ABSTRACT

Title of Thesis: EFFECTS OF REWARD CONTEXT ON FEEDBACK PROCESSING AS INDEXED BY TIME-FREQUENCY ANALYSIS

Adreanna T. Massey, Master of Science, 2017

Thesis Directed By: Dr. Edward M. Bernat, Assistant Professor Department of Psychology

The role of reward context has been investigated as an important factor in feedback processing. Previous work has demonstrated that the amplitude of the feedback negativity (FN) depends on the value of the outcome relative to the range of possible outcomes in a given context, not the objective value of the outcome. However, some research has shown that the FN does not scale with loss magnitude in loss-only contexts, suggesting that some contexts do not show a pattern of context-dependence. Time-frequency decomposition techniques have proven useful for isolating important activity, and have shown that time-domain ERPs can be better represented as separable processes in delta (0-3 Hz) and theta (3-7 Hz). Thus, the current study seeks to assess whether the role of context in feedback processing is better elucidated using time-frequency analysis. Results revealed that theta was more context-dependent and showed a binary response to best-worst differences in the gain and even contexts. Delta was more context-independent: the best outcomes scaled linearly with reward magnitude and best-worst differences scaled with context valence. Our findings reveal that theta and delta are differentially sensitive to context and that context valence may play a critical role in determining how the brain processes good and bad outcomes.

EFFECTS OF REWARD CONTEXT ON FEEDBACK PROCESSING AS INDEXED BY TIME-FREQUENCY ANALYSIS

by

Adreanna T. Massey

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Master of Science
2017

Advisory Committee:
Dr. Edward Bernat
Dr. Andres De Los Reyes
Dr. Tracy Riggins

Acknowledgements

My deepest gratitude to my mentor, Dr. Edward Bernat, for his guidance and support both on this project and in my research training. I would also like to thank my Master's Thesis committee: Dr. Bernat, Dr. De Los Reyes, and Dr. Riggins. Your feedback has been instrumental both for this project and for my future development as a researcher.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

Reward processing and performance monitoring have been widely studied as important factors underlying cognitive and affective processes. Reward and performance monitoring systems are necessary for the adaptation of behavior in the pursuit of goals. Furthermore, these systems have been implicated in various forms of psychopathology, such as depression, anxiety, substance abuse, and behavioral additions. While neuroimaging research has provided some evidence for the neural circuitry of reward processing and performance monitoring and abnormalities of these circuitries in psychopathology, little is known about how the system that determines whether an event is good or bad is influenced by context. For example, is the same outcome processed similarly or differently across varying contexts? Does the value of $100 differ in the Unites States versus a developing country? Most would agree that the value of rewards and losses are dependent on the context in which these outcomes occur, but little is known about the neural mechanisms underlying processing differences across contexts.

The role of reward context has been investigated as an important factor in feedback processing (Holroyd, Larsen, & Cohen, 2004; Kujawa, Smith, Luhmann, & Hajcak, 2013; Nieuwenhuis et al., 2005). Previous work has demonstrated that the amplitude of the feedback negativity (FN), a negative-going event-related potential (ERP) peaking around 250 ms, depends on the value of the outcome relative to the range of possible outcomes in a given context, not the objective value of the outcome (Holroyd et al., 2004). However, some research has shown that the FN does not scale with loss magnitude in loss-only contexts, suggesting that some contexts do not show a pattern of context-dependence reflected in FN amplitude (Holroyd et al., 2004; Kujawa et al., 2013). Time-frequency decomposition techniques have proven useful for isolating important activity, and have shown that time-domain ERPs can be better represented as separable processes in delta (0-3 Hz) and theta (3-7 Hz) (Başar, Başar-Eroglu, Karakaş, & Schürmann, 2001; Bernat, Malone, Williams, Patrick, & Iacono, 2007; Cavanagh, Zambrano-Vazquez, & Allen, 2012; Cohen, Elger, & Ranganath, 2007; Demiralp, Ademoglu, Istefanopulos, Başar-Eroglu, & Başar, 2001). Furthermore, recent work has suggested that differences in FN amplitude are due in large part to the superposition of a reward positivity (RewP) component primarily composed of delta activity and a negative-going deflection consisting of theta activity (Bernat, Nelson, Holroyd, Gehring, & Patrick, 2008a; Holroyd, Pakzad-Vaezi, & Krigolson, 2008). Thus, while there has been important attention on time-frequency decomposition of the FN (Bernat, Nelson, Holroyd, Gehring, & Patrick, 2008b; Foti, Weinberg, Dien, & Hajcak, 2011; Holroyd et al., 2008; Proudfit, 2015), the current study seeks to assess whether the role of context in feedback processing is better elucidated using time-frequency analysis.

Context and Feedback Processing

Holroyd et al. (2004) conducted the initial investigation of the role of context in feedback processing. Reward context was operationalized by employing a gambling task with three blocks of trials: a block in which participants could only win money or break even (gain context), a block in which participants could win, lose, or break even (even context), and a block in which participants could only lose money or break even (loss context). Notably, the data were collected in two separate experiments: Experiment 1 comprised the even context which included best (+10), middle (0), and worst (-10) outcomes, and Experiment 2 comprised gain (+5, +2.5, +0) and loss (-0, -2.5, -5) contexts. Holroyd et al. (2004) evaluated the context-dependence versus independence of the FN. Context-dependence refers to the processing of outcomes in a relative manner within each context, whereas context-independence refers to the processing of outcome values in an absolute manner, independent of context. The FN was considered context-dependent if two criteria were met: 1) identical outcome values were evaluated differently across contexts (e.g., the zero outcomes in the gain and loss contexts) and 2) the same outcome levels (e.g., best outcomes) were processed similarly across contexts where outcome levels scaled with reward magnitude within context (Holroyd et al., 2004). The FN was deemed context-independent if two criteria were met: 1) identical outcome values were evaluated similarly across contexts and 2) the same outcome levels were processed differently across contexts (Holroyd et al., 2004). Results revealed that the FN met criteria for context dependence in some but not all cases (see Figure 1). Both experiments revealed a main effect of outcome level, suggesting that the outcomes were evaluated in a relative, context-dependent manner. The first criterion for context dependence was met for the zero conditions in the gain and loss contexts (i.e., $+0 \neq -0$). In order for the second criterion of context dependence to be met, the same outcome levels should be processed similarly across contexts and the outcome levels should scale with reward magnitude within context. This criterion was only partially met. The middle outcomes across contexts were evaluated similarly but best and worst outcomes were not (see Figure 1a). Additionally, the FN amplitude did not differ for middle and worst outcomes within each context (see Figure 1a). Furthermore, the breaking even outcome accounted for outcome differences in the loss context, as differences were not seen between varying loss magnitudes. Because a breaking even outcome was included in each context, the contexts did not purely reflect one type of outcome valence (either all gains or all losses). If the breaking even outcomes were removed and only valenced outcomes were considered, FN differences would only be seen in gain-possible contexts (i.e., gain and even) but not in the loss context (see Figure 1b). Thus, the loss context does not show a pattern of context-dependence when only the loss-valenced outcomes are considered. Taken together, these findings suggest that the FN may reflect a combination of context-dependent and independent processing which may be influenced by the breaking even condition.

Nieuwenhuis et al. (2005) investigated the role of context on reward-sensitive brain regions indexed by functional magnetic resonance imaging (fMRI). Nieuwenhuis et al. (2005) used a similar task and design as Holroyd et al. (2004), except the even context was excluded. The gain context consisted of best (+60), middle (+30), and worst (+0) outcomes, and the loss context consisted of best (-0), middle (-20), and worst (-40) outcomes. In order to determine which brain regions were sensitive to reward, Nieuwenhuis et al. (2005) analyzed the blood oxygen level dependent (BOLD) difference between the highest magnitude gain outcome (+60) and the lowest magnitude loss outcome (-40). The regions that showed significant BOLD differences were used to evaluate the effects of context. In all regions, the zero conditions in the gain and loss contexts were significantly different, meeting the first criterion for context-dependence (Nieuwenhuis et al., 2005). The second criterion was partially met, with most reward regions showing a similar BOLD response to the best outcome in both contexts (+60 in gain and -0 in loss; Nieuwenhuis et al., 2005). However, the BOLD response did not differ between the middle and worst outcomes, suggesting a binary difference between the best outcomes in each context and the remaining outcomes (Nieuwenhuis et al., 2005). As in the Holroyd et al. (2004) paper, the breaking even condition drove the binary difference between the best outcome (-0) and the remaining outcomes (-20 and -40) in the loss context. The BOLD response differed between varying gain magnitudes (+60 and +30) but not varying loss magnitudes, which is consistent with the FN findings described above and other previous work (Holroyd et al., 2004; Yeung & Sanfey, 2004). These results provide more evidence of context-dependent processing in gain-possible contexts and not loss contexts.

The previous literature on context has separated contexts into blocks of trials within the task (Holroyd et al., 2004; Nieuwenhuis et al., 2005). Kujawa et al. (2013) evaluated the context-dependence of the FN by manipulating outcomes on a trial level (local outcomes) versus a task level (global outcomes). A cued gambling task was utilized, where one cue signaled a gain (+50) or breaking even (+0) and the other cue signaled a loss (-25) or breaking even (-0) (Kujawa et al., 2013). They hypothesized that if the FN was sensitive to local outcomes, then the worst outcomes for each cue would be associated with a relative negativity compared to the best outcomes (Kujawa et al., 2013). If the FN were sensitive to global outcomes, then a relative negativity would be seen for all non-gain outcomes. Results revealed that breaking even for gain (+0) and loss (-0) cues and losing (-25) were associated with similar FNs (Kujawa et al., 2013). Only the gain outcome (+50) was associated with an enhanced positivity and was significantly different than the other outcomes (Kujawa et al., 2013). These results suggest that the FN is more sensitive to global than local contexts, such that a relative negativity was seen for all non-gain outcomes. FN was found to be a binary representation of favorable (+50) compared to unfavorable outcomes (+0, -0, -25) (Kujawa et al., 2013). These results indicate that outcomes are not processed in a relative, context-dependent manner when the context changes at a local, trial level. Rather, context-dependent processing may be differentially elicited when trials are presented in a sustained block representing one context, as shown by the global context in Kujawa et al. (2013) and the block designs in Holroyd et al. (2004) and Nieuwenhuis et al. (2005).

Time-domain vs. Time-frequency Analysis

Conventional FN measures have traditionally been associated with negative feedback because the component is diminished or absent following positive feedback (Gehring & Willoughby, 2002; Miltner, Braun, & Coles, 1997); however, more recent work has suggested modulation of the FN by positive feedback. Bernat et al. (2008) and Holroyd et al. (2008) provided initial evidence of a reward positivity (RewP) component, which is enhanced for positive relative to negative feedback. This research and other recent work has indicated that smaller negative FN amplitude elicited by positive feedback is partially explained by the superposition of a heightened slow, positive waveform, the RewP (Bernat et al., 2008a; Foti et al., 2011; Holroyd et al., 2008; Kujawa et al., 2013; Proudfit, 2015). Time-frequency analysis suggests that the RewP is composed primarily of delta activity, not theta (Bernat et al., 2011; Bernat et al., 2015). Recent work based on temporal-spatial principal component analysis (PCA) and time-frequency PCA of the FN has indexed a positive amplitude component in delta that is increased for gains relative to losses (Bernat, Nelson, & Baskin-Sommers, 2015; Bernat, Nelson, Steele, Gehring, & Patrick, 2011; Carlson, Foti, Mujica-Parodi, Harmon-Jones, & Hajcak, 2011; Foti, Weinberg, Bernat, & Proudfit, 2014; Foti et al., 2011; Weinberg, Riesel, & Proudfit, 2014). Additionally, this RewP component has been shown to be sensitive to reward magnitude, relative outcome, and outcome expectancy (Bernat et al., 2015, 2008a; Cavanagh, 2015; Holroyd, Krigolson, & Lee, 2011; Holroyd et al., 2008; Massey et al., 2015; Massey, Bachman, & Bernat, 2016), but the influence of context on the RewP has not been evaluated.

Delta activity associated with the RewP is partially responsible for modulations in FN, but time-frequency analysis has revealed that the FN is composed of theta activity as well. Regression analyses using time-frequency components as predictors and the FN as the outcome measure have revealed that delta and theta contribute unique sources of variance to the FN, such that increased theta reflects losses and increased delta is associated with gains (Bernat et al., 2015, 2008a, 2011; Cohen et al., 2007; Nelson, Patrick, Collins, Lang, & Bernat, 2011). These findings provide strong evidence that separable neural activity indexing losses and gains contribute to the FN. Foti et al. (2014) extended this work by applying source localization to time-frequency measures of the FN, where two distinct neural generators were identified. Loss-related theta activity was localized in the ACC, while gain-related delta activity was focused in the striatum (Foti et al., 2014). These results indicate that discrepancies regarding the FN and outcome valence can be clarified by time-frequency analytic approaches.

Previous research has shown that when outcome stimuli in a gambling task provide multiple pieces of information, theta is sensitive to the most primary or salient stimulus attributes (often outcome valence – loss vs. gain), while delta is modulated by primary as well as more complex secondary characteristics, such as reward magnitude and expectancy (Bernat, Nelson, & Baskin-Sommers, 2015; Massey et al., 2015; Massey, Bachman, & Bernat, revise and resubmit). This work provides further support that feedback processing is better measured using time-frequency analysis, which indexes separable processes in delta and theta that underlie the FN. Additionally, this work provides insight into how theta and delta may be influenced by context. Because theta has been shown to reflect a more binary evaluation of bad vs. good, we predict that theta will show similar best-worst differences across contexts, indicating a pattern of context-dependent processing. Whereas we expect delta to show an overall sensitivity to best versus worst outcomes across contexts, we predict that this effect will be qualified by an interaction between outcome and context where delta scales with the magnitude of the reward across contexts. Additionally, we expect best-worst differences in delta to scale with context valence (gain > even > loss), which is consistent with previous research showing outcome valence differences in gain-possible contexts but not among varying loss magnitudes in loss contexts (Holroyd et al., 2004; Nieuwenhuis et al., 2005). This pattern of effects in delta would be more consistent with context-independent processing.

Current Study

The current study seeks to build on previous work investigating the role of context in reward processing by utilizing time-frequency analysis of the FN. Previous research using ERP and fMRI methodology has provided important considerations for understanding the context-dependence of reward processing. Namely, the breaking even outcome has led to results showing context-dependent processing in all contexts. However, when only valenced outcomes are considered (i.e., gains and losses), context-dependent processing is only seen in gain-possible contexts, as no differences exist between varying loss magnitudes in the loss context. Furthermore, context-dependent processing seems to be limited to tasks in which trials for a given context are presented in a sustained block. Lastly, the ERP studies investigating context used time-domain analysis of the FN, which is problematic due to research showing that the FN contains separable underlying processes indexed in delta and theta frequency bands.

The current study utilized a modified version of the gambling task created by Holroyd et al. (2004) with one key difference: the current study's task did not include a breaking even outcome in each context due to the problems described above. In removing the breaking even outcome, the current study aims to assess processing variations across contexts that purely reflect one type of outcome (all gain or all losses) or a combination of both types (as in our even context). Based on previous research on context as well as more recent research on time-frequency analysis of feedback processing, the current study aims to assess: 1) whether theta reflects context-dependent processing, 2) whether delta reflects a combination of context-independent and dependent processing, and 3) how delta and theta are related to modulations of the FN by context.

Chapter 2: Methods

Participants

Participants (n = 152) were recruited from undergraduate students at Florida State University. Five participants were excluded due to a problem with the EEG recording (e.g., experimenter error or software malfunction) and fifteen participants were excluded due to an excessive number of EEG artifacts (>33% of trials rejected using methods described below). The final sample contained 132 participants 18 years of age or older (80 females; M age = 19.99, SD = 3.52). The final sample was not significantly different than the original sample on key demographic variables, including gender and age. Participants were screened for visual impairments, neurological conditions, and/or traumatic brain injuries. Participants were provided informed consent before starting the study and were offered monetary compensation ($10/hr) or course credit for participation.

Procedures

EEG data was collected in a sound-attenuated, dimly lit room. Experimental stimuli were presented on a 21-inch Dell high-definition CRT color monitor, centrally placed at a viewing distance of 100 cm, subtending a visual angle of 3.5°. E-Prime version 1.1 was used to present the stimuli, and a PST Serial Response Box (Psychology Software Tools, Inc.) was used to collect responses to the task.

Participants performed a modified version the gambling task used in (Holroyd et al., 2004), as shown in Figure 2. Each trial consisted of two circles presented side-by-side with a black border and white background. Participants were instructed to select one of the circles by pressing the left or right button on the button box. The circles remained on the screen until the participant made a selection, at which time the selected circle turned red. 1000 ms after the selected circle turned red, the participant received monetary feedback inside the circle, which was displayed for 1000 ms. Best and worst outcomes were possible in each of three context: +5 or +15 in the gain context, +5 or -5 in the even context, and -5 or -15 in the loss context. The task was divided into six blocks, with two blocks for each context (i.e., gain, even, and loss). Each block consisted of 24 trials, resulting in 48 trials per context and a total of 144 trials. Blocks were counterbalanced across participants such that no one context was presented in two consecutive blocks. Participants were informed of the context type before each block. They were also told they should respond in a way that maximized their earnings, and that they would be given a monetary reward associated with their performance at the end of the task. Unbeknownst to the participants, the two outcomes in each block were presented at random with an equal probability. All participants were given $5.00 at the end of the task. Before the task began, participants completed a brief practice.

Psychophysiological Data Acquisition

Data were recorded using a Neuroscan 128-channel Quik-Cap (sintered Ag-Ag/Cl; non-standard layout) as well as a 128-channel Synamps RT amplifier (Neuroscan, Inc.). Ten electrodes around the ears were removed from analysis due to inconsistent connection to the scalp across participants, leaving a total of 113

EEG channels. Horizontal electrooculogram activity was recorded from electrodes placed on the outer canthus of both eyes, while vertical electrooculogram activity was recorded from electrodes placed above and below the left eye. Impedances were kept below 10 kΩ. EEG signals were vertex referenced during recording (directly between Cz and CPz), and re-referenced to averaged mastoid signals offline, collected using an analog 0.05 to 200 Hz bandpass filter and digitized at 1000 Hz using Neuroscan Acquire (Neuroscan,Inc.).

Data Preprocessing

Epochs of three seconds were then taken from 1000 ms pre- to 2000 ms post-stimulus onset with a 150 ms pre-stimulus baseline, and were re-referenced to averaged mastoid sites. Ocular artifacts were corrected with a regression-based algorithm developed by Semlitsch, Anderer, Schuster, & Presslich (1986) in the Neuroscan Edit 4.5 software (Neuroscan, Inc.) and downsampled to 128 Hz using the Matlab resample function (Mathworks, Inc.), which utilizes an anti-aliasing filter before resampling. Then, two criteria for data cleaning were used. In the first, trials were rejected if activity at F3 or F4 exceeded ±100 μV in either the pre-stimulus period of -1000 to -1 ms, or the post stimulus period of 1 to 2000 ms, to remove larger face or eye artifacts not appropriately handled by the Semlitsch algorithm. For the second criterion, trials were rejected if activity in any electrode exceeded ±150 μV during the same pre- and post-stimulus time periods. Together, these removed 9.1% of all trials from analysis. Visual analysis of the averaged waveforms indicated that 0.002% of electrodes were disconnected during recording and were replaced with the mean of the nearest neighbors. After preprocessing, data were averaged according to the six different outcomes specified above under *Procedures*.

Subsampling

Although data cleaning improves the quality of the data, it removes several trials, leaving an uneven number of trials across outcome types and participants. Subsampling and bootstrapping during ERP averaging are methods that are particularly useful for extracting the maximum amount of variance possible in situations with limited data. These approaches are helpful when participants have low trial counts for a given outcome type, and they reduce any bias associated with an uneven number of trials across outcome types and participants. Subsets of five trials for each outcome were subsampled 50 times, and then bootstrapped 500 additional times.

Data Reduction

Time-Domain amplitude components.

Time-domain (TD) measures of evoked power were extracted for the feedback negativity (FN). The FN was defined as a negative deflection ranging between 203 to 352 milliseconds post-outcome stimulus (fit to the edges of the FN negative peak in the grand average waveform), consistent with previous work (Gehring & Willoughby, 2002; Holroyd & Coles, 2002; Miltner et al., 1997). This time range was converted from bins of the 128 Hz resampled signal. For statistical analyses, this component was reduced to a group of 9 electrodes (shown in Figure

3) clustered around Cz. While the FN has traditionally been quantified at FCz, time-frequency analysis has revealed that the FN is composed of fronto-central theta activity and central-parietal delta activity. Thus, 9 electrodes centered around Cz were used due to their location in the middle of these contributing regions.

Time-Frequency evoked power.

To evaluate the time-frequency (TF) phase dynamics related to time-domain ERP signals, TF decompositions were performed upon trial-averaged ERPs. This procedure allows phase-consistent evoked ERP activity to be re-represented in the TF domain, and similar methods have been successfully used to evaluate the relationships between time-frequency activity and time-domain ERPs in a number of other reports (Bernat et al., 2011; Harper, Malone, Bachman, & Bernat, 2016; Harper, Malone, & Bernat, 2014; Nelson et al., 2011). First, 3rd order Butterworth filters were used to isolate activity within delta and theta frequency ranges, based on the visual inspection of the unfiltered representation of time-frequency energy following the outcome stimulus for one second. A 4 Hz lowpass filter was employed to isolate delta, and a 3 Hz highpass filter in conjunction with an 8 Hz lowpass filter was used for theta. TF decompositions were produced using a binomial reduced interference distribution (RID) variant of Cohen's class of time-frequency transformations upon the full epoch of the filtered signals, using 32 time bins per second and 2 frequency bins per Hz. The RID was chosen to better represent low-frequency activity and avoid smearing the representation of such activity in time (Bernat, Williams, & Gehring, 2005a).

Principal component analysis (PCA) was then used separately on each filtered TF decomposition, using a post-stimulus time window of 0-1000ms and a 0-12 Hz frequency window. The TF-PCA data matrix contained TF points as vectors and subject/electrode/trial-averaged scores as rows (a more detailed explanation of this process can be found in Bernat et al. (2005)).

Figure 4 displays the grand-averaged TF-PCA decomposition. Four principal components (PCs), explaining 41% of the total variance, were extracted from theta as the best representation of the data. PC1 represented medial frontal theta during the FN, PC2 reflected theta activity after P3, PC3 represented high delta (3 Hz) bilateral occipital activity, and PC4 reflected P2 frontal theta activity. A three PC solution was used as the best representation of delta activity during feedback processing, which explained 68% of the total variance. PC1 reflected low frequency (~1 Hz) slow wave activity, PC2 represented bilateral occipital post-P3 activity, and PC3 reflected parietal delta activity (~2 Hz) during the FN time window, also known as the RewP. Theta PC1 and delta PC3 were chosen for further statistical analysis given a priori hypotheses regarding theta and delta activity during the FN time window (see Figure 4). The mean PC-weighted TF evoked energy of these two PCs were narrowed down to a subset of 9 electrodes each, shown in Figure 3, with delta clustered around a centro-parietal electrode and theta clustered around a fronto-central electrode. These clusters were selected based on the topographical center of best-worst differences.

Specific Aims and Hypotheses

There are three primary aims of this study. The first is to assess whether theta is primarily context dependent. Previous research has suggested that theta represents a more narrow, local process, sensitive to primary task features such as best-worst differences (Bernat et al., 2015; Massey et al., 2015, 2016). We thus hypothesize as follows (see Figure 5):

1a. Theta will be associated with similar best/worst differences across contexts (context-dependence).
1b. Theta activity will be different between outcomes of the same absolute value across contexts, such as +5 in the gain context and +5 in the even context (context-dependence).

The second aim is to evaluate whether delta reflects both context-dependent and context-independent processing. Previous work has shown that delta is increased for best/reward outcomes and can index a number of different processes operating across lower frequencies, including primary best/worst differences and reward magnitude (Bernat et al., 2015). We thus hypothesize as follows (see Figure 5):

2a. Delta in response to best outcomes across contexts will be associated with increases in amplitude linearly related to the amount of reward (context-independence).
2b. Best/worst differences in delta will be attenuated in the loss condition, relative to the even and gain conditions given previous literature showing ERP differences only in gain-possible contexts (context-independence).
2c. Delta activity will be different between outcomes of the same absolute value across contexts, such as +5 in the gain context and +5 in the even context (context-dependence).

A third aim in this study is to assess the conventional time-domain FN amplitude and context-related processing. Based on previous literature, we expect FN to show a combination of context-dependence and independence. Given that the time-domain FN contains multiple frequency bands reflecting separable processes, we do not propose specific hypotheses regarding the relationship between context and the FN. However, we do expect theta and delta to contribute unique sources of variance to the FN as previous research has shown.

Data Analysis Plan

Three measures were analyzed: the FN and the TF principal components representing delta and theta. First, each of the three measures was analyzed in separate 2 x 3 ANOVAs of outcome (best vs. worst) by context (gain, even, or loss). A main effect of outcome is expected if theta is a reflection of context-dependent evaluation of outcomes. An interaction of outcome and context will be seen if delta is a reflection of a combination of context-independent and dependent evaluation of outcomes. Furthermore, this interaction will show the largest delta to best outcomes in the gain context and the smallest in the loss context, and will show little difference in delta across the worst outcomes. Additionally, pairwise comparisons will be assessed between outcomes of the same absolute value across context, a necessary comparison for determining context dependence and independence. Finally, linear regression analysis will be used to assess the contributions of delta and theta to the FN.

Design Considerations

The study design was similar to previous research on reward context with one key exception: there was no breaking even outcome in each condition. Holroyd et al.'s (2004) initial study using this design aimed to assess reward context within a reinforcement learning theory framework. The reinforcement learning theory of the FN suggests that the monitoring system judges an outcome to be favorable or unfavorable based on the range of possible outcomes. Of the possible outcomes in a given context, the monitoring system will predict the middle value and deviations from this expectation will produce relative increases or decreases in FN amplitude. Thus, in order to evaluate reward context under this theoretical framework, there must be at least three possible outcomes. It is unclear why Holroyd et al. (2004) chose to utilize breaking even as the third outcome rather than an additional valenced outcome (i.e., another gain value for the gain context and another loss value for the loss context). However, we wanted the gain and loss contexts to purely reflect one type of outcome and thus we decided to remove the breaking even outcome.

An additional design consideration was whether or not to inform participants about the parameters of each context in advance. Previous research utilizing this type of task has informed participants about what types of outcomes they will receive in each block, which sets expectations about whether participants will be only winning money, only losing money, or a combination of both. Because the primary aim was to investigate the role of context, the current study informed participants about the block parameters in order to set clear expectations about the type of context. Without this expectation set in advance, participants would have to learn the parameters of each block over time and the results may not have been as robust.

Chapter 3: Results

Time-domain FN Analysis

The first row of Figure 3 displays the average time-domain outcome-locked ERP waveform for the average of nine electrodes (see Figure 2). The FN is evident as the negative deflection peaking approximately 230 milliseconds after outcome stimulus onset. The scalp distribution of the FN is parietal, which is common when a difference-wave approach is not used.

To test for the effects of outcome (best vs. worst) and context (gain, even, and loss) on the FN, a repeated-measures 2 X 3 ANOVA was performed. For the FN, an interaction between outcome and context was found ($F(2,130) = 38.36$, $p < .001$), where the negative amplitude of the FN scaled with outcome value across the best outcomes and was largest in the loss context (see Figure 4). No differences were seen between the worst outcomes across contexts (see Figure 4). Best-worst differences scaled with the context, where the greatest difference was seen in the gain context (see Figure 4).

Time-frequency Analysis

The second row of Figure 3 depicts the time-frequency (TF) average waveform and PCA decomposition of the outcome-locked ERP for theta and delta. The distribution of peak activation is medial frontal for theta and parietal for delta. Consistent with prior research (Bernat et al., 2015, 2011) theta activity mirrored the FN in terms of latency and enhanced amplitude to loss relative to gain feedback. Furthermore, it is apparent that theta negative polarity activity contributes to enhanced *negative* FN amplitude while positive delta activity counteracts the negative amplitude of the FN.

To test for the effects of outcome and context on theta and delta, we again utilized a repeated measure 2 X 3 ANOVA design. For theta, an interaction between outcome and context was found ($F(2,130) = 26.10$, $p < .001$), where the gain and even contexts showed similar best-worst differences and the loss context showed no difference (see Figure 4). Thus, there was partial support for hypothesis 1a, such that theta showed a pattern of context dependence in the gain and even contexts but not the loss context. For delta, an interaction was found between outcome and context ($F(2,130) = 43.83$, $p < .001$), where, similar to the FN, delta scaled with reward magnitude and the largest best-worst difference was seen in the gain context. These findings provide strong support for hypotheses 2a and 2b, suggesting that delta partially reflects context-independent processing.

To test the hypotheses that theta and delta amplitudes will be different between outcomes of the same absolute value across contexts (indicating context-dependence), paired sample t-tests were conducted between the +5 outcomes in the gain and even contexts as well as the -5 outcomes in the even and loss contexts. For theta, both +5 ($t(131) = 4.80$, $p < .001$) and -5 ($t(131) = 4.30$, $p < .001$) comparisons were significantly different across contexts, supporting hypothesis 1b and providing more evidence that theta reflects context-dependent processing. For delta, results revealed a significant difference between the +5 outcomes in the gain and even contexts ($t(131) = -7.17$, $p < .001$) but no difference between the -5 outcomes in the even and loss contexts ($t(131) = -.50$, p

= .62). These findings provide partial support for hypothesis 2c, suggesting that delta reflects context-dependent processing in the gain and even contexts but not in the loss context.

Regression analysis predicting the FN with theta and delta, a technique used in previous research (Bernat et al., 2015, 2011), indicated that delta and theta contribute unique variance to the FN (see Table 1). However, delta accounted for more variance in the FN than theta.

The relationship between context and outcome was more similar across delta and theta than expected, especially because there was no difference in theta between the two loss magnitudes in the loss context. However, the pattern of results in the gain and even contexts were more consistent with the proposed hypotheses and appear to show meaningful differences between delta and theta in relation to context. Thus, in order to more directly compare context-related processing between delta and theta in the gain and even contexts, correlations of best-worst differences between the gain and even contexts were performed and compared across frequencies. Results revealed a large correlation between best-worst differences in the gain versus even contexts for theta ($r(130) = .60$, $p < .001$) but a small correlation for delta ($r(130) = .17$, $p = .05$). Fisher's z-test confirms that these correlations are significantly different ($z = 4.19$, $p < .001$), which reveals that outcome processing in the gain and even contexts was more similar in theta than in delta. These findings further support the hypotheses that theta reflects more context-dependent processing and delta reflects a combination of context-independent and context-dependent processing.

Chapter 4: Discussion

In the current study, we investigated the role of context in feedback processing by using time-frequency analysis to evaluate ERPs in a modified gambling task. The primary aims were to: 1) assess whether theta reflects context-dependent processing, 2) assess whether delta reflects a combination of context-independent and context-dependent processing, and 3) assess the contributions of theta and delta to the FN and the relationship between context and the FN.

Neither theta nor delta was completely context-dependent or context-independent; however, theta met more criteria for context-dependence and delta met more criteria for context independence. In support of hypothesis 1a, theta showed similar best-worst differences across the gain and even contexts, suggesting a pattern of context-dependence. These results are consistent with findings suggesting that theta is a binary reflection of best and worst outcomes, where theta is enhanced for the worst outcome (Bernat et al., 2015, 2011). However, similar best-worst differences were not seen in the loss context, suggesting that theta does not reflect context-dependent processing in a loss-only context. In support of hypothesis 1b, theta activity was significantly different between the +5 outcomes in the gain and even contexts and the -5 outcomes in the even and loss contexts, again suggesting that theta reflects context-dependent processing.

Compared to theta, delta met more criteria for context-independence. In support of hypothesis 2a, delta activity to the best outcomes scaled linearly with reward magnitude, such that the largest delta activity was seen in the gain context and the smallest was seen in the loss context. These results are consistent with previous literature indicating that delta, specifically the reward positivity, scales with reward magnitude (Bernat et al., 2015; Proudfit, 2015). Additionally, delta's sensitivity to reward magnitude replicates previous findings suggesting that delta is modulated by complex, secondary stimulus characteristics (Bernat et al., 2015). In support of hypothesis 2b, best-worst differences in delta were evident in the gain and even contexts but not in the loss context, revealing a pattern of context-independence. These results are consistent with previous research showing processing differences in gain-only contexts but not in loss contexts (Holroyd et al., 2004; Kujawa et al., 2013; Nieuwenhuis et al., 2005). Delta activity differed between outcomes of the same absolute value in the gain and even contexts, meeting one criterion for context-dependence. However, this was the only case of context-dependence, suggesting that delta is largely context-independent.

In order to more directly compare context-related processing between theta and delta, correlations of best-worst differences between gain and even contexts were compared across frequencies. Best-worst differences in the gain and even contexts were significantly correlated in theta and delta, but there was a stronger relationship between the gain and even contexts in theta. These findings provide further support for the prediction that theta reflects more context-dependent processing and delta reflects more context-independent processing.

Analysis of the FN revealed a similar pattern to delta, where the FN scaled with the best outcomes across contexts and showed little change in response to the worst outcomes across contexts. Indeed, regression analysis predicting the FN with delta and theta revealed that delta and theta contributed unique sources of variance to the FN, but delta accounted for more variance than theta. These results are consistent with previous work showing that the superposition of a slow, positive wave (i.e., delta) during the FN time-window is partially responsible for modulation of the FN (Bernat et al., 2008a; Holroyd et al., 2011, 2008; Proudfit, 2015). Furthermore, similar processing of losses in loss-only contexts as reflected by FN amplitude was shown in the current study and in previous work (Holroyd et al., 2004; Nieuwenhuis et al., 2005). The lack of differences among varying loss magnitudes in the loss context for delta and theta suggests that both frequencies are contributing the null FN effects in the loss context. These findings highlight the importance of evaluating time-domain components using time-frequency measures, as they isolate frequencies that represent distinct processes that are convoluted in the time-domain.

Our findings also indicate that context valence plays a critical role in reward processing. In gain-possible contexts (i.e., the gain and even contexts), best-worst differences scaled positively with context valence (gain > even) for delta and showed binary differences for theta (gain = even). Alternatively, in the context where a gain was not possible (i.e. the loss context), all outcomes were processed similarly in delta and theta (i.e. there were no outcome valence differences, and amplitudes were the lowest of all contexts). This context valence theory may help explain previous work showing no differences between varying loss magnitudes in the FN and BOLD signal (Holroyd et al., 2004; Kujawa et al., 2013; Nieuwenhuis et al., 2005). Evidence from the adaptive gain theory may provide support for this context valence theory (Aston-Jones & Cohen, 2005). The adaptive gain theory suggests that when task utility is adequate, participants exploit as much reward as possible (Aston-Jones & Cohen, 2005). Thus, in delta, a moderate best-worst difference is observed in the even context when there is the option of a small reward and a larger best-worst difference is seen in the gain context when there is the option of a large reward. However, when task utility wanes and becomes less rewarding, participants enter an exploration mode where they disengage from the task at hand and explore alternative means for reward (Aston-Jones & Cohen, 2005). Thus, it may be that participants became defeated in the loss context and disengaged from the task.

Limitations and Future Directions

The context valence explanation is post-hoc and should be explored in future work. For example, our work suggests that activity in delta and theta is similar for varying amounts of loss in a loss-only context, but future work should investigate if these frequencies are modulated when varying loss magnitudes are presented in the same context as gains. Additionally, previous work on context dependence has employed an additional level – the breaking even outcome. Our task only contained gains and losses because we wanted each context to purely reflect one type of outcome (e.g., only gains in the gain context). The absence of the breaking even outcome in the current study limits the ability to make direct comparisons to previous work.

Future work should also evaluate how one's internal context interacts with reward context as indexed by theta and delta. For example, how might individual differences in psychopathology influence context-related processes in delta and theta? Because theta represents a more basic evaluation of unfavorable vs. favorable outcomes, that is similar across contexts (with the exception of the loss context), it may be more resilient to psychopathology. Individual differences in psychopathology may result in an overall diminution or enhancement of theta, but we would expect the unfavorable vs. favorable differences to stay intact. Indeed, previous research has shown that theta is less susceptible to individual differences in psychopathology, but more work in this area is needed (Bernat et al., 2011). Delta reflects a more nuanced evaluation of outcomes, functioning in a context-dependent and independent way, and perhaps context-related processes in delta would be more sensitive to individual differences in psychopathology. Previous research has shown modulation of delta in relation to psychopathology (Bernat et al., 2011; Foti & Hajcak, 2009; Foti et al., 2014; Proudfit, 2015), but future work should assess how psychopathology influences reward processing in varying contexts.

Conclusions

The results of the current study indicate that the role of context in reward processing is better elucidated using time-frequency analysis. Theta was more context-dependent and showed a binary response to best-worst differences in the gain and even contexts. Delta was more context-independent: the best outcomes scaled linearly with reward magnitude and best-worst differences scaled with context valence. The relationship between context and FN amplitude was similar to that of delta, and delta accounted for more variance in the FN than theta. Our findings reveal that theta and delta are differentially sensitive to context and that context valence may play a critical role in determining how the brain processes good and bad outcomes.

Tables

Table 1.

Multiple regression model of delta and theta predicting FN.

	FN	
	Beta	*t*
Delta	.68	26.24***
Theta	-.21	-8.29***

Notes: $R^2 = 0.48$

*$p < .05$, **$p < .01$, ***$p < .001$

Figures

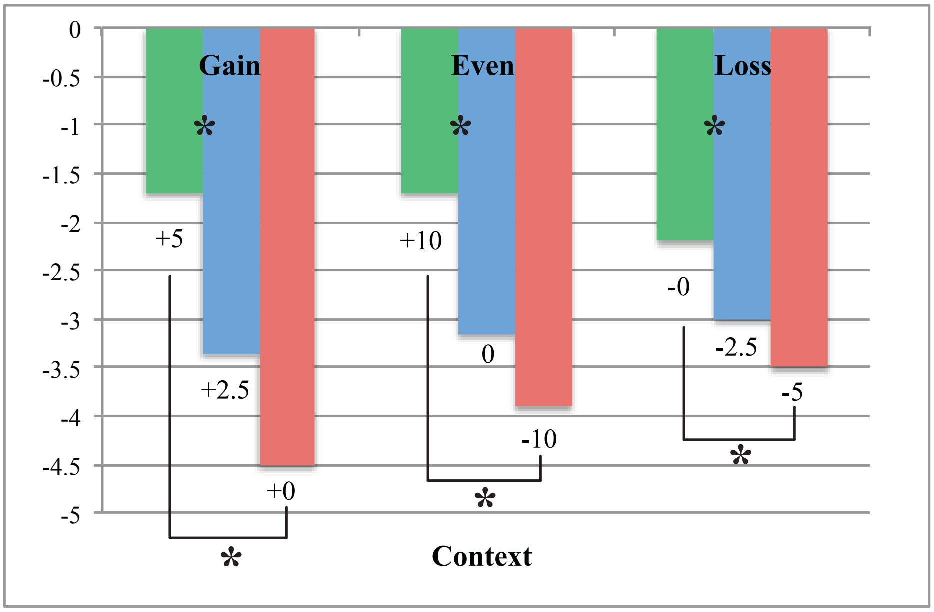

Figure 1a. Results estimated from Holroyd et al. (2004).

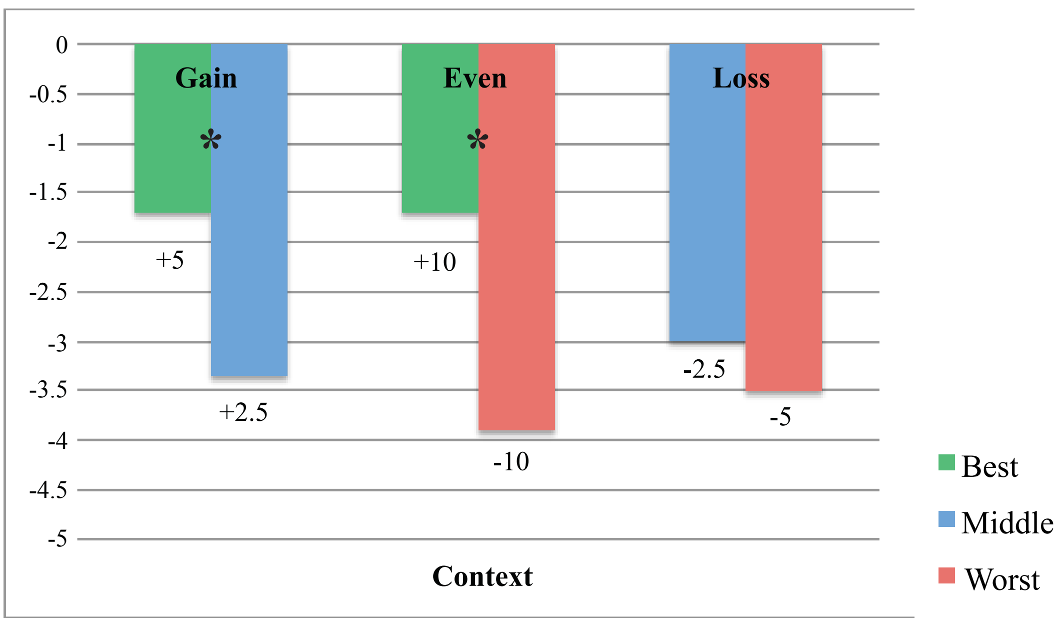

Figure 1b. Results estimated for Holroyd et al. (2004) without the breaking even outcome.

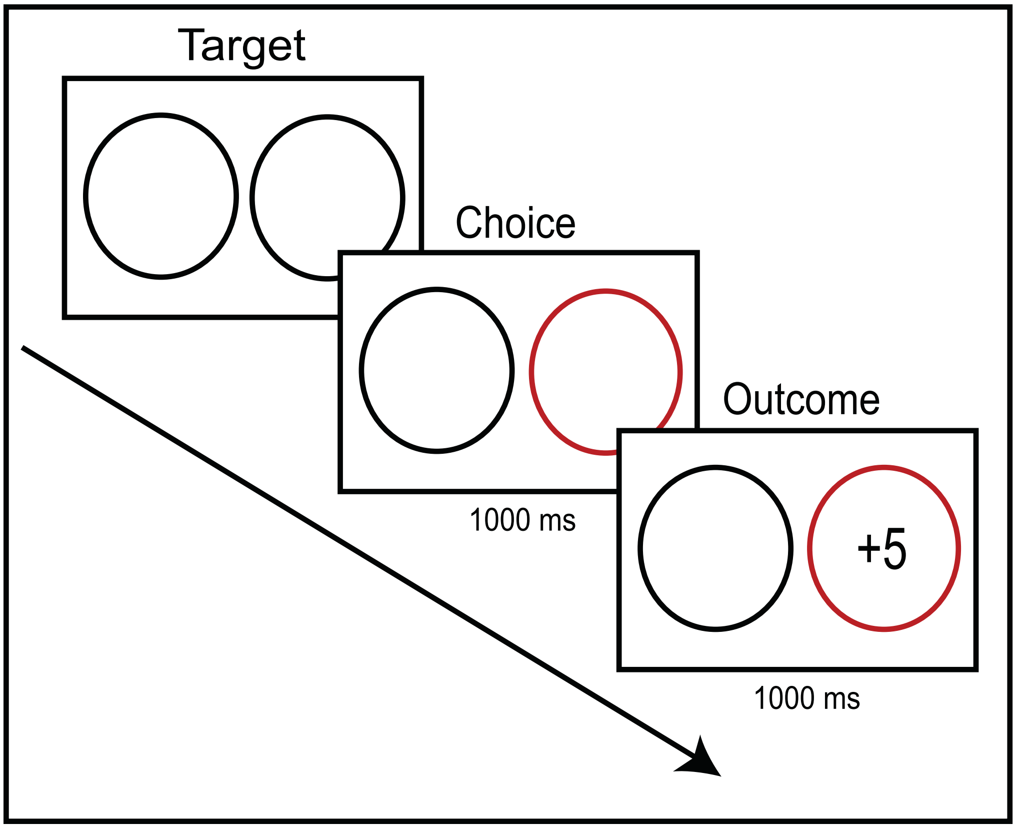

Figure 2. Sequence of stimulus and outcome events in the reward context gambling task.

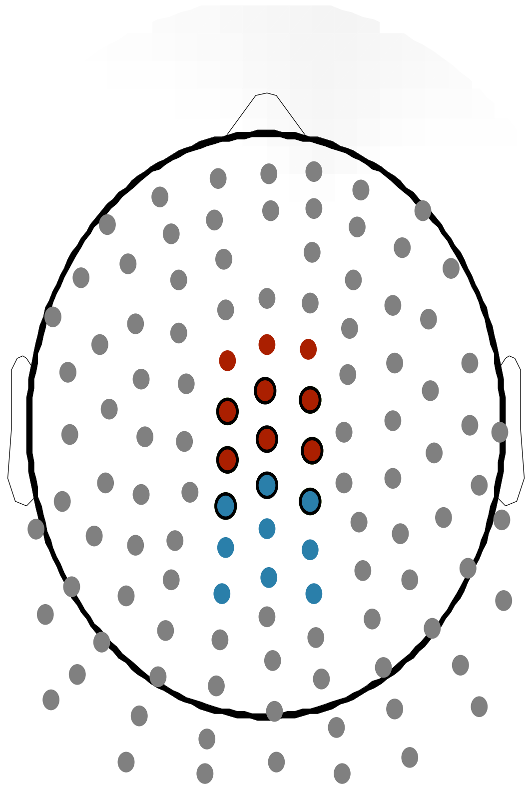

Figure 3. Electrode clusters for analysis of delta (blue), theta (red), and FN (black border).

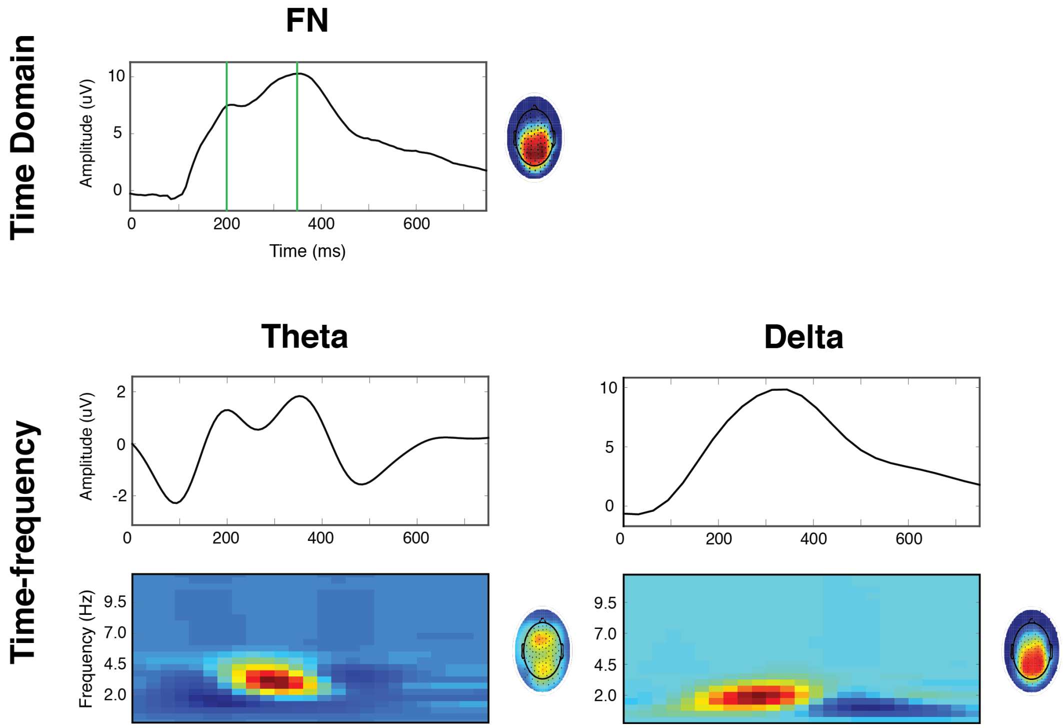

Figure 4. Time domain and time-frequency (TF) decompositions of outcome-locked ERPs. The top panel depicts the average unfiltered time-domain waveform across all trials. FN was quantified as the negative-going deflection between 203-352 ms using peak measurement. The bottom panel shows the filtered ERP waveforms and the TF-PCA decompositions for theta and delta activity during the FN window across all trials. From the topographic maps, theta activity is maximal at fronto-central sites and delta is maximal at parietal sites.

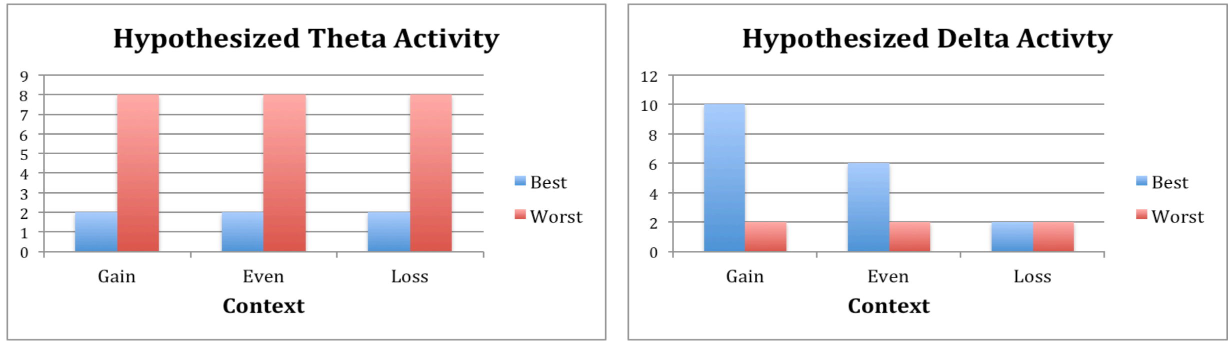

Figure 5. *Hypothesized relationship between context and outcome for delta and theta.*

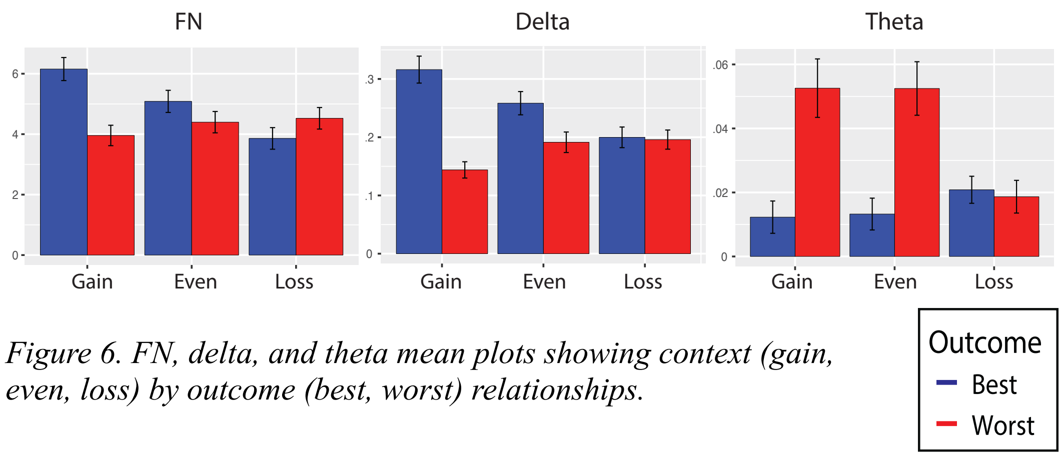

Figure 6. FN, delta, and theta mean plots showing context (gain, even, loss) by outcome (best, worst) relationships.

References:

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.*, *28*, 403–450.

Başar, E., Başar-Eroglu, C., Karakaş, S., & Schürmann, M. (2001). Gamma, alpha, delta, and theta oscillations govern cognitive processes. *International Journal of Psychophysiology*, *39*(2–3), 241–248. https://doi.org/10.1016/S0167-8760(00)00145-8

Bernat, E. M., Malone, S. M., Williams, W. J., Patrick, C. J., & Iacono, W. G. (2007). Decomposing delta, theta, and alpha time–frequency ERP activity from a visual oddball task using PCA. *International Journal of Psychophysiology*, *64*(1), 62–74.

Bernat, E. M., Nelson, L. D., & Baskin-Sommers, A. R. (2015). Time-frequency theta and delta measures index separable components of feedback processing in a gambling task. *Psychophysiology*, *52*(5), 626–637. https://doi.org/10.1111/psyp.12390

Bernat, E. M., Nelson, L. D., Holroyd, C. B., Gehring, W. J., & Patrick, C. J. (2008a). Separating cognitive processes with principal components analysis of EEG time-frequency distributions (Vol. 7074, p. 70740S–70740S–10). https://doi.org/10.1117/12.801362

Bernat, E. M., Nelson, L. D., Holroyd, C. B., Gehring, W. J., & Patrick, C. J. (2008b). Separating cognitive processes with principal components analysis of EEG time-frequency distributions (Vol. 7074, p. 70740S–70740S–10). https://doi.org/10.1117/12.801362

Bernat, E. M., Nelson, L. D., Steele, V. R., Gehring, W. J., & Patrick, C. J. (2011). Externalizing psychopathology and gain–loss feedback in a simulated gambling task: Dissociable components of brain response revealed by time-frequency analysis. *Journal of Abnormal Psychology*, *120*(2), 352.

Bernat, E. M., Williams, W. J., & Gehring, W. J. (2005a). Decomposing ERP time–frequency energy using PCA. *Clinical Neurophysiology*, *116*(6), 1314–1334. https://doi.org/10.1016/j.clinph.2005.01.019

Bernat, E. M., Williams, W. J., & Gehring, W. J. (2005b). Decomposing ERP time–frequency energy using PCA. *Clinical Neurophysiology*, *116*(6), 1314–1334.

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. https://doi.org/10.1016/j.neuroimage.2011.05.037

Cavanagh, J. F. (2015). Cortical delta activity reflects reward prediction error and related behavioral adjustments, but at different times. *NeuroImage*, *110*, 205–216. https://doi.org/10.1016/j.neuroimage.2015.02.007

Cavanagh, J. F., Zambrano-Vazquez, L., & Allen, J. J. (2012). Theta lingua franca: A common mid-frontal substrate for action monitoring processes. *Psychophysiology*, *49*(2), 220–238.

Cohen, M. X., Elger, C. E., & Ranganath, C. (2007). Reward expectation modulates feedback-related negativity and EEG spectra. *NeuroImage*, *35*(2), 968–978. https://doi.org/10.1016/j.neuroimage.2006.11.056

Demiralp, T., Ademoglu, A., Istefanopulos, Y., Başar-Eroglu, C., & Başar, E. (2001). Wavelet analysis of oddball P300. *International Journal of Psychophysiology*, *39*(2–3), 221–227. https://doi.org/10.1016/S0167-8760(00)00143-4

Foti, D., & Hajcak, G. (2009). Depression and reduced sensitivity to non-rewards versus rewards: Evidence from event-related potentials. *Biological Psychology*, *81*(1), 1–8.

Foti, D., Weinberg, A., Bernat, E. M., & Proudfit, G. H. (2014). Anterior cingulate activity to monetary loss and basal ganglia activity to monetary gain uniquely contribute to the feedback negativity. *Clinical Neurophysiology*. Retrieved from http://www.sciencedirect.com/science/article/pii/S1388245714005148

Foti, D., Weinberg, A., Dien, J., & Hajcak, G. (2011). Event-related potential activity in the basal ganglia differentiates rewards from nonrewards: Temporospatial principal components analysis and source localization of the feedback negativity. *Human Brain Mapping*, *32*(12), 2207–2216.

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.

Harper, J., Malone, S. M., Bachman, M. D., & Bernat, E. M. (2016). Stimulus sequence context differentially modulates inhibition-related theta and delta band activity in a go/no-go task. *Psychophysiology*. Retrieved from http://onlinelibrary.wiley.com/doi/10.1111/psyp.12604/full

Harper, J., Malone, S. M., & Bernat, E. M. (2014). Theta and delta band activity explain N2 and P3 ERP component activity in a go/no-go task. *Clinical Neurophysiology*, *125*(1), 124–132.

Holroyd, C. B., & Coles, M. G. (2002). The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, *109*(4), 679.

Holroyd, C. B., Krigolson, O. E., & Lee, S. (2011). Reward positivity elicited by predictive cues. *Neuroreport*, *22*(5), 249–252.

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. https://doi.org/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.

Kujawa, A., Smith, E., Luhmann, C., & Hajcak, G. (2013). The feedback negativity reflects favorable compared to nonfavorable outcomes based on global, not local, alternatives. *Psychophysiology*, *50*(2), 134–138.

Massey, A. T., Bachman, M. D., & Bernat, E. M. (2016). Expectancy effects in feedback processing are explained primarily by time-frequency delta not theta. *Manuscript Submitted*.

Massey, A. T., Bachman, M. D., Bernat, E. M., Tootell, A. V., Ellis, J. S., Kothur, S., & Eckrish, S. (2015). Expectancy effects in feedback processing are explained primarily by time-frequency delta not theta. *Journal of Cognitive Neuroscience*, *27*, S213.

Miltner, W. H., Braun, C. H., & Coles, M. G. (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), 788–798.

Nelson, L. D., Patrick, C. J., Collins, P., Lang, A. R., & Bernat, E. M. (2011). Alcohol impairs brain reactivity to explicit loss feedback. *Psychopharmacology*, *218*(2), 419–428. https://doi.org/10.1007/s00213-011-2323-3

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

Proudfit, G. H. (2015). The reward positivity: From basic research on reward to a biomarker for depression. *Psychophysiology*, *52*(4), 449–459.

Semlitsch, H. V., Anderer, P., Schuster, P., & Presslich, O. (1986). A solution for reliable and valid reduction of ocular artifacts, applied to the P300 ERP. *Psychophysiology*, *23*(6), 695–703.

Weinberg, A., Riesel, A., & Proudfit, G. H. (2014). Show me the Money: The impact of actual rewards and losses on the feedback negativity. *Brain and Cognition*, *87*, 134–139.

Yeung, N., & Sanfey, A. G. (2004). Independent coding of reward magnitude and valence in the human brain. *The Journal of Neuroscience*, *24*(28), 6258–6264.