

Context dependence of the event-related brain potential associated with reward and punishment

CLAY B. HOLROYD,^a JEFF T. LARSEN,^b AND JONATHAN D. COHEN^{a,c}

^aDepartment of Psychology, Princeton University, Princeton, New Jersey 08540, USA

^bDepartment of Psychology, Texas Tech University, Lubbock, Texas 79409, USA

^cCenter for the Study of Mind Brain and Behavior, Princeton University, Princeton, New Jersey 08544, USA

Abstract

The error-related negativity (ERN) is an event-related brain potential elicited by error commission and by presentation of feedback stimuli indicating incorrect performance. In this study, the authors report two experiments in which participants tried to learn to select between response options by trial and error, using feedback stimuli indicating monetary gains and losses. The results demonstrate that the amplitude of the ERN is determined by the value of the eliciting outcome relative to the range of outcomes possible, rather than by the objective value of the outcome. This result is discussed in terms of a recent theory that holds that the ERN reflects a reward prediction error signal associated with a neural system for reinforcement learning.

Descriptors: Error-related negativity, Feedback, Reinforcement learning, Reward, Punishment

When an organism takes an action in pursuit of a goal, the success or failure of that effort is often indicated by an event in the external environment. Such *performance feedback* can consist of physical rewards and punishments and, for people, abstract symbols that convey the degree of correctness or incorrectness of the action. Typically the information conveyed by the feedback is used by the organism to improve its performance at the task (Thorndike, 1911/1970). When the feedback tells only the outcome of the behavior (e.g., correct or incorrect), but not how the behavior should be improved, that process is called *reinforcement learning* (for review, see Sutton & Barto, 1998). Much flexibility in human behavior can be attributed to principles of reinforcement learning (for review, see Davey & Cullen, 1988), and researchers have for decades used animal models to study its neural basis (e.g., see Gabriel, 1993; Reynolds, Hyland, & Wickens, 2001; Schoenbaum, Chiba, & Gallagher, 1998; Shidara & Richmond, 2002; Shima & Tanji, 1998; Waelti, Dickinson, & Schultz, 2001). This interest is reflected in a recent flurry of functional neuroimaging studies investigating how rewards, punishments, and abstract perfor-

mance feedback are processed in the human brain (e.g., Breiter, Aharon, Kahneman, Dale, & Shizgal, 2001; Bush et al., 2002; Delgado, Locke, Stenger, & Fiez, 2003; Elliott, Newman, Longe, & Deakin, 2003; Holroyd et al., 2003; Knutson, Adams, Fong, & Hommer, 2001; McClure, Berns, & Montague, 2003; Monchi, Petrides, Petre, Worsley, & Dagher, 2001; O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003; Rogers et al., 1999; Thut et al., 1997; Ullsperger & von Cramon, 2003; Van Veen, Holroyd, Cohen, Stenger, & Carter, 2002; Zalla et al., 2000).

Several studies have also examined the effect of performance feedback on the event-related brain potential (ERP). Much of this research has focused on the amplitude of the P300, a parietally distributed ERP component sensitive, among many things, to the occurrence of infrequent, task-relevant events (for review, see Donchin & Coles, 1988). Whereas the earliest of these studies suggested that feedback stimuli indicating incorrect performance (*negative feedback*) elicited larger P300s than feedback stimuli indicating correct performance (*positive feedback*; e.g., MacKay, 1984; Picton, Hillyard, & Galambos, 1976; Squires, Hillyard, & Lindsay, 1973), subsequent experiments revealed that, when equated for probability of occurrence, positive and negative feedback elicited equally large P300s (Campbell, Courchesne, Picton & Squires, 1979). Rather, the amplitude of the P300 was shown to be largest when the feedback stimulus disconfirmed the participant's prediction of the outcome. Thus, the largest P300s were elicited by negative feedback when participants thought they made a correct response, and by positive feedback when participants thought they made an incorrect response (Horst, Johnson, & Donchin, 1980). Remarkably few of these studies examined P300 amplitude elicited by rewards and punishments, as opposed to feedback indicating

This research was supported in part by National Institute of Mental Health (NIMH) grant MH62196 and postdoctoral fellowship MH63550. Jeff Larsen was supported by research funding from the Woodrow Wilson School for Public and International Affairs, Princeton University. We are grateful to Jack Gelfand and Dustin Reagan for technical assistance. We also thank Sander Nieuwenhuis and Nick Yeung for their helpful comments on previous drafts of this article.

Address reprint requests to: Clay Holroyd, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08540, USA. E-mail: cholroyd@princeton.edu.

success or failure. The few studies that did examine this question found that P300 amplitude increased in proportion to both the amount of reward received and the amount of penalty incurred (Johnston, 1979; Sutton, Tueting, Hammer, & Hakerem, 1978; see also Yeung & Sanfey, 2003).

In contrast to the P300, a component of the ERP that was differentially sensitive to negative and positive feedback was recently identified by Miltner, Braun, and Coles (1997). This component was distributed over frontal-central regions of the scalp and reached maximum amplitude about 250 ms following feedback presentation. Evidence of this component can be seen in several previous ERP studies (Horst et al., 1980; Johnson & Donchin, 1978; Picton et al., 1976; Squires et al., 1973), some of which even identified the component (Campbell et al., 1979; Johnston, 1979; Ruchkin, Sutton, Munson, Silver, & Macar, 1981) and discussed its involvement in error-related processing at length (Sutton et al., 1978; Takasawa, Takino, & Yamazaki, 1990). In contrast to these studies, however, the study by Miltner and colleagues equated for the probability of occurrence of the positive and negative feedback, ruling out the possibility that the component could be associated with stimulus frequency rather than with response correctness. Miltner and colleagues proposed that the component was related to the “error-related negativity” (ERN), an ERP component that reaches maximum amplitude about 100 ms following error commission in speeded response time tasks (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1990; Gehring, Goss, Coles, Meyer, & Donchin, 1993). The authors suggested that a single error-processing mechanism produced both an ERN associated with error commission and an ERN associated with feedback. This hypothesis was later extended by Holroyd and Coles (2002), who proposed that both the “response ERN” and the “feedback ERN” were elicited by a system for reinforcement learning. They proposed that this system applied information obtained from rewards and punishments, as well as from abstract indicators of success and failure, toward updating the response production system (see also Holroyd, Coles, & Nieuwenhuis, 2002).

One unresolved issue concerns how the system that produces the ERN determines whether an event is good or bad. Is the same event always evaluated by the system in the same way? To examine this issue, in this study we considered the effects of *context* on the amplitude of the feedback ERN. By *context* we mean the set of experimental factors that affect how a particular feedback stimulus is evaluated. If a feedback stimulus is evaluated the same way in different experimental contexts, then that stimulus is evaluated in a *context-independent* fashion. Conversely, if a feedback stimulus is evaluated differently in different experimental contexts, then that stimulus is evaluated in a *context-dependent* fashion. We conducted experiments in which the context associated with particular rewards and punishments was varied by condition. We adopted a pseudo trial-and-error learning task in which participants selected one of several possible actions and, following each choice, received a monetary reward or penalty. Unbeknownst to the participants, however, the outcomes were actually random. The experiments consisted of three conditions, with three possible outcomes in each condition. In a “win” condition, the outcomes consisted of no money gained (nor lost), a small win, and a large win. In a “lose” condition, the outcomes consisted of no money lost (nor gained), a small loss, and a large loss. Last, in an “even” condition, the outcomes consisted of some money gained, some money lost, and no money gained nor lost.

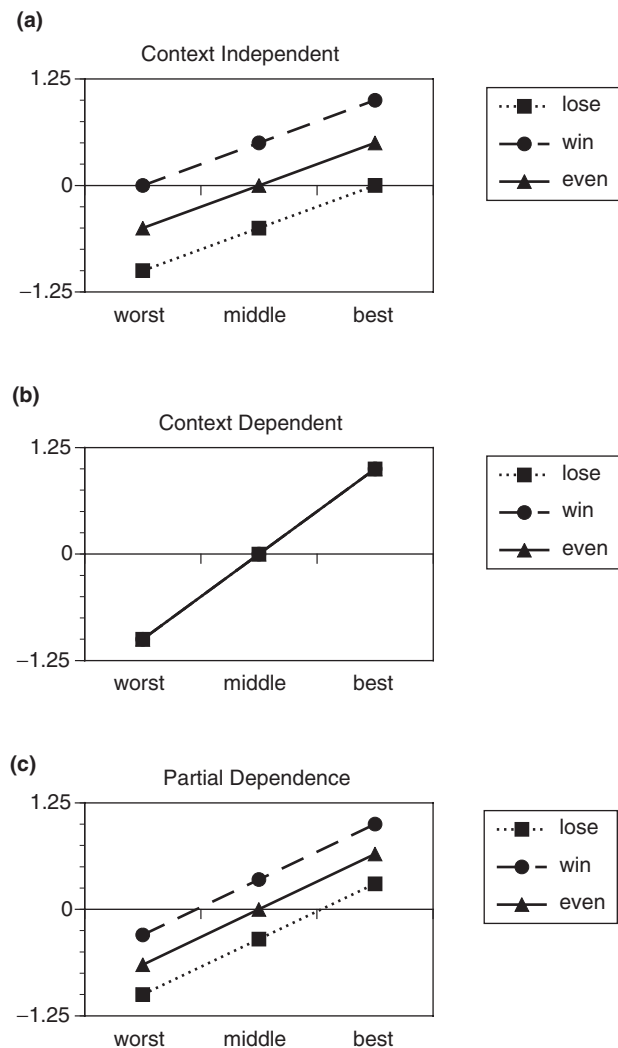


Figure 1. Examples of value system context dependence. (a): Context independent value system. (b): Context dependent value system. (c): Value system exhibiting both context dependence and context independence. Lose, win, even: task conditions. Worst, middle, best: the worst, intermediate, and best of three possible outcomes, respectively, for each condition, as given by their objective values. Ordinate corresponds to the internal value attributed to each outcome, in arbitrary units; positive values indicate desired outcomes and negative values indicate undesired outcomes. See text for details.

Figure 1 illustrates possible ways in which each outcome for the lose, win, and even conditions could be evaluated. For all three conditions, “worst” indicates the least favorable outcome, “best” indicates the most favorable outcome, and “middle” indicates the intermediate outcome. The ordinate shows, in arbitrary units, the internal value associated with each outcome, with positive values indicating desired outcomes and negative values indicating undesired outcomes. If the internal values attributed to the outcomes were context independent, then the value associated with each outcome would reflect the objective values of the outcomes (Figure 1a). In this case, the system would equally value receiving nothing when it could have gained money (the worst outcome in the win condition), and receiving nothing when it could have lost money (the best outcome in the lose condition). In contrast, if the internal values were context dependent, then the system would determine the value of each

outcome relative to the other potential outcomes in the same condition, rather than in terms of the objective value associated with each outcome (Figure 1b). Thus, in this example, the system would value receiving nothing when it could have lost money more than receiving nothing when it could have gained money. Figure 1c illustrates a third possibility involving a combination of context dependence and context independence (“partial dependence”). Note that in this case, the system values a large gain in the win condition more than receiving nothing in the lose condition (context independence), and values receiving nothing in the lose condition more than receiving nothing in the win condition (context dependence). Of course, such a system of preferences could be described by a myriad of value functions, so long as the functions increased monotonically with reward (French, 1988); we adopted the linear functions in Figure 1 simply for ease of exposition.

To investigate how context affects the system that produces the ERN, we conducted two experiments. The first experiment consisted of an even condition, and the second experiment consisted of both a win condition and a lose condition. The results are interpreted in terms of the reinforcement learning theory of the ERN (Holroyd & Coles, 2002).

Method

Participants

Sixteen people participated in Experiment 1 (12 men, 4 women) and 16 other people participated in Experiment 2 (11 men, 5 women). Most participants were undergraduate students at Princeton University (age 26.0 ± 9.7 years) and took part in the experiment either for course credit or for payment (\$20.00 for a 2-h session). In addition, all participants were paid a sum of bonus money (\$5.00 for Experiment 1, \$15.00 for Experiment 2) at the end of the session (see below).

Task

In both experiments, participants sat comfortably about 1.5 m in front of a computer screen in an electromagnetically shielded room. On each trial of both experiments, participants saw an imperative stimulus (0.6° high, 5.0° wide, white color against a black background) consisting of three circles in a row (i.e., “O O O”). Participants were asked to imagine that these circles were balloons, and told that some of the balloons contained an item (see below). The imperative stimulus remained on the screen until the participant selected a balloon by pressing one of three buttons on a response pad. At the time of the response, the imperative stimulus was replaced by a second stimulus (0.6° high, 5.0° wide, blue color, 1 s duration) in which the selected balloon was replaced by an asterisk (e.g., “O * O,” if the participant selected the middle balloon). The purpose of the asterisk was to indicate to the participant which balloon they had selected. Following the offset of the second stimulus, a feedback stimulus appeared (0.6° , red color, 1 s duration) directly above the location associated with the middle balloon. The interstimulus interval (ISI) between the offset of the feedback stimulus and the onset of the next imperative stimulus was 0.5 s.

In Experiment 1 (the “even” condition), participants were told that they would begin the task with \$5.00 in bonus money, and that on each trial one of the balloons contained a dime, that a second balloon contained nothing, and that a third balloon contained a “gremlin” that would steal 10 cents from them. The feedback stimuli consisted of three corresponding types: a “+”

stimulus informed participants that they gained 10 cents, a “-” stimulus informed the participants that they lost 10 cents, and a “o” stimulus informed the participants that they neither gained nor lost any money. (Note that the feedback stimuli were not counterbalanced across participants because the ERN is insensitive to the physical nature of the stimuli, as it can be elicited by arbitrary stimuli in multiple input modalities; Holroyd & Coles, 2002; Miltner et al., 1997.) Participants were told they should respond in such a way that they would maximize the amount of money they would receive at the end of the experiment. They were also told that they would not end the experiment owing the experimenter money. Unbeknownst to the participants, the type of feedback stimulus was selected at random, with equal probability and without replacement, on each trial. The task was divided into three blocks of 99 trials. Thus participants saw each feedback stimulus 33 times per block, and 99 times across the entire experiment. Because the feedback probabilities were balanced across feedback types, participants neither earned nor lost money and at the end of the experiment were given the \$5.00 in bonus money with which they began the task. No information about the participants’ overall performance was provided until the end of the experiment, when they received the bonus.

Experiment 2 consisted of two conditions. In a “win” condition, participants were told that one of the three balloons was empty, that a second balloon contained 2.5 cents, and that the third balloon contained 5 cents. The feedback stimuli consisted of three corresponding types: a “o” stimulus informed participants that they neither gained nor lost any money, a “+” stimulus informed participants that they gained 2.5 cents, and a “++” stimulus informed participants that they gained 5 cents. Participants were told that they should respond in such a way that they would maximize the amount of bonus money obtained. In a “lose” condition, participants were told that they would begin the task with \$15.00 in bonus money, and that one of the three balloons was empty, that a second balloon contained a gremlin that would steal 2.5 cents from them, and that third balloon contained a “big gremlin” that would steal 5 cents from them. The feedback stimuli consisted of three corresponding types: a “o” stimulus indicated to the participant that they neither gained nor lost any money, a “-” stimulus indicated to the participant that they lost 2.5 cents, and a “--” stimulus indicated to the participant that they lost 5 cents. Participants were told to respond in such a way that they would maximize the amount of bonus money saved. As in Experiment 1, the type of feedback stimulus was selected at random, with equal probability and without replacement, on each trial in each condition. Each condition was divided into three blocks of 99 trials, so the participants saw each feedback stimulus 33 times per block, and 99 times across each condition. In the win condition, participants earned \$7.50 in bonus money, and in the lose condition, participants saved \$7.50 in bonus money. Thus the participants received \$15.00 in bonus money at the end of the session. At the end of the first condition, participants were told the amount of bonus money they earned in that condition and were informed that they would receive the money at the end of the session. The order of the two conditions was counterbalanced across participants.¹

¹In addition to the reward tasks, an “oddball” task, in which participants counted the occurrence of an infrequent target stimulus, was also conducted in each experiment. This data will be presented elsewhere (Holroyd, in press).

Data Acquisition

An 87-channel electrode cap with Ag/AgCl electrodes was applied to each participant. The electroencephalogram (EEG) was recorded along the midline according to the 10–20 system from channels FPz, AFz, Fz, FCz, Cz, CPz, Pz, OPz, Oz, and Iz (Jasper, 1958). Other electrodes were placed on the right mastoid, above and below the right eye, and on the outer canthi of both eyes. The electrode common was placed on the chin or on the cheek. All electrode recordings were referenced to an electrode placed on the left mastoid. EEG data were recorded with Sensorium Inc. (Charlotte, VT) EPA-6 128 Channel Electro-Physiology Amplifiers at a sample rate of 250 Hz (bandpass 0.1–300 Hz). Impedances were less than 40 K Ω . Experimental control and data acquisition were controlled by E-Prime (Psychology Software Tools, Inc., Pittsburgh, PA) and Cogniscan (Newfoundland, NJ), respectively. Participants completed a short questionnaire upon completion of the experiment.

Data Analysis

For each feedback stimulus, a 1-s epoch of data (200 ms baseline) was extracted from the continuous data file for analysis. Ocular artifact was corrected with an eye-movement correction algorithm (Gratton, Coles, & Donchin, 1983). The EEG data were rereferenced off-line to linked-mastoid electrodes and baseline corrected by subtracting, from each sample of data recorded at each channel, one-half the activity recorded at the right mastoid and the average activity of that channel during the baseline period. Single trial EEG data were lowpass filtered below 20 Hz with the Interactive Data Language (Research Systems, Inc., Boulder, CO) digital filter algorithm. ERPs were created for each participant by averaging the single-trial EEG according to feedback type.

ERN amplitude was measured base to peak at all EEG channels with an algorithm described in Holroyd, Nieuwenhuis, Yeung, and Cohen (in press). The amplitude of this component was then evaluated at channel FCz, where it reached a maximum. The data were submitted to analysis of variance (ANOVA) with repeated measures, and to paired *t* tests. The Greenhouse–Geisser correction for repeated measures was applied where appropriate (Keselman & Rogan, 1980).

Results

Behavior

Because the feedback stimuli in the reward tasks were delivered at random, the task did not allow for any meaningful analysis of participant performance. However, upon completion of the task, the participants were asked to rate the amount of attention they paid to the feedback stimuli on a scale from 1 (*the feedback stimuli were generally ignored*) to 5 (*the feedback stimuli were evaluated closely; participants were interested whether or not they won or lost money on each trial*). Participants rated their interest in the feedback as 4.0 ± 0.8 on the scale, suggesting that they did, in fact, attend to the task in which they were engaged. Furthermore, debriefing revealed that most participants detected transient patterns in the feedback or believed that they exercised some degree of control over the feedback, suggesting that they evaluated the feedback and tried to use it to guide their performance.

ERPs

Figure 2 illustrates the scalp distribution of the ERN for the even (top) and lose and win (bottom) conditions. As can be seen from the figure, the feedback ERN displayed a frontal/frontal-central scalp distribution (cf. Gehring and Willoughby, 2002; Miltner et al., 1997). For Experiment 1 (the even condition), a two-way repeated measures ANOVA on outcome (worst, middle, best) and electrode (the 10 midline electrodes) revealed main effects of outcome, $F(2,30) = 9.5$, $p < .005$, $\epsilon = 0.81$, and electrode, $F(9, 135) = 10.4$, $p < .0005$, $\epsilon = 0.24$, and an interaction between electrode and outcome, $F(18,270) = 3.05$, $p < .05$, $\epsilon = 0.22$. Similarly, for Experiment 2, a three-way repeated-measures ANOVA on outcome (worst, middle, best), condition (lose, win), and electrode (the 10 midline electrodes) revealed a main effect of outcome, $F(2,30) = 6.54$, $p = .005$, $\epsilon = 0.96$, a main effect of electrode, $F(9,135) = 13.7$, $p < .0001$, $\epsilon = 0.29$, and an outcome by electrode interaction, $F(18,270) = 5.2$, $p < .001$, $\epsilon = 0.23$. In contrast, the main effect of condition and all interactions involving condition were not significant, $p > .05$.

To examine the ERN more closely, we evaluated the ERN at channel FCz, where it was largest. Figure 3 illustrates the ERPs, recorded at channel FCz, elicited by the feedback stimuli in Experiment 1 (a) and Experiment 2 (b and c). The ERN can be seen as the negativity peaking at about 250 ms in response to negative feedback (see for example the “–” condition in Figure 3a). This negativity does not appear in response to stimuli




Figure 2. ERN scalp distributions for the even (Experiment 1, top) and lose and win (Experiment 2, bottom) conditions. Solid, dashed, and dotted lines correspond, respectively, to the worst, middle, and best outcomes within each condition. Ordinate is in microvolts. Abscissa: electrode location.

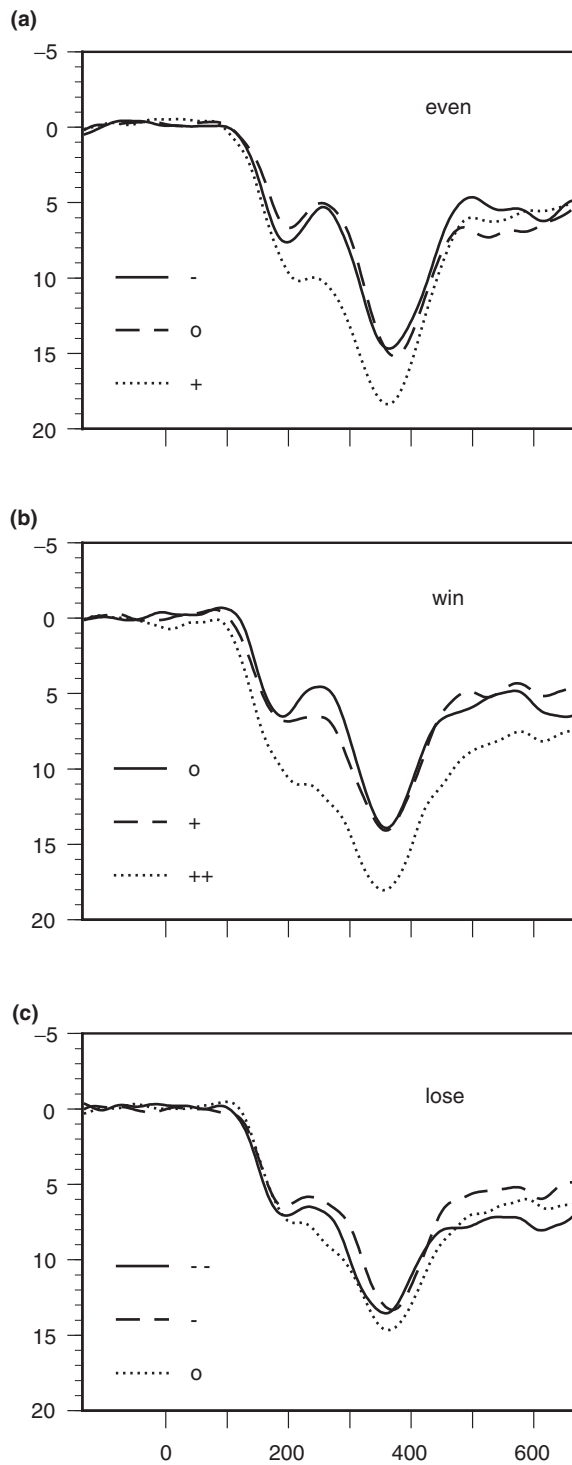


Figure 3. ERPs. (a): Experiment 1. —, o, +: ERPs elicited by the worst, middle, and best outcomes of the even condition, respectively. (b): Experiment 2. o, +, ++: ERPs elicited by the worst, middle, and best outcomes of the win condition, respectively. (c): Experiment 2. --, -, o: ERPs elicited by the worst, middle, and best outcomes of the lose condition, respectively. Ordinate is in microvolts. Abscissa is in milliseconds; 0ms corresponds to time of stimulus onset. Data associated with channel FCz.

indicating that participants received maximum reward (e.g., the “+” condition in Figure 3a). Figure 4 plots the amplitude of the ERN associated with these conditions according to their relative

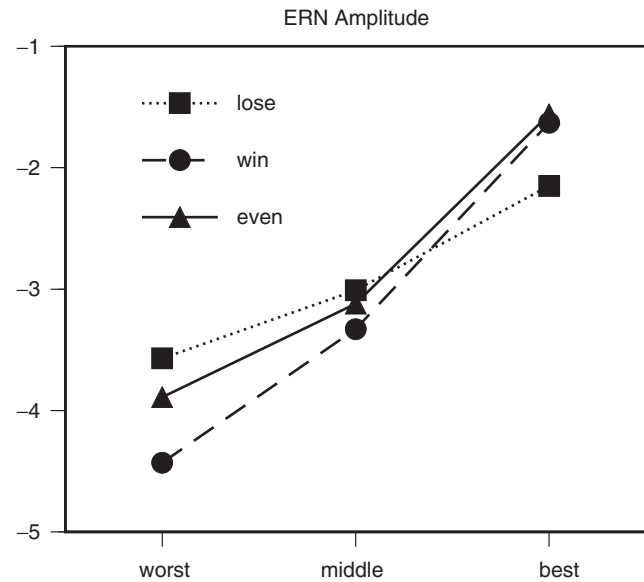


Figure 4. ERN amplitudes recorded at channel FCz, for the lose, win, and even conditions. Worst, middle, best: the worst, intermediate, and best of the three outcomes, respectively, for each condition.

outcome. For Experiment 1 (the even condition), a one-way ANOVA with repeated measures on ERN amplitude as a function of outcome (worst, middle, best) indicated that not all the ERN amplitudes were equal, $F(2,30) = 8.4$, $p < .005$, $\epsilon = 0.87$. For Experiment 2, a two-way repeated-measures ANOVA on condition (win, lose) and outcome (worst, middle, best) on ERN amplitudes revealed a main effect of outcome, $F(2, 30) = 9.0$, $p = .001$, $\epsilon = 0.96$, but no effect of condition, $F(1, 15) = 0.3$, $p = 0.60$, nor a condition by outcome interaction $F(2, 30) = 1.0$, $p = 0.39$. Taken together, these results suggest that the system that produces the ERN is context dependent: The system appears to determine the value associated with each outcome relative to the other potential outcomes in the same condition, rather than in terms of the objective value associated with each outcome (Figure 1b). This interpretation is underscored by a paired t test indicating that the empty balloons in the win condition (which represented the worst outcome) elicited larger ERNs than the empty balloons in the lose condition (which represented the best outcome; Figure 5a), $t(15) = -3.3$, $p < .005$; and by a paired t test suggesting that the middle outcomes in both conditions elicited ERNs of comparable size (Figure 5b), $t(15) = 0.4$, $p = .72$. Thus, the value associated with a particular outcome appears to depend on the context in which the outcome is received: In one case, identical outcomes (empty balloons) are associated with different values, and in the other case, different outcomes (“+” and “-”) are associated with similar values.

For the ERN to reflect a value system, however, the amplitude of the ERN must be ordered by preference, such that the amplitude of the worst outcome is greater than or equal to the amplitude of the middle outcome, which in turn is greater than or equal to the amplitude of the best outcome. In fact, confidence intervals associated with Bonferroni multiple pairwise comparisons indicated that, for both Experiment 1 (the even condition) and Experiment 2 (data collapsed across the win and lose conditions), ERN amplitudes were significantly different from each other between the worst and best outcomes (Experiment 1 [$-3.8 \mu\text{V}$, $-0.9 \mu\text{V}$], Experiment 2 [$-3.4 \mu\text{V}$, $-0.8 \mu\text{V}$]), and




Figure 5. ERPs. (a): ERPs elicited by the “empty balloon” outcomes for the lose and win conditions. (b): ERPs elicited by the middle outcomes for the lose and win conditions. Ordinate is in microvolts. Abscissa is in milliseconds; 0 ms corresponds to time of stimulus onset. Data associated with channel FCz.

between the middle and best outcomes (Experiment 1 [$-3.0 \mu\text{V}$, $-0.9 \mu\text{V}$], Experiment 2 [$-2.5 \mu\text{V}$, $-0.1 \mu\text{V}$]), but not between the worst and middle outcomes (Experiment 1 [$-2.2 \mu\text{V}$, $0.7 \mu\text{V}$], Experiment 2 [$-2.1 \mu\text{V}$, $0.4 \mu\text{V}$]).

Discussion

These results strongly suggest that the system that produces the ERN operates in a context-dependent manner: The system determines whether an outcome is favorable or unfavorable on the basis of the range of possible outcomes it might have encountered in the current setting. As such, the system judges the best possible outcome to be favorable and the worst possible outcome to be unfavorable. For example, a \$1,000 reward would be considered favorable if the alternative reward was \$500, but the same \$1,000 reward would be considered unfavorable if the alternative reward was \$2,000.

We interpret this finding in terms of the reinforcement learning theory of the ERN (Holroyd & Coles, 2002). According to this theory, a monitoring system located in the basal ganglia

continually evaluates ongoing events, including responses and feedback, and predicts whether those events will terminate favorably or unfavorably. When the monitoring system revises its prediction for the better, it produces a positive prediction error signal (indicating that ongoing events are “better than expected”), and when the monitoring system revises its predictions for the worse, it produces a negative prediction error signal (indicating that ongoing events are “worse than expected”). The theory holds that the mesencephalic dopamine system conveys these error signals back to the basal ganglia, which use the error signals to improve their own predictions (for reviews, see Schultz, 1998, 2002), and to motor areas in anterior cingulate cortex, which use the error signals to improve performance on the task at hand. Furthermore, the theory proposes that variation in ERN amplitude is produced by the impact of these phasic dopamine error signals on anterior cingulate cortex, with phasic decreases in dopamine associated with larger ERN amplitudes and phasic increases in dopamine associated with smaller ERN amplitudes (see Holroyd, in press).

The results of the present study fit naturally within this theoretical framework. According to the theory, when presented with three possible outcomes such that the objective value of the middle outcome is halfway between the objective values of the best and worst outcomes, the system that produces the ERN will come to predict the middle value. Furthermore, departures from this expected outcome will produce relative increases or decreases in ERN amplitude, as observed. This computational process establishes context dependence: An outcome is judged to be good or bad relative to the expectation, rather than in terms of its objective value. These results dovetail with those of a recent experiment by Yeung and Sanfey (2003), who showed that ERN amplitudes elicited by large losses in the context of large gains and losses are about the same size as ERN amplitudes elicited by small losses in the context of small gains and losses. Thus, the evaluative system that produces the ERN appears to normalize its evaluations according to the best and worst outcomes judged possible in the task: Losing \$1,000 when \$1,000 is the maximum loss possible appears to be “just as bad” as losing \$2,000 when \$2,000 is the maximum loss possible.

An important unresolved issue concerns the fact that, in both of the experiments in this study, the ERNs associated with the worst and middle outcomes were not significantly different from each other. This finding may have occurred either because the study lacked sufficient statistical power to distinguish the size of the ERN between the two conditions or because the value function associated with the ERN is actually nonlinear, weighting the worst and middle outcomes more or less equally. Moreover, it may have been the case that the participants, believing that they had discovered contingencies between their responses and the feedback stimuli, tended to predict that their responses would achieve maximum reward (rather than the average reward); if so, then the intermediate outcomes would have been worse than they expected and would have produced ERNs. This question remains a subject for future research.

Other predictions of the theory have been tested in a recent series of trial-and-error and pseudo trial-and-error learning experiments (cf. Badgaiyan & Posner, 1998; Luu, Tucker, Derryberry, Reed, & Poulsen, 2003; Ruchow, Grothe, Spitzer, & Kiefer, 2002; Wang, Kong, Tang, Zhuang, & Li, 2000). According to the theory, for example, the monitoring system comes to expect reward when rewards are frequent, and to expect nothing when rewards are infrequent; thus the ERN elicited by

an absence of an expected reward when rewards are frequent should be larger than the ERN elicited by an absence of an expected reward when rewards are infrequent. This prediction was confirmed in a recent experiment (Holroyd et al., in press). In another study, in a condition in which the rewards and punishments were random and equiprobable, the amplitude of the ERN tracked the reward prediction error on a trial-by-trial basis; consistent with the theory, the system appeared to use the feedback information on each trial to update its predictions for future trials (Holroyd & Coles, 2002). Even in a task without feedback, ERN amplitude at the time of error commission increased as the probability of making the error decreased (Holroyd & Coles, 2002), suggesting that infrequent errors are associated with large negative prediction errors. The theory also holds that, when the stimulus–response mappings are fixed and can be learned, the ERN should occur at the time of the error response and not at the time of feedback presentation. According to the theory, on such trials the system detects the error at the time of the response, producing a negative prediction error and a large ERN, and so by the time of feedback presentation the system has already detected the error, thus producing no change in prediction and no ERN—a prediction that has also been confirmed (Holroyd & Coles, 2002; Nieuwenhuis et al., 2002).

Other evidence indicates that the system that produces the ERN does not depend on whether the feedback indicates a gain or loss per se, but rather on any performance-related information conveyed by the feedback. As a case in point, Gehring and Willoughby (2002) recently conducted a gambling task in which the feedback stimuli on each trial conveyed information along two different dimensions: a “gain/loss” dimension indicating monetary gains or losses, and a “correct/error” dimension indicating the degree of correctness or incorrectness of the response. The authors inferred from their data that a “medial frontal negativity” (MFN) was sensitive to the gain/loss information conveyed by the feedback (being larger to feedback indicating loss than to feedback indicating gain), and that the MFN was insensitive to the correct/error information conveyed by the feedback. Although the MFN and the feedback ERN share similar latencies (about 250–300 ms following feedback onset) and scalp distributions (frontal-central), the results of this study seemed to suggest that the MFN was a new component that was distinct from the feedback ERN. In a separate experiment, however, we replicated the original results reported by Gehring and Willoughby and demonstrated that the negativity elicited by the feedback stimuli in the gambling task is sensitive to both the gain/loss information and the correct/error information conveyed by the feedback stimuli, depending on which aspect of the feedback stimuli is made most salient to the participants (Nieuwenhuis, Yeung, Holroyd, Schurger, & Cohen, in press; see also Holroyd, Coles, et al., 2002). These results indicate that the system that produces the ERN can base its evaluations on different types of performance-related

information, and that the type of information can be determined by the environmental context in which the information is provided.

It is worth noting that studies that have investigated the biological mechanisms that give rise to the ERN have also provided support for the theory (see also Holroyd & Coles, 2002). The theory is predicated on the observation that mesencephalic dopaminergic neurons appear to carry a reward prediction error signal (Schultz, 1998, 2002). Importantly, this proposal leads to the straightforward prediction that disruption of the mesencephalic dopamine system should affect the ERN. In fact, several studies have provided evidence consistent with this prediction. For example, administration of d-amphetamine, which releases dopamine and inhibits its reuptake, increases ERN amplitude (De Bruijn, Hulstijn, Verkes, Ruijt, & Sabbe, 2003). Conversely, alcohol consumption reduces ERN amplitude (Ridderinkhof et al., 2002), possibly because alcohol affects dopamine receptors (Holroyd & Yeung, 2003). ERN amplitude also decreases with age, an observation that has been attributed to age-related changes in dopamine function (Nieuwenhuis et al., 2002). In theory, Parkinson’s disease, which damages the mesencephalic dopamine system, should also reduce the amplitude of the ERN, but evidence to this effect has been mixed, with experiments indicating both a reduction (Falkenstein et al., 2001) and no reduction (Holroyd, Praamstra, Plat, & Coles, 2002). The theory also holds that the response ERN and the feedback ERN are produced by the impact of these dopamine error signals on the same area of anterior cingulate cortex. In fact, the results of several equivalent dipole source localization experiments have suggested that anterior cingulate cortex represents a common source for both the response ERN (e.g., Dehaene, Posner, & Tucker, 1994; Holroyd, Dien, & Coles, 1998) and the feedback ERN (Gehring & Willoughby, 2002; Miltner et al., 1997). Furthermore, a recent functional magnetic resonance imaging study has demonstrated that a single area in dorsal anterior cingulate cortex activates for both error responses and error feedback in a trial-and-error learning task (Holroyd et al., 2003). Moreover, activity of this region of dorsal anterior cingulate cortex is greatest in the same conditions in which ERN amplitude is also largest, namely, when the negative prediction error is large (cf. Holroyd & Coles, 2002).

In summary, when taken together, these studies provide support for the hypothesis that the ERN is produced by a predictive error signal conveyed by the mesencephalic dopamine system to anterior cingulate cortex. The present study contributes to that body of knowledge by demonstrating that the system that produces the ERN bases its evaluations not on the objective value of the outcome that was obtained, but rather on the value of the outcome relative to the range of other outcomes possible in the task. This research promises to inform the study of how evaluative mechanisms apply performance-related error information for behavioral adaptation.

REFERENCES

- Badgaiyan, R. D., & Posner, M. I. (1998). Mapping the cingulate cortex in response selection and monitoring. *Neuroimage*, 7, 255–260.
- Breiter, H. C., Aharon, I., Kahneman, D., Dale, A., & Shizgal, P. (2001). Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron*, 30, 619–639.
- Bush, G., Vogt, B. A., Holmes, J., Dale, A. M., Greve, D., Jenike, M. A., & Rosen, B. R. (2002). Dorsal anterior cingulate cortex: A role in reward-based decision making. *Proceedings of the National Academy of Sciences, USA*, 99, 523–528.
- Campbell, K. B., Courchesne, E., Picton, T. W., & Squires, K. C. (1979). Evoked potential correlates of human information processing. *Biological Psychology*, 8, 45–68.
- Davey, G., & Cullen, C. (Eds.). (1988). *Human operant conditioning and behavioral modification*. New York: John Wiley & Sons.

- De Bruijn, E. R. A., Hulstijn, W., Verkes, R. J., Ruigt, G. S. F., & Sabbe, B. G. C. (2003). Drug-induced stimulation and suppression of action monitoring. Manuscript submitted for publication.
- Dehaene, S., Posner, M. I., & Tucker, D. M. (1994). Localization of a neural system for error detection and compensation. *Psychological Science*, *5*, 303–305.
- Delgado, M. R., Locke, H. M., Stenger, V. A., & Fiez, J. A. (2003). Dorsal striatum responses to reward and punishment: Effects of valence and magnitude manipulations. *Cognitive, Affective, and Behavioral Neuroscience*, *3*, 27–38.
- Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences*, *11*, 355–372.
- Elliott, R., Newman, J. L., Longe, O. A., & Deakin, J. F. W. (2003). Differential response patterns in the striatum and orbitofrontal cortex to financial reward in humans: A parametric functional magnetic resonance imaging study. *Journal of Neuroscience*, *23*, 303–307.
- Falkenstein, M., Hielscher, H., Dziobek, I., Schwarzenau, P., Hoormann, J., Sundermann, B., & Hohnsbein, J. (2001). Action monitoring, error detection, and the basal ganglia: An ERP study. *NeuroReport*, *12*, 157–161.
- Falkenstein, M., Hohnsbein, J., Hoormann, J., & Blanke, L. (1990). Effects of errors in choice reaction tasks on the ERP under focused and divided attention. In C. Brunia, A. Gaillard, & A. Kok (Eds.), *Psychophysiological brain research* (pp. 192–195). Tilburg, The Netherlands: Tilburg University Press.
- French, S. (1988). *Models of decision and choice*. New York: John Wiley & Sons.
- Gabriel, M. (1993). Discriminative avoidance learning: A model system. In B. A. Vogt & M. Gabriel (Eds.), *Neurobiology of cingulate cortex and limbic thalamus: A comprehensive handbook* (pp. 478–523). Boston: Birkhauser.
- Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E., & Donchin, E. (1993). A neural system for error detection and compensation. *Psychological Science*, *4*, 385–390.
- Gehring, W. J., & Willoughby, A. R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science*, *295*, 2279–2282.
- Gratton, G., Coles, M. G. H., & Donchin, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography and Clinical Neurophysiology*, *55*, 468–484.
- Holroyd, C. B. (in press). A note on the N200 and the feedback ERN. In M. Ullsperger & M. Falkenstein (Eds.), *Errors, conflicts, and the brain: Current opinions on performance monitoring*. Leipzig: MPI of Cognitive Neuroscience.
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, *109*, 679–709.
- Holroyd, C. B., Coles, M. G. H., & Nieuwenhuis, S. (2002). Medial prefrontal cortex and error potentials. *Science*, *296*, 1610–1611.
- Holroyd, C. B., Dien, J., & Coles, M. G. H. (1998). Error-related scalp potentials elicited by hand and foot movements: Evidence for an output-independent error-processing system in humans. *Neuroscience Letters*, *242*, 65–68.
- Holroyd, C. B., Nieuwenhuis, S., Yeung, N., & Cohen, J. D. (in press). Reward prediction errors are reflected in the event-related brain potential. *NeuroReport*.
- Holroyd, C. B., Nieuwenhuis, S., Yeung, N., Nystrom, L., Mars, R. B., Coles, M. G. H., & Cohen, J. D. (2003). Dorsal anterior cingulate cortex is sensitive to internal and external sources of error formation: A functional magnetic resonance imaging study. Manuscript submitted for publication.
- Holroyd, C. B., Praamstra, P., Plat, E., & Coles, M. G. H. (2002). Spared error-related potentials in mild to moderate Parkinson's disease. *Neuropsychologia*, *40*, 2116–2124.
- Holroyd, C. B., & Yeung, N. (2003). Alcohol and error processing. *Trends in Neuroscience*, *26*, 402–404.
- Horst, R. L., Johnson, R., & Donchin, E. (1980). Event-related brain potentials and subjective probability in a learning task. *Memory & Cognition*, *8*, 476–488.
- Jasper, H. H. (1958). The ten twenty electrode system of the international federation. *Electroencephalography and Clinical Neurophysiology*, *10*, 371–375.
- Johnson, R., & Donchin, E. (1978). On how P300 amplitude varies with the utility of the eliciting stimuli. *Electroencephalography and Clinical Neurophysiology*, *44*, 424–437.
- Johnston, V. S. (1979). Stimuli with biological significance. In H. Begleiter (Ed.), *Evoked brain potentials and behavior* (pp. 1–12). New York: Plenum.
- Keselman, H. J., & Rogan, J. C. (1980). Repeated measures *F* tests and psychophysiological research: Controlling the number of false positives. *Psychophysiology*, *17*, 499–503.
- Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *Journal of Neuroscience*, *21*, RC159, 1–5.
- Luu, P., Tucker, D. M., Derryberry, D., Reed, M., & Poulsen, C. (2003). Electrophysiological responses to errors and feedback in the process of action regulation. *Psychological Science*, *14*, 47–53.
- MacKay, D. M. (1984). Do 'evaluation potentials' reflect cognitive assessment? *Experimental Brain Research*, *55*, 184–186.
- McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, *38*, 339–346.
- Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a "generic" neural system for error detection. *Journal of Cognitive Neuroscience*, *9*, 788–798.
- Monchi, O., Petrides, M., Petre, V., Worsley, K., & Dagher, A. (2001). Wisconsin card sorting revisited: Distinct neural circuits participating in different stages of the task identified by event-related functional magnetic resonance imaging. *Journal of Neuroscience*, *21*, 7733–7741.
- Nieuwenhuis, S., Ridderinkhof, K. R., Talsma, D., Coles, M. G. H., 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. *Cognitive, Affective, and Behavioral Neuroscience*, *2*, 19–36.
- Nieuwenhuis, S., Yeung, N., Holroyd, C. B., Schurger, A., & Cohen, J. D. (in press). Sensitivity of electrophysiological activity from medial frontal cortex to utilitarian and performance feedback. *Cerebral Cortex*.
- O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, *28*, 329–337.
- Picton, T. W., Hillyard, S. A., & Galambos, R. (1976). Habituation and attention in the auditory system. In W. D. Keidel & W. D. Neff (Eds.), *Handbook of sensory physiology: Vol. V/3. Auditory system, clinical and special topics*. Berlin: Springer-Verlag.
- Reynolds, J. N. J., Hyland, B. I., & Wickens, J. R. (2001). A cellular mechanism of reward-related learning. *Nature*, *413*, 67–70.
- Ridderinkhof, K. R., de Vlugt, Y., Bramlage, A., Spaan, M., Elton, M., Snel, J., & Band, G. P. H. (2002). Alcohol consumption impairs detection of performance errors in mediofrontal cortex. *Science*, *298*, 2209–2211.
- Rogers, R. O., Owen, A. M., Middleton, H. C., Williams, E. J., Pickard, J. D., Sahakian, B. J., & Robbins, T. W. (1999). Choosing between small, likely rewards and large, unlikely rewards activates inferior and orbital prefrontal cortex. *Journal of Neuroscience*, *19*, 9029–9038.
- Ruchkin, D. S., Sutton, S., Munson, R., Silver, K., & Macar, F. (1981). P300 and feedback provided by the absence of the stimulus. *Psychophysiology*, *18*, 271–282.
- Ruchow, M., Grothe, J., Spitzer, M., & Kiefer, M. (2002). Human anterior cingulate cortex is activated by negative feedback: Evidence from event-related potentials in a guessing task. *Neuroscience Letters*, *325*, 203–206.
- Schoenbaum, G., Chiba, A. A., & Gallagher, M. (1998). Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature Neuroscience*, *1*, 155–159.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1–27.
- Schultz, W. (2002). Getting formal with dopamine and reward. *Neuron*, *36*, 241–263.
- Shidara, M., & Richmond, B. J. (2002). Anterior cingulate: Single neuronal signals related to degree of reward expectancy. *Science*, *296*, 1709–1711.
- Shima, K., & Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science*, *282*, 1335–1338.

- Squires, K. C., Hillyard, S. A., & Lindsay, P. H. (1973). Cortical potentials evoked by confirming and disconfirming feedback following an auditory discrimination. *Perception & Psychophysics*, *13*, 25–31.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Sutton, S., Tueting, P., Hammer, M., & Hakerem, G. (1978). Evoked potentials and feedback. In D. Otto (Ed.), *Multidisciplinary perspectives in event-related brain potential research* (pp. 184–188). (Publication No. EPA-600/9-77-043). Washington, DC: U.S. Government Printing Office.
- Takasawa, N., Takino, R., & Yamazaki, K. (1990). Event-related potentials derived from feedback tones during motor learning. *Japanese Journal of Physiological Psychology and Psychophysiology*, *8*, 95–101.
- Thorndike, E. L. (1911/1970). Laws and hypotheses for behavior. In E. L. Thorndike (Ed.), *Animal intelligence* (pp. 241–281). Darien, CT: Hafner Publishing Co.
- Thut, G., Schultz, W., Roelcke, U., Nienhusmeier, M., Missimer, J., Maguire, R. P., & Leenders, K. L. (1997). Activation of the human brain by monetary reward. *NeuroReport*, *8*, 1225–1228.
- Ullsperger, M., & von Cramon, D. Y. (2003). Error monitoring using external feedback: Specific roles of the habenular complex, the reward system, and the cingulate motor area revealed by fMRI. *Journal of Neuroscience*, *23*, 4308–4314.
- Van Veen, V., Holroyd, C. B., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2002). Errors without conflict don't activate the anterior cingulate cortex: Implications for performance monitoring theories. Program No. 16.1 2002 *Abstract Viewer/Itinerary Planner*. Washington, DC: Society for Neuroscience 2002. Available at: <http://sfn.scholarone.com/itin2002/>.
- Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature*, *412*, 43–48.
- Wang, Y., Kong, J., Tang, X., Zhuang, D., & Li, S. (2000). Event-related potential N270 is elicited by mental conflict processing in human brain. *Neuroscience Letters*, *293*, 17–20.
- Yeung, N., & Sanfey, A. (2003). Discrete coding of magnitude and valence in the human brain. Manuscript submitted for publication.
- Zalla, T., Koechlin, E., Pietrini, P., Basso, G., Aquino, P., Sirigu, A., & Grafman, J. (2000). Differential amygdala responses to winning and losing: A functional magnetic resonance imaging study in humans. *European Journal of Neuroscience*, *12*, 1764–1770.

(RECEIVED February 12, 2003; ACCEPTED September 19, 2003)