Running head: Task learnability and reinforcement learning

**Task learnability modulates surprise but not valence
processing for reinforcement learning in probabilistic choice
tasks**

Wurm, F.^{1,2,3*}, Walentowska, W.^{4,5,*}, Ernst, B.¹, Severo, M. C.⁵, Pourtois, G.⁵, & Steinhauser, M.¹

¹Catholic University of Eichstätt-Ingolstadt, Germany

²Leiden University, Cognitive Psychology Unit, the Netherlands

³Leiden Institute for Brain and Cognition, Leiden, the Netherlands

⁴Jagiellonian University, Krakow, Poland

⁵Ghent University, Belgium

* These authors contributed equally to this work

Correspondence should be addressed to:

Franz Wurm

Leiden University

Wassenaarseweg 52

2333 AK Leiden

The Netherlands

E-mail: f.r.wurm@fsw.leidenuniv.nl

Phone: +49 (0)1520-5484886

Abstract

The goal of temporal difference (TD) reinforcement learning is to maximize outcomes and improve future decision making. It does so by utilizing a prediction error (PE), which quantifies the difference between the expected and the obtained outcome. In gambling tasks, however, decision making cannot be improved due to the lack of learnability. Based on the idea that TD utilizes two independent bits of information from the PE (valence and surprise), we asked which of these aspects is affected when a task is not learnable. We contrasted behavioral data and event-related potentials (ERPs) in a learning variant and a gambling variant of a simple two-armed bandit task, in which outcome sequences were matched across tasks. Participants were explicitly informed that feedback could be used to improve performance in the learning task but not in the gambling task, and we predicted a corresponding modulation of the aspects of the PE. We used a model-based analysis of ERP data to extract the neural footprints of the valence and surprise information in the two tasks. Our results revealed that task learnability modulates reinforcement learning via the suppression of surprise processing but leaves the processing of valence unaffected. Based on our model and the data, we propose that task learnability can selectively suppress TD learning as well as alter behavioral adaptation based on a flexible cost-benefit arbitration.

(222 words)

Keywords: Reinforcement learning, Event-related potentials, Feedback processing,
Model-based analysis

Introduction

How do agents optimally learn to select actions that lead to the achievement of their long-term goals? Reinforcement learning provides an elegant answer to this question by offering a normative computational framework of decision-making for goal-directed agents interacting with an uncertain environment (Sutton & Barto, 2018). While reinforcement learning spans a very broad range of different algorithms, the most prominent candidate is temporal difference (TD) learning, as it offers a simple yet powerful formalization of trial-and-error learning. More importantly, the existing evidence suggests distinct neural footprints of TD in both animals and humans (Botvinick, Niv, & Barto, 2009; Daw & O'Doherty, 2014; Holroyd & Coles, 2002; D. Lee, Hyojung, & Min Whan, 2012; Niv, 2009). At its core, TD learning revolves around the idea of a prediction error (PE), which is calculated as the difference between the expected and the actual outcome of an action or state in the environment. This PE can then be used as an internal teaching signal as it carries two independent bits of information: (1) information about valence, which describes if an outcome was better (+) or worse (-) than expected, and (2) information about surprise, which quantifies how far off the expectation was. For subsequent computations, this information can be used to update expectations and translate them into future behavior.

While the idea of TD learning emerged as an important principle for explaining human decision making (Botvinick et al., 2009; Daw & O'Doherty, 2014; Holroyd & Coles, 2002; D. Lee et al., 2012; Niv, 2009), little is known on how its application is influenced by explicit knowledge about the environmental structure. Crucial but still open questions are whether TD learning can be suppressed if a task is not learnable (i.e., outcome feedback is uninformative for optimizing behavior), and if so, for which bit of information this suppression occurs. In the present study, the effect of learnability is investigated by contrasting brain activity in a learning task and a gambling task. Considering electrophysiological correlates of feedback

processing in a model-based analysis, we asked which aspect of the PE is affected by learnability – information about valence or information about surprise.

The significance of TD learning for behavior is supported by evidence coming from both animal and human studies (Balleine, 2005; Dayan & Niv, 2008; Dickinson & Balleine, 2002). While early work showed that the firing pattern of midbrain dopamine neurons in monkeys closely resembles a PE as predicted by TD algorithms (Montague, Dayan, & Sejnowski, 1996; Schultz, Dayan, & Montague, 1997), similar findings were also later observed in humans (D'Ardenne, McClure, Nystrom, & Cohen, 2008; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006; Schonberg et al., 2010). Since then, correlates of the PE have been replicated extensively in subpopulations of several neural structures such as the striatum, the amygdala and multiple areas of the (pre)frontal cortex (Schultz, 2016).

In recent years, an increasing number of studies have used event-related potentials (ERPs) to investigate the neural dynamics underlying reinforcement learning. It has consistently been shown that ERPs are sensitive to feedback valence already at around 200 ms after feedback onset in a component called the feedback-related negativity (FRN, Miltner et al., 1997). The FRN manifests as a negative deflection at frontocentral electrodes about 200 to 350 ms after feedback onset. The FRN is typically larger for negative than positive feedback (for reviews, see San Martín, 2012; Walsh & Anderson, 2012), a phenomenon we refer to as the ΔFRN ¹. Crucially, numerous studies could show that the ΔFRN scales with outcome magnitude and likelihood, supporting theoretical assumptions that the FRN amplitudes also reflect information about surprise as suggested by the reinforcement learning theory (Holroyd & Coles, 2002; for reviews, see Sambrook & Goslin, 2015; San Martín, 2012; Walsh & Anderson, 2012). In addition to the ΔFRN , feedback elicits a P3, which has been associated

¹ Please note, that there is no clear consensus about the labelling of the FRN in the literature. To prevent confusion, we refer to the ΔFRN as the difference between positive and negative outcomes. For an insightful discussion on the naming and quantification of the FRN component, refer to a recent methodological review by Krigolson (2017).

with a broad range of cognitive phenomena (for reviews, see Polich, 2020; San Martín, 2012). Manifesting as a parietal positivity around 300 to 600 ms after feedback onset, this feedback-P3 has been mainly reported to reflect surprise (Mars et al., 2008; Kolossa, Kopp, & Fingscheidt, 2015; Kopp et al., 2016; Seer, Lange, Boos, Dengler, & Kopp, 2016) and learning (Fischer & Ullsperger, 2013; Jepma et al., 2018, 2016; Nassar, Bruckner, & Frank, 2019). Importantly, recent evidence suggests that the relationship between P3 and learning is strongly modulated by the task context, thereby supposedly reflecting a mediated response to surprise (Nassar et al., 2019). Despite the obvious link between the P3 and surprise, its role within a reinforcement learning perspective of human decision making is still unclear.

Crucial yet unanswered questions are whether and how top-down processes can modulate TD learning. These questions have recently been addressed in studies manipulating the goal relevance of feedback, that is, the informativeness of feedback for optimizing behavior. These studies revealed rather similar effects for the Δ FRN and the P3. When participants were instructed on the reliability of feedback in probabilistic learning (i.e., the probability that the same action will lead to the same outcome on future trials), the Δ FRN as well as the P3 was increased for reliable feedback as compared to unreliable feedback (Di Gregorio, Ernst, & Steinhauser, 2019; Ernst & Steinhauser, 2017, 2018; Schiffer, Siletti, Waszak, & Yeung, 2017). The modulation of the Δ FRN could partially be attributed to a top-down modulation because the effect was observed even when the objective reliability of feedback was held constant (Di Gregorio et al., 2019; Schiffer et al., 2017). Similar results were obtained in a speeded go/no-go task in which feedback could indicate whether responses were correct and faster than an unknown response time cut-off (Walentowska, Moors, Paul, & Pourtois, 2016) and in a time estimation task (Severo, Paul, Walentowska, Moors, & Pourtois, 2020). Again, the Δ FRN and the P3 were larger when feedback provided valid and relevant information for optimizing task performance.

The results described above could be taken as evidence that all aspects of TD learning are suppressed when feedback is goal irrelevant. However, such a conclusion might be premature for several reasons. First, as we have seen, the two main ERP components of feedback processing (Δ FRN and P3) cannot be unequivocally mapped onto the aspects of the PE. Valence and surprise information might contribute to both components, which raises the possibility that only one of these processes caused the Δ FRN and P3 results. A viable method to deal with this problem is to use computational modeling to provide direct estimates of the PE that can inform model-based analyses of the neural data (Gläscher & O’Doherty, 2010; Mars, Shea, Kolling, & Rushworth, 2012). Second, while the above-mentioned studies manipulated the information content of feedback, behavioral adaptation could still lead to performance improvements in all conditions, at least in some of the studies (e.g., Di Gregorio et al., 2019). A stronger manipulation would be to vary the learnability of tasks, that is, to contrast a probabilistic learning task with a mere gambling task. While performance can be improved by utilizing feedback in a learning task, gambling tasks do not allow for improving performance at all. From a normative perspective, the calculation of a PE has no utility in the latter condition. Therefore, the comparison between learning and gambling is a strong test for revealing top-down processes in TD learning. First evidence for such a modulation stems from a study in which the Δ FRN has been shown to vary with PEs in a learning task but not in a gambling task (Holroyd, Krigolson, Baker, Lee, & Gibson, 2009).

The aim of the present study is to investigate how the calculation of a PE is influenced by the learnability of the environment, which was manipulated by varying the causal structure of a task. Based on the idea that the PE carries information about both valence and surprise of an outcome, we asked which of these aspects is affected when participants are confronted with a learnable task and a non-learnable task, respectively. In both tasks, we used a simple variant of a two-armed bandit problem (Daw, O’Doherty, Dayan, Seymour, & Dolan, 2006).

Whereas participants could adapt to the varying reward probabilities by utilizing feedback in the learning task, feedback was fully unrelated to participants' choice behavior in the gambling task. To prevent the results from reflecting systematic differences in the frequency and order of feedback, the two tasks were matched with respect to feedback sequence (see Methods for details). Crucially, we used a computational modelling approach to separate the effects of learnability on the two aspects of the PE. By modeling behavior in the two tasks using multiple instantiations of the same computational algorithm, we could robustly estimate PEs for each single trial and use its constituent parts to predict neural data in a linear model. Based on the idea that valence information is strongly related to the FRN amplitudes whereas surprise information is reflected by the P3 (Mars et al., 2008; Nassar et al., 2019), we expected to find correlates of valence at earlier frontocentral locations and correlates of surprise at later posterior locations. Importantly, if TD learning is influenced by learnability, we hypothesized to find either a reduced Δ FRN or a reduced P3, or both, for the gambling task relative to the learning task. Moreover, any influence of task learnability on the calculation of prediction errors should be reflected by a modulation of the strength by which its constituent parts predict the respective neural activity.

Methods

Participants

30 participants (24 female) between 18 and 28 years of age ($M = 21.9$; $SD = 2.6$) with normal or corrected-to-normal vision and free from neurological and psychiatric history (based on self-reports) participated in the study. Participants were recruited at Ghent University. They received 20 € for participation and a performance-dependent bonus ($M = 2.16$ €; $SD = 1.14$ €). All participants provided informed consent and the study was approved

by the local ethics committee (Faculty of Psychology and Educational Sciences of Ghent University).

Stimuli

The experimental task was programmed and executed using Presentation software (Neurobehavioral Systems, Albany, CA). The stimulus set consisted of 112 black and white drawings obtained from the online database of the International Picture Naming Project (Székely et al., 2005). Pictures were converted into 300x300 pixel images corresponding to a visual angle of 6.4° in height and 6.4° in width at a viewing distance of 80 cm. For each block, eight pictures were randomly sampled from the stimulus array (without replacement) and then further divided into four stimulus pairs which were shown to the left and right of a centrally presented fixation cross at a distance of 0.7° . The left or right position for each stimulus in a pair was randomly determined for each trial. The feedback stimuli were a white diamond (win) and a white rock (loss) that were centrally presented at a height of 2.8° and a width of 2.5° . All stimuli were presented on a black background.

Task and procedure

We employed two variants (learning task vs. gambling task) of a simple two-armed bandit task in which participants had to repeatedly choose between two actions corresponding to the two stimuli in a stimulus pair (Fig. 1A). Both tasks differed regarding their objective learnability, i.e., the utility of feedback for improving performance. For the *learning task*, the mapping between actions and feedback was highly coherent, making behavioral adaptation based on feedback relevant for optimizing outcome. The two possible actions of each stimulus pair were associated with reward probabilities p and $(1 - p)$, whereas p varied throughout a block according to a Gaussian random walk with reflecting boundaries at 0 and 1 (e.g. Daw, 2011; Sambrook, Hardwick, Wills, & Goslin, 2018). On each repetition of this stimulus pair,

p was constantly updated by adding a value drawn from a Gaussian distribution with $M = 0$ and $SD = 0.1$ (Fig. 1B). If this update led to values below 0 or above 1, p was set to 0 or 1, respectively. For the *gambling task*, the mapping between action and feedback was pseudorandom. More specifically, we used the same feedback sequence as in the preceding learning-task block while presenting a randomly shuffled stimulus pair sequence. Due to this combination of a yoked feedback sequence and a randomized stimulus sequence, participants' actions could not influence subsequent feedback making behavioral adaptation irrelevant for optimizing outcome.

Figure 1.

The procedure of a trial was identical in the two different task variants and is illustrated in Figure 1A. At the beginning of each trial, a stimulus pair was presented for up to 3500 ms. An action had to be made by pressing one of two keys on an English standard keyboard (E with the left index finger for choosing the left stimulus, I with the right index finger for choosing the right stimulus). Following the key press, the stimulus pair disappeared, and a fixation cross was presented for a fixed time of 1000 ms. Afterwards, feedback was centrally presented for 1000 ms. Diamonds indicated wins, whereas rocks indicated losses. After the disappearance of the feedback stimulus, a fixation cross displayed for 1000 ms marked the end of the trial. If no action was carried out during presentation of the stimulus pair, the trial was aborted and counted as a miss. The feedback stimulus for misses was a white question mark. The participants were informed that misses were associated with a loss of 5 cents, whereas positive and negative feedback was associated with a win/loss of +/-2 cents, respectively. Participants worked through ten blocks with 64 experimental trials each. Learning and gambling tasks alternated across blocks, and the task of the first block was

counterbalanced across participants. Prior to each block, participants were informed about which task was used in the upcoming block. In each block, four new stimulus pairs were introduced. Blocks consisted of 16 subblocks in which all four stimulus pairs were presented in random order. To avoid unequal attention allocation to the feedback stimuli between tasks, we included 10 catch trials in each block which were randomly interspersed among the experimental trials. Participants' task in catch trials was to recall the previously shown feedback. Participants were presented a prompt ("What feedback was shown last?") and indicated the feedback by pressing the E key in case of a win, the I key in case of a loss, or the space key if they did not remember (or if the last trial was a miss).

At the beginning of the experiment, participants received written instructions on the task. The instructions included information on decision making and feedback in the task in general, as well as the basic properties of the learning and the gambling task. For the learning task, participants were instructed that they could learn which action is most likely to be followed by a win. Furthermore, they were informed about the implementation of the random walk and the resulting changes in reward probabilities. For the gambling task, participants were told that they are unable to learn which action is most likely to be followed by reward. Afterwards, they worked through two short practice blocks, one from each task, in which two stimulus pairs were presented ten times. These blocks were then followed by two long practice ones, again one from each task, in which four stimulus pairs were presented 16 times. The order of tasks in the practice blocks was the same as in the subsequent experimental blocks. For participants starting the experiment with the gambling task, the long training block for the learning task was used for initial yoking. In practice blocks with the gambling task, the feedback sequence was random and not yoked to a learning block.

After the experiment, participants completed a number of questionnaire items to collect data on subjective information and use of feedback in the two tasks using a computer-

mouse slider on the visual analog scale (VAS). For a detailed description of the items refer to Table 1. Each aspect was rated separately for the learning and gambling tasks on a continuous scale ranging between 0 and 100. For Items 1 to 5, the scale ranged between “not at all” and “a lot”. For Items 6 and 7, the scale ranged between “very unhappy” and “very happy”. Both the catch trials and the subjective ratings were implemented to quantify that the feedback between learning and gambling tasks was equally attended.

Behavioral data analysis

Choice behavior was analyzed in two steps. First, we considered the proportion of correct actions in the learning task. Correct actions were defined as those actions that were associated with a higher reward probability. The total proportion of correct actions was tested against chance level using a two-tailed paired-sample *t*-test. Furthermore, the proportion of correct actions calculated separately for each of the 16 subblocks was submitted to a one-way repeated measurement ANOVA with the variable subblocks. For both significance tests, the proportion of correct action was arcsine-square-root-transformed (Winer, Brown, & Michels, 1991). Second, we analyzed the proportion of stay behavior, which was defined as the probability of repeating the same action as on the previous encounter with the same stimulus pair. The goal of this analysis was to investigate whether we find the commonly observed win-stay lose-shift (WSLS) behavior as an indicator of behavioral adaptation, and whether this pattern was more pronounced in the learning task than in the gambling task. We therefore applied a logistic regression with choice type as criterion and previous outcome, task, and their interaction as predictors. Choice type could either be a stay or a switch (coded as 1 and 0), previous outcome could either be a win or a loss (coded as 1 and -1), and task could either be learning or gambling (coded as 1 and -1). All within-subjects variables (intercept, previous outcome, task, interaction) were implemented as random effects and therefore were allowed to vary across participants (e.g. Daw et al., 2011; Gillan et al., 2015). Performance on catch

trials was analyzed by calculating the proportion of catch trials on which the feedback of the previous trial was correctly indicated. This accuracy was compared between learning-task blocks and the gambling-task blocks using a paired-sample *t*-test. Finally, mean ratings for each item of the final questionnaire were compared between learning and gambling task using paired-sample *t*-tests.

Electrophysiological recordings and analyses

Participants were seated comfortably in a dimly lit, sound-attenuated, and electrically shielded cabin. The electroencephalogram (EEG) was recorded using a BIOSEMI Active-Two system (BioSemi, Amsterdam, The Netherlands) with 64 Ag-AgCl electrodes from channels Fp1, AF7, AF3, F1, F3, F5, F7, FT7, FC5, FC3, FC1, C1, C3, C5, T7, TP7, CP5, CP3, CP1, P1, P3, P5, P7, P9, PO7, PO3, O1, Iz, Oz, POz, Pz, CPz, Fpz, Fp2, AF8, AF4, AFz, Fz, F2, F4, F6, F8, FT8, FC6, FC4, FC2, FCz, Cz, C2, C4, C6, T8, TP8, CP6, CP4, CP2, P2, P4, P6, P8, P10, PO8, PO4, and O2, as well as the left and right mastoid. The horizontal and vertical EOG was monitored by means of four electrodes, placed above and below the right eye and the outer canthi of both eyes. Sampling rate was 512 Hz.

As with the behavioral data, EEG data was analyzed using custom-made routines in MatLab. For the preprocessing, we used EEGLAB 13.5.4b (Delorme & Makeig, 2004), an open source toolbox for EEG data analysis (EEGLAB toolbox for single-trial EEG data analysis, Swartz Center for Computational Neurosciences, La Jolla, CA; <http://www.sccn.ucsd.edu/eeglab>). EEG data were offline re-referenced to averaged mastoids, band-pass filtered to exclude frequencies below 0.1 Hz and above 35 Hz and divided into epochs from 1000 ms before to 1500 ms after feedback onset. Baseline activity from 200 ms before feedback onset was removed. Bad channels were interpolated using spherical spline interpolation if they met the joint probability criterion (threshold 5) as well as the kurtosis criterion (threshold 5) in EEGLAB's channel rejection routine. Epochs were excluded

whenever neural activity in a channel deviated more than $\pm 300 \mu\text{V}$ from the epoch mean. This criterion was not applied to those channels that are typically contaminated by blinks (Fp1, Fpz, Fp2, AF7, and AF8) as this activity was corrected later. In a next step, an infomax-based independent component analysis (ICA, Bell & Sejnowski, 1995) was conducted. After visual inspection of the derived independent components, those representing eye blinks and muscular artifacts were identified and removed from the data. The remaining epochs were averaged separately for each participant and task. On average, this resulted in the following numbers of artifact-free trials in the respective feedback/task conditions: 170.7 ($SD = 17.0$) for win/learning, 127.1 ($SD = 13.3$) for loss/learning, 170.0 ($SD = 17.3$) for win/gambling, and 126.3 ($SD = 14.7$) for loss/gambling.

FRN amplitudes were quantified using the mean amplitude in a time window of 200-400 ms after feedback presentation at electrode FCz (Sambrook & Goslin, 2014). The P3 amplitude was measured using a peak amplitude approach (Ernst & Steinhauser, 2018; Pontifex, Hillman, & Polich, 2009; Sailer, Fischmeister, & Bauer, 2010) because P3 peaks are often shifted across conditions, making a quantification in a fixed time interval difficult. We first identified the maximum amplitude at electrode Pz in a time window of 200-700 ms, and analyzed the peak amplitude as well as the latency of the peak. For statistical analysis, we applied repeated measurement ANOVAs involving the variables outcome (win, loss) and task (learning, gambling) for amplitudes in both the FRN and P3 time window.

Computational modeling

We used computational modeling for two reasons. Firstly, we sought to investigate the mechanisms at play during feedback processing in general. By constructing different instantiations of computational reinforcement learning models, we aimed to identify the mechanisms that are most likely implemented on a group-level. Secondly, we sought to leverage this mechanistic explanation of the behavioral data for a subsequent model-based

analysis of the EEG data. By deriving single-trial estimates of PEs and using their constituent parts (valence and surprise) as predictors for the neural data, we aimed to reveal the specific patterns of TD learning, especially regarding its modulation via the learnability of the environment.

The fundamental rationale which is shared between all computational models detailed below is that learning occurs in two successive steps: the calculation of a (reward) PE and the updating of action values based on this specific PE signal. We therefore implemented a learning policy that can be characterized as a TD learning architecture (Gläscher & O’Doherty, 2010; Sutton & Barto, 2018). The calculation of the PE follows the general approach with the equation

$$PE(t) = [r(t) - Q(a, s, t)], \quad (1)$$

where $r(t)$ denotes the outcome received in that trial, and $Q(a, s, t)$ denotes the expected value of the chosen action in the specific stimulus-pair. Please note that this calculation of the PE is most accurately called Rescorla-Wagner learning rule, as it does not include a term for the expected value of future states, which determines how agents navigate in more complex environments. Based on the PE, the expected value of the chosen action is incrementally updated according to the equation

$$Q(a, s, t + 1) = Q(a, s, t) + \alpha * PE(t), \quad (2)$$

where α is the learning rate, which controls for the speed of updating. The probability of selecting action a in a specific stimulus-pair s is then simply determined by inserting the updated action values in a soft-max decision rule

$$P(a_t = a | s) = \frac{\exp(\beta * Q(a, s, t) + p * rep(a, s))}{\sum_{a'} \exp(\beta * Q(a', s, t) + p * rep(a', s))}, \quad (3)$$

where the inverse temperature β guides the stochasticity of the choices and the perseveration parameter p captures choice perseveration ($p > 0$) or switching ($p < 0$) (Lau &

Glimcher, 2005). The indicator function $rep(a)$ takes a value of 1 if action a was chosen on the last trial of the same stimulus pair and 0 otherwise.

We hypothesized that the instruction on the nature of the task (learning vs. gambling) causes behavioral differences between these tasks by affecting parameter settings. In this study, we identified three possible mechanisms that could lead to differences between learning and gambling behavior. Each putative mechanism is linked to one specific model parameter. For models with a *variable learning rate*, the learning rate α is estimated separately for each task condition. This allows the updating of action values (Equation 2) to be modulated between the two tasks. As this parameter controls for the integration of new experience over trials, we expect to find higher learning rates for the learning task compared to the gambling task. For models with *variable inverse temperature*, the inverse temperature β is estimated separately for each task condition. This allows action selection (Equation 3) to be modulated between the tasks. As this parameter controls for the stochasticity of choice behavior (with higher values increasing experience-based action selection), we expect to find higher inverse temperatures for the learning task compared to the gambling task. For models with *variable policy mixture*, we took into consideration previous findings showing that behavioral and neural data can be best explained by a mixture of different choice policies (e.g., Daw et al., 2011; Keramati, Dezfouli, & Piray, 2011). Core mechanism of this mixture model is the eponymous mixture of action values for two separate and independent choice policies. Here, these policies are TD learning, as described earlier, and random choice or guessing behavior. Formally, the mixture of policies was realized by calculating a net action value Q_{NET} as the weighted combination of the independent action values according to the equation

$$Q_{NET}(a, s, t) = \omega Q_{LEARN}(a, s, t) + (1 - \omega) Q_{GUESS}(a, s, t), \quad (4)$$

where $Q_{LEARN}(a, s, t)$ denotes action values derived from a TD mechanisms (equation 1 and 2) and $Q_{GUESS}(a, s, t)$ denotes action values derived from a guessing mechanism. To allow random action selection, $Q_{GUESS}(a, s, t)$ is initialized as 0 and kept constant during the experiment. For action selection, Q_{NET} values are submitted to Equation 3. The RL-basedness parameter ω describes the contribution of the learning policy, thus guiding the tradeoff between the contribution of guessing and learning policy to overt behavior. If the RL-basedness parameter approaches 1, the learning policy is predominant and drives action selection. If the RL-basedness parameter approaches 0, the guessing policy is predominant and drives action selection. For models with this mechanism, this RL-basedness parameter ω was estimated separately for each task condition. In the gambling task, feedback cannot be predicted, which might make a mere stochastic action selection viable (Wurm, Ernst, & Steinhauser, 2020) and sometimes even advantageous to learning policies (Tervo et al., 2014). We therefore expect to find higher RL-basedness for the learning task compared to the gambling task.

Given the task design and instruction, feedback on a trial cannot only be used to update the value of the chosen action. In addition, the forgone action on each trial can be updated using the reversed prediction error. In contrast to the partial updating mechanism, as detailed in Equation 2, this full updating mechanism should speed up learning. In line with the literature (Burnside, Fischer, & Ullsperger, 2019; Palminteri, Lefebvre, Kilford, & Blakemore, 2016), it is plausible to assume that the participants exploited this feature of the task. Therefore, we incorporated each of the three architectures for both the partial and the full updating mechanism, resulting in six candidate models.

All models were implemented in MatLab v8.6 (The Mathworks Inc., Natick, MA), and the mfit toolbox (Gershman, 2016) was used for parameter fitting. Parameters were fitted to the simulated or observed data using non-uniform priors and specific bounds (Gershman,

2016). The learning rate was bound in $[0; 1]$ and the respective prior was $\alpha \sim \text{beta}(1.2, 1.2)$. The inverse temperature was bound in $[0; 20]$ the prior was $\beta \sim \text{gamma}(2, 1)$. The perseveration parameter was bound in $[-5; 5]$ and the prior was $p \sim \text{gauss}(0, 1)$. Finally, RL-basedness was bound in $[0; 1]$ and the prior was $w \sim \text{beta}(1.2, 1.2)$.

In order to show that the above-described models lead to distinct and identifiable behavioral patterns, we conducted a model recovery analysis. We simulated each of the six candidate models for 200 participants. Consistent with the empirical task, each simulated participant worked through 10 blocks with each 64 trials, task identity alternated across blocks and feedback for the gambling blocks was yoked. The parameters for each model and participant were drawn from [the same priors that are used for model fitting \(see above\)](#). To ensure behavioral modulation in line with our predictions for the empirical data, the variable parameters for each model (learning rate, inverse temperature or RL-basedness) were sampled independently for each task condition. For the variable learning rate models, learning rates for the learning task were uniformly drawn from the interval $[0.5; 1]$, whereas learning rates for the gambling task were drawn from the interval $[0; 0.5]$. For the variable inverse temperature models, inverse temperatures for the learning task were drawn from the interval $[1; 5]$, whereas inverse temperatures for the gambling task were drawn from the interval $[0; 1]$. For the variable policy mixture models, RL-basedness for the learning task were drawn from the interval $[0.5; 1]$, whereas RL-basedness for the gambling task were drawn from the interval $[0; 0.5]$. This simulation procedure ensures that the models generate meaningful behavioral patterns that are generated via quantitative differences in model parameters. The simulated data from each model was then subjected to model comparison in the same way as for the empirical data (see below) to verify that the model that best fit to the data is also the model that generated the data.

In a next step, we performed model selection based on the empirical data. We fitted each of the previously detailed models to the observed behavioral data from the participants. Three measures are reported for model selection: the Bayesian information criterion (BIC), the Akaike information criterion (AIC) and the protected exceedance probability (PXP). In contrast to the two former metrics, the latter is derived using a Bayesian approach with PXP quantifying the probability that any model considered for model comparison is more frequent than all other possible models, given equal prevalence in the population (Rigoux, Stephan, Friston, & Daunizeau, 2014). In contrast to BIC or AIC-informed approximation of log-evidence, and as implemented as default in the mfit toolbox, the PXP in our study is derived under the Laplace approximation (e.g., Friston, Mattout, Trujillo-barreto, Ashburner, & Penny, 2006). Moreover, a Bayesian approach for model selection allows us to derive posterior probabilities for each participant, which quantify the posterior belief that a specific model generated the data for that participant. As the model recovery analysis revealed that the PXP metric leads to a better recoverability than AIC and BIC, we relied our model selection procedure on this metric but also report BIC and AIC for comparison².

Subsequently, we compared parameter estimates from the empirical data between the learning and the gambling task using two-tailed paired-sample tests. Because learning rate and RL-basedness are interpreted as a ratio, their values were arcsine-square-root-transformed (Winer et al., 1991).

For the model-based single-trial analysis, we constructed a general linear model (GLM) to predict single-trial EEG activity at each electrode and time point, separately for each participant. Regressors included the task condition, and the information about valence and surprise. Task could either be learning or gambling (coded as 1 and -1). Valence was

² In contrast to the Laplace-informed Bayesian model selection, BIC-informed Bayesian model selection exhibits similar biases in model recovery as the BIC and AIC metric, namely favoring simpler models over complex models (in terms of number of free parameters).

defined as the sign of the PE (+1 vs. -1), surprise was defined as the absolute value of the PE. PEs for each candidate model and participant were simulated by feeding the estimated model parameters back into the same model that was initially used to calculate the model parameters. Based on the idea of random effects (different models may explain the data from different participants), we utilized the Bayesian model selection approach and weighted PEs from all models according to their subject-specific posterior belief to obtain a weighted average PE. This weighted average PE was used to derive the valence and surprise regressors. Prior to regression, the feedback-locked EEG data were downsampled to 125 Hz. All regressors were z-scored. For the surprise regressor, individual values were z-scored separately for the learning and the gambling task. Moreover, we included all possible interactions between regressors. The resulting linear equation took the following form for each task,

$$\begin{aligned}
 EEG \sim & b_0 + \\
 & b_1 * Task + b_2 * Valence + b_3 * Surprise \\
 & b_4 * TaskValence + b_5 * TaskSurprise + b_6 * ValenceSurprise + \quad (5) \\
 & b_7 * TaskValenceSurprise \\
 & + error,
 \end{aligned}$$

where composite terms (e.g., *TaskValence*) denote the interaction between the respective regressors. To ensure comparability within task conditions and between subjects and to penalize multicollinearity of predictors, the resulting beta values were standardized by their respective standard deviation (Fischer & Ullsperger, 2013). Only after this normalization procedure, the individual beta weights were tested against zero via two-tailed cluster-based permutation tests implemented in the Mass univariate ERP Toolbox (Groppe, Urbach, & Kutas, 2011) conducted across sampling points and electrodes. Corrections for a family-wise alpha level of 0.05 were applied and clusters were identified for all sampling points at which the uncorrected p value fell below 0.05. We used 10^5 permutations to obtain sufficient test

distributions. All sampling points between 0 and 1000 ms after feedback presentation and all electrodes were considered for permutation testing. To investigate how reinforcement learning is modulated between the learning and gambling task, we applied the identical analysis to the difference between beta values from both tasks, separately for RPEs and action values.

Results

Behavioral data

In a first analysis, we investigated whether participants show learning in the learning task by considering the mean percentage of correct actions. An action was defined as correct if the reward probability associated with this action exceeded 50%. Accordingly, a correct decision was not defined by actual positive feedback but by the underlying reward probability. Due to the non-learnable structure of the gambling task, correct actions can be determined only for the learning task. Overall, participants' performance in this task ($M = 62.1\%$; $SEM = 1.3\%$) significantly exceeded chance level, $t(29) = 9.22$, $p < .001$, $d = 1.68$. Comparing performance across subblocks, we found a significant main effect of subblock, $F(15, 29) = 4.47$, $p < .001$, $\eta_p = 0.13$. Figure 2A suggests that no performance improvement occurs beyond the second trial. This reflects that reward probabilities are not only variable but even include reversals of the correct actions, which prevents a monotonically increasing performance. Nevertheless, these results indicate that participants were able to use feedback for improving performance in the learning task.

Figure 2.

Whereas the preceding analysis shows that participants can utilize coherent action-feedback mappings in the learning task to improve performance, we were now interested whether participants adapted behavior based on feedback by considering switch/stay behavior across consecutive encounters with the same stimulus pair. A common finding is that participants stay with their choices after wins but switch choices after losses (e.g., Cohen & Ranganath, 2007; Daw et al., 2011). This pattern of WSLS behavior is also possible in the gambling task, even though it cannot lead to performance improvement due to the non-learnable structure of this task. We hypothesized to find more evidence for win-stay lose-shift behavior in the learning task provided that the learnability influences the strength of behavioral adaptation. As illustrated in Figure 2B, the probability of stay behavior was indeed higher following wins than following losses in both tasks, but this effect was larger for the learning task. Logistic regression analysis showed a significant main effect for previous reward, $\beta = 0.39$, $p < .001$, and task, $\beta = 0.32$, $p < .001$, as well as a significant interaction between both, $\beta = 0.27$, $p < .001$. Separate analyses for each task revealed that previous reward affected stay behavior for both, the learning task, $\beta = 0.67$, $p < .001$, and the gambling task, $\beta = 0.11$, $p = .002$. These results demonstrate that the learnability of a task influences the strength of behavioral adaptation based on feedback.

In a last step, we analyzed catch trial performance and subjective ratings. Analyzing accuracy for catch trials, we found that although mean values were high for both the learning task ($M = 94.0\%$; $SEM = 1.0\%$) and the gambling task ($M = 86.6\%$; $SEM = 1.6\%$), performance was significantly higher in the learning task, $t(29) = 5.89$, $p < .001$, $d = 1.08$, indicating decreased attentional resource allocation in the gambling context. Such an interpretation receives further support when analyzing the subjective ratings obtained after the experiment (Table 2). Here, participants explicitly reported a significantly stronger allocation of attention towards feedback in the learning task compared to the gambling task.

Additionally, the learning task was associated with stronger interest in the task, higher perceived contribution to win, and more blame following losses. Crucially, outcome in the learning task was perceived to be significantly more useful compared to the gambling task. Taken together, the results from both catch trial performance and subjective ratings provide support for the effectiveness of our manipulation of learnability.

Table 1.

ERPs

After finding differences in behavioral adaptation between tasks, we were interested to elucidate the underlying neural mechanisms by analyzing feedback-locked ERPs. If reinforcement learning is influenced by learnability, we hypothesized to find either a reduced Δ FRN or a reduced P3, or both, for the gambling task relative to the learning task.

Analyzing amplitudes in the time window of the FRN at electrode site FCz (Fig. 3A), we found a significant main effect of both outcome, $F(1, 29) = 66.15, p < .001, \eta_p = 0.70$, and task, $F(1, 29) = 15.72, p < .001, \eta_p = 0.35$. The FRN was increased (i.e., more negative) for loss ($M = 2.22 \mu V$, $SEM = 0.80$) relative to win feedback ($M = 6.47 \mu V$, $SEM = 0.72$), and was increased in the gambling ($M = 3.25 \mu V$, $SEM = 0.76$) relative to the learning task ($M = 5.44 \mu V$, $SEM = 0.83$). However, we did not find a significant interaction between task and outcome, $F(1, 29) = 0.70, p = .410, \eta_p = 0.02$. That is, the Δ FRN represented by the difference between win and loss (Fig. 3B-C) was comparable in the learning and gambling tasks.

Figure 3.

Figure 4 shows the results for the analysis for the P3. In contrast to the previous analysis, the size of the P3 was quantified as the amplitude of the peak in the time range 200 – 700 ms at electrode Pz, determined individually for each task and participant³. This method also allows for analyzing possible latency shifts across tasks, which appear to be evident in Figure 4A. For the analysis of peak amplitudes we found a marginally significant main effect of outcome, $F(1, 29) = 4.16, p = .051, \eta_p = 0.13$, and a significant main effect of task, $F(1, 29) = 27.09, p < .001, \eta_p = 0.48$, but these effects were qualified by a significant interaction between both variables, $F(1, 29) = 6.34, p = .018, \eta_p = 0.18$. Win feedback showed a larger P3 than loss feedback in the gambling task, $t(29) = 3.06, p = .005$, but no such difference was obtained in the learning task, $t(29) = 0.21, p = .837$. Notably, however, the P3 was larger in the learning than the gambling task for both win feedback, $t(29) = 3.19, p = .003$, and loss feedback, $t(29) = 5.62, p < .001$. Although Figure 4A suggests that the P3 peaks for loss feedback occurred slightly later than the peaks for win feedback, no significant effects were obtained for the analysis of peak latencies (all $ps > .16$).

Figure 4.

Taken together, these results show that learnability influences feedback processing on the neural level. However, this effect differed according to the stage of feedback processing. While the amplitude was generally more negative for the gambling task in the time window of the FRN, the Δ FRN (i.e., the increased negativity for losses relative to wins), which has previously been associated with reinforcement learning and RPE calculation, was not

³ Please note that a similar pattern of results was obtained when we quantified the P3 as mean amplitude in a time window between 200 – 500 ms at electrode Pz.

modulated by learnability. In contrast, the P3 was generally larger for the learning task but showed an effect of feedback valence only for the gambling task.

After showing the impact of learnability on feedback processing, we also included subsequent choice as a variable for the ANOVA in an exploratory analysis to investigate the relationship between feedback-related brain activity and subsequent behavior. Subsequent choice codes for whether the upcoming choice for the same stimulus pair is the same (stay) or different (switch) to the current trial, thus linking neural feedback processing with WSLS behavior. In order to analyze a sufficient number of trials (>10), we had to exclude three participants from this additional step. For the time window of the FRN, we found similar effects of outcome and task as previously reported. In addition, the interaction between subsequent choice and task condition was significant, $F(1,26) = 9.23, p = .005, \eta_p = 0.26$. Subsequent switching behavior was associated with stronger positive deflections compared to stay behavior only in the learning task, $t(26) = 2.68, p = .013$, but not in the gambling task, $t(26) = 0.37, p = .716$. While amplitudes in the time window of the FRN revealed only a marginally significant interaction between all three involved factors, $F(1, 26) = 3.00, p = .095, \eta_p = 0.10$, amplitudes in the P3 time window clearly indicate the same interactive effect, $F(1, 26) = 7.75, p = .001, \eta_p = 0.23$. Only for the learning task, there was a significant interaction between subsequent choice and outcome, $F(1, 26) = 9.56, p = .005, \eta_p = 0.27$. Only for wins, subsequent switching ($M = 17.46 \mu V, SEM = 0.67$) indicates higher P3 amplitudes compared to stay behavior ($M = 14.56 \mu V, SEM = 0.50$), $t(26) = 4.10, p < .001$. For losses subsequent switching ($M = 15.12 \mu V, SEM = 0.55$) was not different from stay behavior ($M = 14.90 \mu V, SEM = 0.68$), $t(26) = 0.34, p = .735$. There was no comparable effect for the gambling task, $F(1, 26) = 0.12, p = .732, \eta_p = 0.00$, indicating that there is no link between subsequent choice and P3 amplitude for the gambling task. In conclusion, the results from this exploratory analysis reveal that subsequent choice can be predicted from neural activity following

feedback processing. More specifically, following the idea of WSLS, it seems that P3 reflects the overcoming of such behavioral tendencies in the learning task (Daw et al., 2006), thus suggesting a link between activity in the P3 window and behavioral adaptation that goes beyond a simple TD account of behavior.

Model-based analysis

After finding evidence for an influence of learnability on choice behavior and the neural correlates of feedback processing, we applied model-based analysis to reveal the functional mechanisms underlying these effects. In comparison to the previously reported analyses, computational modelling allows us to investigate the mechanistic underpinnings of decision-making and learning by establishing a formally explicit link between model variables and behavioral and neural activity. Based on the assumption that this link resembles evidence for TD, we hypothesized to find modulations of learnability, leading to separable patterns in the two tasks.

In order to validate the identifiability of the proposed mechanisms, we first conducted a model recovery procedure during which we simulated behavioral data from 100 participants for each of the six candidate models. At their core, these models share the rationale for learning but use different variable parameters to explain differences between the learning and gambling tasks. *Variable learning rate* models assume that instructions on learning and gambling tasks lead to distinct learning rates; *variable inverse temperature* models assume that the instruction drives distinct inverse temperature; whereas *variable policy mixture* models assume that behavior can be captured as a mixture of a pure learning and a pure guessing policy. Simulated data from these models were then subjected to model selection. Bayesian model comparison via PXP revealed that our models were highly recoverable from the simulated data, as shown in Figure 5. Interestingly, the BIC and AIC metric were unable

to recover the models which generated the simulated data, which appears to reflect a bias towards simple models due to the strong weighting of model complexity in these measures. Taken together, these results suggest that model recovery can be subject to considerable biases arising from the specific comparison metrics. Given these findings, we conclude that, [for the present task and models](#), PXP seems to be the most reliable metric for model comparison.

Figure 5.

Afterwards, we fitted the six candidate models to the empirical data. Table 2 shows BIC, AIC and PXP for each model. Although the variable policy mixture model with the full updating mechanism has the highest PXP (37%), the overall pattern (see Table 2) indicates that models are similarly prevalent in the population. In line with our model recovery findings, our interpretation based on PXP strongly deviates from the BIC and AIC metrics, which clearly identify the variable learning rate model with full updating as the best model. While this finding from BIC and AIC seems to have a straightforward interpretation, the previous model recovery procedure already pointed towards a considerable shortcoming of those metrics, namely the confusion between models involved and the preference for simple compared to complex model. Although BIC and AIC are still widely used in the literature, their application is increasingly criticized, especially with the emergence of alternative metrics that remedy some common shortcomings (e.g., Rigoux, Stephan, Friston, & Daunizeau, 2014; Stephan et al., 2009).

Table 2.

One of those common shortcomings is that BIC and AIC must be interpreted as fixed-effects metrics, i.e., only one model accounts for all the data simultaneously. On the other hand, Bayesian model selection, as employed in this study, derives random-effects metrics which imply that different models can account for data from different participants. In line with this assumption and based on the PXP, we extracted for each participant the posterior beliefs that a specific model generated its data. As shown in Figure 6, the dominating model varied considerably between participants. In line with the finding that PXP did not identify a winning model on the group-level, this further supports the existence of strong interindividual differences in learning strategies.

Figure 6.

In a next step, we compared the estimated values of the different parameter between the learning and gambling tasks (Table 3). We found that the variable parameters for all models are significantly different in the learning task compared to the gambling task. Crucially, RL-basedness for the winning mixture model with full updating were higher in learning tasks compared to gambling tasks, $t(29) = 11.88, p < .001, d = 2.73$. In line with our hypothesis, this indicates that the dominant driving force for behavior in the learning task was action selection under the (reinforcement) learning policy, whereas the dominant driving force for behavior in the gambling task was action selection under the guessing policy.

Table 3.

In a final step, we used the estimated PE from our model to investigate whether valence and surprise processing is modulated by learnability. To this end, we simulated these

single-trial values for each model and participant. To account for interindividual differences, we calculated the weighted average PE across models for each participant from the posterior probabilities and correlated these values with neural activity using multiple linear regression analyses. We hypothesized to find a stronger connection between model estimates and neural data in the learning task than in the gambling task for those aspects of the PE that are influenced by learnability.

As the model-based regressor for valence perfectly coincides with the outcome variable in the ERP analyses, we expected to find similar modulations in our single-trial analysis. In line with this hypothesis, we found two significant clusters that distinguish between wins and losses (Fig. 6). First, there was a positive correlation between valence and neural activity at frontocentral electrode sites ranging between 92 and 452 ms ($p < .001$). Second, there was a sustained negative correlation at posterior (occipital) electrode sites ranging between 60 and 1000 ms ($p = .009$). Crucially, these valence-driven neural patterns were not modulated between tasks ($p > .083$), suggesting that learnability does not modulate the processing of valence information.

In contrast to valence information, surprise processing was clearly modulated between the two tasks, as indicated by a significant interaction between task and surprise at posterior electrode sites in a time window between 244 and 748 ms ($p = .036$). Based on the visual inspection of the topographies, we selected the time course at electrode site Pz, where the significant positive correlation in the learning task ranged between 372 and 908 ms ($p = .001$). For the gambling task, no comparable correlation was evident (all $ps \geq .592$). In line with our hypothesis, this indicates that learnability impacts TD learning on the neural level, leading to a suppression of surprise processing if the task is not learnable, as is the case in the gambling task.

Interestingly, we also found a significant interaction between valence and surprise ($p < .008$), suggesting that these two aspects of the PE are interrelated on the neural level. The effect covered from 212 to 1000 ms but was most pronounced over central electrode sites in a at around 300 ms. While the correlation between surprise and neural activity was evident when valence was positive ($p = .021$), no such correlation was evident for negative valence ($p = .060$).

Figure 7.

Please note that the qualitative and quantitative patterns of the single-trial regression from the Laplace-informed posterior probabilities can be replicated for each individual model as well as the BIC-informed posterior probabilities. In line with suggestions from the literature (Wilson & Niv, 2015), model-based neural results are often robust across a wide variety of computational models and not specifically conditioned on the choices for computational modelling.

Discussion

Temporal difference (TD) learning revolves around the concept of a prediction error (PE), which carries information about the valence and the surprise of a specific outcome. In the present study, we investigated the influence of learnability on the neural underpinnings of these constituent aspects of the PE by contrasting a learning variant and a gambling variant of a simple two-armed bandit task. The learnable structure of the learning task was realized as a learnable link between actions and subsequent feedback, whereas the random structure of the gambling task was realized as the absence of such a learnable link. Based on the idea of a goal-directed agent, we hypothesized that the difference in learnability should lead to a

corresponding modulation of the behavioral and neural patterns of (TD) learning, and we asked which level of processing could be affected by this modulation.

Using computational modeling, we successfully extracted the neural footprints of both valence and surprise information in our EEG data. On the one hand, information about the surprise of an outcome was reflected in a pronounced central cluster. Due to its spatial and temporal distribution, this effect cannot be easily assigned to a specific component of the human ERPs. Crucially, however, differences between task conditions were evident over posterior sites that strongly resemble the P3. As indexed by the absence of any correlation between surprise and neural activity in the gambling task, its expression appeared to be largely suppressed when feedback processing was random and thus irrelevant for behavioral adaptation. On the other hand, information about the valence of an outcome was reflected by an early frontocentral cluster of activation that is spatially and temporally very similar to the Δ FRN. In contrast to surprise, the manipulation of learnability had no influence on valence calculation. The fully random feedback in the gambling task led to the same valence effect as the feedback in the learning task, even though learning was strongly reduced in the former. In sum, these findings support the idea that the different aspects of the PE are differentially modulated by the learnability of the environment. While early processing of valence in the time range of the FRN remains intact across the two task conditions, subsequent processing of surprise could be suppressed for environments with a random structure, i.e. during the gambling task.

Interestingly, our results of the single-trial analysis also contribute to answering the question on how information about valence and surprise interact on the neural level. In the cognitive neuroscience literature, there is an ongoing debate on this issue. Following the earliest account of the FRN (Holroyd & Coles, 2002), activity in the time window of the FRN should reflect surprise only for negative outcomes (negative and signed PE). Based on new

evidence, which showed that the neural response to feedback was mainly driven by positive feedback, this account was subsequently altered and updated (Holroyd et al., 2009; Holroyd, Pakzad-Vaezi, & Krigolson, 2008). Now, a reward positivity is assumed to be elicited after positive outcomes (positive and signed PE). Alternatively, others have suggested that PEs are reflected in the brain irrespective of the valence of the outcome (unsigned PE or surprise; Alexander & Brown, 2011, 2015). While evidence from a meta-analysis suggests that the FRN is sensitive to an unsigned PE (Sambrook & Goslin, 2015), only few studies actually estimated PEs using computational modelling (Burnside et al., 2019; Fischer & Ullsperger, 2013; Sambrook et al., 2018; Walsh & Anderson, 2011). In our study, surprise was reflected in the neural data more strongly for positive than negative outcomes, thus adding to the idea of a positive and signed PE signal in the brain, which however is not reflected in an FRN or reward positivity, but rather in a more posterior brain activity.

Regarding our central manipulation of learnability, the analysis of ERPs revealed a pattern of results that showed commonalities but also differences as compared to that of the model-based analysis. Although the overall amplitude in the FRN time window was influenced by learnability, the Δ FRN (i.e., the increased negativity for negative feedback as compared to positive feedback) did not differ between the learning and the gambling task, suggesting that the frontocentral cluster reflecting valence information in the model-based analysis is linked to the Δ FRN. In contrast to the FRN, P3 amplitudes revealed a modulation by learnability. Surprisingly, the difference in P3 amplitudes between win and loss outcomes was evident in the gambling task but was absent for the learning task. This pattern could reflect that due to the yoking of feedback sequences, positive feedback was unexpectedly more frequent in the gambling task, thus inducing an additional expectancy effect in the P3 data.

The comparison of results from the ERP and model-based analyses demonstrates that a direct mapping of ERP components to distinct (cognitive) processes is insufficient and premature. The observation that the gambling task shows a generally reduced P3 (possibly reflecting reduced learning), but an increased P3 valence effect suggests that more than one process is reflected in this component. The same could be inferred from the differential effect of learnability on overall FRN amplitudes in general, and on the Δ FRN in particular.

Additional leverage for the idea of multiple parallel neural processes reflected in the distinct ERP components comes from our exploratory analysis, suggesting that the P3 reflects an overcoming of behavioral tendencies imposed by TD learning. Independent support for this idea comes from different lines of research. For example, a recent study showed that while the link between the FRN and behavioral adaptation was driven by reinforcement learning, the link between P3 and behavioral adaptation was triggered by explicit rules that go beyond patterns derived from reinforcement learning (Chase, Swainson, Durham, Benham, & Cools, 2011). Based on the idea that action selection involves switching between exploitative and explorative modes (Daw et al., 2006), the P3 could provide a neural correlate for a switch towards explorative behavior.

On a methodological level, the present study highlights a crucial advantage of model-based analyses, which enabled us to extract functionally meaningful activity from the neural data by combining the strength of multiple methods (computational modelling, regression and cluster-based permutation testing). Computational modelling is particularly beneficial when theoretical variables cannot be controlled sufficiently by experimental design, as is the case with surprise in volatile environments and learning in general. Even under such challenging conditions, we were thus able to derive predictions about neural patterns. In contrast to the discrete nature of the traditional ERP approach, the use of a regression approach in our study allowed us to refrain from splitting surprise into bins (e.g., median or mean), but analyze it

instead as a continuous variable. As the regression method has gained popularity over the last few years, there is now a huge body of literature for further discussion, but also comprehensive introduction of this matter (Gläscher & O'Doherty, 2010; Mars et al., 2012; O'Doherty, Hampton, & Kim, 2007; Wilson & Niv, 2015). While the traditional ERP approach always relies on ad-hoc defined components that often consist of a mixture of different activities (e.g., theta band and delta band in the FRN time window; Bernat et al., 2015; Foti, Weinberg, Bernat, & Proudfit, 2015; Harper, Malone, & Bernat, 2014) and are often a subject to debates regarding their quantification and functional meaning (Krigolson, 2017; Picton et al., 2000), cluster-based permutation testing is not bound to strong constraints in both temporal as well as spatial domain, making it a valuable tool for both confirmatory and explorative analyses (for a helpful discussion on the advantages and pitfalls, see Maris & Oostenveld, 2007; Sassenhagen & Draschkow, 2019). Although each of the methods described above makes a specific contribution to the main question under scrutiny, only their combination allowed us to gain novel and critical insight into feedback processing in the human brain, which would not have been possible using traditional methods. Hence, their combination allowed us to foster our understanding of the putative underpinnings of decision-making and learning.

Our study design and results are complemented by a recent study contrasting feedback processing under an active and a observational learning condition (Burnside et al., 2019). Using model-based single-trial regression analyses, the authors showed that neural activity in the time window of the FRN is driven by a PE and clearly modulated between task conditions. Neural activity in the time window of the P3, however, is driven only by outcome and is not modulated between task conditions. While these results are in stark contrast to the new findings reported here above, it is noteworthy to take a closer look at the paradigm used in that previous study and underlying these effects. In the active condition, participants played

a probabilistic but stationary bandit task, whereas in the observational condition they merely watched another person. While the active condition shares some similarities with our learning condition, outcome information in the observation condition is orthogonally different to our gambling condition. On the one hand, observational learning still necessitates learning from (observed) action, but without the outcome being relevant for the agent's goal achievement. Gambling on the other hand does not necessitate learning, while the outcome is still relevant for goal-achievement. These differences in task design are mirrored on the neural level and indicate the impact of top-down processing selectively modulating TD learning based on environmental demands.

The question emerges whether, rather than the selective suppression of surprise processing, differences in visual attention to stimuli and feedback could have produced the present results. Recent studies highlighted the importance of attention for reinforcement learning using a multidimensional bandit task, in which only one out of multiple stimulus dimensions was relevant for feedback (Leong, Radulescu, Daniel, DeWoskin, & Niv, 2017; Niv et al., 2015). Critically, however, they found attention to unselectively constrain reinforcement learning, whereas only the surprise processing was selectively modulated in our unidimensional bandit task. This could suggest that our results do not primarily reflect differences in attention between learning and gambling tasks. In fact, our task design, especially the implementation of catch trials, was chosen to control for attentional differences between tasks. Although attention to feedback was slightly weaker in the gambling task, as indicated by both catch trial accuracy and subjective ratings, the same measures also showed that attention was clearly allocated to feedback even in the gambling task, and ERP data as well as model-based analysis revealed that the effect of valence was identical between tasks. The interpretation that attention plays only a minor role for our results is further supported by

our computational modelling results. In sum, our findings preclude a simple alternative explanation by mere attentional differences.

The accuracy of model-based analyses strongly relies on multiple factors such as the model selection procedure and the validity of both the applied computational model and experimental manipulation (Nassar & Frank, 2016; Nassar & Gold, 2013; Palminteri, Wyart, & Koehlin, 2017; Rigoux et al., 2014; Stephan et al., 2009; Wilson & Niv, 2015). Therefore, we first evaluated the identifiability of our computational models using a model recovery procedure. This procedure indicated that recoverability differed considerably across metrics. Consequently, we relied model comparison on the metric which is most convincing given the model recovery procedure (i.e., the protected exceedance probability). Independent of model comparison, parameter estimates from each model revealed a significant difference between instructed task types, further validating our experimental manipulation.

The obtained context-dependent arbitration between policies/parameters is highly adaptive. In the learnable environment, behavior was dominated by reinforcement learning principles and showed clear hallmarks of behavioral adaptation. Otherwise, in the random environment, behavior was dominated by stochastic choice (although some behavioral adaptation was still observable). Irrespective of the true underlying mechanism (decreased learning rate, increased temperature, or decreased RL-basedness) stochastic choice is often implemented in computational models of learning. Its purpose is mostly to serve as a baseline or dummy condition that is compared to biologically plausible choice policies (e.g., Doll, Duncan, Simon, Shohamy, & Daw, 2015; Steingroever, Wetzels, & Wagenmakers, 2014; Worthy, Hawthorne, & Otto, 2013; Worthy & Maddox, 2014). The present findings however suggest that stochastic choice behavior can be utilized in a goal-directed way as well. In line with recent animal findings (Tervo et al., 2014) and based on the idea of a cost-benefit-arbitration (Kool, Gershman, & Cushman, 2017, 2018), we suggest that stochastic choice

behavior maximizes outcomes, while simultaneously minimizing computational costs, for example by withholding resources budgeted for the processing of surprise or the updating of expectations.

The idea that learnability exerts its influence by altering the trade-off between alternative internal task representations/parameters is compatible with suggestions in the existing literature. It has already been assumed that learning tasks are characterized by a reliance on external feedback, whereas gambling tasks are characterized by a reliance on internal expectations (Hajcak, Moser, Holroyd, & Simons, 2007; Holroyd, Hajcak, & Larsen, 2006; Holroyd et al., 2009). Similar ideas received attention also within other domains of reinforcement learning research (Daw & O'Doherty, 2014; Dayan & Niv, 2008; Niv, 2009). A frequent assumption is that reinforcement learning is achieved by the interplay between a model-free system that learns from experience and a model-based system in which internal task models drive behavior in a more top-down way. According to one proposal, arbitration between these different systems for decision making is based on the relative uncertainty of their respective estimates with the less uncertain estimate predominantly driving behavior (Daw, Niv, & Dayan, 2005). While such uncertainty-based arbitration is implemented on a trial-to-trial basis (S. W. Lee, Shimojo, & O'Doherty, 2014), related evidence suggests that instructions can influence the arbitration between learning systems via the prefrontal cortex loops that modulate reinforcement learning in the striatum (Doll, Hutchison, & Frank, 2011; Doll, Jacobs, Sanfey, & Frank, 2009). A similar neural pathway between the prefrontal cortex and the striatum may account for the putative top-down modulation of surprise in the present study.

Taken together, our results show that task learnability can modulate reinforcement learning by means of two mechanisms. On the one hand, this modulation occurs via the selective suppression of surprise when a task is unlearnable, leaving the evaluation of valence

unaffected. On the other hand, we suggest that behavioral adaptation on the basis of task learnability is driven by a flexible cost-benefit arbitration.

References

- Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience*, *14*, 1338–1344.
- Alexander, W. H., & Brown, J. W. (2015). Hierarchical error representation: A computational model of anterior cingulate and dorsolateral prefrontal cortex.
- Balleine, B. W. (2005). Neural bases of food-seeking: Affect, arousal and reward in corticostriatolimbic circuits. *Physiology and Behavior*, *86*, 717–730.
- Bell, A. J., & Sejnowski, T. J. (1995). An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, *7*, 1129–1159.
- Bernat, E. M., Nelson, L. D., & Baskin-Sommers, A. R. (2015). Time-frequency theta and delta measures index separable components of feedback processing in a gambling task. *Psychophysiology*, *52*, 626–637.
- Botvinick, M. M., Niv, Y., & Barto, A. G. (2009). Hierarchically organized behavior and its neural foundations: A reinforcement learning perspective. *Cognition*, *113*, 262–280.
- Burnside, R., Fischer, A. G., & Ullsperger, M. (2019). The feedback - related negativity indexes prediction error in active but not observational learning. *Psychophysiology*, e13389.
- 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. *Journal of Cognitive Neuroscience*, *23*, 936–946.
- Cohen, M. X., & Ranganath, C. (2007). Reinforcement learning signals predict future decisions. *Journal of Neuroscience*, *27*.
- D’Ardenne, K., McClure, S. M., Nystrom, L. E., & Cohen, J. D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science*, *319*, 1264–1267.

- Daw, N. (2011). Trial-by-trial data analysis using computational models. In M. R. Delgado, E. A. Phelps, & T. W. Robbins (Eds.), *Decision making, affect, and learning* (Vol. 23, pp. 3–38). New York: Oxford University Press.
- Daw, N., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69, 1204–1215.
- Daw, N., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8, 1704–1711.
- Daw, N., & O'Doherty, J. P. (2014). Multiple systems for value learning. In *Neuroeconomics* (pp. 393–410). Elsevier.
- Daw, N., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876–879.
- Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, 18, 185–196.
- Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134, 9–21.
- Di Gregorio, F., Ernst, B., & Steinhauser, M. (2019). Differential effects of instructed and objective feedback reliability on feedback - related brain activity. *Psychophysiology*, e13399.
- Dickinson, A., & Balleine, B. W. (2002). The role of learning in the operation of motivational systems. In C. R. Gallistel (Ed.), *Steven's handbook of experimental psychology: Learning, motivation and emotion* (Vol. 3, pp. 497–534). New York: Wiley.
- Doll, B. B., Duncan, K. D., Simon, D. A., Shohamy, D., & Daw, N. (2015). Model-based choices involve prospective neural activity. *Nature Neuroscience*, 18, 767–772.
- Doll, B. B., Hutchison, K. E., & Frank, M. J. (2011). Dopaminergic genes predict individual

- differences in susceptibility to confirmation bias. *Journal of Neuroscience*, *31*, 6188–6198.
- Doll, B. B., Jacobs, W. J., Sanfey, A. G., & Frank, M. J. (2009). Instructional control of reinforcement learning: A behavioral and neurocomputational investigation. *Brain Research*, *1299*, 74–79.
- Ernst, B., & Steinhauser, M. (2017). Top-down control over feedback processing: The probability of valid feedback affects feedback-related brain activity. *Brain and Cognition*, *115*, 33–40.
- Ernst, B., & Steinhauser, M. (2018). Effects of feedback reliability on feedback-related brain activity: A feedback valuation account. *Cognitive, Affective and Behavioral Neuroscience*, *18*, 596–608.
- Fischer, A. G., & Ullsperger, M. (2013). Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron*, *79*, 1243–1255.
- Foti, D., Weinberg, A., Bernat, E. M., & Proudfit, G. H. (2015). Anterior cingulate activity to monetary loss and basal ganglia activity to monetary gain uniquely contribute to the feedback negativity. *Clinical Neurophysiology*, *126*, 1338–1347.
- Friston, K., Mattout, J., Trujillo-barreto, N., Ashburner, J., & Penny, W. (2006). Variational free energy and the Laplace approximation, *34*, 220–234.
- Gershman, S. J. (2016). Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology*, *71*, 1–6.
- Gillan, C. M., Otto, A. R., Phelps, E. A., & Daw, N. (2015). Model-based learning protects against forming habits. *Cognitive, Affective, & Behavioral Neuroscience*, *15*, 523–536.
- Gläscher, J. P., & O’Doherty, J. P. (2010). Model-based approaches to neuroimaging: Combining reinforcement learning theory with fMRI data. *Wiley Interdisciplinary Reviews: Cognitive Science*.

- Groppe, D. M., Urbach, T. P., & Kutas, M. (2011). Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review. *Psychophysiology*, 48, 1711–1725.
- Hajcak, G., Moser, J. S., Holroyd, C. B., & Simons, R. F. (2007). It's worse than you thought: The feedback negativity and violations of reward prediction in gambling tasks. *Psychophysiology*, 44, 905–912.
- Harper, J., Malone, S. M., & Bernat, E. M. (2014). Theta and delta band activity explain N2 and P3 ERP component activity in a go/no-go task. *Clinical Neurophysiology*, 125, 124–132.
- Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, 109, 679–709.
- Holroyd, C. B., Hajcak, G., & Larsen, J. T. (2006). The good, the bad and the neutral: Electrophysiological responses to feedback stimuli. *Brain Research*, 1105, 93–101.
- Holroyd, C. B., Krigolson, O. E., Baker, R., Lee, S., & Gibson, J. (2009). When is an error not a prediction error? An electrophysiological investigation. *Cognitive, Affective, & Behavioral Neuroscience*, 9, 59–70.
- Holroyd, C. B., Pakzad-Vaezi, K. L., & Krigolson, O. E. (2008). The feedback correct-related positivity: Sensitivity of the event-related brain potential to unexpected positive feedback. *Psychophysiology*, 45, 688–697.
- Jepma, M., Brown, S. B. R. E., Murphy, P. R., Koelewijn, S. C., de Vries, B., van den Maagdenberg, A. M., & Nieuwenhuis, S. (2018). Noradrenergic and cholinergic modulation of belief updating. *Journal of Cognitive Neuroscience*, 30, 1803–1820.
- Jepma, M., Murphy, P. R., Nassar, M. R., Rangel-Gomez, M., Meeter, M., & Nieuwenhuis, S. (2016). Catecholaminergic regulation of learning rate in a dynamic environment. *PLoS Computational Biology*, 12, 1–24.

- Keramati, M., Dezfouli, A., & Piray, P. (2011). Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Computational Biology*, 7.
- Kolossa, A., Kopp, B., & Fingscheidt, T. (2015). A computational analysis of the neural bases of Bayesian inference. *NeuroImage*, 106, 222–237.
- Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychological Science*, 28, 1321–1333.
- Kool, W., Gershman, S. J., & Cushman, F. A. (2018). Planning complexity registers as a cost in metacontrol. *Journal of Cognitive Neuroscience*, 30, 1391–1404.
- Kopp, B., Seer, C., Lange, F., Kluytmans, A., Kolossa, A., Fingscheidt, T., & Hoijtink, H. (2016). P300 amplitude variations, prior probabilities, and likelihoods: A Bayesian ERP study. *Cognitive, Affective, & Behavioral Neuroscience*, 16, 911–928.
- Krigolson, O. E. (2017). Event-related brain potentials and the study of reward processing: Methodological considerations. *International Journal of Psychophysiology*, 1–9.
- Lau, B., & Glimcher, P. W. (2005). Dynamic response-by-response models of matching behavior in rhesus monkeys. *Journal of the Experimental Analysis of Behavior*, 84, 555–579.
- Lee, D., Hyojung, S., & Min Whan, J. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, 35, 287–308.
- Lee, S. W., Shimojo, S., & O’Doherty, J. P. (2014). Neural computations underlying arbitration between model-based and model-free learning. *Neuron*, 81, 687–699.
- Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron*, 93, 451–463.
- Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, 164, 177–190.

- Mars, R. B., Debener, S., Gladwin, T. E., Harrison, L. M., Haggard, P., Rothwell, J. C., & Bestmann, S. (2008). Trial-by-trial fluctuations in the event-related electroencephalogram reflect dynamic changes in the degree of surprise. *Journal of Neuroscience*, 28, 12539–12545.
- Mars, R. B., Shea, N. J., Kolling, N., & Rushworth, M. F. (2012). Model-based analyses: Promises, pitfalls, and example applications to the study of cognitive control. *Quarterly Journal of Experimental Psychology*, 65, 252–267.
- Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a “generic” neural system for error detection. *Journal of Cognitive Neuroscience*, 9, 788–798.
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *The Journal of Neuroscience*, 16, 1936–1947.
- Nassar, M. R., Bruckner, R., & Frank, M. J. (2019). Statistical context dictates the relationship between feedback-related EEG signals and learning. *ELife*, 8, 1–26.
- Nassar, M. R., & Frank, M. J. (2016). Taming the beast: Extracting generalizable knowledge from computational models of cognition. *Current Opinion in Behavioral Sciences*, 11, 49–54.
- Nassar, M. R., & Gold, J. I. (2013). A healthy fear of the unknown: Perspectives on the interpretation of parameter fits from computational models in neuroscience. *PLoS Computational Biology*, 9, 1–6.
- Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, 53, 139–154.
- Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C. (2015). Reinforcement learning in multidimensional environments relies on attention

- mechanisms. *Journal of Neuroscience*, 35, 8145–8157.
- O’Doherty, J. P., Hampton, A. N., & Kim, H. (2007). Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of Sciences*, 1104, 35–53.
- Palminteri, S., Lefebvre, G., Kilford, E. J., & Blakemore, S. (2016). Confirmation bias in human reinforcement learning. *BioRxiv*, 1–21.
- Palminteri, S., Wyart, V., & Koechlin, E. (2017). The Importance of Falsification in Computational Cognitive Modeling. *Trends in Cognitive Sciences*, 21, 425–433.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, 442, 1042–1045.
- Picton, T. W., Bentin, S., Berg, P., Donchin, E., Hillyard, S. A., Johnson Jr., R., ... Taylor, M. J. (2000). Guidelines for using human event-related potentials to study cognition: recording standards and publication criteria. *Psychophysiology*, 37, 127–152.
- Polich, J. (2020). 50+ years of P300: Where are we now? *Psychophysiology*, 57, 2–3.
- Pontifex, M. B., Hillman, C. H., & Polich, J. (2009). Age, physical fitness, and attention: P3a and P3b. *Psychophysiology*, 46, 379–387.
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies - Revisited. *NeuroImage*, 84, 971–985.
- Sailer, U., Fischmeister, F. P. S., & Bauer, H. (2010). Effects of learning on feedback-related brain potentials in a decision-making task. *Brain Research*, 1342, 85–93.
- Sambrook, T. D., & Goslin, J. (2014). Mediofrontal event-related potentials in response to positive, negative and unsigned prediction errors. *Neuropsychologia*, 61, 1–10.
- Sambrook, T. D., & Goslin, J. (2015). A neural reward prediction error revealed by a meta-analysis of ERPs using great grand averages. *Psychological Bulletin*, 141, 213–235.

- Sambrook, T. D., Hardwick, B., Wills, A. J., & Goslin, J. (2018). Model-free and model-based reward prediction errors in EEG. *NeuroImage*, 178, 162–171.
- San Martín, R. (2012). Event-related potential studies of outcome processing and feedback-guided learning. *Frontiers in Human Neuroscience*, 6, 1–17.
- Sassenhagen, J., & Draschkow, D. (2019). Cluster-based permutation tests of MEG/EEG data do not establish significance of effect latency or location. *Psychophysiology*.
- Schiffer, A. M., Siletti, K., Waszak, F., & Yeung, N. (2017). Adaptive behaviour and feedback processing integrate experience and instruction in reinforcement learning. *NeuroImage*, 146, 626–641.
- Schonberg, T., O’Doherty, J. P., Joel, D., Inzelberg, R., Segev, Y., & Daw, N. (2010). Selective impairment of prediction error signaling in human dorsolateral but not ventral striatum in Parkinson’s disease patients: evidence from a model-based fMRI study. *NeuroImage*, 49, 772–781.
- Schultz, W. (2016). Dopamine reward prediction error coding. *Dialogues in Clinical Neuroscience*, 18, 23–32.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275, 1593–1599.
- Seer, C., Lange, F., Boos, M., Dengler, R., & Kopp, B. (2016). Prior probabilities modulate cortical surprise responses: A study of event-related potentials. *Brain and Cognition*, 106, 78–89.
- Severo, M. C., Paul, K., Walentowska, W., Moors, A., & Pourtois, G. (2020). Neurophysiological evidence for evaluative feedback processing depending on goal relevance. *NeuroImage*, 215, 116857.
- Steingroever, H., Wetzels, R., & Wagenmakers, E.-J. (2014). Absolute performance of reinforcement-learning models for the Iowa Gambling Task. *Decision*, 1, 161–183.

- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*, 1004–1017.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). Cambridge: The MIT Press.
- Székely, A., D’Amico, S., Devescovi, A., Federmeier, K., Herron, D., Iyer, G., ... Bates, E. (2005). Timed action and object naming. *Cortex*, *41*, 7–25.
- Tervo, D. G. R., Proskurin, M., Manakov, M., Kabra, M., Vollmer, A., Branson, K., & Karpova, A. Y. (2014). Behavioral variability through stochastic choice and its gating by anterior cingulate cortex. *Cell*, *159*, 21–32.
- 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*, 1020–1033.
- Walsh, M. M., & Anderson, J. R. (2011). Modulation of the feedback-related negativity by instruction and experience. *Proceedings of the National Academy of Sciences*, *108*, 19048–19053.
- Walsh, M. M., & Anderson, J. R. (2012). Learning from experience: Event-related potential correlates of reward processing, neural adaptation, and behavioral choice. *Neuroscience and Biobehavioral Reviews*, *36*, 1870–1884.
- Wilson, R. C., & Niv, Y. (2015). Is model fitting necessary for model-based fMRI? *PLoS Computational Biology*, *11*, 1–21.
- Winer, B. J., Brown, D. R., & Michels, K. M. (1991). *Statistical principles in experimental design* (3rd ed.). New York: McGraw-Hill.
- Worthy, D. A., Hawthorne, M. J., & Otto, A. R. (2013). Heterogeneity of strategy use in the Iowa gambling task: A comparison of win-stay/lose-shift and reinforcement learning models. *Psychonomic Bulletin and Review*, *20*, 364–371.

- Worthy, D. A., & Maddox, W. T. (2014). A comparison model of reinforcement-Learning and win-stay- lose-shift decision-making processes: A tribute to W.K. Estes. *Journal of Mathematical Psychology*, 59, 41–49.
- Wurm, F., Ernst, B., & Steinhauser, M. (2020). The influence of internal models on feedback-related brain activity. *Cognitive, Affective, & Behavioral Neuroscience*.

Author contributions

FW, WW, BE, MCS, GP and MS designed the study; WW and MCS collected the data; FW analyzed the data; FW, WW, BE, MCS, GP and MS wrote the paper. All authors approved the final version of the manuscript prior to submission.

Acknowledgments

This work is supported by funding from the National Science Centre of Poland (2015/19/B/HS6/01259) and the Polish National Agency for Academic Exchange (Bekker Programme signature: PPN/ BEK/2018/1/00257), awarded to WW. Moreover, this work is supported by research grants (G024716N, G048119N) from the Research Foundation Flanders (FWO), awarded to GP and MCS. The authors declare no conflict of interests.

Author Note

Correspondence concerning this article should be addressed to Franz Wurm, Leiden University, Wassenaarseweg 52, 2333 AK Leiden, the Netherlands. E-mail: f.r.wurm@fsw.leidenuniv.nl.