When is an error not a prediction error?
An electrophysiological investigation

CLAY B. HOLROYD, OLAVE E. KRIGOLSON, ROBERT BAKER, SEUNG LEE, AND JESSICA GIBSON
University of Victoria, Victoria, British Columbia, Canada

A recent theory holds that the anterior cingulate cortex (ACC) uses reinforcement learning signals conveyed by the midbrain dopamine system to facilitate flexible action selection. According to this position, the impact of reward prediction error signals on ACC modulates the amplitude of a component of the event-related brain potential called the *error-related negativity* (ERN). The theory predicts that ERN amplitude is monotonically related to the expectedness of the event: It is larger for unexpected outcomes than for expected outcomes. However, a recent failure to confirm this prediction has called the theory into question. In the present article, we investigated this discrepancy in three trial-and-error learning experiments. All three experiments provided support for the theory, but the effect sizes were largest when an optimal response strategy could actually be learned. This observation suggests that ACC utilizes dopamine reward prediction error signals for adaptive decision making when the optimal behavior is, in fact, learnable.

The anterior cingulate cortex (ACC) constitutes a central component of a neural system for cognitive control (Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004). Theories about ACC function alternatively hold that it is responsible for detecting response conflict (Botvinick, Braver, Barch, Carter, & Cohen, 2001) or for the adaptive modification of behavior by relating actions with their consequences (Holroyd & Coles, 2002, 2008; Rushworth, Walton, Kennerley, & Bannerman, 2004). Notably, a recent theory holds that ACC uses reinforcement learning signals conveyed by the midbrain dopamine system to update the response selection process (Holroyd & Coles, 2002). This theory proposes that the impact of these dopamine signals on the ACC modulates the amplitude of a component of the event-related brain potential (ERP) called the *error-related negativity* (ERN). The ERN appears to exist in at least two forms: The "response ERN" (rERN) reflects a difference in neural processing between error responses and correct responses in speeded response time (RT) tasks (Falkenstein, Hohnsbein, Hoormann, & Blanke, 1990; Gehring, Goss, Coles, Meyer, & Donchin, 1993), and the "feedback ERN" (fERN) reflects a difference in neural processing between error feedback and correct feedback in guessing and trial-and-error learning tasks (Holroyd & Krigolson, 2007; Miltner, Braun, & Coles, 1997; for reviews, see Holroyd, Nieuwenhuis, Mars, & Coles, 2004; Nieuwenhuis, Holroyd, Mol, & Coles, 2004). Furthermore, experiments in the monkey have demonstrated that the midbrain dopamine system carries "reward prediction error signals"; phasic increases and decreases from baseline firing rate follow unexpected rewards and absences of reward, respectively (Schultz, 2002), which are then used by the targets of the dopamine system for the purpose of reinforcement learning (Montague, Hyman, & Cohen, 2004). Because of this property, the dopamine/ERN theory—sometimes called the *reinforcement learning theory of the ERN*, or RL–ERN theory—predicts that fERN amplitude should be monotonically related to the unexpectedness of the outcome: When measured as the difference between the ERPs associated with correct and error trials, relatively large fERNs will be elicited by unexpected outcomes, whereas relatively small fERNs will be elicited by expected outcomes (Holroyd & Coles, 2002; Holroyd & Krigolson, 2007; see also Holroyd, Pakzad-Vaezi, & Krigolson, 2008).

There is now a great deal of evidence that the ERN indexes a reward prediction error signal (Baker & Holroyd, in press; Butterfield & Mangels, 2003; Cohen & Ranganath, 2007; Compton et al., 2007; Donkers & van Boxtel, 2005; Dunning & Hajcak, 2007; Eppinger, Kray, Mock, & Mecklinger, 2008; Hajcak, Moser, Holroyd, & Simons, 2007; Hewig et al., 2007, 2008; Holroyd & Coles, 2002, 2008; Holroyd & Krigolson, 2007; Holroyd, Nieuwenhuis, Yeung, & Cohen, 2003; Krigolson & Holroyd, 2007; Krigolson, Mathewson, Baker, Baker, & Holroyd, 2006; Krigolson, Pierce, Holroyd, & Tanaka, in press; Morris, Heerey, Gold, & Holroyd, 2008; Nieuwenhuis, Nielen, Mol, Hajcak, & Veltman, 2005; Nieuwenhuis et al., 2002; Yasuda, Sato, Miyawaki, Kumano, & Kuboki, 2004). Nevertheless, a study by Hajcak, Holroyd, Moser, and Simons (2005) failed to confirm this prediction (see also Donkers, Nieuwenhuis, & van Boxtel, 2005; Larson, Kelly, Stigge-Kaufman, Schmalfuss, & Perlstein, 2007). In that study, participants were asked to complete two guess-

C. B. Holroyd, holroyd@uvic.ca

ing tasks. In a first experiment, on each trial, participants were presented with a cue that indicated whether one, two, or three out of four doors presented on a computer display hid a reward, and they were then required to choose one of the doors. Thus, the cue indicated a 25%, 50%, or 75% chance of finding a reward on each trial. Contrary to the prediction of the RL–ERN theory, although the amplitude of the fERN was differentially modulated by reward and nonreward feedback, the degree of this modulation was not affected by reward probability. Likewise, in a second experiment, participants selected among four "balloons" displayed on a computer screen. In each block, participants were rewarded at random on 25%, 50%, or 75% of the trials, and could thus infer whether they were relatively more or less likely to find the reward. Nevertheless, as with the first experiment, reward and nonreward feedback differentially affected fERN amplitude, but this effect was not modulated by reward probability. Note that although statistical power could be a factor underlying this result, the number of participants in the Hajcak et al. (2005) study were comparable to that in most fERN experiments.

At first blush, the RL–ERN theory would appear to apply to these guessing tasks, just as it applies to any task involving performance feedback. Thus, this failure to confirm the central prediction of the RL–ERN theory would seem to undermine its validity. Alternatively, given the evidence cited previously in favor of the theory, we think it likely that the theory is correct but that the tasks used by Hajcak et al. (2005) lie outside its domain of applicability.[1] Despite preserving the theory, however, this possibility would curtail the theory's predictive power: Until the theory's domain is better characterized, its continued application would risk yielding still more negative results. For this reason, we have sought to identify the essential factor or factors that distinguish the tasks used by Hajcak et al. (2005) from other tasks that have provided support for the RL–ERN theory.

In the present article, we investigated this issue in a series of three experiments.[2] Our empirical approach began with a task that closely resembled those used by Hajcak et al. (2005); then, we systematically modified the task so that it more closely resembled those that provided support for the RL–ERN theory. In Experiment 1, we tested the hypothesis raised by Hajcak et al. (2005) that fERN amplitude may be nonlinearly related to the size of the reward prediction error. That is, participants may find it easier to discriminate larger probability differences than smaller probability differences, so fERN amplitude may be relatively more exercised by extremely unexpected events than by moderately unexpected events. In Experiment 2, we tested the hypothesis that the reward prediction effect depends on a feature of many successful fERN tasks—namely, that participants learn to associate the reward probabilities with arbitrary predictive stimuli, rather than extract the probabilities from explicitly provided cues (as was the case in the Hajcak et al., 2005 study). Finally, in Experiment 3, we tested the hypothesis that the reward-prediction error effect occurs in tasks with meaningful and consistent stimulus–response mappings—which is also characteristic of other fERN studies—as opposed

to in tasks in which the feedback stimuli are delivered at random.

EXPERIMENT 1

In the Hajcak et al. (2005) study, participants selected between four "doors" displayed on a computer screen to find a reward (USD $0.10) on each trial. At the start of each trial, a predictive cue indicated whether rewards would be hidden behind one, two, or three doors, indicating reward probabilities of 25%, 50%, or 75%, respectively. The fERN was measured across three conditions using the difference wave approach (Holroyd & Krigolson, 2007). For each participant, the ERP associated with unexpected reward trials (rewards following the 25% cue) was subtracted from the ERP associated with unexpected nonreward trials (nonrewards following the 75% cue), creating an "unexpected difference wave." The ERP associated with expected reward trials (rewards following the 75% cue) was subtracted from the ERP associated with expected nonreward trials (nonrewards following the 25% cue), creating an "expected difference wave"; and the ERP associated with intermediate reward trials (rewards following the 50% cue) was subtracted from the ERP associated with intermediate nonreward trials (nonrewards following the 50% cue), creating a "control difference wave." The RL–ERN theory predicts that fERN amplitude should be monotonically related to the expectedness of the outcome, so that the unexpected difference wave should be larger than the control difference wave, which in turn should be larger than the expected difference wave. However, this was not the case: Hajcak et al. (2005) found that the difference waves reached about the same maximum amplitude, irrespective of condition. As was noted above, these unexpected results challenge the validity of the RL–ERN theory, which would appear to apply to this task just as it does to other tasks that have confirmed its predictions.

In the Hajcak et al. (2005) study, it was suggested that the sensitivity of fERN amplitude to reward probability might be highly nonlinear, such that probabilities between 25% and 75% were not extreme enough to modulate fERN amplitude significantly. In particular, participants may be relatively better at detecting extreme differences in probability than they are at detecting small differences in probability. To test this hypothesis, in Experiment 1, we skewed the reward probabilities even further, from 25%–75% to 5%–95%, so that on each trial, a predictive cue indicated whether 1, 10, or 19 out of 20 doors hid rewards (Figure 1, top). As before, the RL–ERN theory predicts that fERN amplitude should increase monotonically with the unexpectedness of the outcome.

Method

Participants. Twenty undergraduate students at the University of Victoria (8 male; mean age = 19.7 ± 2.1 years) participated in the experiment. All of the participants were volunteers who received extra credit in a first- or second-year psychology course for their participation, and who had provided written, informed consent. Participants also earned a small monetary bonus that depended on their performance (about CAN $5; see below). The study was conducted in accordance with the ethical standards prescribed in the Declara-

Experiment 1

| 1,500 msec | 2,000 msec Maximum | 200 msec | 500 msec |

Experiment 2

| 2,000 msec Maximum | 200 msec | 300 msec | 500 msec |

Experiment 3

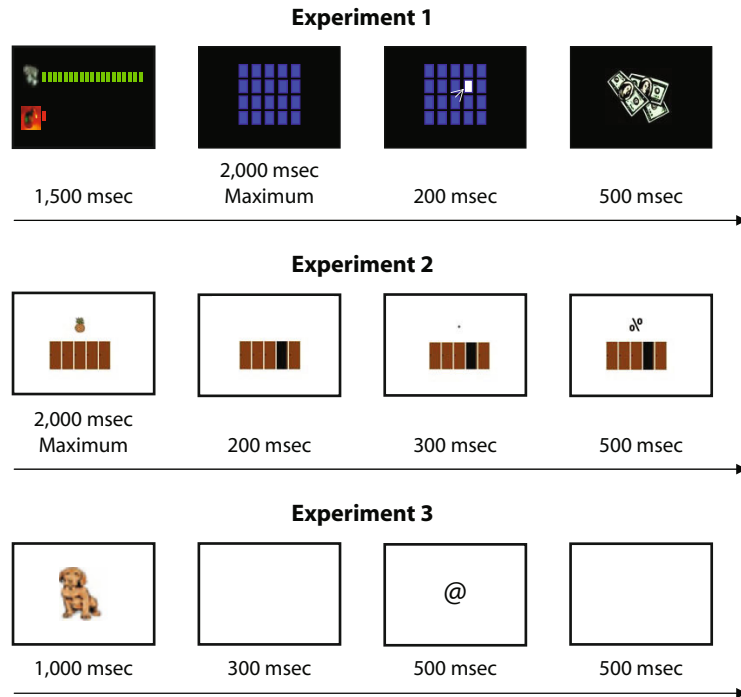| 1,000 msec | 300 msec | 500 msec | 500 msec |

Figure 1. Experimental tasks. An example trial is illustrated for each experiment; stimulus durations are given below each image. The intertrial interval for Experiments 1 and 2 was 1,000 msec (not shown).

tion of Helsinki and was approved by the human subjects review board at the University of Victoria.

Apparatus and Procedure. Participants were seated comfortably in front of a computer monitor in an electromagnetically shielded booth, and they engaged in a modified version of the cued door selection task utilized in Hajcak et al. (2005). On each trial, participants were presented with an image that indicated whether 1, 10, or 19 doors out of 20 hid rewards (1,500 msec; 8.5º of visual angle); these images were displayed as colored scales that indicated the relative probabilities of finding the reward (a picture of a dollar adjacent to a green scale) versus not finding the reward (a picture of a bomb adjacent to a red scale; Figure 1, top). Participants were informed of the meaning of these stimuli at the outset of the experiment; note that 1, 10, and 19 doors correspond to probabilities of 5%, 50%, and 95%, respectively. Then, participants were presented with an image of a 5 × 4 array of blue rectangles representing 20 "doors"; this image remained on the screen for 2,000 msec or until the participant selected a door, whichever came first (9º of visual angle; Figure 1, top). Response selection was carried out by using a mouse to direct a cursor in the form of a white arrow to the desired door and then pressing a mouse button. Immediately upon response selection, the chosen door changed color to white. This image—including the cursor, door array, and selected door—remained on the screen for 200 msec. Then, a visual feedback stimulus appeared (2º of visual angle, 500-msec duration) in the form of several dollar bills for reward, a bomb for no reward, and a turtle if participants responded too slowly, which was presented whenever the participant's RT exceeded 2,000 msec. Participants were told that they would earn CAN $0.25 by finding the reward, that they would be penalized CAN $0.23 if they did not find the reward, and that they would be penalized CAN $0.50 if they responded too slowly. All images were presented against a black background, and the intertrial interval (ITI) was 1,000 msec. Participants completed 10 blocks of 100 trials. Within each block, 4 trials corresponded to 5% probability (2 reward, 2 nonreward), 20 trials corresponded to 50% probability (10 reward, 10 nonreward), and 76 trials corresponded to 95% probability (38 reward, 38 nonreward);

reward and nonreward feedback were delivered at random without replacement, in accordance with the probability of the predictive cue on each trial, but independent of the response selected.

Data acquisition. The electroencephalogram (EEG) was recorded from 36 electrode locations using BrainVision Recorder software (Version 1.3, Brainproducts, Munich, Germany). The electrodes were mounted in a fitted cap with a standard 10–20 layout and were referenced to the average voltage across channels. The vertical and horizontal electrooculogram were recorded from electrodes placed above and below the right eye and on the outer canthi of the left and right eyes, respectively. Electrode impedances were kept below 10 kΩ. The EEG data were sampled at 250 Hz, amplified (Quick Amp, Brainproducts, Munich), and filtered through a passband of 0.017 Hz–67.5 Hz (90 dB/octave roll-off).

Data analysis. The EEG data were filtered offline through a 0.1 Hz–20 Hz passband phase-shift-free Butterworth filter and re-referenced to linked mastoids. Ocular artifacts were removed using the algorithm described by Gratton, Coles, and Donchin (1983). Trials in which the change in voltage at any channel exceeded 35 μVs per sampling point were also discarded. In total, fewer than 5% of the data were discarded. An 800-msec epoch of data (from 200 msec before the feedback stimulus onset to 600 msec following feedback stimulus onset) was extracted from the continuous EEG for each trial, channel, and participant. These epochs were baseline corrected relative to the 200-msec segment preceding feedback stimulus onset by subtracting the average value of the baseline from the potential at each sample in the epoch. ERPs were created by averaging the EEG data by condition for each electrode channel and participant.

To isolate variance in the ERP associated with the fERN from other overlapping ERP components, we created difference waves by subtracting each reward ERP from its appropriate corresponding nonreward ERP (Hajcak et al., 2005; Holroyd & Krigolson, 2007). Specifically, for each participant and channel, we created three fERN difference waves by (1) subtracting the reward ERP in the 50% condition from the nonreward ERP in the 50% condition, creating a "control" difference wave; (2) subtracting the reward ERP in the 5% re-

ward condition from the nonreward ERP in the 95% reward condition, creating an "unexpected" difference wave; and (3) subtracting the reward ERP in the 95% reward condition from the nonreward ERP in the 5% reward condition, creating an "expected" difference wave. The amplitude of each difference wave was measured for each participant and electrode as the most negative deflection within the 0–600 msec following feedback stimulus onset. Consistent with previous studies, fERN amplitude was evaluated at channel FCz, where it is normally maximal (Holroyd & Krigolson, 2007; Miltner et al., 1997).

Results

Figure 2 (left column) illustrates the ERPs recorded at channel FCz elicited by unexpected, control, and expected reward and nonreward feedback stimuli, respectively, and Figure 3 (top) illustrates the difference waves associated with these ERPs. Each of the difference waves reached maximum amplitude at channel FCz; Figure 4 (left column) shows their associated scalp distributions. An ANOVA performed across the three conditions indicated that fERN amplitude was modulated by condition [$F(2,19) = 7.57$,

$p < .01$, $\eta_p^2 = 0.28$]. The amplitude of the fERN was larger in the unexpected condition ($-6.8 \pm 0.6 \mu V$) than in the expected condition ($-4.3 \pm 0.3 \mu V$) [$t(19) = 3.5$, $p < .005$, Cohen's $d = 1.6$], and larger in the unexpected condition than in the control condition ($-4.8 \pm 0.4 \mu V$) [$t(19) = 2.4$, $p < .05$, Cohen's $d = 1.1$], but the control and expected conditions were not significantly different from each other [$t(19) = 1.2$, $p = .23$, Cohen's $d = 0.6$]. fERN latencies are indicated in Table 1 for each condition.

EXPERIMENT 2

The results of Experiment 1 confirmed our hypothesis that fERN amplitude would be modulated by the large reward prediction errors associated with the relatively extreme reward probabilities in this task (5%–95%), providing further support for the RL–ERN theory. However, the effect size associated with this finding was rather small ($\eta_p^2 = 0.28$), despite the fact that the study included 20 partici-
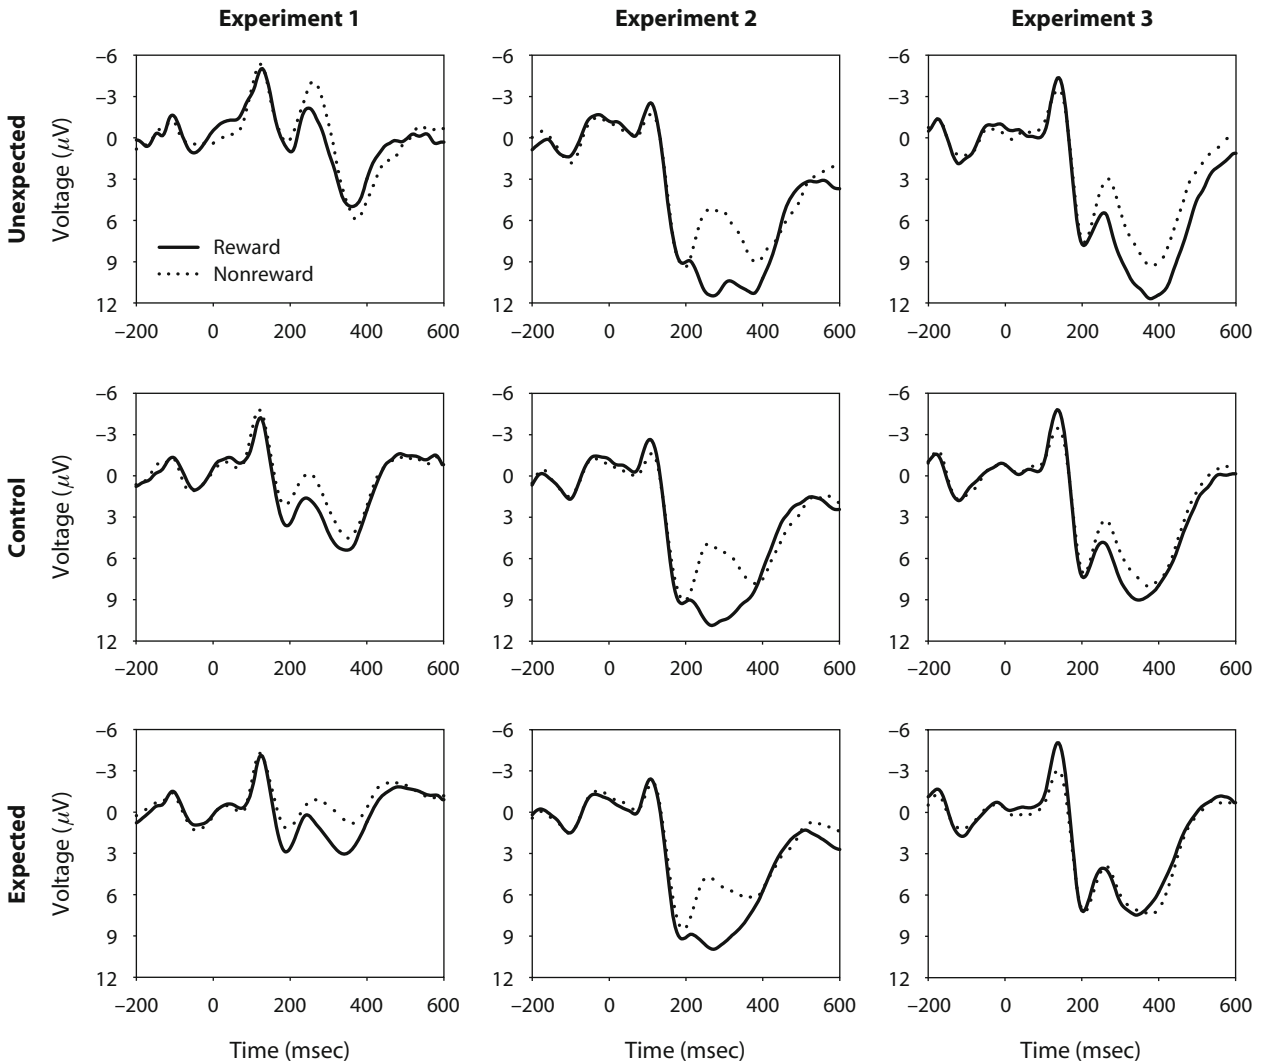

Figure 2. Event-related brain potentials elicited by reward and nonreward feedback stimuli for the unexpected, control, and expected conditions. Zero on each abscissa indicates time of feedback stimulus onset. Negative is plotted up by convention. Data recorded at channel FCz.
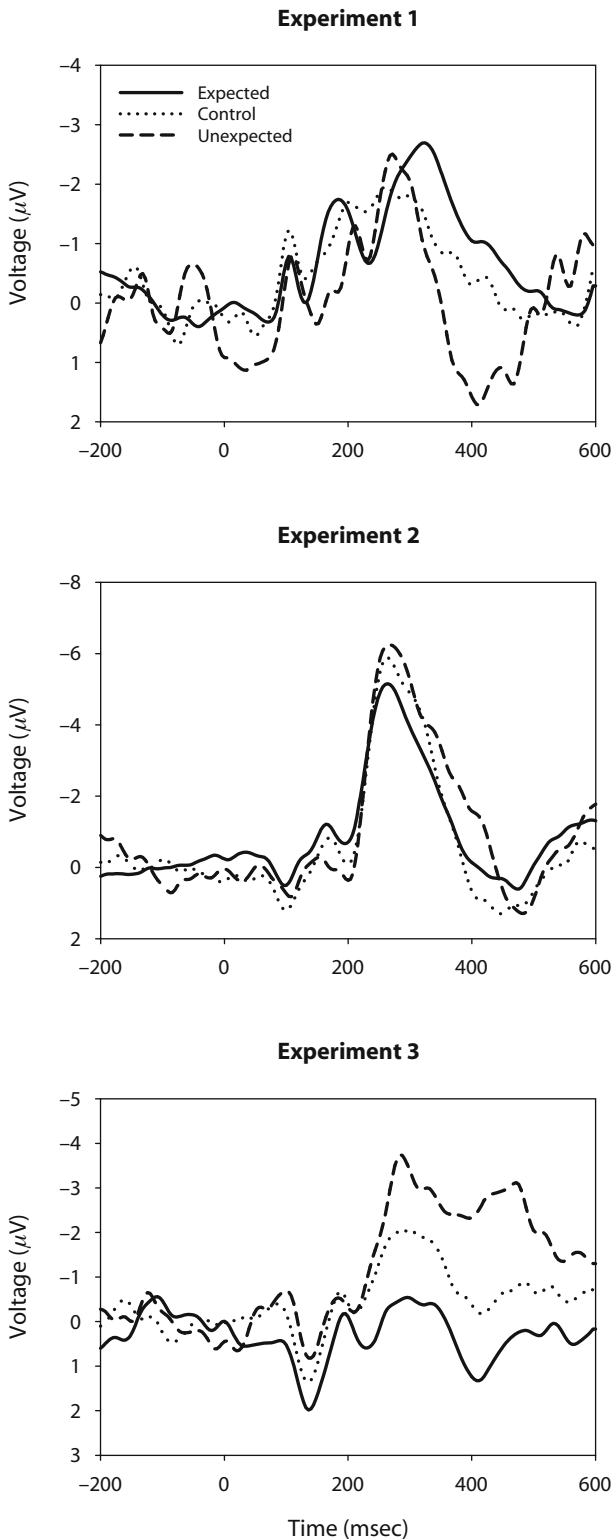
Experiment 1

Experiment 2

Experiment 3

Figure 3. Feedback error–related negativity difference waves associated with the unexpected, control, and expected conditions. Zero on each abscissa indicates time of feedback stimulus onset. Negative is plotted up by convention. Data are associated with channel FCz. Note the different scales on the ordinate for each experiment.

pants, more than the 10–15 participants typically utilized in other fERN studies. Moreover, although fERN amplitude was larger in the unexpected condition than in the expected condition, and was larger in the unexpected condition than in the control condition, it was not larger in the control condition than in the expected condition, so fERN amplitude was not monotonically related to reward probability, as predicted by the RL–ERN theory. In short, a study that used extreme probabilities together with a relatively large number of participants produced only modest results. These findings suggest that participants are not particularly sensitive to extremely unexpected outcomes in comparison with less unexpected outcomes; thus, the failure to confirm the predictions of the RL–ERN theory in the study by Hajcak et al. (2005) did not result from this possibility.

Hence, in Experiment 2, we tested a second hypothesis that might explain the failure in the Hajcak et al. (2005) study to replicate the standard fERN reward prediction error result. This hypothesis was predicated on the observation that in several previous fERN experiments, participants learned by trial and error how to respond to visual images of arbitrary objects (see, e.g., Eppinger et al., 2008; Holroyd & Coles, 2002; Morris et al., 2008; Nieuwenhuis, Nielen et al., 2005; Nieuwenhuis et al., 2002). We speculated that fERN amplitude would be modulated by reward probability when participants learn to associate reward probabilities with predictive cues, as opposed to when the probabilities are explicitly indicated by the predictive cues on the start of each trial (as in Experiment 1; Figure 1, top). As such, in Experiment 2, participants were shown visual images of arbitrary objects (specifically, a cherry, pineapple, banana, piano, guitar, trumpet, plane, bus, and boat; Figure 1, middle) at the start of each trial and, rather than being told explicitly the probabilities associated with each cue, the participants were required to infer the probabilities by trial and error. Additional, relatively minor task changes are described below.

Method

Except where noted, all of the procedures were identical to those in Experiment 1.

Participants. Fifteen undergraduate students at the University of Victoria (7 male; mean age = 20.5 ± 2.6 years) participated in the experiment.

Apparatus and Procedure. The following changes were made to the task used in Experiment 1.

1. As was noted previously, rather than show participants cues that explicitly indicated the probability of reward on that trial as in Experiment 1 (Figure 1, top), in Experiment 2, participants were shown visual images of arbitrary objects (cherry, pineapple, banana, piano, guitar, trumpet, plane, bus, and boat; Figure 1, middle). Participants were not told the probabilities associated with each cue and instead had to infer them by trial and error. The mappings between the objects and the reward conditions were fixed throughout each block so that the participants could, in fact, learn to associate each cue with a probability of reward. The experiment consisted of three blocks of trials, with 300 trials per block; a different set of three probability cues was associated with each block.

2. Because the extreme reward probabilities utilized in Experiment 1 were only modestly successful in modulating fERN amplitude, we assumed that this factor was not critical and thus relaxed this constraint. Accordingly, in Experiment 2, we adopted a narrower
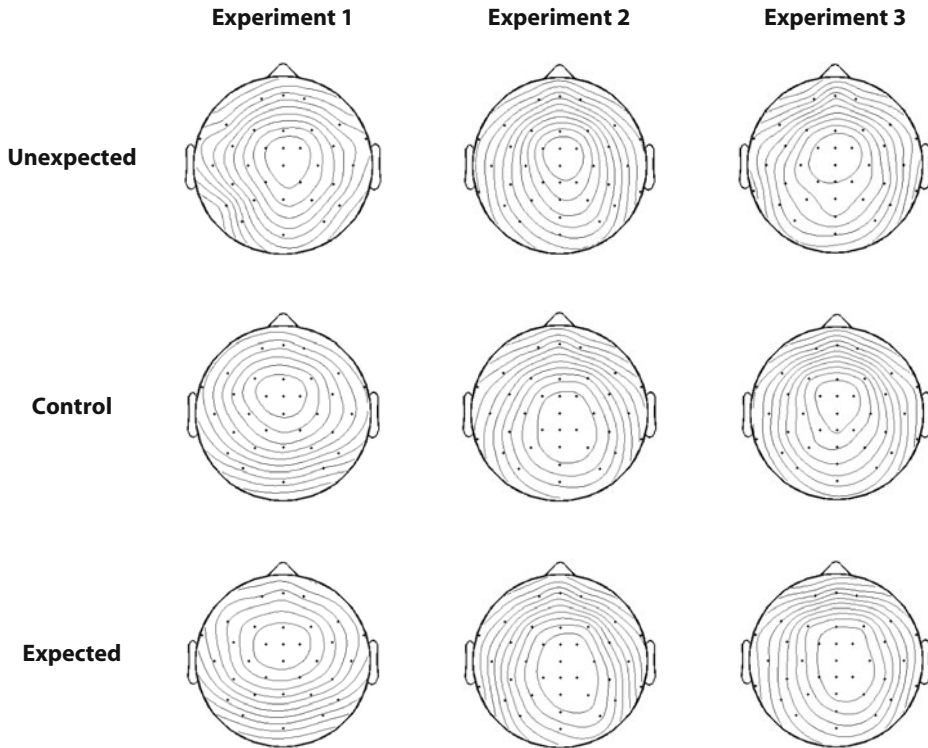
Experiment 1   **Experiment 2**   **Experiment 3**

Unexpected

Control

Expected

Figure 4. Scalp distributions associated with the feedback error–related negativity for the unexpected, control, and expected conditions.

range of reward probabilities (20%, 50%, 80%) that allowed more trials to be obtained per condition. Thus, on each trial, participants selected 1 of 5 doors rather than 1 of 20 doors (Figure 1, middle).

3. To ensure that the participants did not simply ignore the cues, the images of the objects and doors appeared together and remained on the display until either the participant made a response or 2,000 msec had elapsed, whichever occurred first (Figure 1, middle). This design contrasts with that of Experiment 1, where the predictive cue was removed before the door image appeared (Figure 1, top).

4. In order to not confound the feedback stimuli with the cue images, reward, nonreward, and "too slow" feedback consisted of the typographical symbols "%", "&", and "@", the mappings of which were systematically varied across participants. (Note that the response deadline was the same as in Experiment 1: 2,000 msec.) Participants earned CAN$0.02 for each reward feedback stimulus and earned nothing for each nonreward feedback stimulus; they were penalized CAN$0.10 for each too-slow feedback stimulus. On average, each participant earned about CAN$8.50 in bonus money.

5. To summarize, each trial began with the predictive cue image (1.3º of visual angle) presented directly above an array of five brown rectangles signifying five doors (8.5º of visual angle; Figure 1, middle). This image remained on the screen for 2,000 msec, or until the participant made a response. Response selection occurred by pressing one of five keys on a standard keyboard ("a," "s," "d," "f," and

"g"). By doing so, the selected door turned black and the predictive cue disappeared; this image remained on the screen for 200 msec. Then, a fixation cue (0.5º of visual angle) appeared above the doors, which together remained on the screen for a further 300 msec. Finally, the feedback stimulus was presented in the location of the previous fixation cue; the image of the feedback and doors remained on the screen for 500 msec. All images were presented against a white background, and the ITI was 1,000 msec. The experiment consisted of three 300-trial blocks. In each block, one of the three imperative stimuli was associated with reward on 20 trials and with nonreward on 80 trials (20% reward probability), a second stimulus was associated with reward on 50 trials and with nonreward on 50 trials (50% reward probability), and the third stimulus was associated with reward on 80 trials and with nonreward on 20 trials (80% reward probability).

Data acquisition. Data acquisition was the same as in Experiment 1.

Data analysis. ERPs were created in the same way as in Experiment 1. For each participant and channel, three fERN difference waves were created by (1) subtracting the reward ERP in the 50% condition from the nonreward ERP in the 50% condition, creating a control difference wave; (2) subtracting the reward ERP in the 20% reward condition from the nonreward ERP in the 80% reward condition, creating an unexpected difference wave; and (3) subtracting the reward ERP in the 80% reward condition from the nonreward ERP in the 20% reward condition, creating an expected difference wave.

Results

Figure 2 (middle column) illustrates the ERPs elicited by the unexpected, control, and expected reward and nonreward feedback, respectively. Figure 3 (middle) illustrates the difference waves associated with these ERPs, and Figure 4 (middle column) illustrates the associated scalp distributions. The difference waves reached maximum

Table 1
Means and Standard Deviations (in Milliseconds)
of Feedback ERN Latencies

	Experiment 1		Experiment 2		Experiment 3	
	M	SD	M	SD	M	SD
Frequent	291	17	275	10	298	34
Control	281	24	311	16	351	25
Infrequent	249	32	327	28	352	25

amplitude at channel FCz in the unexpected condition ($-8.2 \,\mu V$) and channel Cz in the control ($-7.5 \,\mu V$) and the expected ($-6.1 \,\mu V$) conditions; the difference between the voltages recorded at FCz and Cz was significant for the control condition [$t(14) = 3.25, p = .005$], but not for the expected condition ($p > .05$). An ANOVA performed across the three conditions indicated that fERN amplitude was modulated by condition [$F(2,14) = 4.6, p < .05, \eta_p^2 = 0.24$]. The amplitude of the fERN was larger in the unexpected condition ($-8.2 \pm 1.0 \,\mu V$) than in the expected condition ($-5.9 \pm 0.5 \,\mu V$) [$t(14) = 2.3, p < .05$, Cohen's $d = 1.2$], and it was larger in the control condition ($-6.9 \pm 0.6 \,\mu V$) than in the expected condition [$t(14) = 2.4, p < .05$, Cohen's $d = 1.3$], but the control and unexpected conditions were not significantly different from each other [$t(14) = 1.7, p = .11$, Cohen's $d = 0.9$]. fERN latencies are indicated in Table 1 for each condition.

EXPERIMENT 3

The results of Experiment 2 confirmed our hypothesis that fERN amplitude would depend on reward probability when participants are required to learn by trial and error to associate the probabilities with visual images of objects, as was found in previous experiments (e.g., Eppinger et al., 2008; Holroyd & Coles, 2002; Morris et al., 2008; Nieuwenhuis, Nielen et al., 2005; Nieuwenhuis et al., 2002). Nevertheless, as with Experiment 1, the effect size associated with this finding was modest ($\eta_p^2 = 0.24$). Furthermore, fERN amplitude was not monotonically related to reward probability, once again failing to confirm fully all of the predictions of the RL–ERN theory. Thus, the critical ingredient for modulating fERN amplitude as a function of reward probability does not appear to be associated with any of the factors that were changed in Experiments 1 and 2—namely, the degree of reward expectancy and whether reward probabilities are learned by trial and error.

Hence, in Experiment 3, we examined whether yet another factor might contribute to modulating fERN amplitude in the predicted manner. An important difference between the first two experiments and many other fERN studies is that the feedback was delivered entirely at random: Although individual participants may have believed that they discovered the optimal response strategy, no such strategy in fact existed. By contrast, in many other experiments, the type of feedback stimulus presented on each trial has depended partially on the participant's own behavior. For example, following certain stimulus cues in Holroyd and Coles's (2002) study, participants were always rewarded if they pressed one of two buttons and were always punished if they pressed the other button ("100% condition"), whereas following other cues in this study, participants were rewarded at random (50% probability), irrespective of the response actually made ("50% condition"). Consistent with the prediction of the RL–ERN theory, it was found that in the 100% condition, fERN amplitude decreased as participants learned the stimulus responses mappings, whereas in the 50% condition, fERN amplitude remained large throughout each block because the mappings could not be learned. Furthermore, in a vari-

ation of this paradigm, Nieuwenhuis et al. (2002) included an "80% condition" whereby the appropriate response was mapped to each cue with an 80% probability of reward. If, on these trials, the participants always responded with the optimal response, then they would be rewarded 80% of the time. In the Nieuwenhuis et al. (2002) study also, fERN amplitude was consistent with the predictions of the theory: In particular, the fERN was especially large in the condition in which the reward prediction error was biggest—namely, when participants selected the optimal response in the 80% condition but nevertheless received no-reward feedback.

Considering these observations, we hypothesized that a key ingredient missing from the designs of Experiments 1 and 2 was the dependence of the feedback on the participants' actual behavior. Although debriefing indicated that many participants believed that they had discovered an appropriate strategy for maximizing reward, we expected that such beliefs would be even stronger if the feedback did, in fact, depend on their behavior. Accordingly, we altered the task in Experiment 2 so that, as with previous studies, the appropriate response mappings were associated probabilistically with individual stimulus cues (Holroyd & Coles, 2002). Our design resembled that of Nieuwenhuis et al. (2002), which included an 80% condition in addition to the 100% and 50% conditions that were utilized in the study of Holroyd and Coles (2002).[3] Importantly, the 80% condition provides a strong test of the theory, since the feedback stimuli in this condition elicit relatively large reward prediction errors. Thus, we modified our experiments so that for each stimulus, an appropriate response could actually be learned, in the expectation that this change would increase the sensitivity of fERN amplitude to reward probability. As with Experiments 1 and 2, the RL–ERN theory predicts that fERN amplitude should increase monotonically with the unexpectedness of the outcome.

Method

The method was the same as that used in Experiment 2, except where noted.

Participants. Fifteen undergraduate students at the University of Victoria (9 male; mean age = 20.6 ± 1.5 years) participated in the experiment.

Apparatus and Procedure. On each trial, an imperative stimulus consisting of a visual image of an object (or animal) was displayed for 1,000 msec (1.3º of visual angle; Figure 1, bottom); participants were required to respond within these 1,000 msec. This stimulus was followed by a blank screen for 300 msec. Then, a feedback stimulus was displayed for 500 msec. Feedback consisted of "%", "@", and "&" symbols (1.3º of visual angle), indicating that participants earned CAN$0.02, lost CAN$0.00, or responded too slowly and lost CAN$0.10, respectively. The stimulus–feedback mappings were systematically varied across participants. All images were presented against a white background, and the ITI was 500 msec. On average, each participant earned about CAN$9 in bonus money.

Participants completed three blocks of 300 trials, and on each block, they were exposed to six novel visual stimuli. Participants responded by pressing a left mouse button for one response, and a right mouse button for the other. Critically, for each stimulus, one response never yielded a reward, whereas the other response yielded a reward on some proportion of the trials. Thus, participants could learn an optimal response for each stimulus. Furthermore, the prob-

ability of reward associated with each optimal stimulus–response combination depended on the stimulus; the participants could learn that some stimuli were more likely than other stimuli to yield reward if the optimal response was selected. Specifically, for each block, following one stimulus, the left button was rewarded on 80% of the trials, and the right button was never rewarded; and, following a second stimulus, the right button was rewarded on 80% of the trials, and the left button was never rewarded, so optimal responses following these stimuli were usually rewarded. Conversely, following a third stimulus, the left button was rewarded on 20% of the trials, and the right button was never rewarded; and, following a fourth stimulus, the right button was rewarded on 20% of the trials, and the left button was never rewarded, so optimal responses following these stimuli were rarely rewarded. Finally, following a fifth and a sixth stimulus, the left and the right buttons were rewarded, respectively, on 50% of the trials, whereas the alternative buttons were never rewarded, so optimal responses following these stimuli were rewarded on about half of the trials. Note that for each stimulus, one of the two possible responses was always partially rewarded (on 20%, 50%, or 80% of the trials), whereas the complementary response was never rewarded. Thus, participants could learn to choose the optimal response even in the 20% reward condition, in which the rewards were unlikely.

Data acquisition. Data acquisition was the same as in Experiments 1 and 2.

Data analysis. ERPs were created in the same way as in Experiments 1 and 2; only trials in which the optimal response was issued were analyzed. Three difference waves were formed by (1) subtracting the reward ERP in the 50% reward condition from the nonreward ERP in the 50% reward condition, creating a control difference wave; (2) subtracting the reward ERP in the 20% reward condition from the nonreward ERP in the 80% reward condition, creating an unexpected difference wave; and (3) subtracting the reward ERP in the 80% reward condition from the nonreward ERP in the 20% reward condition, creating an "expected" difference wave.

Results

Behavioral data. Participants chose the optimal response on a majority of trials in every condition: The optimal response was selected on 85.7% of the 80% reward trials [$t(14) = 19.6, p < .001$], on 80.0% of the 50% reward trials [$t(14) = 11.5, p < .001$], and on 68.0% of the 20% reward trials [$t(14) = 5.8, p < .001$]. RTs differed across the three conditions [$F(2,28) = 16.0, p < .001$] and were faster on 80% reward trials (499 msec) than on 50% reward trials (516 msec) [$t(14) = 2.55, p < .05$], which in turn were faster than on 20% reward trials (537 msec) [$t(14) = 2.77, p < .05$].

Electrophysiological data. Figure 2 (right column) shows the ERPs elicited by unexpected, control, and expected reward and nonreward feedback, respectively, and Figure 3 (bottom) illustrates the difference waves associated with these ERPs. Figure 4 (right column) depicts the scalp distributions of these difference waves. The difference waves reached maximum amplitude at channel FCz in the unexpected (-6.8μV) and control (-4.2μV) conditions, and at channel CP2 in the expected condition (-2.8μV), but the amplitude of the difference wave for the expected condition was not significantly different between channels FCz and CP2 ($p > .05$). An ANOVA performed across the three conditions indicated that fERN amplitude was modulated by condition [$F(2,14) = 22.7, p < .0001, \eta_p^2 = 0.62$]. The amplitude of the fERN was larger in the unexpected condition ($-6.8 \pm 0.6 \mu$V) than in the expected condition ($-2.7 \pm 0.4 \mu$V) [$t(14) = 6.1,$

$p < .0001$, Cohen's $d = 3.2$], larger in the unexpected condition than in the control condition ($-4.2 \pm 0.4 \mu$V) [$t(14) = 3.8, p < .005$, Cohen's $d = 2.0$], and larger in the control condition than in the expected condition [$t(14) = 3.3, p = .005$, Cohen's $d = 1.8$]. fERN latencies are indicated in Table 1 for each condition.

GENERAL DISCUSSION

The RL–ERN theory predicts that the amplitude of the fERN, measured as the maximal difference in the ERPs elicited by reward and nonreward feedback stimuli, will be larger for unexpected outcomes than for expected outcomes (Holroyd & Coles, 2002; see also Holroyd & Krigolson, 2007; Holroyd et al., 2008). Although the results of several experiments have supported this prediction, the results of other experiments have not (see, e.g., Donkers et al., 2005; Hajcak et al., 2005; Larson et al., 2007). These negative findings challenge the validity of the theory, curtail its predictive power, and undermine the position that the anterior cingulate cortex uses dopamine reward prediction error signals for the purpose of adaptive decision making (Holroyd & Coles, 2002, 2008). In the present article, we sought to better characterize the empirical domain to which the RL–ERN theory applies, specifically by investigating three possible reasons that might underlie the negative results reported by Hajcak et al. (2005).

First, Hajcak et al. (2005) suggested that fERN amplitude might be more sensitive to larger differences in reward probability than to smaller differences in reward probability, especially if participants fail to detect smaller differences. We tested this hypothesis in Experiment 1 by adopting a task that was modeled after that in Hajcak et al. (2005), but that included many more doors (20 instead of 4), enabling us to explore the consequence of relatively extreme reward probabilities (5%–95%) on the ERP. We reasoned that this modification would increase the variance in the size of the reward prediction error signals and concomitantly increase the variance in fERN amplitude. Nevertheless, our efforts were only partially successful. On the one hand, larger fERNs were, in fact, elicited by the unexpected outcomes as compared with the expected outcomes. On the other hand, the effect size associated with this result was relatively small ($\eta_p^2 = 0.28$), and the fERN was not significantly larger in a control condition than in the expected condition, which is inconsistent with the predictions of the RL–ERN theory. This result suggests that fERN amplitude is not particularly sensitive to extreme probabilities in comparison with moderate probabilities.

We then tested a second hypothesis, which was that fERN amplitude would be more sensitive to reward probabilities if these became associated with arbitrary visual stimuli by an associative learning process, as opposed to their being conveyed to the participants explicitly at the start of each trial. This hypothesis was motivated by the observation that many previous experiments that reported modulation of fERN amplitude by reward probability were constructed in this manner. For this reason, in Ex-

periment 2, we modified the task used in Experiment 1 so that instead of presenting predictive cues that expressly indicated the probability of reward (with scales that indicated the number of doors hiding a reward), we presented visual images of arbitrary objects that participants could learn to associate with particular reward probabilities. Nevertheless, this change was also only partially successful. Consistent with the RL–ERN theory, the fERN was larger following unexpected outcomes than it was following expected outcomes, but again, the statistical effect was relatively small ($\eta_p^2 = 0.24$). Furthermore, contrary to the prediction of the RL–ERN theory, the fERN was not larger for unexpected outcomes than for control outcomes. Finally, for the control condition, fERN amplitude was actually significantly larger at channel Cz than at channel FCz, which deviates somewhat from the typical fERN scalp distribution (Holroyd & Krigolson, 2007; Miltner et al., 1997). These results indicate that an associative learning process tying reward probability with arbitrary predictive cues is not sufficient for modulating fERN amplitude as a function of reward probability, despite the fact that such a learning process is characteristic of many previous fERN experiments.

Finally, in Experiment 3, we tested the hypothesis that fERN amplitude would be sensitive to reward probability in tasks in which stimulus–response contingencies could actually be learned. Thus, we modified the feedback contingencies so that the participants could truly learn an optimal response for every predictive cue. In contrast with the first two experiments, in which the type of feedback was delivered randomly and independently of each participant's choice, in Experiment 3, one of two possible response options for each stimulus was never rewarded. Thus, the alternative response for each stimulus was always the optimal choice, even though those choices were rewarded on only a subset of trials (from 20% to 80%, depending on the predictive cue). We found that participants learned to choose the optimal response across all three stimulus conditions. Further, the fERN results conformed to the predictions of the RL–ERN theory, since its amplitude increased monotonically with the unexpectedness of the outcome. In addition, the effect size for this result was more than twice as large as that observed in the first two experiments ($\eta_p^2 = 0.62$). These findings strongly suggest that fERN amplitude is relatively sensitive to reward probability in tasks in which an optimal response can actually be learned, as opposed to tasks in which the feedback is delivered independently of the participant's behavior. Note that this issue appears to be independent of the participants' certainty of whether each trial would end in success or failure, because—insofar as the participants learned to choose the optimal response in Experiment 3—the probability of reward was the same in Experiment 3 as it was in Experiment 2 (20%, 50%, and 80%), so both experiments should have been associated with comparable subjective probabilities.[4]

In a recent follow-up to their original door study, Hajcak et al. (2007) asked participants on each trial whether they believed that they would find the reward. In this way, fERN amplitude could be determined not only as a function of the objective reward probabilities indicated by the external cues, but also as a function of the subjective reward probabilities internalized by each participant. Interestingly, Hajcak et al. (2007) found that fERN amplitude did not correspond to the participants' own predictions when the participants were asked before they made their responses, but it did correspond to their predictions when they were asked after they made their responses. Thus, fERN amplitude was seen to vary with subjective expectation only when the prediction was made following a choice. The authors suggested that the predictions solidified on each trial after the participant committed to a particular action, perhaps because of increased attention to or confidence in the action–outcome associations. Similarly, a recent study of rat midbrain dopamine neurons demonstrated that reward predictions can be revised by the response-selection process (Roesch, Calu, & Schoenbaum, 2007)—an observation that seems consistent with the fact that the response ERN is seen immediately following error responses (Holroyd, Yeung, Coles, & Cohen, 2005). Taken with the present results, these findings suggest that the degree to which reward expectation modulates fERN amplitude depends strongly on whether each outcome can be attributed to a meaningful causal behavior. To the extent that participants do not perceive any action–outcome contingencies, fERN amplitude appears to be less sensitive to reward probability.

Note that this issue is different from the related question of whether fERN amplitude depends on behavior per se. In fact, several experiments have demonstrated that the fERN can be elicited by outcomes that are not immediately preceded by an overt behavior (Donkers & van Boxtel, 2005; Potts, Martin, Burton, & Montague, 2006; Yeung, Holroyd, & Cohen, 2005). Our present results extend this finding by indicating that variance in fERN amplitude as a function of reward probability depends on the degree to which the outcomes are perceived to be elicited by causal behaviors. Of course, it does not follow that perceived causality is a sufficient or even a necessary condition for modulating fERN amplitude; rather, it is undoubtedly one of several related factors that conspire to determine fERN amplitude in any given context. For example, fERN amplitude is correlated with participants' interest in the outcomes (Yeung et al., 2005). One might suppose that participants take greater interest in tasks in which they perceive a causal connection between their actions and outcomes, so that in these instances, fERN amplitude would be relatively sensitive to reward probability.

These results support the theory that the ACC uses reward prediction error signals conveyed by the midbrain dopamine system for the adaptive modification of behavior (Holroyd & Coles, 2002). Although the results of fMRI studies have been suggestive in this regard (Bush et al., 2002; Holroyd, Nieuwenhuis, Yeung, et al., 2004; Klein et al., 2007; Mars et al., 2005; Ullsperger & von Cramon, 2003; see also Behrens, Woolrich, Walton, & Rushworth, 2007; Nieuwenhuis, Slagter, von Geusau, Heslenfeld, & Holroyd, 2005; Paulus, Hozack, Frank, & Brown, 2002; Van Veen, Holroyd, Cohen, Stenger, & Carter, 2004), those of intracranial studies in both monkeys and humans have been particularly illuminating. This research has

highlighted the role played by the ACC in goal-directed decision making and in mediating flexible behavior associated with trial-and-error learning (Isomura, Ito, Akazawa, Nambu, & Takada, 2003; Isomura & Takada, 2004; Matsumoto & Tanaka, 2004; Procyk, Tanaka, & Joseph, 2000). Consistent with the RL–ERN theory, ACC neurons are sensitive to the degree of reward expectancy (Shidara & Richmond, 2002) and appear to be especially important for motor selection on the basis of anticipation of reward (Akkal, Bioulac, Audin, & Burbaud, 2002; Matsumoto, Suzuki, & Tanaka, 2003; Quilodran, Rothe, & Procyk, 2008). Furthermore, ACC neurons are sensitive to the omission of rewards (Ito, Stuphorn, Brown, & Schall, 2003; Niki & Watanabe, 1979; Shima & Tanji, 1998; Williams, Bush, Rauch, Cosgrove, & Eskandar, 2004), respond differentially to stimuli that predict rewarding and aversive events (Nishijo et al., 1997), code for reward prediction errors associated with action selection (Amiez, Joseph, & Procyk, 2005; Matsumoto, Matsumoto, Abe, & Tanaka, 2007), and appear to produce fERN-like extracellular potentials (Emeric et al., 2008; Halgren, Boujon, Clarke, Wang, & Chauvel, 2002; Wang, Ulbert, Schomer, Marinkovic, & Halgren, 2005). Moreover, disruption of ACC activity following the reduction of an expected reward impairs the animal's ability to switch to an alternative, more task-appropriate behavior (Shima & Tanji, 1998; Williams et al., 2004). Most significantly with respect to the present study, motor-related neurons in ACC are more active when rewards are presented following a correct response than when the rewards are obtained passively, suggesting that this neural area uses reward information for the purpose of reinforcing adaptive behaviors associated with meaningful stimulus–response contingencies (Michelet, Bioulac, Guehl, Escola, & Burbaud, 2007). By contrast, there appears to be less evidence that the monkey ACC is involved in monitoring for response conflict as predicted by the conflict monitoring theory of ACC (Emeric et al., 2008; Nakamura, Roesch, & Olson, 2005). Taken together, the human ERN and primate single-cell findings are broadly consistent with the position that ACC integrates reinforcement history over time to guide voluntary behavior (Holroyd & Coles, 2008; Kennerley, Walton, Behrens, Buckley, & Rushworth, 2006).

REFERENCES

AKKAL, D., BIOULAC, B., AUDIN, J., & BURBAUD, P. (2002). Comparison of neuronal activity in the rostral supplementary and cingulate motor areas during a task with cognitive and motor demands. *European Journal of Neuroscience*, **15**, 887-904.

AMIEZ, C., JOSEPH, J.-P., & PROCYK, E. (2005). Anterior cingulate error-related activity is modulated by predicted reward. *European Journal of Neuroscience*, **21**, 3447-3452.

BAKER, R., KRIGOLSON, O. E., & HOLROYD, C. B. (2006). Examining the feedback error-related negativity using predictive stimuli. *Psychophysiology*, **43**, S22.

BAKER, T. E., & HOLROYD, C. B. (in press). Which way do I go? Neural activation in response to feedback and spatial processing in a virtual T-maze. *Cerebral Cortex*.

BEHRENS, T. E. J., WOOLRICH, M. W., WALTON, M. E., & RUSHWORTH, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, **9**, 1214-1221.

BOTVINICK, M. M., BRAVER, T. S., BARCH, D. M., CARTER, C. S., & COHEN, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, **108**, 624-652.

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*, **99**, 523-528.

BUTTERFIELD, B., & MANGELS, J. A. (2003). Neural correlates of error detection and correction in a semantic retrieval task. *Cognitive Brain Research*, **17**, 793-817.

COHEN, M. X., & RANGANATH, C. (2007). Reinforcement learning signals predict future decisions. *Journal of Neuroscience*, **27**, 371-378.

COMPTON, R. J., CARP, J., CHADDOCK, L., FINEMAN, S. L., QUANDT, L. C., & RATLIFF, J. B. (2007). Anxiety and error monitoring: Increased error sensitivity or altered expectations? *Brain & Cognition*, **64**, 247-256.

DONKERS, F. C. L., NIEUWENHUIS, S., & VAN BOXTEL, G. J. M. (2005). Mediofrontal negativities in the absence of responding. *Cognitive Brain Research*, **25**, 777-787.

DONKERS, F. C. L., & VAN BOXTEL, G. J. M. (2005). Mediofrontal negativities to averted gains and losses in the slot-machine task: A further investigation. *Journal of Psychophysiology*, **19**, 256-262.

DUNNING, J. P., & HAJCAK, G. (2007). Error-related negativities elicited by monetary loss and cues that predict loss. *NeuroReport*, **18**, 1875-1878.

EMERIC, E. E., BROWN, J. W., LESLIE, M., POUGET, P., STUPHORN, V., & SCHALL, J. D. (2008). Performance monitoring local field potentials in the medial frontal cortex of primates: Anterior cingulate cortex. *Journal of Neurophysiology*, **99**, 759-772.

EPPINGER, B., KRAY, J., MOCK, B., & MECKLINGER, A. (2008). Better or worse than expected? Aging, learning, and the ERN. *Neuropsychologia*, **46**, 521-539.

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, and A. Kok (Eds.), *Psychophysiological Brain Research* (pp. 192-195). Tilburg, The Netherlands: Tilburg University Press.

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.

GIBSON, J., KRIGOLSON, O. E., & HOLROYD, C. B. (2006). Sensitivity of the feedback error–related negativity to reward probability. *Psychophysiology*, **43**, S41-S42.

GRANT, D. A., HAKE, H. W., & HORNSETH, J. P. (1951). Acquisition and extinction of verbal expectations in a situation analogous to conditioning. *Journal of Experimental Psychology*, **42**, 1-5.

GRATTON, G., COLES, M. G. H., & DONCHIN, E. (1983). A new method for off-line removal of ocular artifact. *Electroencephalography & Clinical Neurophysiology*, **55**, 468-484.

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. (2007). It's worse than you thought: The feedback negativity and violations of subjective expectancy. *Psychophysiology*, **44**, 905-912.

HALGREN, E., BOUJON, C., CLARKE, J., WANG, C., & CHAUVEL, P. (2002). Rapid distributed fronto-parieto-occipital processing stages during working memory in humans. *Cerebral Cortex*, **12**, 710-728.

HEWIG, J., TRIPPE, R., HECHT, H., COLES, M. G. H., HOLROYD, C. B., & MILTNER, W. H. R. (2007). Decision-making in Blackjack: An electrophysiological analysis. *Cerebral Cortex*, **17**, 865-877.

HEWIG, J., TRIPPE, R. H., HECHT, H., COLES, M. G. H., HOLROYD, C. B.,

& MILTNER, W. H. R. (2008). An electrophysiological analysis of coaching in Blackjack. *Cortex*, **44**, 1197-1205.

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. (2008). Dorsal anterior cingulate cortex integrates reinforcement history to guide voluntary behavior. *Cortex*, **44**, 548-559.

HOLROYD, C. B., & KRIGOLSON, O. E. (2007). Reward prediction error signals associated with a modified time estimation task. *Psychophysiology*, **44**, 913-917.

HOLROYD, C. B., NIEUWENHUIS, S., MARS, R., & COLES, M. G. H. (2004). Anterior cingulate cortex, selection for action, and error processing. In M. Posner (Ed.), *Cognitive neuroscience of attention* (pp. 219-231). New York: Guilford Publishing, Inc.

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.

HOLROYD, C. B., NIEUWENHUIS, S., YEUNG, N., NYSTROM, L., MARS, R., COLES, M. G. H., & COHEN, J. D. (2004). Dorsal anterior cingulate cortex shows fMRI response to internal and external error signals. *Nature Neuroscience*, **7**, 497-498.

HOLROYD, C. B., PAKZAD-VAEZI, K. L., & KRIGOLSON, O. E. (2008). The feedback correct–related positivity: Sensitivity of the event-related brain potential to unexpected positive feedback. *Psychophysiology*, **45**, 688-697.

HOLROYD, C. B., YEUNG, N., COLES, M. G. H., & COHEN, J. D. (2005). A mechanism for error detection in speeded response time tasks. *Journal of Experimental Psychology: General*, **134**, 163-191.

ISOMURA, Y., ITO, Y., AKAZAWA, T., NAMBU, A., & TAKADA, M. (2003). Neural coding of "attention for action" and "response selection" in primate anterior cingulate cortex. *Journal of Neuroscience*, **23**, 8002-8012.

ISOMURA, Y., & TAKADA, M. (2004). Neural mechanisms of versatile functions in primary anterior cingulate cortex. *Reviews in the Neurosciences*, **15**, 279-291.

ITO, S., STUPHORN, V., BROWN, J. W., & SCHALL, J. D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science*, **302**, 120-122.

KENNERLEY, S. W., WALTON, M. E., BEHRENS, T. E. J., BUCKLEY, M. J., & RUSHWORTH, M. F. S. (2006). Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience*, **9**, 940-947.

KLEIN, T. A., NEUMANN, J., REUTER, M., HENNIG, J., VON CRAMON, D. Y., & ULLSPERGER, M. (2007). Genetically determined differences in learning from errors. *Science*, **318**, 1642-1645.

KRIGOLSON, O. E., & HOLROYD, C. B. (2007). Predictive information and error processing: The role of medial-frontal cortex during motor control. *Psychophysiology*, **44**, 586-595.

KRIGOLSON, O. E., MATHEWSON, K., BAKER, T. E., BAKER, R., & HOLROYD, C. B. (2006). The role of medial-frontal cortex in sequence learning. *Psychophysiology*, **43**, S54-S55.

KRIGOLSON, O. E., PIERCE, L. J., HOLROYD, C. B., & TANAKA, J. W. (in press). Learning to become an expert: Reinforcement learning and the acquisition of perceptual expertise. *Journal of Cognitive Neuroscience*.

LARSON, M. J., KELLY, K. G., STIGGE-KAUFMAN, D. A., SCHMALFUSS, I. M., & PERLSTEIN, W. M. (2007). Reward context sensitivity impairment following severe TBI: An event-related brain potential experiment. *Journal of the International Neuropsychological Society*, **13**, 615-625.

LEE, K., KRIGOLSON, O. E., & HOLROYD, C. B. (2006). Modulation of fERN amplitude by reward expectancy. *Psychophysiology*, **43**, S57-S58.

MARS, R. B., COLES, M. G. H., 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.

MATSUMOTO, M., MATSUMOTO, K., ABE, H., & TANAKA, K. (2007). Medial prefrontal cell activity signaling prediction errors of action values. *Nature Neuroscience*, **10**, 647-656.

MATSUMOTO, K., SUZUKI, W., & TANAKA, K. (2003). Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science*, **301**, 229-232.

MATSUMOTO, K., & TANAKA, K. (2004). The role of medial prefrontal

cortex in achieving goals. *Current Opinion in Neurobiology*, **14**, 178-185.

MICHELET, T., BIOULAC, B., GUEHL, D., ESCOLA, L., & BURBAUD, P. (2007). Impact of commitment on performance evaluation in the rostral cingulate motor area. *Journal of Neuroscience*, **27**, 7482-7489.

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.

MONTAGUE, P. R., HYMAN, S. E., & COHEN, J. D. (2004). Computational roles for dopamine in behavioral control. *Nature*, **431**, 760-767.

MORRIS, S. E., HEEREY, E. A., GOLD, J. M., & HOLROYD, C. B. (2008). Learning-related changes in brain activity following errors and performance feedback in schizophrenia. *Schizophrenia Research*, **99**, 274-285.

NAKAMURA, K., ROESCH, M. R., & OLSON, C. R. (2005). Neuronal activity in macaque SEF and ACC during performance of tasks involving conflict. *Journal of Neurophysiology*, **93**, 884-908.

NIEUWENHUIS, S., HOLROYD, C. B., MOL, N., & COLES, M. G. H. (2004). Reinforcement-related brain potentials from medial frontal cortex: Origins and functional significance. *Neuroscience & Biobehavioral Reviews*, **28**, 441-448.

NIEUWENHUIS, S., NIELEN, M. M., MOL, N., HAJCAK, G., & VELTMAN, D. J. (2005). Performance monitoring in obsessive-compulsive disorder. *Psychiatry Research*, **134**, 111-122.

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, & Behavioral Neuroscience*, **2**, 19-36.

NIEUWENHUIS, S., SLAGTER, H. A., VON GEUSAU, N. J. A., HESLENFELD, D., & HOLROYD, C. B. (2005). Knowing good from bad: Differential activation of human cortical areas by positive and negative outcomes. *European Journal of Neuroscience*, **21**, 3161-3168.

NIKI, H., & WATANABE, M. (1979). Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Research*, **171**, 213-224.

NISHIJO, H., YAMAMOTO, Y., ONO, T., UWANO, T., YAMASHITA, J., & YAMASHIMA, T. (1997). Single neuron responses in the monkey anterior cingulate cortex during visual discrimination. *Neuroscience Letters*, **227**, 79-82.

PAULUS, M. P., HOZACK, N., FRANK, L., & BROWN, G. G. (2002). Error rate and outcome predictability affect neural activation in prefrontal cortex and anterior cingulate during decision-making. *NeuroImage*, **15**, 836-846.

POTTS, G. F., MARTIN, L. E., BURTON, P., & MONTAGUE, P. R. (2006). When things are better or worse than expected: The medial frontal cortex and the allocation of processing resources. *Journal of Cognitive Neuroscience*, **18**, 1112-1119.

PROCYK, E., TANAKA, Y. L., & JOSEPH, J. P. (2000). Anterior cingulate activity during routine and nonroutine sequential behaviors in macaques. *Nature Neuroscience*, **3**, 502-508.

QUILODRAN, R., ROTHE, M., & PROCYK, E. (2008). Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron*, **57**, 314-325.

RIDDERINKHOF, K. R., ULLSPERGER, M., CRONE, E. A., & NIEUWENHUIS, S. (2004). The role of medial frontal cortex in cognitive control. *Science*, **306**, 443-447.

ROESCH, M. R., CALU, D. J., & SCHOENBAUM, G. (2007). Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nature Neuroscience*, **10**, 1615-1624.

RUSHWORTH, M. F. S., WALTON, M. E., KENNERLEY, S. W, & BANNERMAN, D. M. (2004). Action sets and decisions in the medial frontal cortex. *Trends in Cognitive Sciences*, **8**, 410-417.

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*, **31**, 1709-1711.

SHIMA, K., & TANJI, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science*, **282**, 1335-1338.

ULLSPERGER, M., & VON CRAMON, Y. (2003). Error monitoring using external feedback: Specific roles of the habenular complex, the reward

system, and the cingulate motor area revealed by functional magnetic resonance imaging. *Journal of Neuroscience*, **23**, 4308-4314.

VAN VEEN, V., HOLROYD, C. B., COHEN, J. D., STENGER, A., & CARTER, C. (2004). Errors without conflict: Implications of performance monitoring theories for anterior cingulate cortex. *Brain & Cognition*, **56**, 267-276.

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. *Journal of Neuroscience*, **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. *Nature Neuroscience*, **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., HOLROYD, C. B., & COHEN, J. D. (2005). ERP correlates of feedback and reward processing in the presence and absence of response choice. *Cerebral Cortex*, **15**, 535-544.

NOTES

1. It is not surprising that the predictions of this or any theory do not obtain in all situations. As is well known, even Newton's laws of motion do not hold at very small scales, very high speeds, or in very strong gravitational fields.

2. These results have been presented previously in poster format (Baker, Krigolson, & Holroyd, 2006; Gibson, Krigolson, & Holroyd, 2006; Lee, Krigolson, & Holroyd, 2006). Since then, we have modified the window size for determining fERN amplitude, which has altered their statistical significance somewhat. However, the overall pattern of results across the three experiments remains the same.

3. Critically, in doing so, we sought to correct for an inelegance associated with the original Nieuwenhuis et al. (2002) experiment. In the 80% condition of that study, participants received feedback that was either consistent (i.e., valid) or inconsistent (i.e., invalid) with the selected response. Thus, participants received valid correct feedback and invalid error feedback following the optimal response, and invalid correct feedback and valid error feedback following the suboptimal response. To examine the fERN, expected difference waves were created by subtracting the ERPs associated with valid correct feedback following optimal responses from the ERPs associated with valid error feedback following suboptimal responses, and unexpected difference waves were created by subtracting the ERPs associated with invalid correct feedback following suboptimal responses from the ERPs associated with invalid error feedback following optimal responses. Although it was found that the unexpected difference wave was larger than the expected difference wave, this analysis was complicated by the fact that the number of trials associated with each ERP differed markedly across the four conditions, since participants were exposed to relatively many valid correct trials and relatively few invalid correct trials. For example, participants who selected the optimal response on 80% of the trials (i.e., who probability matched; Grant, Hake, & Hornseth, 1951) would have received valid correct feedback on 64% of the trials, invalid error feedback on 16% of the trials, invalid correct feedback on 4% of the trials, and valid error feedback on 16% of the trials. Furthermore, participants learned to favor the optimal response over the suboptimal response as each block progressed. Thus, the prediction errors associated with the suboptimal responses would have tended to occur early in learning before the reward predictions had fully developed, whereas the prediction errors associated with the optimal responses would have tended to occur late in learning after the reward predictions had more fully developed. A consequence of this design was that the conditions were not truly counterbalanced. Because P300 amplitude is larger for unexpected than for expected outcomes, it is possible that the difference wave approach would not have successfully removed P300 overlap from the fERN measure (Holroyd & Krigolson, 2007). For this reason, in Experiment 3, we modified the task so as to eliminate this confound (see the Method section for Experiment 3).

4. At the same time, it should be remembered that other task changes could have also contributed to this result. For example, whereas Experiment 3 required participants to select between two choices, Experiments 1 and 2 required them to select between 20 and 5 choices, respectively; perhaps participants have increased confidence in their predictions when their options are limited. Likewise, in contrast with the first two experiments, Experiment 3 did not use the door stimuli, which may have increased attentional resources devoted to the visual cue. It is also possible that in Experiment 2, participants would have developed stronger expectations about the stimulus–reward associations had they been exposed to more trials.