

The Importance of Failure: Feedback-Related Negativity Predicts Motor Learning Efficiency

Jurjen van der Helden¹, Maarten A. S. Boksem² and
Jorian H. G. Blom¹

¹Department of Cognitive Psychology and Ergonomics,
University of Twente, 7522 NB Enschede, the Netherlands and
²Department of Social Psychology, Tilburg University, 5000 LE
Tilburg, the Netherlands

Jurjen van der Helden and Maarten A. S. Boksem have
contributed equally to this work.

Learning from past mistakes is of prominent importance for successful future behavior. In the present study, we tested whether reinforcement learning signals in the brain are predictive of adequate learning of a sequence of motor actions. We recorded event-related potentials (ERPs) while subjects engaged in a sequence learning task. The results showed that brain responses to feedback (the feedback-related negativity [FRN]) predicted whether subjects learned to avoid an erroneous response the next time this action had to be performed. Our findings add to a growing literature on feedback-based performance adjustment, by showing that FRN amplitudes may reflect the acquisition of motor skill and the consolidation of contingencies between stimuli or cues and their associated responses, providing evidence that learning efficiency and future performance can be predicted by the neural response to current feedback: FRN amplitude associated with a mistake is predictive of whether this mistake will be repeated, or learned from.

Keywords: ACC, ERP, FRN, reinforcement learning

Introduction

Although there are examples of exceptions, in general the best decisions are made by people who actually know what they are doing. These people have learned from (often bitter) experience the consequences of their actions and are therefore now able to select the actions that they know will have the greatest probability of success. Therefore, learning from past mistakes is of prominent importance for successful future behavior. Reinforcement learning (RL) theory has been developed to describe how organisms are able to learn these action–outcome associations (Barto and Sutton 1997). In a typical RL model, behavioral options that have a high expected value are preferred over options with lower expected values. Whenever the expected outcome differs from the actual outcome, this is coded as a reward prediction error. This error signal is then used to update the expected reward value of the chosen behavioral option so that it better reflects the observed reward value. That is, the system learns which actions result in desirable outcomes.

The work of Schultz (Schultz et al. 1997; Schultz 2002, 2004) suggests that these reward prediction errors are encoded in midbrain dopamine (DA) neurons. These neurons have been shown to respond with increased activity when outcomes are better than expected (positive reward prediction error), whereas decreases in activity occur when outcomes are below expectations (negative reward prediction error). Holroyd and Coles (2002) suggested that this negative RL error is conveyed to the anterior cingulate cortex (ACC), where it produces an error signal that can be measured as a negative event-related

potential (ERP) on the scalp, called the feedback-related negativity (FRN).

The FRN consists of a negative shift in the ERP occurring 200–400 ms after the presentation of feedback informing the subjects about the outcome of their performance (Miltner et al. 1997; Luu et al. 2003; Nieuwenhuis et al. 2004). Consistent with the proposal by Holroyd and Coles (2002), source localization studies have suggested that the FRN is indeed generated in the ACC (Gehring and Willoughby 2002; Nieuwenhuis et al. 2004). In addition, processing of negative feedback in the ACC has been reported in studies using functional magnetic resonance imaging (fMRI; Nieuwenhuis et al. 2005, 2007).

If neural prediction error signals, reflected in the FRN, are indeed indicators of RL processes, one would predict that the FRN actually predicts RL. That is, high-amplitude FRNs in response to feedback on a given action should indicate adequate updating of action–outcome contingencies and should therefore be associated with good performance when the subject performs this action on a future occasion. Thus far, however, the evidence for this is limited.

Learning has been defined as the act, process, or experience of gaining knowledge or skill. Although there have been reports of FRN (or error-related negativity [ERN]) being related to corrective actions or performance adjustments (e.g., Gehring et al. 1993; Ridderinkhof et al. 2003; Holroyd and Krigolson 2007), which can be considered a basic form of learning, evidence that these ERP components reflect the process of actually gaining knowledge or skill remains forthcoming. To be able to show that the FRN is indeed related to the acquisition or learning of a skill requires that subjects need to consolidate contingencies between cues and their associated responses, or actions and their associated outcomes. Therefore, to show that the FRN reflects the process of learning, one would have to show that the amplitude of this ERP component is associated with consolidating these contingencies.

We designed a task to specifically evaluate feedback-contingent motor learning. In this task, participants had to learn a sequence of button presses by trial and error (see Fig. 1). The task was such that every time subjects chose the correct button press out of a possible 4 options, they progressed to the next item in the sequence of 12 button-presses; if, however, they chose the wrong button, the sequence would restart at item 1 of that sequence. Restarting the sequence allowed us to relate FRN amplitude elicited by feedback on a particular item to performance on that same item when it was encountered for a subsequent time. The RL theory of the FRN would predict that increased FRN amplitudes elicited by feedback to a specific choice will be

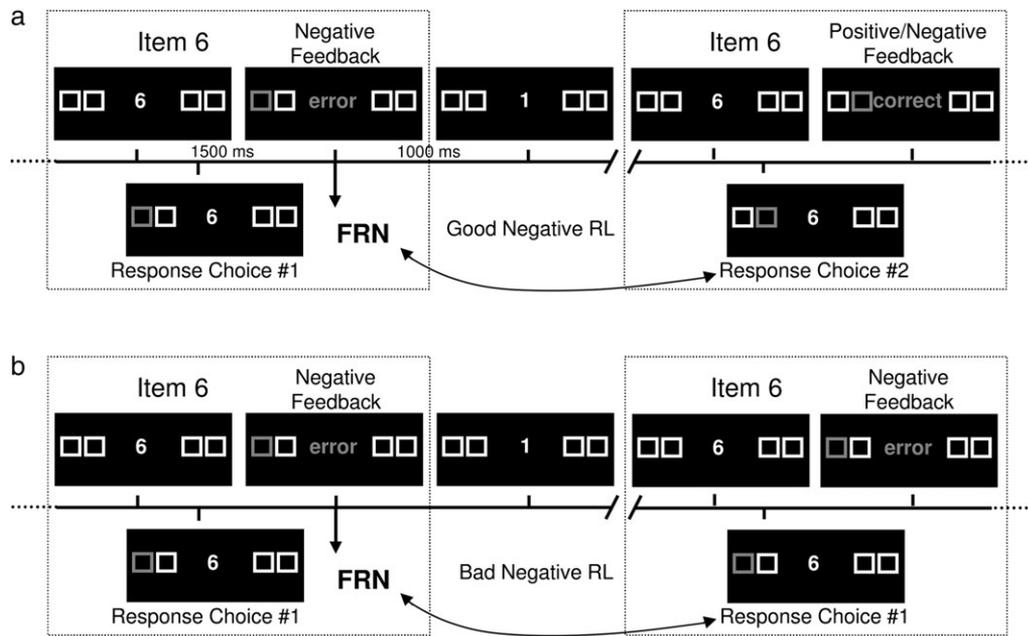


Figure 1. (a) An example of a good negative RL trial. (b) An example of bad negative RL. In both examples, the participant first chooses the leftmost button on item 6. In example (a), the next choice is the second button from the left. This choice is, in this example, followed by positive feedback after which the participant can proceed with item 7 but can also be followed by negative feedback (not shown) after which the participant has to restart with item 1. In both cases, however, this trial is classified as “good learning” because the participant correctly refrained from choosing the same button as on the previous trial. In example (b), the participant’s choice is the same as on the first encounter of this item (and for which he/she received negative feedback). This indicates that the subject has not learned from the previous negative feedback and this trial is therefore referred to as a “bad learning” trial.

associated with good performance on future instantiations of this same choice (i.e., subjects learned the correct response), whereas less pronounced FRN amplitudes would be associated with bad future performance (i.e., subjects did not learn the correct response).

In addition, to rule out changes in learning efficiency and FRN amplitude resulting from fluctuations in attention, we also measured P3 amplitudes, which have been shown to be indicative of the efficiency of attentional processes (see Herrmann and Knight 2001 for a review).

Materials and Methods

Participants

Nineteen participants (8 men) between 17 and 23 (mean [M] = 19.8, standard deviation [SD] = 1.8) years of age were recruited from the university population and received course credit for their participation. Handedness was indexed by the Annett Handedness Inventory (Annett 1970). Twelve participants described themselves as being right-handed. One participant was classified as being ambidextrous and 6 participants were classified as left-handed. All participants had normal or corrected-to-normal vision. Written informed consent was obtained prior to the experiment.

Stimuli

Figure 1 shows a schematic representation of the task. Stimuli were presented on a black background. Four white squares ($1.4^\circ \times 1.4^\circ$ each; 10.4° wide in total) that represented the 4 response buttons were presented on screen for the entire duration of the task. Each trial began with the presentation of the current item number in white ($0.4^\circ \times 0.6^\circ$) on fixation. Participants were asked to choose between 4 response buttons as soon as possible after presentation of the current item number. If their choice exceeded the time limit (1500 ms), participants were presented visual feedback “too late” ($3.3^\circ \times 0.6^\circ$, in blue font) on

fixation, indicating they responded too late on the current trial. If participants chose, for example, to push the leftmost button (pushed with their left middle finger), the leftmost square on the screen instantly changed from white to blue. About 1000 ms after the participant’s response, feedback was presented visually on fixation. If the participant’s choice was correct, they received positive feedback (i.e., “correct”) presented in green; if the participant’s choice was incorrect, they received negative feedback (i.e., “error”) presented in red. The visual feedback remained on screen for 1000 ms, until the start of the following trial.

Procedure

Stimuli were presented with the E-Prime package (version 1.2; Psychology Software Tools, Inc., Pittsburgh, PA; <http://www.psnet.com>) on a 17-inch monitor and responses were collected through an E-Prime-compatible PST Serial Response Box.

Participants were asked to learn a sequence of 12 specific responses. They were instructed to learn this sequence by trial and error. When subjects chose the correct response, the task proceeded with the next item in the sequence. If they chose an incorrect response or if they did not respond in time, the sequence restarted at item number 1. Participants thus only successfully completed a sequence if all 12 items were responded to correctly in a row. When a sequence was completed, participants received overall feedback and a short break (30 s) before they proceeded with the next sequence. In total, participants completed a maximum of 10 sequences.

Unbeknownst to the subjects, for each item in the sequence, it was manipulated how many response choices were considered incorrect before a choice would be positively reinforced and considered as correct. In every sequence (of 12 items), there were 3 items for which responses were not considered correct until the first, second, third, and fourth encounter of that item choice. Thus, it was predetermined how many attempts were required to get positive feedback for a particular item in the sequence. The order in which these items occurred was randomized within every sequence. This ensured that there were no differences in performance between subjects and sequences due to better guessing.

We distinguished between 3 types of feedback-performance contingencies (see Fig. 1). First, we labeled negative feedback that was followed by a novel response choice on that same item in the sequence as “good negative RL” as this indicates that the participant has learned from the feedback and chose a response he or she had not tried before. Second, we labeled positive feedback that was followed by the same response choice on the same item as “good positive RL” as this indicates that the participant consolidated the appropriate response to this item in accordance with the feedback signal (because positive feedback is only informative the first time a trial is performed correctly, only these trials were included in the good positive RL ERPs). Third, we labeled negative feedback that was followed by a response choice that the subject had tried before on that same item as “bad negative RL.” Finally, we labeled positive feedback that was followed by an alternative response choice on the same item as “bad positive RL.” This however happened too infrequently to warrant ERP analyses, so this category will not be considered further.

Electroencephalography Recording and Data Reduction

Electroencephalography (EEG) was recorded from 61 standard channels (10-20 system; Pivek et al. 1993), using Ag/AgCl ring electrodes mounted on an electrocap (EasyCap), with a forehead ground and an online average reference. The vertical and horizontal electro-oculograms were measured from electrodes above and below the left eye and from the outer canthi of both eyes, respectively. Electrode impedance was kept below 5 k Ω . Signals were passed through a BrainAmp amplifier (Brain Products GmbH, Munich, Germany; www.brainproducts.com), recorded online at a sample rate of 500 Hz, offline filtered with a 200-Hz low-pass filter and a notch filter of 50 Hz, and amplified with BrainAmp amplifiers (Brain Products).

Data Analysis

EEG segments containing artifacts ($\pm 100 \mu\text{V}$) and eye movements ($\pm 100 \mu\text{V}$) were rejected. EEG artifact detection resulted in rejection of 2.6% of the segments. Additional eye movement detection resulted in rejection of a further 15.1% of the segments. In addition, remaining ocular artifacts were corrected with the Gratton-Coles algorithm (Gratton et al. 1983). ERPs of feedback signals associated with good negative RL, bad negative RL, and good positive RL were analyzed and averaged separately. A baseline voltage averaged over the 100-ms interval preceding the onset of the feedback signal was subtracted from the averages.

To minimize the effects of overlap between the FRN and other ERP components, most notably the P3, we created difference waves (see Holroyd and Krigolson 2007; Holroyd et al. 2008) by subtracting the ERPs associated with good positive RL from 1) the ERPs associated with good negative RL, creating a “good learning” difference wave, and 2) the ERPs associated with bad negative RL, creating a “bad learning” difference wave. Finally, we created a “learning” difference wave by 3) subtracting ERPs associated with good negative RL from ERPs associated with bad negative RL. Visual inspection of grand-averaged difference waveforms and their scalp distributions (Fig. 3) indicated an FRN that reached its maximum at a latency around 285 ms after feedback presentation on FCz. Because peak detection proved to be unreliable in individual subject data (especially for FRNs generated by positive feedback), we submitted the average ERP difference wave amplitude in a time window of 270–300 ms after feedback to statistical analyses. To further rule out potential contamination of the FRN by the P3, we also submitted difference wave data from Pz at a latency of 300–350 ms, where visual inspection showed this component to have its maximum, to further statistical analyses.

Finally, we separately analyzed FRNs recorded at the first, second, third, and fourth attempt at a particular item in the sequence. Feedback received after the third attempt may be much more informative (there is only one remaining response option) than feedback received after the first attempt (when there are still 3 options remaining). In addition, negative feedback after the first attempt may not be as unexpected as negative feedback received after the fourth attempt. Indeed, it has been shown that the FRN is also affected by the expectedness of an outcome (Hajcak et al. 2007). Therefore, we decided to run an additional analysis

that included the factor attempt. However, because of an insufficient number of bad negative RL trials, this analysis was only possible for good positive RL and good negative RL trials.

Results

Performance

Because the type of feedback subjects received when they encountered a certain trial was manipulated in the experiment (it was predetermined that items required 4, 3, 2, or 1 attempts before positive feedback was given), the lowest number of errors that subjects could make on a particular sequence was 18. In addition to these 18 errors, participants made 12.4 errors (SD = 4.8) on average per sequence. These errors consisted of failures to refrain from repeating negatively reinforced responses (negative RL failures; $M = 3.5$, $SD = 1.8$) and of failures in reproducing positively reinforced responses (positive RL failures; $M = 8.8$, $SD = 7.3$). The negative RL failures consisted of choosing the same incorrect response on the next encounter of that item ($M = 1.6$, $SD = 0.7$) and of choosing this incorrect response on a later encounter of that item ($M = 2.0$, $SD = 1.3$). These 2 types of negative RL failures occurred equally often, $t(18) = 1.6$, not significant (NS), and were both used to compute the bad negative RL ERPs. Of the positive RL failures, only a small proportion ($M = 1.5$, $SD = 0.6$) consisted of “true” positive RL failures (i.e., a failure to respond correctly to an item while a positive feedback had been received on the previous encounter of that item); the remaining failures to reproduce positively reinforced responses involved errors that were made after the subject had already responded correctly to that trial repeatedly ($M = 7.3$, $SD = 5.0$). Because these errors are difficult to relate to learning processes, they will not be considered further.

In addition, we analyzed if repetitions of negatively reinforced responses depended on the number of successive attempts at that item. For each subject, we scored the number of immediate repetitions of negatively reinforced responses after the first, second, and third attempt at that item and also for later repetitions of negatively reinforced responses. Because all items required a first attempt and more second attempts are made than third attempts, we corrected for these differences in frequency. We analyzed these data in a 2×3 design with bad negative RL type (“immediate” repetition and “later” repetition) and attempt (first, second, and third) as factors. RL type ($F_{1,18} = 2.5$), attempt ($F_{2,36} = 1.0$), and their interaction ($F_{2,36} = 0.1$) were all NS, showing that the number of learning errors was independent from the number of attempts made at a particular item.

Finally, we calculated the number of learning errors for every sequence for each subject. The results show that subjects gradually made fewer errors. We tested this statistically by averaging the number of errors in the first and the last 3 sequences for each subject and t -tested these averages. Subjects made significantly more errors at the start of the experiment as compared with the last part of the experiment, $t(18) = 2.6$, $P < 0.05$.

Event-Related Potentials

Negative feedback (averaged over all error trials) elicited a negative deflection in the ERP in the latency range of interest that was not observed for positive feedback (averaged over all

correct trials), resulting in a negative-going FRN difference wave that was significantly different from zero, $t(18) = -6.7, P < 0.001$ (see Fig. 2a). Figure 2b shows that this FRN has a fronto-central distribution, in accordance with previous studies reporting FRN data. Importantly, when we distinguish between negative feedback that is followed by a novel response (i.e., good negative RL) and negative feedback that is followed by a response that had already been tried before (i.e., bad negative

RL), we find that whereas FRNs elicited by both types of feedback differ from that elicited by positive feedback, difference wave $-3.8 \mu\text{V}$, $t(18) = -6.7, P < 0.001$, and $-2.4 \mu\text{V}$, $t(18) = -3.5, P < 0.005$, respectively (see Figs 3 and 4), the FRN elicited by good negative RL was significantly enhanced compared with that elicited by bad negative RL, $-1.3 \mu\text{V}$, $t(18) = -3.0, P < 0.01$. In addition, the “learning difference wave” (good negative RL - bad negative RL; Fig. 3b) was significantly different from zero between 150 and 500 ms on FCz, $t(18) < -2.7, P < 0.05$). This indicates that a large FRN following negative feedback is predictive of not selecting a response that has already been tried before, the next time this same item in the sequence has to be performed.

Turning to the P3 data collected from Pz (300–350 ms), we found that negative feedback is associated with a reduction in P3 amplitude, $-2.4 \mu\text{V}$, $t(18) = 5.7, P < 0.001$ (see Figs 3 and 4). In contrast to the data on the FRN, however, we found no difference between P3 amplitudes associated with good negative RL and bad negative RL, $t(18) = -5, \text{NS}$.

In Figure 5, the ERPs are shown as a function of the successive attempts. We found that FRNs on good positive RL trials were significantly different depending on the number of attempts, on both FCz ($F_{3,54} = 5.8, P < 0.005$) and Pz ($F_{3,54} = 4.8, P < 0.05$). Contrast analysis showed that this effect was caused by a reduced positivity when positive feedback was given after the fourth attempt, compared with the other attempts, $t(18) = 2.46\text{--}4.95, P < 0.025$ (see Fig. 5a). None of the other contrasts reached significance. FRNs associated with good negative RL only showed a marginally significant effect of attempt on FCz ($F_{2,36} = 3.3, P = 0.07$) and not at all on Pz ($F_{2,36} = 0.416, \text{NS}$; see Fig. 5b). Also, none of the contrasts reached significance, $t(18) > -2.1, P > 0.05$. For the good learning difference wave (Fig. 5c), no effect of attempt was observed on the FRN ($F_{2,36} = 2.1, \text{NS}$), nor on the P3 ($F_{2,36} = 0.76, \text{NS}$).

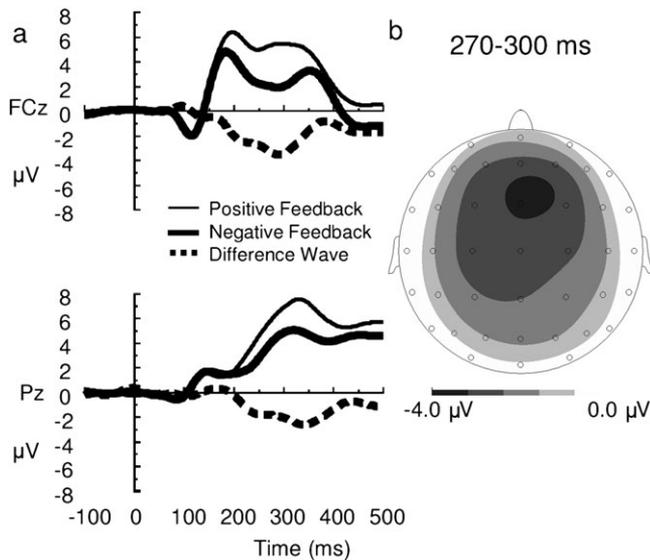


Figure 2. The left graph (a) shows the ERP to positive feedback (thin line), negative feedback (strong line), and the associated difference wave (dashed line) on FCz and Pz. In the right graph (b), the topographical distribution of this difference wave in the 270- to 300-ms latency range is presented. As can be seen, negative feedback compared with positive feedback elicits a fronto-central negativity. The change in potential between adjacent contours is $0.8 \mu\text{V}$.

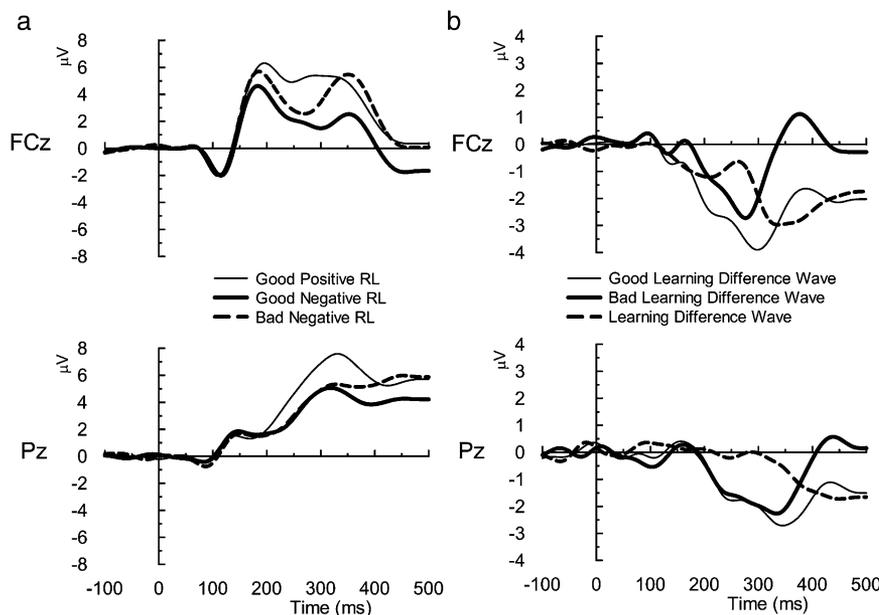


Figure 3. (a) FRN waveforms associated with the different feedback–performance contingencies at channels FCz and Pz. The thin line represents good positive RL; the strong line, good negative RL; and the dashed line, bad negative RL. Difference waves are presented in the graphs on the right (b). Here, the thin line (good learning difference wave) represents difference waves obtained by subtracting good positive RL from good negative RL, and the strong line (bad learning difference wave) represents difference waves obtained by subtracting good positive RL from bad negative RL. The dashed line (learning difference wave) represents the difference wave obtained by subtracting bad negative RL from good negative RL. As can be seen in these graphs, negative feedback resulted in more negative-going deflections than positive feedback, particularly on FCz. Furthermore, good negative RL elicited a larger FRN than bad negative RL, only on FCz.

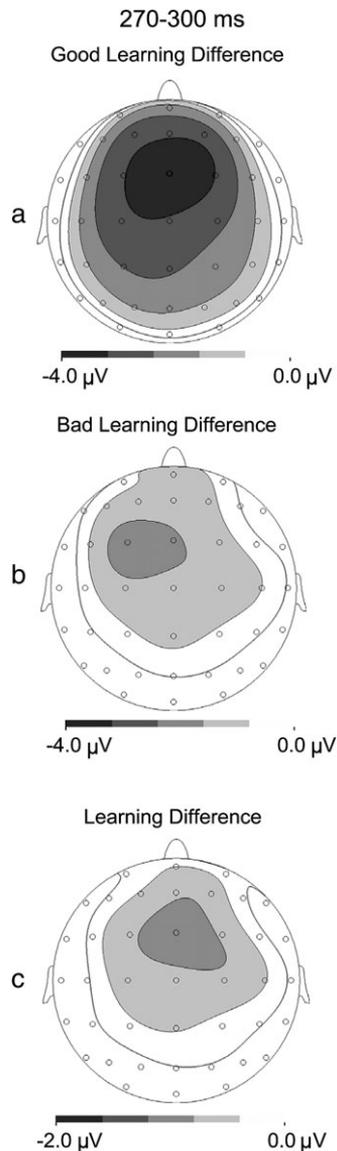


Figure 4. The topographical distributions of the difference waves are sampled at 280-ms after feedback. (a) Topographical distribution of the difference wave of good negative RL–good positive RL; (b) topographical distribution of the difference wave of bad negative RL–good positive RL; (c) topographical distribution of the difference wave of good negative RL–bad negative RL. In the top and middle graphs, the change in potential between adjacent contours is 0.8 μ V. In the lower graph, the change in potential between adjacent contours is 0.4 μ V.

Discussion

Learning from past mistakes is of prominent importance for successful future behavior. In recent years, RL theory has been developed to describe how organisms are able to learn which actions result in desirable outcomes. Holroyd and Coles (2002) suggested that an RL (reward prediction error) signal from the midbrain DA system is conveyed to the ACC, where it produces an error signal that can be measured as a negative ERP on the scalp, called the FRN.

If neural prediction error signals, reflected in the FRN, are indeed indicators of RL processes, one would expect that the FRN actually predicts learning from reinforcement. That is, high-amplitude FRNs in response to negative feedback on a given choice should indicate adequate updating of action-

outcome contingencies and should therefore be associated with good performance when the subject is confronted with this same choice on a future occasion. Here, we provide evidence that this is indeed the case. In the present experiment, participants had to learn a sequence of button presses by trial and error. This paradigm allowed us to relate FRN amplitude elicited by feedback on a particular response choice to performance when this same choice was encountered a subsequent time. When, after negative feedback, subjects again pressed a button that they could have known was incorrect, we classified this as bad negative RL; when subjects chose a button they had not tried before on the following occasion the item was presented, we classified this as good negative RL. When subjects received positive feedback and chose that same button on the next occasion the item was presented, we classified this as good positive RL. The RL theory of the FRN would predict that good RL is associated with increased FRN amplitudes compared with bad RL.

The results strongly support the RL theory of the FRN. That is, when receiving negative feedback for a certain action, FRN amplitude is more negative when behavior is subsequently successfully adjusted (i.e., when subjects have learned from the feedback and choose a response they have not tried before), compared with when after negative feedback an erroneous response is repeated (i.e., subjects have not learned from the feedback and choose a response they have already tried). In other words, the FRN amplitude reflects whether action–outcome associations are adequately updated, thus predicting future performance. These findings closely resemble previous results obtained using fMRI, showing that activity in the posterior medial prefrontal cortex during errors was predictive of whether future responses would be correct or incorrect (Hester et al. 2008).

Importantly, the difference in ERP amplitudes associated with good and bad RL was shown to be limited to the FRN and was not observed for P3 amplitudes. This is in itself an important finding: P3 amplitudes have been shown to be indicative of attentional processes (see Herrmann and Knight 2001 for a review). Thus, the bad learning performance associated with attenuated FRNs observed in the present study cannot be explained by lapses of attention on these trials. Instead, it appears that specifically the RL processes reflected by the FRN are involved in the adequate updating of action–outcome associations in the present study. This interpretation is supported by our finding that error rates did not depend on the number of attempts made at a particular item and also by the fact that error rates did not increase with time on task (which may be accompanied by reduced vigilance and levels of attention, and has been associated with reduced ERN amplitudes; see Boksem, Meijman, et al. 2006, Boksem, Tops, et al. 2006).

When we analyze our data based on whether it is the first, second, or third attempt at that particular item in the sequence, we observe no differences in the FRNs elicited by negative feedback. Although negative feedback after the first attempt (when there are still 3 remaining response options) may be considered less informative than negative feedback after attempt 3 (when there is only one possible correct option left), this is not reflected in the FRN in the present experiment. This may indicate that subjects do not use this negative feedback to update the contingency between the stimulus and the possible correct response, but rather, they update which response not to make to that particular item in the sequence.

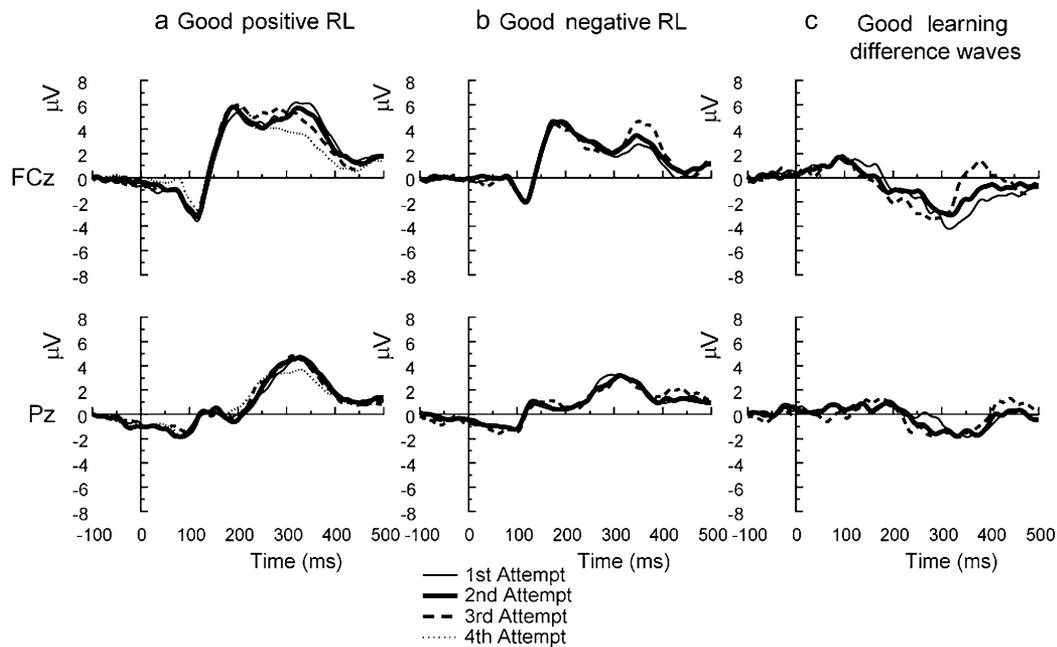


Figure 5. ERP waveforms recorded for the first, second, third, and fourth attempt at a particular item in the sequence, for trials associated with (a) good positive RL and (b) good negative RL. The associated good learning difference waves are shown in (c).

Indeed, this is how we operationalized good learning in the present study: not choosing an erroneous response again. However, separately analyzing FRNs for the different attempts subjects made leaves us with a very low number of trials per attempt, especially in the third and fourth attempt conditions, so these data should be interpreted with caution.

The only difference we found between the attempts is that the ERP associated with good positive RL on the fourth attempt was less positive on FCz than the ERP associated with the other attempts. Because with positive feedback, there is no real difference in how informative this feedback is, we suggest that this difference may be related to the expectancy of the feedback. After having received negative feedback on 3 of the 4 possible response options, there is only one possible correct response remaining, so only positive feedback on the fourth attempt at that item is fully expected. Compared with the more unexpected positive feedback, the ERP associated with the fourth attempt mainly shows a reduced positivity on the P3-like deflection following the negative deflection in the FRN latency range (see Fig. 5a). Indeed, the P3 has been shown to be sensitive to expectancy violations (e.g., Duncan-Johnson and Donchin 1977). Again, because of the low number of trials per attempt, this interpretation has to remain speculative. Importantly, the difference wave approach followed in the present study avoids the confounds of expectancy associated with overlapping P3 activity in the FRN latency range.

In line with the original proposal by Holroyd and Coles (2002), some previous studies have already suggested that the FRN reflects an RL process. However, the present findings add to these previous results in 2 important ways. First, the present study reports an association between FRN amplitude and future performance: Although previous studies have shown that ERN and FRN amplitudes are related to whether a particular stimulus-response association has been learned, we show here that FRN amplitude actually predicts whether such an association will be learned. For example, using a probabilistic

reward task (one stimulus, the rich stimulus, is differentially more rewarded than the other), Santesso et al. (2008) found that as the task progressed, those subjects that showed a bias toward selecting the rich stimulus, displayed a more positive FRN upon receiving positive feedback. Although providing valuable insight in the processes reflected in the FRN, this effect is somewhat difficult to interpret in terms of learning: At the time the FRN was measured in this study, subjects had (or had not) already developed a bias toward responding toward the rich stimulus. In other words, some subjects had already learned that the rich stimulus was more valuable than the other stimulus, whereas the others had not learned this association. Then, after selecting the rich stimulus, subjects who already knew that this was the more valuable of the 2 stimuli showed a positive ERP deflection, whereas subjects who did not know this showed a more negative ERP deflection.

These results bring to mind recent results from Oliveira et al. (2007). These authors demonstrated that the FRN not only reflects that outcomes are below expectations but can also be elicited by positive outcomes, when subjects are expecting a negative outcome, indicating that FRN does not reflect that outcomes are worse than expected but that outcomes are simply different than expected. Thus, for subjects in the Santesso et al. (2008) study who had not learned the proper stimulus-reward association, receiving positive feedback was more unexpected than for those who already knew that this particular stimulus was the rich stimulus, eliciting a negative-going FRN deflection in the former but not in the latter subjects (see also Hajcak et al. 2007). Therefore, we would suggest that the FRN effects in the Santesso et al. (2008) study reflect a post hoc measure of whether the learning process has been successful, whereas our FRN effects reflect an online measure of whether learning will be successful. The same point can be made regarding other previous studies relating FRN amplitude to learning performance (e.g., Holroyd and Coles 2002; Frank et al. 2005; Bellebaum and Daum 2008).

Second, although previous studies have provided valuable insight into the role of the FRN in strategic performance adjustment, we argue that actual learning entails more than performance adjustment and requires the consolidation of contingencies, for example, between stimuli and their associated value or cues and their associated responses, or actions and their associated outcomes. Indeed, Kennerley et al. (2006) showed that lesions to the ACC (the putative source of the FRN) did not impair performance on trials immediately following negative feedback: Lesioned monkeys adjusted their performance after feedback equally well compared with control monkeys. What these authors found, however, was that these monkeys were less likely to repeat a response that had previously been rewarded, suggesting that the ACC may not be involved in simply using negative feedback to adjust performance but rather to develop a representation of the value of each response option. Therefore, the authors proposed that the function of the ACC may be to build or modify action–outcome contingencies to develop value estimates for the available options, which can be used to guide future optimal choice behavior (Kennerley et al. 2006). Indeed, Holroyd and Coles (2008) recently showed that ERN amplitude may reflect this ACC function of integrating the recent history of reinforcements, guiding future choice behavior. Therefore, to show that the FRN indeed reflects learning, one would have to show that the amplitude of this ERP component is associated with the building or consolidating of these contingencies. This evidence however has remained forthcoming, with most studies focusing on post-feedback performance adjustments.

Most notable in the present context is a recent study by Cohen and Ranganath (2007). In this experiment, subjects performed a “matching pennies” task, requiring them to select 1 of 2 possible responses: When subjects selected the same response as a computer opponent, they lost a point; when they selected the opposite response they earned a point. The results showed that FRN amplitude after receiving negative feedback (i.e., the subject chose the same response as the computer) predicted whether subjects would alter their response pattern for the following trial. Although this study provides important data on the involvement of FRN in post-feedback performance adjustments, the drawback of the design employed in this study is that performance is never contingent on particular cues or stimuli. So, this task does not require the learning of stimulus–response or action–outcome associations; feedback only serves to immediately change the response strategy and is no longer useful after this performance adjustment and so does not have to be consolidated. In contrast, in the present study subjects were required to learn (on the basis of the feedback provided) and remember which response should be associated with which particular stimulus (i.e., the item number of the present sequence). This allowed us to show that FRN amplitudes indeed reflect the process of learning and skill acquisition as predicted by the RL model of the FRN.

To conclude, our findings add to a growing literature on the neural correlates of learning (e.g., Frank et al. 2005; Cohen and Ranganath 2007; Klein et al. 2007), by showing that FRN amplitudes not only are related to performance adjustments but may also reflect the acquisition of motor skill and the consolidation of contingencies between stimuli or cues and their associated responses. In addition, we provide evidence that learning efficiency and future performance can be predicted by the neural response to current feedback: FRN

amplitude after we have made a mistake is predictive of whether a mistake will be repeated or whether we will learn from our mistake.

Notes

Conflict of Interest: None declared.

Address correspondence to Maarten A. S. Boksem, Department of Social Psychology, Tilburg University, Warandelaan 2, PO Box 90153, 5000 LE Tilburg, the Netherlands. Email: Maarten@Boksem.nl.

References

- Annett M. 1970. A classification of hand preference by association analysis. *Br J Psychol.* 61:303–321.
- Barto AG, Sutton RS. 1997. Reinforcement learning in artificial intelligence. In: Donahoe JW, Packard Dorsal V, editors. *Neural network models of cognition.* Amsterdam, The Netherlands: Elsevier. p. 358–386.
- Bellebaum C, Daum I. 2008. Learning-related changes in reward expectancy are reflected in the feedback-related negativity. *Eur J Neurosci.* 27:1823–1835.
- Boksem MAS, Meijman TF, Lorist MM. 2006. Mental fatigue, motivation and action monitoring. *Biol Psychol.* 72:123–132.
- Boksem MAS, Tops M, Wester AE, Meijman TF, Lorist MM. 2006. Error-related ERP components and individual differences in punishment and reward sensitivity. *Brain Res.* 1101:92–101.
- Cohen MX, Ranganath C. 2007. Reinforcement learning signals predict future decisions. *J Neurosci.* 27:371–378.
- Duncan-Johnson CC, Donchin E. 1977. On quantifying surprise: the variation of event-related potentials with subjective probability. *Psychophysiology.* 14:456–467.
- Frank MJ, Woroch BS, Curran T. 2005. Error-related negativity predicts reinforcement learning and conflict biases. *Neuron.* 47:495–501.
- Gehring WJ, Goss B, Coles MGH, Meyer DE, Donchin E. 1993. A neural system for error detection and compensation. *Psychol Sci.* 4: 385–390.
- Gehring WJ, Willoughby AR. 2002. The medial frontal cortex and the rapid processing of monetary gains and losses. *Science.* 295: 2279–2282.
- Gratton G, Coles MGH, Donchin E. 1983. A new method for off-line removal of ocular artifact. *Electroenceph Clin Neurophysiol.* 55:468–484.
- Hajcak G, Moser JS, Holroyd CB, Simons RF. 2007. It’s worse than you thought: the feedback negativity and violations of reward prediction in gambling tasks. *Psychophysiology.* 44:905–912.
- Herrmann CS, Knight RT. 2001. Mechanisms of human attention: event-related potentials and oscillations. *Neurosci Biobehav Rev.* 25: 465–476.
- Hester R, Barre N, Murphy K, Silk TJ, Mattingley JB. 2008. Human medial frontal cortex activity predicts learning from errors. *Cereb Cortex.* 18:1933–1940.
- Holroyd CB, Coles MGH. 2002. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev.* 109:679–709.
- Holroyd CB, Coles MGH. 2008. Dorsal anterior cingulate cortex integrates reinforcement history to guide voluntary behavior. *Cortex.* 44:548–559.
- Holroyd CB, Krigolson OE. 2007. Reward prediction error signals associated with a modified time estimation task. *Psychophysiology.* 44:913–917.
- Holroyd CB, Pakzad-Vaezi KL, Krigolson OE. 2008. The feedback correct-related positivity: sensitivity of the event-related brain potential to unexpected positive feedback. *Psychophysiology.* 45:688–697.
- Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF. 2006. Optimal decision making and the anterior cingulate cortex. *Nat Neurosci.* 9:940–947.
- Klein TA, Neumann J, Reuter M, Hennig J, von Cramon YD, Ullsperger M. 2007. Genetically determined differences in learning from errors. *Science.* 318:1642–1645.

- Luu P, Tucker DM, Derryberry D, Reed M, Poulsen C. 2003. Electrophysiological responses to errors and feedback in the process of action regulation. *Psychol Sci.* 14:47-53.
- Miltner WHR, Braun CH, Coles MGH. 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.
- Nieuwenhuis S, Holroyd CB, Mol N, Coles MGH. 2004. Reinforcement-related brain potentials from medial frontal cortex: origins and functional significance. *Neurosci Biobehav Rev.* 28: 441-448.
- Nieuwenhuis S, Schweizer TS, Mars RB, Botvinick MM, Hajcak G. 2007. Error-likelihood prediction in the medial frontal cortex: a critical evaluation. *Cereb Cortex.* 17:1570-1581.
- Nieuwenhuis S, Slagter HA, von Geusau NJA, Heslenfeld DJ, Holroyd CB. 2005. Knowing good from bad: differential activation of human cortical areas by positive and negative outcomes. *Eur J Neurosci.* 21:3161-3168.
- Oliveira FT, McDonald JJ, Goodman D. 2007. Performance monitoring in the anterior cingulate is not all error related: expectancy deviation and the representation of action-outcome associations. *J Cogn Neurosci.* 19:1994-2004.
- Pivek RT, Broughton RJ, Coppola R, Davidson RJ, Fox N, Nuwer MR. 1993. Guidelines for the recording and quantitative analysis of electroencephalographic activity in research contexts. *Psychophysiology.* 30:547-558.
- Ridderinkhof RK, Nieuwenhuis S, Bashore TR. 2003. Errors are foreshadowed in brain potentials associated with action monitoring in cingulate cortex in humans. *Neurosci Lett.* 348:1-4.
- Santesso DL, Dillon DG, Birk JL, Holmes AJ, Goetz E, Bogdan R, Pizzagalli DA. 2008. Individual differences in reinforcement learning: behavioral, electrophysiological, and neuroimaging correlates. *Neuroimage.* 42:807-816.
- Schultz W. 2002. Getting formal with dopamine and reward. *Neuron.* 36:241-263.
- Schultz W. 2004. Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. *Curr Opin Neurobiol.* 14:139-147.
- Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science.* 275:1593-1599.