Hippocampal Contribution to Probabilistic Feedback Learning: Modeling Observation- and Reinforcement-based Processes

Virginie M. Patt[1,2], Daniela J. Palombo[3], Michael Esterman[1,2], and Mieke Verfaellie[1,2]

Abstract

■ Simple probabilistic reinforcement learning is recognized as a striatum-based learning system, but in recent years, has also been associated with hippocampal involvement. This study examined whether such involvement may be attributed to observation-based learning (OL) processes, running in parallel to striatum-based reinforcement learning. A computational model of OL, mirroring classic models of reinforcement-based learning (RL), was constructed and applied to the neuroimaging data set of Palombo, Hayes, Reid, and Verfaellie [2019. Hippocampal contributions to value-based learning: Converging evidence from fMRI and amnesia. *Cognitive, Affective & Behavioral Neuroscience, 19*(3), 523–536]. Results suggested that OL processes may indeed take place concomitantly to reinforcement learning and involve activation of the hippocampus and central orbitofrontal cortex. However, rather than independent mechanisms running in parallel, the brain correlates of the OL and RL prediction errors indicated collaboration between systems, with direct implication of the hippocampus in computations of the discrepancy between the expected and actual reinforcing values of actions. These findings are consistent with previous accounts of a role for the hippocampus in encoding the strength of observed stimulus–outcome associations, with updating of such associations through striatal reinforcement-based computations. In addition, enhanced negative RL prediction error signaling was found in the anterior insula with greater use of OL over RL processes. This result may suggest an additional mode of collaboration between the OL and RL systems, implicating the error monitoring network. ■

INTRODUCTION

The basal ganglia and hippocampal memory systems are traditionally viewed as separate, with the basal ganglia (and in particular the ventral striatum) enabling trial-and-error reinforcement-based learning and the hippocampus facilitating efficient and flexible learning of single instance events (Squire, 2004; Knowlton, Mangels, & Squire, 1996). Recent studies have uncovered interactions between these systems (Packard & Goodman, 2013; Gold, 2004; White & McDonald, 2002), with the hippocampus proposed to contribute enhanced flexibility and efficiency to striatum-based probabilistic reinforcement learning (Gershman & Daw, 2017). Such collaboration has been especially evident in tasks that involve multistep sequential decisions, in which the hippocampus is thought to facilitate learning of an internal model of the task, through simulation of the possible sequences of states and rewards that may follow any chosen action (i.e., model-based learning; e.g., Wang, Schoenbaum, & Kahnt, 2020; Bornstein & Daw, 2012, 2013). It has also been evidenced in reinforcement learning tasks involving contingencies that drift over time, with suggestion that the hippocampus may contribute to decisions through sampling of previous outcomes (e.g., Bornstein, Khaw, Shohamy, & Daw, 2017).

Hippocampal involvement in reinforcement learning has also been demonstrated in simpler probabilistic learning tasks involving nonsequential states and constant contingency rates, notably when reinforcement-based feedback was delayed by a few seconds (Foerde, Race, Verfaellie, & Shohamy, 2013; Foerde & Shohamy, 2011) or when learning involved acquisition of value-based representations (Palombo, Patt, Hunsberger, & Verfaellie, 2021; Palombo, Hayes, Reid, & Verfaellie, 2019; Dickerson & Delgado, 2015; Dickerson, Li, & Delgado, 2011). The mechanism by which the hippocampus contributes to probabilistic reinforcement learning in these simpler tasks remains to be elucidated.

In a number of studies, hippocampal activation has been demonstrated to correlate with the trial-by-trial reinforcement learning prediction error (e.g., Davidow, Foerde, Galván, & Shohamy, 2016; Dickerson et al., 2011; Foerde & Shohamy, 2011; see supplemental materials in the work of Schonberg et al., 2010), suggesting a role for the hippocampus in computing the discrepancy between actual and expected reinforcing outcomes. These findings, however, rather than supporting a direct contribution of the hippocampus to striatal learning, could be attributed to another learning pathway relying on observation-based processes. Indeed, in most probabilistic reinforcement learning tasks,

[1]VA Boston Healthcare System, MA, [2]Boston University School of Medicine, MA, [3]University of British Columbia, Vancouver, Canada

observation-based learning can be present concomitantly with reinforcement-based learning. That is, the feedback provided at the end of each trial can simultaneously serve to validate the selected action (e.g., selecting the red flower to go with the blue butterfly in Foerde and Shohamy [2011]; selecting the "greater than five" button to go with the circle shape in the work of Dickerson and Delgado [2015]) and to validate the observed stimulus association that results from the action (e.g., the blue butterfly–red flower association; or the circle shape–"greater than five" association). In other words, the prediction error in these tasks may simultaneously qualify as a reinforcement-based and as an observation-based signal. Thus, if reinforcement and observation-based processes take place conjointly during probabilistic feedback learning, it is possible that striatal activation relates to prediction error via the former process, and hippocampal activation relates to prediction error via the latter process.

Past research has implicated the hippocampus in the construction of episodic predictions (Buckner, 2010; Johnson, van der Meer, & Redish, 2007; Schacter, Addis, & Buckner, 2007) and in signaling novelty or surprise when these predictions are violated (c.f. mnemonic prediction error or mismatch detection; e.g., Sinclair, Manalili, Brunec, Adcock, & Barense, 2021; Bein, Duncan, & Davachi, 2020; Hindy, Avery, & Turk-Browne, 2019; Chen, Cook, & Wagner, 2015; Duncan, Ketz, Inati, & Davachi, 2012; Kumaran & Maguire, 2006). Much of this evidence involves representations that pertain to deterministic information, visuospatial maps, or single episodic events, for which the role of the hippocampus has been well established (Burgess, Maguire, & O'Keefe, 2002; Eichenbaum, 2000; Cohen & Eichenbaum, 1993; Squire, Knowlton, & Musen, 1993). In the current study, we explore the role of the hippocampus in progressive learning based on the gradual acquisition of probabilistic contingencies, a domain usually attributed to the striatum. Although that literature is less extensive, some studies have implicated the hippocampus in the coding of observed events that are uncertain or probabilistic (Harrison, Duggins, & Friston, 2006). For example, hippocampal involvement has been demonstrated when learning about temporal regularities of observed events (Schapiro, Kustner, & Turk-Browne, 2012; Turk-Browne, Scholl, Chun, & Johnson, 2009) and when forming associative perceptual predictions (Kok & Turk-Browne, 2018). Furthermore, in tasks requiring probabilistic learning, hippocampal activation has been demonstrated to scale with the strength of stimulus–outcome probabilistic associations independent of choice (Boorman, Rajendran, O'Reilly, & Behrens, 2016) and with a form of observation-based prediction error, scaled with respect to contextual expectation (Bunzeck, Dayan, Dolan, & Duzel, 2010).

Evidence for hippocampal involvement in reinforcement learning on the one hand and in encoding the association among observed events more broadly raises the possibility that the role of the hippocampus in simple probabilistic feedback learning tasks may reflect observation-based learning processes, running in parallel to reinforcement-based learning. To examine this possibility, we developed a computational model of observation-based learning and applied it, along with a model of classic reinforcement learning, to the neuroimaging data set of Palombo et al. (2019). The probabilistic feedback learning task of Palombo et al. was selected for its unique feature of providing reinforcement-based and observation-based feedback that are simultaneous and yet distinct and uncorrelated. In that task, participants learned the status of a series of visual stimuli, depicted as stick figures distinguished by different-colored patterns, that won money with predetermined probabilistic contingencies. On each trial, participants judged whether they believed that a stimulus figure would win by making yes/no responses and received subsequent feedback consisting of the winning outcome for the stimulus figure on that trial (i.e., "the man won $1.00" or "the man didn't win money"). Feedback is therefore directly relevant to the observed winning outcome associated with a stimulus figure, independent of choice; however, it also indirectly provides information about the accuracy of the response. In whole-brain and ROI general linear model analyses, Palombo et al. (2019) reported bilateral activation of the anterior hippocampus and of the nucleus accumbens, a key area of the ventral striatum. Here, we aim to further characterize the nature of the observed hippocampal involvement by using computational modeling.

Modeling Observation- and Reinforcement-based Probabilistic Feedback Learning

Classic models of reinforcement learning (RL) focus on computing the expected reinforcing values $V(s,a)$ of all possible actions a in a specific context or state s (Sutton & Barto, 2018; Rescorla & Wagner, 1972). As the participant gains experience through trial and error, the reinforcing values of these state–action pairs are updated using a delta rule:

$$V(s,a) \leftarrow V(s,a) + \alpha_{RL}(r - V(s,a)) \quad (1)$$

where r is the reinforcement outcome equal to 1 if the participant's choice was reinforced on that trial and to 0 if it was not reinforced, the term $PE_{RL} = (r - V(s,a))$ is the reinforcement prediction error, and α_{RL} is the updating parameter, which represents the extent to which feedback is used to update the expected values of options. The probability of choosing an action a given a specific state s may then be computed as a function of the reinforcing values of the state–action pairs using a softmax rule to account for nonsystematic behavior:

$$prob_{RL}(a/s) = \frac{e^{\beta_{RL}V(s,a)}}{\sum_{\text{all possible actions}} e^{\beta_{RL}V(s,a')}} \quad (2)$$

where β_{RL} is the exploit–explore parameter, with larger β_{RL} corresponding to more systematic choice of the option

with greater expected reinforcing value, and smaller β_{RL} corresponding to more random choice behavior.

Mirroring classic RL modeling, a model of observation-based probabilistic feedback learning (OL), is proposed here, which focuses on computing the expected probabilities of events $P_{ev}(s)$, given a specific context or state s, independent of participants' actions. As the participant gains information through observation of probabilistic events, the expected probabilities of those events are updated using a delta rule, similar to that of the RL model:

$$P_{ev}(s) \leftarrow P_{ev}(s) + \alpha_{OL}(event - P_{ev}(s)) \quad (3)$$

where *event* represents the outcome of the event on that trial (e.g., set to 1 if the event led to one type of outcome or to 0 otherwise), the term $PE_{OL} = (event - P_{ev}(s))$ represents the prediction error in terms of observed occurrence of that outcome, and α_{OL} is the updating parameter. In OL modeling, responses can also be modeled, but their goals must be directly related to the observation of the outcome event—for example, selecting which event is most likely to occur. Similar to RL modeling, the probability of a response given presentation of a specific state s is computed with a softmax rule to account for nonsystematic behavior, but with a denominator covering all possible outcome events rather than all possible actions:

$$Prob_{OL}(response = ev/s) = \frac{e^{\beta_{OL}P_{ev}(s)}}{\sum_{\text{all possible events}} e^{\beta_{OL}P_{ev'}(s)}} \quad (4)$$

where β_{OL} represents the exploit–explore parameter, with larger β_{OL} corresponding to more systematic selection of the option with greater expected probability of occurrence, and smaller β_{OL} corresponding to more random choice behavior.

This Study

This study aims to determine whether hippocampal involvement during probabilistic feedback learning may reflect the activity of observation-based learning processes, running in parallel to striatal reinforcement learning. To test this hypothesis, we examined within- and between-individuals differences in the relative use of one learning process over the other via comparison of overall model fit. We hypothesized that preferential use of observation-based over reinforcement-based learning would be associated with greater activation of the hippocampus and that preferential use of reinforcement-based over observation-based learning would be associated with greater activation of the striatum. Here, as in the work of Palombo et al. (2019), we focus on the nucleus accumbens, a key structure of the ventral striatum that integrates dopamine and limbic inputs and motor effector outputs (Floresco, 2007; Mogenson, Jones, & Yim, 1980) and has intrinsic connectivity with the hippocampus (Kahn & Shohamy, 2013). The brain correlates of the OL and RL prediction errors (PE_{RL}s) were also examined to shed light on the updating mechanisms underlying both learning processes—that is, to determine the brain regions involved in computing discrepancies between expected and actual events, defined in terms of event occurrence in OL and action reinforcement in RL. We hypothesized that the OL prediction error (PE_{OL}) would correlate with hippocampal activity and that the PE_{RL} would correlate with striatal activity.

METHODS

Data Set

The current study makes use of the Palombo et al. (2019) data set, collected on 30 healthy right-handed college students (15 men, 15 women) as they performed a probabilistic feedback learning task while undergoing functional neuroimaging in an MRI scanner. The participants had a mean age of 19.6 years (SD = 1.0 year) and mean education of 13.2 years (SD = 1.1 year). They had no prior history of major psychiatric or neurological condition. The sample size of Palombo et al. (2019) was gauged adequate for the computational modeling needs of this study, based on previous RL computational modeling work that used a similar or smaller sample size (ns between 17 and 30) and reported significant brain correlates of the model, notably in the hippocampal region (Dickerson et al., 2011; Foerde & Shohamy, 2011; Schonberg et al., 2010).

Paradigm

During the task, participants were presented with stimulus figures, one at a time, along with the written question "Does the man win money?" They responded by pressing the "Yes" or "No" buttons on an MRI-compatible button box. Decision screen duration was fixed (2134 msec) and was followed first by a short display of the stimulus figure alone (400 msec) and then by the outcome screen (1067 msec), which revealed "the man wins $1.00" along with a picture of a dollar bill, or "the man does not win money" along with an opaque gray rectangle labeled with "$0.00" (see Figure 1). Trials were separated by a fixation cross, with jittered duration (M = 2801 msec; range: 667–9203 msec). The task required learning the winning status of two sets of six experimental stimuli, depicted as stick figures with different fractal visual patterns. Each set comprised three stimulus figures associated with winning outcomes 75% of the time, and three with not-winning outcomes 75% of the time. Control trials were included involving two additional stimulus figures in each set. These trials were identical to experimental trials, but without the learning component, with correct response displayed on the stimulus figure during the decision screen ("Yes" or "No") and yielding winning or not-winning outcome 100% of the time. The two sets of stimuli were presented over the course of four runs, two consecutive runs per set. Each run comprised 48 experimental trials and 16 control trials, with eight presentations of each stimulus figure randomly interspersed. The presentation order

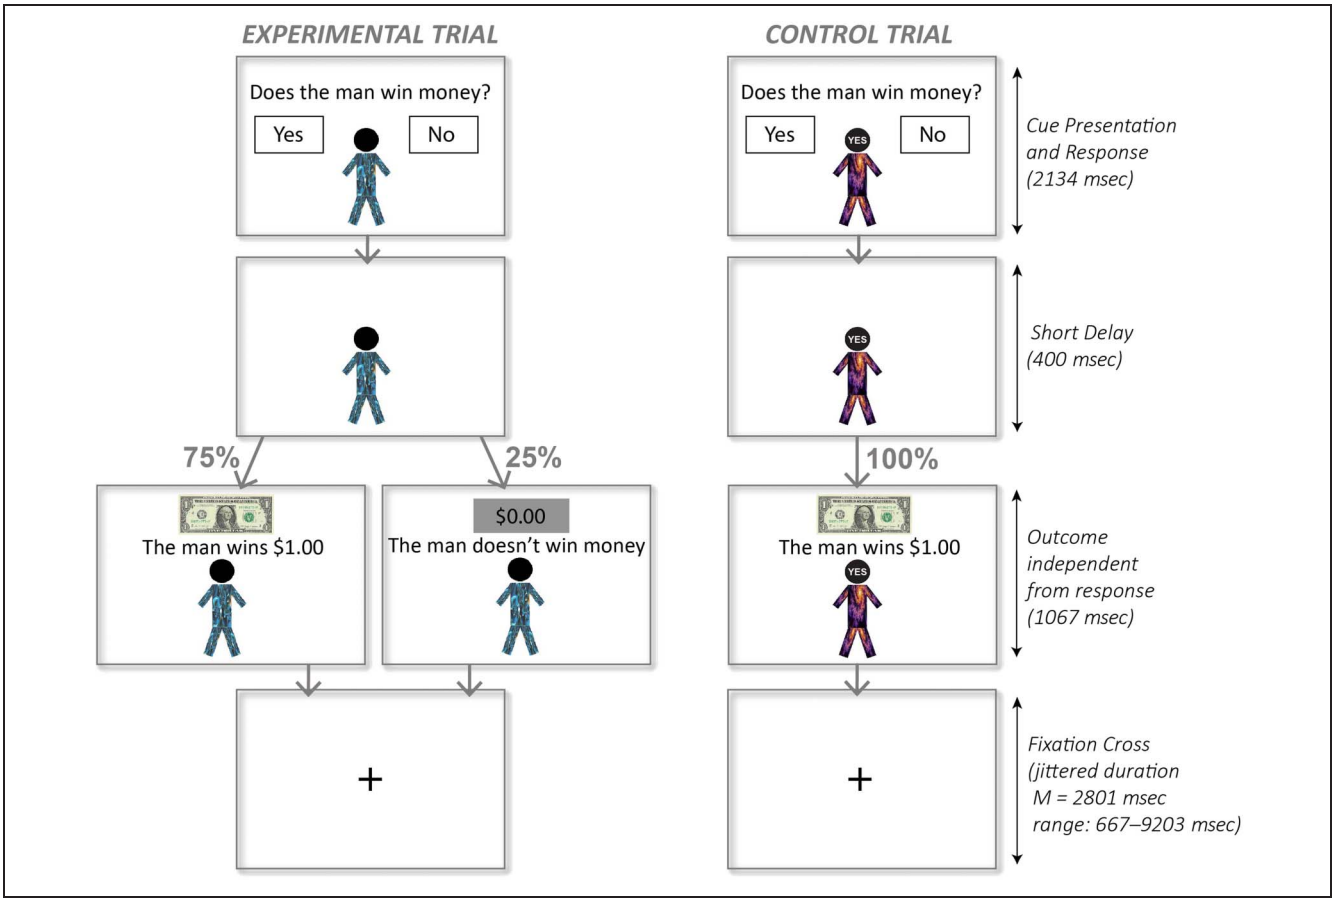

Figure 1. Illustration of an experimental and control trial from the Palombo et al. (2019) paradigm, depicting stimulus figures with winning status. Outcome probabilities are reversed for stimulus figures with no-win status.

of the runs was quasirandomized for each participant, and the assignment of winning status to the stimulus figures was counterbalanced across participants (see the work of Palombo et al. [2019] for a more detailed description of the procedures.)

After the scan, a test phase was carried out, in which, on each trial, two stimulus figures from the experimental trials were presented side by side, and participants had to decide which of the two was most likely to win. This phase comprised 36 trials, 18 trials in which the two stimulus figures were from the same learning set and 18 trials in which they were from different learning sets. Performance on each trial was coded as 1 for correct and 0 for incorrect.

Participants were paid $60 for their participation in the study. There was no additional monetary payment contingent upon participant responses. The rewards available to participants during the task were therefore experienced through two types of intrinsic events: (1) the observation of a stimulus figure winning $1 and (2) the experience of being correct when guessing an outcome.

Brain Imaging Acquisition and Preprocessing

Images were acquired on a 3.0 Tesla Siemens Prisma scanner with a 64-channel head coil. Sequences included a T1-weighted magnetization prepared rapid gradient echo sequence (sagittal plane acquisition, repetition time = 2530 msec, echo time = 3.35 msec, inversion time = 1100 msec, flip angle = 7°, sections = 176, slice thickness = 1 mm, matrix = 2562, field of view = 256 mm, voxel size = 1 mm^3), four functional scans with acquisition parallel to the anterior–posterior commissural plane (multiband = 6; repetition time = 1067 msec, echo time = 34.80 msec, flip angle = 65°, slices = 72, slice thickness = 2 mm, field of view = 208, matrix = 1042, voxel size = 2 mm^3, volumes = 388, phase encoding = anterior–posterior), and a brief functional scan with the same parameters but posterior–anterior encoding direction for correction of distortions. Functional imaging data processing was carried out using tools from the FMRIB Software Library v6.0 (Smith et al., 2004). Preprocessing included motion correction (using motion correction using FMRIB's Linear Image Registration Tool), susceptibility field correction (applytopup), skull stripping (Brain Extraction Tool), and bias-field correction (FMRIB's Automated Segmentation Tool). Prestatistic processing included spatial smoothing (Gaussian kernel of FWHM 5 mm), grand-mean intensity normalization, removal of additional motion artifacts (ICA-based Automatic Removal Of Motion Artifacts), and high-pass temporal filtering

(Gaussian-weighted least squares straight line fitting, sigma = 30 sec). Within-subject registration was carried out to the T1-weighted structural image using FMRIB Linear Image Registration Tool, and between-subjects registration to the MNI152 standard-space template using FMRIB Nonlinear Image Registration Tool (see the work of Palombo et al. [2019] for additional details).

Computational Modeling

The classic RL model and novel OL model described in the Introduction were applied to the Palombo et al. task. The RL model included six possible states, s, corresponding to the presentation of the six stimulus figures, and two possible actions, a, corresponding to "Yes" or "No" responses. The expected reinforcing value of each response given presentation of a specific stimulus figure, $V(s,a)$, was initially set to 0.5 and updated during subsequent trials $t + 1$ whenever that response was chosen with the following delta rule (Sutton & Barto, 2018):

$$\begin{cases} V(s,a)_0 = 0.5 \\ V(s,a)_{t+1} = V(s,a)_t + \alpha_{RL}\left(r_{t+1} - V(s,a)_t\right) \end{cases} \quad (5)$$

where α_{RL} is the updating parameter, which represents the extent to which feedback is used to update the expected values of options, varying between 0 (no updating) and 1 (extreme updating to the value of the previous trial); and r_{t+1} is the reinforcement outcome at trial $t + 1$, set to 1 if the participant's choice was correct on that trial (i.e., if the participant answered "Yes" and the stimulus figure won or if they answered "No" and the stimulus figure did not win) and to 0 if it was not correct. The probability of selecting the "Yes" response at trial $t + 1$ given the presentation of a specific stimulus figure s was computed using a softmax rule to account for nonsystematic behavior:

$$Prob_{RL}(a = \text{"Yes"}/s)_{t+1} = \frac{e^{\beta_{RL} V(s,a=\text{"Yes"})_t}}{e^{\beta_{RL} V(s,a=\text{"Yes"})_t} + e^{\beta_{RL} V(s,a=\text{"No"})_t}} \quad (6)$$

where β_{RL} is the exploit–explore parameter, with larger β_{RL} corresponding to more systematic choice of the option with greater expected value, and smaller β_{RL} to more random choice behavior.

In the OL model, instead of tracking values for state–action pairs, the probability $P_{ev}(s)$ that a stimulus figure s will win was modeled, independent of participant response. The probability of winning for each stimulus figure was initially set to 0.5 and updated during subsequent trials $t + 1$ using a delta rule similar to that of the RL model:

$$\begin{cases} P_{ev}(s)_0 = 0.5 \\ P_{ev}(s)_{t+1} = P_{ev}(s)_t + \alpha_{OL}\left(event_{t+1} - P_{ev}(s)_t\right) \end{cases} \quad (7)$$

where α_{OL} is the updating parameter and $event_{t+1}$ represents the outcome of the event on trial $t + 1$, set to 1 if the stimulus figure won (i.e., picture of a dollar bill) or to 0 if it did not win (i.e., gray rectangle labeled with "$0.00"). To account for nonsystematic behavior, the probability of selecting the "Yes" response at trial $t + 1$ given the presentation of a specific stimulus figure s was computed with the following softmax rule:

$$Prob_{OL}(a = \text{"Yes"}/s)_{t+1} = \frac{e^{\beta_{OL} P_{ev}(s)t}}{e^{\beta_{OL} P_{ev}(s)_t} + e^{\beta_{OL}(1-P_{ev}(s)t)}} \quad (8)$$

where β_{OL} is the exploit–explore parameter, and $P_{ev}(s)_t$ and $1 - P_{ev}(s)_t$ are the expected probability that the stimulus figure will win or not win, respectively, computed from information collected up to the previous trial.

The RL and OL models, each comprising a pair of α and β parameters, were fit to the Palombo et al. data set. Parameters were estimated for each model, each participant, and each set of stimuli separately with Bayesian inference, using an implementation of the affine invariant ensemble Markov Chain Monte Carlo sampler of Goodman and Weare (2010) proposed by Foreman-Mackey, Hogg, Lang, and Goodman (2013) and computed in MATLAB by Grinsted (2015). Flat priors were used for the two models, spanning the intervals (0–1) and (0–20) for the α and β parameters, respectively. Final estimates of parameters were selected as the maximum of the posterior distributions and were used to compute prediction error trial series for each participant, each set of stimuli, and each model. Model evidence was calculated as the marginal likelihood, that is, the likelihood of the data evaluated for each possible pair of α and β parameters, weighted by the prior, and integrated over the entire parameter space (e.g., Friel & Pettitt, 2008). This calculation was carried out using the MATLAB *integral2* function. The OL versus RL model comparison was then quantified by calculating a Bayes factor for each participant and each stimulus set as the quotient of the OL and RL marginal likelihoods. Natural logarithm of the Bayes factor (*logBF*) was calculated and signed so that positive *logBF* would correspond to a superior fit of the OL model and negative *logBF* to a superior fit of the RL model. Model comparison was also carried out for the OL and RL models separately compared with a one-parameter random response model. This model allowed for response bias, using a parameter varying between 0 and 1 to model participant's probability of responding yes to all trials regardless of the stimuli presented. Natural *logBF* was used to evaluate the fit of the OL and RL models separately compared with this basic model.

We also contemplated another method for comparing the OL and RL models, based on the construction of one large dual-process model (for similar modeling, see Collins & Frank, 2012). This model comprised five parameters: the two RL model parameters (α_{RL} and β_{RL}), the two OL parameters (α_{OL} and β_{OL}), and a weight parameter (w) representing the probability that behavior may be governed by RL (w = 0) or by OL (w = 1). Using this latter parameter for comparison of the OL and RL model was deemed to have limited reliability, both because of the

complexity of the model and the general challenge posed by parameter recovery in computational modeling (Wilson & Collins, 2019). The computation of Bayes factors, which is more robust because it involves integrating likelihood over the entire parameter space, was therefore used instead.

To explore whether differential use of one form of learning over the other had an impact on overall retention, logistic mixed effect modeling was carried out with postscan memory performance as dependent variable, *logBF* as fixed effect, and participant as random effect, allowing for different random intercepts for each participant. Because the computation of *logBF* was carried out separately for each set, only the postscan performance data that involved within-set choices was analyzed (i.e., nine trials per set and per participant). Significance of the fixed effect of *logBF* was evaluated with a t test using Satterthwaite's method, as implemented using the R *lme4* package (Bates, Mächler, Bolker, & Walker, 2015).

Finally, to assess learning progression over the course of the task, the OL and RL computational models were run again separately for Run 1 and Run 2 of each stimulus set. For each run, the following measures were computed: *logBF*, the average absolute value of the prediction errors, and the average response accuracy (defined as the proportion of responses corresponding to the majority status of a stimulus figure, e.g., responding "Yes, the man will win money" for a stimulus figure that wins money on 75% of the trials). A 2 × 2 within-subject ANOVA was carried out on each of these measures, including Run and Stimulus Set as independent variables. Reflecting progressive learning, we expected an increase in response accuracy and a decrease in absolute prediction error across runs. Comparison of *logBF* across stimulus set and run was carried out in an exploratory manner to examine the possibility of a shift in type of learning over the course of the task. Because these ancillary analyses are based on measures computed from only 48 trials (instead of 96 for the full stimulus set), their results should be interpreted with caution.

Brain Imaging Analyses

The preprocessed and preregistered BOLD activity time series of Palombo et al. were analyzed in a three-level whole-brain analysis using FMRIB's Improved Linear Model (Woolrich, Ripley, Brady, & Smith, 2001).

At the first level of analysis, general linear modeling (GLM) was carried out on each run and each participant. The GLM model comprised regressors for the mean brain activity during the experimental trials, control trials, and trials of no interest (i.e., inaccurate control trials and trials with no or late responses). The key contrast of interest compared brain activation during the experimental and control trials. Two parametric modulators that modeled the mean centered PE_{OL} and PE_{RL} were also included. To verify lack of correlation across trials between these modulators, linear mixed modeling was carried out that used participant as a random variable, PE_{OL} as the dependent variable, and PE_{RL} as fixed and random effects. Verification of the lack of correlation between PE_{OL} and PE_{RL} (see Results section) enabled their simultaneous inclusion into the GLM model. All trial regressors were modeled using trial onset times convolved with a double gamma hemodynamic response function, with duration comprising the entire trial, including decision and feedback phases (3.6 sec). As noted in the work of Palombo et al. (2019), separating decision and feedback would have required inclusion of a jittered delay between the two phases in the experimental design, which had been deemed undesirable. At the second level of analysis, the contrasts of interest (i.e., the mean activity during the experimental compared with the control trials and the parametric modulators) were combined for Runs 1 and 2 to compute the signal corresponding to the first stimulus set, and for Runs 3 and 4 to compute the signal corresponding to the second stimulus set for each participant.

At the third level, brain correlates of the differential fit of the OL versus RL model were explored by analyzing intra- and interindividual variability in *logBF* and in whole-brain functional signal. Mean trial activation in experimental versus control trials obtained for each participant and each stimulus set was modeled using a fixed effect GLM with 31 regressors: 1 regressor for modeling variability in *logBF* and 30 regressors for modeling mean brain activation in each participant (for similar models, see "Experimental Designs - Repeated measures" in the FMRIB Software Library wiki user guide, https://fsl.fmrib.ox.ac.uk/fsl/fslwiki/GLM). Two contrasts were computed, examining the mean trial activation in experimental versus control trials (Contrast 1) and its correlation with *logBF* (Contrast 2). The same model was used to determine the brain correlates of the OL and PE_{RL}s, using second-level contrasts as input. The third-level contrasts consisted of the mean prediction errors and their correlation with *logBF*. For all contrasts, the cluster-defining (or voxel-wise) threshold was set to $z = 3.09$ ($p = .001$) and the corrected cluster significance threshold to $p = .05$ (Eklund, Nichols, & Knutsson, 2016).

To further test the hypothesis that superior fit of the OL model over the RL model involves increased hippocampal activation and decreased ventral striatal activation, linear mixed modeling was conducted, with participant as random factor on brain signal (experimental trials – control trials) averaged over the right and left hippocampus and over the right and left nucleus accumbens. ROIs were defined using lateralized masks based on the Harvard–Oxford structural atlas (e.g., Desikan et al., 2006) with a threshold of 50%. For both analyses, the model used participant as a random factor and was defined as follows:

$$ROI_{ij} = \beta_{0i} + \beta_1 StimulusSet_j + \beta_2 BrainSide + \beta_3 logBF \quad (9)$$

where ROI_{ij} represents brain activation averaged over lateralized regions of the hippocampus or nucleus

accumbens for participant *i* and stimulus set *j*, $StimulusSet_j$ represents the stimulus set with possible values 1 or 2 corresponding to administration order, and *BrainSide* represents left or right brain lateralization (left used as reference). The model included fixed and random intercepts, β_{0i}, and fixed effect regression coefficients β_1, β_2, and β_3. Model fit was carried out using maximum likelihood as implemented in the *lme4* package (Bates et al., 2015) of R (R Core Team, 2019). Model fit was tested against a null model that was similar but did not include the effect of *logBF* and against a more complex model that also included a *logBF* × *BrainSide* interaction. Model fit was evaluated using the Akaike's Information Criterion (AIC; Akaike, 1974) and Bayesian Information Criterion (BIC; Schwarz, 1978). R^2 effect size was estimated using the method developed by Nakagawa and Schielzeth (2013), implemented with the *piecewiseSEM* R package. Model comparison was carried out using a Likelihood Ratio Test with χ^2 distribution. Significance of the fixed effects was evaluated with a *t* test using Satterthwaite's method, as implemented using the R *lme4* package (Bates et al., 2015).

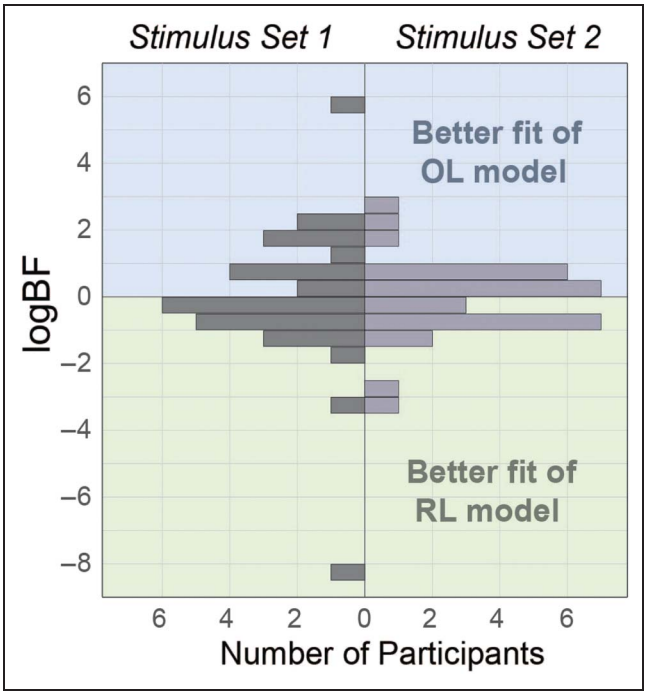

Figure 2. Distribution of *logBF* (i.e., the logarithm of the Bayesian Factor) obtained for stimulus Sets 1 and 2. Positive *logBF* values correspond to a better fit of the OL model over the RL model, and negative *logBF* values to a better fit of the RL model over the OL model.

RESULTS

Differential Use of OL and RL

Modeling

Log-Bayes factors that compared the OL and RL models separately to the random response model with bias were significantly greater than 0 (OL vs. random: $logBF = 5.85$, $SD = 7.79$, $t(29) = 4.11$, $p < .001$; RL vs. random: $logBF = 5.87$, $SD = 7.60$, $t(29) = 4.23$, $p < .001$), with effect size suggesting "decisively better fit" (Jeffreys, 1961) of each learning model compared with the no learning model. The log-Bayes factors computed to compare the OL and RL models to each other were not significantly different from zero (Set 1: $logBF = -0.025$, $SD = 2.25$, $t(29) = -0.06$, $p = .952$; Set 2: $logBF = -0.017$, $SD = 1.27$, $t(29) = -0.07$, $p = .942$), indicating an equivalent overall fit of the OL and RL models (Figure 2). Interestingly, *logBF* signs and values displayed considerable within-subject variability across stimulus sets, as indicated by a change in the sign of *logBF* in 43% of participants and by a nonsignificant Spearman rank correlation of *logBF* values across sets ($\rho = 0.063$, $p = .739$). These findings suggest that a participant who differentially used one learning process when presented with the first stimulus set did not necessarily display that same pattern when presented with the second set. Within the distribution of *logBF* of Set 1, there were two outlier data points with values that were more than 2 *SD*s from the mean. They were removed from all subsequent calculations so that they would not disproportionally impact correlational findings. Analysis of postscan performance data using logistic mixed effect modeling did not reveal any significant effect of *logBF* ($\beta = -0.037$, $SE = 0.104$, $z = -0.358$, $p = .721$), suggesting equivalently effective memory for stimulus–outcome contingencies regardless of differential preference for the observation versus RL strategy.

Results of the ancillary analyses carried out on Run 1 and Run 2 of each stimulus set confirmed that learning indeed took place over the course of the task: Results of the within-subject ANOVA on response accuracy indicated the presence of a main effect of Run, $F(1, 29) = 5.22$, $p = .030$, $\eta^2 = .152$, with accuracy increasing modestly between Run 1 ($M = 0.611$, $SD = 0.080$) and Run 2 ($M = 0.643$, $SD = 0.099$). The fact that response accuracy did not reach 0.75 during the second run suggests that most participants were still learning and that the task was reasonably difficult, likely because of a combination of the large number of stimulus figures (six per set), the complexity of the visual fractal patterns that distinguished them, and their small number of presentations (eight per run). There was no significant effect of Stimulus Set on response accuracy, $F(1, 29) = 3.89$, $p = .058$, $\eta^2 = .118$, and no Stimulus Set × Run interaction, $F(1, 29) = 0.88$, $p = .356$, $\eta^2 = .029$. The results of the within-subject ANOVA on *logBF* indicated no significant main effect of Run, $F(1, 29) = 0.030$, $p = .863$, $\eta^2 = .001$; Stimulus Set, $F(1, 29) = 0.013$, $p = .911$, $\eta^2 < .001$; or Stimulus Set × Run interaction, $F(1, 29) = 0.534$, $p = .471$, $\eta^2 = .018$, providing no evidence for a systematic shift in type of learning across runs.

Brain Correlates

As expected, results of the GLM examining whole-brain correlates of intra- and interindividual variability in *logBF*

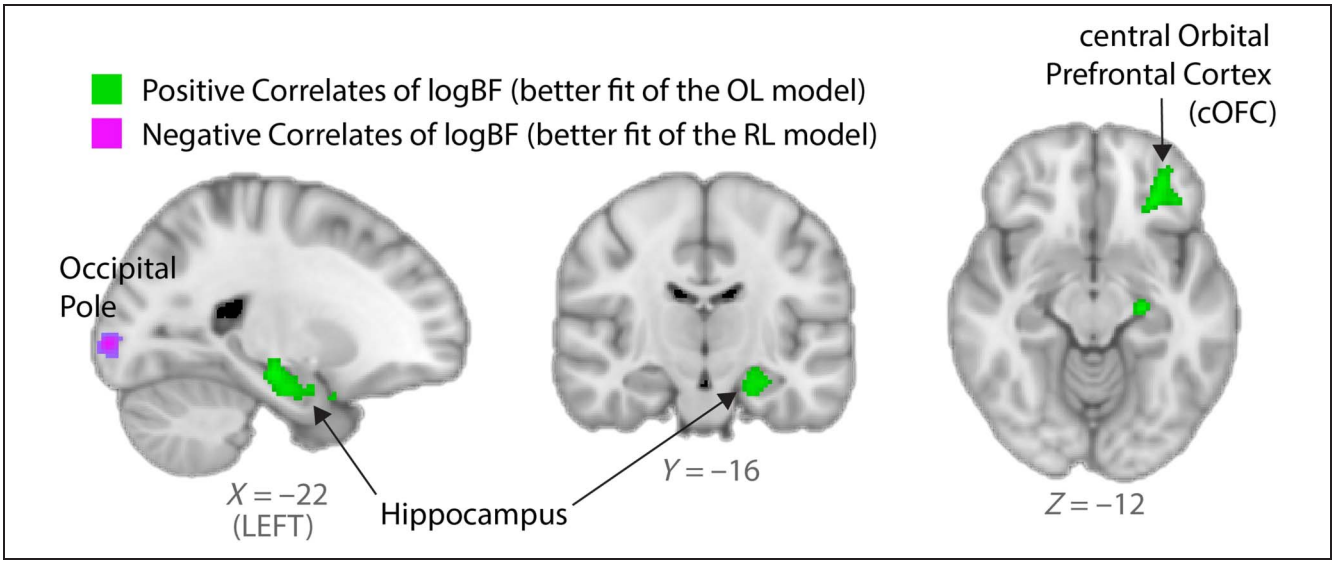

Figure 3. Whole-brain correlates of intra- and interindividual variability in *logBF*, with increased *logBF* indicated in green corresponding to a better fit of the OL model over the RL model, and decreased *logBF* indicated in purple corresponding to a better fit of the RL model over the OL model. (Voxel-wise threshold $p < .001$; cluster-wise threshold $p < .05$).

revealed increased activation in the hippocampus (left) with more positive *logBF* (i.e., better fit of the OL model over the RL model; see green area in Figure 3). *logBF* was also found to positively correlate with activation in one other brain area, the left central orbitofrontal cortex (cOFC). Against expectation, *logBF* (i.e., better fit of the RL model over the OL model) did not negatively correlate with increased striatal activation but correlated instead with activation in the bilateral occipital poles (see purple area in Figure 3).

Contrasts corresponding to areas of general mean activation (orange) and deactivation (blue-green) during the learning trials compared with the control trials are presented for additional information in Figure 4. Areas of relative activation included the bilateral caudate nucleus, thalamus, midbrain, frontal pole, an area at the junction of the frontal orbital cortex, insular cortex and frontal operculum, the middle frontal gyrus, an area at the junction of the superior parietal lobule and angular gyrus, the dorsal posterior precuneus, and a large area encompassing the temporal fusiform cortex, occipital fusiform cortex, and occipital pole. Areas of relative deactivation implicated the bilateral hippocampus; the central orbital pFC; the middle and posterior sections of the dorsal cingulate gyrus; the ventral precuneus; an area at the junction of the frontal pole, subcallosal cortex, and paracingulate gyrus; another area at the junction of the precentral gyrus and superior parietal lobule; and a large area covering the insular cortex, central opercular cortex, anterior middle temporal gyrus, anterior and posterior superior temporal gyrus, angular gyrus, supramarginal gyrus, and superior division of the lateral occipital cortex. Interestingly,

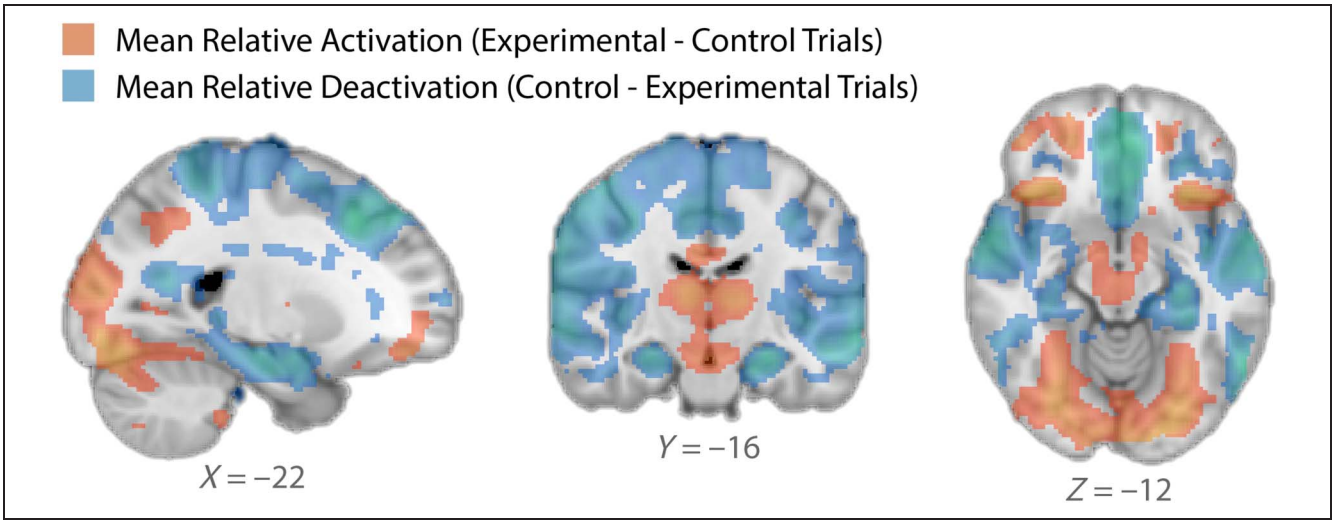

Figure 4. Areas of relative mean activation (orange) and deactivation (blue-green) during the learning trials compared with the control trials. (Voxel-wise threshold $p < .001$; cluster-wise threshold $p < .05$).

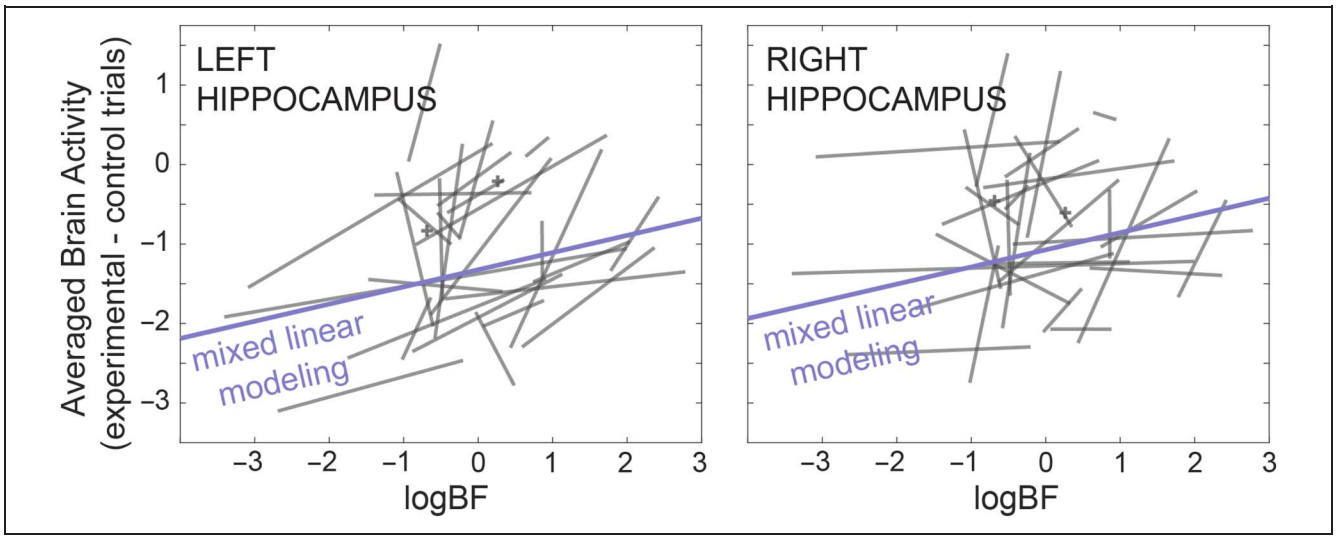

Figure 5. Relation between *logBF* and brain activity averaged over the left and right hippocampus (learning trials minus control trials). Gray solid lines represent pairs of data points obtained for each participant (except for the two participants who had one outlier datapoint, in which case only a small cross is displayed). The purple solid lines represent the linear mixed modeling results.

positive correlates of *logBF* were located in areas of relative mean deactivation and negative correlates of *logBF* in areas of relative mean activation.

Consistent with the whole-brain findings, linear mixed modeling of the brain activation averaged over the hippocampus ROI as a function of *StimulusSet*, *BrainSide*, and *logBF* revealed a model fit (AIC = 285.5, BIC = 302.0, R^2 = .59) that was significantly better than the same model without *logBF* (AIC = 295.4, BIC = 309.1, R^2 = .49, $\chi^2(1)$ = 11.9, $p < .001$). There was a significant fixed effect of *logBF* ($\beta = 0.216$, $SE = 0.060$, $t(107.7) = 3.60$, $p < .001$), indicating increased hippocampal activation with increasingly better fit of the OL model compared with the RL model (Figure 5). The fixed effect of *BrainSide* was also significant, suggesting overall greater activation in the right hippocampus ($\beta = 0.251$, $SE = 0.117$, $t(86.1) = 2.15$, $p = .035$). The effect of *StimulusSet* was not significant ($\beta = 0.187$, $SE = 0.119$, $t(87.5) = 1.58$, $p = .117$). Examination of a model that also included a *logBF* × *BrainSide* interaction did not yield significantly better fit (AIC = 285.7, BIC = 304.9, R^2 = .60, $\chi^2(1) = 1.84$, $p = .175$), and the interaction term was not significant ($\beta = -0.126$, $SE = 0.093$, $t(86.1) = -1.36$, $p = .176$). Thus, although the whole-brain GLM analysis evidenced brain correlates of *logBF* only in the left hippocampus, the linear mixed modeling analysis suggests the presence of a similar but subthreshold signal also in the right hippocampus.

Linear mixed modeling of the brain activation averaged over the nucleus accumbens with regressors *StimulusSet*, *BrainSide*, and *logBF* had a model fit (AIC = 351.4, BIC = 367.9, R^2 = .28) that was not better than the model without *logBF* (AIC = 349.6, BIC = 363.3, R^2 = .28, $\chi^2(1)$ = 0.21, $p = .645$). There was no significant fixed effect of *logBF* ($\beta = 0.038$, $SE = 0.082$, $t(116.0) = 0.47$, $p = .642$); *BrainSide* ($\beta = 0.003$, $SE = 0.173$, $t(86.4) = 0.015$, $p = .988$) ; or *StimulusSet* ($\beta = -0.199$, $SE = 0.174$, $t(88.8) = -1.14$, $p = .256$). Inclusion of the *logBF* × *BrainSide* interaction into the model did not yield better fit (AIC = 352.7, BIC = 372.0, R^2 = .29, $\chi^2(1) = 0.69$, $p = .408$), and the interaction term was not significant ($\beta = 0.114$, $SE = 0.138$, $t(86.4) = 0.83$, $p = .409$).

Prediction Errors Associated with OL and RL

Modeling

As expected, linear mixed modeling confirmed an absence of significant relation between the PE_{OL} and PE_{RL} ($\beta = -0.003$, $SE = 0.030$, $t(29.9) = -0.104$, $p = .918$; see Figure 6). Of note, a significant positive relation is

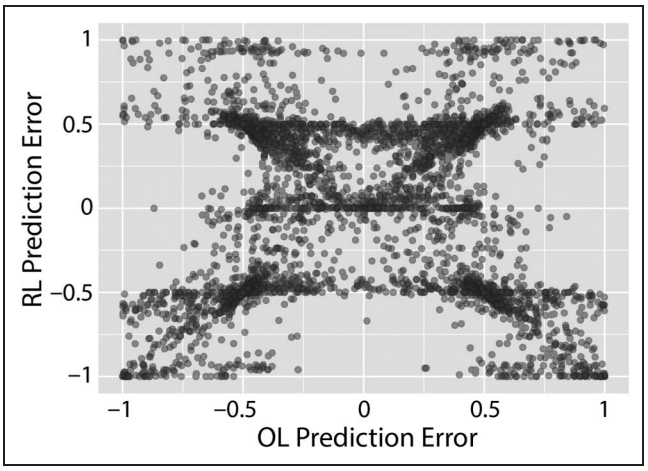

Figure 6. Illustration of the relation between PE_{OL} and PE_{RL} with data points presented for all participants and all trials. These data points are not independent and are presented here for illustrative purpose only (see linear mixed modeling statistics in the text). Results confirmed a lack of relation between the signed signals. A strong relation was, however, found between their absolute values, reflecting general effects of surprise and learning over time.

apparent in Figure 6 between the absolute values of the prediction errors, $|PE_{OL}|$ and $|PE_{RL}|$ (linear mixed modeling: $\beta = 0.505$, $SE = 0.012$, $t(122.6) = 41.5$, $p < .001$). This relation is consistent with progressive learning occurring concomitantly in both models, with overall amount of surprise decreasing similarly over time. Indeed, the absolute values of the OL and PE_{RL}s were both found to decrease significantly between Run 1 and Run 2 of each stimulus set. Results of the within-subject ANOVA on $|PE_{OL}|$ revealed a large main effect of Run ($M_{|PEOL1|} = 0.465$, $M_{|PEOL2|} = 0.431$, $F(1, 29) = 76.2$, $p < .001$, $\eta^2 = .724$), but no significant effect of Stimulus Set, $F(1, 29) = 1.04$, $p = .317$, $\eta^2 = .035$, or Stimulus Set × Run interaction, $F(1, 29) = 0.122$, $p = .730$, $\eta^2 = .004$. Similarly, results of the within-subject ANOVA on $|PE_{RL}|$ revealed a large main effect of Run ($M_{|PERL1|} = 0.477$, $M_{|PERL2|} = 0.445$, $F(1, 29) = 42.1$, $p < .001$, $\eta^2 = .592$, but no significant effects of Stimulus Set, $F(1, 29) = 1.03$, $p = .319$, $\eta^2 = .034$, or Stimulus Set × Run interaction, $F(1, 29) = 3.55$, $p = .070$, $\eta^2 = .109$.

Brain Correlates

Results of the GLM applied to the prediction error for the observation-based model suggested positive correlates of the prediction error ($PE_{OL}+$)—that is, trials when the stimulus figure wins more than predicted through mental representation—that involved the occipito-temporal ventral stream pathway, including the bilateral temporal-occipital fusiform gyrus, occipital fusiform gyrus, inferior lateral occipital cortex, right lingual cortex, and right occipital pole (yellow area in Figure 7). To ensure that this brain activation correlate was not simply because of differences in visual display between the winning (dollar bill) and losing (dark box labeled with $0) stimulus figure outcomes, the analysis was repeated limited to trials with winning outcomes. Although smaller, brain activation correlates of $PE_{OL}+$ were found again in this analysis, overlapping with the previous findings (orange area in Figure 7). These results confirmed brain signals corresponding to a discrepancy between expectations and outcomes rather than to simple differences in visual input. There was no detected activation for negative correlates of the $PE_{OL}-$, corresponding to trials when the stimulus figure wins less than predicted. There was also no detected area of enhanced $PE_{OL}+$ or $PE_{OL}-$ signals with increasing or decreasing *logBF*.

GLM results applied to the reinforcement-based prediction error suggested positive correlates of the PE_{RL} ($PE_{RL}+$), corresponding to responses that were reinforced more

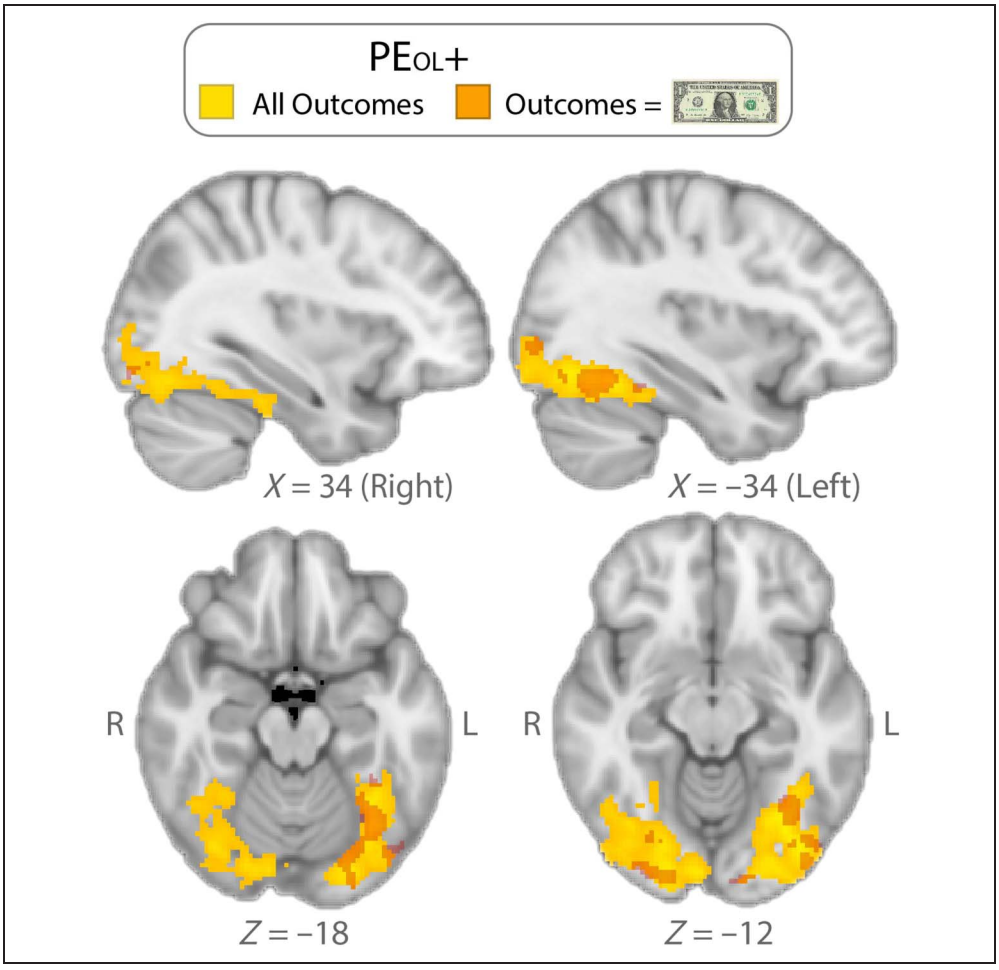

Figure 7. Whole-brain $PE_{OL}+$ for the observation-based model ($PE_{OL}+$), corresponding to trials when the stimulus figure wins more than predicted through mental representation, are shown in yellow. These areas involved the occipito-temporal ventral stream pathway, including the bilateral temporal-occipital fusiform gyrus, occipital fusiform gyrus, inferior lateral occipital cortex, right lingual cortex, and right occipital pole. Areas in orange represent the same construct when the analysis was restricted to trials resulting in observed outcomes featuring a picture of a dollar sign. There was no detected activation correlating with more negative prediction error ($PE_{OL}-$), corresponding to trials when the stimulus figure wins less than predicted. There was also no detected area of enhanced $PE_{OL}+$ or $PE_{OL}-$ related activation with increasing or decreasing *logBF*. (Voxel-wise threshold $p < .001$; cluster-wise threshold $p < .05$).

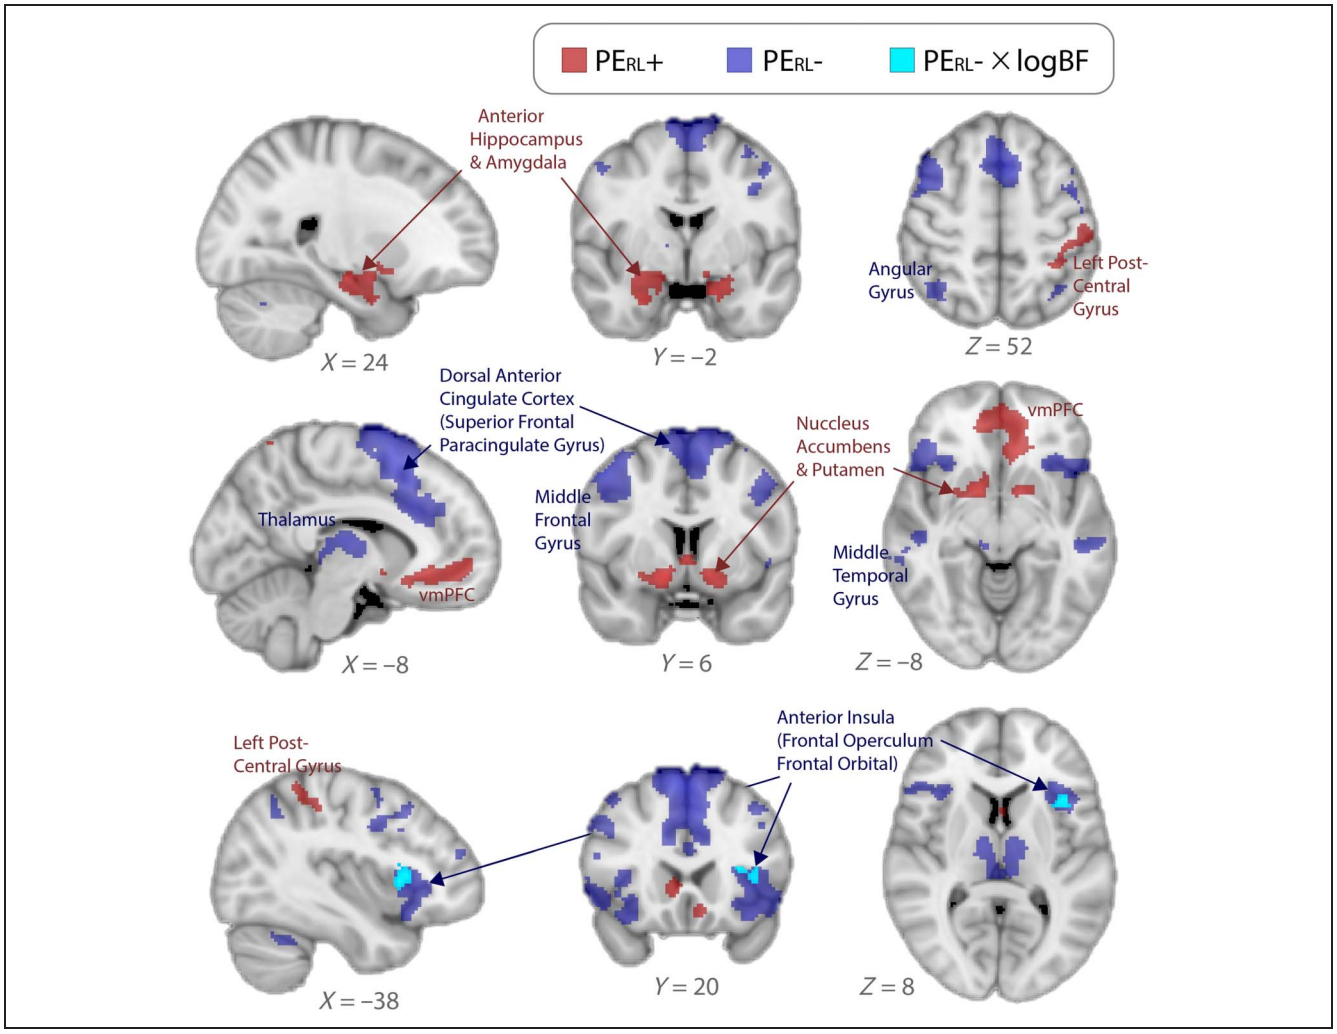

Figure 8. Whole-brain positive correlates (PE_{RL}+, red) and negative correlates (PE_{RL}−, dark blue) of the PE_{RL}, corresponding to trials that were reinforced more and less, respectively, than predicted through mental representation. Areas of increased PE_{RL}− related activation with increased *logBF* are shown in pale blue. There was no detected area of enhanced PE_{RL}+ related activation with increasing *logBF* and no detected area of enhanced PE_{RL}− related activation with decreasing *logBF*. (Voxel-wise threshold $p < .001$; cluster-wise threshold $p < .05$).

than predicted by current mental computations of the options' values that involved the bilateral ventral striatum, including the nucleus accumbens and putamen, bilateral ventromedial prefrontal cortex (vMPFC), bilateral anterior hippocampus, bilateral amygdala, and left posterior central gyrus (red area in Figure 8). Negative correlates of the PE_{RL}s (PE_{RL}−), which correspond to responses that are reinforced less (or are more wrong) than mentally predicted, involved bilateral activation in the anterior insula (frontal orbital/frontal operculum), dorsal anterior cingulate cortex (superior frontal/paracingulate gyrus), middle frontal gyrus, middle temporal gyrus, angular gyrus, and thalamus (dark blue area in Figure 7). One area of enhanced PE_{RL}− signal with increased *logBF* was detected in the left anterior insula (pale blue area in Figure 8). There was no area of enhanced PE_{RL}+ signal with increased *logBF* and no area of enhanced PE_{RL}− signal with decreasing *logBF*.[1]

DISCUSSION

This study examined whether the role of the hippocampus demonstrated in some probabilistic reinforcement learning tasks may be attributed to observation-based learning processes. A computational model of observation-based learning was constructed, mirroring classic models of reinforcement-based learning, and was applied to the neuroimaging data set of Palombo et al. (2019). The processing of separate and possibly concomitant OL and RL learning mechanisms was supported by uncorrelated OL and PE_{RL}s and by intra- and interindividual variability in the preferential use of the OL versus RL model, with no model systematically winning over the other, neither in terms of overall model fit, nor in terms of relation with postscan performance.

As predicted, differential fit of the OL model over the RL model correlated with increased mean activity in the hippocampus. The latter finding is consistent with

hippocampal implication in tasks involving OL with uncertain outcomes (Kok & Turk-Browne, 2018; Boorman et al., 2016; Schapiro et al., 2012; Bunzeck et al., 2010; Turk-Browne et al., 2009), extending it to a simple probabilistic feedback learning task typically thought to reflect reinforcement learning. Of note, the hippocampus was found to be a zone of relative mean deactivation during probabilistic feedback learning trials compared with control trials. Such deactivation has been noted before during reinforcement learning tasks (Poldrack et al., 2001). Thus, to be more precise, the greater use of OL processes corresponded to less deactivation of the hippocampus during learning trials and the differential use of RL processes to more deactivation.

Although individual differences in mean hippocampal activation as a function of differential model fit support involvement of the hippocampus in OL, the mechanisms underlying that involvement remain unclear. Indeed, against expectations, trial-wise hippocampal activation did not correlate with the trial-wise OL prediction error, suggesting that the mechanism of hippocampal involvement in OL is not the computation of discrepancy between expected and actual observed outcome. Although this study was not designed to evaluate other potential mechanisms of involvement of the hippocampus in OL, we note that our results are compatible with the findings of Boorman et al. (2016), which implicate the hippocampus in computing the strength of stimulus–outcome probabilistic associations. Boorman et al. demonstrated such a role using "suppression blocks," involving the presentation of stimuli in a pseudorandom order, hence "suppressing" the probabilistic contingencies that were learned in previous blocks. They found that hippocampal suppression varied as a function of stimulus–outcome association strength derived via computational modeling during the learning trials. Also consistent with a mechanism similar to Boorman et al., another key region associated with differential fit of the OL over RL model was the cOFC. This region was also highlighted in the work of Boorman et al. as critical to learning outcome type during probabilistic feedback learning. Specifically, whereas the hippocampus was shown to be involved in computing the strength of stimulus–outcome associations, the cOFC was implicated in updating these associations in a goal-directed manner.

Hippocampal involvement was not limited to observation-based learning but was also evident in reinforcement-based learning through examination of the RL prediction error. Specifically, a significant relation was found between more $PE_{RL}+$ and brain activation in an area centered on the bilateral anterior hippocampus and amygdala. This finding is consistent with previous reports of hippocampal activation correlating with the $PE_{RL}+$ in simple probabilistic reinforcement learning tasks (Davidow et al., 2016; Dickerson et al., 2011; Foerde & Shohamy, 2011; see also supplemental material in the work of Schonberg et al., 2010). Other brain correlates of $PE_{RL}+$ included areas of dopamine projections in the ventral striatum, consistent with well-established literature (e.g., Jocham, Klein, & Ullsperger, 2011; den Ouden, Daunizeau, Roiser, Friston, & Stephan, 2010; Seymour et al., 2004; McClure, Berns, & Montague, 2003; O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003), and in the vMPFC, a brain region implicated in computing the subjective valuation of choices (Kable & Glimcher, 2009). The vMPFC has also been related to prediction error signals in macaques (Matsumoto, Matsumoto, Abe, & Tanaka, 2007) and with reinforcement prediction error when someone else makes the decisions (Burke, Tobler, Baddeley, & Schultz, 2010).

Although the $PE_{RL}+$ correlates confirmed direct implication of the hippocampus in the computation of discrepancy between actual and expected reinforcing outcomes, the specific hippocampal contribution to this computation remains to be elucidated, possibly involving interactions between OL and RL processes. As proposed by Boorman et al. (2016), reinforcement-based signal from the striatum may be fed into the hippocampus and contribute to updating computation of the strength of stimulus–outcome probabilistic associations. By this view, the RL prediction error signal in the hippocampus would serve as input to observation-based learning. This proposal is compatible with recent findings that the strength of the RL prediction error in the ventral striatum predicts subsequent episodic memory retrieval accuracy (Calderon et al., 2021; Ergo, De Loof, & Verguts, 2020), and with findings of a stronger relation between reinforcement learning and episodic memory performance with increased hippocampus–striatum functional connectivity (Davidow et al., 2016). It is also consistent with previous accounts of a collaboration between the mesolimbic dopamine system and the hippocampus, resulting in enhanced memory for events with motivational relevance (Shohamy & Adcock, 2010).

A possible collaborative relation between the OL and RL systems could also be taking place in the reverse direction. The hippocampus has been suggested to have a role in mediating inferences and expectancies based on the probabilistic structure of observed events (Harrison et al., 2006; Eichenbaum, Dudchenko, Wood, Shapiro, & Tanila, 1999). Thus, as in model-based learning (Gershman & Daw, 2017; Johnson et al., 2007), the hippocampus could contribute through constructing OL-based predictive representations that are then fed into the ventral striatum RL system. This proposal is compatible with demonstrations that other types of hippocampal representations (i.e., conjunctive associations) can be fed into the striatal-based learning system (Ballard, Wagner, & McClure, 2019; Duncan, Doll, Daw, & Shohamy, 2018), and with findings of a relation between hippocampal activation and the accuracy of information learned via reinforcement-based feedback (Dickerson & Delgado, 2015).

In addition to the hippocampus, other brain areas may also be involved in a possible collaboration between the OL and RL systems, as suggested by examination of the $PE_{RL}-$. Consistent with previous literature (e.g., Hauser, Iannaccone, Walitza, Brandeis, & Brem, 2015; Meder,

Madsen, Hulme, & Siebner, 2016; Garrison, Erdeniz, & Done, 2013; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006), $PE_{RL}-$ implicated the error monitoring or salience network, especially the anterior insula and dorsal anterior cingulate, as well as areas in the middle frontal gyrus, middle temporal gyrus, angular gyrus, and thalamus. Interestingly, the $PE_{RL}-$ signal in the left anterior insula was found to be amplified with differential use of OL over RL learning, suggesting that the more one uses OL processes, the stronger their neural signal in response to errors. Individual differences in people's propensity to learn from their errors versus from reinforcing outcomes have previously been demonstrated (e.g., Frank, Woroch, & Curran, 2005). The current results suggest that those who learn more from their errors may also tend to differentially use OL strategies, and error monitoring may be a key mechanism enabling interactions between OL-based learning and action outcomes.

Rather than reflecting collaboration between OL and RL, could the current results reflect a shift from one type of learning to the other over the course of the task? Such a view would accord with findings of a transition from early hippocampal involvement to later basal ganglia involvement during feedback-based paired associate learning (Poldrack et al., 2001). Our behavioral findings provided no evidence of a shift from OL to RL (or vice versa), as *logBF* values did not differ across the two runs. However, average response accuracy in Run 2 did not reach 0.75, and it is likely that participants were still learning during the second half of the task, perhaps lacking time for displaying such a shift. As mentioned before, the computations of *logBF* for each run separately were likely imprecise as they were based on data from 48 test trials, including only eight presentations per stimulus figure. Further research is thus necessary to examine more decisively the presence of a systematic shift between OL and RL during probabilistic feedback learning, with a study design perhaps involving more trials and runs to learn the contingencies of each stimulus set.

Contrary to our prediction, we did not find increased mean trial activity in the striatum with the greater use of RL over OL. This absence of finding is consistent with previous reports that ventral striatal activity does not necessarily correlate with mean trial activity during RL tasks but rather positively correlates with the RL prediction error ($PE_{RL}+$; Jocham et al., 2011; McClure et al., 2003; O'Doherty et al., 2003; Schultz, 1998). This well-established ventral striatal correlate of the RL prediction error was also observed in the current study, but no signal enhancement was found with preferential use of RL over OL. These results may suggest that RL-related activity in the striatum tends to be equivalent within and across participants, and variability primarily comes into play in the extent to which OL strategy is employed to augment a basic RL strategy. Of note, differential use of OL over RL (or vice versa) during learning of the first stimulus set did not necessarily yield the same pattern during the second set. These findings argue against the presence of systematic differences in learning styles across individuals when they perform probabilistic feedback learning, but instead indicate that OL and RL are processes that are both readily available during this type of task.

Interestingly, although greater use of RL over OL was not associated with increased mean trial activation in the striatum, such a pattern was observed in the occipital poles. That is, one effective way to experience reinforcement in the current task may have involved holding in mind a mental image of one outcome (i.e., the dollar bill) for an effective match or nonmatch visual signal to be generated at outcome presentation. Support for this interpretation was provided by ancillary analyses where we examined the RL prediction errors separately for trials with winning and nonwinning outcomes (see Figure 9). In these analyses, occipital pole activation was found to correlate with $PE_{RL}+$ but not $PE_{RL}-$ during trials with winning outcomes, and with $PE_{RL}-$ but not $PE_{RL}+$ during trials with nonwinning outcomes. Activation in the occipital poles thus related to amount of discrepancy between actual and predicted outcome regardless of correct status whenever the "dollar bill" outcome was mentally expected, suggesting operation of a visual–perceptual prediction error similar to signals previously reported in the visual cortex and visual ventral stream (Alink, Schwiedrzik, Kohler, Singer, & Muckli, 2010; den Ouden, Friston, Daw, McIntosh, & Stephan, 2009; Turk-Browne et al., 2009). These findings support the use of predictive mental visual imagery as a means to subserve RL processes in the current task; however, because post hoc, this interpretation must remain tentative.

In addition to examining hippocampal contributions to probabilistic feedback learning, an important contribution of the current study concerns the construction of the OL model, mirroring classic RL modeling and enabling direct examination of the brain correlates of the OL prediction error. In particular, the OL prediction error implicated regions along the occipito-temporal ventral stream pathway, and specifically including the bilateral temporal-occipital fusiform gyrus, occipital fusiform gyrus, inferior lateral occipital cortex, right lingual cortex, and right occipital pole. Similar results were found when restricting the data set to observed winning outcomes (i.e., the dollar bill), confirming attribution to prediction error signaling rather to simple differences in visual input across winning and nonwinning outcomes. This pattern of brain activation is consistent with documented neural correlates of top–down "perceptual prediction errors," where predictions concern the occurrence of perceptual stimuli. Specifically, perceptual prediction errors have been demonstrated to include neural structures along the ventral visual stream, from the primary visual cortex to the inferotemporal cortex and then hippocampus, with more upstream structures involved with increasing complexity of perceptual representations (see the work of den Ouden et al., 2010, for a review). In the present work, the OL prediction error did not implicate the

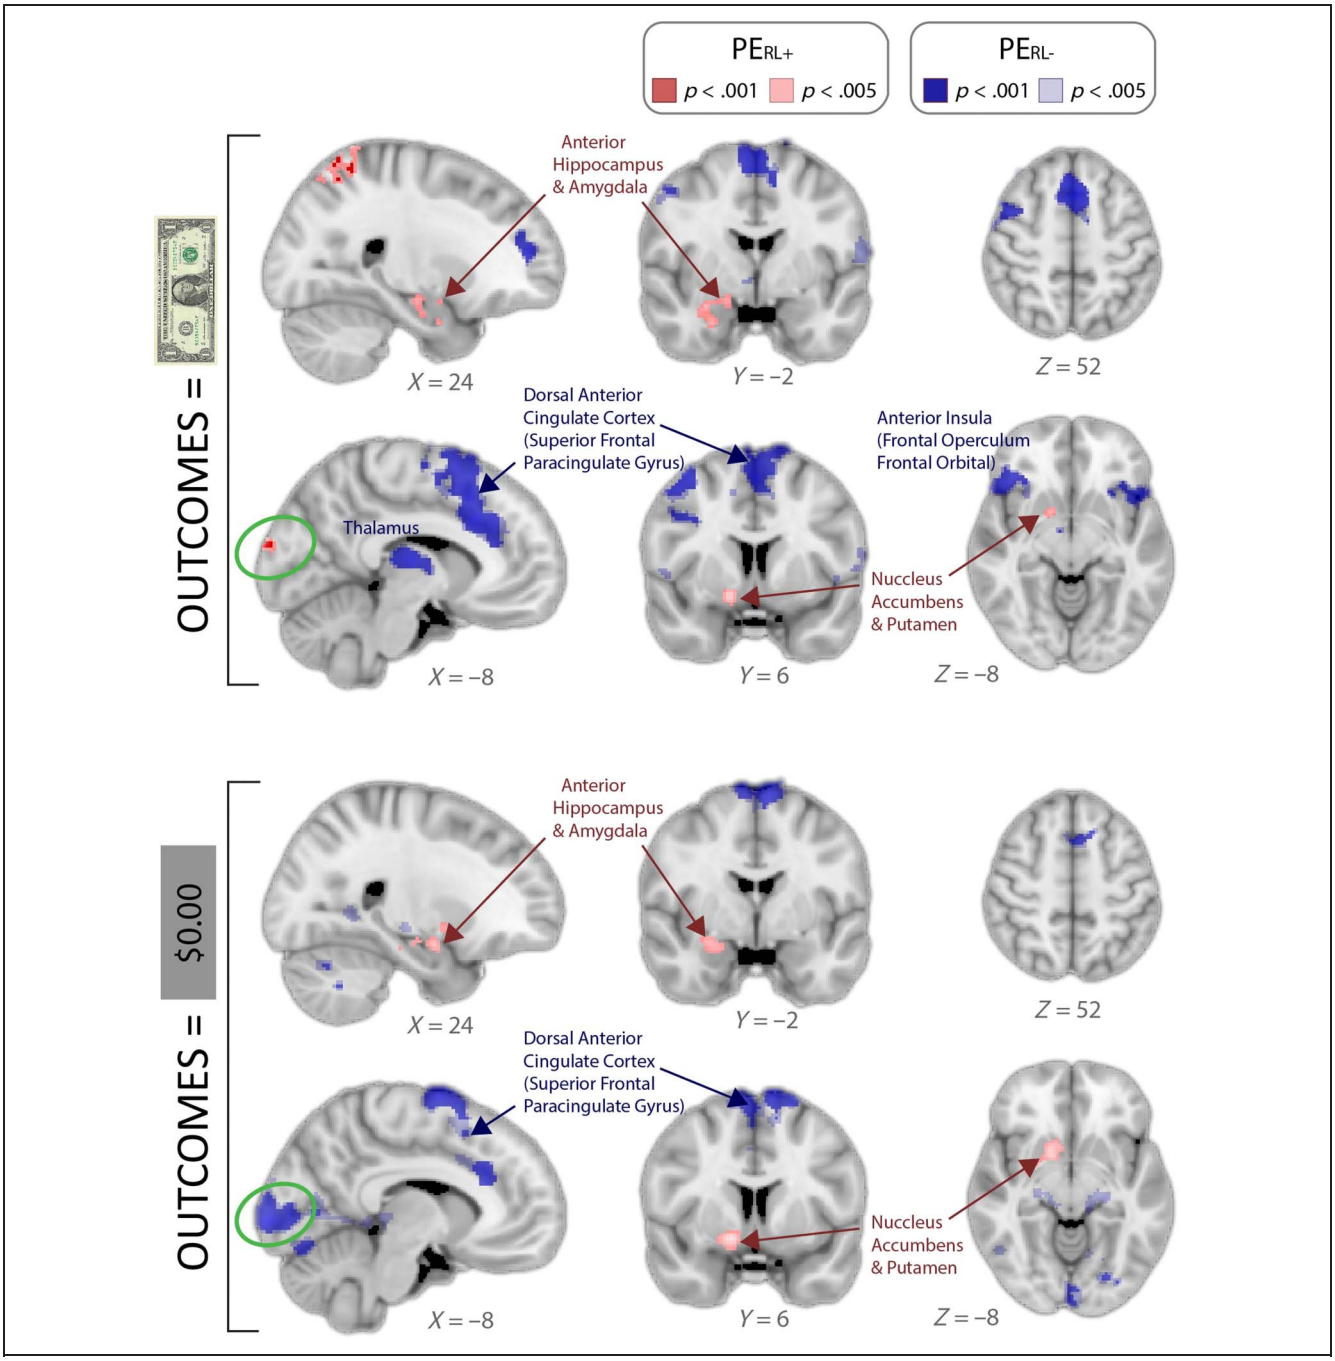

Figure 9. Ancillary analyses examining the correlates of $PE_{RL}+$ and $PE_{RL}-$ obtained when considering separately the observed outcomes featuring a picture of a dollar bill and the observed outcomes featuring a gray box labeled $0.00. (Voxel-wise threshold: $p < .001$ and $p < .005$, see legend; cluster-wise threshold: $p < .05$). The green circles highlight occipital pole activation that correlated with PERL+ but not PERL− during trials with winning outcomes, and with PERL− but not PERL+ during trials with non-winning outcomes. Occipital pole activation thus related to the amount of discrepancy between actual and predicted outcome regardless of correct status whenever the "dollar bill" outcome was mentally expected.

hippocampus and involved structures up to the fusiform gyrus. These results may suggest involvement of perceptual representations of moderate complexity in the current task, perhaps not complex enough to require hippocampal recruitment in terms of construction of perceptual representations. Interestingly, unlike for the OL prediction error, there was no separate area of brain activation that was negatively associated with the OL prediction error ($PE_{OL}-$). This finding suggests the presence of a single network for coding discrepancy between actual and expected observed outcomes, with a continuum of activation signed in the direction of one particular outcome (i.e., the stimulus figure winning). Such a pattern is consistent with prior evidence in studies involving probabilistic outcomes of neural signals signed in the direction of increasing outcome level (e.g., monetary amounts or novelty in the work of Bunzeck et al., 2010) or outcome goal (Boorman et al., 2016).

Of note, the pattern of brain activation associated with the OL prediction error did not overlap with the neural correlates of the state prediction error described in model-based RL, which comprised the intraparietal sulcus and lateral pFC (Gläscher, Daw, Dayan, & O'Doherty, 2010). This study used a task conducive to model-based learning in which each decision yielded probabilistic transition between two situations (or states) and reinforcement was given at the last step of that series of states. The state prediction error in Gläscher et al. thus referred to the discrepancy between modeled and observed situational outcomes. Although the prediction errors across the work of Gläscher et al. and present studies both involved observed outcomes, the identity and probability of these outcomes were dependent upon actions in the work of Gläscher et al., but were independent from actions in the current task. This important difference across paradigms may explain the lack of overlap in the brain imaging correlates of the prediction errors.

Conclusions

This study examined whether the role of the hippocampus in probabilistic reinforcement learning tasks may be attributed to observation-based processes. A computational model of observation-based learning (OL) was constructed, mirroring classic models of reinforcement-based learning (RL), and was applied to the neuroimaging data set of Palombo et al. (2019). Consistent with our prediction, model fit suggested that OL processes may indeed take place concomitantly with reinforcement learning, with differential use of OL involving activation of the hippocampus as well as of the cOFC. However, contrary to predictions, striatal activation did not track with differential use of RL over OL. Furthermore, hippocampal activation did not scale with the OL prediction error but scaled instead with the striatal RL prediction error. Taken together, these findings suggest a role for the hippocampus in probabilistic feedback learning, possibly through collaboration between the systems that mediate OL and RL. In particular, the hippocampus may be involved in encoding the strength of observed stimulus–outcome associations, with updating of these associations through striatal reinforcement-based computations.

Acknowledgments

M. V. was supported by a Senior Research Career Scientist award and M. E. by a Merit Review Award (I01CX001653), both from the Clinical Science Research and Development Service of the Department of Veterans Affairs. The contents of this article do not represent the views of the U.S. Department of Veterans Affairs or the United States Government. The authors have no conflicts of interest to report.

Reprint requests should be sent to Virginie M. Patt, VA Boston Health Care System Jamaica Plain Campus, 150 S. Huntington Avenue, Boston MA, 02130, or via e-mail: virginie.patt@va.gov or to Mieke Verfaellie, VA Boston Health Care System Jamaica Plain Campus, 150 S. Huntington Avenue, Boston MA, 02130, or via e-mail: mieke.verfaellie@va.gov.

Author Contributions

Virginie M. Patt: Conceptualization; Formal analysis; Methodology; Writing—Original draft. Daniela J. Palombo: Conceptualization; Formal analysis; Methodology; Writing—Original draft. Michael Esterman: Conceptualization; Formal analysis; Methodology; Writing—Original draft. Mieke Verfaellie: Conceptualization; Formal analysis; Funding acquisition; Methodology; Supervision; Writing—Original draft.

Diversity in Citation Practices

Retrospective analysis of the citations in every article published in this journal from 2010 to 2021 reveals a persistent pattern of gender imbalance: Although the proportions of authorship teams (categorized by estimated gender identification of first author/last author) publishing in the *Journal of Cognitive Neuroscience* (*JoCN*) during this period were M(an)/M = .407, W(oman)/M = .32, M/W = .115, and W/W = .159, the comparable proportions for the articles that these authorship teams cited were M/M = .549, W/M = .257, M/W = .109, and W/W = .085 (Postle and Fulvio, *JoCN*, 34:1, pp. 1–3). Consequently, *JoCN* encourages all authors to consider gender balance explicitly when selecting which articles to cite and gives them the opportunity to report their article's gender citation balance.

Note

1. Two ancillary analyses were carried using either the unsigned OL prediction error $|PE_{OL}|$ or the unsigned RL prediction error $|PE_{RL}|$ as parametric modulator. (Because of their high correlation, they were not included in the same analysis.) Using a voxel-wise threshold $p = .001$ and cluster-wise threshold $p = .05$, the results showed no area of significant mean $|PE_{OL}|$ or $|PE_{RL}|$ activation and no area of enhanced $|PE_{OL}|$ or $|PE_{RL}|$ activation with *logBF*. These findings are consistent with previous studies showing brain correlates for signed but not necessarily unsigned prediction errors in paradigms involving reward prediction (Ergo et al., 2020; Garrison et al., 2013).

REFERENCES

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*, 716–723. https://doi.org/10.1109/TAC.1974.1100705

Alink, A., Schwiedrzik, C. M., Kohler, A., Singer, W., & Muckli, L. (2010). Stimulus predictability reduces responses in primary visual cortex. *Journal of Neuroscience*, *30*, 2960–2966. https://doi.org/10.1523/JNEUROSCI.3730-10.2010, PubMed: 20181593

Ballard, I. C., Wagner, A. D., & McClure, S. M. (2019). Hippocampal pattern separation supports reinforcement learning. *Nature Communications*, *10*, 1–12. https://doi.org/10.1038/s41467-019-08998-1, PubMed: 30842581

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*, 1–48. https://doi.org/10.18637/jss.v067.i01

Bein, O., Duncan, K., & Davachi, L. (2020). Mnemonic prediction errors bias hippocampal states. *Nature Communications*, *11*, 3451. https://doi.org/10.1038/s41467-020-17287-1, PubMed: 32651370

Boorman, E. D., Rajendran, V. G., O'Reilly, J. X., & Behrens, T. E. (2016). Two anatomically and computationally distinct learning signals predict changes to stimulus–outcome associations in hippocampus. *Neuron*, *89*, 1343–1354. https://doi.org/10.1016/j.neuron.2016.02.014, PubMed: 26948895

Bornstein, A. M., & Daw, N. D. (2012). Dissociating hippocampal and striatal contributions to sequential prediction learning. *European Journal of Neuroscience*, *35*, 1011–1023. https://doi.org/10.1111/j.1460-9568.2011.07920.x, PubMed: 22487032

Bornstein, A. M., & Daw, N. D. (2013). Cortical and hippocampal correlates of deliberation during model-based decisions for rewards in humans. *PLoS Computational Biology*, *9*, e1003387. https://doi.org/10.1371/journal.pcbi.1003387, PubMed: 24339770

Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2017). Reminders of past choices bias decisions for reward in humans. *Nature Communications*, *8*, 1–9. https://doi.org/10.1038/ncomms15958, PubMed: 28653668

Buckner, R. L. (2010). The role of the hippocampus in prediction and imagination. *Annual Review of Psychology*, *61*, 27–48. https://doi.org/10.1146/annurev.psych.60.110707.163508, PubMed: 19958178

Bunzeck, N., Dayan, P., Dolan, R. J., & Duzel, E. (2010). A common mechanism for adaptive scaling of reward and novelty. *Human Brain Mapping*, *31*, 1380–1394. https://doi.org/10.1002/hbm.20939, PubMed: 20091793

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, *35*, 625–641. https://doi.org/10.1016/s0896-6273(02)00830-9, PubMed: 12194864

Burke, C. J., Tobler, P. N., Baddeley, M., & Schultz, W. (2010). Neural mechanisms of observational learning. *Proceedings of the National Academy of Sciences, U.S.A.*, *107*, 14431–14436. https://doi.org/10.1073/pnas.1003111107, PubMed: 20660717

Calderon, C. B., Loof, E. D., Ergo, K., Snoeck, A., Boehler, C. N., & Verguts, T. (2021). Signed reward prediction errors in the ventral striatum drive episodic memory. *Journal of Neuroscience*, *41*, 1716–1726. https://doi.org/10.1523/JNEUROSCI.1785-20.2020, PubMed: 33334870

Chen, J., Cook, P. A., & Wagner, A. D. (2015). Prediction strength modulates responses in human area CA1 to sequence violations. *Journal of Neurophysiology*, *114*, 1227–1238. https://doi.org/10.1152/jn.00149.2015, PubMed: 26063773

Cohen, N. J., & Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system* (p. 326). Cambridge, MA: MIT Press.

Collins, A. G., & Frank, M. J. (2012). How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *European Journal of Neuroscience*, *35*, 1024–1035. https://doi.org/10.1111/j.1460-9568.2011.07980.x, PubMed: 22487033

Davidow, J. Y., Foerde, K., Galván, A., & Shohamy, D. (2016). An upside to reward sensitivity: The hippocampus supports enhanced reinforcement learning in adolescence. *Neuron*, *92*, 93–99. https://doi.org/10.1016/j.neuron.2016.08.031, PubMed: 27710793

den Ouden, H. E., Daunizeau, J., Roiser, J., Friston, K. J., & Stephan, K. E. (2010). Striatal prediction error modulates cortical coupling. *Journal of Neuroscience*, *30*, 3210–3219. https://doi.org/10.1523/JNEUROSCI.4458-09.2010, PubMed: 20203180

den Ouden, H. E., Friston, K. J., Daw, N. D., McIntosh, A. R., & Stephan, K. E. (2009). A dual role for prediction error in associative learning. *Cerebral Cortex*, *19*, 1175–1185. https://doi.org/10.1093/cercor/bhn161, PubMed: 18820290

Desikan, R. S., Ségonne, F., Fischl, B., Quinn, B. T., Dickerson, B. C., Blacker, D., et al. (2006). An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage*, *31*, 968–980. https://doi.org/10.1016/j.neuroimage.2006.01.021, PubMed: 16530430

Dickerson, K. C., & Delgado, M. R. (2015). Contributions of the hippocampus to feedback learning. *Cognitive, Affective, & Behavioral Neuroscience*, *15*, 861–877. https://doi.org/10.3758/s13415-015-0364-5, PubMed: 26055632

Dickerson, K. C., Li, J., & Delgado, M. R. (2011). Parallel contributions of distinct human memory systems during probabilistic learning. *Neuroimage*, *55*, 266–276. https://doi.org/10.1016/j.neuroimage.2010.10.080, PubMed: 21056678

Duncan, K., Doll, B. B., Daw, N. D., & Shohamy, D. (2018). More than the sum of its parts: A role for the hippocampus in configural reinforcement learning. *Neuron*, *98*, 645–657. https://doi.org/10.1016/j.neuron.2018.03.042, PubMed: 29681530

Duncan, K., Ketz, N., Inati, S., & Davachi, L. (2012). Evidence for area CA1 as a match/mismatch detector: A high-resolution fMRI study of the human hippocampus. *Hippocampus*, *22*, 389–398. https://doi.org/10.1002/hipo.20933, PubMed: 21484934

Eichenbaum, H. (2000). A cortical–hippocampal system for declarative memory. *Nature Reviews Neuroscience*, *1*, 41–50. https://doi.org/10.1038/35036213, PubMed: 11252767

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, *23*, 209–226. https://doi.org/10.1016/s0896-6273(00)80773-4, PubMed: 10399928

Eklund, A., Nichols, T., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences, U.S.A.*, *113*, 7900–7905. https://doi.org/10.1073/pnas.1602413113, PubMed: 27357684

Ergo, K., De Loof, E., & Verguts, T. (2020). Reward prediction error and declarative memory. *Trends in Cognitive Sciences*, *24*, 388–397. https://doi.org/10.1016/j.tics.2020.02.009, PubMed: 32298624

Floresco, S. B. (2007). Dopaminergic regulation of limbic-striatal interplay. *Journal of Psychiatry & Neuroscience*, *32*, 400–411, PubMed: 18043763

Foerde, K., Race, E., Verfaellie, M., & Shohamy, D. (2013). A role for the medial temporal lobe in feedback-driven learning: Evidence from amnesia. *Journal of Neuroscience*, *33*, 5698–5704. https://doi.org/10.1523/JNEUROSCI.5217-12.2013, PubMed: 23536083

Foerde, K., & Shohamy, D. (2011). Feedback timing modulates brain systems for learning in humans. *Journal of Neuroscience*, *31*, 13157–13167. https://doi.org/10.1523/JNEUROSCI.2701-11.2011, PubMed: 21917799

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. (2013). Emcee: The MCMC hammer. *Publications of the Astronomical Society of the Pacific*, *125*, 306. https://doi.org/10.1086/670067

Frank, M. J., Woroch, B. S., & Curran, T. (2005). Error-related negativity predicts reinforcement learning and conflict biases. *Neuron*, *47*, 495–501. https://doi.org/10.1016/j.neuron.2005.06.020, PubMed: 16102533

Friel, N., & Pettitt, A. N. (2008). Marginal likelihood estimation via power posteriors. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *70*, 589–607. https://doi.org/10.1111/j.1467-9868.2007.00650.x

Garrison, J., Erdeniz, B., & Done, J. (2013). Prediction error in reinforcement learning: A meta-analysis of neuroimaging studies. *Neuroscience & Biobehavioral Reviews*, *37*, 1297–1310. https://doi.org/10.1016/j.neubiorev.2013.03.023, PubMed: 23567522

Gershman, S. J., & Daw, N. D. (2017). Reinforcement learning and episodic memory in humans and animals: An integrative framework. *Annual Review of Psychology*, *68*, 101–128. https://doi.org/10.1146/annurev-psych-122414-033625, PubMed: 27618944

Gläscher, J., Daw, N., Dayan, P., & O'Doherty, J. P. (2010). States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron*, *66*, 585–595. https://doi.org/10.1016/j.neuron.2010.04.016, PubMed: 20510862

Gold, P. E. (2004). Coordination of multiple memory systems. *Neurobiology of Learning and Memory*, *82*, 230–242. https://doi.org/10.1016/j.nlm.2004.07.003, PubMed: 15464406

Goodman, J., & Weare, J. (2010). Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, *5*, 65–80. https://doi.org/10.2140/camcos.2010.5.65

Grinsted, A. (2015). *Grinsted/gwmcmc* [MATLAB]. https://github.com/grinsted/gwmcmc.

Harrison, L. M., Duggins, A., & Friston, K. J. (2006). Encoding uncertainty in the hippocampus. *Neural Networks*, *19*, 535–546. https://doi.org/10.1016/j.neunet.2005.11.002, PubMed: 16527453

Hauser, T. U., Iannaccone, R., Walitza, S., Brandeis, D., & Brem, S. (2015). Cognitive flexibility in adolescence: Neural and behavioral mechanisms of reward prediction error processing in adaptive decision making during development. *Neuroimage*, *104*, 347–354. https://doi.org/10.1016/j.neuroimage.2014.09.018, PubMed: 25234119

Hindy, N. C., Avery, E. W., & Turk-Browne, N. B. (2019). Hippocampal–neocortical interactions sharpen over time for predictive actions. *Nature Communications*, *10*, 3989. https://doi.org/10.1038/s41467-019-12016-9, PubMed: 31488845

Jeffreys, H. (1961). *The theory of probability* (3rd ed.). Clarendon Press.

Jocham, G., Klein, T. A., & Ullsperger, M. (2011). Dopamine-mediated reinforcement learning signals in the striatum and ventromedial prefrontal cortex underlie value-based choices. *Journal of Neuroscience*, *31*, 1606–1613. https://doi.org/10.1523/JNEUROSCI.3904-10.2011, PubMed: 21289169

Johnson, A., van der Meer, M. A., & Redish, A. D. (2007). Integrating hippocampus and striatum in decision-making. *Current Opinion in Neurobiology*, *17*, 692–697. https://doi.org/10.1016/j.conb.2008.01.003, PubMed: 18313289

Kable, J. W., & Glimcher, P. W. (2009). The neurobiology of decision: Consensus and controversy. *Neuron*, *63*, 733–745. https://doi.org/10.1016/j.neuron.2009.09.003, PubMed: 19778504

Kahn, I., & Shohamy, D. (2013). Intrinsic connectivity between the hippocampus, nucleus accumbens, and ventral tegmental area in humans. *Hippocampus*, *23*, 187–192. https://doi.org/10.1002/hipo.22077, PubMed: 23129267

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, *273*, 1399–1402. https://doi.org/10.1126/science.273.5280.1399, PubMed: 8703077

Kok, P., & Turk-Browne, N. B. (2018). Associative prediction of visual shape in the hippocampus. *The Journal of Neuroscience*, *38*, 6888–6899. https://doi.org/10.1523/JNEUROSCI.0163-18.2018, PubMed: 29986875

Kumaran, D., & Maguire, E. (2006). An unexpected sequence of events: Mismatch detection in the human hippocampus. *PLoS Biology*, *4*, e424. https://doi.org/10.1371/journal.pbio.0040424, PubMed: 17132050

Matsumoto, M., Matsumoto, K., Abe, H., & Tanaka, K. (2007). Medial prefrontal cell activity signaling prediction errors of action values. *Nature Neuroscience*, *10*, 647–656. https://doi.org/10.1038/nn1890, PubMed: 17450137

McClure, S. M., Berns, G. S., & Montague, P. R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron*, *38*, 339–346. https://doi.org/10.1016/S0896-6273(03)00154-5, PubMed: 12718866

Meder, D., Madsen, K. H., Hulme, O., & Siebner, H. R. (2016). Chasing probabilities—Signaling negative and positive prediction errors across domains. *Neuroimage*, *134*, 180–191. https://doi.org/10.1016/j.neuroimage.2016.04.019, PubMed: 27083529

Mogenson, G. J., Jones, D. L., & Yim, C. Y. (1980). From motivation to action: Functional interface between the limbic system and the motor system. *Progress in Neurobiology*, *14*, 69–97. https://doi.org/10.1016/0301-0082(80)90018-0, PubMed: 6999537

Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, *4*, 133–142. https://doi.org/10.1111/j.2041-210x.2012.00261.x

O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, *38*, 329–337. https://doi.org/10.1016/S0896-6273(03)00169-7, PubMed: 12718865

Packard, M. G., & Goodman, J. (2013). Factors that influence the relative use of multiple memory systems. *Hippocampus*, *23*, 1044–1052. https://doi.org/10.1002/hipo.22178, PubMed: 23929809

Palombo, D. J., Hayes, S. M., Reid, A. G., & Verfaellie, M. (2019). Hippocampal contributions to value-based learning: Converging evidence from fMRI and amnesia. *Cognitive, Affective, & Behavioral Neuroscience*, *19*, 523–536. https://doi.org/10.3758/s13415-018-00687-8, PubMed: 30767129

Palombo, D. J., Patt, V. M., Hunsberger, R., & Verfaellie, M. (2021). Probabilistic value learning in medial temporal lobe amnesia. *Hippocampus*, *31*, 461–468. https://doi.org/10.1002/hipo.23317, PubMed: 33638580

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. https://doi.org/10.1038/nature05051, PubMed: 16929307

Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., et al. (2001). Interactive memory systems in the human brain. *Nature*, *414*, 546–550. https://doi.org/10.1038/35107080, PubMed: 11734855

R Core Team. (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory* (Vol. 2).

Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience*, *8*, 657–661. https://doi.org/10.1038/nrn2213, PubMed: 17700624

Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, *22*, 1622–1627. https://doi.org/10.1016/j.cub.2012.06.056, PubMed: 22885059

Schonberg, T., O'Doherty, J. P., Joel, D., Inzelberg, R., Segev, Y., & Daw, N. D. (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. https://doi.org/10.1016/j.neuroimage.2009.08.011, PubMed: 19682583
Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*, 1–27. https://doi.org/10.1152/jn.1998.80.1.1, PubMed: 9658025
Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, *6*, 461–464. https://doi.org/10.1214/aos/1176344136
Seymour, B., O'Doherty, J. P., Dayan, P., Koltzenburg, M., Jones, A. K., Dolan, R. J., et al. (2004). Temporal difference models describe higher-order learning in humans. *Nature*, *429*, 664–667. https://doi.org/10.1038/nature02581, PubMed: 15190354
Shohamy, D., & Adcock, R. A. (2010). Dopamine and adaptive memory. *Trends in Cognitive Sciences*, *14*, 464–472. https://doi.org/10.1016/j.tics.2010.08.002, PubMed: 20829095
Sinclair, A. H., Manalili, G. M., Brunec, I. K., Adcock, R. A., & Barense, M. D. (2021). Prediction errors disrupt hippocampal representations and update episodic memories. *Proceedings of the National Academy of Sciences, U.S.A.*, *118*, e2117625118. https://doi.org/10.1073/pnas.2117625118, PubMed: 34911768
Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E. J., Johansen-Berg, H., et al. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, *23*, S208–S219. https://doi.org/10.1016/j.neuroimage.2004.07.051, PubMed: 15501092
Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*, 171–177. https://doi.org/10.1016/j.nlm.2004.06.005, PubMed: 15464402
Squire, L. R., Knowlton, B., & Musen, G. (1993). The structure and organization of memory. *Annual Review of Psychology*, *44*, 453–495. https://doi.org/10.1146/annurev.ps.44.020193.002321, PubMed: 8434894
Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, *21*, 1934–1945. https://doi.org/10.1162/jocn.2009.21131, PubMed: 18823241
Wang, F., Schoenbaum, G., & Kahnt, T. (2020). Interactions between human orbitofrontal cortex and hippocampus support model-based inference. *PLoS Biology*, *18*, e3000578. https://doi.org/10.1371/journal.pbio.3000578, PubMed: 31961854
White, N. M., & McDonald, R. J. (2002). Multiple parallel memory systems in the brain of the rat. *Neurobiology of Learning and Memory*, *77*, 125–184. https://doi.org/10.1006/nlme.2001.4008, PubMed: 11848717
Wilson, R., & Collins, A. (2019). Ten simple rules for the computational modeling of behavioral data. *eLife*, *8*, e49547. https://doi.org/10.7554/eLife.49547, PubMed: 31769410
Woolrich, M. W., Ripley, B. D., Brady, M., & Smith, S. M. (2001). Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage*, *14*, 1370–1386. https://doi.org/10.1006/nimg.2001.0931, PubMed: 11707093