

Error-Related Negativity Predicts Reinforcement Learning and Conflict Biases

Report

Michael J. Frank,* Brion S. Woroach, and Tim Curran
Department of Psychology and Center
for Neuroscience
University of Colorado at Boulder
345 UCB
Boulder, Colorado 80309

Summary

The error-related negativity (ERN) is an electrophysiological marker thought to reflect changes in dopamine when participants make errors in cognitive tasks. Our computational model further predicts that larger ERNs should be associated with better learning to avoid maladaptive responses. Here we show that participants who avoided negative events had larger ERNs than those who were biased to learn more from positive outcomes. We also tested for effects of response conflict on ERN magnitude. While there was no overall effect of conflict, positive learners had larger ERNs when having to choose among two good options (win/win decisions) compared with two bad options (lose/lose decisions), whereas negative learners exhibited the opposite pattern. These results demonstrate that the ERN predicts the degree to which participants are biased to learn more from their mistakes than their correct choices and clarify the extent to which it indexes decision conflict.

Introduction

The anterior cingulate (ACC) is critically involved in demanding cognitive tasks and is most reliably activated with increasing task difficulty (Paus et al., 1998; Botvinick et al., 2001). Event-related potential studies have localized an ACC component called the error-related negativity (ERN) (Gehring et al., 1993), as it is typically more negative after participants make incorrect responses, compared to correct choices. A prevailing hypothesis holds that the ERN reflects ACC activity caused by dips in dopamine (DA) following incorrect responses (Holroyd and Coles, 2002), which is supported by evidence of these dips from animal experiments (Schultz, 2002; Satoh et al., 2003) and by observations of a strikingly similar ERP component following error feedback in trial-and-error tasks (the feedback-related negativity [FRN]) (Miltner et al., 1997; Gehring and Willoughby, 2002; Luu et al., 2003). However, the fundamental computational nature of the ERN and its implications for behavior are still a matter of intense debate. In particular, a major competing account suggests that the ERN is an index of response or decision conflict, of which errors are simply a special case (Yeung et al., 2004). Each of these ERN theories (error processing and conflict monitoring) accounts for some data that the other does not, suggesting that an inte-

grative approach is in order (Botvinick et al., 2004). Here we show that the relative size of the ERN predicts the degree to which participants learn more about the negative, as compared to positive, consequences of their decisions. Further, while ERN magnitude was not associated with overall response conflict, an effect of conflict was revealed that depended on the participants learning bias. Positive learners had larger ERNs when faced with high-conflict win/win decisions among two good options than during lose/lose decisions among two bad options, whereas negative learners showed the opposite pattern. Thus, we show that the ERN is not only an error detection mechanism, but that its relative magnitude actually predicts the degree to which participants learn from errors and that its indexing of conflict depends on the type of decision faced by the particular learner.

We employed a modified version of a reinforcement learning paradigm previously shown to be sensitive to dopaminergic manipulation (Frank et al., 2004; M.J.F. and R.C. O'Reilly, unpublished data). Three different stimulus pairs (AB, CD, EF) are presented in random order, and participants have to learn to choose one of the two stimuli (Figure 1A). Feedback follows the choice to indicate whether it was correct or incorrect, but this feedback is probabilistic. A choice of stimulus A leads to visual positive feedback in 80% of AB trials, whereas a B choice leads to negative feedback in these trials (and vice versa for the remaining 20% of trials). CD and EF pairs are less reliable: stimulus C is correct in 70% of CD trials, while E is correct in 60% of EF trials. Over the course of training, participants learn to choose stimuli A, C, and E more often than B, D, or F. Note that learning to choose A over B could be accomplished either by learning that A leads to positive feedback or that B leads to negative feedback (or both). To evaluate whether participants learned more about positive or negative outcomes of their decisions, we subsequently tested them with all novel combinations of training stimuli (e.g., AC, BD, etc.). If participants learned more from positive feedback, they should reliably choose stimulus A in all novel test pairs in which it is present. On the other hand, if they learned more from negative feedback, they should more reliably avoid stimulus B.

Differences in positive/negative feedback learning in this task were previously found to depend on different levels of dopamine in the basal ganglia (Frank et al., 2004; M.J.F. and R.C. O'Reilly, unpublished data), as predicted by our computational model (Frank, 2005). By our model's account, dopamine dips during negative feedback support NoGo learning to avoid selecting the same response in the future. Following this same logic, we predicted that the ERN, which is thought to arise from these same dopamine dips (Holroyd and Coles, 2002), would be large for participants who tended to learn more from the negative consequences of their decisions. On the other hand, if the conflict monitoring hypothesis is correct, we expected that larger ERNs would be observed during novel test pairs that involved selecting between two stimuli with similar reinforce-

*Correspondence: frankmj@psych.colorado.edu




Figure 1. Probabilistic Cognitive Reinforcement Learning Task
(A) Example stimulus pairs, which minimize explicit verbal encoding. Each pair is presented separately in different trials in random order, and participants have to select among the two stimuli; correct choices are determined probabilistically. (B) Behavioral results. Positive learners were characterized by better performance at choosing stimulus A (which depends on learning from positive reinforcement), whereas negative learners were better at avoiding stimulus B (which depends on learning from negative reinforcement). Error bars reflect SEM.

ment histories (e.g., 80% against 70%), relative to pairs that should elicit less conflict (e.g., 80% against 30%).

Results

As expected, the base-to-peak ERN magnitude following erroneous test responses was larger than that of the CRN following correct choices ($-3.04 \mu V$ compared with $-2.29 \mu V$, $t[48] = 4.4$, $p < 0.0001$, Figure 2). Notably, the ERN magnitude was predictive of the degree to which participants learned more from the negative than positive consequences of their decisions. That is, more negative ERNs were correlated with a propensity to avoid the most negative stimulus B in novel test pairs ($r = -0.46$, $p < 0.001$; Figure 3A) and with a bias to more reliably avoid B than to choose the most positive stimulus A ($r = -0.35$, $p = 0.01$). In contrast, the CRN magnitude was not predictive of either positive ($r = -0.19$, n.s.) or negative ($r = -0.2$, n.s.) learning. Finally, relatively larger ERN than CRN magnitudes were correlated with better negative than positive learning ($r = -0.35$, $p = 0.01$; Figure 3B).

To determine whether these correlational effects were robust, we categorized participants as either positive or negative reinforcement learners. Positive learners ($n = 24$) were selected as those participants who performed better at choosing A than avoiding B in novel test pairs, whereas negative learners ($n = 25$) were selected based on better performance at avoiding B (Figure 1A). Group comparisons revealed that positive learners were better than negative learners at




Figure 2. ERP Predictors of Positive and Negative Reinforcement Learning

(A and B) Response-locked ERPs during correct and erroneous choices in the test phase of the probabilistic reinforcement task. Larger ERNs were observed in negative learners. (C and D) Scalp topographies are shown for CRNs/ERNs. Gray dots show electrode cluster from which voltages were averaged across to generate all ERP waveforms in this paper. (E and F) Feedback-locked ERPs following correct and incorrect feedback during training, synchronized to each individual's FRN (shown at 0 ms, feedback was presented 250 ms prior on average; see Experimental Procedures). Negative learners had larger relative FRNs to negative than positive feedback.

choosing A ($F[1,47] = 7.4$, $p = 0.009$), while negative learners were better than positive learners at avoiding B ($F[1,47] = 20.0$, $p < 0.0001$; Figure 1B). Analysis of variance (ANOVA) results were consistent with the correlational results described above. Specifically, negative learners had significantly larger ERN amplitudes than positive learners ($F[1,47] = 6.8$, $p = 0.01$; Figure 2), with no group differences in CRN magnitude ($F[1,47] = 0.2$). Moreover, negative learners showed a greater differentiation between ERN and CRN magnitudes than did positive learners ($F[1,47] = 9.5$, $p = 0.0035$). These learning biases and associated ERP correlates were found despite no overall test accuracy difference between groups, either in the tendency to choose A over B in the AB pair ($F[1,47] = 1.0$, n.s.) or in overall performance among novel test pairs ($F[1,47] = 0.85$).

That negative learners had larger ERNs suggests that these individuals are more affected by, and therefore learn more from, their errors. This notion makes the strong prediction that the feedback negativity should also be relatively larger in these participants to negative




Figure 3. ERP Correlates of Learning and Conflict

(A and B) Across all participants, better overall negative reinforcement learning was associated with larger ERN magnitudes, and preferential biases to learn more from negative than positive feedback were associated with relatively more negative ERNs than CRNs. (C and D) Error versus conflict effects on ERN magnitude. There was a large main effect of error, but not conflict. ERNs were somewhat larger on low-conflict error trials, in which participants would be more likely to detect that they had made an error.

compared with positive feedback, which could potentially reflect the neural mechanism causing them to be more sensitive to their mistakes.

To test this idea, we recorded feedback-locked ERPs during training. Indeed, while groups did not differ in their FRN magnitudes to either negative feedback ($F[1,47] = 0.02$) or positive feedback ($F[1,47] = 2.4$, n.s.), negative learners had relatively larger FRNs to negative compared with positive feedback than did positive learners ($F[1,47] = 3.93$, $p = 0.05$; Figures 2E and 2F). In the feedback-locked ERP, there was also a late positive P300 component that was larger for negative than positive feedback ($t[48] = 6.6$, $p < 0.0001$). However, unlike the FRN, relative P300 magnitude to negative feedback was not greater in negative than positive learners ($F[1,47] = 0.0$). Thus, this P300 may relate to conscious awareness that an error has occurred but is not predictive of learning from these errors. This is consistent with reports that the P300 responds greater to negatively affective stimuli than to positive stimuli (Ito et al., 1998) and may reflect disappointment after having made an incorrect choice (Yeung and Sanfey, 2004). Overall, these results are consistent with the hypothesis that the ERN and FRN reflect common neural processes within the ACC (Miltner et al., 1997; Gehring and Willoughby, 2002; Luu et al., 2003; Nieuwenhuis et al., 2004) and are both associated with dips in dopamine release (Holroyd and Coles, 2002).

Finally, we were also interested in evaluating the role of response conflict on ERN magnitude. According to the conflict monitoring hypothesis, rather than detecting errors per se, the ACC (and by extension, the ERN)

is sensitive to the coactivation of mutually incompatible responses (Yeung et al., 2004). By this account, the ERN is observed after errors as participants tend to internally correct their erroneous choice by activating the response they *should* have made, and this conflicts with the incorrect response already activated. Note that when selecting among two stimuli that have similar reinforcement histories, participants should have more difficulty making choices and should be more likely to activate both competing responses. Thus, the conflict monitoring hypothesis predicts that the ERN should be larger in these trials in which multiple responses compete for the control of action (Yeung et al., 2004). To address this issue, we divided novel test pairs into low- and high-conflict decisions. Low-conflict decisions were defined as test trials that required choosing an overall positive stimulus (A, C, or E) over an overall negative stimulus (B, D, or F), whereas high-conflict decisions required choosing among two stimuli that have similar overall associations (both positive or both negative).

To determine whether there were differential contributions of conflict versus error in the ERN magnitude, we performed a $2 \times 2 \times 2$ repeated-measures ANOVA on group, conflict (low, high), and error (correct, incorrect). There was a highly significant main effect of error ($F[1,47] = 22.2$, $p < 0.0001$, Figures 3C and 3D), which interacted with the positive/negative learning group ($F[1,47] = 9.5$, $p = 0.0035$), again showing bigger ERNs in negative learners. Notably, there was no main effect of conflict ($F[1,47] = 2.5$, $p > 0.1$), no interaction with learning group ($F[1,47] = 0.4$), and no error by conflict interaction ($F[1,47] = 0.5$). [The nonsignificant trend for a conflict effect was actually in the reverse direction, with low-conflict decisions eliciting numerically larger ERNs than high-conflict decisions. Strictly speaking, Yeung et al. (2004) argue that the high-conflict effect should be observed *prior* to the response in correct trials (when conflict is resolved). However, we did not observe a consistent negativity prior to the response, despite using preresponse baseline corrections that should have enabled these to be observed if they existed. We also found no effect of conflict on the stimulus-locked N2.]

Nevertheless, we reasoned that the lack of conflict effect across all participants could be explained by the possibility that distinct sorts of decisions would evoke conflict in positive and negative learners. To test this idea, we further divided high-conflict decisions into win/win decisions (choosing among two positive stimuli) and lose/lose decisions (two negative stimuli). We hypothesized that both during and following win/win decisions, positive learners would be more likely to activate the alternative choice for which they had also developed a positive association. In contrast, negative learners should be more likely to activate both responses during lose/lose decisions as they attempt to avoid both negative stimuli. Thus, the conflict hypothesis predicts that positive learners should have larger ERNs during win/win than lose/lose decisions, whereas negative learners should show the opposite pattern.

Indeed, while positive and negative learners did not differ in ERN magnitudes to overall high-conflict decisions ($F[1,47] = 1.94$, n.s.), there was a significant in-




Figure 4. Response-Locked ERPs during Win/Win and Lose/Lose Decisions, Across Both Correct and Error Trials

(A and B) Positive learners had larger ERNs during high-conflict win/win than lose/lose decisions, whereas negative learners showed the reverse pattern. (C and D) Bar graphs show the mean of the per-subject peak amplitudes within each group, again showing an interaction between decision conflict type and learning bias. Note that the discrepancy between ERN magnitudes in the bar graphs and those observed in the waveforms is due to the reduction of grand average ERP amplitude by averaging across individual participants ERPs, which occur at different latencies. (E and F) Scalp topographies of voltage differences between lose/lose and win/win decisions.

interaction between positive/negative group and win/win versus lose/lose decision type ($F[1,47] = 7.6, p = 0.008$; Figure 4). That is, ERNs were larger in positive learners during win/win compared with lose/lose decisions ($t[23] = -2.3, p = 0.03$), while the opposite pattern was observed in negative learners ($t[24] = 1.53, p = 0.14$), a trend that was significant in just error trials ($t[24] = 2.6, p = 0.016$). In other words, positive reinforcement learners appear to have experienced greater conflict when choosing between two stimuli that were each previously associated with positive (compared with negative) feedback, whereas negative reinforcement learners may have experienced greater conflict when choosing among negative stimuli.

Discussion

Taken together, our findings provide valuable constraints toward advancing theoretical perspectives on the role of the ERN in reinforcement learning and decision making. First, our results are consistent with a

large literature pointing to a role for the ERN in error detection (Falkenstein et al., 1991; Holroyd and Coles, 2002; Holroyd et al., 2003; Yasuda et al., 2004; Pailing and Segalowitz, 2004) but go beyond these studies to demonstrate that the ERN also reflects error correction. Previous attempts to link the ERN with error correction failed to show a relationship with the behavioral slowing that is typically observed following errors (Gehring and Fencsik, 2001; Hacıak et al., 2004). We find that, rather than having an effect on response speed, larger ERNs were associated with a bias to learn to avoid negative events more than to seek positive events. Second, because negative learners also had relatively larger FRNs following negative feedback, our results suggest that the FRN may index a trait or state variable associated with whether participants are more responsive to positive or negative reinforcement and provide corroboratory evidence for the hypothesis that the FRN and ERN reflect common neural processes (Luu et al., 2003; Nieuwenhuis et al., 2004; Miltner et al., 1997; Gehring and Willoughby, 2002; but see van Veen et al., 2004; Nieuwenhuis et al., 2005). Third, our findings shed light on competing theories of error and conflict monitoring within the ACC. While there was no overall conflict effect on ERN magnitude, positive and negative learners may have experienced differential conflict depending on whether they were making win/win or lose/lose decisions. Finally, our results provide insight into the consistent observation that both ERNs and harm avoidance are enhanced in participants with negative affect, anxiety, and depression (Luu et al., 2000; Hacıak et al., 2004; Abrams et al., 2004), suggesting that increased ACC activity can lead to enhanced sensitivity to one's mistakes and subsequent avoidance behavior.

Despite not manipulating or measuring dopamine (DA) function directly, our findings also lend support to the general notion that ERNs are related to DA dips that occur during error processing (Holroyd and Coles, 2002). Specifically, our computational model of the basal ganglia suggests that dips in DA enable NoGo learning to avoid maladaptive decisions, whereas DA bursts support Go learning from positive decision outcomes (Frank, 2005). We recently confirmed a central prediction of this model by showing that Parkinson's patients, who have low levels of DA, learned more from negative than positive reinforcement (Frank et al., 2004). Further, dopaminergic medication reversed this bias, improving positive but impairing negative reinforcement learning, consistent with the notion that medication increases DA levels and blocks the effects of DA dips that would normally support NoGo learning (Frank, 2005). We also observed this same crossover interaction on positive/negative learning in young, healthy participants taking low doses of two opposing DA medications (M.J.F. and R.C. O'Reilly, unpublished data).

Moreover, in light of this model, our findings also reconcile conflicting ERN results previously found in Parkinson's patients (Holroyd et al., 2002; Falkenstein et al., 2001). Specifically, the model predicts that, due to low DA levels and intact NoGo learning, patients should have spared ERNs. However, DA medication should block these DA dips during errors and should therefore reduce the ERN. Indeed, spared ERNs were found in one study with nonmedicated patients (Holroyd et al.,

2002), while reduced ERNs were found in another with medicated patients (Falkenstein et al., 2001). Reduced ERNs were also found in healthy participants taking haloperidol (Zirnheld et al., 2004), which at low doses increases DA release (Wu et al., 2002) and therefore may also block the effects of DA dips needed to learn NoGo (Frank, 2005; M.J.F. and R.C. O'Reilly, unpublished data). Notably, this drug also increased errors of commission (Zirnheld et al., 2004), consistent with impaired NoGo learning. Finally, this same logic explains the reduced ERNs observed in patients with schizophrenia (Bates et al., 2004), who have abnormally high levels of subcortical DA (McGowan et al., 2004).

It is tempting to suggest that, in the present study, healthy participants who differed in ERN magnitudes and associated reinforcement learning biases may also have underlying differences in DA levels. However, it is more than likely that other neurotransmitters interacting with dopamine also play an essential role in producing the observed dissociations. In particular, large ERNs and enhanced harm avoidance have both been associated with elevated serotonin levels (Fallgatter et al., 2004; Abrams et al., 2004; Moresco et al., 2002). Notably, serotonin may exert its effects indirectly due to opponent interactions with dopamine (Daw et al., 2002): high levels of serotonin may inhibit DA release, causing dopamine dips to be more effective (Nocjar et al., 2002). Future research will more directly test the involvement of DA and serotonin using imaging and genetic analysis methods. Nevertheless, our present ERP results show that, whatever the neuromodulatory mechanism, individual differences in learning biases are associated with differential recruitment of the ACC during reinforcement.

Our data also revealed some support for the conflict monitoring hypothesis of ACC function (Botvinick et al., 2001; Yeung et al., 2004). While the ERN magnitude was not overall larger for high-conflict decisions, this could be explained by the notion that positive and negative learners may experience conflict under different decision situations. Indeed, positive learners had relatively larger ERNs for win/win decisions, whereas negative learners had larger ERNs for lose/lose decisions. Nevertheless, this same pattern of results is predicted by a recent alternative account suggesting that the ACC predicts error likelihood more than it detects response conflict (Brown and Braver, 2005). In particular, positive learners may be more likely to perceive that they are making an error when faced with a win/win decision, whereas they may feel more uncertain about lose/lose decisions (and vice versa for negative learners). This notion is consistent with recent reports that ERNs are attenuated under conditions of uncertainty (Pailing and Segalowitz, 2004). Future studies are therefore required to distinguish between these alternatives.

In summary, our results provide insight into the underlying computational function of the ERN. As noted above, various accounts suggest that DA dips occur when outcomes are worse than expected. In the Holroyd and Coles (2002) model, these dips activate ACC neurons that subsequently modify behavior. In our own modeling, we have suggested that these DA dips drive NoGo learning in the basal ganglia (Frank, 2005; O'Reilly and Frank, 2005), which then modulate response selec-

tion processes in cortex. This model accounts for various other findings of BG involvement in cognition. Nevertheless, we do not discount the likely possibility that multiple pathways are involved in avoidance learning, including both the BG and the ACC. It is also not entirely clear that ACC activity simply *reflects* the effects of DA dips, as is assumed by most models; it seems plausible that rather this activity *causes* the observed dips in midbrain DA. This is consistent with the existence of projections from the ACC to striosomes of the basal ganglia (Eblen and Graybiel, 1995) which in turn send inhibitory projections to DA cells in the ventral tegmental area (Joel and Weiner, 2000). Further, ACC lesions in animals have been shown to have reciprocal effects on the subcortical DA system (Ventura et al., 2004). Thus, further research is required to disentangle the directional effect of the involvement of the ACC and DA in error detection, correction, and cognitive decision making.

Experimental Procedures

Sample

Sixty-five young healthy individuals participated in our study. Five participants were excluded due to technical problems (one computer crash, one fell asleep, three had problems with EEG net setup/unable to achieve minimum impedance values). In addition, we filtered out participants who did not satisfy global performance measures during the test sessions prior to statistical analysis. Thus, we filtered out data from participants who did not perform better than chance (50%) on either the A or B test pairs across all segments (Frank et al., 2004). This amounted to 11 out of 60 participants whose data were not analyzed further. The remaining 49 participants (female:male = 28:21; between the ages of 18 and 29, mean 20.0, SEM 0.39) were included in the behavioral and electrophysiological analysis described above.

Task Procedures

Procedures were similar those described previously (Frank et al., 2004) but were modified for ERPs. Participants sit in front of a computer screen and view pairs of visual stimuli that are not easily verbalized (Japanese Hiragana characters, Figure 1A), presented in white on a black background, in 44 pt font. On each trial, the stimulus events consisted of a fixation period (randomly sampled from the interval 250–750 ms, green plus sign), followed by stimulus pair presentation for 750 ms, followed by a blank screen for 350 ms, followed by visual feedback for 600 ms.

Three different stimulus pairs (AB, CD, EF) are presented in random order. Participants press the left key on the button box to select the stimulus on the left or press the right key to select the stimulus on the right. Visual feedback is provided following each choice (a yellow smiley face for correct responses or a red crossout symbol for incorrect responses). If no response is made within 1000 ms, the words “no response detected” are printed in red. The position of the correct stimulus was randomized across trials, and the assignment of Hiragana character to hierarchical element A–F was randomized across participants.

We enforced a performance criterion (evaluated after each training block of 60 trials) to ensure that all participants were at the same performance level before advancing to each test segment. Because of the different probabilistic structure of each stimulus pair, we used a different criterion for each (65% A in AB, 60% C in CD, 50% E in EF). After reaching this criterion, participants were subsequently tested with the same training pairs, in addition to all novel combinations of stimuli, in random sequence. They were instructed (prior to the test phase) to use gut instinct if they did not know how to respond to these novel pairs. Each test pair was presented six times for a maximum of 1000 ms duration, and no feedback was provided.

In order to obtain sufficient trial counts for the ERP measures,

participants performed three segments of the same task, each time with new stimuli. To determine whether participants were positive or negative reinforcement learners, performance on A test pairs (AC, AD, AE, EF) and B pairs (BC, BD, BE, BF) was averaged across the three segments. If accuracy in choosing stimulus A was greater than that of avoiding stimulus B, the participant was classified as a positive learner, and vice versa for negative learners.

Electrophysiological Recording and Analysis

Scalp voltages were collected with a 128 channel Geodesic Sensor Net (Tucker, 1993) connected to an AC-coupled, 128 channel, high-input impedance amplifier (200 M Ω , Net Amps, Electrical Geodesics Inc., Eugene, OR). Amplified analog voltages (0.1–100 Hz band-pass, –3 dB) were digitized at 250 Hz. Individual sensors were adjusted until impedances were less than 50 k Ω . The feedback-locked EEG was digitally low-pass filtered at 40 Hz, while a 15 Hz low-pass filter was applied to remove bifurcations in the response-locked EEG, permitting more reliable peak amplitude measures. Trials were discarded from analyses if they contained incorrect responses, eye movements (EOG over 70 V), or more than 20% of channels were bad (average amplitude over 100 V or transit amplitude over 50 V). Across all the present analyses, subjects had at least 29 trials per condition. Individual bad channels were replaced on a trial-by-trial basis with a spherical spline algorithm (Srinivasan et al., 1996). EEG was measured with respect to a vertex reference (Cz), but an average-reference transformation was used to minimize the effects of reference site activity and accurately estimate the scalp topography of the measured electrical fields (Dien, 1998; Picton et al., 1995). The average reference was corrected for the polar average reference effect (Junghofer et al., 1999).

Following Yeung et al. (2004), response-locked ERPs (ERN, CRN) were computed within epochs starting 800 ms prior to the response and lasting 200 ms after the response and were baseline corrected with respect to the first 100 ms of these epochs. Feedback-locked ERPs (FRN) were computed within epochs starting 100 ms prior to the feedback and lasting 1000 ms afterward. In all cases, ERPs were baseline corrected with respect to the first 100 ms of the epoch. Following other ERN studies using similar recording procedures (Tucker et al., 2003), analyses focused on a cluster of sensors surrounding the standard FCz location, as depicted with the sensor locations marked in Figures 2 and 4. Following Yeung and Sanfey (2004) and Holroyd et al. (2002), we defined the ERN/CRN and FRN as the peak-to-peak voltage difference between the first negative peak following the response (mean latency = 58 ms) and the preceding positive peak. Similarly, FRN amplitude was calculated as the difference between the negative FRN peak within a window of 190–300 ms (mean latency = 249 ms) and the preceding positive peak. In the FRN plots (Figure 4), waveforms were synchronized to each individual's FRN, which appears at 0 ms. This was done because better positive than negative learning was correlated with longer FRN latencies ($r = 0.41$, $p = 0.003$), and these latency differences cause smearing of FRN magnitudes when displayed relative to the feedback. Synchronizing to each individual's FRN alleviates this problem and permits displaying of actual FRN magnitudes across positive and negative learners.

Acknowledgments

This research was supported by NIH grants MH64812-01 and MH069597-01, ONR grant N00014-03-1-0428, and the James S. McDonnell Foundation. We thank C. Debusse, D. Blum, J. Capps, D. Chase, D. Germain, J. Hancock, M. Hauser, C. Taylor, E. Walsky, and B. Young for help with subject testing; and Randy O'Reilly for helpful discussion.

Received: March 24, 2005
Revised: June 7, 2005
Accepted: June 20, 2005
Published: August 17, 2005

References

Abrams, K.Y., Yune, S.K., Kim, S.J., Jeon, H.J., Han, S.J., Hwang, J., Sung, Y.H., Lee, K.J., and Lyoo, I.K. (2004). Trait and state as-

pects of harm avoidance and its implication for treatment in major depressive disorder, dysthymic disorder, and depressive personality disorder. *Psychiatry Clin. Neurosci.* 58, 240–248.

Bates, A.T., Liddle, P.F., Kiehl, K.A., and Ngan, E.T. (2004). State dependent changes in error monitoring in schizophrenia. *J. Psychiatr. Res.* 38, 347–356.

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

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

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

Daw, N.D., Kakade, S., and Dayan, P. (2002). Opponent interactions between serotonin and dopamine. *Neural Netw.* 15, 603–616.

Dien, J. (1998). Issues in the application of the average reference: Review, critiques, and recommendations. *Behav. Res. Methods Instrum. Comput.* 30, 34–43.

Eblen, F., and Graybiel, A.M. (1995). Highly restricted origin of prefrontal cortical inputs to striosomes in the macaque monkey. *J. Neurosci.* 15, 5999–6013.

Fallgatter, A.J., Herrmann, M.J., Roemmler, J., Ehlis, A.C., Wagners, A., Heidrich, A., Ortega, G., Zeng, Y., and Lesch, K.P. (2004). Allelic variation of serotonin transporter function modulates the brain electrical response for error processing. *Neuropsychopharmacology* 29, 1506–1511.

Falkenstein, M., Hohnsbein, J., Hoormann, J., and Blanke, L. (1991). Effects of cross-modal divided attention on late ERP components: II. error processing in choice reaction tasks. *Electroencephalogr. Clin. Neurophysiol.* 78, 447–455.

Falkenstein, M., Hielscher, I., Dziobek, P., Schwarzenau, J., and Hoormann, B. (2001). Action monitoring, error detection, and the basal ganglia: an ERP study. *Neuroreport* 12, 157–161.

Frank, M.J. (2005). Dynamic dopamine modulation in the basal ganglia: A neurocomputational account of cognitive deficits in medicated and non-medicated Parkinsonism. *J. Cogn. Neurosci.* 17, 51–72.

Frank, M.J., Seeberger, L.C., and O'Reilly, R.C. (2004). By carrot or by stick: Cognitive reinforcement learning in Parkinsonism. *Science* 306, 1940–1943.

Gehring, W.J., and Fencsik, D.E. (2001). Functions of the medial frontal cortex in the processing of conflict and errors. *J. Neurosci.* 21, 9430–9437.

Gehring, W.J., and Willoughby, A.R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science* 295, 2279–2281.

Gehring, W.J., Goss, B., Coles, M.G.H., Meyer, D.E., and Donchin, E. (1993). A neural system for error detection and compensation. *Psychol. Sci.* 4, 385–390.

Hacjak, G., McDonald, N., and Simons, R.F. (2004). Error-related psychophysiology and negative affect. *Brain Cogn.* 56, 189–197.

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

Holroyd, C.B., Praamstra, P., Plat, E., and Coles, M.G. (2002). Spared error-related potentials in mild to moderate Parkinson's disease. *Neuropsychologia* 40, 2116–2124.

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

Ito, T., Larsen, J.T., Smith, N.K., and Cacioppo, J.T. (1998). Negative information weighs more heavily on the brain: the negativity bias in evaluative categorizations. *J. Pers. Soc. Psychol.* 75, 887–900.

Joel, D., and Weiner, I. (2000). The connections of the dopaminergic system with the striatum in rats and primates: an analysis with respect to the functional and compartmental organization of the striatum. *Neuroscience* 96, 451–474.

- Junghofer, M., Elbert, T., Tucker, D.M., and Braun, C. (1999). The polar average reference effect: a bias in estimating the head surface integral in EEG recording. *Clin. Neurophysiol.* *110*, 1149–1155.
- Luu, P., Collins, P., and Tucker, D.M. (2000). Mood, personality, and self-monitoring negative affect and emotionality in relation to frontal lobe mechanisms of error monitoring. *J. Exp. Psychol. Gen.* *129*, 43–60.
- Luu, P., Tucker, D.M., Derryberry, D., Reed, M., and Poulsen, C. (2003). Electrophysiological responses to errors and feedback in the process of action regulation. *Psychol. Sci.* *14*, 47–53.
- McGowan, S., Lawrence, A.D., Sales, T., Quedest, D., and Graby, P. (2004). Presynaptic dopamine dysfunction in schizophrenia: a positron emission tomographic [¹⁸F]fluorodopa study. *Arch. Gen. Psychiatry* *61*, 134–142.
- Miltner, W.H., Baum, C.H., and 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. *J. Cogn. Neurosci.* *9*, 788–798.
- Moresco, F.M., Dieci, M., Vita, A., Messa, C., Gobbo, C., Galli, L., Rizzo, G., Panzacchi, A., De Peri, L., Invernizzi, G., and Fazio, F. (2002). In vivo serotonin 5HT(2A) receptor binding and personality traits in healthy subjects: a positron emission tomography study. *Neuroimage* *17*, 1470–1478.
- Nieuwenhuis, S., Holroyd, C.B., Mol, N., and Coles, M.G. (2004). Reinforcement-related brain potentials from medial frontal cortex: origins and functional significance. *Neurosci. Biobehav. Rev.* *28*, 441–448.
- Nieuwenhuis, S., Slagter, H.A., von Geusau, N.J., Heslenfeld, D.J., and Holroyd, C.B. (2005). Knowing good from bad: Differential activation of human cortical areas by positive and negative outcomes. *Eur. J. Neurosci.* *21*, 3161–3168.
- Nocjar, C., Roth, B.L., and Pehek, E.A. (2002). Localization of 5-HT(2A) receptors on dopamine cells in subnuclei of the midbrain A10 cell group. *Neuroscience* *111*, 163–176.
- O'Reilly, R.C., and Frank, M.J. (2005). Making working memory work: A computational model of learning in the frontal cortex and basal ganglia. *Neural Computat.*, in press.
- Pailing, P.E., and Segalowitz, S.J. (2004). The effects of uncertainty in error monitoring on associated ERPs. *Brain Cogn.* *56*, 215–233.
- Paus, T., Koski, L., Caramanos, Z., and Westbury, C. (1998). Regional differences in the effects of task difficulty and motor output on blood flow response in the human anterior cingulate cortex: A review of 107 PET activation studies. *Neuroreport* *9*, R37–R47.
- Picton, T.W., Lins, O.G., and Scherg, M. (1995). The recording and analysis of event-related potentials. In *Handbook of Neuropsychology*, Volume 10, F. Boller and J. Grafman, eds. (Amsterdam: Elsevier), pp. 373.
- Satoh, T., Nakai, S., Sato, T., and Kimura, M. (2003). Correlated coding of motivation and outcome of decision by dopamine neurons. *J. Neurosci.* *23*, 9913–9923.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron* *36*, 241–263.
- Srinivasan, R., Nunez, P.L., Tucker, D.M., Silberstein, R.B., and Cadusch, P.J. (1996). Spatial sampling and filtering of EEG with spline laplacians to estimate cortical potentials. *Brain Topogr.* *8*, 355–366.
- Tucker, D.M. (1993). Spatial sampling of head electrical fields: The geodesic sensor net. *Electroencephalogr. Clin. Neurophysiol.* *87*, 154–163.
- Tucker, D.M., Luu, P., Frishkoff, G., Quiring, J., and Poulsen, C. (2003). Frontolimbic response to negative feedback in clinical depression. *J. Abnorm. Psychol.* *112*, 667–678.
- van Veen, V., Holroyd, C.B., Cohen, J.D., Stenger, V.A., and Carter, C.S. (2004). Errors without conflict: implications for performance monitoring theories of anterior cingulate cortex. *Brain Cogn.* *56*, 267–276.
- Ventura, R., Alcaro, A., Madolesi, L., and Puglisi-Allegra, S. (2004). In vivo evidence that genetic background controls impulse-dependent dopamine release induced by amphetamine in the nucleus accumbens. *J. Neurochem.* *89*, 494–502.
- Wu, Q., Reith, M.E., Walker, Q.D., Kuhn, C.M., Carroll, F.I., and Garris, P.A. (2002). Concurrent autoreceptor-mediated control of dopamine release and uptake during neurotransmission: an in vivo voltammetric study. *J. Neurosci.* *22*, 6272–6281.
- Yasuda, A., Sato, A., Miyawaki, K., Kumano, H., and Kuboki, T. (2004). Error-related negativity reflects detection of negative reward prediction error. *Neuroreport* *15*, 2561–2565.
- Yeung, N., and Sanfey, A.G. (2004). Independent coding of reward magnitude and valence in the human brain. *J. Neurosci.* *24*, 6258–6282.
- Yeung, N., Botvinick, M.M., and Cohen, J.D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol. Rev.* *111*, 931–959.
- Zirnheld, P.J., Carroll, C.A., Kieffaber, P.D., O'Donnell, B.F., Shekhar, A., and Hetrick, W.P. (2004). Haloperidol impairs learning and error-related negativity in humans. *J. Cogn. Neurosci.* *16*, 1098–1112.