ELSEVIER

BRAIN RESEARCH

Research Report

The good, the bad and the neutral: Electrophysiological responses to feedback stimuli

Clay B. Holroyd[a,*], Greg Hajcak[b], Jeff T. Larsen[c]

[a]Department of Psychology University of Victoria, PO Box 3050 STN CSC Victoria, Canada BC V8W 3P5
[b]Department of Psychology, University of Delaware, Newark, DE 19716, USA
[c]Department of Psychology, Texas Tech University, Lubbock, TX 79430, USA

ARTICLE INFO

ABSTRACT

The feedback error-related negativity (fERN) is a component of the event-related brain potential elicited in gambling and trial-and-error learning tasks by negative, but not positive, feedback stimuli. Here, we present the results of a series of five experiments that investigated the response of the fERN to the presentation of neutral feedback stimuli. In three of the experiments, the neutral feedback stimuli indicated that the participants did not receive a potential reward nor incur a potential penalty (i.e., they received nothing); and in the remaining two experiments, the neutral feedback stimuli did not convey any meaningful information (i.e., the participants were either successful or unsuccessful on those trials, but the feedback stimuli were uninformative about the outcomes). Across the five experiments, we found that neutral feedback stimuli elicited a fERN about as large as that elicited by negative feedback stimuli. This result is consistent with recent proposals that the evaluative system that produces the fERN classifies outcomes into two categories: those outcomes that indicate that a goal has been satisfied and those that do not.

* Corresponding author. Fax: +1 250 721 7525.
  E-mail address: holroyd@uvic.ca (C.B. Holroyd).
  URL: http://www.web.uvic.ca/psyc/holroyd.html (C.B. Holroyd).

1.   Introduction

In recent years, a surge of interest has focused on the neural mechanisms that underlie how humans learn to pursue reward and avoid punishment (Cohen and Blum, 2002). A window onto this system has been provided by the discovery of a component of the event-related brain potential (ERP) called the "feedback error-related negativity" (fERN or "feedback negativity"). The fERN is a negative deflection in the ERP that is distributed over frontal areas of the scalp and that reaches maximum amplitude about 250 ms following the onset of negative feedback stimuli. Negative feedback stimuli are stimulus events that indicate that an undesirable outcome has occurred, such as an error or a loss of money. Importantly, presentation of positive feedback stimuli (that indicate that a "good" event has occurred) does not elicit the fERN, indicating that the system that produces this component is differentially sensitive to positive and negative feedback (Miltner et al., 1997). According to the "reinforcement learning theory of the error-related negativity" (RL-ERN theory), the evaluative information indexed by the fERN is conveyed via the midbrain dopamine system from a monitoring system located in the basal ganglia to a decision making system located in anterior cingulate cortex, where that information can be applied for the adaptive modification of behavior. Further, the theory proposes that the amplitude of the fERN depends on both the valence and the expectedness of the outcome, such that the largest fERNs are elicited by relatively unexpected negative events (Holroyd and Coles, 2002) (see also Holroyd, 2004). Thus, for example, infrequent penalties elicit larger fERNs than are elicited by infrequent rewards, frequent rewards, and frequent penalties (Holroyd et al., 2003) (but see Hajcak et al., 2005). This position has received some preliminary support in a series of recent studies (reviewed in Holroyd et al., 2004b; Nieuwenhuis et al., 2004).

In one of these studies, participants engaged in a pseudo trial-and-error learning task in which the amount of reward possible on each trial was varied by condition (Holroyd et al., 2004a). On each trial, the participants selected one of three balloons that appeared on a computer screen and were presented with a feedback stimulus that indicated the receipt of a small financial reward or penalty. In each condition, the feedback stimuli were comprised of three types. In a "win" condition, the feedback indicated that the participants either received no reward, received a small reward, or received a large reward, whereas in a "lose" condition, the feedback indicated that they either lost nothing, incurred a small penalty, or incurred a large penalty. Further, in an "even" condition, the feedback indicated that the participants either incurred a penalty, received nothing, or received a reward. Unbeknownst to the participants, the feedback stimuli in each condition were random and equiprobable. The results of the experiment indicated that the system that produced the fERN was sensitive to the task context in which each outcome occurred. For example, feedback stimuli indicating that the participant had neither gained nor lost anything elicited a fERN in the win condition but not in the lose condition. In general, the "worst" possible outcome in each condition elicited the fERN, and the "best" possible outcome did not.

These results were interpreted in terms of the RL-ERN theory, which predicts for this task that the system will come to expect an intermediate outcome in each condition such that the best and worst outcomes reflect deviations from this expectation.

However, an apparent inconsistency with this theoretical prediction was that in this experiment, the amplitudes of the fERNs elicited by the intermediate and worst outcomes in each condition were not significantly different from each other. At face value, the RL-ERN theory would appear to predict that the amplitude of the fERN associated with the intermediate outcome should be midway between the amplitudes associated with the worst and best outcomes. This finding was considered in the light of two possibilities: (1) the study lacked the statistical power necessary to resolve any difference in amplitude between the fERNs elicited by the intermediate and worst outcomes; and 2) the evaluative system that produces the fERN may be non-linear, such that the system weights the worst and middle outcomes about equally. Importantly, two recent studies have provided some support for the latter possibility. Both studies found that the fERN was insensitive to gradations of reward magnitude and appeared to reflect a binary categorization of the outcomes as either good or not good (Hajcak et al., in press; Yeung and Sanfey, 2004). Although two other recent studies also included neutral feedback stimuli, in one of these experiments, the neutral feedback stimuli occurred relatively infrequently (Muller et al., 2005) and so may have confounded fERN amplitude with other ERP components (Holroyd, 2004); and in the other experiment, the fERN elicited by the neutral feedback stimuli was not examined (Nieuwenhuis et al., 2005b).

In the present study, we explored the effect of neutral feedback stimuli on the amplitude of the fERN in a series of five experiments (each associated with a different group of participants). Our purpose was two-fold. First, we wished to determine whether the observed lack of a difference between the fERNs elicited by the neutral and worst outcomes in Holroyd et al. (2004a) would be replicated. Second, we wished to find out if this phenomenon is robust to a variety of experimental manipulations, including certain particulars of the task context. To start, Experiment 1 replicated the "even" condition described in Holroyd et al. (2004a). As with the previous experiment, the feedback stimuli on each trial indicated that participants received either a loss (–10¢), a gain (+10¢), or nothing (0¢). In Experiment 2, we wished to see if the system that produces the fERN evaluates uninformative feedback as it does intermediate outcomes in general. Thus, Experiment 2 was identical to Experiment 1, except that the participants in Experiment 2 were told that the neutral feedback stimulus was uninformative: they were told that although they either received or lost money on such trials, the feedback did not indicate which outcome actually occurred. In Experiment 3, because the evaluative system's assessments can be "anchored" relative to a reference point (Tversky and Kahneman, 1974), we examined whether the amount of the starting bonus affected the amplitude of the fERN elicited by the neutral feedback stimuli. At the start of both Experiments 1 and 2, participants were provided with $0 in bonus money; they were told that although they could accumulate

(and possibly lose) a bonus, they could not in fact finish the task owing the experimenter money. For this reason, in Experiment 3 we replicated Experiment 1, with the exception that participants began the task with a $5 bonus and were told that they would finish the experiment with between $0 and $10. In Experiment 4, we asked whether the large fERN elicited by neutral feedback stimuli was specific to gambling tasks. We conducted a time estimation task in which, on each trial, participants were asked to press a button when they believed that a second had elapsed following a warning cue (cf., Miltner et al., 1997). The feedback stimuli in this task indicated that participants' responses were either on-time, not on-time, or (on one third of the trials selected at random) were uninformative. Finally, in Experiment 5, we asked whether the fERN elicited by the neutral feedback stimuli could be affected by a predictive cue presented at the start of the trial that indicated the potential size of the forthcoming outcome.

These experiments were conducted at two separate institutions using similar, but not identical, methods. For ease of exposition, we present Experiments 1–4 (Princeton University) together first, followed by Experiment 5 (University of Delaware).

2. Experiments 1–4

2.1. Materials and methods

2.1.1. Participants and tasks

2.1.1.1. Experiment 1. In this pseudo trial-and-error learning task, participants used feedback stimuli to select between three "balloons" presented on a computer screen (cf., Holroyd et al., 2004a). 10 participants (7 male, 3 female), age 20 ± 1 years, participated for pay ($20) or course credit and were also given a $5 bonus at the end of the session (see below). The participants sat comfortably in an electromagnetically shielded room about 1 m from a computer screen. On each trial, the imperative stimulus consisted of three circles in a row (i.e., "O O O" 0.6° high, 5° wide, blue color against a black background), and were asked to imagine that each circle was a "balloon". Before the task began, participants were told that one of the three balloons contained 10¢, another balloon contained a "gremlin" that would steal 10¢ from them, and a third balloon was empty. The imperative stimulus remained on the screen until the participants selected a balloon by pressing one of three buttons on a response pad. At the time of the response, the imperative stimulus was replaced with a second stimulus (0.6° high, 5° wide, blue color, 1.0-s duration) in which the selected balloon was replaced by an asterisk (e.g., "O O * ", if the participant selected the rightmost balloon). This was followed by a feedback stimulus that appeared directly above the location associated with the middle balloon (0.6°, red color, 1.0-s duration): either a '+' (positive feedback, +10¢), '–' (negative feedback, –10¢), or 'o' (neutral feedback, +0¢) stimulus. A black screen was then presented for 0.5 s before the onset of the following imperative stimulus. Participants were told that they should respond in a way

that maximized the total amount of money earned, that at the end of the experiment, they would be rewarded the money they found, and that they would not end the experiment owing money. Unbeknownst to them, the feedback stimuli were selected at random (with replacement) and were equiprobable (1/3 probability for each of the three feedback types). The experiment consisted of 3 blocks of 100 trials. At the end of the experiment, participants were informed about the feedback probabilities and told that they did not make any money. However, they were given a $5 consolation bonus.

2.1.1.2. Experiment 2. 10 participants (5 male, 5 female), age 21 ± 1 years, participated for pay ($20) or for course credit. Participants also received a $5.00 bonus at the end of the session. The experiment was identical to Experiment 1, except that before the task began, the participants were told that the neutral feedback stimulus ('o') was uninformative, i.e., although they either gained or lost money on that trial, the stimulus did not indicate the outcome. Participants engaged in 3 blocks of 99 trials; the three feedback types were selected at random (without replacement) and were equiprobable.

2.1.1.3. Experiment 3. 10 participants (4 male, 6 female), age 20 ± 1 years, participated for pay ($20) or for course credit. This experiment was identical to Experiment 1, except that participants began it with $5.00 in bonus money, and were told that they would leave the experiment with a bonus between $0 and $10.

2.1.1.4. Experiment 4. This was a "time estimation task" in which, on each trial, participants estimated when they believed a 1-s interval had elapsed following presentation of a warning tone (cf., Miltner et al., 1997). 10 participants (4 male, 6 female), age 21 ± 2 years, participated for pay ($20) or for course credit. On each trial, participants were presented with a tone (800 Hz, 50-ms duration, ~65 db) and were required to press a button on a response pad when they believed that 1.0 s had elapsed following the onset of the tone. 0.6 s following the response, participants were presented with a feedback stimulus (0.6°, red color on a black background, 1.0-s duration) that indicated that their response was either "on time" (positive feedback, '+') or "not on time" (negative feedback, '–'). Responses were considered on time if the RT occurred within a window from $1.0 - \alpha$ s to $1.0 + \alpha$ s following tone onset, where α equaled 0.1 s at the start of each block of trials and, at the end of each trial, was incremented by 10 ms if the proportion of errors in the block was greater than 50% and was decremented by 10 ms if that proportion was less than 50%. Participants engaged in 1 block of 225 trials. On 1/3 of the trials, selected at random (with replacement), a neutral feedback stimulus ('o') was substituted in lieu of the actual feedback. At the start of the task, participants were told that the neutral feedback stimulus was uninformative and thus did not indicate whether their response was on time or not. (In a "2-feedback" condition, participants also engaged in 1 block of 150 trials in which only positive and negative feedback stimuli were presented, without any uninformative feedback. This condition, the order of which was

counterbalanced across participants with the three-feed-back condition, replicated previous observations of the fERN (Miltner et al., 1997) and will not be discussed here further).

2.1.2. Data acquisition
An 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, and Pz (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. 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.

2.1.3. Data analysis
For each feedback stimulus, a 1-s epoch of data (0.2 s baseline) was extracted from the continuous data file for analysis. Ocular artifact was corrected with an eye move-ment correction algorithm (Gratton et al., 1983). The EEG data were re-referenced offline to linked-mastoid electrodes by subtracting from each sample of data recorded at each channel one-half the activity recorded at the right mastoid. The data were baseline corrected by subtracting from each sample the average activity of that channel during the baseline period. The 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. FERN amplitude was measured base to peak at channel FCz, where it reached a maximum (cf., Holroyd et al., 2004a) using an algorithm described in Holroyd et al. (2003). First, the algorithm identified the sample associated with the most positive value of the ERP within a 160–240 ms window following the presentation of the feedback stimulus. The latency of this sample was taken as the time of onset of the negativity. Then the algorithm identified the sample associated with the most negative value of the ERP within a window extending from the onset of the negativity to 325 ms following the presentation of the feedback stimulus. If the latency of this sample was 325 ms (i.e., at the edge of the window), then the ERP component was considered to be a positivity, and the amplitude of the negativity was taken as 0 mV. Otherwise, the latency of the sample was taken as the time of maximum component amplitude. The ampli-tude of the negativity was then defined by the difference in the ERP values associated with the component maximum and the component onset.

The data were submitted to an analysis of variance (ANOVA) with repeated measures and to paired t tests. The Greenhouse–Geisser correction for repeated mea-sures was applied where appropriate (Keselman and Rogan, 1980).

2.2. Results

Fig. 1 shows ERPs elicited by the positive, neutral, and negative feedback stimuli recorded at channel FCz for Experiment 1–4. A two-way ANOVA on experimental group and feedback type, with repeated measures on feedback type, revealed a main effect of feedback type on fERN amplitude, $F(2,72) = 23.6$, $P < 0.0001$, $\in = 0.83$, but no main effect of experimental group, $P = 0.42$, and no interaction between experimental group and feedback type, $P = 0.14$. Confidence intervals associated with Bonferroni multiple comparisons indicated that negative outcomes elicited larger fERNs than positive outcomes [–3.8 μV, –1.4 μV], and that neutral outcomes elicited larger fERNs than positive out-comes [–4.2 μV, –1.8 μV]. In contrast, fERN amplitudes associated with negative and neutral outcomes did not differ [–0.8 μV, 1.6 μV], Cohen's $d = 0.16$. Note that these confidence intervals, taken together with the absence of a significant interaction between experimental group and feedback type, indicate that fERN amplitudes associated with negative and neutral outcomes were not significantly different for each experimental group.

3. Experiment 5

3.1. Materials and methods

3.1.1. Participants and task
Twenty-three (5 male, 18 female) University of Delaware psychology students participated in the current experiment for course credit. In addition, participants were informed that they could receive between $0.00 and $10.00 based on their performance in the task. All participants were told that they began the task with $5.00, and that at the conclusion of the experiment, they would keep the money they found. Partici-pants were also told that they would not end the experiment owing the experimenters money.

The task was administered on a Pentium I class computer using Presentation software (Neurobehavioral Systems, Inc.) to control stimulus presentation and timing. All participants were told that, on each trial, they could win, break even, or lose money. At the start of each trial a cue ('5' or '25', 1.0 s, white color against a black background) presented in the center of the screen indicated to the participant whether the upcoming trial was worth 5¢ or 25¢. A graphic representing three doors in a horizontal line appeared immediately following cue offset and remained on the screen until the participant responded (cf., Hajcak et al., 2005). Before the start of the task, the participants were told that each trial's specified reward (+5¢ or +25¢), penalty (–5¢ or –25¢), and nothing (±0¢) were hidden behind the three doors. Participants were instructed to choose a door by pressing the left 'ctrl', spacebar, and right 'ctrl' keys (corresponding to the left, center, and right door, respectively). Following their decision, a black screen appeared for 0.5 s, followed by the feedback stimulus (1.0 s, green color against a black background). The

reward, penalty, and neutral feedback were indicated by '+', '–', and 'o' stimuli, respectively. The interval between offset of the feedback stimulus and the onset of the following cue was 1.0 s. All cue and feedback stimuli occupied approximately 2° of visual angle horizontally and 2° vertically. A fixation mark (approximately 2° by 2° of visual angle, white color on a black background) appeared immediately following the participant's decision and remained on the screen until the onset of the feedback stimulus.

After a practice block of 40 trials in which no money could be won or lost, participants performed 240 trials of the experiment. Unknown to the participants, the feedback was pseudo-random and predetermined such that each feedback stimulus was presented with equal frequency; additionally, exactly 50% of all trials were preceded by the '5' cue. Thus, there were a total of 40 trials of each type ('5' followed by '+', '5' followed by 'o', '5' followed by '–', '25' followed by '+', and so on). Each participant ended the experiment with $5.00.

3.1.2. Data analysis

The EEG was recorded using a Neurosoft Quik-Cap (El Paso, TX). Recordings were taken from 3 locations along the midline: frontal (Fz), central (Cz), and parietal (Pz). In addition, Med-Associates tin electrodes (St. Albans, VT) were placed on the left and right mastoids (M1 and M2, respectively). During the recording, all activity was referenced to Cz. The electrooculogram (EOG) was recorded using Med-Associates miniature electrodes placed approximately 1 cm above and below the participant's right eye. The right earlobe served as a ground site. All EEG/EOG electrode impedances were below 10 K, and the data from all channels were recorded by a Grass Model 7D polygraph with Grass Model 7P1F preamplifiers (bandpass 0.05–35 Hz).

All bioelectric signals were digitized on a laboratory microcomputer using VPM software (Cook, 1992). The EEG was sampled at 200 Hz. Data collection began with the participants' response (0.5 s prior to feedback) and continued for 1.5 s. Off-line, the EEG for each trial was corrected for vertical EOG artifacts using the method developed by Gratton et al. (1983) and then re-referenced to the average activity of the mastoid electrodes. Trials were rejected and not counted in subsequent analysis if there was excessive physiological artifact (i.e., 25 ms of invariant analog data on any channel or A/D values on any channel that equaled that converter's minimum or maximum values). Single trial EEG data were lowpass filtered at 20 Hz with an FIR digital filter as per Cook and Miller (1992).

Finally, stimulus-locked ERPs were averaged based on trial value (5¢ and 25¢) and feedback type (negative, neutral, positive). The fERN was quantified at Fz (where it reached maximum amplitude; note that channel FCz was not used in this experiment) using the algorithm described in Experiments 1–4 and Holroyd et al. (2003). Following the convention of this laboratory (Hajcak et al., in press), the window for determining fERN onset extended from 150 ms to 350 ms after feedback stimulus presentation, and the window for determining the peak of the fERN extended from the time of fERN onset to 350 ms after feedback stimulus presentation.
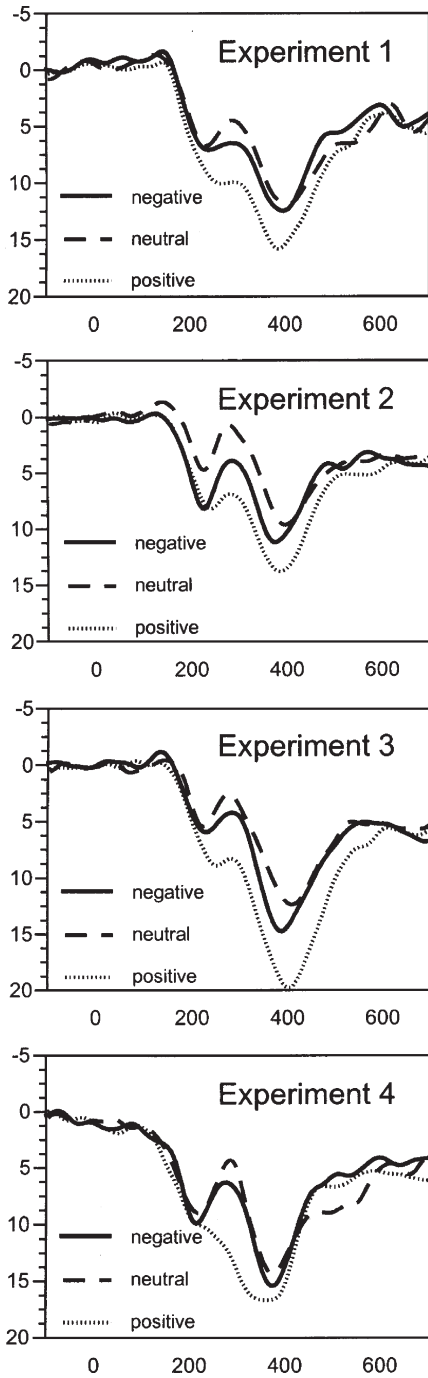
3.2. Results

Fig. 2 presents the ERPs (recorded at channel Fz) elicited by the negative, neutral, and positive feedback stimuli on 5¢ (top) and 25¢ (bottom) trials. As suggested by inspection of the figure, a 3 (feedback type) × 2 (trial value) repeated measures ANOVA on fERN amplitude revealed a main effect feedback type, $F(2,46) = 7.08$, $P < 0.01$, $\in = 0.74$. In contrast, the main effect of trial value, $F(1,23) = 1.14$, $P > 0.25$, and the interaction of trial value and feedback type, $F(2,46) = 1.75$, $P > 0.15$, $\in = 0.96$, did not reach significance. Post hoc comparisons between negative, neutral, and positive conditions indicated the fERN was larger following penalties than following rewards, $t(23) = 2.85$, $P < 0.01$, and larger following neutral outcomes than following rewards, $t(23) = 2.97$, $P < 0.01$, but was not significantly different in amplitude between penalties and neutral outcomes, $t(23) = 1.23$, $P > 0.20$, Cohen's $d = -0.15$. These data indicate that the fERNs elicited by penalties and neutral outcomes were both significantly larger than that elicited by rewards but were not significantly different from each other. Additionally, the trial value did not appear to affect the amplitude of the fERN.

4. Discussion

Across five experiments, negative and neutral feedback stimuli elicited fERNs that were larger than that elicited by positive feedback stimuli but that were comparable in amplitude to one another. This result was sustained despite significant variations in task design, including the participants' starting bonuses and the type of information conveyed by the neutral feedback stimuli. Given that the findings were consistent across these five experiments as well as with the previous experiments reported by Holroyd et al. (2004a), which collectively have involved 95 participants, and given that typical fERN experiments involve samples sizes of 10–15 participants, it seems unlikely that the failure to find a relatively small fERN elicited by neutral feedback stimuli could be due to the sample sizes employed. Indeed, inspection of Figs. 1 and 2 seems to suggest that the neutral feedback stimuli elicited *larger* fERNs than did the negative feedback stimuli, but this finding was not significant—even when tested using an omnibus ANOVA across four experiments. Thus, it appears that the evaluative system that produces the fERN treats neutral and negative feedback stimuli in much the same way. This finding is consistent with the results of recent studies that have indicated that the evaluative system that produces the fERN categorizes outcomes in a binary manner (Hajcak et al., in press; Yeung and Sanfey, 2004). Evidently, the system classifies outcomes into two distinct categories: those outcomes that indicate that the task goal has been satisfied, and everything else. According to this view, the evaluative system that produces the fERN lumps the neutral and negative feedback stimuli together into a single category containing any outcomes that leave open the possibility that the goal has not been obtained.

Can the results of this study be reconciled with the RL-ERN theory? A face-value interpretation of the theory suggests that
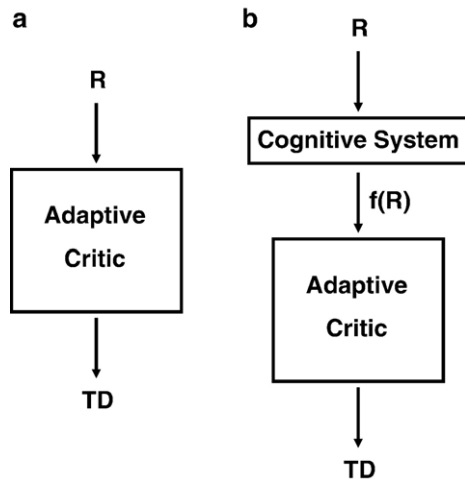
Fig. 1 – Event-related potentials elicited by negative, neutral, and positive feedback stimuli in Experiments 1–4. Data recorded from channel FCz. Abscissa: time (ms); 0 ms denotes time of stimulus onset. Ordinate: voltage (μV). Note that negative is plotted up by convention.

to external events by the evaluative system that produces the fERN is linearly related to the objective value of those events. Fig. 3a illustrates the adaptive critic model that implements this assumption, as described in the artificial intelligence literature (Sutton and Barto, 1998). Here, the reward variable R ranges continuously from −1 (punishment) to +1 (reward) and is converted by the adaptive critic into the temporal difference error TD. This theoretical framework was reflected in the original computational models of phasic midbrain dopamine activity in the monkey (Barto, 1995; Montague et al., 1996) and has been the mainstay of computational models of this system since (see Suri, 2002 for review). As originally proposed, the theory holds that primary reward signals are carried via limbic inputs into the basal ganglia, which implement the adaptive critic (Barto, 1995). The basal ganglia compute errors in reward prediction and broadcast these signals to other neural areas as phasic changes in the activity of the midbrain dopamine system. Note that according to this position, the basal ganglia constitute a rather dumb computational workhorse: They do not determine which outcomes are "good" or "bad" in and of themselves (which is the responsibility of the limbic inputs to the basal ganglia) but rather learn through a conditioning process which ongoing events predict those outcomes. Unstated in these models is the assumption that the limbic reward input is more or less linearly related to the "objective" value of the primary rewards. Thus, for example, two drops of juice would be worth about twice as much as one drop of juice to a thirsty monkey.
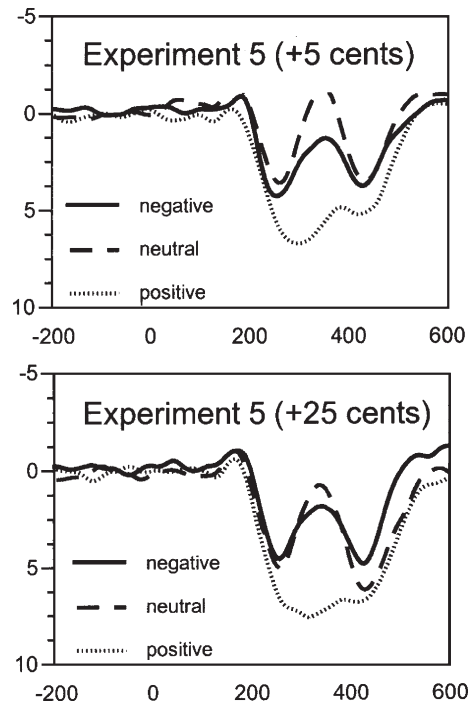
Fig. 2 – Event-related potentials elicited by negative, neutral, and positive feedback stimuli in Experiment 5 on 5¢ trials (top) and 25¢ trials (bottom). Data recorded from channel Fz. Abscissa: time (ms); 0 ms denotes time of stimulus onset. Ordinate: Voltage (μV). Note that negative is plotted up by convention.

the evaluative system should come to predict the intermediate value in each of these experiments, such that the best and worst outcomes reflect positive and negative deviations, respectively, from this expectation. If such were the case, then the neutral feedback stimuli should elicit fERNs that are intermediate in amplitude between that elicited by negative and positive feedback stimuli (Holroyd et al., 2004a). This prediction is based on the assumption that the value assigned

Fig. 3 – Adaptive critic models. (a) Standard adaptive critic model that computes the temporal difference error TD from the primary reward value R. (b) Modified adaptive critic model that computes the temporal difference error TD from the output ƒ(R) of a cognitive preprocessing system.

However, human participants in cognitive psychology experiments are rarely required to work for primary rewards. Instead, they are typically exposed to abstract feedback stimuli that function as secondary reinforcers for other things, often money, which itself is a secondary reinforcer. In other tasks, the stimuli constitute abstract performance feedback that indicate success and failure without any implied financial gain or loss. In all these cases, the human cognitive system must attribute the significance to these abstract external events: The system determines which events are "good" and which are "bad". The results of its determination cannot be a linear function of "objective" reward value because reward preference is intrinsically subjective. For example, consider the outcomes of obtaining correct feedback on the one hand (which presumably would please an experimenter) vs. obtaining error feedback on the other hand (which, if associated with speeded responses, may enable a participant to finish the experimental session sooner). Which of the two outcomes can be said to be objectively "better"? Clearly, people adopt short-term goals that they use to navigate the particular environmental contexts in which they find themselves, but which may be only loosely related to the apparent "objective" values associated with the delivery of specific rewards and punishments. Our previous results have suggested that the manner in which particular task outcomes are categorized depends sensitively on the experimental context, such that even small variations in the range of possible outcomes can affect how the outcomes are classified (Hajcak et al., in press).

If the adaptive critic does not receive reward signals that are linearly related to an objective reward value, then the scenario presented in Fig. 3a is an oversimplification. Instead, the apparent situation is illustrated in Fig. 3b, which shows that the external input R is first evaluated by a cognitive system. This cognitive system establishes the goals for the present task context – for example, to earn as much bonus money as possible – and produces as output some function ƒ(R) that indicates whether or not the goal has been satisfied. In turn, this information is communicated to the adaptive critic, which computes the TD error on this binary quantity. With respect to the present series of experiments, if negative and neutral feedback stimuli are grouped into a single category by the cognitive system, then the RL-ERN theory would predict that negative and neutral feedback stimuli should elicit equally large fERNs, both of which are larger than that elicited by positive feedback stimuli. To be concrete: if the cognitive system assigns a value of "0" to both negative and neutral feedback stimuli, and a value of "1" to positive feedback stimuli, and if each of these feedback types occur with equal probability, then the adaptive critic will come to predict positive feedback with a probability of 1/3. In this case, both negative and neutral feedback stimuli would be equally unexpected (eliciting a TD error of –1/3), and positive feedback stimuli would be better than expected (eliciting a TD error of +2/3).

The obvious question is, how does the cognitive system determine the goals? This question is beyond the scope of these adaptive critic models, which simulate the function of a simple "bottom up" predictive mechanism (Egelman et al., 1998). In fact, it is safe to say that this is the outstanding issue in cognitive control research, as it relates to how individuals come to adopt particular task sets (Miller and Cohen, 2001). This question notwithstanding, ample evidence indicates that the basal ganglia operate in accordance with cognitive input that they receive from cortex. In particular, we have suggested that task-related goals are communicated to basal ganglia via the recurrent "loops" that link prefrontal cortex and the striatum via the thalamus (Holroyd et al., 2005). Consistent with this position, both prefrontal lesions and basal ganglia lesions disrupt the ERN, suggesting that these fronto-striato-thalamo-cortical loops are necessary to generate this ERP component (Ullsperger and von Cramon, in press). Further, the results of a hemodynamic neuroimaging study have indicated that the activity of the human midbrain dopamine system is sensitive to abstract performance feedback (Aron et al., 2004). These observations suggest that, whether or not the basal ganglia indeed perform a "bottom-up" computation of predicted reward, their function is also guided by "top-down" influences from cortex.

Nevertheless, it should be noted that other aspects of the RL-ERN theory have also been challenged recently, including whether or not the fERN in fact reflects a prediction error (Hajcak et al., 2005); whether it depends on a behavioral response (Yeung et al., 2005); whether it is generated by the anterior cingulate cortex (Nieuwenhuis et al., 2005b; Van Veen et al., 2004); whether it is functionally identical with the "response error-related negativity" (response ERN), an ERP component associated with error commission (Gehring and Willoughby, 2004); whether the response ERN is related to error processing (Vidal et al., 2003); and whether the reinforcement process implemented by anterior cingulate cortex depends on midbrain dopamine (Walton et al., 2005). These issues are further complicated by the problem of component overlap of the fERN with other ERP components that are also elicited by feedback stimuli, namely the N200 and P300 (Holroyd, 2004). Answers to these questions await further research.

It is interesting to consider these results in the light of other biophysical evidence related to task context and reward processing. The RL-ERN theory holds that the value of ongoing events are evaluated in the basal ganglia, and that the fERN is elicited by the impact of phasic activity of the midbrain dopamine system on motor-related areas in anterior cingulate cortex (Holroyd and Coles, 2002). Consistent with this position and with the evidence presented in this study, a replication of the Holroyd et al. (2004a) study in a human fMRI experiment indicated that reward-sensitive areas in the basal ganglia – especially the globus pallidus and caudate nucleus – appear to respond to rewards and punishments in a context-sensitive and binary manner (Nieuwenhuis et al., 2005a). In monkeys, it has been demonstrated that the midbrain dopamine system carries phasic signals that indicate when ongoing events are better or worse than the system expects; these errors signals are made available to other neural systems for the purpose of reinforcement (reviewed in Montague et al., 2004; Schultz, 2002). Further, these phasic dopamine signals are, like the fERN, sensitive to the ongoing experimental context (Nakahara et al., 2004; Tobler et al., 2005). Unlike the fERN, however, the dopamine signals appear to be sensitive to gradations in the reward value of the outcome, with intermediate outcomes eliciting intermediate levels of phasic dopamine activity. Likewise, single-unit error-related activity in the monkey cingulate motor areas is modulated by the magnitude of the expected reward (Amiez et al., 2005). It is interesting to speculate that whereas in the monkey, this evaluative system may be relatively sensitive to the apparently objective value of the rewards, in the human, this system may be predisposed to classify outcomes into two categories: as those events that clearly indicate that a task goal has been satisfied, and as those events that do not.

Acknowledgments

REFERENCES

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

Aron, A.R., Shohamy, D., Clark, J., Myers, C., Gluck, M.A., Poldrack, R.A., 2004. Human midbrain sensitivity to cognitive feedback and uncertainty during classification learning. J. Neurophysiol. 92, 1144–1152.

Barto, A.G., 1995. Adaptive critics and the basal ganglia. In: Houk, J., Davis, J., Beiser, D. (Eds.), Models of Information Processing in the Basal Ganglia. MIT Press, Cambridge, MA, pp. 215–232.

Cohen, J.D., Blum, K.I., 2002. Reward and decision. Neuron 36, 1–20.

Cook III, E.W., 1992. VPM Reference Manual. Author, Birmingham, AL.

Cook, E.W., Miller, G.A., 1992. Digital filtering-background and tutorial for psychophysiologists. Psychophysiology 29, 350–367.

Egelman, D.M., Person, C., Montague, P.R., 1998. A computational role for dopamine delivery in human decision-making. J. Cogn. Neurosci. 10, 623–630.

Gehring, W.J., Willoughby, A.R., 2004. All are medial frontal negativities created equal? Toward a richer empirical basis for theories of action monitoring. In: Ullsperger, M., Falkenstein, M. (Eds.), Errors, Conflicts, and the Brain: Current Opinions on Performance Monitoring. MPI of Cognitive Neuroscience, Leipzig, pp. 14–20.

Gratton, G., Coles, M.G.H., Donchin, E., 1983. A new method for off-line removal of ocular artifact. Electroencephalogr. Clin. Neurophysiol. 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., Holroyd, C.B., Simons, R.F., in press. The feedback-related negativity reflects the binary evaluation of good versus bad outcomes, Biol. Psychol.

Holroyd, C.B., 2004. A note on the N200 and the feedback ERN. In: Ullsperger, M., Falkenstein, M. (Eds.), Errors, Conflicts, and the Brain: Current Opinions on Performance Monitoring. MPI of Cognitive Neuroscience, Leipzig, pp. 211–218.

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

Holroyd, C.B., 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., Larsen, J.T., Cohen, J.D., 2004a. Context dependence of the event-related brain potential associated with reward and punishment. Psychophysiology 41, 245–253.

Holroyd, C.B., Nieuwenhuis, S., Mars, R., Coles, M.G.H., 2004b. Anterior cingulate cortex, selection for action, and error processing. In: Posner, M. (Ed.), Cognitive Neuroscience of Attention. Guilford Publishing, Inc, New York, pp. 219–231.

Holroyd, C.B., Yeung, N., Coles, M.G.H., Cohen, J.D., 2005. A mechanism for error detection in speeded response time tasks. J. Exp. Psychol. Gen. 134, 163–191.

Jasper, H.H., 1958. The ten twenty electrode system of the international federation. Electroencephalogr. Clin. Neurophysiol. 10, 371–375.

Keselman, H.J., Rogan, J.C., 1980. Repeated measures F tests and psychophysiological research: Controlling the number of false positives. Psychophysiology 17, 499–503.

Miller, E.K., Cohen, J.D., 2001. An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. 24, 167–202.

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. J. Cogn. Neurosci. 9, 788–798.

Montague, P.R., Dayan, P., Sejnowski, T.J., 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. J. Neurosci. 16, 1936–1947.

Montague, P.R., Hyman, S.E., Cohen, J.D., 2004. Computational roles for dopamine in behavioral control. Nature 431, 760–767.

Muller, S.V., Moller, J., Rodriguez-Fornells, A., Munte, T.F., 2005. Brain potentials related to self-generated and external information used for performance monitoring. Clin. Neurophysiol. 116, 63–74.

Nakahara, H., Itoh, H., Kawagoe, R., Takikawa, Y., Hikosaka, O., 2004. Dopamine neurons can represent context-dependent prediction error. Neuron 41, 269–280.

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

Nieuwenhuis, S., Heslenfeld, D., Alting von Geusau, N.J., Mars, R.B., Holroyd, C.B., Yeung, N., 2005a. Activity in human reward-sensitive brain areas is strongly context dependent. NeuroImage 25, 1302–1309.

Nieuwenhuis, S., Slagter, H.A., von Geusau, N.J.A., Heslenfeld, D., Holroyd, C.B., 2005b. Knowing good from bad: differential activation of human cortical areas by positive and negative outcomes. Eur. J. Neurosci. 21, 3161–3168.

Schultz, W., 2002. Getting formal with dopamine and reward. Neuron 36, 241–263.

Suri, R.E., 2002. TD models of reward predictive responses in dopamine neurons. Neural Netw. 15, 523–533.

Sutton, R.S., Barto, A.G., 1998. Reinforcement Learning: An introduction. MIT Press, Cambridge, MA.

Tobler, P.N., Fiorillo, C.D., Schultz, W., 2005. Adaptive coding of reward value by dopamine neurons. Science 307, 1642–1645.

Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: heuristics and biases. Science 185, 1124–1131.

Ullsperger, M., von Cramon, D.Y., in press, The role of intact frontostriatal circuits in error processing, J. Cogn. Neurosci.

Van Veen, V., Holroyd, C.B., Cohen, J.D., Stenger, V.A., Carter, C.S., 2004. Errors without conflict: implication for performance monitoring theories of anterior cingulate cortex. Brain Cogn. 56, 267–276.

Vidal, F., Burle, B., Bonnet, M., Grapperon, J., Hasbroucq, T., 2003. Error negativity on correct trials: a reexamination of available data. Biol. Psychol. 64, 265–282.

Walton, M.E., Croxson, P.L., Rushworth, M.F.S., Bannerman, D.M., 2005. The mesocortical dopamine projection to anterior cingulate cortex plays no role in guiding effort-related decisions. Behav. Neurosci. 119, 323–328.

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

Yeung, N., Holroyd, C.B., Cohen, J.D., 2005. ERP correlates of feedback and reward processing in the presence and absence of response choice. Cereb. Cortex 15, 535–544.