I knew that! Confidence in outcome prediction and its impact on feedback processing and learning

Frömer, R.[1,2], Nassar, M.R.[2], Stürmer, B.[3], Sommer, W.[1], & Yeung, N.[4]

[1] Humboldt-Universität zu Berlin, [2] Brown University, [3] International Psychoanalytic University, [4] University of Oxford

Romy Frömer

Brown University

Cognitive, Linguistic and Psychological Sciences

190 Thayer Street Providence, RI, USA

romy_fromer@brown.edu

Abstract

Influential theories emphasize the importance of predictions in learning: We learn from response outcomes and feedback to the extent that they are surprising, and thus convey new information. Here we investigated how individuals learn to predict response outcomes based on the subjective confidence and objective accuracy with which these predictions are made. We hypothesized that both of these aspects modulate how feedback is processed and that they are reflected in event-related potentials (ERPs) as measured using EEG. Participants performed a time estimation task with graded, performance-contingent feedback. With this design we could distinguish reward prediction errors (RPE), indexing outcome valence with regard to the goal, and output prediction errors (OPE), indexing the absolute mismatch between predicted motor outcome and actual performance. As we expected, predictions made with higher confidence were more accurate (smaller OPE), and more so as learning progressed. Further, individuals with a better correspondence between confidence judgments and prediction accuracy learned more quickly. Outcome valence, as indexed by RPE was reflected in the feedback-related negativity (FRN). In contrast, P3a amplitude increased with OPE and confidence, that is with the degree of surprise about the outcome. Finally, performance-relevant information converged in the P3b component with confidence modulating RPE effects in early trials while learning took place. Taken together, the results underline the significance of different aspects of predictions and suggest a role of confidence in learning.

Keywords: confidence, reinforcement learning, graded feedback, motor learning

I knew that! Confidence in outcome prediction and its impact on feedback processing and learning

Feedback is an inherent feature of our natural environment and crucial to learning and adaptation. Humans use feedback to obtain and refine their skills, e.g., using feedback from vision to adjust movement, feedback about the consequences of decisions to correct them when outcomes are not as intended, or feedback from instructors and teachers to learn and correct rules underlying skills as diverse as high jumping and algebra. However, even simple forms of feedback have multiple dimensions. Feedback informs about the objective outcome, but can also be evaluated as good or bad (valence), more or less surprising, and more or less useful for learning. Across domains it is thought that feedback supports learning to the degree that it is unexpected, hence provides new information.

The degree to which feedback is expected depends on performance monitoring, the evaluation of one's own actions, supporting the detection of incorrect actions and anticipation of their consequences. Sometimes we cannot tell with certainty whether an action was correct or incorrect, but we can still express a degree of confidence in our choice that can enable effective regulation of the decision process both within and between individuals (Bahrami et al., 2010; Shea et al., 2014). Both of these forms of performance evaluation – error detection and confidence (Yeung & Summerfield, 2012) – require sufficient knowledge about the task at hand. However, little is currently known about how people learn to make these evaluations as they master new skills, and this is particularly true in the context of tasks involving continuous response parameters—as in motor control, e.g., when throwing at a target—where errors are inevitable and graded, as compared with the simpler case of a binary, categorical choice.

In the present study we used this framework of confidence and error detection as a form of outcome prediction to investigate how individuals learn to accurately evaluate responses along two key dimensions: 1) predicting action outcomes – i.e., how people learn to accurately predict the direction and magnitude of a potential error, and 2) accurately indicating confidence – i.e., how people learn to calibrate their confidence judgments such that higher confidence correlates with more accurate predictions. We further studied how these internal evaluations affect the encoding of feedback information as it is reflected in event-related potentials (ERPs) of the EEG.

Internal performance evaluation and its impact on the processing of external feedback

The discrepancy between expected and actual outcomes, prediction-errors, is widely accepted to underlie learning across domains. However, prediction errors are defined very differently across domains and reflect different evaluations of feedback. In the motor domain, outcomes of movements are physical states of the body or the environment (e.g., an object that is being manipulated, such as a ball landing at a certain location) and prediction errors reflect the discrepancy between the predicted and actual outcome of a motor command (Faisal, Selen, & Wolpert, 2008). These errors are used to learn response-outcome associations and abstract them into internal models underlying outcome prediction (forward models; Flanagan, Vetter, Johansson, & Wolpert, 2003; Wolpert, Diedrichsen, & Flanagan, 2011; Wolpert & Ghahramani, 2000) and movement production (Wolpert & Flanagan, 2001). The magnitude of the discrepancy between prediction and outcome (henceforth output prediction error; OPE) signals the quality of the forward model and the necessity to adapt it.

In contrast to the motor domain, in reinforcement learning (RL; Sutton & Barto, 1998), the prediction errors supporting learning are defined in terms of the *signed* difference between the reward value predicted and the value actually obtained (reward prediction errors, RPE; but see also: Behrens, Woolrich, Walton, & Rushworth, 2007; Diederen & Schultz, 2015; Diederen, Spencer, Vestergaard, Fletcher, & Schultz, 2016; McGuire, Nassar, Gold, & Kable, 2014; Pearce & Hall, 1980; Yu & Dayan, 2005). These RPEs can be used to update action value representations according to whether experienced outcomes are better or worse than expected (respectively making actions performed more or less likely to be repeated in the future). RPEs are therefore inherently valenced as good or bad (desirable or undesirable): A large negative prediction error has an entirely different meaning than a large positive prediction error (desirable ore undesirable, i.e. a large negative prediction error has an entirely different meaning than a large positive prediction error). RPEs therefore primarily provide information relevant to action selection (which action to choose) rather than action specification (how to perform a chosen action).

Thus, although motor learning and RL theories share an emphasis on prediction errors as the crucial driver of learning, the meanings of those prediction errors are very different. To illustrate the role of predictions and the difference between RPE and OPE, consider first the case of throwing a dart at a target. While initially aiming for the bullseye, one might miss it 20 cm to the right. If the prediction at the time of feedback was hitting the bullseye, as planned, OPE would equal the objective error magnitude, and RPE would be of the same magnitude but negative—both the forward model and action values would need to be revised in light of these worse-than-predicted outcomes. However, consider a case with the same objective outcome (missing 20 cm to the right) but a predicted outcome of missing by 20 cm to the left (e.g., because efference copy of the motor command indicated suboptimal execution). In this case the RPE is zero because the predicted outcome is exactly as favorable as the actual outcome, such that no update of action values should occur. However, OPE is twice the objective error magnitude, because the direction of error was predicted incorrectly, and a substantial update of the forward model should follow. As a final example, consider the case of expecting a thrown dart to miss 20 cm to the left, but finding it actually hits the target. Here the outcome is much better than expected (positive RPE), which through RL should increase the value of the chosen action, but the outcome was still incorrectly predicted (large OPE) and the underlying forward model must be updated to ensure successful performance in the future. These examples illustrate how feedback can multiplex different kinds of information—in relation to distinct types of prediction—that drive different forms of learning. To date, however, research has rarely studied and contrasted these distinct forms of prediction and learning as they relate to feedback processing. This was the first aim of the present study.

Confidence and Predictions

It seems intuitive that not all predictions are made with the same degree of confidence such that after throwing darts for 10 years one would be both more accurate and more certain in the prediction of a dart's impact location compared to a novice playing at a bar for the first time. In turn, the confidence in the prediction should affect how the actual impact of the darts is being processed (Körding & Wolpert, 2004; cf. Therrien, Wolpert, & Bastian, 2018, for the impact of noise on reinforcement learning; van den Berg et al., 2016; cf. Wu, Miyamoto, Castro, Ölveczky, & Smith, 2014, who propose a link between motor variability during learning and confidence in the internal model). Indeed, people are often robustly aware of their response errors (Maier, di Pellegrino, & Steinhauser, 2012; Maier, Yeung, & Steinhauser, 2011; Riesel, Weinberg, Endrass, Meyer, & Hajcak, 2013; Yeung, Botvinick, & Cohen, 2004), such that negative feedback can be entirely predicted and thus not drive learning. The present study therefore considered the role of confidence in predictions during learning. The formalization of confidence is an ongoing matter of debate, but it is commonly thought to reflect integration of evidence from multiple sources into a probabilistic estimate of performance accuracy (Pouget, Drugowitsch, & Kepecs, 2016; van den Berg et al., 2016). Recently, the interplay of confidence and error detection has received attention, particularly in the domain of perceptual decision making. For example, Boldt and Yeung (2015) demonstrated that confidence and error detection are closely related and share a neural correlate. Similarly, Murphy, Robertson, Harty, and O'Connell (2015) showed that the magnitude of that same neural correlate varied gradually for detected and undetected errors, and its latency predicted reaction times of error detection. Crucially, in both studies, participants received no trial-by-trial feedback, but relied entirely on internal performance evaluation; that is, their outcome predictions, which could be correct or incorrect. From this point of view, confidence can be considered as a property of outcome prediction (Meyniel, Schlunegger, & Dehaene, 2015; Nassar, Wilson, Heasly, & Gold, 2010; Vaghi et al., 2017), in line with the suggestion that confidence is a “second-order” inference process (Fleming & Daw, 2017). Recent research shows ‘Confidence weightings’ in learning transition probabilities (Meyniel & Dehaene, 2017; Meyniel et al., 2015) and confidence based regulation of information seeking/exploration in decision making (Boldt, Blundell, & De Martino, 2017; Desender, Boldt, & Yeung, 2018). Here we test the hypothesis that post-response confidence judgments regulate decision processes and learning, by supporting adaptive learning from feedback.

Neural correlates of feedback processing

Feedback disambiguates response outcomes and provides information about performance accuracy, RPE, OPE, as well as about the accuracy of confidence judgments. Therefore, all of the above should be encoded during feedback processing. Here we used time-resolved ERP methods to disentangle these multiplexed sources of information in performance feedback. We considered three distinct ERP components that are commonly used to investigate feedback processing.

The feedback-related negativity (FRN) is an error-sensitive ERP component with a fronto-central scalp distribution that peaks between 230 to 330 ms following feedback onset (Miltner, Braun, & Coles, 1997). It is commonly thought to index neural encoding of RPE (Holroyd & Coles, 2002): Its amplitude increases with the degree to which an outcome is worse than expected, and conversely decreases to the extent that outcomes are better than expected (Hajcak, Moser, Holroyd, & Simons, 2006; Holroyd, Hajcak, & Larsen, 2006; Holroyd, Nieuwenhuis, Yeung, & Cohen, 2003; Sambrook & Goslin, 2015; Walsh & Anderson, 2012).

Other aspects of predictions appear to be reflected in sub-components of the P3, a positive-going deflection between 250 and 500 ms following stimulus onset. The P3 consists of two topographically and functionally dissociable sub-components (Polich, 2007). The function of the frontocentral P3a can be summarized as signaling the recruitment of attention for action to motivationally relevant stimuli (Nieuwenhuis, De Geus, & Aston-Jones, 2011). It has been shown to increase with increasing processing demands (Frömer, Stürmer, & Sommer, 2016), with larger prediction errors in probabilistic learning tasks (Fischer & Ullsperger, 2013), higher goal relevance in a go/no-go task (Walentowska, Moors, Paul, & Pourtois, 2016), and with meta-memory mismatch (feedback about incorrect responses given with high confidence; Butterfield & Mangels, 2003). The parietally distributed P3b appears to scale with the degree that feedback is useful for future behavioral adaptation (Chase, Swainson, Durham, Benham, & Cools, 2011; Sailer, Fischmeister, & Bauer, 2010; Ullsperger, Fischer, Nigbur, & Endrass, 2014; Yeung & Sanfey, 2004) . It has been found to increase with feedback salience (reward magnitude irrespective of valence; Yeung & Sanfey, 2004), behavioral relevance (choice vs. no choice;

Yeung, Holroyd, & Cohen, 2005), with more negative going RPE (Fischer & Ullsperger, 2014; Ullsperger et al., 2014), but also with better outcomes in more complex tasks (Pfabigan, Zeiler, Lamm, & Sailer, 2014), and to predict subsequent behavioral adaptation (Chase et al., 2011; Fischer & Ullsperger, 2013).

The present study

The aim of the present study was to investigate the interplay of outcome predictions and concurrent confidence in motor learning and control, both behaviorally and in terms of their impact on feedback processing as reflected in ERP waveforms. Broadly, we expected improving performance evaluation as learning progresses, both in outcome prediction and corresponding confidence, as well as distinct neural correlates for RPE, OPE and confidence. To test these predictions, we used a time estimation task, which limits the degrees of freedom of the response, precludes concurrent visual feedback that might affect performance evaluation, and is well established for ERP analyses (e.g. Luft, Takase, & Bhattacharya, 2014; Miltner et al., 1997). In the task, participants pressed a button terminating an initially unknown fixed time interval that began with a tone stimulus. After each response, participants predicted the size and direction of their deviation from the correct time point and rated their confidence in this prediction. Finally, they received feedback about their actual performance. By contrasting objective feedback with participants' predictions, we computed an estimate of RPE—indexing feedback valence with respect to goal achievement to determine whether the outcome was better or worse than expected—and OPE, the valence-free absolute difference between predicted and actual outcome.

We expected improvements in both performance and outcome predictions, including a tightening relationship between OPE (prediction accuracy) and confidence over time. We further expected RPE to be reflected in FRN, with amplitudes increasing to the degree that outcomes are worse than expected. As OPE and confidence should affect surprise about the outcome we predicted it to relate to larger P3a amplitudes. Finally we expected that P3b amplitude reflects the convergence of performance-relevant, that is goal-related, information, i.e., RPE and error magnitude. Based on the literature reviewed above, we expected opposite-going effects of RPE and error magnitude. In line with previous findings, P3b should increase with more negative prediction errors, and for smaller error magnitudes, such that worse than predicted outcomes and better performance should result in larger amplitudes.

Method

Participants

The study included 40 participants (13 males) whose average age was 25.8 years (*SD* = 4.3) and whose mean handedness score (Oldfield, 1971) was 63.96 (*SD* = 52.09; i.e., most participants were right-handed). Participants gave informed consent to the experiment and were remunerated with course credits or 8 € per hour.

Task and procedure

Participants were seated at a table in front of a monitor. We used an adapted time-estimation task (Luft et al., 2014; Miltner et al., 1997), alternating with accuracy and confidence ratings. The task consisted of four parts on each trial, illustrated in Figure 1. After a fixation cross lasting for a random interval of 300-900 ms, a tone (600 Hz, 200 ms duration) was presented. Participants' task was to terminate an initially unknown target interval (of 1504 ms from tone onset), by pressing a response key with their left hand. We chose a supra-second duration to make the task sufficiently difficult (Luft et al., 2014). Following the response, a fixation cross was presented for 800 ms. Participants were then asked to rate the accuracy of the interval they had just produced on a scale presented at the center of the screen (*too short – too long*; ± 125 pixel, 3.15 ° visual angle) by moving an arrow to the place on the scale corresponding to the estimate using a mouse cursor with their right hand. Then, on a scale of the same size, participants rated their confidence in the previous estimate, i.e., the prediction (*very unconfident – very confident*). The confidence rating was followed by a blank screen for 800 ms. Finally, participants received feedback about their performance with a red square (0.25° visual angle) placed on a scale identical to the accuracy estimation scale but without any labels. The placement of the square on the scale visualized the interval produced, with undershoots shown to the left and overshoots on the right side of the center mark on the scale (indicating the correct estimate). Feedback was presented for only 150 ms to avoid eye movements. The trial ended with an inter-trial interval of 1500 ms.

The experiment comprised five blocks of 50 trials, with self-paced rest between blocks. We used presentation software (Neurobs.) for stimulus presentation, event and response logging. Visual stimuli were presented on a 4/3 17'' BenQ Monitor (resolution: 1280x1024, refresh rate: 60Hz) placed in 60 cm distance from the participant. A standard computer mouse and a button (customized, accuracy 2 ms, response latency 9 ms) were used for response registration.

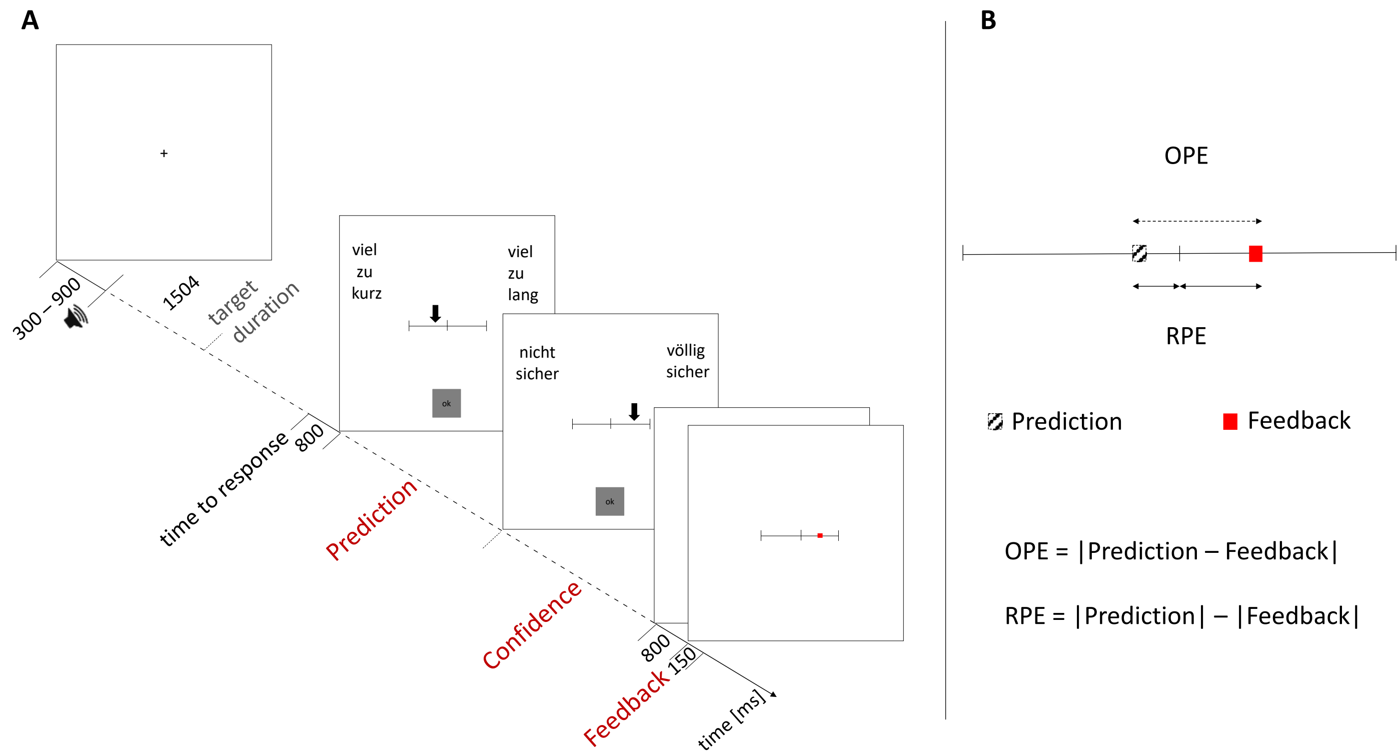

Figure 1. A Trial schema. Participants estimated a time interval by pressing a button following a tone. Subsequently, they used an arrow slider to rate their accuracy (scale: “viel zu kurz” = “much too short” to “viel zu lang” = “much too long”) then their confidence in that prediction (scale: “nicht sicher” = “not certain” to “völlig sicher” = “completely certain”). Finally, feedback was provided. B Prediction errors and their computation.

Prior to the experiment, participants filled in demographic and personality questionnaires (Neuroticism and Conscientiousness Scale of NEO PI-R; Costa & McCrae, 1992; and the BIS/BAS scale; Strobel, Beauducel, Debener, & Brocke, 2001) as well as a subset of the Raven’s progressive matrices as an index for figural-spatial intelligence (Raven, 2000). These measures were registered as potential control variables and for different projects not reported here. Participants were then seated in a shielded EEG cabin, where the experiment including EEG recording was conducted. Prior to the experiment proper, participants performed three practice trials.

Psychophysiological recording and processing

Using BrainVision recorder (Brain Products, München, Germany) with a sampling rate of 500 Hz, we recorded EEG data from 64 Ag/AgCl electrodes, mounted in an electrode cap (ECI Inc.), referenced against Cz. Electrodes below the eyes (IO1, IO2) and at the outer canthi (LO1, LO2) recorded vertical and horizontal ocular activity. We kept electrode impedance below 5 kΩ and applied a 100 Hz low pass filter, a time constant of 10 s and a 50 Hz notch filter. At the beginning of the session we recorded 20 trials of prototypical eye movements (up, down, left, right) for offline ocular artifact correction.

EEG data were processed using Matlab (The MathWorks Inc.) and the EEGlab toolbox (Delorme & Makeig, 2004). We re-referenced to average reference and retrieved the Cz channel. The data were band pass filtered between 0.5 and 40 Hz. Ocular artifacts were corrected using BESA (Ille, Berg, & Scherg, 2002). We segmented the ongoing EEG from -200 – 800 ms relative to feedback onset. Segments containing artifacts were excluded from analyses, based on values exceeding ± 150 µV and gradients larger than 50 µV between two adjacent sampling points. Baselines were corrected to the 200 ms pre-stimulus interval (feedback onset).

The FRN was quantified in single-trial ERP waveforms as a peak-to-peak amplitude at FCz, specifically as the difference between the minimum voltage between 200 and 300 ms post-feedback onset and the preceding positive maximum between -100 and 0 ms relative to the detected negative peak. To define the time windows for P3a and P3b single-trial amplitude analyses, we first determined the average peak latencies at FCz and Pz, respectively, and exported 100 ms time windows centered on the respective latencies. The P3a was quantified on single trials as the average voltage within an interval 330 – 430 ms post-feedback onset in a fronto-central region of interest (ROI: F1, Fz, F2, FC1, FCz, FC2, C1, Cz, C2). P3b amplitude was quantified in single trials as the average voltage within a 416 – 516 ms interval post-feedback in a parietally-focused region of interest (ROI: CP1, CPz, CP2, P1, Pz, P2, PO3, POz, PO4).

Analyses

Outlier inspection of the behavioral data identified one suspicious participant (average RT > 10 s) and single trials in four additional participants (RTs > 6 s). These data were excluded from further analyses. Two kinds of prediction errors were computed (Fig. 1). OPE (dashed arrow) was determined as the absolute difference between predicted and actual interval length: |Prediction – Feedback|. RPE was computed as the difference between the absolute predicted error and the absolute actual error as revealed by feedback: |Prediction| – |Feedback|. The example in Figure 1 depicts a negative RPE, as the actual outcome is worse than the predicted outcome. OPE is even larger, as not only the magnitude of the error was predicted incorrectly but also its direction: Whereas an undershoot was predicted, an overshoot was in fact produced.

Statistical analyses were performed by means of linear mixed models (LMMs) using R (R Core Team, 2014) and the lme4 package (Bates, Maechler, Bolker, & Walker, 2014a). We chose LMMs, similar to linear multiple regression models, as they allow for parametric analyses of single trial measures. Further, LMMs are robust to unequally distributed numbers of observations across participants, and simultaneously estimate fixed effects and random variance between participants in both intercepts and slope estimates. For all dependent variables, full models, including all predictors, were reduced step-wise until model comparisons indicated significantly decreased fit.

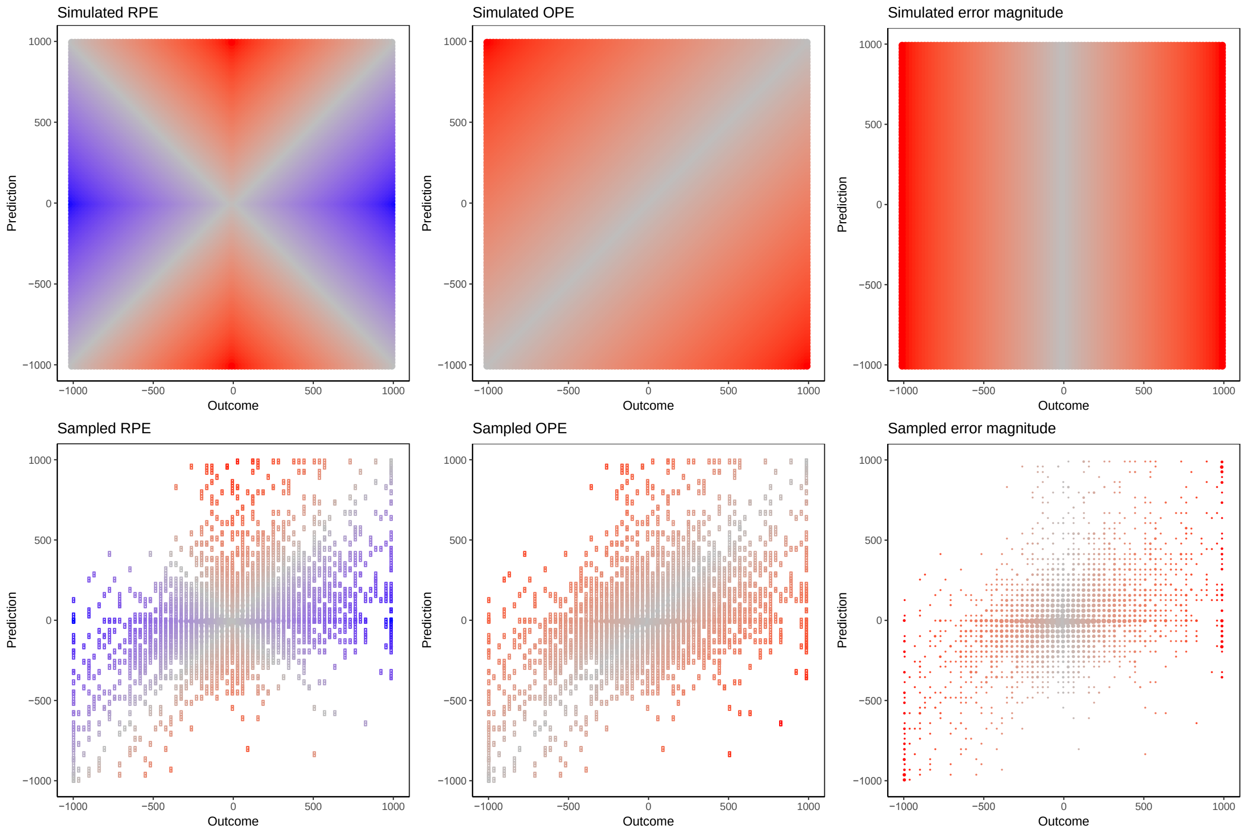

Figure 2. RPE, OPE and error magnitude can dissociate. Distribution of RPE, OPE and error magnitude across the outcome-prediction space in simulated (upper row) and sampled data (lower row). Grey indicates zero, while red indicates positive values and blue indicates negative values, note that RPE is signed, while OPE and Error Magnitude are unsigned.

RPE, OPE and Error Magnitude vary differentially for a given outcome and prediction, both theoretically, as well as in our observed data, such that they are correlated, but dissociable (Fig. 2). However, our subsample does not cover the entire range of these variables, which leads to substantial variance restriction with regard to interactions, which we therefore cannot robustly estimate and hence omit. To quantify the degree to which the variance restriction in our sample affects the interrelationships between the variables of interest, we computed RPE, OPE and Error magnitude on a simulated dataset spanning the entire range of possible predicted and actual outcomes. While independent in the simulated data, in the sampled data, where participants rarely misjudged large overshoots for large undershoots and vice versa, OPE and RPE were moderately negatively correlated, $r = -.44$; however, absolute RPE and OPE correlated highly with $r = .86$, compared to $r = .25$ in the simulated data. Thus, positive RPE was strongly positively correlated with OPE, whereas negative RPE was strongly negatively correlated with OPE. Also, both prediction errors were strongly correlated with performance (absolute error magnitude), OPE: $r = .65$ and RPE: $r = -.69$. The latter is consistent in magnitude with the correlation expected based on the simulated data ($r = -.71$), whereas the correlation between OPE and error magnitude was substantially smaller in the simulated data ($r = .31$), which includes the large misjudgments of error direction that we did not observe in our empirical data.

We report model comparisons and fit indices: Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC), which decrease with improving model fit. Random effect structures were determined using singular value decomposition. Variables explaining zero variance were removed from the random effects structure.

Prior to the analyses, error magnitude, RPE and OPE were scaled to seconds (by dividing them by 1000) and confidence and block were also scaled to a range of $\pm$ 1 for similar scaling of all predictors. Further, block, error magnitude, confidence, and OPE were centered on their medians for accurate intercept computation and to reduce collinearity. RPE was not centered, as zero represents a meaningful value on the scale (predicted and actual error magnitude are the same), and positive and negative values are qualitatively different (negative and positive values represent outcomes that are, respectively, worse or better than expected).

To establish that task performance improved with practice, single-trial error magnitude (absolute deviation from the target interval) was submitted to an LMM with block as a continuous predictor with a linear and a quadratic component. Subsequently, we tested the assumption that outcome prediction improves with learning and that the relationship is increasingly moderated by confidence. To this end, we analyzed signed error magnitude (negative = undershoot, positive = overshoot) with the predictors signed predicted error magnitude (negative = undershoot, positive = overshoot), block, and confidence.

To test directly whether accuracy of the prediction per se (OPE) relates to confidence, single-trial confidence ratings were submitted to LMMs with OPE, as well as block, and error magnitude as predictors. To investigate the functionality of accurate confidence judgments, we computed individual correlations between OPE and confidence, controlling for covariance with performance across blocks (confidence accuracy) and correlated this index with the individual slopes of block extracted from the simple learning model on error magnitude.

To assess the interplay of confidence and prediction errors on ERP markers of feedback processing, we analyzed FRN, P3a, and P3b amplitudes with RPE, OPE, confidence, block, and error magnitude as predictors.

Results

Behavioral results

Before testing the effects of confidence and prediction errors on learning and feedback processing, we demonstrate that a) learning takes place, such that performance improves over time, b) predictions relate meaningfully to behavior, and c) that confidence in predictions meaningfully relates to prediction accuracy, and the accuracy of predictions and corresponding confidence judgments increased with learning.

Performance improves across blocks. Performance (*absolute* error magnitude) significantly improved across blocks (Table 1). This effect is strongest in the beginning of the session with relative stability in performance thereafter (Fig. 3A, quadratic coef = 35.14, t = 3.88).

Table 1

Learning effects on performance

		Absolute error magnitude		
Variable	*b*	*SE*	*t*	*p*
Intercept	181.45	11.03	18.05	<.001
Block linear	-35.65	6.88	-5.18	<.001
Block quadratic	35.14	9.57	3.88	<.001

Variance Components	*SD*	*Goodness of fit*	
Participants	68.56	Log likelihood	-67318
Block linear	39.66	REML deviance	134636
Block quadratic	48.55		
Residuals	201.03		

Outcome predictions reflect actual performance and their accuracy increases over time and with confidence. A critical assumption of our study is that participants are aware of their errors and are able to predict the outcomes of their actions, rather than expecting correct performance on each trial. Our results confirm that participants were aware of their errors prior to receiving feedback, as their predictions broadly tracked the actual outcomes of their responses. Their predictions further improved over time and were more accurate when made with higher confidence. We regressed *signed* actual error (signed error magnitude, negative values are undershoots, positive are overshoots) on *signed* outcome predictions (predicted outcome, negative undershoots, positive overshoots). Simultaneously we tested how this relationship changes with learning, and whether confidence modulates the prediction-outcome-relationship by including block and confidence, as well as their interactions with predicted outcomes in the model. The results confirm an average signed error (Intercept) close to zero, indicating that performance varied around the correct time interval (Table 2). Increasing values of predicted outcomes were associated with increases in actual outcomes, indicating that participants could broadly indicate the direction of their errors. However, outcome prediction was far from perfect (Fig. 3B). The reliable interaction between block and predicted outcome (PO) reflects the fact that in Block 1, when little information is available to form a reasonable prediction, participants made many prediction errors in which large actual errors were predicted to be zero, resulting in a steep slope for predicted outcome on actual outcome, whereas these prediction errors became dramatically less frequent in subsequent blocks, when better predictions could be made. To test this interpretation, we ran a simple LMM, regressing OPE (the absolute difference between actual and predicted outcome) on block, which confirmed an overall decrease of prediction errors across blocks, $b = -18.82$, $t = -3.49$, $p = .001$, with significant improvement from block 1 to 2, $b = -42.54$, $p < .001$, and a numerical improvement in subsequent blocks that did not reach statistical significance, $bs = -11.22$, $ps > .05$. In line with our interpretation, participants on average underestimated error magnitude and OPE decreased across blocks (Fig. 3A). The consistent underestimation of errors is likely due to the fact that participants would bias their predictions towards zero, and particularly so when they were less certain about the direction of their predictions. Overall, undershoots were predicted more confidently than overshoots (signed actual errors decreased with increasing confidence), and this effect increased across blocks. Crucially, prediction accuracy (i.e., the slope of predicted outcome on actual outcome) increased with increasing confidence (Fig. 3B) and more so across blocks (Fig. 3C). Thus, with practice, confidence offered a better calibrated reflection of prediction accuracy.

Table 2

Learning effects on the relationship between predicted and actual outcomes and its modulation by confidence

	Signed error magnitude			
Variable	*b*	*SE*	*t*	*p*
Intercept	4.13	10.38	0.40	.693
PO	522.27	28.93	18.05	<.001
block	29.16	8.15	3.58	<.001
Conf	-24.55	10.98	-2.24	.031
PO: block	-143.69	22.10	-6.50	<.001
PO: Conf	308.00	36.15	8.52	<.001
Block: Conf	-23.84	9.18	-2.60	.009
PO: block: Conf	78.70	34.60	2.27	.023

Variance Components	*SD*	*Goodness of fit*	
Participants	62.12	Log likelihood	-68807
Confidence	52.42	REML deviance	137615
PO	141.87		
block	44.33		
Confidence: PO	132.86		
Residuals	233.10		

Note: PO = predicted outcome (in seconds, signed negative = undershoot, positive = overshoot), Conf = Confidence; ":" indicates interactions between predictors

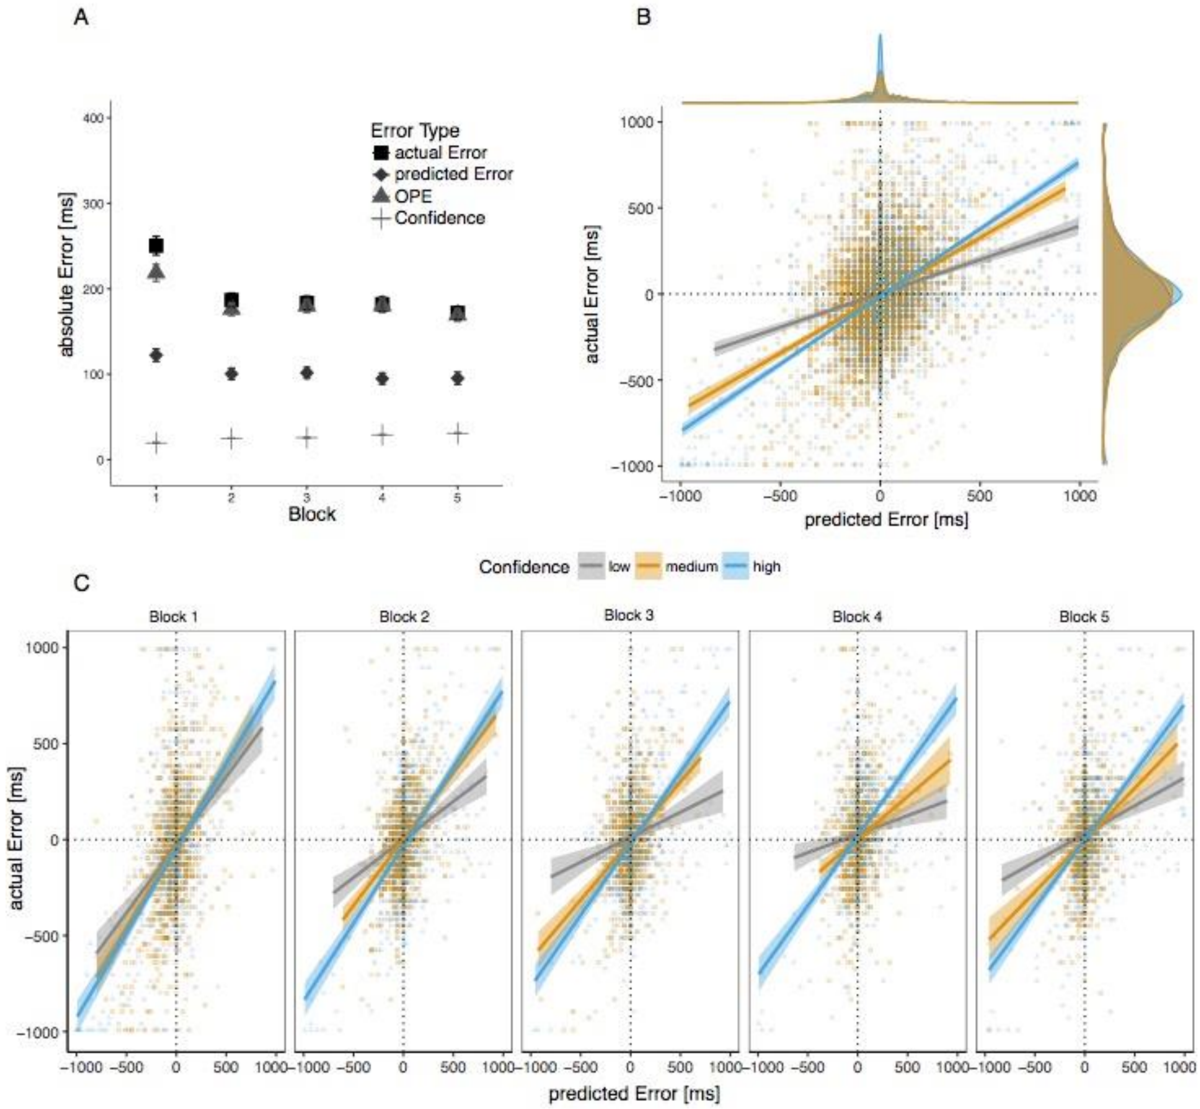

Figure 3. Relationship between outcome predictions and actual outcomes. A. Average predicted errors vs. average actual errors, OPE (in ms), and confidence (rescaled to -100 to 100 for visualization), across the experiment (note that each of the 5 blocks comprised 50 trials). B. Relationship between predicted and actual errors. Each data point corresponds to one trial of one participant; all trials of all participants are plotted together. Regression lines are local linear models visualizing the relationship between predicted and actual error separately for high, medium and low confidence. At the edges of the plot, the marginal distributions of actual and predicted errors are depicted by confidence levels. C. Change in the relationship between actual and predicted outcomes across blocks and for the three levels of confidence.

Confidence varies as a function of prediction accuracy, error magnitude and task experience. Having established that confidence scales prediction accuracy, our next analysis focused on the hypothesis that confidence primarily relates to prediction accuracy, but also increases with learning (i.e., experience with the task). Therefore, we tested whether confidence was predicted by OPE, as an indicator of prediction accuracy, and how it changed across blocks. As larger error magnitudes are likely more easily predicted (it might be easier to tell apart overshoots and undershoots when errors are large) and therefore may result in predictions made with higher confidence, we further modeled a quadratic component of error magnitude in addition to the linear component.

As a proof of concept of our design, participants were able to differentiate more accurate from less accurate predictions with their confidence ratings, such that confidence first and foremost significantly decreased with increasing OPEs (cf. Fig. 4A). Participants were further more confident in identifying larger errors and more so over time. That is, confidence in outcome predictions increased with *increasing* error magnitude with more extreme error magnitudes leading to stronger increases, and more so with learning (cf. Fig. 4B). Thus, our measure of confidence in the prediction did not increase with improving performance as decision confidence should, but was indeed specifically related to prediction accuracy. Participants further became more confident overall, such that confidence increased across blocks independent of error magnitude and OPE. LMM statistics and goodness of fit parameters are summarized in Table 3.

Table 3

Effects of learning, OPE and error magnitude on confidence

		Confidence		
Variable	*b*	*SE*	*t*	*p*
Intercept	7.39	1.24	5.98	<.001
OPE	-13.69	1.17	-11.69	<.001
EM linear	2.94	1.50	1.96	.053
EM quadratic	8.50	2.06	4.13	<.001
block	1.39	0.48	2.87	.006
EM quadratic: block	6.63	2.31	2.87	.004

Variance Components	*SD*	*Goodness of fit*	
Participants	7.85	Log likelihood	-37950
OPE	5.12	REML deviance	75901
EM	6.81		
block	2.93		
OPE:block	6.87		
EM:block	10.18		
Residuals	10.57		

Note: EM = absolute error magnitude; ":" indicates interactions between predictors

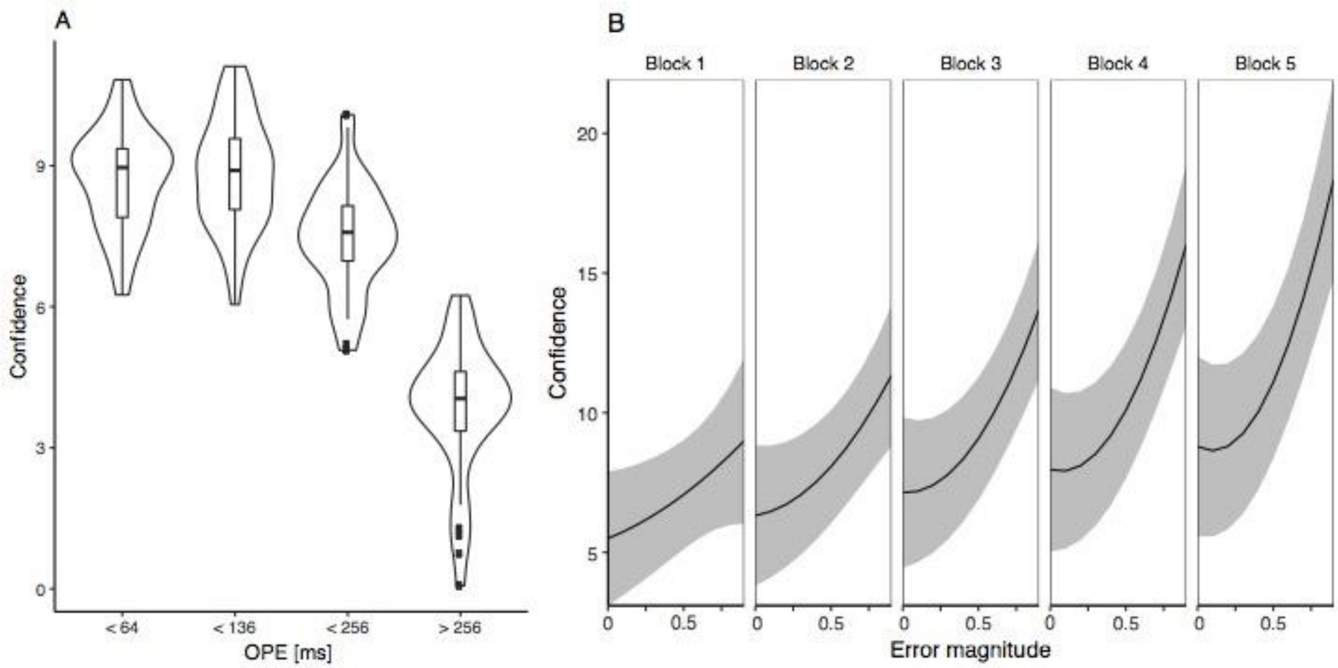

Figure 4: A. Relationship between OPE and Confidence. Confidence was averaged within OPE quartiles and participants. Violin plots illustrate the distribution of Confidence between participants, while bar plots show the medians, and 25th and 75th percentiles, respectively, while whiskers extend to the highest and lowest value, or 1.5 times the inter quartile range, in which case outliers are displayed as dots. For the computation of the OPE effect on Confidence, average error magnitude was assumed. B. Predicted Confidence as a function of error magnitude across blocks, where average OPE is assumed for the computation. Error bars are 95% confidence intervals around the predicted values.

In summary, these results provide a proof of concept of our task: participants learned, such that their error magnitudes decreased across blocks. Predicted outcomes meaningfully related to actual outcomes, and more so over time. The accuracy of those predictions was reflected in confidence judgments and increasingly so across blocks. Hence, learning improved accuracy at all levels: performance, outcome prediction and confidence judgments.

The accuracy of confidence judgments relates to individual differences learning

We hypothesized that metacognitive monitoring is important for learning, therefore we predicted that individuals with a better ability to *accurately* differentiate between correct and incorrect predictions with confidence (here termed confidence accuracy) would learn better. As these individuals have a better sense for the reliability of their predictions and therefore potentially a better credit assignment of their errors, they might know better when errors are caused by noisy execution, and when feedback is useful to update their representation of the correct time interval/response.

To test this hypothesis, we computed individual correlations of confidence and OPE across all trials as an estimate of overall confidence accuracy. To account for shared changes in our confidence accuracy measure with performance, we partialled out each participant's average error magnitude per block. We sign reversed the correlations, such that higher values correspond to higher confidence accuracy, to ease interpretation. All but two participants showed positive confidence accuracy values, meaning that more accurate predictions were associated with higher confidence judgments (Fig. 5). Figure 5A shows an emerging relationship between confidence accuracy and performance across individuals and blocks. Confidence accuracy did not depend on overall performance (Fig. 5B), supporting the assumption that confidence accuracy relates to learning, rather than performance. To directly test the relationship between confidence accuracy and learning, we extracted the linear block effects for each participant from the simple learning model (Table 1) and computed the correlation with confidence accuracy. As expected, confidence accuracy correlated significantly with the individual block estimates (average change in error across blocks), i.e., with learning ($r = -.37$, $p = .02$; Fig. 5C). Thus, participants with higher confidence accuracy showed greater performance improvements across the course of the experiment. Importantly, this confidence accuracy effect on learning is independent of overall performance—participants with higher confidence accuracy did not perform the task more or less accurately on average.

To summarize the behavioral results, participants improved their performance and their outcome predictions over time. Further, as learning progressed, confidence in the predictions was increasingly reflective of the accuracy of the predictions, and individual differences in confidence accuracy predicted individual differences in learning.

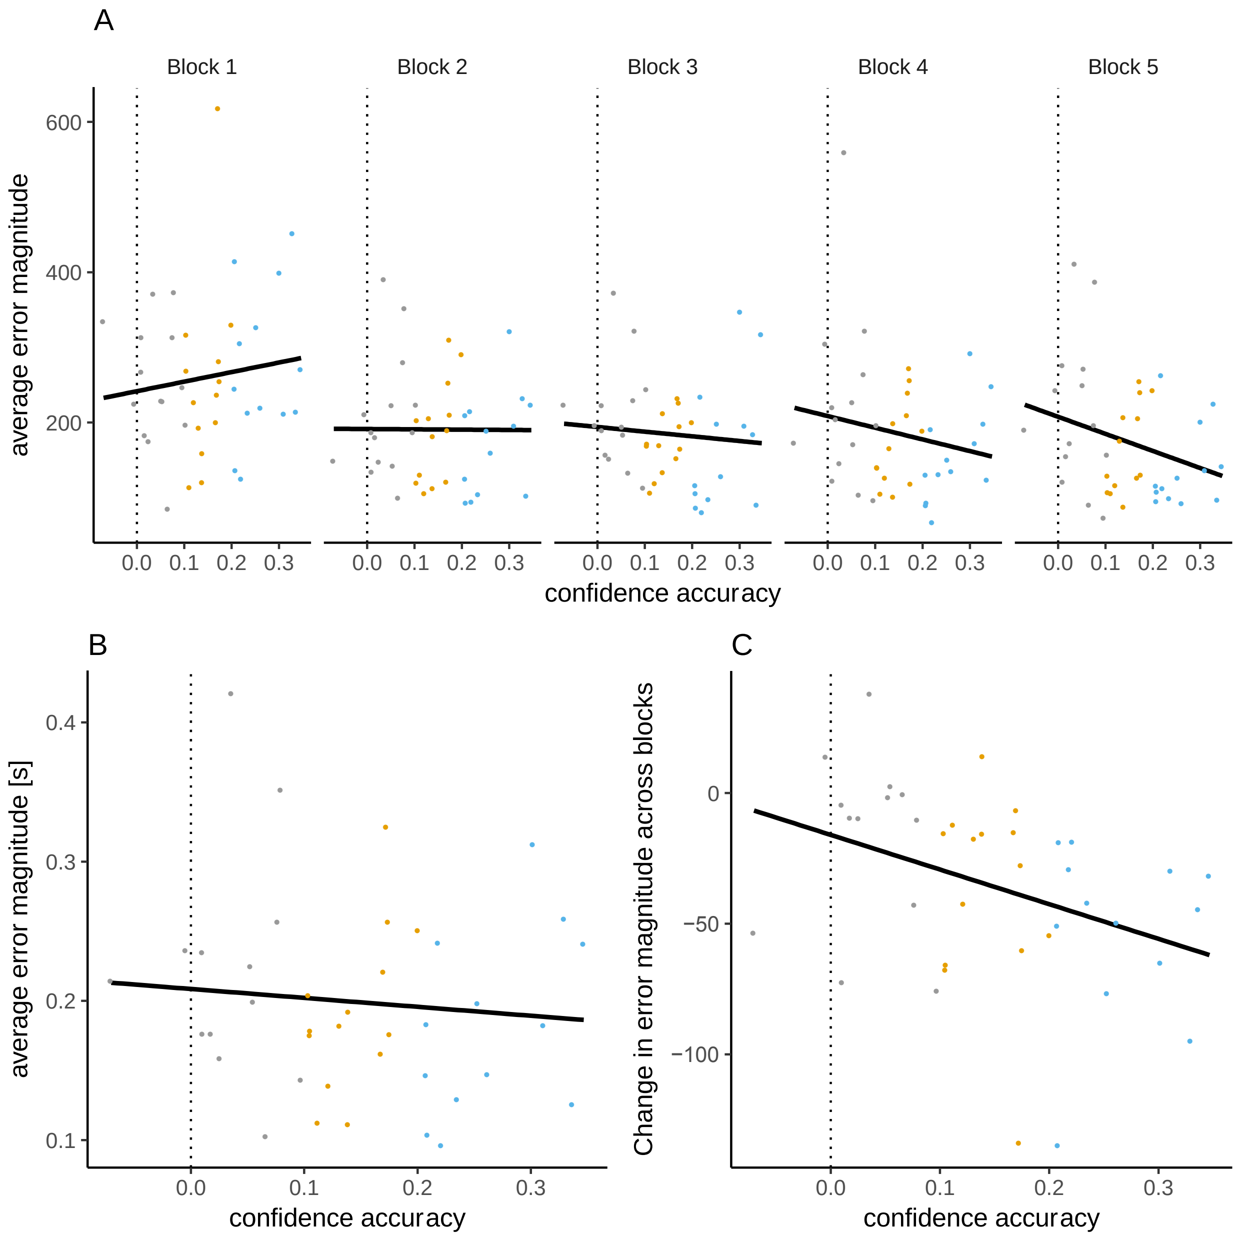

Figure 5. A. Observed performance averaged within blocks for each participant is plotted as a function of confidence accuracy (sign reversed correlation between Confidence and OPE across all trials controlling for improvement related changes across blocks) of that participant. B. Average error magnitude for each participant is plotted as a function of confidence accuracy of that participant. C. Correlation between confidence accuracy and learning across individuals. LMM-estimated change in error magnitude across blocks for each participant is plotted as a function of confidence accuracy of that participant. For all plots, confidence accuracy terciles are color coded (grey < orange < blue) for illustration.

EEG results

In these analyses we sought to demonstrate (a) that internal evaluations affect how feedback is processed, and (b) that different aspects of these internal evaluations would be reflected in distinct feedback-related ERP components. Specifically, we predicted that FRN amplitude would scale inversely with RPE (Holroyd & Coles, 2002; Sambrook & Goslin, 2015), and that in line with previously reported effects of surprise and metacognitive mismatch (Butterfield & Mangels, 2003; Fischer & Ullsperger, 2013), P3a amplitude would increase with confidence and OPE. Finally, we predicted convergence of information supporting successful future performance on P3b (Fischer & Ullsperger, 2013; Ullsperger et al., 2014). The full models for all ERP components included RPE, OPE and error magnitude, as well as main effects of Confidence and Block and their interactions.

FRN. As outlined above, we predicted FRN amplitude to scale with RPE, such that more negative RPEs would produce larger FRN amplitudes. However, we included OPE, Confidence, error magnitude and Block in the model to control for these factors, starting with a full model including interactions of RPE, OPE and Error magnitude with Confidence and Block. This full model on peak-to-peak FRN amplitude revealed no significant interactions among the predictors. Step-wise exclusion of the interaction terms did not lead to a significant drop in goodness of fit, and the final model did not fit significantly worse than the full model, $\Delta X^2(10) = 10.98$, $p = .359$, but showed better model fit indices (AIC: 59844 vs. 59853, BIC: 59937 vs. 60018). It is therefore justified to draw inferences from that reduced model, of which LMM statistics and goodness of fit parameters are summarized in Table 4.

Table 4

LMM statistics of learning effects on FRN

	Peak-to-Peak FRN amplitude			
Variable	*b*	*SE*	*t*	*p*
Intercept	-12.67	0.49	-26.03	<.001
Conf	-0.19	0.15	-1.25	.212
RPE	1.43	0.41	3.47	<.001
OPE	-0.67	0.42	-1.61	.108
EM	0.51	0.55	0.92	.357
block	-0.15	0.11	-1.43	.159

Variance Components	*SD*	*Goodness of fit*	
Participants	3.03	Log likelihood	-29909
block	0.47	REML deviance	59818
EM	1.49		
Residuals	5.26		

Note: EM = absolute error magnitude, Conf = Confidence

The significant negative Intercept term reflects the stability of the FRN [across participants/trials] as defined in a peak-to-peak measure. Importantly, consistent with our prediction and previous work (Holroyd & Coles, 2002; Sambrook & Goslin, 2015), the model revealed a significant main effect of RPE, with FRN amplitudes decreasing with more positive-going RPEs (Fig. 6B). None of the other predictors had significant effects. In particular, the objective outcome valence as indicated by error magnitude did not significantly affect FRN amplitude above and beyond RPE. Interestingly, the RPE effect was not further modulated by Confidence, indicating that feedback valence was processed largely independently of the confidence with which the prediction was made. One could have expected larger FRNs on trials with higher Confidence, or an interaction of Confidence with RPE, in the sense that RPE effects would be amplified for incorrect predictions made with high Confidence, but we did not observe evidence for that. We conclude that while other variables might potentially contribute as well, RPE based on subjective outcome predictions dominates variability in peak-to-peak FRN amplitude.

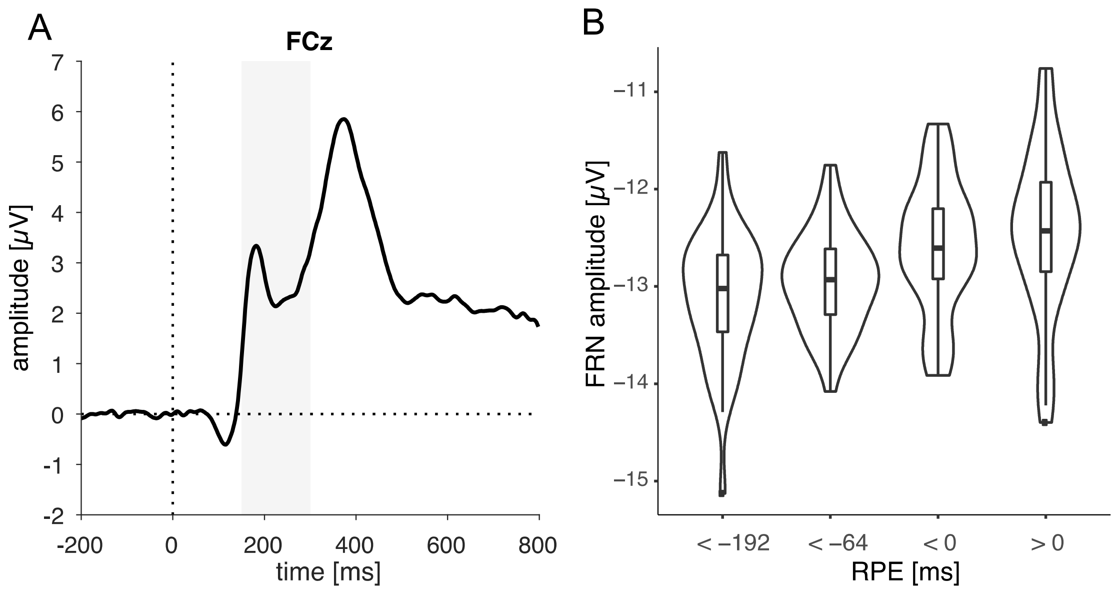

Figure 6. A. FRN, Grand mean, the shaded area marks the time interval for peak-to-peak detection of FRN. Negative peaks between 200 and 300 ms post feedback were subtracted from positive peaks in the preceding 100 ms time window. B. Partial effects of RPE on peak-to-peak FRN amplitude as estimated with the remef package. RPE is split into quartiles for visualization.

P3a. We predicted increasing P3a amplitude with increasing OPE and Confidence. As with the FRN analysis, we started out with a full LMM and stepwise excluded non-significant interaction terms, to base our inferences on the most parsimonious model explaining our data. This reduction did not diminish goodness of fit, $\Delta X^2(10) = 10.36$, $p = .410$, but both AIC and BIC were smaller for the reduced model (AIC: 58016 vs. 58026, BIC: 58088 vs. 58169). Table 5 summarizes this model's LMM statistics and goodness of fit indicators.

As predicted, P3a amplitude significantly increased with increasing OPEs, in line with stronger violations of expectations by less accurately predicted (i.e., more surprising) outcomes (Fig. 7B, C). Moreover, there was a significant increase of P3a amplitude with increasing confidence (Fig. 7B, D). Due to the presence of prediction errors in the vast majority of trials in our task, we interpret this as an amplification of the prediction error in high confidence trials. This empirically additive effect of surprise (OPE) and Confidence is further in line with theoretical analysis for a confidence-weighted updating mechanism reported by Meyniel and Dehaene (2017).

P3a amplitude decreased across blocks, likely because overall the feedback became less relevant as participants improved their performance and predictions (e.g. Walentowska et al., 2016 for effects of goal relevance on P3a). Further, we observed a significant P3a decrease with increasing error magnitude. However, as can be seen in Figure 7B, this error magnitude effect shows a more posterior distribution than those of OPE and confidence. The same is true for the significant effect of RPE on P3a amplitude with larger amplitudes for more negative prediction errors. As P3a temporally overlaps with more posteriorly distributed P3b, these effects are likely a spillover of the error magnitude effect on P3b described below. We will hence discuss them in that context.

Table 5

LMM statistics of learning effects on P3a

	P3a amplitude			
Variable	*b*	*SE*	*t*	*p*
Intercept	4.10	0.42	9.79	<.001
Conf	0.97	0.13	6.96	<.001
OPE	2.05	0.48	4.27	<.001
EM	-1.95	0.44	-4.43	<.001
Block	-0.91	0.07	-12.93	<.001
RPE	-0.74	0.38	-1.98	.048

Variance Components	*SD*	*Goodness of fit*	
Participants	2.61	Log likelihood	-28998
OPE	1.74	REML deviance	57996
Residuals	4.79		

Note: EM = absolute error magnitude, Conf = Confidence

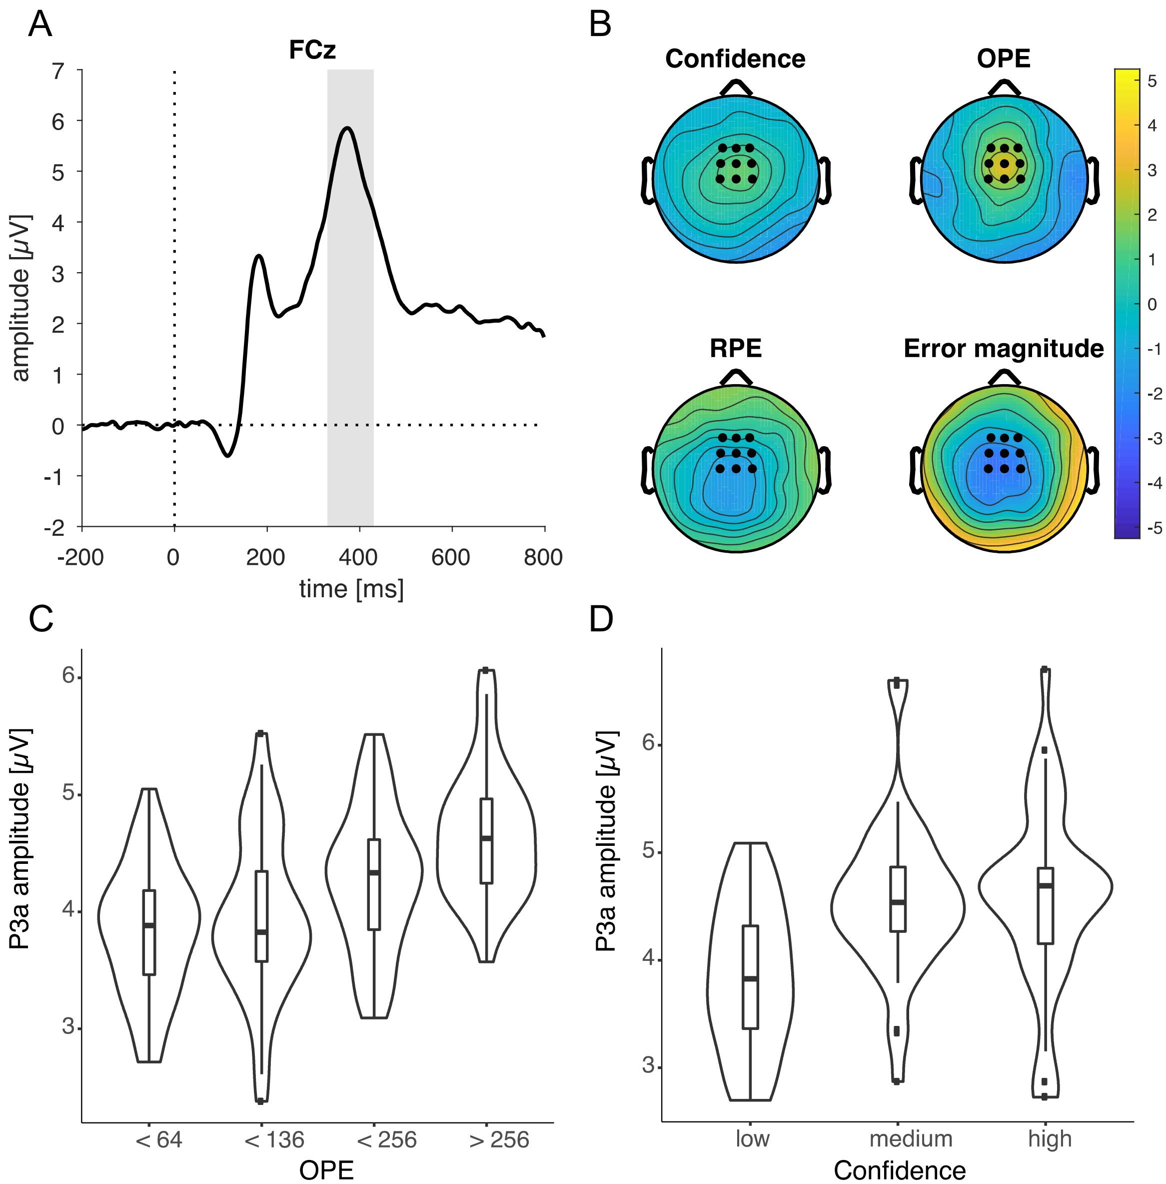

Figure 7. A. Grand mean ERP with the time-window for quantification of P3b highlighted. B. LMM effect distributions of Confidence, OPE, RPE and error magnitude. C. Partial effects of OPE on P3a amplitude as obtained with the remef package. OPE is split by quartiles. D. Partial effects of Confidence on P3a amplitude as obtained with the remef package. Confidence is summarized in terciles.

P3b. As P3b amplitude increases to the degree that feedback supports successful future performance, and based on previous literature, we predicted that P3b amplitude would increase with decreasing error magnitude (Pfabigan et al., 2014), supporting better performance, and to the degree that RPEs are more negative (Ullsperger et al., 2014), supporting predictions that reduce negative prediction errors in the future. We again started out with a full model including all interactions of Block and Confidence with RPE, Error magnitude and OPE. This full model was reduced, excluding non-significant interaction terms. The reduction did not significantly diminish goodness of fit, $\Delta X^2(5) = 10.443$, $p = .064$. AIC was identical for both models, while BIC favored the reduced model (AIC: 58440 vs. 58440, BIC: 58598 vs. 58633). Therefore, Table 6 summarizes LMM statistics and goodness of fit indicators of the more parsimonious reduced model.

In line with our predictions, we observed significant increases of P3b amplitude with decreasing error magnitude, thus, for better outcomes (Fig. 8B, C). P3b amplitude further increased with negative going RPE (Fig. 8B, D), hence, for worse than expected outcomes. However, RPE interacted significantly with confidence and block, indicating that the main effect needs to be interpreted with caution, and the relationship between P3b and RPE is more nuanced that we originally envisaged. In the first block, P3b increased with negative going RPE when confidence was high, but increased with more positive going RPE when confidence was low (Fig. 8F). Hence, negative RPEs appear to have a stronger impact under high confidence, while positive RPEs appear to have a stronger impact under low confidence, specifically early on in learning when little is known about the task yet. While the RPE effect levels off for higher confidence levels, it reverses for low confidence levels across blocks. This could be interpreted such that participants learned more in earlier blocks when they made large negative prediction errors with high confidence, and therefore expected to be accurate in their predictions, whereas in low confidence trials, when they were less certain, they learned more from outcomes that were better than predicted and were thus suitable to guide learning at the production level in the desired direction. In later trials, when higher confidence is associated with more accurate predictions and performance is usually accurate, the necessity to update in these trials may be reduced, whereas for low confidence trials worse than predicted outcomes may have a higher motivational salience and therefore lead to increased updating. In summary, as we predicted, we identified a negative effect of RPE on P3b amplitude, but this effect was only reliable following the first block, whereas in the first block, the effect was depended on Confidence as described above.

Beyond those effects, the model revealed increasing P3b amplitudes with increasing OPE (Fig. 8B, E), indicating that participants extracted more information to update when outcomes were more surprising. However the OPE effect decreased across blocks and so did overall P3b amplitude, suggesting that participants made less use of the feedback as they improved on the task. This finding is consistent with the idea that participants learn less from feedback over time as their performance saturates, and similar results have been reported by Fischer and Ullsperger (2013).

Thus, to summarize the ERP findings overall, FRN amplitude increased with more negative going RPE, P3a amplitude additively increased with OPE and Confidence, and P3b amplitude increased for smaller errors, and overall larger OPEs. Overall, P3b amplitude also increased with more negative RPEs, however, RPE effects were modulated by Confidence and Block, which might indicate adaptive feedback use over time.

Table 6

LMM statistics of learning effects on P3b

	P3b amplitude			
Variable	*b*	*SE*	*t*	*p*
Intercept	4.12	0.29	14.12	<.001
Conf	0.08	0.20	0.42	.677
OPE	1.75	0.46	3.76	<.001
EM	-2.35	0.46	-5.14	<.001
RPE	-1.12	0.46	-2.43	.018
Block	-0.48	0.09	-5.20	<.001
Conf: RPE	-0.51	0.55	-0.92	.356
Conf: Block	0.07	0.18	0.41	.682
RPE: Block	-0.52	0.44	-1.18	.236
OPE: Block	-0.98	0.46	-2.12	.034
Conf: RPE: Block	2.22	0.72	3.08	.002

Variance Components	*SD*	*Goodness of fit*	
Participants	1.78	Log likelihood	-29198
OPE	1.47	REML deviance	58396
RPE	1.29		
Conf	0.79		
Residuals	4.89		

Note: EM = absolute error magnitude, Conf = Confidence

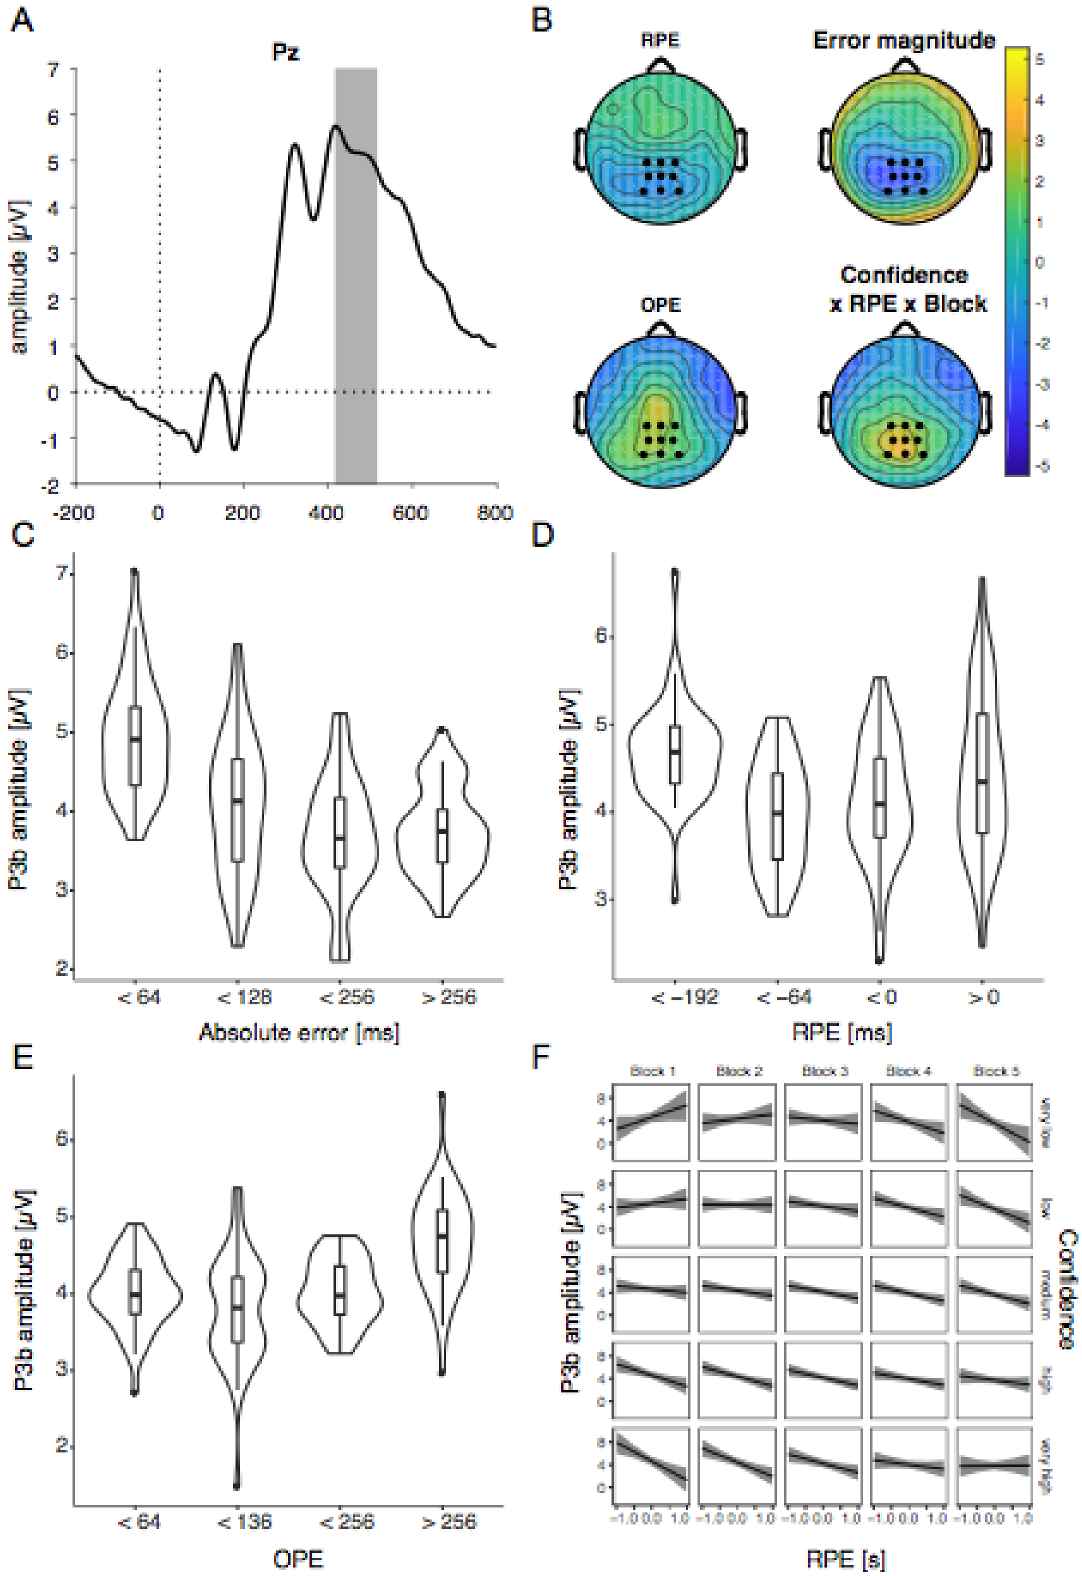
A
Pz
amplitude [μV]
B
RPE
Error magnitude
OPE
Confidence
x RPE x Block
C
P3b amplitude [μV]
< 64
< 128
< 256
> 256
Absolute error [ms]
D
P3b amplitude [μV]
< −192
< −64
< 0
> 0
RPE [ms]
E
P3b amplitude [μV]
< 64
< 136
< 256
> 256
OPE
F
Block 1
Block 2
Block 3
Block 4
Block 5
P3b amplitude [μV]
very low
low
medium
high
very high
Confidence
RPE [s]

Figure 8. A. Grand average ERP waveform at Pz with the time window for quantification highlighted. B. Effect topographies as predicted by LMMs for RPE, error magnitude, OPE and the RPE by Confidence by Block interaction. C. Partial effects of error magnitude on P3b amplitude visualized per error magnitude quartiles. D. Partial effects of RPE visualized by RPE quartiles. Note the interaction effects with Block and Confidence (F), that modulate the main effect. E. Partial effects of OPE visualized by OPE quartile. F. 3-way interaction of RPE, Confidence and Block as estimated by the LMM and visualized with the effects package for R.

Discussion

Here we explored the hypothesis that feedback is multiplexed and its processing is dependent on internal evaluations of performance as reflected in predictions and confidence. To test this hypothesis, we employed a time estimation task with continuous errors and feedback, allowing us to dissociate contributions of value-based and quality-based aspects of predictions to feedback processing, as indicated by RPE (valence; difference between predicted and actual error magnitude) and OPE (prediction accuracy; absolute deviation between predicted and actual outcome), respectively. We further tested the role of confidence in the predictions underlying these different prediction errors in learning and feedback processing. Our results show that as learning progresses, participants improved on all levels of motor control: execution, outcome prediction, and corresponding confidence judgments. They further show that participants' subjective predictions, as well as their confidence in those predictions, influenced feedback processing as revealed by feedback-related potentials, such that distinct aspects of feedback processing—RPE, OPE, and behavioral adaptation—were reflected in distinct ERP components.

Learning on all levels: Performance, Outcome Prediction and Confidence

As expected, performance (error magnitude) and outcome prediction (OPE) improved with practice, primarily across the first two blocks. Further, participants' confidence differentiated increasingly well between accurate and inaccurate predictions as learning progressed, such that over time participants were less confident in their outcome prediction when their predictions were likely wrong. This result dovetails with evidence from the perceptual domain, showing that predictions are employed to construct confidence (Sherman, Seth, & Kanai, 2016), and taken together these findings support the proposal that error monitoring and confidence are two sides of the same coin (Boldt & Yeung, 2015; Yeung & Summerfield, 2012).

Beyond accuracy of predictions, confidence was higher for larger errors (likely because participants found those errors easier to detect) and more so as learning progressed. These results conform with the notion of dependence between metacognitive sensitivity, the correspondence between confidence and performance, and performance per se (Fleming & Lau, 2014): Responses that were clearly worse than current average performance were more easily classified as such than responses representative for current ability. In line with that notion, average error in performance predictions (OPE) decreased to the level of average performance (error magnitude) across participants as learning progressed (cf. Fig. 3A). Alternatively the results could be accounted for by the recent finding that extreme values are stored more accurately in memory (Bays & Dowding, 2017). Indeed, the relationship between confidence and absolute error magnitude was not entirely linear but showed also a U-shaped component, with higher confidence for very large errors (cf. Fig. 4B). Thus, the subjective ease of detecting larger errors might be grounded in more stable representations of outcomes with larger error magnitude.

The functionality of accurate confidence representations is shown in its contribution to differences in learning across participants. Participants with overall more appropriate confidence judgments improved more than participants with lower confidence accuracy. This finding is even more striking given that overall confidence accuracy was not higher for participants with smaller average error magnitude or vice versa. Therefore the confidence accuracy effect on learning is unlikely to be merely an artifact of better overall ability as described in the “unskilled and unaware effect” (Kruger & Dunning, 1999) or the dependence of confidence accuracy (or metacognitive sensitivity) on performance (Fleming & Lau, 2014). As we can rule out this alternative explanation, we conclude that higher confidence accuracy supports learning, potentially via optimized feedback processing and credit assignment. Our results provide further evidence to the growing literature on the role of confidence in learning and behavioral adaptation (Boldt et al., 2017; Desender et al., 2018; Meyniel & Dehaene, 2017).

These analyses and conclusions were made possible by the use of a motor task with continuous errors and feedback. This is an advance on previous studies of feedback processing that have commonly used categorical responses and corresponding binary (correct/incorrect) feedback. More generally, by reinterpreting error detection as outcome prediction our results shed new light on the well-supported claim that error monitoring and confidence are tightly intertwined (Boldt & Yeung, 2015; Yeung & Summerfield, 2012) and forge valuable links between research on decision-making with categorical outcomes to findings from studies of motor control with continuous outcomes.

Evaluating Feedback: Neural Correlates of Errors, Prediction Errors and Confidence

We proposed that subjective predictions and corresponding confidence judgments would affect feedback processing. Therefore feedback-related potentials would reflect the different kinds of information decodable from feedback: valence reflected in RPE, surprise as indexed by OPE and weighted by confidence, and other performance-relevant information such as error magnitude. Our results suggest a sequential decoding of valence, surprise and task-relevant information: Fast decoding of outcome valence (RPE) was reflected in the FRN, followed by surprise related processing of the outcome (OPE and Confidence) reflected in P3a, and subsequent accumulation of task relevant information (RPE; OPE; and Error magnitude) reflected in P3b.

In line with our expectation, feedback valence, as indexed by RPE, was reflected in the FRN. FRN amplitude increased with the degree that outcomes were worse than expected. We did not observe any other effects on the FRN, notably including OPE or error magnitude effects. Thus, neither surprise nor objective outcome valence modulated FRN amplitude. Taken together, this pattern of findings is consistent with the reward prediction error account of the FRN (Holroyd & Coles, 2002; Sambrook & Goslin, 2015; Walsh & Anderson, 2012).

Further in line with our expectations, P3a amplitude increased with the absolute mismatch between outcome expectations and actual outcomes, reflected in OPE, and with increasing confidence, thus to the degree that the outcome was surprising. One might also have expected an interaction between OPE and confidence on P3a amplitude, with larger OPE effects for high compared to low confidence predictions. Instead we observed additive effects of OPE and confidence. However, our additive empirical findings are consistent with predicted near-additive effects of surprise and confidence in confidence-weighted updating (Meyniel & Dehaene, 2017). Taken together, our results align with the proposed role of P3a in mobilization of attention for action to motivationally relevant stimuli (Nieuwenhuis et al., 2011). This interpretation of the P3a also parsimoniously accounts for previous findings of metacognitive mismatch (Butterfield & Mangels, 2003). Errors committed with high confidence attract more attention, leading to more in-depth processing of feedback and therefore hypercorrection compared to low confidence errors (Butterfield & Metcalfe, 2001, 2006). Hence, inaccurately predicted outcomes and more so with higher confidence might signal the need for adaptation and therefore result in more in-depth processing of those outcomes.

This in-depth processing should hence temporarily follow the detection of task relevant information. The P3b component has been described as a candidate for such a process, indexing the accumulation or storage of task relevant information (O'Connell, Dockree, & Kelly, 2012; Sailer et al., 2010; Ullsperger et al., 2014). Fischer and Ullsperger (2013) further demonstrated that the P3b scales with the learning rate in a reinforcement learning paradigm and that it integrates information from factual and counterfactual feedback and predicts behavioral change, extending the literature implicating P3b in updating and supporting its functional role in the extraction of task relevant information (Chase et al., 2011; Ullsperger et al., 2014; Yeung & Sanfey, 2004). Indeed, in our data, this component increased for less accurate predictions (OPE and RPE), specifically negative going RPE, and more accurate performance (decreasing error magnitude), that is to the degree that feedback provided information relevant to optimizing future performance. The effects of OPE and RPE support our hypothesis that these are used to update the forward model, supporting accurate future predictions, and specifically avoiding negative RPEs. In our task, the avoidance of negative reward prediction errors does not support performance per se because positive RPEs are equally indicative of an inaccurate internal model and both can vary across a range of error magnitudes, but it changes subjective experience of the task, reducing the aversiveness of feedback and can therefore be considered as adaptive. However, an alternative, but not mutually exclusive account is that given its aversive nature, worse than expected feedback is more salient (Yeung & Sanfey, 2004) and therefore processed preferentially. Unexpectedly, confidence and block modulated RPE effects. This finding suggests that when participants knew little about the task and their confidence was low, feedback that was better than expected was processed preferentially. This observation again hints to a role of confidence in learning via adaptive feedback use.

In addition to prediction error effects on the P3b, we observed increasing P3b amplitude for more accurate performance, replicating previous findings (Ernst & Steinhauser, 2018; Pfabigan et al., 2014). In contrast to previous studies, correct and error feedback were equally informative in our task, yet our findings suggest that positive outcomes were more useful for adaptation. This result perhaps reflects the fact that successful parameters (or responses) can be promoted by reinforcement mechanisms, whereas negative feedback, even if it is technically equally informative, informs about the problem, but not the solution, which remains to be derived from the information provided. Thus, in contexts where response parameters need to be learned, positive performance feedback is more easily deployed for learning. This idea is not new, but favorable effects of feedback following successful trials have been reported previously in the motor domain (Chiviacowsky & Wulf, 2002, 2005, 2007). Here we provide further evidence for the assumption of preferential processing of positive performance feedback, with a signature in a neural correlate of performance-relevant information processing. It is important to point out that of course positive performance feedback can only support future performance when there is still ambiguity about the underlying performance, that is it should not be relevant to learning in the absence of prediction errors.

In summary, our ERP findings demonstrate that participants' trial-by-trial predictions were indeed used to inform feedback processing, and that different aspects of outcomes affected by these predictions (valence, surprise and adaptive information) were sequentially decoded. Crucially, confidence in the predictions affected feedback processing, in a manner suggestive of adaptive use of feedback in support of learning.

Implications and Future Directions

Our findings demonstrate the relevance of internal evaluations—predictions and confidence—for feedback processing, and further show that these evaluations guide the interpretation of multiple dimensions of feedback (valence, surprise, information for adaptation). These 'hidden' representations may not be the only internal representations that affect feedback processing. For example, it has recently been shown that subjective beliefs about the reliability of feedback affect feedback processing irrespective of objective feedback reliability (Schiffer, Siletti, Waszak, & Yeung, 2016). Across domains, uncertainty is assumed to weight sensory information from feedback against predictions derived from internal models, i.e., the forward model in the motor domain (Franklin & Wolpert, 2011; Tan, Wade, & Brown, 2016). An interesting future direction is thus how individuals weigh internal and external information as a function of their subjective reliability to support optimal feedback processing via credit assignment. Another exciting avenue to explore is how the emphasis of certain aspects of feedback (e.g., valence vs. usefulness for adaptation) alters sub-components of feedback processing and whether such alterations in feedback evaluation are systematically linked to individual differences and prevalent processing biases in clinical populations.

Conclusion

To summarize, we demonstrated that feedback processing is rich and complex, being sensitive to multiple external informational aspects (RPE, OPE, Error Magnitude) but also to internal evaluations (outcome predictions and confidence). While learning, individuals use their subjective outcome predictions and confidence judgments to modulate how they process feedback. Based on feedback, individuals improve these outcome predictions and corresponding confidence weightings with learning, along with their behavioral performance. Based on EEG results, we demonstrated that different aspects of outcomes relative to these predictions (valence, surprise, task relevant information) are sequentially decoded from feedback, supporting the notion that feedback processing is an active constructive process that is fundamentally affected by an individual's internal state at the time of feedback. Importantly, the accuracy of subjective confidence in the predictions affects how well people learn, suggesting a crucial role of confidence in credit assignment and optimal feedback use.

Acknowledgments

We thank Lena Fliedner and Lara Montau for support in data acquisition and helpful discussions during the setup of the task, Rainer Kniesche for advice on programming the stimulus government, and Rasmus Bruckner, Markus Ullsperger, Martin Maier, and Rasha Abdel Rahman for valuable discussion. RF is further grateful for the continuous scientific and personal support by her office mates Benthe Kornrumpf and Florian Niefind at Humboldt-University, who made her life and work a lot more fun and happened to also have inspired the title of this paper.

References

Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science, 329*(5995), 1081-1085. doi:10.1126/science.1185718

Bates, D., Maechler, M., Bolker, B. M., & Walker, S. C. (2014a). lme4: Linear mixed-effects models using Eigen and S4. R package version 1.1-8. Retrieved from http://lme4.r-forge.r-project.org/

Bays, P. M., & Dowding, B. A. (2017). Fidelity of the representation of value in decision-making. *PLoS computational biology, 13*(3), e1005405. doi:10.1371/journal.pcbi.1005405

Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nat Neurosci, 10*(9), 1214-1221. doi:10.1038/nn1954

Boldt, A., Blundell, C., & De Martino, B. (2017). Confidence modulates exploration and exploitation in value-based learning. *bioRxiv*. doi:10.1101/236026

Boldt, A., & Yeung, N. (2015). Shared neural markers of decision confidence and error detection. *J Neurosci, 35*(8), 3478-3484. doi:10.1523/JNEUROSCI.0797-14.2015

Butterfield, B., & Mangels, J. A. (2003). Neural correlates of error detection and correction in a semantic retrieval task. *Brain Res Cogn Brain Res, 17*(3), 793-817. doi:http://dx.doi.org/10.1016/S0926-6410(03)00203-9

Butterfield, B., & Metcalfe, J. (2001). Errors committed with high confidence are hypercorrected. *J Exp Psychol Learn Mem Cogn, 27*(6), 1491-1494. doi:10.1037/0278-7393.27.6.1491

Butterfield, B., & Metcalfe, J. (2006). The correction of errors committed with high confidence. *Metacognition and Learning, 1*(1), 69-84. doi:10.1007/s11409-006-6894-z

Chase, H. W., Swainson, R., Durham, L., Benham, L., & Cools, R. (2011). Feedback-related negativity codes prediction error but not behavioral adjustment during probabilistic reversal learning. *J Cogn Neurosci, 23*(4), 936-946. doi:10.1162/jocn.2010.21456

Chiviacowsky, S., & Wulf, G. (2002). Self-controlled feedback: Does it enhance learning because performers get feedback when they need it? *Research Quarterly for Exercise and Sport, 73*(4), 408-415.

Chiviacowsky, S., & Wulf, G. (2005). Self-controlled feedback is effective if it is based on the learner's performance. *Res Q Exerc Sport, 76*(1), 42-48. doi:10.1080/02701367.2005.10599260

Chiviacowsky, S., & Wulf, G. (2007). Feedback after good trials enhances learning. *Res Q Exerc Sport, 78*(2), 40-47. doi:10.1080/02701367.2007.10599402

Costa, P. T., & McCrae, R. R. (1992). *Revised NEO personality inventory (NEO-PI-R) and NEO Five-Factor inventory (NEO-FFI): Professional Manual*. Odessa: Psychological Assessment Resources.

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*(1), 9-21. doi:10.1016/j.jneumeth.2003.10.009

Desender, K., Boldt, A., & Yeung, N. (2018). Subjective Confidence Predicts Information Seeking in Decision Making. *Psychol Sci, 0*(0), 956797617744771. doi:10.1177/0956797617744771

Diederen, K. M., & Schultz, W. (2015). Scaling prediction errors to reward variability benefits error-driven learning in humans. *J Neurophysiol, 114*(3), 1628-1640. doi:10.1152/jn.00483.2015

Diederen, K. M., Spencer, T., Vestergaard, M. D., Fletcher, P. C., & Schultz, W. (2016). Adaptive Prediction Error Coding in the Human Midbrain and Striatum Facilitates Behavioral Adaptation and Learning Efficiency. *Neuron, 90*(5), 1127-1138. doi:10.1016/j.neuron.2016.04.019

Ernst, B., & Steinhauser, M. (2018). Effects of feedback reliability on feedback-related brain activity: A feedback valuation account. *Cogn Affect Behav Neurosci, 18*(3), 596-608. doi:10.3758/s13415-018-0591-7

Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nat Rev Neurosci, 9*(4), 292-303. doi:10.1038/nrn2258

Fischer, A. G., & Ullsperger, M. (2013). Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron, 79*(6), 1243-1255. doi:10.1016/j.neuron.2013.07.006

Fischer, A. G., & Ullsperger, M. (2014). When is the time for a change? Decomposing dynamic learning rates. *Neuron, 84*(4), 662-664. doi:10.1016/j.neuron.2014.10.050

Flanagan, J. R., Vetter, P., Johansson, R. S., & Wolpert, D. M. (2003). Prediction precedes control in motor learning. *Curr Biol, 13*(2), 146-150. doi:http://dx.doi.org/10.1016/S0960-9822(03)00007-1

Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychol Rev, 124*(1), 91-114. doi:10.1037/rev0000045

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Front Hum Neurosci, 8*, 443. doi:10.3389/fnhum.2014.00443

Franklin, D. W., & Wolpert, D. M. (2011). Computational mechanisms of sensorimotor control. *Neuron, 72*(3), 425-442. doi:10.1016/j.neuron.2011.10.006

Frömer, R., Stürmer, B., & Sommer, W. (2016). (Don't) Mind the effort: Effects of contextual interference on ERP indicators of motor preparation. *Psychophysiology, 53*(10), 1577-1586. doi:10.1111/psyp.12703

Hajcak, G., Moser, J. S., Holroyd, C. B., & Simons, R. F. (2006). The feedback-related negativity reflects the binary evaluation of good versus bad outcomes. *Biol Psychol, 71*(2), 148-154. doi:10.1016/j.biopsycho.2005.04.001

Holroyd, C. B., & Coles, M. G. (2002). The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol Rev, 109*(4), 679-709. doi:10.1037/0033-295X.109.4.679

Holroyd, C. B., Hajcak, G., & Larsen, J. T. (2006). The good, the bad and the neutral: electrophysiological responses to feedback stimuli. *Brain Res, 1105*(1), 93-101. doi:10.1016/j.brainres.2005.12.015

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*(18), 2481-2484. doi:10.1097/01.wnr.0000099601.41403.a5

Ille, N., Berg, P., & Scherg, M. (2002). Artifact correction of the ongoing EEG using spatial filters based on artifact and brain signal topographies. *Journal of Clinical Neurophysiology, 19*(2), 113-124. doi:10.1097/00004691-200203000-00002

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature, 427*, 244. doi:10.1038/nature02169

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments. *J Pers Soc Psychol, 77*(6), 1121-1134.

Luft, C. D., Takase, E., & Bhattacharya, J. (2014). Processing graded feedback: electrophysiological correlates of learning from small and large errors. *J Cogn Neurosci, 26*(5), 1180-1193. doi:10.1162/jocn_a_00543

Maier, M. E., di Pellegrino, G., & Steinhauser, M. (2012). Enhanced error-related negativity on flanker errors: error expectancy or error significance? *Psychophysiology, 49*(7), 899-908. doi:10.1111/j.1469-8986.2012.01373.x

Maier, M. E., Yeung, N., & Steinhauser, M. (2011). Error-related brain activity and adjustments of selective attention following errors. *NeuroImage, 56*(4), 2339-2347. doi:10.1016/j.neuroimage.2011.03.083

McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally dissociable influences on learning rate in a dynamic environment. *Neuron, 84*(4), 870-881. doi:10.1016/j.neuron.2014.10.013

Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proc Natl Acad Sci U S A, 114*(19), E3859-E3868. doi:10.1073/pnas.1615773114

Meyniel, F., Schlunegger, D., & Dehaene, S. (2015). The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLoS Comput Biol, 11*(6), e1004305. doi:10.1371/journal.pcbi.1004305

Miltner, W. H., Braun, C. H., & Coles, M. G. (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*(6), 788-798. doi:10.1162/jocn.1997.9.6.788

Murphy, P. R., Robertson, I. H., Harty, S., & O'Connell, R. G. (2015). Neural evidence accumulation persists after choice to inform metacognitive judgments. *eLife, 4*. doi:10.7554/eLife.11946

Nassar, M. R., Wilson, R. C., Heasly, B., & Gold, J. I. (2010). An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *The Journal of Neuroscience, 30*(37), 12366-12378. doi:10.1523/jneurosci.0822-10.2010

Nieuwenhuis, S., De Geus, E. J., & Aston-Jones, G. (2011). The anatomical and functional relationship between the P3 and autonomic components of the orienting response. *Psychophysiology, 48*(2), 162-175. doi:10.1111/j.1469-8986.2010.01057.x

O'Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat Neurosci, 15*(12), 1729-+. doi:10.1038/nn.3248

Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia, 9*(1), 97-113.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87*(6), 532-552. doi:10.1037/0033-295x.87.6.532

Pfabigan, D. M., Zeiler, M., Lamm, C., & Sailer, U. (2014). Blocked versus randomized presentation modes differentially modulate feedback-related negativity and P3b amplitudes. *Clin Neurophysiol, 125*(4), 715-726. doi:10.1016/j.clinph.2013.09.029

Polich, J. (2007). Updating P300: an integrative theory of P3a and P3b. *Clin Neurophysiol, 118*(10), 2128-2148. doi:10.1016/j.clinph.2007.04.019

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nat Neurosci, 19*(3), 366-374. doi:10.1038/nn.4240

R Core Team. (2014). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org/

Raven, J. (2000). The Raven's Progressive Matrices: Change and Stability over Culture and Time. *Cognitive Psychology, 41*(1), 1-48. doi:10.1006/cogp.1999.0735

Riesel, A., Weinberg, A., Endrass, T., Meyer, A., & Hajcak, G. (2013). The ERN is the ERN is the ERN? Convergent validity of error-related brain activity across different tasks. *Biol Psychol, 93*(3), 377-385. doi:10.1016/j.biopsycho.2013.04.007

Sailer, U., Fischmeister, F. P., & Bauer, H. (2010). Effects of learning on feedback-related brain potentials in a decision-making task. *Brain Res, 1342*, 85-93. doi:10.1016/j.brainres.2010.04.051

Sambrook, T. D., & Goslin, J. (2015). A neural reward prediction error revealed by a meta-analysis of ERPs using great grand averages. *Psychol Bull, 141*(1), 213-235. doi:10.1037/bul0000006

Schiffer, A. M., Siletti, K., Waszak, F., & Yeung, N. (2016). Adaptive behaviour and feedback processing integrate experience and instruction in reinforcement learning. *NeuroImage*. doi:10.1016/j.neuroimage.2016.08.057

Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., & Frith, C. D. (2014). Supra-personal cognitive control and metacognition. *Trends Cogn Sci, 18*(4), 186-193. doi:10.1016/j.tics.2014.01.006

Sherman, M. T., Seth, A. K., & Kanai, R. (2016). Predictions Shape Confidence in Right Inferior Frontal Gyrus. *J Neurosci, 36*(40), 10323-10336. doi:10.1523/JNEUROSCI.1092-16.2016

Strobel, A., Beauducel, A., Debener, S., & Brocke, B. (2001). Eine deutschsprachige Version des BIS/BAS-Fragebogens von Carver und White. *Zeitschrift für Differentielle und Diagnostische Psychologie, 22*(3), 216-227. doi:10.1024//0170-1789.22.3.216

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction* (Vol. 1). Cambridge: MIT press.

Tan, H., Wade, C., & Brown, P. (2016). Post-Movement Beta Activity in Sensorimotor Cortex Indexes Confidence in the Estimations from Internal Models. *J Neurosci, 36*(5), 1516-1528. doi:10.1523/JNEUROSCI.3204-15.2016

Therrien, A. S., Wolpert, D. M., & Bastian, A. J. (2018). Increasing Motor Noise Impairs Reinforcement Learning in Healthy Individuals. *eneuro*. doi:10.1523/eneuro.0050-18.2018

Ullsperger, M., Fischer, A. G., Nigbur, R., & Endrass, T. (2014). Neural mechanisms and temporal dynamics of performance monitoring. *Trends Cogn Sci, 18*(5), 259-267. doi:10.1016/j.tics.2014.02.009

Vaghi, M. M., Luyckx, F., Sule, A., Fineberg, N. A., Robbins, T. W., & De Martino, B. (2017). Compulsivity Reveals a Novel Dissociation between Action and Confidence. *Neuron, 96*(2), 348-354 e344. doi:10.1016/j.neuron.2017.09.006

van den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D. M. (2016). A common mechanism underlies changes of mind about decisions and confidence. *eLife, 5*, e12192. doi:10.7554/eLife.12192

Walentowska, W., Moors, A., Paul, K., & Pourtois, G. (2016). Goal relevance influences performance monitoring at the level of the FRN and P3 components. *Psychophysiology, 53*(7), 1020-1033. doi:10.1111/psyp.12651

Walsh, M. M., & Anderson, J. R. (2012). Learning from experience: event-related potential correlates of reward processing, neural adaptation, and behavioral choice. *Neurosci Biobehav Rev, 36*(8), 1870-1884. doi:10.1016/j.neubiorev.2012.05.008

Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nat Rev Neurosci, 12*(12), 739-751. doi:10.1038/nrn3112

Wolpert, D. M., & Flanagan, J. R. (2001). Motor prediction. *Curr Biol, 11*(18), R729-732.

Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nat Neurosci, 3 Suppl*, 1212-1217. doi:10.1038/81497

Wu, H. G., Miyamoto, Y. R., Castro, L. N. G., Ölveczky, B. P., & Smith, M. A. (2014). Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci, 17*, 312. doi:10.1038/nn.3616

https://www.nature.com/articles/nn.3616#supplementary-information

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol Rev, 111*(4), 931-959. doi:10.1037/0033-295X.111.4.939

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*(5), 535-544. doi:10.1093/cercor/bhh153

Yeung, N., & Sanfey, A. G. (2004). Independent coding of reward magnitude and valence in the human brain. *J Neurosci, 24*(28), 6258-6264. doi:10.1523/JNEUROSCI.4537-03.2004

Yeung, N., & Summerfield, C. (2012). Metacognition in human decision-making: confidence and error monitoring. *Philos Trans R Soc Lond B Biol Sci, 367*(1594), 1310-1321. doi:10.1098/rstb.2011.0416

Yu, A. J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron, 46*(4), 681-692. doi:10.1016/j.neuron.2005.04.026