Brief Communications

Updating Expected Action Outcome in the Medial Frontal Cortex Involves an Evaluation of Error Type

Martin E. Maier[1,2] **and Marco Steinhauser**[1]

[1]Catholic University of Eichstätt-Ingolstadt, 85072 Eichstätt, Germany, and [2]Centro Studi e Ricerche in Neuroscienze Cognitive, 47023 Cesena, Italy

Forming expectations about the outcome of an action is an important prerequisite for action control and reinforcement learning in the human brain. The medial frontal cortex (MFC) has been shown to play an important role in the representation of outcome expectations, particularly when an update of expected outcome becomes necessary because an error is detected. However, error detection alone is not always sufficient to compute expected outcome because errors can occur in various ways and different types of errors may be associated with different outcomes. In the present study, we therefore investigate whether updating expected outcome in the human MFC is based on an evaluation of error type. Our approach was to consider an electrophysiological correlate of MFC activity on errors, the error-related negativity (Ne/ERN), in a task in which two types of errors could occur. Because the two error types were associated with different amounts of monetary loss, updating expected outcomes on error trials required an evaluation of error type. Our data revealed a pattern of Ne/ERN amplitudes that closely mirrored the amount of monetary loss associated with each error type, suggesting that outcome expectations are updated based on an evaluation of error type. We propose that this is achieved by a proactive evaluation process that anticipates error types by continuously monitoring error sources or by dynamically representing possible response–outcome relations.

Introduction

Anticipating the outcome of a particular action is an important prerequisite not only for choosing future actions but also for learning from previous actions. Prediction errors reflecting violations or changes in expected outcome have been assumed to form the basis of reinforcement learning (Sutton and Barto, 1998). A crucial role in computing outcome expectation has been ascribed to the medial frontal cortex (MFC) (for review, see Rushworth and Behrens, 2008). Animal studies provided evidence that outcome expectations in the MFC are updated not only by evaluating external cues (Matsumoto et al., 2003; Amiez et al., 2005) but also by evaluating own actions (Amiez et al., 2006; Matsumoto et al., 2007; Seo and Lee, 2007). In humans, errors in choice tasks are accompanied by a negative deflection in the event-related potential, called the error-related negativity (Ne or ERN) (Falkenstein et al., 1990; Gehring et al., 1993). The Ne/ERN peaks ~50 ms after the error response and is generated in the MFC (Ridderinkhof et al., 2004). Reinforcement-learning frameworks proposed that the Ne/ERN is linked to an update of outcome expectations resulting from the detection of an error (Holroyd and Coles, 2002; Alexander and Brown, 2011; Silvetti et al., 2011).

The present study aimed to investigate whether updating expected outcome after errors involves evaluative processes that go beyond mere error detection. In everyday behavior, actions can fail in various ways and different types of errors can have different consequences. For instance, pouring salt instead of sugar into a coffee has more severe consequences than confusing milk with sugar. In such a case, an adequate computation of expected outcome cannot rely solely on error detection but requires that the type of error is identified. However, although the Ne/ERN indicates that updating outcome expectation based on error detection occurs almost simultaneously with response execution, it is unclear whether information about the error type is already available at this early time point.

Our approach was to use a speeded choice task in which two types of errors can occur (Maier et al., 2008, 2011, 2012; see also van Driel et al., 2012) and to manipulate the amount of monetary loss associated with each error type. Because each error type could occur on each trial, updating the expected outcome on error trials required not only the detection of an error but also an evaluation of the error type. If updating expected outcome were based on an evaluation of error type, the Ne/ERN should be sensitive to the specific error type (i.e., larger Ne/ERNs should be obtained for errors associated with larger monetary loss). Although previous studies have shown that the Ne/ERN varies with the amount of monetary loss (Gehring et al., 1993; Hajcak et al., 2005; Ganushchak and Schiller, 2008; Potts, 2011), in none of these studies was an evaluation of error type necessary to update outcome expectations. Therefore, such a result would demonstrate that complex performance monitoring processes beyond mere error detection inform expected outcome computation already at a very early time point.

Received July 1, 2013; revised Aug. 23, 2013; accepted Aug. 27, 2013.

Author contributions: M.E.M. and M.S. designed research; M.E.M. performed research; M.E.M. and M.S. contributed unpublished reagents/analytic tools; M.E.M. analyzed data; M.E.M. and M.S. wrote the paper.

This work was supported by the Deutsche Forschungsgemeinschaft Young Researchers Scholarship MA 4864/1-1 to M.E.M. We thank Barbara Zarrillo for assistance in conducting the experiment and Giuseppe di Pellegrino for helpful comments on an earlier version of this manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Martin E. Maier, Catholic University Eichstätt-Ingolstadt, Ostenstraße 25, 85072 Eichstätt, Germany. E-mail: martin.maier@ku.de.

DOI:10.1523/JNEUROSCI.2785-13.2013

Copyright © 2013 the authors 0270-6474/13/3315705-05$15.00/0

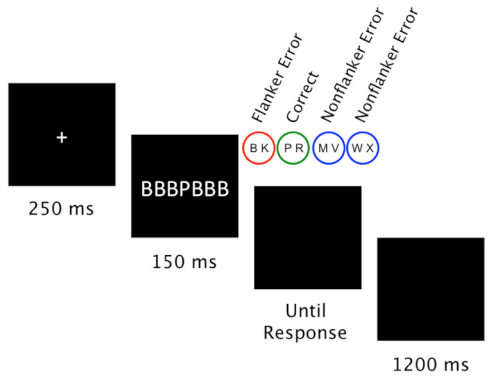

Figure 1. Sequence of events on one trial of the flanker task. Participants had to classify the central target while ignoring the flankers. Because four response alternatives were used, two different types of errors could occur for incongruent stimuli. If participants responded by pressing the button associated with the flankers (button BK in the depicted example), an FE resulted. If participants responded by pressing a button not associated with an element in the stimulus array (buttons MV and WX in the depicted example), an NFE resulted.

Materials and Methods

Participants. Thirty participants (20 female) between 17 and 27 years of age (mean 21.6 years) were recruited at the Centro Studi e Ricerche in Neuroscienze Cognitive Cesena. They were paid 10 € plus a performance-dependent bonus of up to 10 €. The study was conducted in accordance with institutional guidelines, and informed consent was acquired from all participants.

Behavioral task and procedure. We used a flanker task (Eriksen and Eriksen, 1974) with four response alternatives. On each trial, participants had to respond to a central target letter while ignoring three identical flanker letters on each side of the target. Eight target letters were mapped onto four responses: the left middle finger (B, K), the left index (P, R), the right index (M, V), and the right middle finger (W, X). A total of 48 incongruent stimuli were constructed by combining each target letter with each of the six letters not associated with the target response. A total of 48 neutral stimuli were constructed by combining each target letter with each of the six neutral symbols (§, $, %, &, #, and ?). Letters and symbols subtended a visual angle of 0.26° × 0.33° (width × height) at a viewing distance of 70 cm. The width of the whole stimulus array was 2.10°. A sample trial is depicted in Figure 1. After a 250 ms fixation cross, the stimulus was presented for 150 ms followed by a black screen. When a response was given, this triggered an interval of 1200 ms before the next trial started. This interval was restarted whenever further responses (e.g., spontaneous error corrections) occurred.

With four response alternatives, one can distinguish between two types of errors on incongruent trials. An error response could consist of either the response associated with the flankers (flanker errors [FEs]) or one of the remaining responses not associated with the flankers (nonflanker errors [NFEs]). To manipulate the outcome associated with each error type, committing an error led to a monetary loss, which differed depending on condition and error type. Participants were randomly assigned to one of two experimental groups. In the high FE group, committing an FE was associated with a high monetary loss, whereas committing an NFE was associated with a low monetary loss. In the high NFE group, the assignment of error type and monetary loss was reversed.

The experiment comprised a practice session and a test session, distributed over 2 consecutive days. In each session, participants worked through blocks of 96 trials (one trial for each possible stimulus). At the beginning of the experiment, participants were instructed that (1) they start with 10,000 points; (2) they will lose some of these points depending on their performance during the experiment; and (3) the remaining points will be changed into money (with 1000 points corresponding to 1 €) at the end of the experiment. The practice session started with two blocks without speed pressure to learn the stimulus-response mapping. After this, participants were instructed that the following blocks serve to train their response speed and that, in the test session, they will lose 20 points for each slow response. During blocks 3–10, the error rate was adjusted by an adaptive deadline procedure. At the end of each block, participants were informed about the number of "slow responses" (i.e., the number of responses exceeding a deadline). The deadline was initialized at 1000 ms. After each block, it was decreased by 50 ms whenever the error rate in this block was <15% and was increased by 50 ms whenever the error rate was >25%. The deadline from the last practice block was applied to all blocks of the test session. The test session started with two practice blocks. Then, participants were instructed about the two error types (using the same wording as in our definition described above) and the amount of loss associated with each type. In the test session, participants lost 20 points for each slow response, 1 point for each error on a neutral trial, 10 points for each error type with high loss, and 1 point for each error type with low loss. After each block, participants were informed about the number of errors of each type and the number of slow responses. Twelve test blocks with a total number of 1152 trials were conducted.

Psychophysiological recording. During the test session, the electroencephalogram (EEG) was recorded with Ag/AgCl electrodes (Fast'n'Easy-Electrodes, Easycap) from 27 electrode sites (Fp1, F3, F7, FC1, FC5, C3, CP1, CP5, P3, P7, O1, AFz, Fz, FCz, Cz, Pz, Fp2, F4, F8, FC2, FC6, C4, CP2, CP6, P4, P8, and O2) and from the right mastoid. The left mastoid was used as reference, and the ground electrode was placed on the right cheek. The electro-oculogram (EOG) was recorded from above and below the left eye and from the outer canthi of both eyes. EEG and EOG were recorded with a bandpass filter of 0.01–100 Hz, amplified by a BrainAmp DC amplifier (Brain Products), digitized at a sampling rate of 1000 Hz, and resampled to 500 Hz offline.

Data analyses. Trials with response times (RTs) deviating >4 SDs from the condition mean were excluded from RT analyses (<1% of all trials). Frequency data were arcsine-transformed for statistical testing (Winer et al., 1991).

EEG data were analyzed using EEGLAB version 6.03b (Delorme and Makeig, 2004) and custom routines written in MATLAB version 7.0.4 (MathWorks). ERP data were rereferenced off-line to the average of both mastoids and filtered with a 1–20 Hz pass-band. Epochs of 200 ms before and 800 ms after the response were extracted from the continuous EEG and baseline-corrected using a 150 to 50 ms prereresponse window. Epochs were excluded if the voltage on an EOG channel exceeded an individually adjusted threshold to remove trials with large EOG peaks or if an epoch contained data that deviated >5 SDs from the mean (joint probability criterion in EEGLAB version 6.03b). The mean percentage of excluded trials was 14.2% (SE 3.42%). Remaining EOG artifacts were corrected using a regression approach (Gratton et al., 1983). The Ne/ERN was quantified as the difference between the most positive peak in an interval of −100 and 0 ms and the most negative peak in an interval of 0 and 120 ms relative to the response. Statistical analysis was restricted to data from channel FCz at which the Ne/ERN was maximal. In addition to mixed-model ANOVAs, *post hoc* tests between low-loss and high-loss errors on Ne/ERN amplitudes were conducted using one-sided *t* tests because results from previous studies (e.g., Potts, 2011) suggest the strong directional hypothesis that the Ne/ERN is larger for high-loss relative to low-loss errors. *Post hoc t* tests on behavioral data were two-sided.

RT matching. To investigate whether RT differences between error types and conditions influenced the results, we conducted a control analysis on a subset of RT-matched trials. To this end, we matched incongruent trials with correct responses and NFEs to the FE condition, which was the condition with the fewest trials in all participants, using the following algorithm (Steinhauser and Yeung, 2012): First, an FE trial was randomly drawn without replacement. Second, from each of the remaining conditions, the trial providing the closest match to the RT of the drawn trial was selected without replacement and assigned to the RT-matched sample. These steps were repeated until all FE trials were drawn. Because only artifact-free trials were included, mean RTs on RT-matched trials deviate slightly from the results of the initial analysis.

Table 1. Behavioral performance in the flanker task

	All trials				Subset of RT-matched trials			
	High NFE group		High FE group		High NFE group		High FE group	
	Mean	SE	Mean	SE	Mean	SE	Mean	SE
RT neutral correct (ms)	591	11.9	572	11.1	—	—	—	—
RT incongruent correct (ms)	606	13.0	589	11.5	580	21.8	554	12.3
Relative frequency[a] FE (%)	10.1	1.20	7.98	0.62	—	—	—	—
Relative frequency[a] NFE (%)	14.2	1.71	12.3	0.99	—	—	—	—
RT FE (ms)	575	15.7	538	8.83	576	23.3	553	12.5
RT NFE (ms)	574	19.4	559	8.98	573	23.2	557	12.8

[a]Relative frequency is relative to all trials in the incongruent condition.

Results

Behavioral data

Task performance

Error rates and correct RTs were subjected to ANOVAs with the variables group (high FE, high NFE) and congruency (neutral, incongruent). Error rates showed a main effect of congruency ($F_{(1,28)} = 44.5$, $p < 0.001$), denoting higher error rates on incongruent (22.3%, SE 1.6%) than on neutral trials (18.5%, SE 1.4%) (all other p values > 0.27). The ANOVA on RTs (Table 1) revealed a main effect of congruency ($F_{(1,28)} = 119.9$, $p < 0.001$; all other p values > 0.29). Thus, task performance did not differ between the two groups of participants.

FE versus NFE

In a second step, we examined proportions of FE and NFE relative to all incongruent trials. To this end, we entered relative frequencies of FE and NFE in the incongruent condition into an ANOVA with the variables group (high FE, high NFE) and error type (FE, NFE). A main effect of error type was revealed ($F_{(1,28)} = 80.9$, $p < 0.001$), denoting that NFE occurred more frequently (13.2%, SE 1.0%) than FE (9.04%, SE 0.69%). Crucially, neither the main effect of group ($p > 0.28$) nor the interaction between group and error type ($p > 0.86$) was significant, indicating that the two groups differed neither with respect to the overall error rate nor with respect to the proportions of FE and NFE (Table 1). Because of three possible error responses on each trial, the relative proportion of FE among all errors (40.9%, SE 0.92%) was larger than that of each NFE (29.6%, SE 0.46% each; $t_{(28)} = 8.24$, $p < 0.001$), thus making FE disproportionally more frequent than NFE (for discussion, see Maier et al., 2011).

An analogous ANOVA on RTs of FE and NFE (Table 1) revealed a main effect of error type ($F_{(1,28)} = 4.27$, $p < 0.05$), which was qualified by a significant interaction between error type and group ($F_{(1,28)} = 5.18$, $p < 0.04$). RT was shorter for FE than for NFE in the high FE group ($t_{(14)} = 3.03$, $p < 0.01$) but comparable in the high NFE group ($p > 0.88$).

ERP data

To test the effects of incentives on the Ne/ERN, we compared the Ne/ERN for FE and NFE in both groups of participants. Grand-average waveforms of response-locked ERPs at channel FCz and scalp topographies of peak-to-peak Ne/ERN amplitudes are depicted in Figure 2. Clear Ne/ERNs and maxima of peak-to-peak Ne/ERN amplitudes over channel FCz are visible. An ANOVA with the variables group (high FE, high NFE) and error type (FE, NFE) revealed an interaction between both variables ($F_{(1,28)} = 9.45$, $p < 0.01$). In the high FE group, the Ne/ERN for FE (7.63 μV, SE 1.28 μV) was larger than the Ne/ERN for NFE (6.10 μV, SE 1.08 μV) ($t_{(14)} = 2.74$, $p < 0.05$). By contrast, in the high NFE group, the Ne/ERN was numerically larger for NFE (7.45 μV, SE 1.25 μV) than for FE (6.66 μV, SE 1.01 μV), although this difference was not significant ($t_{(14)} = 1.56$, $p = 0.07$, all other p values > 0.17).

The same analyses were conducted on an RT-matched subset of data to ensure that effects on the NE/ERN amplitude are not the result of differences in RT. An ANOVA on RTs in the RT-matched data (Table 1) with the variables group (high FE, high NFE) and trial type (correct response, FE, NFE) revealed no significant effects ($p > 0.26$), indicating that RT matching was successful. Grand-average waveforms of response-locked ERPs at channel FCz for RT-matched correct incongruent, FE and NFE trials are depicted in Figure 3. An ANOVA with the variables group (high FE, high NFE) and error type (FE, NFE) on RT-matched peak-to-peak amplitudes yielded an interaction between group and error type ($F_{(1,28)} = 11.3$, $p < 0.01$). The Ne/ERN was again larger for FE (7.63 μV, SE 1.28 μV) than for NFE (6.28 μV, SE 0.98 μV) in the high FE group ($t_{(14)} = 1.92$, $p < 0.04$). In contrast, in the high NFE group, Ne/ERN amplitudes were now significantly larger for NFE (8.05 μV, SE 1.19 μV) than for FE (6.66 μV, SE 1.01 μV) ($t_{(28)} = 3.31$, $p < 0.01$). Thus, analyses of RT-matched trials replicated the findings obtained in the analyses of all trials, with the exception that, in the high NFE group, the difference between FE and NFE was now significant.

Discussion

The present study demonstrates that MFC activity is sensitive to the amount of monetary loss in a task in which updating expected outcome requires that the type of error is evaluated. Previous studies manipulating the amount of monetary loss for errors used only a single error type, although the amount of monetary loss was known before the error was committed. In these studies, monetary loss was manipulated between blocks of trials (Gehring et al., 1993; Ganushchak and Schiller, 2008) or was indicated by a cue or the stimulus (Hajcak et al., 2005; Potts, 2011). As a consequence, detection of an error was sufficient to update expected outcome. In the present task, however, updating expected outcome required that the type of error response had to be evaluated because neither the stimulus alone nor error detection was predictive of the outcome. Our results therefore demonstrate that outcome computation receives input from performance monitoring processes that are more complex than mere error detection.

Recent models of reinforcement learning and the MFC (Holroyd and Coles, 2002; Alexander and Brown, 2011; Silvetti et al., 2011) assume that the Ne/ERN directly or indirectly reflects the generation of a negative prediction error and, thus, an update of expected outcome. Although prediction errors are sensitive to both the probability and the value of possible actions, these models were mainly used to account for effects of error probability on

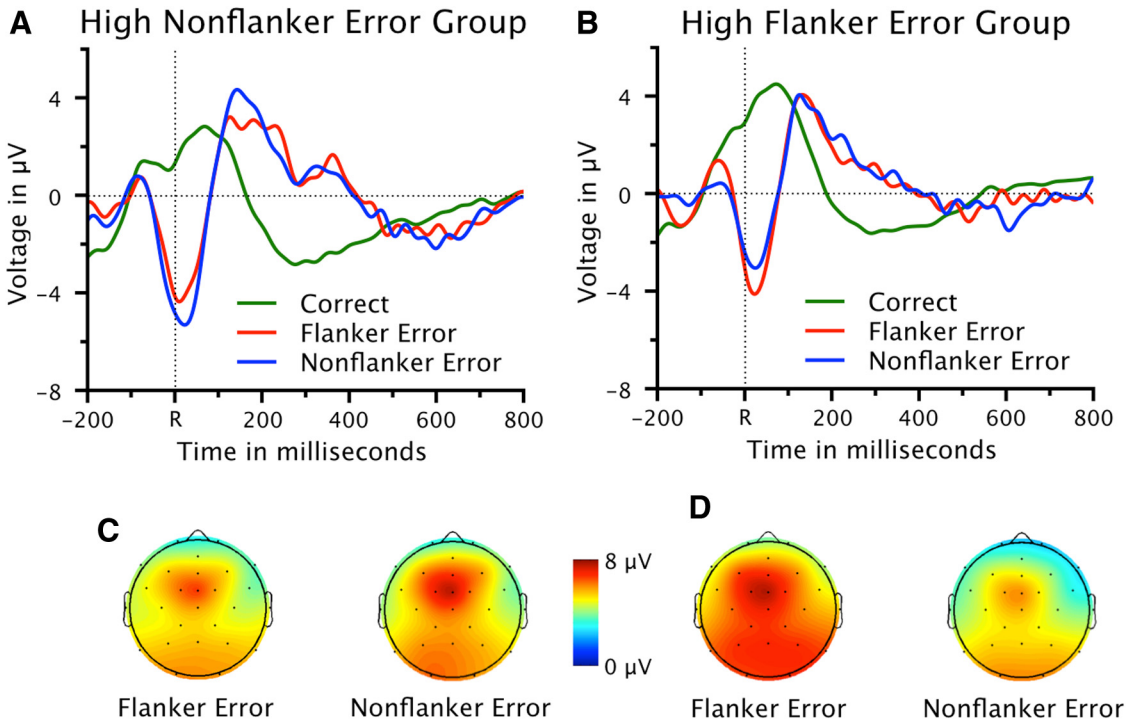

Figure 2. ERP data. ***A***, ***B***, Response-locked waveforms at electrode FCz for correct trials, FEs, and NFEs, separately for each group. ***C***, ***D***, Scalp topographies of the peak-to-peak Ne/ERN amplitudes, separately for each group. The mean latency of the Ne/ERN peak was 26 ms. R = time point of the button press.

the Ne/ERN (i.e., the inverse relation between Ne/ERN and error likelihood) (Gehring et al., 1993). In the present study, however, the pattern of Ne/ERN amplitudes was unaffected by relative frequencies of error types. NFEs were consistently more frequent than FEs and yet showed larger Ne/ERNs when associated with larger monetary loss. Furthermore, the amount of monetary loss had no effect on the relative frequency of each error type across groups. It seems that, in the present paradigm, Ne/ERN amplitude varied more strongly according to error value than according to error probability (see also Maier et al., 2012), possibly reflecting the increased salience of value resulting from the manipulation of monetary loss.

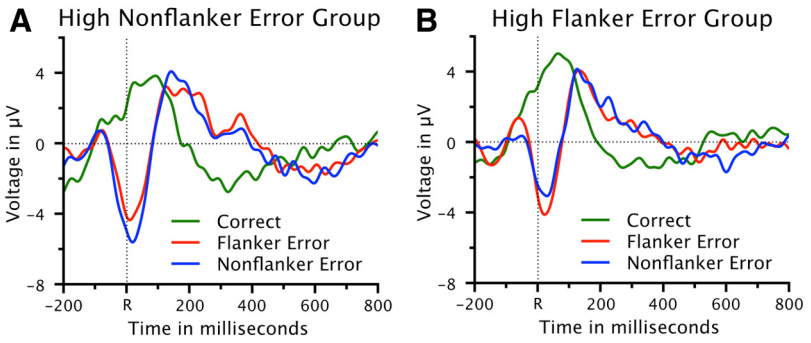

Figure 3. RT-matched ERP data. Response-locked waveforms at electrode FCz for correct trials, FEs, and NFEs for a subset of RT-matched trials, separately for each group. ***A***, High NFE group. ***B***, High FE group. R = time point of the button press.

The question arises how information about the error type can be available immediately after the response. An intuitive solution would be a retroactive evaluation process that operates after the detection of an error. This process could evaluate whether the just executed error response corresponds to an FE by comparing the response with a representation of the flanker in working memory. Although such a process is not impossible, it seems rather implausible given the short period between the response and the onset of the Ne/ERN. Instead, we propose a proactive evaluation process that anticipates the error type before a response is executed. Such a process could operate in one or both of two ways:

1. The evaluation process could monitor possible error sources to anticipate the probability that a given error corresponds to an FE, p_{FE}, or NFE ($1 - p_{FE}$). Using these probabilities, error detection would allow for an update of expected outcome according to $p_{FE} \times V_{FE} + (1 - p_{FE}) \times V_{NFE}$, with V_{FE} *and* V_{NFE} reflecting monetary loss associated with flanker and NFEs, respectively. Because FEs are more likely when attention is erroneously directed to the flankers, the state of selective attention could be a valid estimator of p_{FE}. Such an account would imply a monitoring mechanism that learns to map specific error sources to action outcomes in the course of an experiment.

2. The evaluation process could anticipate the outcomes of all possible responses given the stimulus (response–outcome relations) (Alexander and Brown, 2011). The occurrence of a specific response would allow for an update of expected outcome according to the respective response–outcome relation. Although such a mechanism is computationally expensive, it allows for a very parsimonious account of error detection and error type evaluation; both result implicitly as a byproduct of the representation of response–outcome relations. Indeed, the PRO model by Alexander and Brown (2011) implements such a mechanism and assumes that the context-dependent computation of response–outcome relations is a core function of the MFC.

The two possibilities are not mutually exclusive, and further research is necessary to determine which of these two mechanisms is actually involved. A viable approach to test the idea of an error source monitor might be to ask whether Ne/ERN differences across error types are linked to specific error sources, such as impaired selective attention. First evidence for this can be derived from a recent study in which we used multinomial modeling to determine the proportion of FEs that were the result of attention errors (Maier et al., 2012). We found larger Ne/ERN amplitudes on FEs when the estimated proportion of attention errors among FEs was increased. This could reflect that impaired selective attention is used to anticipate that an error will be an FE.

Together, we have shown that the Ne/ERN is sensitive to the amount of monetary loss, even when updating the expected outcome requires an evaluation of error type. Our results add to the literature showing that the MFC is crucially involved in the encoding of action outcomes (Rushworth and Behrens, 2008; Rushworth et al., 2012) and demonstrate that updating expected outcome is not only guided by an evaluation of environmental cues but also by a fast and yet complex evaluation of own actions.

References

Alexander WH, Brown JW (2011) Medial prefrontal cortex as an action-outcome predictor. Nat Neurosci 14:1338–1344. CrossRef Medline

Amiez C, Joseph JP, Procyk E (2005) Anterior cingulate error-related activity is modulated by predicted reward. Eur J Neurosci 21:3447–3452. CrossRef Medline

Amiez C, Joseph JP, Procyk E (2006) Reward encoding in the monkey anterior cingulate cortex. Cereb Cortex 16:1040–1055. CrossRef Medline

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

Eriksen BA, Eriksen CW (1974) Effects of noise letters upon the identification of a target letter in a nonsearch task. Percept Psychophys 16:143–149. CrossRef

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: Psychophysiological brain research, Vol. 1 (Brunia CHM, Gaillard AWK, Kok A, eds), pp 192–195). Tilburg, The Netherlands: UP.

Ganushchak LY, Schiller NO (2008) Motivation and semantic context affect brain error-monitoring activity: an event-related brain potentials study. Neuroimage 39:395–405. CrossRef Medline

Gehring WJ, Goss B, Coles MGH, Meyer DE, Donchin E (1993) A neural system for error detection and compensation. Psychol Sci 4:385–390. CrossRef

Gratton G, Coles MG, Donchin E (1983) A new method for off-line removal of ocular artifact. Electroencephalogr Clin Neurophysiol 55:468–484. CrossRef Medline

Hajcak G, Moser JS, Yeung N, Simons RF (2005) On the ERN and the significance of errors. Psychophysiology 42:151–160. CrossRef Medline

Holroyd CB, Coles MG (2002) The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. Psychol Rev 109:679–709. CrossRef Medline

Maier ME, Yeung N, Steinhauser M (2011) Error-related brain activity and adjustments of selective attention following errors. Neuroimage 56:2339–2347. CrossRef Medline

Maier ME, di Pellegrino G, Steinhauser M (2012) Enhanced error-related negativity on flanker errors: error expectancy or error significance? Psychophysiology 49:899–908. CrossRef Medline

Maier M, Steinhauser M, Hübner R (2008) Is the error-related negativity amplitude related to error detectability? Evidence of effects from different error types. J Cogn Neurosci 20:2263–2273. CrossRef Medline

Matsumoto K, Suzuki W, Tanaka K (2003) Neural correlates of goal-based motor selection in the prefrontal cortex. Science 301:229–232. CrossRef Medline

Matsumoto M, Matsumoto K, Abe H, Tanaka K (2007) Medial prefrontal cell activity signaling prediction errors of action values. Nat Neurosci 10:647–656. CrossRef Medline

Potts GF (2011) Impact of reward and punishment motivation on behavior monitoring as indexed by the error-related negativity. Int J Psychophysiol 81:324–331. CrossRef Medline

Ridderinkhof KR, Ullsperger M, Crone EA, Nieuwenhuis S (2004) The role of the medial frontal cortex in cognitive control. Science 306:443–447. CrossRef Medline

Rushworth MF, Behrens TE (2008) Choice, uncertainty and value in prefrontal and cingulate cortex. Nat Neurosci 11:389–397. CrossRef Medline

Rushworth MFS, Kolling N, Sallet J, Mars RB (2012) Valuation and decision-making in frontal cortex: one or many serial or parallel systems? Curr Opin Neurobiol 22:946–955. CrossRef Medline

Seo H, Lee D (2007) Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. J Neurosci 27:8366–8377. CrossRef Medline

Silvetti M, Seurinck R, Verguts T (2011) Value and prediction error in medial frontal cortex: integrating the single-unit and systems levels of analysis. Front Hum Neurosci 5:75. CrossRef Medline

Steinhauser M, Yeung N (2012) Error awareness as evidence accumulation: effects of speed-accuracy trade-off on error signaling. Front Hum Neurosci 6:240. CrossRef Medline

Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. Cambridge, MA: Massachusetts Institute of Technology.

van Driel J, Ridderinkhof KR, Cohen MX (2012) Not all errors are alike: θ and α EEG dynamics relate to differences in error-processing dynamics. J Neurosci 32:16795–16806. CrossRef Medline

Winer BJ, Brown DR, Michels KM (1991) Statistical principles in experimental design. London: McGraw-Hill.