors used a decision task in which participants were presented with two squares on a computer monitor and asked to select between them by pressing a key. Each box displayed a number (either 5 or

25) indicating the number of points that could be won or lost depending on whether the box was later revealed to represent a gain or loss. Following participants' selection, the numbers in the boxes turned either red or green, with red indicating a loss and green indicating a gain. For example, if the boxes contained values 5 and 25, the participant chose the box containing 5, and the selected box turned red, the participant lost 5 points from their total score. Conversely, if the selected box turned green, the participant gained 5 points. At the conclusion of the experiment, participants were paid the value of their final score at \$.01/point. Importantly, the design of this experiment allowed participants to observe not only the valence and magnitude of the actual outcome, but also the same parameters of the alternative outcome (i.e., the amount they would have won or lost had they chosen the alternative box). For instance, in the above example, if the participants chose 5 instead of 20 only to have both boxes turn green, then they had won 5 points, but had chosen the worse of the two possible outcomes. Losses, then, were not necessarily the same as errors. However, the authors observed that the FRN only indexed losses and gains, with losses producing more negative FRNs than gains. In contrast, their analyses did not reveal any significant effect of alternative outcomes on the FRN, such as a reduction in amplitude when a loss was the better of two possible losses. Gehring and Willoughby thus concluded that the FRN was sensitive to loss-gain status, and not to errors as such. A variety of researchers have arrived at a similar conclusion. For instance, Hajcak, Moser, Holroyd, & Simmons (2006) varied multiple outcome parameters in a gambling task and observed modulation of the FRN only in response to the valence (i.e., win vs. loss) of the outcome (see also Holroyd, Krigolson, Baker, Lee, & Gibson, 2009).

The FRN and the Probability of Outcomes

Other research has focused on more subtle features of feedback to which the FRN may be sensitive. A review by Walsh and Anderson (2012) compared over 200 articles published on the FRN to gain a more coherent picture of the factors that affect its presence and size. To this end, they identified 25 articles within the larger body of literature that seemed to present a consistent picture of the FRN as being larger for improbable outcomes. Notable among these is a study by Bellebaum and Daum (2008) in which the researchers employed a probabilistic selection task to measure participants' response to likely and unlikely outcomes. Participants were presented with a screen containing 24 boxes: 12 on the left side of the screen, 12 on the right. Their task was to indicate which of the two arrays of 12 boxes was more likely to contain a box hiding a 5-cent token. On each trial, the computer would reduce the number of possible winning boxes on each side by "pre-selecting" either 4 or 8 of the 12 boxes, leading subjects to believe that they would be better off choosing the array that contained more pre-selected boxes. However, the game was programmed so that subjects would eventually learn that the token was always located on the right side of the 12-box array that contained it, such that the optimal strategy on any given trial was to select the array with more pre-selected boxes in the right column. Reward probability was thus manipulated by varying the relative number of pre-selected boxes in the right-hand column on either side of the screen.

Bellebaum and Daum (2008) found that following (but not prior to) rule acquisition, the amplitude of the FRN was significantly larger for low-probability non-rewards than high-probability non-rewards. Conversely, no such effect was observed between probability conditions on trials where participants received rewards, which is consistent with the idea that

the FRN is sensitive to outcome probability, but only in the case of negative outcomes. The explanation offered for why the FRN showed differential coding of reward probabilities only following rule acquisition was that “Only the [participants who learned the rule] developed specific reward predictions, the violations of which were reflected in FRN amplitude modulations” (Bellebaum & Daum, 2008, p. 1833). This is probably not precisely correct, since the number and arrangement of pre-selected boxes in the array should still lead participants to form probabilistic expectations of rewards even before the rule is learned – just not the correct expectations. However, given that a reliable difference in FRN amplitude was found between conditions differing in their likelihood of reward/nonreward, this study seems to provide strong evidence that the FRN is sensitive to the probability of outcomes.

While Walsh and Anderson (2012) state that the bulk of published work points to an inverse relationship between reward probability and FRN amplitude, they are also quick to identify alternative perspectives. Representative among these is a paper by Hajcak, Moser, Holroyd and Simons (2007) that described multiple experiments using a gambling task in which reward probability was manipulated by varying how many of 4 boxes displayed on a screen contained a winning token. In this study, subjective appraisal of reward probability was directly measured by asking participants to indicate whether or not they would win on each trial following the cue indicating reward probability. Contrary to other studies, and despite finding that subjects’ appraisal was sensitive to the cued reward probability, the researchers found no effect of reward probability on FRN amplitude in their initial experiment. Thus, even though subjects could verbally report appropriate expectations of decision outcomes, the manipulation failed to produce a measurable effect on the FRN. The authors were able to observe

modulation of FRN amplitude in response to reward probability only when they instructed participants to make their predictions of winning or losing immediately after they made their selection. The researchers' account of this effect was that "Participants' predictions fluctuated over the course of each trial, solidifying only after participants were committed to a response" (Hajcak et al., 2007, p. 911). Given the extent to which the task had to be constrained to produce a reliable effect of reward probability on FRN amplitude, it is not obvious that the FRN can be taken as a reliable index of violations of subjects' expectations across tasks. Thus, the role of expectancy in modulating the FRN remains an open question.

The FRN and the Magnitude of Wins and Losses

Another variable of interest in the feedback processing literature is the magnitude of gains and/or losses. While evidence for magnitude effects are not consistently found (Hajcak, Holroyd, Moser, & Simons, 2005) or found only under highly constrained conditions (Hajcak et al., 2006; Holroyd, Larsen, & Cohen, 2004) the results of several studies indicate that the FRN can reliably code the size of potential rewards (Bellebaum, Polezzi, & Daum, 2010; West et al., 2014). Bellebaum et al. (2010) report using a task in which - similar to that used in their study discussed above - participants were told to guess which of two arrays of boxes contained a box concealing a winning token. The probability of winning was again manipulated by "pre-selecting" boxes and requiring subjects to learn a rule that would give them an advantage (i.e., the coin was always hidden in the lower row of boxes), but in this version, the token could be worth either 5¢, 20¢, or 50¢. The results showed that when (and only when) feedback was negative, 20¢ and 50¢ produced a reliably larger FRN than 5¢ trials. The authors speculated that a possible reason that reliable magnitude effects had not been observed elsewhere was

because the reported study required participants to learn a rule, “inducing the expectation of a potential reward which was then obtained or not” (Bellebaum et al., 2010, p. 3348). Given the tentative nature of this proposed requirement, as well as the fact that a variety of other studies have failed to observe reliable effects of reward magnitude on the FRN, further research should be conducted to establish the effects of outcome magnitude on this component.

The Relationship Between the FRN and Other ERP Components

A final point of inquiry regarding the FRN warrants brief mention. Some researchers have postulated that the FRN is best understood not in isolation, but rather as an element of a larger complex. Specifically, it has been suggested that the FRN, along with the P2 (an earlier, positive-going component displaying greater amplitude for gains vs. losses) and the P3a (a later, positive-going component displaying greater amplitude for wins vs. losses and unexpected vs. expected outcomes) are actually peaks and trough in phase-locked theta-band activity triggered by feedback stimuli (Potts, Martin, Burton, & Montague, 2006; Holroyd, Pakzad-Vaezi, & Krigolson, 2008; West et al., 2014). While the current study is not set up to test this idea directly, it is worth noting that the concept of the FRN as an independent entity is not universally recognized. This having been acknowledged, the point that is critical to the present discussion, namely that there exists a negative-going peak in ERP data recorded at mediofrontal electrodes which is known to code aspects of feedback stimuli, is not in dispute.

The Role of the Anterior Cingulate Cortex in Feedback Processing

There is reasonable consensus in the neuroscientific literature that the neural generators of the FRN lie within the ACC (though see Carlson, Foti, Mujica-Parodi, Harmon-Jones, & Hajcak, 2011), a region of the medial prefrontal cortex (mPFC) that surrounds the

rostral region of the corpus callosum (Devinsky, Morell, & Vogt, 1995; Allman, Hakeem, Erwin, Nimcinsky, & Hof, 2006.) The ACC is part of the rostral limbic system, along with the amygdala, peraqueductal grey, ventral striatum, orbital frontal cortex, and anterior insula; a network implicated in the motivational appraisal of stimuli and control of context-dependent behaviors. (Devinsky et al., 1995). The ACC itself appears to be an evolutionarily recent specialization of the neocortex characterized by a distinct class of spindle cells observed only in humans and great apes, and not in humans until roughly the age of four (Allman et al., 2006).

The ACC has been functionally and cytoarchitecturally divided into dorsal and ventral aspects, sometimes referred to as “cognitive” and “affective” aspects, respectively (Devinsky et al., 1995; Bush, Luu, & Posner, 2000). The dorsal aspect is reciprocally connected with prefrontal and parietal regions of the cortex, as well as motor regions and the spinal cord (Devinsky et al., 1995). The ventral aspect, on the other hand, has dense connections with the periaqueductal grey and amygdala, though it too has motor afferents. This network of interconnectivity has the potential to provide the ACC with a great deal of information concerning behavior (e.g., via motor connections), motivational states (e.g., via the amygdala), and executive processes (e.g., via the prefrontal cortex). As further support for the characterization of the ACC as an intersection of numerous processing sources, it is noteworthy that lesions to the ACC have been reported to produce an appreciable variety of emotional and behavioral disorders, ranging from inattention and emotional instability to akinetic mutism and dysregulation of autonomic functions (Bush et al., 2000).

One emerging theory of the mPFC is that it functions as an action-outcome predictor that serves to update outcome predictions in response to unexpected feedback. Early research

on ACC functioning demonstrated that this area was active in situations where participants were required to select from among multiple incompatible response options (Botvinick, Nystrom, Fissel, & Carter, 1999; MacDonald, Cohen, Stenger, & Carter, 2000) and also in situations where actual outcomes differed from participants' intended outcomes (Ito, Stuphorn, Brown, & Schall, 2003). In an attempt to reconcile these two observations, Brown and Braver (2005) formulated the Error Likelihood Model of the ACC. The theory posited that increased activity in the ACC corresponded to an increase in the perceived likelihood of making an error, and that this activation served as a signal to recruit increased cognitive control. As support, the authors showed that their model was more accurate than a similarly designed response conflict model in simulating the pattern of ACC responses observed in humans during a simulation of the stop-signal task (Brown & Braver, 2005).

A more recent extension of the error likelihood model, known as the predicted response-outcome (PRO) model was described by Alexander and Brown (2011). According to this framework, learned predictions of the probability and timing of outcomes are coded at the neuronal level and multiple predictions are separately but simultaneously activated. The authors theorize that these prediction-related activations are suppressed when a predicted outcome is actually observed, so that activity in this region will be greatest when actual outcomes deviate significantly from predictions. The authors ran a series of simulations using a neural network instantiation of this model and found that its activity captured a variety of known feedback-related effects in humans.

First, they evaluated the PRO model during a simulated version of a signal-change task, a paradigm used to study response conflict. In this task, an on-screen stimulus cues a

participant to make one of two behavioral responses, but this cue is countermanded on a subset of trials by a subsequent cue instructing the participant to make the alternate response. During this task, the authors reported that the PRO model produced an error signal that reliably coded error likelihood and response conflict, both of which have been reported to modulate the FRN. Second, they evaluated the PRO model during a simulated version of the Eriksen flanker task and found that the error signal generated by the PRO model reliably simulated the increase in FRN amplitude observed in response to increasingly improbable outcomes.

Most relevant to the current study, in response to a previous finding that participants in gambling tasks tend to alter their strategies more quickly in situations where outcome probabilities are more volatile (see Behrens, Woolrich, Walton, & Rushworth, 2007), the authors observed the PRO model during a simulation of a two-armed bandit task designed to manipulate environmental volatility. This task required participants to select from two different response options whose likelihood of reward vary throughout the course of the experiment. The rate at which reward probability fluctuates between the two options was manipulated to test the speed at which participants adapt their strategy. When the PRO model was run through this simulated task, it produced a larger error signal as reward probability shifted more quickly in response to its predictions being more frequently violated. The authors thus concluded that modelling the mPFC as an action-outcome predictor effectively simulates a broad range of phenomena from the behavioral and electrophysiological literature on feedback processing.

Previous Research from Our Lab

Paralleling the two-armed bandit task above, our lab has also sought to investigate feedback processing in an environment with a complex system of reward probabilities.

Accordingly, West et al. (2012) made electrophysiological recordings of 40 participants during a computerized version of the card game blackjack (see Figure 1) and compared ERPs for four possible outcomes (Win, Loss, Bust, and Tie). In blackjack, a dealer and at least one other player draw cards in an attempt to get the highest score without exceeding 21 points. Once the initial cards are dealt, the player can request to be dealt additional cards (“hit”) or remain with the cards that were initially dealt (“stay”). Thus, in playing against the computerized dealer, participants could either Win (i.e., draw cards summing to more points than the dealer without exceeding 21), Lose (i.e., draw cards summing to fewer points in the dealer while not exceeding 21), Bust (i.e., draw cards with points summing in excess of 21, which resulted in a loss independent of the cards held by the dealer), or Tie (i.e., draw cards summing to the same number of points as the dealer).




Figure 1. Screen shots taken from the feedback stage of the blackjack task used in both West et al., 2012 and West et al., 2014. The displays were designed to be as similar as possible across outcome conditions, as seen here with the Win (left) and Loss (right) feedback stimuli. The task was programmed to resemble a typical video blackjack game as closely as possible.

In their analysis of the FRN, West et al. (2012) made the unexpected observation that the amplitude of the FRN differed significantly between the two negative outcomes, with Busts ($M = 8.99 \mu\text{V}$, $SD = 5.48$) being reliably greater in amplitude than Losses ($M = 4.85 \mu\text{V}$, $SD = 3.51$)

(see Figure 2). Differences in ERPs for Losses and Busts had not been predicted, and a theoretical explanation for the discrepancy was not immediately obvious. From the standpoint of the existing literature, whether one subscribed to the view that the FRN indexes wins and losses (Hajcak et al., 2006) or failure to achieve a desired goal (Gentsch et al., 2009), there was no basis for distinguishing between different types of losses. As for the task itself, Losses and Busts seemed largely similar on several relevant dimensions: The feedback displays were the same across conditions (see Figure 1). Busts ($M = 43.59$, $SD = 23.74$) were more frequent than Losses ($M = 31.13$, $SD = 14.01$), so the observed amplitude difference between the two conditions did not fit neatly into the PRO model of Alexander and Brown (2011).

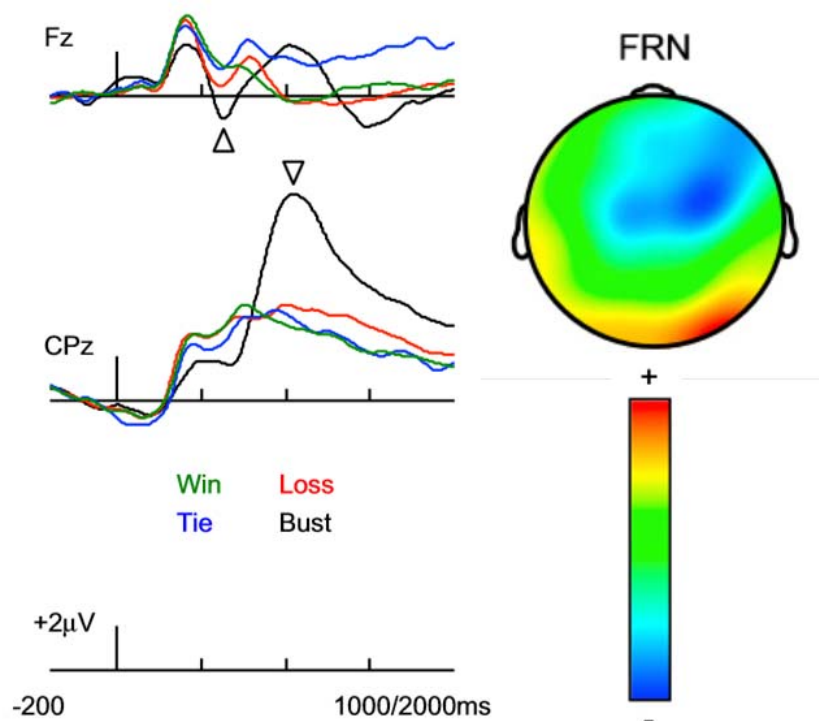


Figure 2. Grand-averaged ERPs recorded at two midline electrodes and topography map of participants receiving feedback in the blackjack task reported in West et al., 2012. Note the significant difference in amplitude of the FRN between the two negative outcomes: Loss and Bust.

Finally, the implications of each type of trial for participants seemed to be the same: both resulted in a loss of the amount wagered by the participant. The current study seeks, in part, to clarify the basis for the FRN's differential coding of negative outcomes.

Theoretical Justification

This thesis examines two hypotheses which might account for the differences in FRN amplitude observed between the different types of negative outcomes in West et al., 2012. First, if one subscribes to the PRO account of ACC activity (i.e., the FRN) as reflecting the processing of a system computing differences between actual and predicted outcomes, and if one assumes that busts and losses are sufficiently similar across other relevant dimensions (e.g., outcome frequency and visual properties of the feedback displays), it follows that the observed differences in FRN amplitude may arise principally from differences in expectancy. Within the context of blackjack, if FRN amplitude codes some quantitative difference between real and expected outcomes, then greater amplitude for busts implies a greater expectancy violation. Second, there is literature beginning to emerge suggesting that feedback processing may show effects of individual responsibility. Thus, in blackjack, the smaller FRN in response to losses may be a result of shared responsibility with the dealer for the outcome of the hand, whereas busts are entirely the result of the player's actions. The goal of this study was to examine the roles of expectancy and agency in feedback processing and, more specifically, their ability to modulate the amplitude of the FRN in a well-controlled binary selection task.

CHAPTER 2

EXPERIMENT 1

Hypothesis:

The amplitude of the FRN will be greater for unexpected than expected losses.

Introduction

This experiment was designed to test the effect of reward expectancy on the FRN while tightly controlling as many other experimental factors as possible (e.g., the complexity of the rules and the detail of the display). Therefore a simple, binary response task was used in which subjects were asked to guess which of two boxes on a computer screen contained a single gold token. In the present study, expectancy was examined by utilizing three different colors of boxes and varying the number of Win and Loss trials associated with each color. Of the three colors (orange, blue or purple), one was designated High Probability of winning (i.e., selecting either box resulted in a win 70% of the time), one was designated Low Probability of winning (i.e., selecting either box resulted in a win 30% of the time), and the third was designated Even Probability (i.e., selecting either box resulted in a win 50% of the time). If it is correct that the FRN is sensitive to violations of expectancy, then this effect should reflect a more negative FRN for unexpected outcomes (e.g., Losses in the High Probability condition) than expected outcomes (e.g., Losses in the Low Probability condition).

As noted previously, the PRO model of Alexander and Brown (2011) is concerned only with violations of expectation, and thus makes no distinction between wins and losses so long as their probability is perceived to be similar. However, since the balance of the literature

seems to suggest a disproportionate increase in FRN amplitude following losses compared to wins (Gehring & Willoughby, 2002; Miltner et al., 1997; Bellebaum et al., 2010) and given the ACC's reciprocal connections with the amygdala and other limbic structures (Devinsky et al., 1987), the prediction for the current study is that the FRN will be more negative for unexpected vs. expected outcomes, but this effect will be moderated by the outcome valence, with amplitude for unexpected losses being more negative than for unexpected wins.

Method

Participants

Thirty-six students enrolled in an undergraduate psychology course at Iowa State University completed the study. Eighteen of these were female and ages ranged between 18 and 28 ($M = 19.69$, $SD = 2.69$). The Edinburgh Handedness Inventory was used to assess handedness, and 29 participants were found to be right-handed, 4 were left-handed, and the remaining 3 were ambidextrous (Oldfield, 1971). All students received course credit for their participation in this study.

Materials and Design

Decision-Making Task. The trial sequence for Experiment 1 is shown in Figure 3. EEG was recorded for a single block of 300 trials, during which participants were asked to guess which of two boxes on a computer screen contained a winning token. As stated above, expectancy was manipulated using three different colors of boxes that were associated with different probabilities of winning. Both boxes that appeared on the screen were always the same color so that the color did not provide the participant with any basis for favoring one box over the other. The color assigned to each probability level was counterbalanced across subjects.

Previous studies have indicated that explicitly stating probabilistic rules to participants can be less effective than having subjects learn these rules for themselves (Holroyd & Coles, 2002; Holroyd et al., 2009, Bellebaum et al., 2010). Thus, prior to the onset of experimental trials, the present study included a learning phase of 60 trials in which the High Probability color resulted in a win on 20 of 20 trials (100% of the time) and the Low Probability color resulted in a loss on 20 of 20 trials (100% of the time), while the probability of winning on the Even Probability color was 50%. This allowed subjects to acquire the rules implicitly prior to the onset of experimental trials.

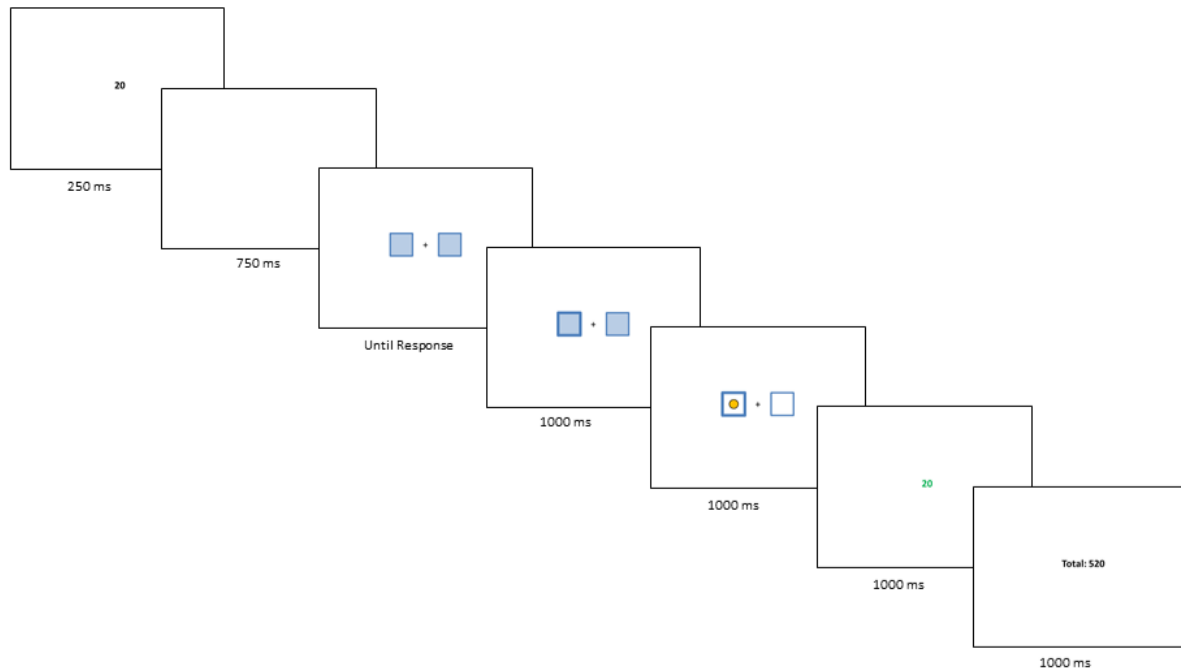


Figure 3. Trial sequence for Experiment 1.

Points in this study were assigned a monetary value (1 US cent per point). Participants started the task with a score of 500 points (5 dollars) and were told that they would be paid the value of their final score at the end of the task. Unbeknownst to participants, trial order and outcome were all determined prior to the beginning of the task, such that Win trials would

display a winning outcome irrespective of the box chosen, and likewise for Loss trials. Of the 300 experimental trials, 150 were Wins and 150 were losses. At the end of the experiment, subjects were paid the value of their final score, which was always \$5.

The effect of magnitude was examined by varying the number of points that could be won or lost on each trial. Low Magnitude trials were worth 5 points and High Magnitude trials were worth 20 points. These values were selected based on procedures from other studies of this type (Gehring & Willoughby, 2002; Mennes, Wouters, van den Bergh, Lagae, & Stiers, 2008; Bellebaum et al., 2010), so that both would be motivating to participants, yet still be perceived as differentially valuable. To make trial magnitude salient to participants, the value of the trial was displayed at both the beginning and end of each trial.

On each trial of the task, participants were first presented with a value at the center of the screen (either 5 or 20) displayed in black indicating the number of points that could be won or lost on that trial. The value remained on the screen for 1000 ms, after which participants saw a small fixation cue appear at the center of the screen with a small colored box displayed on either side of it. Both boxes were always the same color on a given trial, and the color could be either orange, blue, or purple depending on the trial. Participants were then asked to select which of the two boxes they believed contained a gold token. This selection was made using a standard computer keyboard, where the “v” key was used to select the box on the left and the “n” key was used to select the box on the right. Following participants’ selection, the selected box was highlighted on the screen for 1000 ms by increasing the width of its outline, after which both boxes were opened and the location of the gold token was displayed for 1000 ms. Next, a feedback screen appeared for 1000 ms which displayed the number of points won or

lost on the trial. When a correct selection was made, the point value of the trial was displayed at the center of the screen in green, while incorrect selections resulted in the point value of the trial being displayed in red. Finally, a screen appeared showing that participants' total score for the task in black numbers. A blank white screen then appeared for 1000 ms prior to the beginning of the next trial.

The trial list for this experiment was sorted into a quasi-random order to conceal the deterministic nature of the study. Because of the equal numbers Win vs Loss, High Magnitude vs Low Magnitude, and High Probability vs Low Probability trials throughout the experiment, all subjects received the same final score of 500 pts and were paid \$5 at the conclusion of the experiment. The entire experiment contained 360 trials (including the 60-trial learning phase) and lasted approximately 90 minutes. All stimuli were presented using E-Prime 2.0 Software (Psychology Software Tools, Pittsburgh, PA).

Procedure

Upon arriving at the lab, participants were given a description of the EEG electrode cap and its application to insure their full understanding of the procedure involved. Participants were told that they could leave the experiment at any time without forfeiture of credit if they were uncomfortable with the recording equipment or its application. If they agreed to participate in the experiment, they were then asked to read and sign the informed consent and to complete the modified Edinburgh Handedness Inventory (Oldfield, 1971). They were subsequently seated in a chair while the electrode cap and gel were applied, and moved to a smaller room with minimum illumination to complete the Decision-Making task. At the

conclusion of the experiment, all participants were debriefed, paid, and thanked for their participation.

Electrophysiological Recording and Analysis

The electroencephalogram (EEG, filter .02-150Hz, gain 1000, 16-bit A/D conversion) was recorded from an array of 68 tin electrodes sewn into an Electro-cap (Electro-Cap International, Eaton, OH) or affixed to the skin with an adhesive patch (ocular and mastoid electrodes). The Electro-cap was interfaced to a DBPA-1 (Sensorium Inc, Charlotte VT) that amplified and digitized the data. Vertical and horizontal eye movements were recorded from electrodes placed next to and below the right and left eyes. During recording all electrodes were referenced to electrode Cz, then re-referenced to an average reference for data analysis. A .1 to 12 Hz IIR bandpass filter was applied to the ERP data. Ocular artifacts were corrected using a covariance-based technique including empirically derived estimates of the EEG associated with artifact and artifact-free data (Source-Signal Imaging, San Diego). Feedback-locked ERPs were obtained offline and represented -200 to 1000 ms around the onset of the trial outcome display, including a -200-0 ms baseline. Measurements of mean voltage were taken at electrode FCz from 300-400 ms. Electrode FCz was chosen based on previous experience with feedback-processing tasks and inspection of the grand-averaged ERPs. Partial eta-squared (η_p^2) is reported as a measure of effect size for main effects and interactions (Bonett, 2008) and 95% confidence intervals are included for post hoc comparisons.

Results

The grand-averaged ERPs recorded at electrode FCz for the four conditions are portrayed in Figure 4. The analyses of mean voltage focused on how the amplitude of the FRN




Figure 4. Grand-averaged waveforms measured at electrode FCz for the four expectancy x outcome combinations in the Low (top) and High (center) magnitude conditions. The yellow bar below the x-axis on each graph denotes the epoch of interest (300-400ms post-feedback).

was affected by whether or not trial outcome conformed to the probability structure of the task, as well as effects of outcome valence and magnitude. To address the question of whether the FRN was sensitive to violations of expectancy, mean amplitude at electrode FCz was examined in a 2 (violation: violation or no violation) x 2 (valence: win or loss) x 2 (magnitude: 5-point or 20-point) ANOVA. The analysis revealed a main effect of valence, $F(1,34) = 9.63$, $p = .004$, $\eta_p^2 = .22$, with Losses eliciting a significantly more negative FRN ($M = -.13 \mu\text{V}$, 95% CI $[-.70, .44]$) than Wins ($M = .63 \mu\text{V}$, 95% CI $[-.17, 1.09]$). The main effects of violation, $F(1,34) = .22$, $p = .224$, and magnitude, $F(1,34) = 1.22$, $p = .277$, were not significant. There was a significant violation x magnitude interaction, $F(1,34) = 5.81$, $p = .021$, $\eta_p^2 = .14$. A comparison of violation conditions at each level of magnitude (Figure 5) showed a significant effect of violation for 5-

point trials, $F(1,34) = 5.18$, $p = .029$, $\eta_p^2 = .129$, with expectancy violations eliciting significantly more negative FRNs ($M = -.18 \mu\text{V}$, 95% CI $[-.74, .37]$) than trials which resulted in the expected outcome ($M = .41 \mu\text{V}$, 95% CI $[-.13, .95]$). In contrast, FRN amplitude did not differ significantly between violation conditions for 20-point trials, ($F(1,34) = 1.27$, $p = .268$). No other significant interactions were observed.




Figure 5. The effect of Violation vs. No Violation at each level of magnitude. Amplitude differences in the FRN are only significant for 5-point trials.

Because the interaction between magnitude and violation conditions was unexpected and lacked a clear explanation, behavioral data was also examined for differences between Low and High magnitude trials (see Figure 6). While this analysis was exploratory in nature, finding a behavioral measure that varied between magnitude conditions would provide converging evidence that there was some difference in the way that participants were processing low- and high-point trials, and that this difference might be influencing the amplitude of the FRN in the two magnitude conditions. Initially, reaction times (RT) were compared between magnitude

conditions by collapsing across valence and violation conditions (note that valence conditions are meaningless in the context of reaction time, since information about outcome is unavailable at the time of the participants' selection) but the analysis failed to show a significant difference between 5- and 20-point trials, $F(1,34) = .33, p = .571$). Since the only other information available at the time of the participants' decision was the outcome probability indicated by the color of the boxes, a 2 (magnitude: 5-point or 20-point) x 2 (probability: expect win or expect loss) ANOVA was performed. This analysis revealed a significant magnitude x probability interaction, $F(1,34) = 6.68, p = .014, \eta_p^2 = .16$. Subsequent comparisons showed an effect of probability on 20-point trials approaching significance, $F(1,34) = 3.37, p = .075, \eta_p^2 = .09$, with trials where wins were expected ($M = 588$ ms, 95% CI [479, 698]) producing faster RTs than trials where losses were expected ($M = 654$ ms, 95% CI [546, 761]). There were no significant differences in RT observed between probability conditions in 5-point trials, $F = 1.194, p = .28$.




Figure 6. Reaction time data from Experiment 1. Left: Comparison of RT between all 5- and 20-point trials. Right: Comparisons of RT between 5- and 20-point trials where participants either expected to win (blue) or lose (red).

As mentioned previously, the PRO model posits a pure expectancy account of the FRN that would not predict any effect of valence across equivalent expectancy conditions in the present task. However, based on previous findings from our lab and others, the current prediction was that Losses should produce more negative FRNs across expectancy conditions.

Accordingly, mean amplitude at electrode FCz for was compared between all conditions in which an expectancy violation occurred, resulting in a 2 (valence: win or loss) x 2 (magnitude: 5-point vs. 20-point) ANOVA (see Figure 7). In this analysis, the main effect of valence was significant, $F(1,34) = 7.14, p = .011, \eta_p^2 = .17$, with unexpected losses producing a more negative FRN ($M = -0.29 \mu\text{V}$, 95% CI [-1.00, .42]) than unexpected wins ($M = 0.68 \mu\text{V}$, 95% CI [.17, 1.19]). Interestingly, the main effect of magnitude was significant in this analysis, $F(1,34) = 5.70, p = .023, \eta_p^2 = .14$, with low-magnitude expectancy violations producing a more negative FRN ($M = -0.19 \mu\text{V}$, 95% CI [-.74, .37]) than high-magnitude violations ($M = 0.58 \mu\text{V}$, 95% CI [-.05, 1.20]). No significant valence x magnitude interaction was observed ($F(1,34) = 2.943, p = .095$).




Figure 7. Mean amplitude of the FRN for conditions in which an expectancy violation occurred for Experiment 1. Analysis revealed main effects of both valence and magnitude. The error bars represent the standard error of the mean.

Discussion

The finding that unexpected outcomes produced reliably more negative FRNs in the low-magnitude condition of this experiment provides support for an effect of expectancy in the neural activity that underlies this component. This remains true despite the failure to find a similar effect in the high-magnitude condition, although further research is needed to explain

the discrepancy. The fact that an interaction was also observed in the RT data between outcome magnitude and reward probability conditions suggests that processing may change as the potential value of trials increases, and that this may occur independently of trial feedback, though nothing in the present data justifies speculation on what these processing differences might be.

The additional finding of a main effect of valence in Experiment 1 fits well with the corpus of research on the FRN, going all the way back to Gehring and Willoughby (2002). That Losses produced more negative FRNs than Wins across expectancy conditions suggests that the response profile of the FRN is more nuanced than either a pure expectancy or pure gain-loss account allows. In addition, the fact that unexpected losses produce a more negative FRN than unexpected wins does not fit tidily into the PRO model of ACC function proposed by Alexander and Brown (2011). Since the error signal in the PRO model corresponds to prediction-related activity that is not suppressed because the predicted outcome fails to occur, the implication of a more negative FRN for unexpected losses seems to be that there is initially more prediction-related activity that requires suppressing when wins are anticipated than when losses are anticipated. Given the fairly consistent finding that losses tend to produce more negative FRNs than gains across tasks and conditions, it may be more parsimonious to infer that the loss outcome itself is what causes the increase in FRN amplitude, as opposed to the antecedent prediction-related activity. However, given that the analysis for Experiment 1 did reveal a significantly more negative FRN for unexpected outcomes in the 5-point condition does suggest that this component is sensitive to expectancy violations under some circumstances. Further research will be required to elucidate this effect.

CHAPTER 3

EXPERIMENT 2

Hypothesis:

The amplitude of the FRN will be smaller for losses when the outcome can be attributed to the action of an opponent.

Introduction

Despite the functional similarities between busts and losses in the blackjack task (e.g., same display, similar outcome frequencies, same amount lost) there may be reason to question whether an optimally designed feedback processing system would give equal weight to these two results. Specifically, it may be reasonable to suppose that such a system would give greater consideration to outcomes over which the agent has more control, since adaptive processing could more directly inform similar actions in the future. This scheme maps nicely onto the bust vs loss distinction: In blackjack, if a player loses because her hand contained fewer points than the dealer's, then part of the responsibility for the loss is mitigated by the dealer's fortunes: the player did the best she could, but in the end the dealer did better. With a bust, however, the dealer cannot be construed as affecting the outcome: the player has beaten herself. Additionally, because there is no other agent in this latter scenario to serve as an imponderable variable, information about this outcome may provide more utility to a learning system concerned with predicting and evaluating outcomes.

A modulatory effect of agency on the amplitude of the FRN does have precedent in the neurophysiological literature. For instance, researchers have demonstrated that the amplitude

of the FRN is greater during tasks requiring an active response than during similar tasks requiring only passive viewing. In one study, Yeung, Holroyd, and Cohen, (2004) had subjects perform two gambling tasks which were identical except that in one task (the *Choice task*) subjects were instructed to press a button to indicate which of four balloons contained a prize, while in the other (the *No-choice task*) subjects were told that pressing a key would simply direct the computer to randomly choose one of the four balloons on the player's behalf (Yeung et al., 2004). Under these conditions, the researchers observed an FRN in response to losses in both the *Choice* and *No-choice* tasks, but reported that FRN amplitude was greater for the *Choice* task (see also Zhou, Yu, & Zhou, 2010).

In another study, Li, Han, Lei, Holroyd, & Li (2011) reported that the FRN can be attenuated in cases where perceived responsibility for an outcome is reduced. The researchers manipulated subjects' sense of control in a gambling task by manipulating whether or not participants believed there was a pattern of correct responses that could be learned (Li et al., 2011). On each trial, subjects were presented with four colored balls on a computer screen, two of which were said to contain a winning coin worth .5 Yuan and two of which were said to contain a "thief" which would deduct .5 Yuan from the participants' scores. Critically, while the relative position of the colored balls and coins were random in both conditions, in the *High-responsibility* condition, participants were told that there was a rule that would allow them to reliably locate the coins according to the color of the balls, while subjects in the *Low-responsibility* condition received no such information. When subjects believed that the location of the coins was non-random and predictable, the researchers observed significantly more

negative FRN in response to loss trials, a result interpreted as showing that subjects' perceived sense of responsibility for the outcome of each trial was being coded by their FRN.

In another study that seems to speak most directly to the role of agency in a gambling task, Li et al. (2010) instructed groups of 3 participants to play a dice-rolling game in which the objective was to roll three virtual die to accumulate a score in excess of 10 points. In the *Self-execution* condition, three dice at a time were rolled by a single participant, whereas in the *Cooperating* condition each of the three participants rolled a single dice and their score was cumulative. Importantly, each subject gained or lost the same amount of money per trial (.5 Yuan) regardless of whether there were one or three players, so the potential value of a trial was the same across conditions. The researchers found that the FRN was smaller in the Cooperating condition than in the Self-execution condition, leading the authors to conclude that participants in the Cooperating condition were discounting their losses in response to their shared responsibility. The combined evidence from these three studies supports the idea that a smaller FRN observed for loss trials in the blackjack task could be the result of a sense of mitigated responsibility brought on by the dealer's role in the outcome.

In Experiment 2, agency was manipulated by simulating gameplay between the participant and the computer. This was accomplished by using a modified version of the binary selection task from Experiment 1 with two different agency conditions: one in which the outcome was the result of the participant's selection, and another in which the outcome was the result of a virtual opponent's selection. Given the available evidence suggesting an effect of agency on the FRN, I predicted that FRN amplitude following a loss resulting from the

computer's selection would be smaller relative to a loss resulting from the participant's selection.

Method

Participants

Twenty-six students enrolled in an undergraduate psychology course at Iowa State University completed the study. Seventeen of these were female and ages ranged between 18 and 27 years ($M = 20.04$, $SD = 2.46$). The Edinburgh Handedness Inventory (Oldfield, 1971) was used to assess handedness, and 25 participants were found to be right-handed, while the remaining participant was ambidextrous. All students received \$5 and course credit for their participation in this study.

Materials and Design

Decision-Making Task

The trial sequence for Experiment 2 is shown in Figure 8. As mentioned previously, this version of the binary selection task utilized two different agency conditions. Thus, in Experiment 2, EEG was recorded for 144 trials, each of which consisted of a Player Select (PS) condition in which the participant made a selection and a Computer Select (CS) condition where the computer made a selection. For the PS condition, each trial began with the prompt "Your Turn!" presented at the center of the screen for 250 ms. This was followed by a 750 ms baselining period, after which two colored boxes appeared on the screen. As in the first experiment, the boxes remained on the screen until the participant chose the left or the right box, again using the "v" key to indicate left and "n" key to indicate right. Once the participant made a selection, the border around the box thickened to highlight his/her selection, after




Figure 8. Trial sequence for Experiment 2.

which the token's true location was revealed. By contrast, the CS condition in each trial began with the prompt "Computer's Turn!" presented at the center of the screen, followed by the same baselining period and two boxes appearing at the center of the screen. The computer's choice was simulated by leaving the boxes on the screen for 500 ms (to create the appearance of deliberation), after which the border around the computer's "selected" box was highlighted, and the actual location of the token appeared on the screen, along with the participant's score for that trial in place of the fixation cue. If the outcome favored the participant, the score appeared in green to indicate a gain, and an outcome favoring the computer resulted in the score appearing in red, indicating a loss. Both PS and CS conditions in each trial concluded with the player's total score appearing on the screen for 1000 ms. Because each trial contained one PS and one CS condition, selection alternated between the participant and the computer throughout the game, and whether the trials began with the PS or CS condition was

counterbalanced across participants. As in Experiment 1, the outcome of each trial was predetermined to allow the trial outcomes to be controlled, ensuring an equal number of trials in each cell.

Several other modifications of the task from Experiment 1 bear mentioning. First, given the theoretical questions concerning processing differences for 20-point trials in Experiment 1, and given the relatively small effect sizes in the first experiment, all trials in Experiment 2 were worth 5 points. Also, the screenshots indicating the token's location and the participant's points for the trial were combined in Experiment 2 such that all of the relevant information about trial outcome was displayed on a single screen. This change was implemented to minimize any tendency by participants to attend to different aspects of trial feedback on different displays, which could make their data difficult to interpret. Finally, trial order in Experiment 2 was generated by allowing E-Prime to randomly select trials from a list so that participants did not all have their trials presented in the same order.

As in Experiment 1, participants were again paid out the value of their final score (\$.01/point) at the end of the study. Participants began the experiment with 500 points and, because there were equal numbers of wins and losses, all participants ended with a score of 500 (\$5 US).

Procedure

The procedure for Experiment 2 was the same as the first experiment.

Electrophysiological Recording and Analysis

Data collection and analysis was performed using the same methods as in Experiment 1. Based on previous experience with feedback-processing tasks and inspection of the grand-

averaged ERPs, measurements of mean voltage were taken at electrode FCz from 230-250 ms and analyzed in a 2 (agency: computer select, player select) x 2 (valence of outcome: win, loss) ANOVA. Partial eta-squared (η_p^2) is reported as a measure of effect size for main effects and interactions (Bonett, 2008) and 95% confidence intervals included for post hoc comparisons.

Results

The grand-averaged ERPs recorded at electrode FCz for the two agency conditions are portrayed in Figure 9. The analyses of mean voltage focus on how the amplitude of the FRN was affected by whether or not the trial outcome was the result of the participant's decision or that of the computer, as well as outcome valence.




Figure 9. Grand-averaged waveforms measured at electrode FCz for Wins and Losses in the Player Select (top) and Computer Select (center) conditions. The yellow bar below the x-axis on each graph denotes the epoch of interest (230-250 ms post-feedback).

To address the question of whether the FRN was sensitive to the agent responsible for an outcome, mean amplitude at electrode FCz was examined in a 2 (agency: computer select,

player select) x 2 (valence: win, loss) ANOVA. The results revealed a main effect of valence, $F(1,25) = 9.27, p = .005, \eta_p^2 = .27$, with Losses eliciting a more negative FRN ($M = 1.83 \mu\text{V}$, 95% CI [.55, 3.11]) than Wins ($M = 2.89 \mu\text{V}$, 95% CI [1.43, 4.36]). While no main effect of agency was observed, $F(1,25) = .97, p = .334$, the agency x valence interaction was significant, $F(1,25) = 7.58, p = .011, \eta_p^2 = .23$. A comparison of valence conditions for each agency condition showed a more negative FRN for Losses than Wins in the Player Select condition, $F(1,25) = 15.43, p = .001, \eta_p^2 = .38$, but no difference in amplitude for the Computer Select condition, $F(1,25) = .00, p = .985$.

While the previous analysis provides support for the hypothesis that the FRN differentially codes valence of outcomes primarily in the Player Select condition, it leaves open the possibility that activity relevant to valence in the Computer Select condition may occur elsewhere in the brain. Visual inspection of the ERP waveforms revealed an observable difference between Wins and Losses in the Computer Select condition at electrode CPz, caudal to FCz but still on the midline. Based on inspection of the grand average waveforms, a larger epoch from 300-600 ms post-feedback was selected for analysis (see Figure 10 for grand-averaged ERPs) and the same 2 x 2 ANOVA was performed. This analysis revealed a main effect of agency, $F(1,25) = 16.15, p < .001, \eta_p^2 = .39$, with Player Select trials eliciting a significantly more positive mean amplitude ($M = 4.82 \mu\text{V}$, 95% CI [3.23, 6.41]) than Computer Select trials ($M = 2.17 \mu\text{V}$, 95% CI [1.34, 2.99]). No main effect of valence observed, $F(1,25) = 1.46, p = .238$, and the interaction between agency and valence approached significance, $F(1,25) = 3.99, p = .057$. A comparison between mean amplitude for Wins and Losses in the Computer Select condition revealed a significant effect of valence, $F(1,25) = 4.36, p = .047, \eta_p^2 = .15$, while no

similar valence effect was found for Player Select conditions, $F(1,25) = .20, p = .657$. This finding is critical for showing that while the FRN may not differentially code valence when the computer determines trial outcome, the distinction between Wins and Losses resulting from another agent's actions may still be coded elsewhere in the brain. To the best of my knowledge this is a novel finding and should be further investigated in future studies.

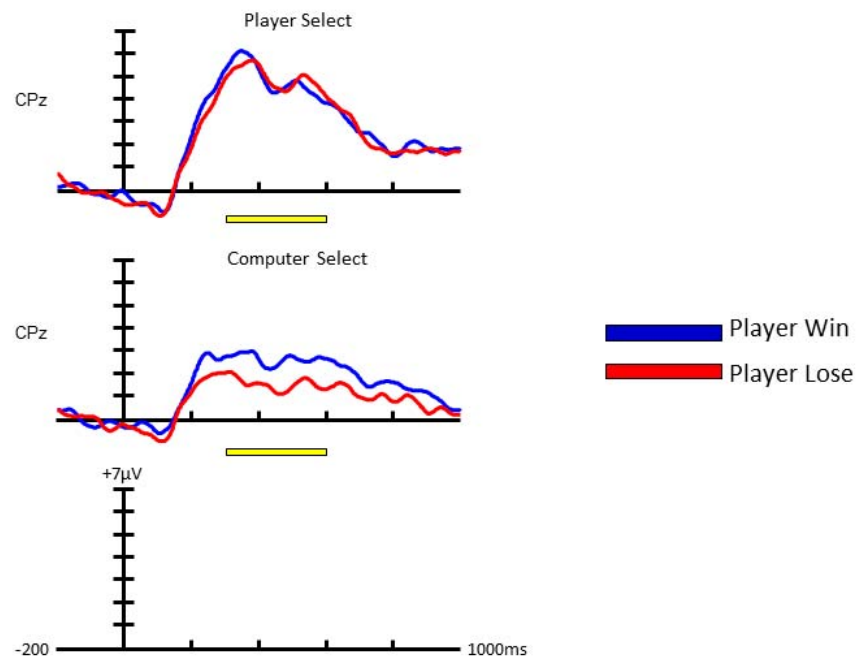


Figure 10. Grand-averaged waveforms measured at electrode CPz for Wins and Losses in the Player Select (top) and Computer Select (center) conditions. The yellow bar below the x-axis on each graph denotes the epoch of interest (300-600 ms post-feedback). Note that at this more posterior location, the previously observed effect reverses, with the valence conditions differentially coded in the Computer Select condition, but not the Player Select condition.

Discussion

The results of the Experiment 2 provide support for an effect of agency on the FRN. This experiment was designed so that the only difference between agency conditions was whether or not the outcome was the result of the player's action or the action of the computer, yet this

was enough to eliminate the normally observed differences in the FRN between Wins and Losses in the Computer Select condition.

One potential objection to this interpretation of the data is that the decrease in the amplitude of the FRN in the Computer Select condition may actually be due to inattention. Based simply on the data from electrode FCz, it is difficult to rule out the possibility that the lack of an FRN for Losses in the Computer Select condition might simply be due to the participants failing to focus on the task during the computer's turn. However, the fact that differences between valence conditions are observed in the Computer Select condition at a more posterior location on the scalp provides some evidence that participants are attending to the computer's selection and coding the result. What sort of neural process this activity at CPz represents presents an interesting avenue for future study.

Finally, it should be noted that while an effect of agency on the FRN was supported by the data of Experiment 2, the main effect of valence observed in the FRN data replicates the findings of numerous other studies. This data lends further support to an account of the FRN as an index of a neural system primarily concerned with negative outcomes.

CHAPTER 4

GENERAL DISCUSSION

The goal of the current study was to examine the potential for two aspects of decision feedback to influence the size of the FRN: the participants' subjective expectations of winning or losing in a task, and the participants' perceived role in task outcomes. Violations of expectancy were generated by teaching participants to associate different stimulus colors with winning and losing outcomes and creating trials that did not conform to those associations. Such violations were found result in an increase in the amplitude of the FRN for losses, but only for trials with smaller point values. Different agency conditions were generated by using a selection task where trial outcomes could be the result of either the player's decision or that of a computerized opponent. This manipulation was found to result in a decrease in the amplitude of the FRN for losses in cases where losses the trial result was beyond the participant's control. Together, these findings paint a picture of the FRN that is somewhat more nuanced than is typically presented in the feedback processing literature, and provide a strong counterargument to the purely win/loss account advocated by some researchers (see Hajcak et al., 2006). While these two experiments replicate previous findings that outcome valance has a strong effect on FRN amplitude (Gehring & Willoughby, 2002; Walsh & Anderson, 2012), this component also appears to be influenced by the perceived likelihood of an outcome, as well as an individual's perceived responsibility for an outcome.

The observed increases in the amplitude of the FRN for unexpected losses support the idea that participants' expectations about event outcomes may influence feedback processing

in the very early stages; in this case, only 300-400 ms after feedback is presented. The fact that outcome valence seems to have a much stronger influence on FRN amplitude than expectancy means that this result is not a total vindication of the PRO model of mPFC function (Alexander & Brown, 2011), since this model does not predict any effects of the relative goodness or badness of an outcome. That said, the fact that there may be any influences of participants' outcome expectations on early feedback processing make clear that the FRN is more complex than the simple win/loss signal originally imagined by Gehring and Willoughby (2002).

The finding that the FRN was sensitive to expectancy only in low-magnitude trials does not lend itself to an obvious explanation. This is particularly true since the few studies which have demonstrated effects of reward magnitude on the FRN have shown the opposite trend, such as Bellebaum et al. (2010) who observed larger FRNs for losses on 20¢ and 50¢ trials relative to 5¢ trials. One way to further elucidate the role of magnitude might be to use a broader array of point values and attempt to observe any trends in FRN amplitude across finer point gradations. It is possible, for instance, that adding 10-point trials to the procedure for Experiment 1 might reveal an increase in FRN amplitude for 10-point losses relative to 5-point losses, but that 20-point losses are sufficiently extreme to recruit different response processes in the ACC. While the finding that expectancy violations do modulate FRN amplitude for 5-point trials should not be trivialized, further research is needed to discover why this effect does not hold across all values of reward magnitude.

The finding from Experiment 2 that the effect of outcome valence on the FRN is mitigated in instances when the losses result from the actions of another player provides still further evidence that ostensibly higher-order cognitive processes may influence feedback

processing at a relatively early stage. It would be instructive to conduct additional studies to see if the observed effect of agency on the FRN is simply the result of mitigated responsibility for an outcome or whether the effect might vary depending on the nature of the other decision-maker. For instance, Sanfey, Rilling, Aronson, Nystrom, & Cohen (2003) found that participants playing an ultimatum game would behave differently and produce a distinctly different Blood-Oxygen-Level Dependent (BOLD) signal in an fMRI scanner depending on whether they were told that the ultimatum was offered by a computer or another player. Specifically, the researchers observed that when participants were told that the ultimatum came from another person, they would reject more equitable offers and show more activity in the anterior insula and dorsolateral prefrontal cortex than when they were told the offer came from the computer, though both conditions showed activation of the ACC. Experiment 2 in the current study could thus be altered to employ an additional agency condition where participants were told that it was in fact another player making the selection and not the computer. If the attenuation of the FRN in the Computer Select condition was simply due to lack of responsibility for the outcome, there should be no difference in the data observed for a computer opponent and human opponent. This might also be a way to gain further insight into the nature of the more posterior component observed in Experiment 2 that coded differences between Wins and Losses in the Computer Select conditions.

To my knowledge, the identification of an ERP that is sensitive to the valence of outcomes resulting from the actions of a non-self agent is a novel finding of this study. Other researchers have reported modulations of the P3b, a positive-going component with a similar topography that is observed when subjects are presented with unlikely task-relevant stimuli, in

response to outcome valence in gambling tasks (Hajcak et al., 2007; Li et al., 2010), and the results of at least one study have been used as support for the idea that this component has some sensitivity to social context (i.e., larger when participants win less than another player) (Wu, Zhang, Elieson, & Zhou, 2012). However, if the component observed in Experiment 2 is the P3b, the observation that it indexes outcome valence preferentially in situations where the outcome is the result of another player's actions would itself be an important finding, and one which has not been reported elsewhere. Further studies should be conducted investigate the nature of the observed component, including further tests of its response properties and a source analysis of to identify the location in the brain where it is generated.

This study was designed in part as a tool for better understanding the observation that busts in the game of blackjack produce more negative FRNs than losses (West et al., 2012; West et al., 2014). It was postulated that this effect might be due either to busts being perceived as a less likely outcome by participants, or the fact that responsibility for losses can be partially attributed to the actions of the dealer while busts cannot. In the end, the results of this study provide support for both hypotheses, and indeed the observed effects in the blackjack studies may represent some combination of both expectancy and agency effects. However, care should be taken in generalizing the results of Experiment 1 until the interaction between expectancy violation and reward magnitude can be better explained, since a similar interaction was not observed in participants playing blackjack. Future research should focus on using similar manipulations to those used in Experiments 1 and 2 in a scaled-down version of the blackjack task to see if analogous results are found using a well-controlled task with a more complex probability and outcome structure.

REFERENCES

- Alexander, W.H., & Brown, J.W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience*, *14*(10), 1338-1344. doi:10.1038/nn.2921
- Allman, J.M., Hakeem, A., Erwin, J.M., Nimcinsky, E., & Hof, P. (2006). The anterior cingulate cortex: The evolution of an interface between emotion and cognition. *Annals of the New York Academy of Sciences*, *935*, 107-117. doi:10.1111/j.1749-6632.2001.tb03476.x
- Baron, J., & Ritov, I. (1994). Reference points and omission bias. *Organizational Behavior and Human Decision Processes*, *59*(3), 475-498. doi:10.1006/obhd.1994.1070
- Baron, J., & Ritov, I. (2004). Omission bias, individual differences, and normality. *Organizational Behavior and the Human Decision Process*, *94*, 74-85. doi:10.1016/j.obhdp.2004.03.003
- Behrens, T.E., Woolrich, M.W., Walton, M.E., & Rushworth, M.F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*, 1214-1221. doi:10.1038/nn1954
- Bellebaum, C. & Daum, I. (2008). Learning-related changes in reward expectancy are reflected in the feedback-related negativity. *European Journal of Neuroscience*, *27*(7), 1823-1835. doi: 10.1111/j.1460-9568.2008.06138.x
- Bellebaum, C., Pomezzi, D., & Daum, I. (2010). It is less than your expected: The feedback-related negativity reflects violations of reward magnitude expectations. *Neuropsychologia*, *48*(11), 3343-3350. doi: 10.1016/j.neuropsychologia.2010.07.023
- Botvinick, M.M., Nystrom., L., Fissel, K., & Carter, C.S. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, *402*(8), 179-181. doi:10.1038/46035

- Brown, J.W., & Braver, T.S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science*, *307*(5712), 1118-1121. doi:10.1126/science.1105783
- Bush, G., Luu, P., & Posner, M.I. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences*, *4*(6), 215-221. doi:10.1016/S1364-6613(00)01483-2
- Carlson, J.M., Foti, D., Mujica-Parodi, L.R., Harmon-Jones, E., & Hajcak, G. (2011). Ventral striatal and medial prefrontal BOLD activation is correlated with reward-related electrocortical activity: A combined ERP and fMRI study. *NeuroImage*, *57*(4), 1608-1616. doi:10.1016/j.neuroimage.2011.05.037
- Carter, C.S., Botvinick, M.M., & Cohen, J.D. (1999). The contribution of anterior cingulate cortex to executive processes in cognition. *Reviews in the Neurosciences*, *10*, 49-57. doi:10.1515/REVNEURO.1999.10.1.49
- Devinsky, O., Morrell, M.J., & Vogt, B.A. (1995). Contributions of anterior cingulate cortex to behavior. *Brain*, *118*(1), 279-306. doi:10.1093/brain/118.1.279
- Gehring, W.J., & Willoughby, A.R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science*, *295*(5563), 2279-2282. doi:10.1126/science.1066893
- Gentsch, A., Ullsperger, P., & Ullsperger, M. (2009). Dissociable medial frontal negativities from a common monitoring system for self- and externally caused failure of goal achievement. *NeuroImage*, *47*(4), 2023-2030. doi:10.1016/j.neuroimage.2009.05.064
- Hajcak, G., Holroyd, C.B., Moser, J.S., & Simons, R.F. (2005). Brain potentials associated with expected and unexpected good and bad outcomes. *Psychophysiology*, *42*(2), 161-170. doi:10.1111/j.1469-8986.2005.00278.x

- Hajcak G., Moser, J.S., Holroyd, C.B., & Simons, R.F. (2006). The feedback-related negativity reflects the binary evaluation of good versus bad outcomes. *Biological Psychology*, 71(2), 148-154. doi:10.1016/j.biopsycho.2005.04.001
- Hajcak, G., Moser, J.S., Holroyd, C.B., & Simons, R.F. (2007). It's worse than you thought: The feedback negativity and violations of reward prediction in gambling tasks. *Psychophysiology*, 44(6), 905-912. doi:10.1111/j.1469-8986.2007.00567.x
- 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(4), 679-709. doi:10.1037/0033-295X.109.4.679
- Holroyd, C.B., Larsen, J.T., & Cohen, J.D. (2004). Context dependence of the event-related brain potential associated with reward and punishment. *Psychophysiology*, 41(2), 245-253. doi:10.1111/j.1469-8986.2004.00152.x
- Holroyd, C.B., Pakzad-Vaezi, K.L., & Krigolson, O.E. (2008). The feedback correct-related positivity: Sensitivity of the event-related brain potential to unexpected positive feedback. *Psychophysiology*, 45(5), 688-697. doi:10.1111/j.1469-8986.2008.00668.x
- Holroyd, C.B., Krigolson, O.E., Baker, R., Lee, S., & Gibson, J. (2009). When is an error not a prediction error? An electrophysiological investigation. *Cognitive, Affective, and Behavioral Neuroscience*, 9(1), 59-70. doi:0.3758/CABN.9.1.59
- Ito, S., Stuphorn, V., Brown, J.W., & Schall, J.D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science*, 302(5642), 120-122. doi:10.1126/science.1087847
- Janer, K.W., & Pardo, J.V. (1991). Deficits in selective attention following bilateral anterior

cingulotomy. *Journal of Cognitive Neuroscience*, 3(3), 231-241.

doi:10.1162/jocn.1991.3.3.231

Kahneman, D., & Miller, D.T. (1986). Norm theory: Comparing reality to its alternatives.

Psychological Review, 93(2), 136-153. doi:10.1037/0033-295X.93.2.136

Li, P., Jia, S., Feng, T., Liu, Q., Suo, T. & Li, H. (2010). The influence of the diffusion of responsibility effect on outcome evaluation: Electrophysiological evidence from an ERP study. *NeuroImage*, 52(4), 1727-1733. doi:10.1016/j.neuroimage.2010.04.275

Li, P., Han, C., Lei, Y., Holroyd, C.B., Li, H. (2011). Responsibility modulates neural mechanisms of outcome processing: An ERP study. *Psychophysiology*, 48(8), 1129-1133.

doi:10.1111/j.1469-8986.2011.01182.x

MacDonald, A.W., Cohen, J.D., Stenger, V.A., & Carter, C.S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science*, 288(5472), 1835-1838. doi:10.1126/science.288.5472.1835

Mennes, M., Wouters, H., van den Bergh, B., Lagae, L., & Stiers, P. (2008). ERP correlates of complex human decision making in a gambling paradigm: Detection and resolution of conflict. *Psychophysiology*, 45(5), 714-720. doi:10.1111/j.1469-8986.2008.00678.x

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(6), 449-451.

doi:10.1162/jocn.1997.9.6.788

Potts, G.F. (2004). An ERP index of task relevance evaluation of visual stimuli. *Brain and Cognition*, 56(1), 5-13. doi:10.1016/j.bandc.2004.03.006

- Potts, G.F., Martin, L.E., Burton, P., & Montague, P.R. (2006). When things are better or worse than expected: The medial frontal cortex and the allocation of processing resources. *Journal of Cognitive Neuroscience, 18(7)*, 1112-1119. doi:10.1162/jocn.2006.18.7.1112
- Sanfey, A.G., Rilling, J.K., Aronson, J.A., Nystrom, L.E., & Cohen, J.D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science, 300(5626)*, 1755-1758. doi:10.1126/science.1082976
- Walsh, M.M., & Anderson, J.R. (2012). Learning from experience: Event-related potential correlates of reward processing, neural adaptation, and behavioral choice. *Neuroscience and Biobehavioral Reviews, 36(8)*, 1870-1884. doi:10.1016/j.neubiorev.2012.05.008
- West, R., Bailey, K., Tiernan, B.N., Boonsuk, W., & Gilbert, S. (2012). The temporal dynamics of medial and lateral frontal activity related to proactive cognitive control. *Neuropsychologia, 50(14)*, 3450-3460. doi:10.1016/j.neuropsychologia.2012.10.011
- West, R., Bailey, K., Anderson, S., & Kieffaber, P.D. (2014). Beyond the FN: A spatio-temporal analysis of the neural correlates of feedback processing in a virtual blackjack game. *Brain and Cognition, 86*, 104-115. doi:10.1016/j.bandc.2014.02.003
- Wu, Y., Zhang, D., Elieson, B., & Zhou, X. (2012). Brain potentials in outcome evaluation: When social comparison takes effect. *International Journal of Psychophysiology, 85(2)*, 145-152. doi:10.1016/j.ijpsycho.2012.06.004
- Yeung, N., Holroyd, C.B., & Cohen, J.D. (2004). ERP correlates of feedback and reward processing in the presence and absence of response choice. *Cerebral Cortex, 15(5)*, 535-544. doi:10.1093/cercor/bhh153
- Zhou, Z., Yu, R., & Zhou, X. (2010). To do or not to do? Action enlarges the FRN and P300 effects

on outcome evaluation. *Neuropsychologia*, 48(12), 3606-3613.

doi:10.1016/j.neuropsychologia.2010.08.010