Reward expectation modulates feedback-related negativity and EEG spectra

Michael X. Cohen,[a,b,*] Christian E. Elger,[a] and Charan Ranganath[b]

[a]Department of Epilepsy and Center for Life and Brain, Sigmund-Freud-Str. 25, Bonn 53105, Germany
[b]Department of Psychology and Center for Neuroscience, University of California, Davis, CA 95616, USA

The ability to evaluate outcomes of previous decisions is critical to adaptive decision-making. The feedback-related negativity (FRN) is an event-related potential (ERP) modulation that distinguishes losses from wins, but little is known about the effects of outcome probability on these ERP responses. Further, little is known about the frequency characteristics of feedback processing, for example, event-related oscillations and phase synchronizations. Here, we report an EEG experiment designed to address these issues. Subjects engaged in a probabilistic reinforcement learning task in which we manipulated, across blocks, the probability of winning and losing to each of two possible decision options. Behaviorally, all subjects quickly adapted their decision-making to maximize rewards. ERP analyses revealed that the probability of reward modulated neural responses to wins, but not to losses. This was seen both across blocks as well as within blocks, as learning progressed. Frequency decomposition via complex wavelets revealed that EEG responses to losses, compared to wins, were associated with enhanced power and phase coherence in the theta frequency band. As in the ERP analyses, power and phase coherence values following wins but not losses were modulated by reward probability. Some findings between ERP and frequency analyses diverged, suggesting that these analytic approaches provide complementary insights into neural processing. These findings suggest that the neural mechanisms of feedback processing may differ between wins and losses.
© 2006 Elsevier Inc. All rights reserved.

Keywords: Reward prediction error; ERP; Decision-making; EEG oscillations; Reinforcement learning

Introduction

To optimize behavior, organisms must evaluate outcomes of their actions, and use these evaluations to guide decision-making. The neural mechanisms of feedback evaluation are receiving increasing attention in cognitive neuroscience. In particular, researchers using event-related potentials (ERPs) have identified a component of the feedback-locked ERP that is sensitive to the valence of the feedback. This feedback-related negativity (FRN) is a relatively negative deflection in the ERP following losses or error feedback compared to wins or positive feedback. The FRN peaks at around 300 ms and is maximal at fronto-central scalp electrode sites (Hajcak et al., 2005; Holroyd et al., 2003; Yasuda et al., 2004). Convergent findings from source modeling, fMRI, and single-unit recording studies suggest that the FRN is generated in the medial frontal cortex, and probably in the anterior cingulate cortex (Amiez et al., 2005; Brown and Braver, 2005; Mars et al., 2005; Miltner et al., 2003; Niki and Watanabe, 1979; Paulus et al., 2004; Ridderinkhof et al., 2004; Shidara and Richmond, 2002; Tsujimoto et al., 2006; van Schie et al., 2004; Williams et al., 2004). Topographically and functionally similar feedback-locked ERP modulations have been called the medial frontal negativity and feedback error-related negativity (Gehring and Willoughby, 2002; Holroyd et al., 2003). These effects also share many similarities with the error-related negativity (ERN), a negative-going mid-frontally distributed potential elicited by erroneous responses on speeded response tasks. These potentials are thought to reflect activation of a reinforcement learning system that rapidly evaluates outcomes of decisions to guide reward-seeking behavior (Holroyd and Coles, 2002; Nieuwenhuis et al., 2004). This system is capable of rapidly determining whether feedback is better or worse than expected, and encodes this difference between expectations and actual outcomes as a reward prediction error. The anterior cingulate cortex might use these prediction errors to improve performance due to its role in cognitive control and action monitoring (Barber and Carter, 2005; Bokura et al., 2001; Botvinick et al., 2004; Kerns et al., 2004).

Given that a reward prediction error is the difference between an expected and received reward, differences in expectations of rewards should modulate the size of prediction error signals. Single-unit recording studies in nonhuman primates suggest that this is indeed the case, with more unexpected outcomes yielding larger neural responses in midbrain dopamine neurons (Fiorillo et al., 2003). It is unclear whether the magnitude of the FRN is also modulated by reward expectation, because previous studies have yielded inconsistent findings. In two studies (Holroyd et al., 2003; Yasuda et al., 2004), the magnitude of the FRN was larger when outcomes were unexpected. In another study, no statistically

* Corresponding author. Fax: +49 228 2871 1560.
E-mail address: mcohen@ucdavis.edu (M.X. Cohen).

Available online on ScienceDirect (www.sciencedirect.com).

significant modulation was observed (Hajcak et al., 2005), although from visual inspection, it appears that the FRN was larger for unexpected than expected outcomes. Of the two studies that found a significant modulation, Yasuda and colleagues (2004) found that ERPs following both losses and wins were enhanced. In the Holroyd et al. (2003) study, however, it appears from visual inspection that only the win-related ERPs were modulated, although a statistical test of this asymmetry was not reported. We designed an experiment to investigate this issue further by examining not only how reward probability might modulate outcome-locked ERPs, but also how changes in reward expectation that occur during learning might further modulate ERPs.

Because the FRN (and ERPs in general) is measured by averaging single-trial EEG traces, this potential will not reflect oscillatory activity that varies in phase from trial-to-trial (particularly in high frequencies, such as gamma). Such event-related oscillations can be assessed using time–frequency decomposition analyses such as complex wavelet convolutions, from which one can obtain estimates of instantaneous power (i.e., energy at different frequencies) and inter-trial phase coherence (i.e., consistency of oscillation onset across trials). Recent findings using this approach have revealed novel insights into task-related cognitive processes beyond what is evident in averaged ERPs (Fell et al., 2004; Makeig et al., 2002; Salinas and Sejnowski, 2001). Although the frequency characteristics of feedback processing are largely unknown, research into the frequency characteristics of the response-related ERN (Bernat et al., 2005; Luu and Tucker, 2001; Luu et al., 2004; Trujillo and Allen, in press) suggests it reflects enhanced theta (i.e., 4–8 Hz) activity following incorrect compared to correct responses. Based on the idea that the ERN and FRN reflect similar mechanisms of monitoring and controlling behavior (Holroyd and Coles, 2002), we hypothesized that feedback processing would therefore induce increased EEG theta activity for losses compared to wins.

In the present study, we sought to investigate the effects of reward probability on ERP and oscillatory correlates of neural feedback processing. Subjects chose one of two targets on each trial, and received positive or negative feedback (±10 cents) following each choice. In blocks of 80–150 trials, we manipulated the probability of winning and losing such that subjects had to learn which of the two targets rewarded more often in order to maximize their winnings. This design allowed us to examine neural responses to winning and losing as a function of the probability of wins and losses, using both conventional ERP and time–frequency analyses.

Materials and methods

Subjects

Seventeen (6 males) subjects aged 20–30 from the University of Bonn community participated in this experiment. Subjects were paid the amount they earned in the experiment or 10 Euros per hour (whichever was higher), and typically earned around 25 Euros. Informed consent documents were signed prior to the start of the experiment, which was approved by the local ethics committee.

Experiment

On each of 1200 trials during the experiment, subjects saw two small targets on the left and right side of the screen, and had to choose one on each trial, after which they either won or lost 0.10 Euros

(see Fig 1a). Visual feedback indicating the subjects' response (a green box behind the target) was shown for 600 ms, followed by a reward (10 cents) or punishment (−10 cents) that was shown for 1000 ms. A variable, 1500- to 2500-ms inter-trial-interval separated each trial. If subjects did not select a target after 2000 ms, they lost on that trial. Self-paced rest breaks were given every 30 trials, during which time subjects saw how much money they had earned so far. Within each block of trials, the probability of wins and losses associated with each target was surreptitiously manipulated. For example, on some blocks, one target would reward with 75% probability and the other target would reward with 25% probability. We refer to the "25 target" as trials when subjects chose the target that rewarded 25% of the time, and the "75 target" as trials when subjects chose the target that rewarded 75% of the time. In a third condition, both targets rewarded with 50% probability, and we refer to all trials in these blocks as the "50 target". Probabilities were fixed across trials within blocks, which lasted 80–150 trials (block length was randomly chosen by the computer at the start of each block; see Fig. 1b for examples). Block order and length was randomized and interchanged across subjects. Subjects were instructed to choose one of the two targets on each trial, and to try to win as much money as possible. The entire experiment, including time to set up, was approximately 2.5 h.

EEG recording and analyses

EEG data were recorded at 1000 Hz (with an anti-aliasing low-pass filter set at 300 Hz) from 21 scalp electrodes spread out across the scalp, and 4 ocular (two HEOG and two VEOG) electrodes. All EEG channels were referenced to the left mastoid, and were re-referenced online to the average of the left and right mastoids by the acquisition software. Scalp channels were Fp1, Fp2, F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T7, T8, P7, P8, Cz, Fz, Pz, F9, and F10. Data were resampled to 250 Hz and bandpass filtered from .1 to 40 Hz off-line for ERP analyses, and were left unfiltered for frequency analyses. Trials containing blink or other artifacts were removed prior to averaging (mean: 11%; standard error: 4%). Data were analyzed with the EEGLAB toolbox (Delorme and Makeig, 2004) (available at http://sccn.ucsd.edu/eeglab/) in Matlab. In the present study, we separately analyzed ERPs to wins and losses, rather than the difference between losses and wins, because we were interested in testing whether reward probability might have asymmetric effects depending on the feedback valence. Statistical analyses were performed by entering average ERP values from windows of 300- to 400-ms (when loss−win difference was maximal) and 400- to 600-ms (when the probability effects were maximal) post-feedback window into repeated-measures ANOVAs using SPSS 11 software. Greenhouse–Geisser corrections were used, and adjusted degrees of freedom are reported. ERPs were averaged across a time window to minimize potential noise fluctuations in ERP waveforms. To confirm that our findings reflected modulations of the FRN instead of the P300, we conducted additional ANOVAs on the difference between the most positive and most negative peak (i.e., the "peak-to-peak" amplitude) in a 200- to 600-ms time window (e.g., Yeung and Sanfey, 2004). We included midline electrodes Fz, Cz, and Pz in the ANOVAs.

Frequency analyses

Induced spectral EEG activity was assessed by creating event-related spectral perturbations using a complex sinusoidal wavelet transform procedure as implemented in EEGLAB (Delorme and Makeig, 2004). Briefly, the procedure involves convolving the time

domain signal with a complex sinusoidal wavelet with increasing cycles as frequencies increase (Delorme and Makeig, 2004). The resulting complex signal provides an estimate of instantaneous power and phase for each time point at frequencies of 4–48 Hz. This procedure is done on each trial, and then power values are averaged across trials. Power values were normalized with respect to a −200- to 0-ms prestimulus baseline and transformed into decibel scale ($10*\log_{10}$ of the signal). Phase coherence is calculated across trials using circular variance, and phase-coherence values vary from 0 to 1, with 0 indicating completely randomized phases across trials and 1 indicating completely locked phases across trials. We used an EEG epoch window of −500 to +1000 ms from each event to ensure that edge effects would not contaminate our windows of interest, and visual inspection confirmed that edge effects did not extend into our time windows of analyses. Statistical analyses were conducted by entering average power and phase coherence values from 300 to 400 and from 400 to 600 ms post-feedback in 4- to 8-Hz (theta band) and 21- to 29-Hz (lower gamma band) windows for each condition into 2 (frequency band)×2 (feedback: win or lose)×3 (probability: 25, 50, 75) repeated measures ANOVAs. For these analyses, we focused on data from electrode Fz because topographic distributions of power and phase exhibited a fronto-central peak that was maximal around Fz in both time windows (see Figs. 5 and 6 for examples in the 300–400 ms time window). The Greenhouse–Geisser correction was used to adjust the degrees of freedom. We selected theta and gamma band activity based on previous EEG research, which suggests strong task-related activity in the theta range (Kirk and Mackay, 2003; Klimesch, 1999; Luu et al., 2004) and lower gamma range (Keil et al., 2001), and based on visual inspection of time–frequency plots, which show predominant energy in these ranges (e.g., Fig. 5c). We note that 21–29 Hz is sometimes referred to as fast or upper beta

(Steriade, 2006) and sometimes as lower gamma (Keil et al., 2001; Shibata et al., 1999); we use the latter term without intention of distinguishing lower gamma from upper beta.

Results

Behavior

Although subjects were not told about changes in probabilities of rewards, they quickly adapted their behavior to find the optimal strategy: During blocks when the right-hand target rewarded 25%, 50%, and 75% of the time, subjects selected the right-hand target on 36.1%, 53.4%, and 71.4% of trials (SEM: 2.1%, 1.4%, 1.7%), respectively (Fig. 1b). A 3-way ANOVA revealed a main effect of probability ($F_{2,32}=97.70$, $p<0.0001$), and planned comparisons of the simple effects confirmed that each condition was significantly different from each other (75>50: $t_{16}=9.22$; 75>25: $t_{16}=10.50$; 50>25: $t_{16}=8.63$; all p values<0.001). Response times, however, did not differ according to condition (496, 489, 474 ms for 75, 50, and 25 conditions, respectively; $F_{2,32}=2.00$, $p=0.15$). In Fig. 1c we display the trial-by-trial choices and reward probability blocks for two subjects to demonstrate how they adapted their behavior to changes in reward contingency.

ERP (time domain) responses to feedback

Valence effects

We observed a prominent FRN (e.g., loss–win ERPs) in each reward probability condition (see Fig. 2). This effect had a similar fronto-central topographical distribution in each condition (left panel of Fig. 3), and peaked at around 320 ms. A 3 (target: 25, 50, 75)×2
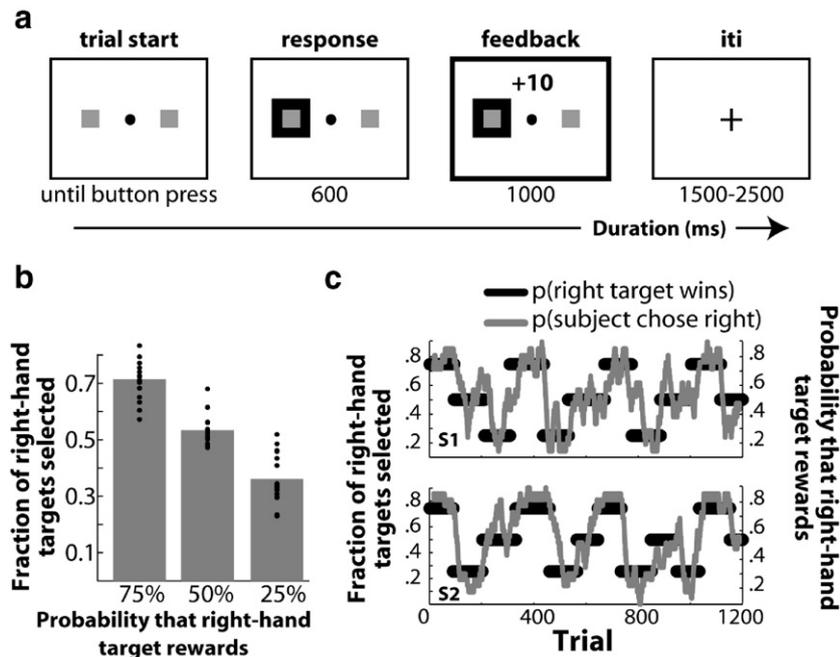
Fig. 1. Overview of experiment and behavioral results. (a) Visual display and timing of one trial. Time flows from left to right, numbers indicate the duration of each event in ms. All analyses were time-locked to the onset of the feedback stimulus (bolded box). iti=inter-trial-interval. (b) Average (gray bars) and single-subject (black dots) proportion of right-hand target selections as a function of probability manipulation. (c) Behavioral results from two subjects. The *x*-axis represents trials during the experiment and the *y*-axis represents the probability that the right-hand target rewards and the local fraction of right-hand responses selected by the subject. Thick black lines represent the reward probability and thinner gray lines represent the local fraction of right-hand responses (choices smoothed with a 20-trial kernel running average filter).
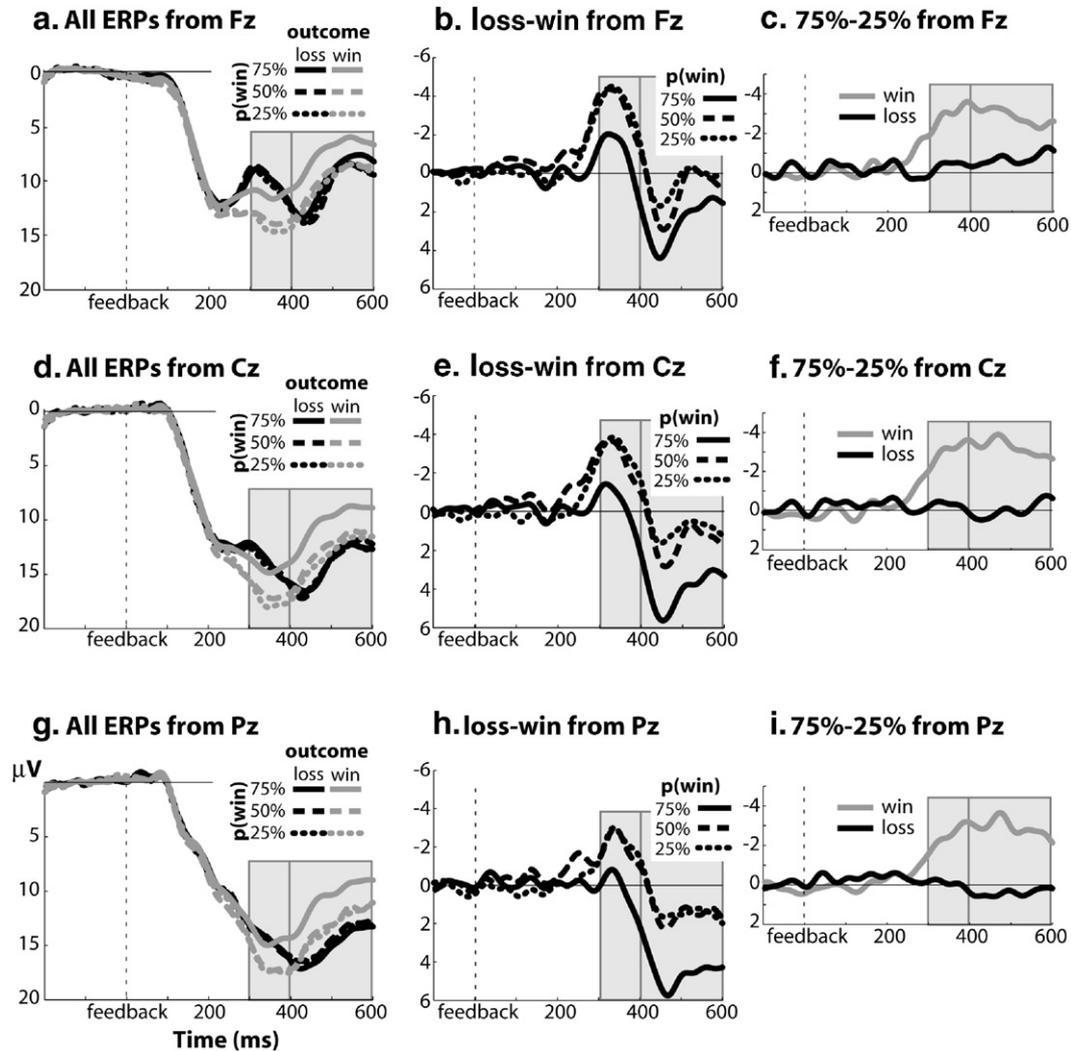
Fig. 2. Grand-averaged ERP responses to feedback. Left column depicts all ERPs following losses (black lines) and wins (gray lines) as a function of probability for electrodes Fz (a), Cz (d) and Pz (g). Middle column displays FRN (i.e., the difference wave for losses compared to wins) for each probability condition for electrodes Fz, Cz, and Pz. Right column displays the probability effect (75 target − 25 target) separately for losses and wins. Light gray bars indicate the windows used for statistical analyses. Note that b, c, e, f, h, and i have the same y-axis scaling, and a, d, and g have a different y-axis scaling.

(feedback: win, loss) repeated-measures ANOVA on ERP responses from a 300- to 400-ms post-feedback window revealed a highly significant effect of feedback ($F_{1,16}$=17.18, $p<0.001$), which was larger at Fz than at Cz or Pz ($F_{1,8,14}$=15.9, $p<0.001$). These effects were also statistically significant when we analyzed the peak-to-peak amplitude differences (see Materials and methods). In a later time window of 400–600 ms, there was a significant feedback valence effect ($F_{1,16}$=9.23, p=0.003), although it was in the opposite direction (i.e., loss-related ERPs were more positive than were win-related ERPs; see Figs. 2 and 3). There was also a valence effect × electrode site interaction ($F_{1.5,24.3}$=8.88, p=0.003). Follow-up analyses revealed that, in this time window, the valence effect was largest at Pz, and was not significant at Fz ($F_{1,16}$=3.66, p=0.075).

Reward probability effects

From 300 to 400 ms, there was a significant main effect of probability ($F_{1.6,26.8}$=7.68, p=0.004), although it is qualified by a feedback valence × probability interaction ($F_{1.9,31.5}$=12.11, p=0.001). Specifically, ERPs following losses were not modulated

by probability ($F_{1.7,27.9}$=0.29, p=0.716), whereas ERPs following wins were significantly affected by probability with a 25>50>75 profile ($F_{1.8,29.2}$=22.28, $p<0.001$). As seen in Fig. 2a, d, and g, this interaction was driven by the response to 75 target wins. All of these effects were also statistically significant when we analyzed the peak-to-peak amplitude differences. Additional ANOVAs confirmed that the probability × feedback valence interaction extended later in time (400–600 ms window: $F_{1.9,12.1}$=6.31, p=0.005). There were no interactions involving electrode site and probability in either time window (all p values>0.15).

ERP changes during learning

The previous set of analyses focused on changes in ERPs across blocks. In our next set of analyses, we examined how feedback potentials might change within blocks, as learning progressed. To do this, we quantified differences in the ERPs between the first third and the last third of trials in each block, when subjects were learning and had learned, respectively, the changes in optimal response patterns. As seen in Fig. 4a–b, ERP responses to rewards, but not to losses,
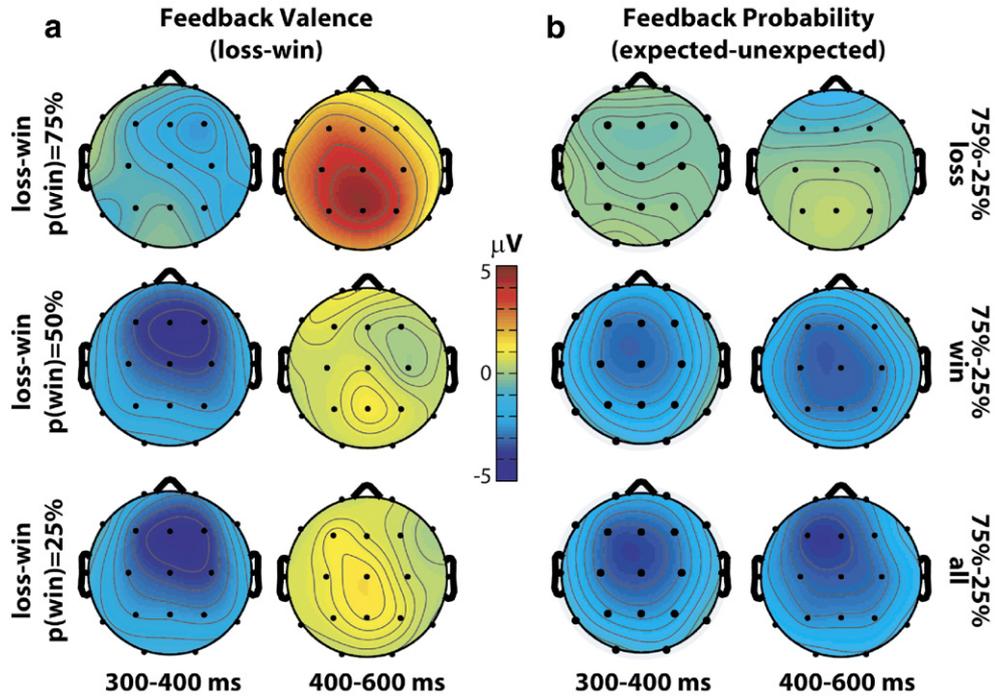
Fig. 3. Topographical distribution maps taken from a 300- to 400-ms post-feedback window. (a) Maps of the feedback valence effect (loss–win) as a function of reward probability during the 300–400 ms window (left column) and 400–600 ms window (right column). (b) The probability effect (75% vs. 25% probability of reward) as following losses (top), wins (middle), and losses and wins averaged together (bottom). Topographic differences for 75%–25% are displayed, but all probability differences (e.g., 75%–50%) for wins have nearly identical topographic distributions.
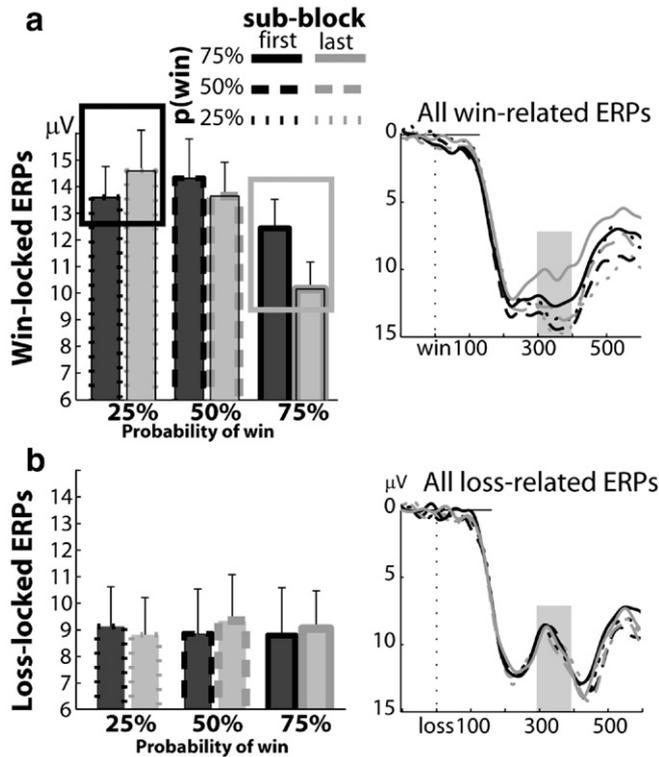
Fig. 4. Changes in ERP responses to feedback during learning. ERPs were separated according to the probability condition and epoch (first vs. last third of trials) within blocks for wins (top panel) and losses (bottom panel). On the right are all ERPs, and on the left are bar graphs showing averaged ERP responses from 300 to 400 ms (light gray bar in ERP plots). The interaction ("learning effect") is highlighted with boxes. In the 25 target condition, as subjects learn that rewards are infrequent, ERP responses to rewards become more positive (black box). Conversely, in the 75 target condition, as subjects learn that rewards are likely, ERP responses to rewards become less positive (gray box). This pattern of effects was not seen for losses (bottom panel).

changed with learning. Specifically, ERP amplitudes became *less* positive as subjects learned that rewards become more likely for the 75 target. Conversely, ERP responses became *more* positive as subjects learned that rewards become less likely for the 25 target. To quantify this effect, we conducted a 3 [target: 25, 50, 75] × 2 [time: first vs. last third of block] × 3 [electrode: Fz, Cz, Pz] ANOVA on win responses, which revealed a significant target × time interaction ($F_{1.5,23.8} = 3.84$, $p = 0.04$). The interaction (see top panel of Fig. 4a) is especially striking when comparing ERPs elicited by the 25 target at the last third of the block (i.e., when wins are the most unexpected) and those elicited by the 75 target at the last third of the block (i.e., when wins are the most expected). A direct contrast between these conditions revealed that ERPs were significantly more positive in the 25 target condition as compared with the 75 target condition (14.46 and 10.42 μV, respectively, at Fz; post-hoc *t*-test: $t_{16} = 4.74$, $p < 0.01$). The learning effect was specific for wins: As shown in Fig. 4b, responses to losses did not differ as a function of learning ($F_{1.8,30.1} = 0.841$, $p = 0.435$). Even when comparing losses that presumably were the least and most expected (75 target at last third and 25 target at last third, respectively), ERPs were of nearly identical magnitude (9.38 and 9.14 μV at Fz, $t_{16} = 0.31$). Thus, learning about reward probabilities within each block was associated with changes in ERPs following wins, whereas learning did not modulate ERPs following losses.

Because the overall probability effect peaked later in time (see Fig. 2), we additionally conducted ANOVAs on ERPs from a 400- to 600-ms time window. The learning effects continued to be significant following wins ($F_{1.5,24.7} = 8.45$, $p = 0.003$), but not following losses ($F_{1.8,29.3} = 0.20$, $p = 0.79$). This effect was not significant when using peak-to-peak amplitude. There were no significant interactions involving electrode site in either time window (all *p* values > 0.11).

Frequency decomposition of EEG

Effects of feedback valence

Figs. 5a and 6a display the loss–win spectrograms, topographical maps, and bar graphs for power and phase, respectively. As shown in these figures, feedback valence modulated power and phase values differently in the two frequency bands. Specifically, for power values, there was a significant frequency band × valence cross-over interaction from 300 to 400 ms in Fz ($F_{1,16} = 37.94$, $p < 0.001$), such that in the theta band, there was more power following losses than wins ($F_{1,16} = 28.65$, $p < 0.001$), whereas in the gamma band, there was more power following wins than losses ($F_{1,16} = 4.91$, $p = 0.042$) (see bar graphs in Fig. 5a–b). ANOVAs from a time window of 400–600 ms revealed that the loss > win effect remained significant ($F_{1,16} = 36.66$, $p < 0.001$). For phase coherence values, we found a significant frequency band × valence interaction ($F_{1,16} = 8.47$, $p = 0.010$), such that in the theta band there was greater phase coherence following losses compared to wins ($F_{1,16} = 8.84$, $p = 0.009$), but in the gamma band there was no effect of feedback valence ($F_{1,16} = 0.44$, $p = 0.516$). ANOVAs from a time window of 400–600 ms revealed no effects of feedback valence (all *p* values > 0.3). We additionally plotted the time course of power and phase coherence separately for each condition in Fig. 7. For illustrative purposes, we calculated a *t*-test of the difference between theta power following losses versus that following wins at each time point, and time points with a *t*-value of $p < 0.01$ have a bolded *x*-axis.

Effects of reward probability

In our next set of analyses, we tested whether feedback probability affected power and phase values. These effects largely paralleled those of the time domain, such that reward probability modulated EEG power following wins but not following losses. Specifically, for power values, there was a significant valence × probability interaction
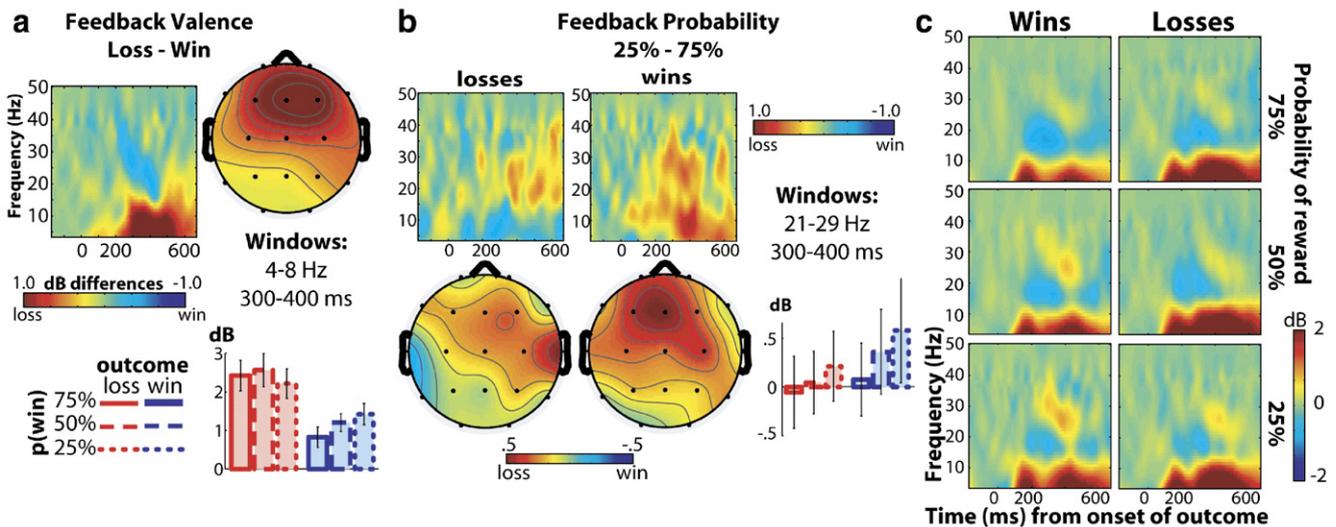
Fig. 5. Effects of feedback valence and probability on frequency power. (a) Difference in power between losses and wins. The time–frequency map shows the subtraction of power values of losses from wins over time (*x*-axis; 0 is onset of feedback stimulus) and frequency (*y*-axis). Red colors indicate more power for losses and blue colors indicate more power for wins. Topographical map and bar graphs show data taken from a 300- to 400-ms, 4- to 8-Hz window. (b) Effect of reward probability on power values. Time–frequency and topographical distributions are separately plotted for losses (left) and wins (right) from a 300- to 400-ms, 21- to 29-Hz window. (c) Time–frequency plots separately for all combinations of feedback valence (columns) and reward probability (rows). All data are taken from electrode Fz, where the effects were maximal.
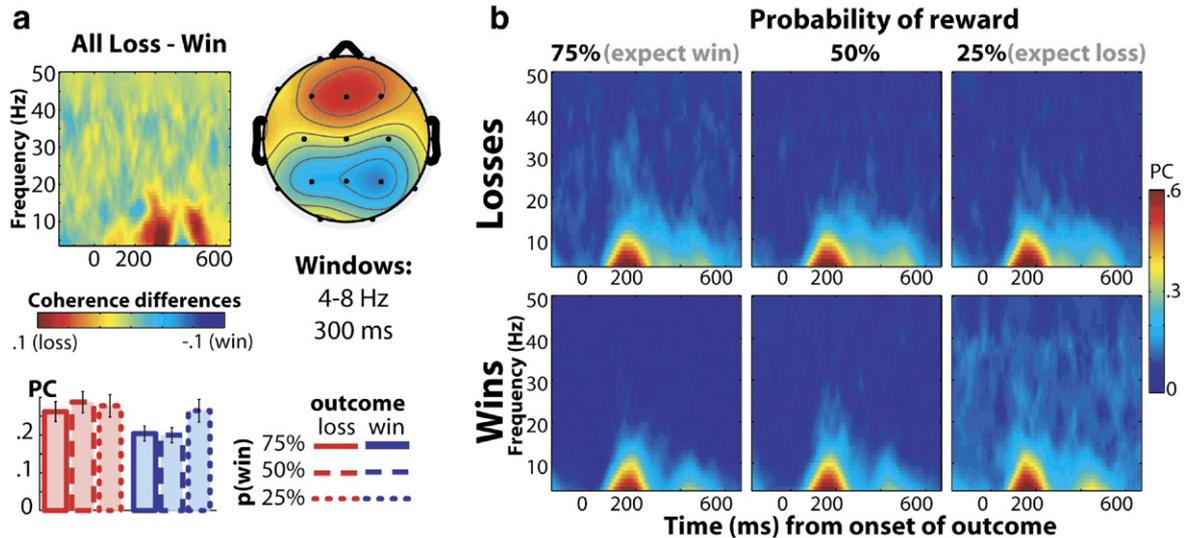
Fig. 6. Effects of feedback valence and probability on inter-trial phase coherence. (a) Difference in phase coherence values between losses and wins. The time–frequency map shows the subtraction of coherence values of losses from those of wins over time (x-axis; 0 is onset of feedback stimulus) and frequency (y-axis). Red colors indicate enhanced coherence for losses and blue colors indicate enhanced coherence for wins. Topographical map and bar graphs show data taken from a 300- to 400-ms, 4- to 8-Hz window. (b) Time–frequency plots separately for all combinations of feedback valence (rows) and reward probability (columns). All data are taken from electrode Fz.

($F_{1.2,20.3}=4.03$, $p=0.050$), such that power values following wins increased with decreasing reward probability ($F_{1.5,24.5}=6.99$, $p=0.007$), but they were not significantly altered following losses ($F_{1.5,25.3}=0.468$, $p=0.587$). There were no interactions involving probability and frequency band (p values >0.3), suggesting that reward probability did not differentially affect power in the two frequency bands. ANOVAs from a time window of 400–600 ms revealed the same pattern of results (see Table 1 for statistical values).

For phase coherence values, we found a significant valence × probability interaction ($F_{1.7,27.1}=14.67$, $p=0.020$), such that phase coherence following wins increased with decreasing reward probability ($F_{1.4,23.0}=22.89$, $p>0.001$), whereas phase coherence following losses was unaffected by reward probability ($F_{1.8,29.8}=1.36$, $p=0.270$). As with power, there were no interactions involving probability and frequency band interaction (p values >0.5), suggesting that the impact of probability was not different in the two frequency bands. ANOVAs from a time window of 400–
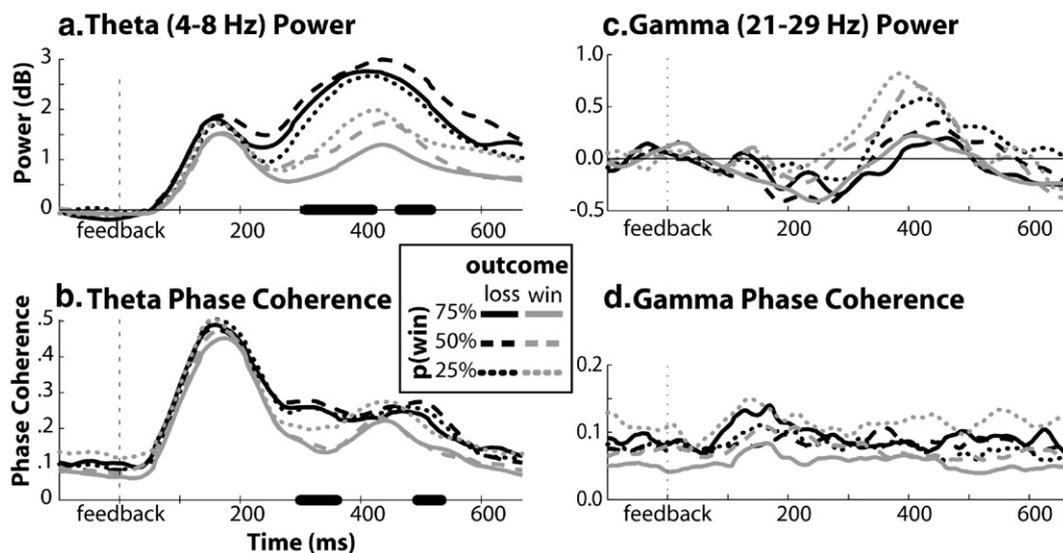
Fig. 7. Time course of power (top panel) and phase coherence (bottom) in theta (left) and gamma (right) frequency bands, plotted separately for combinations of feedback valence and probability. A paired-sample t-test between values for losses and wins (averaging over probability condition) at each time point was run, and values that were significant at $p<0.01$ have bolded x-axes. Full ANOVAs of these differences are reported in the text. Data are taken from electrode Fz.

Table 1
Summary of effects

Effect	Dependent variable	Time window (ms)	F-value	p-value	Direction
Feedback	ERP	300–400	17.18[a]	<0.001[a]	Win>Loss
valence	ERP	400–600	3.66	0.075	Win>Loss
Feedback	Power	300–400	28.65[a]	<0.001[a]	Loss>Win
valence:	Power	400–600	36.66[a]	<0.001[a]	Loss>Win
theta	Phase	300–400	8.84[a]	0.009[a]	Loss>Win
	Phase	400–600	6.51[a]	0.021[a]	Loss>Win
Feedback	Power	300–400	4.91[a]	0.042[a]	Win>Loss
valence:	Power	400–600			
gamma	Phase	300–400			
	Phase	400–600			
Reward	ERP	300–400	22.28[a]	<0.001[a]	25>50>75
probability:	ERP	400–600	21.77[a]	<0.001[a]	25>50>75
wins	Power	300–400	6.99[a]	0.007[a]	25>50>75
	Power	400–600	6.52[a]	0.006[a]	25>50>75
	Phase	300–400	22.89[a]	<0.001[a]	25>50>75
	Phase	400–600	24.86[a]	<0.001[a]	25>50>75
Reward	ERP	300–400			
probability:	ERP	400–600			
losses	Power	300–400			
	Power	400–600	3.53[a]	0.043[a]	50>75>25
	Phase	300–400			
	Phase	400–600			
Learning:	ERP	300–400	3.84[a]	0.040[a, b]	25_3>25_1; 75_1>73_3
wins	ERP	400–600	8.45[a]	0.003[a]	25_3>25_1; 75_1>73_3
	Power (theta)	300–400			
	Power (theta)	400–600			
	Phase (theta)	300–400			
	Phase (theta)	400–600			
Learning:	ERP	300–400			
losses	ERP	400–600			
	Power (theta)	300–400			
	Power (theta)	400–600			
	Phase (theta)	300–400	2.06	0.152	
	Phase (theta)	400–600	2.02	0.155	

Subscripts indicate position within block (first third or last third). See text for degrees of freedom.

[a] Statistically significant effects.

[b] Results from all ERP analyses in the 300–400 ms window were also significant when using peak-to-peak measurements except this one. Only statistics from Fz are reported. Entries with an F-value less than 2 are omitted.

600 ms revealed the same pattern of findings (statistical values in Table 1).

Learning effects

Finally, we examined whether there were changes in power or phase as a function of learning. There were no main effects or interactions involving position within block for wins or losses

Table 2
Correlation coefficients among average ERP amplitude- and frequency-domain effects from 300 to 400 ms in Fz

	Loss			Win		
	75%	50%	25%	75%	50%	25%
Theta power	−0.06	−0.21	−0.11	0.27	0.28	−0.01
Theta phase	−0.11	−0.04	0.07	0.35	0.34	0.27
Gamma power	0.65**	0.50*	0.72**	0.62**	0.55*	0.72**
Gamma phase	0.32	−0.03	0.09	−0.29	−0.15	−0.04

Correlation coefficient is statistically significant at *p<.05, **p<0.01.

for power or phase values in theta or gamma bands (all F values<2).

Comparison between time- and frequency-domain effects

To more formally compare information from ERPs and frequency decomposition, we conducted two analyses to correlate the magnitude of the ERPs with the average power and phase coherence values from theta and gamma bands. First, we correlated the average ERP magnitudes from 300 to 400 ms from Fz with theta and gamma power and phase coherence values from the same time window. As seen in Table 2, there were significant correlations between ERPs and power in the gamma band in all conditions, but no significant correlations between ERPs and theta band activity. However, increased power could simultaneously decrease the magnitude of ERPs in the trough of the wave and increase the magnitude of the peak of the wave, thus causing no change in the average amplitude in a time window. Thus, in a second analysis, we correlated power and phase coherence values with the peak-to-peak amplitude measure (see Materials and methods). In this analysis, a reverse pattern of findings was observed: ERPs were strongly correlated with theta power and phase coherence, but not with gamma activity (Table 3). Interestingly, peak-to-peak measures of ERPs in nearly all conditions were significantly correlated with theta phase coherence, but were strongly correlated with theta power only following losses. Finally, there were no significant correlations among theta and gamma power and phase values in any condition.

Discussion

In the present study, we examined whether and how expectations of rewards and losses affected ERP and oscillatory correlates of feedback processing. We found that ERPs, theta, and gamma activity following wins, but not losses, were modulated by the

Table 3
Correlation coefficients among ERP peak-to-peak differences from 200 to 600 ms and theta and gamma band activity from 300 to 400 ms average in electrode Fz

	Loss			Win		
	75	50	25	75	50	25
Theta power	0.72**	0.61**	0.70**	0.09	0.14	0.19
Theta phase	0.76**	0.82**	0.84**	0.61**	0.68**	0.36
Gamma power	0.15	0.32	0.45	0.13	0.23	0.55*
Gamma phase	0.12	0.23	0.27	−0.24	−0.29	−0.32

Correlation coefficient is statistically significant at *p<.05, **p<0.01.

feedback probability manipulation. This was seen both across and within (i.e., learning effects) blocks of trials. Additionally, we found enhanced power and cross-trial phase coherence in the theta frequency band (4–8 Hz) for losses compared to wins, but opposite effects in the lower gamma frequency band (21–29 Hz). Together, these findings provide strong evidence for a valence-related asymmetry in EEG correlates of feedback processing.

Win-related, but not loss-related, ERPs were modulated by feedback probability

We observed modulations in the magnitude of the loss–win ERP difference wave according to the probability of reward vs. loss feedback. Closer inspection revealed that this effect was driven by win-related ERPs. Indeed, loss-related ERPs were unaffected by the probability manipulation. This interaction was present regardless of analysis approach (average amplitude values or peak-to-peak amplitude differences) or time window. Within the win conditions, ERPs following 50 and 25 target wins differed only slightly, whereas ERPs following the 75 target wins showed a larger and more sustained effect. This may have been due to subjects overestimating the 25 target condition as being closer to the 50 target condition. This interpretation is also consistent with the asymmetry in behavioral responses: When the right-hand target rewarded 25% of the time, subjects chose that target on 36% of trials (compare with 53% and 71% when the right-hand target rewarded 50% and 75% of the time).

To our knowledge, no previous study has reported asymmetric effects of reward probability on feedback ERPs. However, close examination suggests that a similar effect might have been present in previous studies. For example, from visual inspection, it appears that Holroyd et al. (2003) (Fig. 1) found modulations of win-related, but not loss-related, ERPs, although this asymmetry was not tested. Such an asymmetry can also be seen in two recent studies by Hajcak and colleagues. In the first (Hajcak et al., 2005), it appears that the FRN was larger following unexpected compared to expected outcomes (Figs. 2 and 4), although these effects were reported to be nonsignificant. In the second study (Hajcak et al., 2006), in which they manipulated the size (rather than probability) of rewards and losses, differently sized losses (5 vs. 25 cents) yielded indistinguishable ERPs, but larger wins elicited larger ERPs than smaller wins. Here as well the effects were visually observable but not statistically significant. It is possible that there were not enough trials to statistically identify these modulations. Finally, in a study by Yasuda et al. (2004), ERPs following both wins and losses appear to be modulated by the level of expectation.

A more sensitive measure of expectation-related modulations in the ERPs comes from inspection of changes in ERPs that occur as learning progresses (i.e., differences between ERPs during the first and last third of trials in each block). We found that changes in win-related ERPs were consistent with predictions of reinforcement learning theory. Specifically, in reinforcement learning theory, reward prediction errors should increase or decrease as rewards become less or more expected, respectively. Consistent with this prediction, we found that in the 25 target condition, as subjects learned that rewards are relatively infrequent, ERP magnitudes became more positive, whereas in the 75 target condition, as subjects learned that rewards are likely, ERP magnitudes became less positive. This pattern of changes in the magnitude of neural feedback responses is similar to results showing that the magnitude of the FRN decreased over the course

of probabilistic learning experiments (Mars et al., 2005; Nieuwenhuis et al., 2002). In the Nieuwenhuis et al. (2002) study, for example, subjects learned to select particular stimuli that were associated with either 20%, 50%, 80%, or 100% probability of reward. Consistent with our findings, they found that FRN magnitudes decreased as a function of reward probability across the experiment. Similar results were obtained by Holroyd and Coles (2002). Given that this modulation was specific to wins in our study, it is possible that the changes in FRN magnitude in the studies of Nieuwenhuis et al. (2002) and Holroyd and Coles (2002) were driven solely by responses to wins.

One might wonder whether our ERP effects are more related to the P300, a complex of ERP modulations that has been related to stimulus probability, subjective confidence, and target detection (Donchin and Coles, 1988; Johnson, 1993; Kugler et al., 1993; Polich and Kok, 1995; Squires et al., 1976). Although the later temporal peak of the probability effect is consistent with this speculation, several considerations suggest that the probability-related ERP modulations are more closely related to the FRN than to the P300. Specifically, if the effects were P300 modulations, one would make at least two predictions: First, the P300 should be larger following all low-probability events, independent of the valence of the outcome (Yeung and Sanfey, 2004). This might reflect stimulus probability estimations and/or a context-updating process (Donchin and Coles, 1988; Johnson, 1986; Johnson and Donchin, 1980). Second, the effects should have a posterior spatial distribution with a peak around Pz. In contrast to these predictions, the probability manipulation affected only win-related ERPs, and there were no probability effects on responses to losses (see Fig. 2), regardless of the time window, measurement approach (average ERP amplitude or peak-to-peak difference), or data domain (time vs. frequency). Further, the probability effects we observed had an anterior topography (see Fig. 3). Indeed, the probability effects were numerically smaller at Pz than at electrodes Fz and Cz. Even from 400 to 600 ms, when the probability effects were largest in time, the effects were smaller at Pz than at Cz (see Fig. 3). These observations, combined with previous dissociations of the FRN and P300 (Sato et al., 2005; Yeung and Sanfey, 2004), suggest that our findings are more related to the FRN than the P300.

Oscillations in feedback-locked potentials

Medial frontal regions including the anterior cingulate exhibit large theta oscillations that are present at rest, during sleep, and during cognitive tasks, and are modulated by task demands such as working memory, attention, and action selection (Basar-Eroglu et al., 1992; Basar-Eroglu and Demiralp, 2001; Ishii et al., 1999; Kubota et al., 2001; Onton et al., 2005; Tsujimoto et al., 2006; Wang et al., 2005). Little is known about frequency characteristics of feedback processing, but researchers have investigated the frequency characteristics of response error processing. In these cases, response errors, compared with correct responses, is associated with enhanced theta power (Gevins et al., 1997; Luu et al., 2004; Trujillo and Allen, in press; Yordanova et al., 2004). To the extent that similar mechanisms underlie processing of error responses and error feedback (Holroyd and Coles, 2002), these findings are consistent with our finding of enhanced theta power and phase-locking following losses compared to wins. More directly relevant to our findings, Luu and Tucker (2001) observed an FRN in response to error feedback after band-pass filtering

their data in the theta range, suggesting that much of the energy in the FRN comes from theta frequencies.

We also found that power and phase coherence were modulated by the probability of rewards. Consistent with what we observed in the ERPs, following wins, theta power and phase coherence values tracked the probability of rewards, such that the less likely the reward, the larger the power and phase coherence values. Following losses, however, no such modulation was observed. Previous studies have reported that decreased stimulus probability increases theta power at around 300 ms (Spencer and Polich, 1999; Yordanova et al., 2000). However, these previous experiments used tones that did contain any valence or motivational value, and, to our knowledge, an asymmetry between responses to losses and to wins has not been previously reported.

Correspondence between time domain and frequency analyses

Many analyses revealed similar patterns of results across the time and frequency domains. For example, the probability manipulation affected neural responses to wins, but not to losses, in ERPs, theta power, and phase coherence. We also observed some important differences between the time and frequency domain results. Most notably, we observed a cross-over interaction between feedback valence in theta and gamma bands, such that there was more theta power following losses compared to wins, but more gamma power following wins compared to losses. Lower gamma has been associated with thalamocortical and cortico-cortical coupling (Steriade, 2006), suggesting enhanced connectivity during wins compared to losses. On the other hand, the changes in ERPs as a function of learning were not observed in the frequency domain. Thus, time- and frequency-domain analyses provided some overlapping, but some independent, insights into task-induced EEG activity.

When directly correlating ERPs and frequency information, we found differences in how the ERPs related to gamma and theta according to our ERP measurement approach. Specifically, average ERP amplitudes were correlated with gamma power but not with theta activity, whereas peak-to-peak amplitudes were correlated with theta activity but not with gamma. This finding, the lack of correlations between theta and gamma activity, and the cross-over interaction between the loss–win effect in theta and gamma bands together suggest that independent neural processes drive theta and gamma activity in this task.

The present findings raise the question of whether the FRN (and ERPs in general), arise from bursts of increased neural activity or from phase-locking of ongoing oscillations. The oscillatory nature of loss-related ERPs is apparent from visual inspection of the time domain ERPs, both in our study and in other FRN and response-locked error negativity papers. Conceptualizing the loss-related ERP as a theta oscillation also provides a parsimonious interpretation for the fact that there are three to four peaks of the loss-related ERP (around 200, 300, 420, and possibly 550 ms in Fig. 2): If the neural response to losses is conceptualized as a phase-locked theta oscillation, these peaks may reflect the same neural process. Finally, the correlations between frequency information and ERP measures suggest that whereas all feedback-related ERPs comprise phase-locking in the theta band and enhanced gamma power, only responses to losses are additionally driven by power in the theta band. One caution, however, is that any signal, even if there are no oscillations present, can be perfectly represented in frequency space, given fine enough sampling between frequencies. Yeung et al. (2004) provide an example of this applied to EEG data. They created an artificial EEG dataset that contained only activity bursts and no oscillations. When the data were transformed into frequency space using complex sinusoidal wavelets, enhanced power in the theta range was observed following "error" compared to "correct" trials. Thus, although we feel that loss-related ERPs may be more easily conceptualized as a theta wave than as a series of phasic neural bursts, our findings do not provide irrefutable evidence either way (see also Mazaheri and Jensen, 2006).

In conclusion, we have provided strong evidence for an asymmetry between EEG responses to losses and wins. These findings additionally suggest that future research might benefit from considering win and loss responses separately, rather than considering their differences as a single measure.

References

Amiez, C., Joseph, J.P., Procyk, E., 2005. Anterior cingulate error-related activity is modulated by predicted reward. Eur. J. Neurosci. 21, 3447–3452.

Barber, A.D., Carter, C.S., 2005. Cognitive control involved in overcoming prepotent response tendencies and switching between tasks. Cereb. Cortex 15, 899–912.

Basar-Eroglu, C., Demiralp, T., 2001. Event-related theta oscillations: an integrative and comparative approach in the human and animal brain. Int. J. Psychophysiol. 39, 167–195.

Basar-Eroglu, C., Basar, E., Demiralp, T., Schurmann, M., 1992. P300-response: possible psychophysiological correlates in delta and theta frequency channels. A review. Int. J. Psychophysiol. 13, 161–179.

Bernat, E.M., Williams, W.J., Gehring, W.J., 2005. Decomposing ERP time–frequency energy using PCA. Clin. Neurophysiol. 116, 1314–1334.

Bokura, H., Yamaguchi, S., Kobayashi, S., 2001. Electrophysiological correlates for response inhibition in a Go/NoGo task. Clin. Neurophysiol. 112, 2224–2232.

Botvinick, M.M., Cohen, J.D., Carter, C.S., 2004. Conflict monitoring and anterior cingulate cortex: an update. Trends Cogn. Sci. 8, 539–546.

Brown, J.W., Braver, T.S., 2005. Learned predictions of error likelihood in the anterior cingulate cortex. Science 307, 1118–1121.

Delorme, A., Makeig, S., 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J. Neurosci. Methods 134, 9–21.

Donchin, E., Coles, M.G., 1988. Is the P300 component a manifestation of context updating? Behav. Brain Sci. 11, 357–427.

Fell, J., Dietl, T., Grunwald, T., Kurthen, M., Klaver, P., Trautner, P., Schaller, C., Elger, C.E., Fernandez, G., 2004. Neural bases of cognitive ERPs: more than phase reset. J. Cogn. Neurosci. 16, 1595–1604.

Fiorillo, C.D., Tobler, P., Schultz, N., 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299, 1898–1902.

Gehring, W.J., Willoughby, A.R., 2002. The medial frontal cortex and the rapid processing of monetary gains and losses. Science 295, 2279–2282.

Gevins, A., Smith, M.E., McEvoy, L., Yu, D., 1997. High-resolution EEG mapping of cortical activation related to working memory: effects of task difficulty, type of processing, and practice. Cereb. Cortex 7, 374–385.

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, 161–170.

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. Biol. Psychol. 71, 148–154.

Holroyd, C.B., Coles, M.G., 2002. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. Psychol. Rev. 109, 679–709.

Holroyd, C.B., Nieuwenhuis, S., Yeung, N., Cohen, J.D., 2003. Errors in reward prediction are reflected in the event-related brain potential. NeuroReport 14, 2481–2484.

Ishii, R., Shinosaki, K., Ukai, S., Inouye, T., Ishihara, T., Yoshimine, T., Hirabuki, N., Asada, H., Kihara, T., Robinson, S.E., Takeda, M., 1999. Medial prefrontal cortex generates frontal midline theta rhythm. NeuroReport 10, 675–679.

Johnson Jr., R., 1986. A triarchic model of P300 amplitude. Psychophysiology 23, 367–384.

Johnson Jr., R., 1993. On the neural generators of the P300 component of the event-related potential. Psychophysiology 30, 90–97.

Johnson Jr., R., Donchin, E., 1980. P300 and stimulus categorization: two plus one is not so different from one plus one. Psychophysiology 17, 167–178.

Keil, A., Muller, M.M., Gruber, T., Wienbruch, C., Elbert, T., 2001. Human large-scale oscillatory brain activity during an operant shaping procedure. Brain Res. Cogn. Brain Res. 12, 397–407.

Kerns, J.G., Cohen, J.D., MacDonald III, A.W., Cho, R.Y., Stenger, V.A., Carter, C.S., 2004. Anterior cingulate conflict monitoring and adjustments in control. Science 303, 1023–1026.

Kirk, I.J., Mackay, J.C., 2003. The role of theta-range oscillations in synchronising and integrating activity in distributed mnemonic networks. Cortex 39, 993–1008.

Klimesch, W., 1999. EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. Brain Res. Brain Res. Rev. 29, 169–195.

Kubota, Y., Sato, W., Toichi, M., Murai, T., Okada, T., Hayashi, A., Sengoku, A., 2001. Frontal midline theta rhythm is correlated with cardiac autonomic activities during the performance of an attention demanding meditation procedure. Brain Res. Cogn. Brain Res. 11, 281–287.

Kugler, C.F., Taghavy, A., Platt, D., 1993. The event-related P300 potential analysis of cognitive human brain aging: a review. Gerontology 39, 280–303.

Luu, P., Tucker, D.M., 2001. Regulating action: alternating activation of midline frontal and motor cortical networks. Clin. Neurophysiol. 112, 1295–1306.

Luu, P., Tucker, D.M., Makeig, S., 2004. Frontal midline theta and the error-related negativity: neurophysiological mechanisms of action regulation. Clin. Neurophysiol. 115, 1821–1835.

Makeig, S., Westerfield, M., Jung, T.P., Enghoff, S., Townsend, J., Courchesne, E., Sejnowski, T.J., 2002. Dynamic brain sources of visual evoked responses. Science 295, 690–694.

Mars, R.B., Coles, M.G., Grol, M.J., Holroyd, C.B., Nieuwenhuis, S., Hulstijn, W., Toni, I., 2005. Neural dynamics of error processing in medial frontal cortex. NeuroImage 28, 1007–1013.

Mazaheri, A., Jensen, O., 2006. Posterior alpha activity is not phase-reset by visual stimuli. Proc. Natl. Acad. Sci. U. S. A. 103, 2948–2952.

Miltner, W.H., Lemke, U., Weiss, T., Holroyd, C., Scheffers, M.K., Coles, M.G., 2003. Implementation of error-processing in the human anterior cingulate cortex: a source analysis of the magnetic equivalent of the error-related negativity. Biol. Psychol. 64, 157–166.

Nieuwenhuis, S., Ridderinkhof, K.R., Talsma, D., Coles, M.G., Holroyd, C.B., Kok, A., van der Molen, M.W., 2002. A computational account of

altered error processing in older age: dopamine and the error-related negativity. Cogn. Affect. Behav. Neurosci. 2, 19–36.

Nieuwenhuis, S., Holroyd, C.B., Mol, N., Coles, M.G., 2004. Reinforcement-related brain potentials from medial frontal cortex: origins and functional significance. Neurosci. Biobehav. Rev. 28, 441–448.

Niki, H., Watanabe, M., 1979. Prefrontal and cingulate unit activity during timing behavior in the monkey. Brain Res. 171, 213–224.

Onton, J., Delorme, A., Makeig, S., 2005. Frontal midline EEG dynamics during working memory. NeuroImage 27, 341–356.

Paulus, M.P., Feinstein, J.S., Simmons, A., Stein, M.B., 2004. Anterior cingulate activation in high trait anxious subjects is related to altered error processing during decision making. Biol. Psychiatry 55, 1179–1187.

Polich, J., Kok, A., 1995. Cognitive and biological determinants of P300: an integrative review. Biol. Psychol. 41, 103–146.

Ridderinkhof, K.R., Ullsperger, M., Crone, E.A., Nieuwenhuis, S., 2004. The role of the medial frontal cortex in cognitive control. Science 306, 443–447.

Salinas, E., Sejnowski, T.J., 2001. Correlated neuronal activity and the flow of neural information. Nat. Rev., Neurosci. 2, 539–550.

Sato, A., Yasuda, A., Ohira, H., Miyawaki, K., Nishikawa, M., Kumano, H., Kuboki, T., 2005. Effects of value and reward magnitude on feedback negativity and P300. NeuroReport 16, 407–411.

Shibata, T., Shimoyama, I., Ito, T., Abla, D., Iwasa, H., Koseki, K., Yamanouchi, N., Sato, T., Nakajima, Y., 1999. Event-related dynamics of the gamma-band oscillation in the human brain: information processing during a GO/NOGO hand movement task. Neurosci. Res. 33, 215–222.

Shidara, M., Richmond, B.J., 2002. Anterior cingulate: single neuronal signals related to degree of reward expectancy. Science 296, 1709–1711.

Spencer, K.M., Polich, J., 1999. Poststimulus EEG spectral analysis and P300: attention, task, and probability. Psychophysiology 36, 220–232.

Squires, K.C., Wickens, C., Squires, N.K., Donchin, E., 1976. The effect of stimulus sequence on the waveform of the cortical event-related potential. Science 193, 1142–1146.

Steriade, M., 2006. Grouping of brain rhythms in corticothalamic systems. Neuroscience 137, 1087–1106.

Trujillo, L.T., Allen, J., in press. Theta EEG Dynamics of the Error-Related Negativity. Clin. Neurophysiol.

Tsujimoto, T., Shimazu, H., Isomura, Y., 2006. Direct recording of theta oscillations in primate prefrontal and anterior cingulate cortices. J. Neurophysiol. 95, 2987–3000.

van Schie, H.T., Mars, R.B., Coles, M.G., Bekkering, H., 2004. Modulation of activity in medial frontal and motor cortices during error observation. Nat. Neurosci. 7, 549–554.

Wang, C., Ulbert, I., Schomer, D.L., Marinkovic, K., Halgren, E., 2005. Responses of human anterior cingulate cortex microdomains to error detection, conflict monitoring, stimulus–response mapping, familiarity, and orienting. J. Neurosci. 25, 604–613.

Williams, Z.M., Bush, G., Rauch, S.L., Cosgrove, G.R., Eskandar, E.N., 2004. Human anterior cingulate neurons and the integration of monetary reward with motor responses. Nat. Neurosci. 7, 1370–1375.

Yasuda, A., Sato, A., Miyawaki, K., Kumano, H., Kuboki, T., 2004. Error-related negativity reflects detection of negative reward prediction error. NeuroReport 15, 2561–2565.

Yeung, N., Sanfey, A.G., 2004. Independent coding of reward magnitude and valence in the human brain. J. Neurosci. 24, 6258–6264.

Yeung, N., Bogacz, R., Holroyd, C.B., Cohen, J.D., 2004. Detection of synchronized oscillations in the electroencephalogram: an evaluation of methods. Psychophysiology 41, 822–832.

Yordanova, J., Devrim, M., Kolev, V., Ademoglu, A., Demiralp, T., 2000. Multiple time–frequency components account for the complex functional reactivity of P300. NeuroReport 11, 1097–1103.

Yordanova, J., Falkenstein, M., Hohnsbein, J., Kolev, V., 2004. Parallel systems of error processing in the brain. NeuroImage 22, 590–602.