

Hierarchical Prediction Errors in Midbrain and Basal Forebrain during Sensory Learning

Sandra Iglesias,^{1,2,*} Christoph Mathys,^{1,2} Kay H. Brodersen,^{1,2} Lars Kasper,^{1,2} Marco Piccirelli,² Hanneke E.M. den Ouden,³ and Klaas E. Stephan^{1,2,4}

¹Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich and Swiss Federal Institute of Technology (ETH), 8032 Zurich, Switzerland

²Laboratory for Social and Neural Systems Research (SNS), University of Zurich, 8091 Zurich, Switzerland

³Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen, 6500 HE, The Netherlands

⁴Wellcome Trust Centre for Neuroimaging, University College London, London WC1N 3BG, UK

*Correspondence: iglesias@biomed.ee.ethz.ch

<http://dx.doi.org/10.1016/j.neuron.2013.09.009>

SUMMARY

In Bayesian brain theories, hierarchically related prediction errors (PEs) play a central role for predicting sensory inputs and inferring their underlying causes, e.g., the probabilistic structure of the environment and its volatility. Notably, PEs at different hierarchical levels may be encoded by different neuromodulatory transmitters. Here, we tested this possibility in computational fMRI studies of audio-visual learning. Using a hierarchical Bayesian model, we found that low-level PEs about visual stimulus outcome were reflected by widespread activity in visual and supramodal areas but also in the midbrain. In contrast, high-level PEs about stimulus probabilities were encoded by the basal forebrain. These findings were replicated in two groups of healthy volunteers. While our fMRI measures do not reveal the exact neuron types activated in midbrain and basal forebrain, they suggest a dichotomy between neuromodulatory systems, linking dopamine to low-level PEs about stimulus outcome and acetylcholine to more abstract PEs about stimulus probabilities.

INTRODUCTION

The notion that the brain has evolved to implement a predictive machinery for anticipation of future events has existed since early cybernetic theories (Ashby, 1952). The mechanisms by which the brain learns the probabilistic structure of the world have been examined primarily from the perspective of reinforcement learning (RL), with a focus on how reward learning is driven by prediction errors (PEs) (Fletcher et al., 2001; McClure et al., 2003; O'Doherty et al., 2003; Pessiglione et al., 2006; Wunderlich et al., 2011). Another perspective is provided by theories that view the brain as approximating optimal Bayesian inference (Dayan et al., 1995; Doya et al., 2011; Friston, 2009; Knill and Pouget, 2004; Körding and Wolpert, 2006). These theories go beyond reward learning and have been applied to many aspects of perception as, for example, in theories of “predictive coding”

(Rao and Ballard, 1999) and the “free energy principle” (Friston et al., 2006).

A central postulate of these Bayesian perspectives is that the brain continuously updates a hierarchical generative model of its sensory inputs to predict future events and infer on the causal structure of the world. This belief updating process rests on multiple, hierarchically related PEs that are weighted by their precision. Notably, these PEs are not restricted to reward, but concern all types of sensory events as well as their underlying “laws,” e.g., probabilistic associations and how these change in time (volatility; Behrens et al., 2007). Simply speaking, estimates of environmental volatility are updated in proportion to PEs about stimulus probabilities; in turn, estimates of stimulus probabilities are updated by PEs about stimulus occurrences.

While several empirical studies have examined human behavior and brain activity from this Bayesian perspective, the hierarchical nature of PEs has received little attention so far. This is a significant gap, not only because hierarchically related PEs are at the heart of the Bayesian formalism, but also because PEs at different hierarchical levels may be linked to different neuromodulatory transmitter systems. While dopamine (DA) has long been related to the encoding of PEs about reward (Daw and Doya, 2006; Schultz et al., 1997), other modulatory neurotransmitters have been linked to more abstract roles, such as encoding of “expected uncertainty” by acetylcholine (ACh) (Yu and Dayan, 2002, 2005). Notably, this was (implicitly) operationalized as a higher-level PE in that it represents the difference between a conditional probability (degree of cue validity) and certainty. Other computational concepts of ACh suggested that it may be representing the learning rate (Doya, 2002). Again, this notion can be related to hierarchical Bayesian accounts where the learning rate at any given level is proportional to the precision of predictions and evolves under the influence of the next higher level in the hierarchy (Mathys et al., 2011). This weighting by precision (a form of adaptive scaling) is crucial and has been described for DA responses to reward (Tobler et al., 2005) and novelty (Bunzeck et al., 2010). Such a function may generalize across neuromodulators: it has been suggested that both DA and ACh may be involved in the precision-weighting of PEs (Friston, 2009; Friston et al., 2012).

Here, we present behavioral and fMRI studies that examine possible links between neuromodulatory systems and hierarchical precision-weighted PEs during associative learning. The

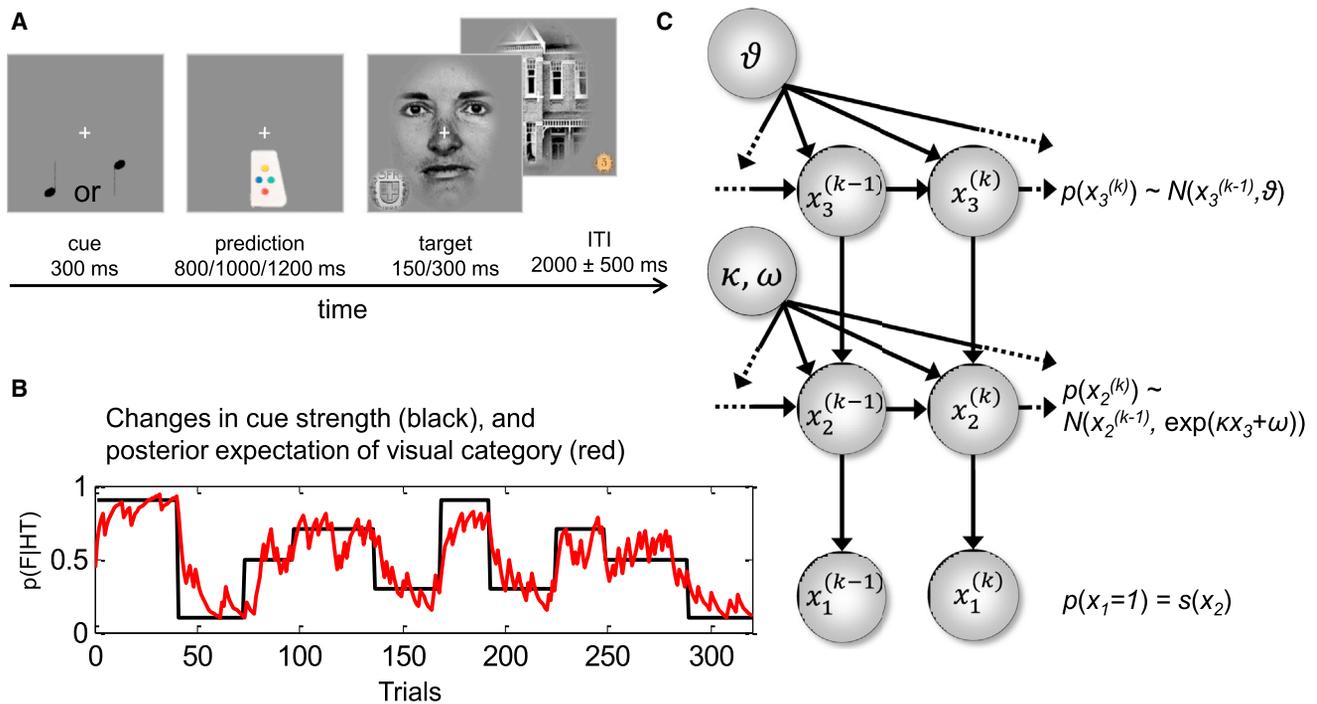


Figure 1. Task Design and Model

(A) Task design. Subjects had to predict within 800 ms (behavioral study), 1,000 ms (first fMRI study), or 1,200 ms (second fMRI study) which visual stimulus (face or house) followed an auditory cue (high or low tone). In the behavioral study and first fMRI study, a monetary reward (0.05 or 5.00 Swiss Francs coin) was randomly presented in one of the four corners. The type of coin presented was uncorrelated to visual stimulus outcome and was omitted in the second fMRI study. (B) Black: time-varying cue–outcome contingency, including strongly predictive cues (probabilities of 0.9 and 0.1), moderately predictive cues (0.7, 0.3) and nonpredictive cues (0.5); red: example of a subject-specific trajectory of the posterior expectation of visual category. (C) HGF: generative model. x_1 represents the stimulus identity (category), x_2 the cue–outcome contingency (the conditional probability of the visual stimulus given the auditory cue) in logit space, and x_3 represents the log–volatility of the environment. See Equations 2, 3, and 4 and Table S2. See also Figures S1, S2, and S3 and Tables S1, S2, S4, S5, and S6.

analyses rest on a recently developed hierarchical Bayesian model, the Hierarchical Gaussian Filter (HGF) (Mathys et al., 2011), which does not assume fixed “ideal” learning across subjects but contains subject-specific parameters that couple the hierarchical levels and allow for individual expression of (approximate) Bayes-optimal learning. Using the subject-specific learning trajectories, we examined whether activity in neuromodulatory nuclei could be explained by precision-weighted PEs, and if so, at which hierarchical level. In particular, we focused on dopaminergic and cholinergic nuclei, using anatomical masks specifically developed for these regions. Importantly, we examined 118 healthy volunteers from three separate samples, two of which underwent fMRI ($n = 45$ and $n = 27$, respectively). This enabled us to verify the robustness of our results and test which of them would replicate across samples.

RESULTS

We report findings obtained from three separate samples of healthy volunteers undergoing purely behavioral assessment ($n = 46$) or combined fMRI-behavior ($n = 45$ and $n = 27$). All three studies used a simple associative audio-visual learning task where participants had to learn the time-varying predictive strengths of auditory cues and predict upcoming visual stimuli

(faces or houses) by button press (Figure 1). This task required hierarchical learning about stimulus occurrences, stimulus probabilities, and volatility that we modeled as a hierarchical Bayesian belief updating process, using a standard HGF with three levels (Mathys et al., 2011); see Experimental Procedures for details.

Modeling of Behavioral Data

In a first step, we used random effects Bayesian model selection (BMS) (Stephan et al., 2009) to examine the possibility that our subjects might have engaged in a different cognitive process than intended, or may have used a different model than hypothesized. In the behavioral study and first fMRI study, we tried to ensure constant motivation of our participants by associating each trial with a monetary reward whose potential pay-out at the end of the experiment depended on successful prediction of the visual outcome (face or house). Even though subjects were explicitly instructed that these reward were random and orthogonal to the visual outcomes, one may wonder whether subjects’ learning might nevertheless have been driven by (implicit) prediction of these trial-wise reward. To exclude this possibility, we compared a three-level HGF assuming that audio-visual associations were learned and guided subjects’ behavior (HGF₁; Figure 1C) to a second HGF

that assumed that participants attempted to learn and predict trial-wise reward (HGF₂).

A second question was whether our participants were indeed engaging in hierarchical learning and updating their learning rate dynamically, as our Bayesian model assumed, or used a simpler learning mechanism. To clarify this, we added two more models to our comparison set. The models were a Bayesian model with reduced hierarchical depth (HGF₃) in which the third level was eliminated from the hierarchy, and a standard Rescorla-Wagner (RL) model with a fixed learning rate. Finally, we implemented a RL model with dynamic learning rate (Sutton, 1992) that was recommended by one of the reviewers as a non-Bayesian alternative to HGF₁. See the [Supplemental Experimental Procedures section C](#) (available online) for more information on these models.

Comparing these five models, we found that, across studies, HGF₁ was the superior model in 86 out of our 118 participants. Examining each study separately, random effects BMS yielded posterior model probabilities of 84% (behavioral study), 74% (first fMRI study), and 72% (second fMRI study) for HGF₁, which was five to ten times higher than for the next best model in each case (Table S1). As a consequence, in each study, the exceedance probability in favor of HGF₁ (i.e., the probability that its posterior probability was higher than that of any other model considered) (Stephan et al., 2009) was indistinguishable from 100%. These results provide strong evidence that our participants did learn the task-relevant conditional probabilities of visual stimuli (instead of predicting the incidental reward) and were capable of updating their learning rate dynamically.

We next examined the estimates of the free parameters (κ , ϑ , ζ) from the winning model (Table S2). These estimates were comparable across the three studies, as demonstrated by ANOVA: none of the model parameters showed significant differences across studies (κ : $F(2,115) = 1.04$, $p = 0.358$; ϑ : $F(2,115) = 0.91$, $p = 0.405$; ζ : $F(2,115) = 2.98$, $p = 0.055$). Additionally, we used multiple regression to evaluate how well our model explained subjects' behavior (percentage of correct responses). This quantified model performance in terms of variance explained, complementary to the relative model comparison by BMS above. This analysis showed that the linear combination of the three model parameters predicted subjects' task performance well (behavioral study: $R^2 = 0.64$, $F(3,42) = 25.3$, $p < 0.001$; first fMRI study: $R^2 = 0.59$, $F(3,41) = 20.1$, $p < 0.001$; second fMRI study: $R^2 = 0.63$, $F(3,23) = 13.2$, $p < 0.001$).

fMRI Data Analysis

As detailed in the [Experimental Procedures](#) section, our fMRI analysis focused on precision-weighted PEs and uncertainty estimates across the hierarchical levels of the HGF. For each of these variables, our analysis proceeded in three steps (see [Experimental Procedures](#)): first, we performed whole-brain analyses; second, we focused on our anatomically defined regions of interest (ROIs), using a combined mask of dopaminergic and cholinergic nuclei in the brain stem and subcortex; finally, we conducted these fMRI analyses separately in two independent samples of $n = 45$ and $n = 27$ volunteers. Note that we only report those findings that survived stringent family-wise error (FWE) peak-level correction for multiple tests ($p < 0.05$) and

that could be replicated across studies. Replication was assessed using a voxel-wise "logical AND" operation on the FWE-thresholded activation maps from both fMRI studies, and only those activations are being reported in which this procedure showed an overlap of significant activations in both fMRI studies.

Low-Level Precision-Weighted Prediction Errors

Initially, we examined the precision-weighted PE about visual stimulus outcome, ε_2 (for mathematical details, see [Experimental Procedures](#) and the [Supplemental Experimental Procedures, section A](#)). In both fMRI studies, our whole-brain analyses demonstrated significant activations in a widely distributed set of regions (Table 1; Figure 2). In addition to the visual cortex (around the calcarine sulcus), the activity of numerous supramodal regions correlated positively with trial-wise estimates of ε_2 , including the middle and inferior frontal gyri, anterior cingulate cortex (ACC), intraparietal sulcus (IPS), and anterior insula, all located bilaterally. Perhaps the most notable finding, however, was a significant activation of the midbrain (ventral tegmental area [VTA]/substantia nigra [SN]). In both fMRI studies, this VTA/SN activation not only survived FWE correction within our anatomically defined mask, but also across the whole brain ($p < 0.05$; Figure 3). This finding is remarkable because the precision-weighted PE ε_2 concerns a purely sensory event: the visual stimulus category predicted by the auditory cue. This conclusion is supported by the BMS analysis of the behavioral data described above that demonstrated that in the first fMRI study subjects were not trying to predict reward but visual outcomes. Furthermore, in the second fMRI, study rewards were omitted entirely while keeping sensory stimulation and task demands identical.

Interestingly, as implied by predictive coding theories (cf. Friston, 2005), regions whose activity correlated positively with PEs about visual inputs considerably overlapped with regions that activated on each trial, regardless of the computational state and stimulus category ("task execution per se"). Figure 4 shows the results of a nested conjunction analysis: this combined the conjunction analyses of contrasts testing for task execution per se (i.e., a statistical contrast on the base regressor encoding trial events, not the parametric modulators) and for ε_2 , respectively, across both fMRI studies. These results indicated that in both studies, primary visual cortex (calcarine sulcus), bilateral IPS, right dorsolateral prefrontal cortex (DLPFC), and right anterior insula were activated by the task per se and by precision-weighted PEs about stimulus category. Please note that this is an extremely conservative analysis: all conjunction analyses tested the conjunction null hypothesis, i.e., a "logical AND" (Nichols et al., 2005), with all contrasts thresholded at $p < 0.05$ (FWE whole-brain corrected), and the combination of these conjunctions across both studies corresponded to a double logical AND.

The results reported so far refer to the outcome prediction error ε_2 ; this is the (precision-weighted) difference between the actual visual stimulus outcome and its a priori probability (i.e., before trial outcome observation). However, we can also use the predictions from our model to examine activations reflecting choice prediction error ε_{ch} ; this is the difference between the correctness of the subject's choice and the a priori probability of this

Table 1. Whole-Brain Activations by ϵ_2

fMRI study 1	Hemisphere	x	y	z	t Score	fMRI Study 2	Hemisphere	x	y	z	t Score
ϵ_2 : Positive Correlation						ϵ_2 : Positive Correlation					
Middle frontal gyrus/ Anterior/ middle cingulate cortex	R	34	8	57	10.25	Middle frontal gyrus	R	34	14	55	7.95
Insula	R	33	24	-3	10.13	Anterior/middle cingulate cortex	R	2	30	40	8.91
Inferior parietal cortex	R	39	-49	45	9.49	Insula	R	32	24	-3	10.85
Precuneus	R	8	-69	49	9.00	Inferior parietal cortex	R	38	-46	46	8.98
Intraparietal sulcus/ inferior parietal cortex	L	-28	-61	43	8.53	Precuneus	R	4	-70	46	8.70
Inferior frontal gyrus	L	-44	26	31	8.25	Intraparietal sulcus/ inferior parietal cortex	L	-28	-61	39	7.59
Insula	L	-30	24	-0	7.96	Inferior frontal gyrus	L	-44	24	33	9.30
Middle frontal gyrus	L	-28	5	63	7.52	Insula	L	-28	24	-3	9.20
Middle frontal gyrus	L	-27	50	15	6.30	Middle frontal gyrus	L	-28	11	60	7.92
Lingual gyrus	L	-8	-78	3	5.55	Middle frontal gyrus	L	-28	53	13	6.88
Lingual gyrus	R	2	-78	3	5.36	Lingual gyrus	L	-12	-81	4	5.29
Supramarginal gyrus	R	48	-48	27	5.40	Lingual gyrus	R	2	-82	4	5.09
Cerebellum	L	-30	-57	-32	5.35	Cerebellum	L	-30	-55	-32	6.16
Middle temporal gyrus	R	58	-30	-8	5.21	Supramarginal gyrus	R	45	-46	25	6.59
VTA / substantia nigra	R	3	-24	-18	5.12	Middle temporal gyrus	R	56	-30	-8	6.18
Prefrontal cortex	L	-16	14	64	5.00	VTA / substantia nigra	R	2	-21	-18	5.06
						Prefrontal cortex	L	-18	18	66	8.30

All results: $p < 0.05$ FWE whole-brain corrected. MNI coordinates and t values for regions activated by ϵ_2 , the precision-weighted PE about stimulus outcome, in the first and second fMRI study. Only those activations are listed that were replicated across studies. The activation in the first row constituted a single cluster in the first study, whereas it was split into two separate clusters in the second study.

choice being correct (see the [Supplemental Experimental Procedures, section B](#), for formal definitions of both PEs).

In both fMRI studies, choice PEs evoked prominent activations ($p < 0.05$ FWE whole-brain corrected; [Figure 5](#)) in numerous regions, including the bilateral ventral striatum, ventromedial prefrontal cortex, OFC and ACC (for a complete list, see [Table S7](#)). Activations of these regions are commonly found for reward PEs, and it is remarkable that we obtain a similar activation pattern even though in our studies learning was orthogonal to reward (fMRI study 1) and reward were absent (fMRI study 2). Finally, it is notable that the activation of the ventral striatum also extended into the basal forebrain, as delineated by our anatomical mask ($p < 0.05$ FWE corrected for the entire mask volume).

High-Level Precision-Weighted Prediction Errors

Subsequently, we investigated precision-weighted PEs at the next higher level of the hierarchy in our Bayesian model. This PE, ϵ_3 , concerns the cue-outcome contingency, i.e., the probability (in logit space) of the visual stimulus category given the auditory cue, and is used to update estimates of log-volatility at the third level of the HGF. We found that the trial-wise expression of this PE correlated positively with activity in the septal part of the cholinergic basal forebrain ([Table 2](#); [Figure 6](#)). In both fMRI studies, this activation was significant ($p < 0.05$) when corrected for multiple comparisons across the volume of our anatomically defined mask (that included all cholinergic and dopaminergic nuclei in brain stem and subcortex).

DISCUSSION

In this study, three independent groups of healthy volunteers ($n = 118$ in total) performed an audio-visual associative learning task that required explicit predictions about an upcoming visual stimulus category (face or house) given a preceding auditory cue. Because the cue-outcome contingencies were varying unpredictably in time, optimal performance required hierarchical learning about conditional stimulus probabilities and their change in time.

Our analyses showed that participants were indeed likely to engage in such a hierarchical learning process. Formal statistical comparison of five alternative models indicated that a hierarchical Bayesian model (a three-level HGF) best explained the observed behavioral data. Applying the computational trajectories from this model to fMRI data, we found that precision-weighted PEs about visual outcome, ϵ_2 , were not only encoded by numerous cortical areas, including dopaminoceptive regions like DLPFC, ACC, and insula, but also by the dopaminergic VTA/SN. Notably, we verified both statistically and experimentally that these PE responses concerned visual stimulus categories and not reward. At the higher level of the model's hierarchy, precision-weighted PEs about cue-outcome contingencies (conditional probabilities of the visual outcome given the auditory cue), ϵ_3 , were reflected by activity in the cholinergic basal forebrain.

Our findings have two important implications. First, our results are in accordance with a central notion in Bayesian theories of

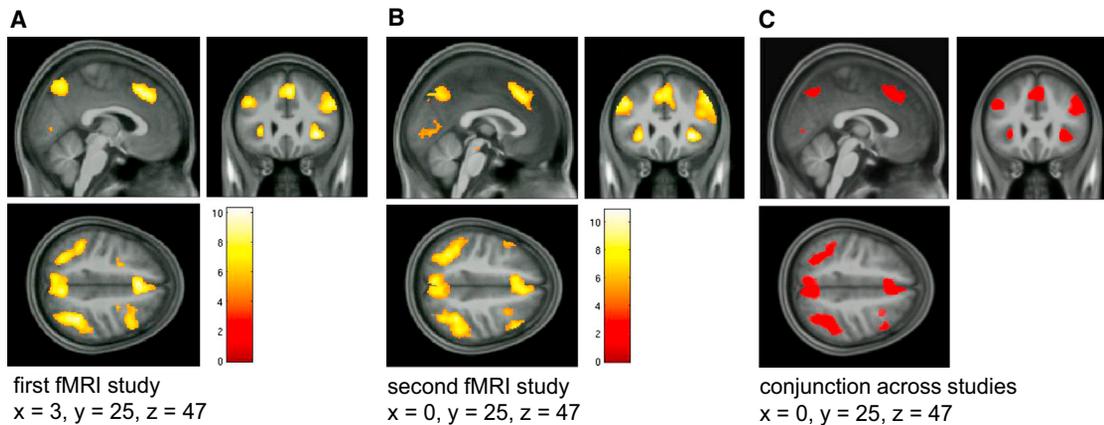


Figure 2. Whole-Brain Activations by ϵ_2

Activations by precision-weighted prediction error about visual stimulus outcome, ϵ_2 , in the first fMRI study (A) and the second fMRI study (B). Both activation maps are shown at a threshold of $p < 0.05$, FWE corrected for multiple comparisons across the whole brain. To highlight replication across studies, (C) shows the results of a “logical AND” conjunction, illustrating voxels that were significantly activated in both studies. See Table S3 for deactivations.

brain function, such as predictive coding (Friston, 2005; Rao and Ballard, 1999): even seemingly simple processes of perceptual inference and learning do not rest on a single PE but rely on hierarchically related PE computations. As a corollary, one would expect a widespread expression of PEs within the neuronal system engaged by a particular task. Indeed, we found a remarkable overlap of areas involved in the execution of the task and areas expressing PEs (Figure 4). Second, our findings suggest a potential dichotomy with regard to the computational roles of DA and ACh. According to our results, the midbrain may be encoding outcome-related PEs, independent of extrinsic reward. In contrast, the basal forebrain may be signaling more abstract PEs that do not concern sensory outcomes per se but their probabilities. In the following, we will discuss these two implications in the context of the previous literature.

Since early accounts of general systems theory and cybernetics (Ashby, 1952), the notion of PE as a teaching signal for adaptive behavior has taken an increasingly central place in theories of brain function. In contemporary neuroscience, PEs play a pivotal role in two frameworks, reinforcement learning (RL) and Bayesian theories. Studies inspired by RL have largely focused on the role of reward PEs, suggesting that these are encoded by phasic dopamine release from neurons in VTA/SN (Montague et al., 2004; Schultz et al., 1997). In humans, this has been supported by fMRI studies that have demonstrated the presence of reward PE signals in the VTA/SN (e.g., D’Ardenne et al., 2008; Diuk et al., 2013; Klein-Flügge et al., 2011) or in regions targeted by its projections, such as the striatum (Gläscher et al., 2010; McClure et al., 2003; Murray et al., 2008; O’Doherty et al., 2003; Pessiglione et al., 2006; Schonberg et al., 2010).

While RL models have also been used to study PE-dependent learning in the sensory domain (den Ouden et al., 2009; Law and Gold, 2009), a more prevalent framework to study perception has been the “Bayesian brain hypothesis” that the brain constructs and updates a generative model of its sensory inputs (Doya et al., 2011). One particular formulation of this hypothesis is pre-

dictive coding (Friston, 2005; Rao and Ballard, 1999) that postulates that PEs are weighted by their precision and are computed at any level of hierarchically organized information processing cascades, as in sensory systems. This has been examined by several fMRI studies that contrasted predictable versus unpredictable visual stimuli, finding PE responses in visual areas specialized for the respective stimuli used (Harrison et al., 2007; Summerfield and Koechlin, 2008) and precision-weighting under attention (Kok et al., 2012). Other studies have used an explicit model of trial-wise PEs, using visual (Egner et al., 2010) or audio-visual associative learning (den Ouden et al., 2010; den Ouden et al., 2009) paradigms. Notably, these studies did not have explicit readouts of subjects’ predictions and used relatively simple modeling approaches: they either described implicit learning processes (in the absence of behavioral responses) using a delta-rule RL model (den Ouden et al., 2009; Egner et al., 2010), or dealt with indirect measures of prediction (e.g., reaction times) using an ideal Bayesian observer with a fixed learning trajectory across subjects (den Ouden et al., 2010).

Our present study goes beyond these previous attempts by (1) requiring explicit trial-by-trial predictions, and (2) characterizing learning via a hierarchical Bayesian model that provides subject- and trial-specific estimates of precision-weighted PEs at different hierarchical levels of computation. Based on these advances, the present study shows much more widespread sensory PE responses than previously reported. Replicated in two separate groups, these responses were not only found in the visual cortex, but also in many supramodal areas in prefrontal, cingulate, parietal, and insular cortex (Figure 2). Whereas a distribution of reward (Vickery et al., 2011) and value signals (FitzGerald et al., 2012) across the whole brain have recently been demonstrated in humans, this has not yet been shown, to our knowledge, for PEs; in this case, precision-weighted PEs about the sensory outcome (visual stimuli).

Perhaps the most interesting aspect of our findings on sensory outcome PEs, ϵ_2 , was the significant activation of the midbrain. In humans, strong empirical evidence exists for DA involvement

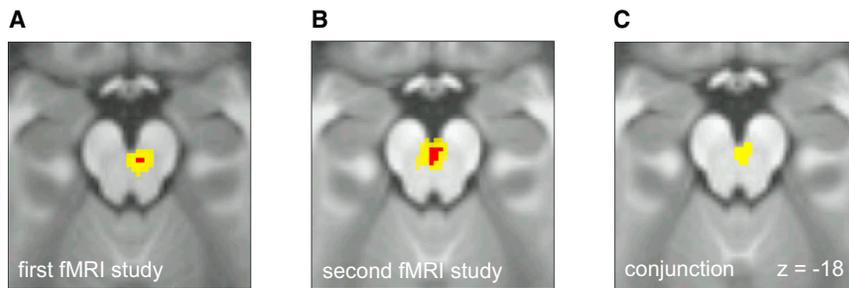


Figure 3. Midbrain Activation by ε_2

Activation of the dopaminergic VTA/SN associated with precision-weighted prediction error about stimulus category, ε_2 . This activation is shown both at $p < 0.05$ FWE whole-brain corrected (red) and $p < 0.05$ FWE corrected for the volume of our anatomical mask comprising both dopaminergic and cholinergic nuclei (yellow).

(A) Results from the first fMRI study.

(B) Second fMRI study.

(C) Conjunction (logical AND) across both studies.

in processing reward PEs (Montague et al., 2004; Schultz et al., 1997) and novelty (Bunzeck and Düzzel, 2006). In animal studies, dopaminergic midbrain responses to visual stimuli have been reported in the absence of reward; however, this required that the stimuli were novel, arousing or physically similar to reward-related stimuli (Horvitz, 2000; Redgrave and Gurney, 2006; Schultz, 1998). In contrast, in our study the VTA/SN responses scaled with trial-by-trial precision-weighted PE about the stimulus category; these were neither reward-related, arousing nor novel (we kept repeating two to four face and house stimuli in each study). One could think of VTA/SN activity reflecting conditional novelty (Bayesian surprise); however, this is not a tight link because ε_2 is only related but not identical to Bayesian surprise (see Supplemental Experimental Procedures).

An important caveat is that we cannot claim with certainty that the midbrain activation we found specifically reflects the activity of DA neurons in VTA/SN because this region is not homogenous in its cellular composition and also contains glutamatergic and GABAergic neurons (Nair-Roberts et al., 2008). In particular, our anatomical mask does not distinguish pars compacta and pars reticularis of the SN; the latter contains GABAergic neurons whose contribution to the blood oxygen level-dependent (BOLD) signal is not well understood (Logothetis, 2008). While multimodal investigations have demonstrated good correspondence between striatal DA release and BOLD signal in response to reward PEs or novel stimuli (see Düzzel et al., 2009 for review), this relation still remains to be established for sensory PEs. Similar caveats apply to our findings on the basal forebrain, which also contains other neurons than only cholinergic ones (Zaborszky et al., 2008).

With this caveat in mind, our study suggests that in humans the dopaminergic midbrain may not only encode PEs about reward, but also precision-weighted PEs about purely sensory outcomes. To our knowledge, similar midbrain activations have not been reported in previous studies on reward-unrelated learning (e.g., d'Acromont et al., 2013; Gläscher et al., 2010). Notably, our experiments were designed to detect brainstem activations, including an optimized fMRI sequence and careful correction for physiological (cardiac and respiratory) noise. Last but not least, our studies had considerably larger sample sizes, and consequently higher statistical power, than previous fMRI studies on reward-unrelated learning.

It is worth mentioning that the recent study by Ide et al. (2013), which reports activity for unsigned PEs (Bayesian surprise) in

ACC during a Go/NoGo task, does show a midbrain activation (their Figure 3); however, this is not a sensory PE but reflects a main effect of stop versus go trials. Another recent fMRI study (Payzan-LeNestour et al., 2013) on neuromodulatory mechanisms during learning focused on different forms of uncertainty and on the noradrenergic system but did not report any findings related to PEs, nor to DA or ACh, as in this study.

In animal studies, disentangling responses to sensory and reward aspects of stimuli is often difficult because stimulus-bound reward are required to maintain motivation (Maunsell, 2004). In our study, however, the finding of a sensory PE response in the midbrain cannot easily be explained by any (hidden) reward effect since we controlled for the potential influence of reward in two ways. In the first fMRI study, we orthogonalized reward delivery to the task-relevant predictions about visual stimuli; additionally, we verified by model comparison that our subjects' decisions were unlikely to be driven by reward predictions. In our second fMRI study, we entirely omitted any reward, yet found exactly the same VTA/SN response to PEs about visual stimuli as in the first fMRI study (Figure 3).

Beyond PEs about visual stimulus category, our hierarchical model also enabled us to examine higher-level PEs. Specifically, in both fMRI studies, we found a significant activation of the cholinergic basal forebrain by the precision-weighted PE ε_3 about conditional probabilities (of the visual stimulus given the auditory cue) or, equivalently, cue-outcome contingencies. This finding provides a new perspective on possible computational roles of ACh. In the previous literature, the release of acetylcholine has been associated with a diverse range of functions, including working memory (Hasselmo, 2006), attention (Demeter and Sarter, 2013), or learning (Dayan, 2012; Doya, 2002).

A recent influential proposal was that ACh levels may encode the degree of "expected uncertainty" (EU) (Yu and Dayan, 2002, 2005). Operationally, EU was defined (in slightly different ways across articles) in reference to a hidden Markov model representing the relation between contextual states, cue validity, and sensory events. Notably, Yu and Dayan (2002, 2005) implicitly defined EU as a high-level PE, in the sense that it represents the difference between a conditional probability (degree of cue validity) and certainty. Despite clear differences in the underlying models, this definition is conceptually related to ε_3 in our model (see Supplemental Experimental Procedures, section A, for details) that we found was encoded by activity in the basal forebrain. Our empirical findings thus complement the previous theoretical arguments by Yu and Dayan (2002, 2005), offering a

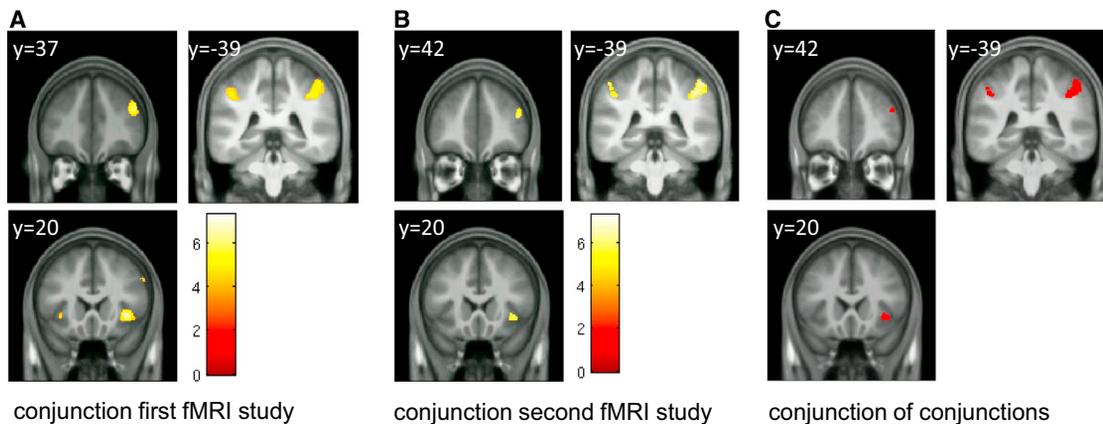


Figure 4. Overlap of Activations by Task Execution Per Se and ϵ_2

Conjunction analysis (“logical AND,” conjunction null hypothesis) of the contrasts testing for trial events and for the precision-weighted prediction error about stimulus visual outcome, ϵ_2 .

(A) First fMRI study.

(B) Second fMRI study.

(C) Results of a double conjunction, i.e., the conjunction of the results from (A) and (B) across both studies.

related perspective on ACh function by conceptualizing it as a precision-weighted PE about conditional probabilities (cue-outcome contingencies). The precision-weighting of this PE also relates our results on basal forebrain activation to the previous suggestion of a link between ACh and learning rate (Doya, 2002). This is because, in its numerator, ψ_3 (the precision weight of ϵ_3) contains an equivalent to a dynamic learning rate (Preuschhoff and Bossaerts, 2007) for updating cue-outcome contingencies (see Equation A.10 in the Supplemental Experimental Procedures, section A and Equation 27 in Mathys et al., 2011).

In summary, our findings are important in two ways. First, they provide empirical support for the importance of precision-weighted PEs as postulated by the Bayesian brain hypothesis. Furthermore, they contribute to the ongoing debate about the computational roles of neuromodulatory transmitters (Dayan, 2012), suggesting a more general role for DA than only encoding reward-related PEs and providing empirical evidence for ACh involvement in representing higher-order PEs (about conditional probabilities). Our results are compatible with the notion that multiple neuromodulators may be involved in the precision-weighting of PEs (Friston, 2009), but suggest separable roles for DA and ACh at different hierarchical levels of learning.

In future analyses, we will focus on elucidating how these PEs may be used as “teaching signals” for synaptic plasticity (expressed through changes in effective connectivity; cf. den Ouden et al., 2010). We hope that, eventually, this work will contribute to establishing neurocomputational assays that allow for inference on neuromodulatory function in the brains of individual patients. If successful, this could have far-reaching implications for diagnostic procedures in psychiatry and neurology (Maia and Frank, 2011; Moran et al., 2011; Stephan et al., 2006).

EXPERIMENTAL PROCEDURES

Subjects

This article reports findings obtained from three separate samples of healthy volunteers. The three studies used nearly identical experimental paradigms,

enabling us to test which results would survive replication, both in the presence of monetary reward (behavioral study and first fMRI study) and in their absence (second fMRI study).

The first sample containing 63 male volunteers (mean age \pm SD: 21 ± 2.2 years) was examined behaviorally only. The second sample (48 male volunteers; 23 ± 3.1 years) and third sample (27 male volunteers; 21 ± 2.2 years) underwent both behavioral assessment and fMRI (the third sample corresponded to the placebo group from a pharmacological study whose results will be reported elsewhere). We only employed male participants to exclude variations of hormonal effects on the BOLD signal during the menstrual cycle. The participants were all nonsmokers, without any psychiatric or neurological disorders in their past medical history and were not taking any medication.

All three studies employed a near-identical audio-visual associative learning task (see below). Prior to data analysis, each subject’s data was examined for invalid trials. These were defined as missed responses or as trials with excessively long reaction times (late responses; $>1,100$ ms in the behavioral study, $>1,300$ ms in the first fMRI study, and $>1,500$ ms in the second fMRI study). Subjects with more than 20% invalid trials or less than 65% correct responses were excluded from further analyses. These criteria led to the exclusion of 17 participants in the behavioral study and three participants in the first fMRI study; no participants were excluded from the second fMRI study. As a consequence, the final data analysis included 46 subjects from the behavioral study (21 ± 2.3 years), 45 subjects from the first fMRI study (23 ± 3.0 years), and 27 subjects from the second fMRI study (21 ± 2.2 years). All participants gave written informed consent before the study, which had received ethics approval by the local responsible authorities (Kantonale Ethikkommission, KEK 2010-0312/3 for the behavioral and first fMRI study, KEK 2011-0101/3 for the second fMRI study).

Experimental Design: Associative Learning Task

A cross-modal associative learning task (audio-visual stimulus-stimulus learning [SSL]) was used in all three studies (Figure 1) where participants had to learn the predictive strength of auditory cues and predict a subsequent visual stimulus. Notably, this prediction was explicit and indicated by button press before the visual stimulus appeared. The task design was near-identical in all three studies; the only variations concerned: (1) response interval (800 ms in the behavioral study, 1,000 ms and 1,200 ms in the first and second fMRI studies), (2) duration of the visual outcome presentation (150 ms in the behavioral and first fMRI study, 300 ms in the second fMRI study), and (3) the presence or absence of trial-wise monetary reward (see below).

Stimuli were presented using Cogent2000 (<http://www.vislab.ucl.ac.uk/Cogent/index.html>). Trials were presented with a randomized intertrial interval

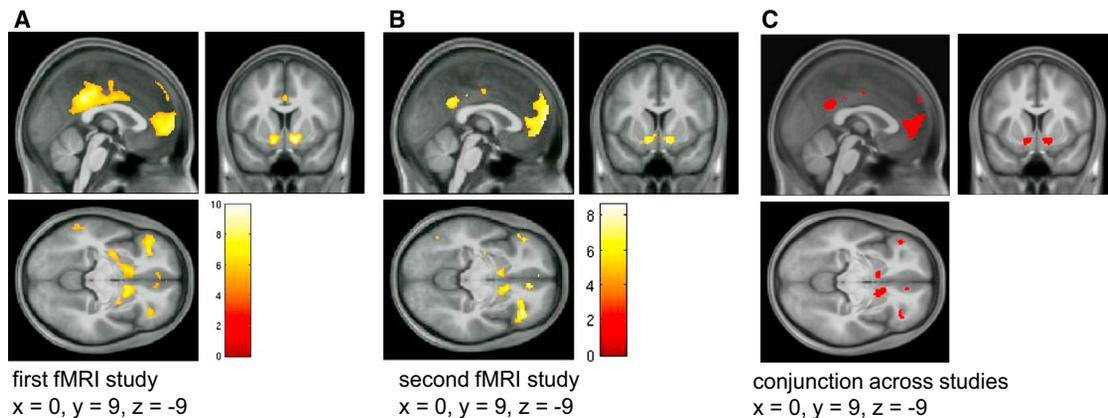


Figure 5. Choice Prediction Error

Activations by choice prediction error, ϵ_{ch} , in the first (A) and the second fMRI study (B). Both activation maps are shown at a threshold of $p < 0.05$, FWE corrected for multiple comparisons across the whole brain. To highlight replication across studies, (C) shows the results of a “logical AND” conjunction, illustrating voxels that were significantly activated in both studies.

See also Table S7.

(ITI) of 1.5–2.5 s. At the beginning of each trial, participants heard one of two possible auditory cues for 300 ms, a high (576 Hz) or a low tone (352 Hz). To ensure that both tones were perceived equally loudly, subjects performed an initial psychophysical matching task in which they had to adapt the volumes until they perceived both cues as equally loud (cf. den Ouden et al., 2010).

Following the cue, participants had to signal their prediction by button press (right index and middle finger), as quickly and as accurately as possible, which of two possible visual outcome categories (houses and faces) would follow. These comprised a small subset of stimuli (two to four) from our previous work (den Ouden et al., 2010).

Critically, in our task the cue–outcome association strength changed over time (i.e., reversal learning), including strongly predictive (probabilities of 0.9 and 0.1), moderately predictive (0.7, 0.3), and nonpredictive cues (0.5). Each subject completed 320 trials, divided into ten blocks of different association strengths. Our stimulus sequence (Figure 1B) had two key features: both block length (24 to 40 trials) and magnitude of changes in cue–outcome contingency varied unpredictably across blocks. Over the experiment, this led to changes in two related variables of interest: (1) volatility, and (2) precision-weighted prediction error about cue–outcome contingency ϵ_3 (a proxy to “expected uncertainty”; see Discussion). Please note that in our modeling framework, there is a formal connection between the concepts of volatility and expected uncertainty: ϵ_3 depends on the previous estimate of log-volatility μ_3 ; in turn, ϵ_3 determines the updating of μ_3 (see Equations A.10 and A.11 in the Supplemental Experimental Procedures).

The probability sequence was pseudorandom and fixed across subjects to ensure comparability of the induced learning process and thus model parameter estimates. Subjects were informed in which range the probabilities could change but not about their order or possible values. Also, as in previous work (den Ouden et al., 2010), they were explicitly instructed that the conditional probabilities were coupled as follows (f : face; h : house; \uparrow : high tone; \downarrow : low tone):

$$p(f|\uparrow) = 1 - p(h|\uparrow) = p(h|\downarrow) = 1 - p(f|\downarrow) \quad (\text{Equation 1})$$

We ensured that the marginal probabilities of face and house outcomes were identical across the experiment and could thus not bias the participants’ predictions. This was achieved by requiring that (1) the probability of one outcome given a particular cue was the same as the probability of the other outcome given the other cue (Equation 1), and (2) in each block, both cue types appeared equally often and in random order. With these two manipulations, we ensured that, on average, before the cue was presented, the a priori probability of a face or a house occurring was 50% each. Thus, on any given trial, it was

not possible to make an informed prediction about the outcome before having heard the cue.

In the behavioral study and first fMRI study, each trial was associated with a potential monetary reward. Specifically, at the end of each trial the visual outcome was presented for 150 ms in the center of the image, together with a coin (5 CHF or 0.05 CHF) randomly located in one of the corners (Figure 1A). Critically, reward size was uncorrelated to the visual outcome to be predicted. In other words, high and low reward appeared randomly on 50% of the trials each, ensuring that any cue would predict any reward with 50% probability. At the end of the experiment, we applied a simple pay-out rule: 100 low-rewarding trials and one high-rewarding trial were randomly chosen, and the summed reward from correct trials only was paid out (note that the maximal possible net value for both low- and high-reward trials was identical, i.e., 5 CHF). This procedure was used to motivate the participants to deliver constantly high performance throughout the experiment: by minimizing the number of incorrect predictions about the visual outcome, participants would maximize their expected total reward.

Although we instructed our participants explicitly that the reward sequence was random and could not be learned, one might wonder whether some subjects might nevertheless have tried to predict upcoming reward instead of visual outcomes. We therefore also modeled any putative learning of the orthogonal reward and performed model comparison to quantify whether predictions of visual outcomes or reward would better explain the subjects’ observed behavior (see below). Finally, in the second fMRI study, we omitted reward. This enabled us to examine experimentally whether behavior and fMRI activations would remain identical when monetary reward were absent.

Hierarchical Gaussian Filter

For behavioral data analysis, we applied a Hierarchical Gaussian Filter (HGF) that describes learning at multiple levels and allows for inference on an agent’s belief about the causes of its sensory inputs (Mathys et al., 2011). The HGF rests on a variational approximation to ideal hierarchical Bayes, which conveys two major advantages. First, the HGF allows for individualized Bayesian learning: it contains subject-specific parameters that couple the different levels of the hierarchy and determine the individual learning process. Second, the update equations are analytic and contain reinforcement learning as a special case, with precision-weighted prediction errors (PEs) driving belief updating at the different levels of the hierarchical model (see below).

Here, we implemented a three-level HGF as described by Mathys et al. (2011) and summarized by Figure 1C, using the HGF Toolbox v2.1 that is available as open source code (<http://www.translationalneuromodeling.org/tapas>). The first level of this model represents a sequence of environmental states x_1 (here: whether a face or house was presented), the second level represents the

Table 2. Basal Forebrain Activations by ε_3

fMRI Study 1	X	y	z	t Score	fMRI Study 2	x	Y	z	t Score
ε_3 : Positive Correlation					ε_3 : Positive Correlation				
Basal forebrain	0	10	-8	4.22	Basal forebrain	0	10	-8	5.02

MNI coordinates and t values for regions activated by ε_3 , the precision-weighted PE about stimulus probability in the first and second fMRI study. Only those activations are listed that were replicated across studies.

cue-outcome contingency x_2 (i.e., the conditional probability, in logit space, of the visual target given the auditory cue), and the third level the log-volatility of the environment x_3 . Each of these hidden states is assumed to evolve as a Gaussian random walk, such that its variance depends on the state at the next higher level (Figure 1C):

$$p(x_1|x_2) = s(x)^{x_1} (1 - s(x))^{1-x_1} = \text{Bernoulli}(x_1; s(x_2)), \quad (\text{Equation 2})$$

$$p(x_2^{(k)} | x_2^{(k-1)}, x_3^{(k)}) = N(x_2^{(k)}; x_2^{(k-1)}, \exp(\kappa x_3^{(k)} + \omega)), \quad (\text{Equation 3})$$

$$p(x_3^{(k)} | x_3^{(k-1)}, \vartheta) = N(x_3^{(k)}; x_3^{(k-1)}, \vartheta), \quad (\text{Equation 4})$$

where $s(\cdot)$ is a sigmoid function.

In Equations 2, 3, and 4, ϑ determines the speed of learning about the log-volatility of the environment; κ determines how strongly the second and third levels are coupled and thus how much the estimated environmental volatility affects the learning rate at the second level; and ω is a constant component of the step size at the second level. Finally, the predicted probability of a visual target given the auditory cue (i.e., the posterior mean of x_2) is linked to trial-wise predictions of visual stimulus category by means of a softmax function with parameter ζ (encoding decision noise). Our three-level HGF for categorical outcomes thus has four parameters. In our implementation, three of them were free (ϑ , κ , ζ), whereas ω was fixed to -4 in our analyses in order to ensure model identifiability.

Importantly, the variational approximation underlying the HGF provides analytic update equations that share a general form: At any level i of the hierarchy, the update of the belief on trial k (i.e., posterior mean $\mu_i^{(k)}$ of the state x_i) is proportional to the precision-weighted prediction error (PE) $\varepsilon_i^{(k)}$. This weighted PE is the product of the PE $\delta_{i-1}^{(k)}$ from the level below and a precision ratio $\psi_i^{(k)}$:

$$\mu_i^{(k+1)} - \mu_i^{(k)} \propto \psi_i^{(k)} \delta_{i-1}^{(k)} = \varepsilon_i^{(k)}, \quad (\text{Equation 5})$$

$$\psi_i^{(k)} = \frac{\hat{\pi}_{i-1}^{(k)}}{\pi_i^{(k)}}, \quad (\text{Equation 6})$$

where $\hat{\pi}_{i-1}^{(k)}$ represents the precision of the prediction about input from the level below and $\pi_i^{(k)}$ encodes the precision of the belief at the current level. The form of this general update equation is reminiscent of RL models. Specifically, the precision-weighting can be understood as (component of) a dynamic learning rate (cf. Preuschoff and Bossaerts, 2007); see Mathys et al. (2011) and section A of the Supplemental Experimental Procedures for details.

In our three-level HGF, two precision-weighted PEs ε_i occur. At the second level, ε_2 is the precision-weighted PE about visual stimulus outcome that serves to update the estimate of x_2 (the cue-outcome contingency in logit space). At the third level, ε_3 is the precision-weighted PE about cue-outcome contingency that is proportional to the update of x_3 (environmental log-volatility). These are the two quantities of interest that the fMRI analyses in this article focus on. For the exact equations, see the Supplemental Experimental Procedures, section A.

fMRI Data Acquisition and Analysis

The experiment was conducted on a 3T Philips Achieva MR Scanner at the SNS Lab, using an eight channel SENSE head-coil. Structural images were ac-

quired using a T₁-weighted sequence. For functional imaging, 500 whole-brain images were acquired in the first fMRI study and 550 images in the second fMRI study, using a T₂*-weighted echo-planar imaging sequence that had been optimized for brain stem imaging (slice thickness: 3 mm; in-plane resolution: 2 × 2 mm; interslice gap: 0.6 mm; ascending continuous in-plane acquisition; TR = 2,500 ms; TE = 36 ms; flip angle = 90°; field of view = 192 × 192 × 118 mm; SENSE factor = 2; EPI factor = 51). In order to reduce field inhomogeneities a second order pencil-beam volume shim (provided by Philips) was applied during the functional acquisition. Functional data acquisition lasted ~21 min. During fMRI data acquisition, respiratory and cardiac activity was acquired using a breathing belt and an electrocardiogram, respectively.

fMRI data were analyzed using statistical parametric mapping (SPM8). Following motion correction of the functional images and coregistration to the structural image, we warped both functional and structural images to MNI space using the “New Segment” toolbox in SPM; see Appendix A in Ashburner and Friston (2005). The functional images were smoothed applying a 6 mm full-width at half maximum Gaussian kernel and resampled to 1.5 mm isotropic resolution. In order to optimize signal-to-noise ratio for critical regions such as the brain stem, we corrected for physiological noise using RETROICOR (Glover et al., 2000) based on an in-house implementation (Kasper et al., 2009) (open source code available at <http://www.translationalneuromodeling.org/tapas>).

For fMRI data analysis, we specified a voxel-wise general linear model (GLM) for each participant. In the first fMRI study, this GLM reflected a 2 × 2 factorial design with visual outcome category (face, house) and incidental reward stimulus (high, low) as factors. In the second fMRI study, reward stimuli were absent; therefore, the GLM only contained the two visual outcome conditions. Additionally, we modeled missed and late responses, respectively, by separate regressors. All regressors were convolved with a canonical hemodynamic function and its temporal derivative. The subject-specific belief trajectories, obtained from the HGF, were used in the GLM as parametric modulators. These variables included (cf. Equations 2, 3, 4, 5, and 6; Figures S1 and S2):

- (1) ε_2 , the precision-weighted PE about visual stimulus outcome (that serves to update the estimate of visual stimulus probabilities in logit space);
- (2) ε_3 , the precision-weighted PE about cue-outcome contingency (that serves to update the estimate of log-volatility);
- (3) ψ_2 , precision weight at the second level; this corresponds to the learning rate by which estimates of cue-outcome contingency are updated;
- (4) ψ_3 , precision weight at the third level; this is proportional to the learning rate by which log-volatility estimates are updated;
- (5) μ_3 , the predicted log-volatility; and
- (6) ε_{ch} , the choice prediction error.

Importantly, choice PE ε_{ch} and precision-weighted outcome PE ε_2 have distinct definitions (see sections A and B of the Supplemental Experimental Procedures for mathematical details). The choice PE ε_{ch} is the difference between the correctness of the subject’s choice (1 if choice was correct, 0 otherwise) and the a priori probability of this choice being correct. This PE is positive when the subject’s choice was correct and negative when it was wrong. In contrast, ε_2 multiplies two components (Equations 5 and 6): (1) the precision weight $\psi_i^{(k)}$ (that is always positive), and (2) δ_1 , the difference between the actual visual stimulus outcome and its a priori probability (also always positive); the latter corresponds to Bayesian surprise and is bounded between 0 and 1.

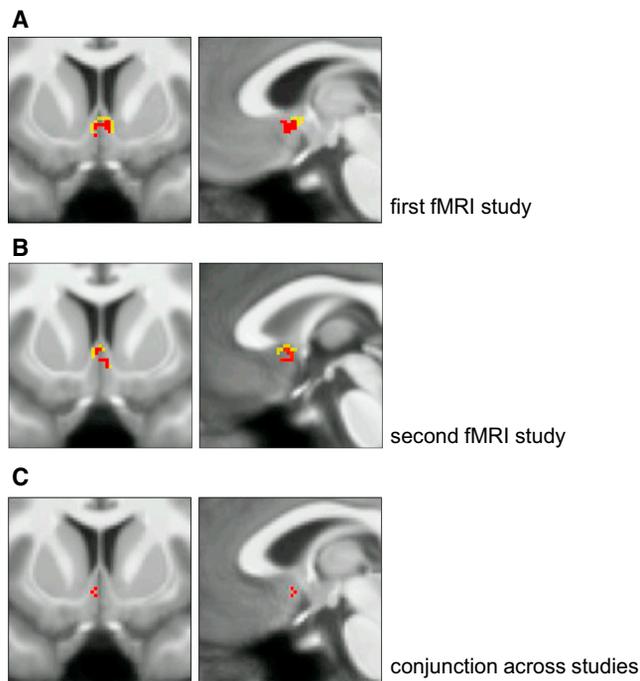


Figure 6. Basal Forebrain Activations by ε_3

Activation of the cholinergic basal forebrain associated with precision-weighted prediction error about stimulus probabilities ε_3 within the anatomically defined mask. For visualization of the activation area we overlay the results thresholded at $p < 0.05$ FWE corrected for the entire anatomical mask (red) on the results thresholded at $p < 0.001$ uncorrected (yellow) in the first (A: $x = 3, y = 9, z = -8$) and the second fMRI study (B: $x = 0, y = 10, z = -8$). (C) The conjunction analysis (“logical AND”) across both studies ($x = 2, y = 11, z = -8$).

Importantly, the GLM used all computational trajectories in their original form, without any orthogonalization. Thus, we did not impose any judgment on the relative importance of regressors for explaining the fMRI data. Also, the timings of our events were chosen such that PE estimates were time-locked to the visual outcome at the end of the trial; prediction and precision regressors spanned the entire trial and changed at outcome, according to the update induced by the PE.

Our subject-specific (first-level) GLM also included regressors representing potential confounds. This included the realignment parameters (encoding head movements) and their first derivative, a regressor marking scans with >1 mm scan-to-scan head movement, and physiological confound variables (cardiac activity and breathing), provided by RETROICOR.

In addition to whole-brain analyses, we performed ROI analyses based on anatomical masks of dopaminergic and cholinergic nuclei. These included (1) the dopaminergic midbrain (SN and VTA), (2) the cholinergic basal forebrain, (3) cholinergic nuclei in the tegmentum of the brainstem, i.e., the pedunculopontine tegmental (PPT) and laterodorsal tegmental (LDT) nuclei. For the VTA/SN, we used an anatomical atlas based on magnetization transfer-weighted structural MR images (Bunzeck and Düzel, 2006). The basal forebrain was defined using the maximum probability map from a probabilistic cytoarchitectonic atlas warped into MNI space (Eickhoff et al., 2005; Zaborzky et al., 2008). This map included the different compartments of the basal forebrain with cholinergic neurons (septum, the diagonal band of Broca, and subpallidal regions including the basal nucleus of Meynert). Given the lack of a published atlas for PPT and LDT, we used MRICron to manually trace the region of these nuclei according to anatomical landmarks from the literature (Naidich et al., 2009; Zrinzo et al., 2011). Note that we did not use these anatomical masks separately to test for activations; instead, all regions mentioned above were combined into a single mask image, and each ROI analysis used this combined mask for multiple comparison correction.

Contrasts of interest testing for each of the parametric modulators specified above were defined at the first level and entered into second level ANOVAs to allow for inference at the group level. We tested for both positive and negative effects of our parametric modulators. Please note that we only report results that (1) survived stringent family-wise error correction (FWE) at the voxel level ($p < 0.05$), based on Gaussian random field theory (Worsley et al., 1996), across the whole brain and within ROIs, respectively, and (2) were replicated in both fMRI studies. Replicability was assessed by testing the conjunction null hypothesis, i.e., a voxel-wise “logical AND” analysis (Nichols et al., 2005). In the main text of this article, we focus on activations related to prediction errors; for other findings related to the remaining regressors, see Supplemental Experimental Procedures (Figure S3; Tables S3, S4, S5, and S6).

Bayesian Model Selection

To disambiguate alternative explanations (models) for the participants’ behavior, we used Bayesian model selection (BMS). BMS is a standard approach in machine learning and neuroimaging (MacKay, 1992; Penny et al., 2004) for comparing competing models that describe how neurophysiological or behavioral responses were generated. BMS evaluates the relative plausibility of competing models in terms of their log-evidences. The log-evidence of a model corresponds to the negative surprise about the data, given the model, and quantifies the trade-off between accuracy (fit) and complexity of a model. Here, we used a recently developed random effects BMS method to account for potential interindividual variability in our sample (Penny et al., 2010; Stephan et al., 2009), quantifying the posterior probabilities of five competing models (see Results and Supplemental Experimental Procedures for details).

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, three figures, and seven tables and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2013.09.009>.

ACKNOWLEDGMENTS

We acknowledge support by the Zurich Neuroscience Centre (S.I., K.E.S.), the René and Susanne Braginsky Foundation (K.E.S.), KFSP “Molecular Imaging,” and SystemsX.ch (K.E.S.). We are very grateful to Simon Eickhoff and Emrah Düzel for providing us with the anatomical masks for delineating the basal forebrain and VTA/SN, respectively.

Accepted: September 3, 2013

Published: October 16, 2013

REFERENCES

- Ashburner, J., and Friston, K.J. (2005). Unified segmentation. *Neuroimage* 26, 839–851.
- Ashby, W.R. (1952). *Design for a Brain*. (London: Chapman & Hall).
- Behrens, T.E., Woolrich, M.W., Walton, M.E., and Rushworth, M.F. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221.
- Bunzeck, N., and Düzel, E. (2006). Absolute coding of stimulus novelty in the human substantia nigra/VTA. *Neuron* 51, 369–379.
- Bunzeck, N., Dayan, P., Dolan, R.J., and Düzel, E. (2010). A common mechanism for adaptive scaling of reward and novelty. *Hum. Brain Mapp.* 31, 1380–1394.
- d’Acremont, M., Fornari, E., and Bossaerts, P. (2013). Activity in inferior parietal and medial prefrontal cortex signals the accumulation of evidence in a probability learning task. *PLoS Comput. Biol.* 9, e1002895.
- D’Ardenne, K., McClure, S.M., Nystrom, L.E., and Cohen, J.D. (2008). BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science* 319, 1264–1267.

- Daw, N.D., and Doya, K. (2006). The computational neurobiology of learning and reward. *Curr. Opin. Neurobiol.* *16*, 199–204.
- Dayan, P. (2012). Twenty-five lessons from computational neuromodulation. *Neuron* *76*, 240–256.
- Dayan, P., Hinton, G.E., Neal, R.M., and Zemel, R.S. (1995). The Helmholtz machine. *Neural Comput.* *7*, 889–904.
- Demeter, E., and Sarter, M. (2013). Leveraging the cortical cholinergic system to enhance attention. *Neuropharmacology* *64*, 294–304.
- den Ouden, H.E., Friston, K.J., Daw, N.D., McIntosh, A.R., and Stephan, K.E. (2009). A dual role for prediction error in associative learning. *Cereb. Cortex* *19*, 1175–1185.
- den Ouden, H.E., Daunizeau, J., Roiser, J., Friston, K.J., and Stephan, K.E. (2010). Striatal prediction error modulates cortical coupling. *J. Neurosci.* *30*, 3210–3219.
- Diuk, C., Tsai, K., Wallis, J., Botvinick, M., and Niv, Y. (2013). Hierarchical learning induces two simultaneous, but separable, prediction errors in human basal ganglia. *J. Neurosci.* *33*, 5797–5805.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Netw.* *15*, 495–506.
- Doya, K., Ishii, S., Pouget, A., and Rao, R.P. (2011). *Bayesian Brain: Probabilistic Approaches to Neural Coding.* (Cambridge, MA: MIT Press).
- Düzel, E., Bunzeck, N., Guitart-Masip, M., Wittmann, B., Schott, B.H., and Tobler, P.N. (2009). Functional imaging of the human dopaminergic midbrain. *Trends Neurosci.* *32*, 321–328.
- Egner, T., Monti, J.M., and Summerfield, C. (2010). Expectation and surprise determine neural population responses in the ventral visual stream. *J. Neurosci.* *30*, 16601–16608.
- Eickhoff, S.B., Stephan, K.E., Mohlberg, H., Grefkes, C., Fink, G.R., Amunts, K., and Zilles, K. (2005). A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *Neuroimage* *25*, 1325–1335.
- FitzGerald, T.H., Friston, K.J., and Dolan, R.J. (2012). Action-specific value signals in reward-related regions of the human brain. *J. Neurosci.* *32*, 16417–16423a.
- Fletcher, P.C., Anderson, J.M., Shanks, D.R., Honey, R., Carpenter, T.A., Donovan, T., Papadakis, N., and Bullmore, E.T. (2001). Responses of human frontal cortex to surprising events are predicted by formal associative learning theory. *Nat. Neurosci.* *4*, 1043–1048.
- Friston, K. (2005). A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* *360*, 815–836.
- Friston, K. (2009). The free-energy principle: a rough guide to the brain? *Trends Cogn. Sci.* *13*, 293–301.
- Friston, K., Kilner, J., and Harrison, L. (2006). A free energy principle for the brain. *J. Physiol. Paris* *100*, 70–87.
- Friston, K.J., Shiner, T., FitzGerald, T., Galea, J.M., Adams, R., Brown, H., Dolan, R.J., Moran, R., Stephan, K.E., and Bestmann, S. (2012). Dopamine, affordance and active inference. *PLoS Comput. Biol.* *8*, e1002327.
- Gläscher, J., Daw, N., Dayan, P., and 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.
- Glover, G.H., Li, T.Q., and Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magn. Reson. Med.* *44*, 162–167.
- Harrison, L.M., Stephan, K.E., Rees, G., and Friston, K.J. (2007). Extra-classical receptive field effects measured in striate cortex with fMRI. *Neuroimage* *34*, 1199–1208.
- Hasselmo, M.E. (2006). The role of acetylcholine in learning and memory. *Curr. Opin. Neurobiol.* *16*, 710–715.
- Horvitz, J.C. (2000). Mesolimbocortical and nigrostriatal dopamine responses to salient non-reward events. *Neuroscience* *96*, 651–656.
- Ide, J.S., Shenoy, P., Yu, A.J., and Li, C.S. (2013). Bayesian prediction and evaluation in the anterior cingulate cortex. *J. Neurosci.* *33*, 2039–2047.
- Kasper, L., Marti, S., Vannesjö, S.J., Hutton, C., Dolan, R., Weiskopf, N., Prüssmann, K.P., and Stephan, K.E. (2009). Cardiac artefact correction for human brainstem fMRI at 7T. *Neuroimage* *47*(Supplement 1), S100.
- Klein-Flügge, M.C., Hunt, L.T., Bach, D.R., Dolan, R.J., and Behrens, T.E. (2011). Dissociable reward and timing signals in human midbrain and ventral striatum. *Neuron* *72*, 654–664.
- Knill, D.C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* *27*, 712–719.
- Kok, P., Rahnev, D., Jehee, J.F., Lau, H.C., and de Lange, F.P. (2012). Attention reverses the effect of prediction in silencing sensory signals. *Cereb. Cortex* *22*, 2197–2206.
- Körding, K.P., and Wolpert, D.M. (2006). Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* *10*, 319–326.
- Law, C.T., and Gold, J.I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nat. Neurosci.* *12*, 655–663.
- Logothetis, N.K. (2008). What we can do and what we cannot do with fMRI. *Nature* *453*, 869–878.
- MacKay, D.J.C. (1992). Bayesian interpolation. *Neural Comput.* *4*, 415–447.
- Maia, T.V., and Frank, M.J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nat. Neurosci.* *14*, 154–162.
- Mathys, C., Daunizeau, J., Friston, K.J., and Stephan, K.E. (2011). A bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* *5*, 39.
- Maunsell, J.H. (2004). Neuronal representations of cognitive state: reward or attention? *Trends Cogn. Sci.* *8*, 261–265.
- McClure, S.M., Berns, G.S., and Montague, P.R. (2003). Temporal prediction errors in a passive learning task activate human striatum. *Neuron* *38*, 339–346.
- Montague, P.R., Hyman, S.E., and Cohen, J.D. (2004). Computational roles for dopamine in behavioural control. *Nature* *431*, 760–767.
- Moran, R.J., Symmonds, M., Stephan, K.E., Friston, K.J., and Dolan, R.J. (2011). An in vivo assay of synaptic function mediating human cognition. *Curr. Biol.* *21*, 1320–1325.
- Murray, G.K., Corlett, P.R., Clark, L., Pessiglione, M., Blackwell, A.D., Honey, G., Jones, P.B., Bullmore, E.T., Robbins, T.W., and Fletcher, P.C. (2008). Substantia nigra/ventral tegmental reward prediction error disruption in psychosis. *Mol. Psychiatry* *13*, 267–276.
- Naidich, T.P., Duvernoy, H.M., Delman, B.N., Sorensen, A.G., Kollias, S.S., and Haacke, E.M. (2009). *Duvernoy’s Atlas of the Human Brain Stem and Cerebellum: High-Field MRI, Surface Anatomy, Internal Structure, Vascularization and 3 D Sectional Anatomy.* (Vienna: Springer Verlag).
- Nair-Roberts, R.G., Chatelain-Badie, S.D., Benson, E., White-Cooper, H., Bolam, J.P., and Ungless, M.A. (2008). Stereological estimates of dopaminergic, GABAergic and glutamatergic neurons in the ventral tegmental area, substantia nigra and retrorubral field in the rat. *Neuroscience* *152*, 1024–1031.
- Nichols, T., Brett, M., Andersson, J., Wager, T., and Poline, J.-B. (2005). Valid conjunction inference with the minimum statistic. *Neuroimage* *25*, 653–660.
- O’Doherty, J.P., Dayan, P., Friston, K., Critchley, H., and Dolan, R.J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron* *38*, 329–337.
- Payzan-LeNestour, E., Dunne, S., Bossaerts, P., and O’Doherty, J.P. (2013). The neural representation of unexpected uncertainty during value-based decision making. *Neuron* *79*, 191–201.
- Penny, W.D., Stephan, K.E., Mechelli, A., and Friston, K.J. (2004). Comparing dynamic causal models. *Neuroimage* *22*, 1157–1172.
- Penny, W.D., Stephan, K.E., Daunizeau, J., Rosa, M.J., Friston, K.J., Schofield, T.M., and Leff, A.P. (2010). Comparing families of dynamic causal models. *PLoS Comput. Biol.* *6*, e1000709.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R.J., and Frith, C.D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature* *442*, 1042–1045.

- Preusschoff, K., and Bossaerts, P. (2007). Adding prediction risk to the theory of reward learning. *Ann. N Y Acad. Sci.* 1104, 135–146.
- Rao, R.P.N., and Ballard, D.H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87.
- Redgrave, P., and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? *Nat. Rev. Neurosci.* 7, 967–975.
- Schonberg, T., O'Doherty, J.P., Joel, D., Inzelberg, R., Segev, Y., and Daw, N.D. (2010). Selective impairment of prediction error signaling in human dorso-lateral but not ventral striatum in Parkinson's disease patients: evidence from a model-based fMRI study. *Neuroimage* 49, 772–781.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80, 1–27.
- Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Stephan, K.E., Baldeweg, T., and Friston, K.J. (2006). Synaptic plasticity and dysconnection in schizophrenia. *Biol. Psychiatry* 59, 929–939.
- Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., and Friston, K.J. (2009). Bayesian model selection for group studies. *Neuroimage* 46, 1004–1017.
- Summerfield, C., and Koehlin, E. (2008). A neural representation of prior information during perceptual inference. *Neuron* 59, 336–347.
- Sutton, R.S. (1992). Gain adaptation beats least squares? *Proceedings of the Seventh Yale Workshop on Adaptive and Learning Systems* 161–166.
- Tobler, P.N., Fiorillo, C.D., and Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science* 307, 1642–1645.
- Vickery, T.J., Chun, M.M., and Lee, D. (2011). Ubiquity and specificity of reinforcement signals throughout the human brain. *Neuron* 72, 166–177.
- Worsley, K.J., Marrett, S., Neelin, P., Vandal, A.C., Friston, K.J., and Evans, A.C. (1996). A unified statistical approach for determining significant signals in images of cerebral activation. *Hum. Brain Mapp.* 4, 58–73.
- Wunderlich, K., Symmonds, M., Bossaerts, P., and Dolan, R.J. (2011). Hedging your bets by learning reward correlations in the human brain. *Neuron* 71, 1141–1152.
- Yu, A.J., and Dayan, P. (2002). Acetylcholine in cortical inference. *Neural Netw.* 15, 719–730.
- Yu, A.J., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692.
- Zaborszky, L., Hoemke, L., Mohlberg, H., Schleicher, A., Amunts, K., and Zilles, K. (2008). Stereotaxic probabilistic maps of the magnocellular cell groups in human basal forebrain. *Neuroimage* 42, 1127–1141.
- Zrinzo, L., Zrinzo, L.V., Massey, L.A., Thornton, J., Parkes, H.G., White, M., Yousry, T.A., Strand, C., Revesz, T., Limousin, P., et al. (2011). Targeting of the pedunculopontine nucleus by an MRI-guided approach: a cadaver study. *J. Neural Transm.* 118, 1487–1495.

Neuron, Volume *80*

Supplemental Information

Hierarchical Prediction Errors in Midbrain and Basal Forebrain during Sensory Learning

Sandra Iglesias, Christoph Mathys, Kay H. Brodersen, Lars Kasper, Marco Piccirelli,
Hanneke E.M. den Ouden, and Klaas E. Stephan

Supplemental Figures

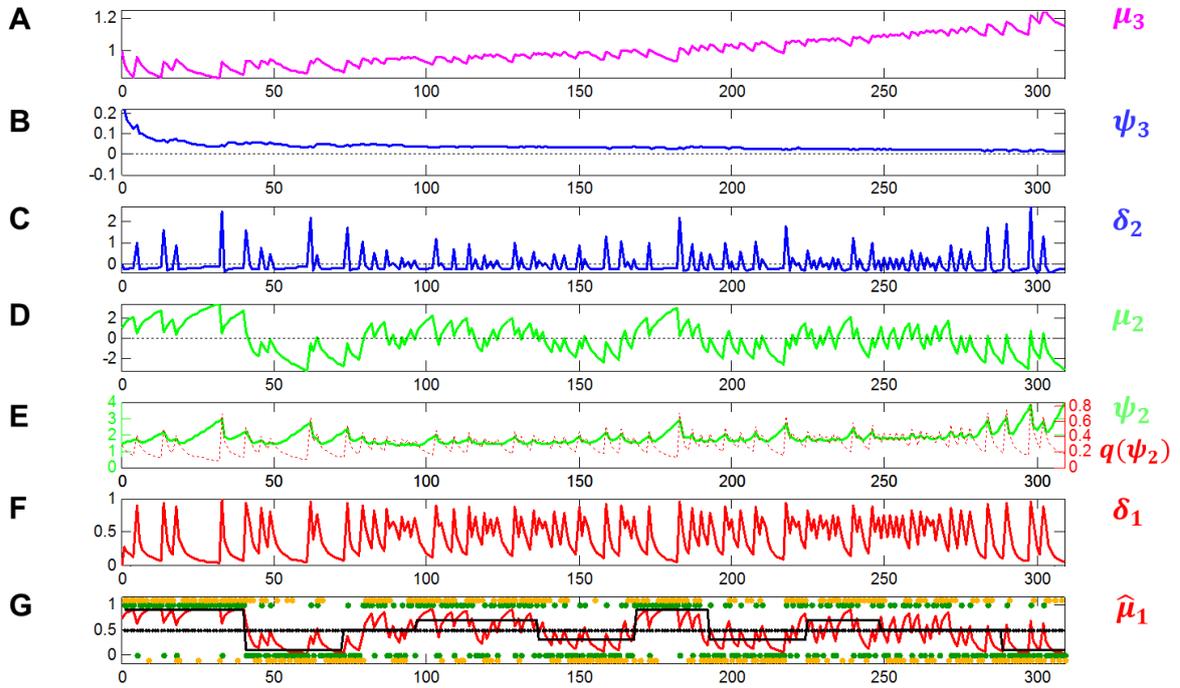


Figure S1: Computational trajectories from subject GE28, with parameter estimates $\vartheta = 4.69 \times 10^{-3}$ and $\kappa = 3.77$. In panels, B-D the dotted line represents the zero line. Related to Figure 1.

(A) Posterior expectation μ_3 of log-volatility x_3 . (B) Precision weight ψ_3 which modulates the impact of prediction error δ_2 on log-volatility updates. (C) Prediction error δ_2 of the cue-outcome contingency (the conditional probability of the visual outcome, given the auditory cue). (D) Posterior expectation μ_2 of cue-outcome contingency x_2 . (E) Precision weight ψ_2 (in green) which modulates the impact of outcome prediction error δ_1 on μ_2 , the conditional probability of the visual outcome given the auditory cue. Since μ_2 is in logit space (cf. Equation A.9), the function of ψ_2 as a dynamic learning rate is more easily visible if transformed according to Equation D.1 (this results in the red line labeled $q(\psi_2)$). (F) Prediction error δ_1 about the trial outcome (visual stimulus category). (G) **Black:** true cue-outcome association strengths. **Red:** posterior expectation of trial outcome, $\hat{\mu}_1$ (i.e., posterior expectation of $p(f|\mathcal{J}=\uparrow)$, see Section B above); this corresponds to a sigmoid transformation of μ_2 in Panel E (cf. Equation A.9). **Green:** trial outcome (or sensory input). **Orange:** subject's observed predictions y .

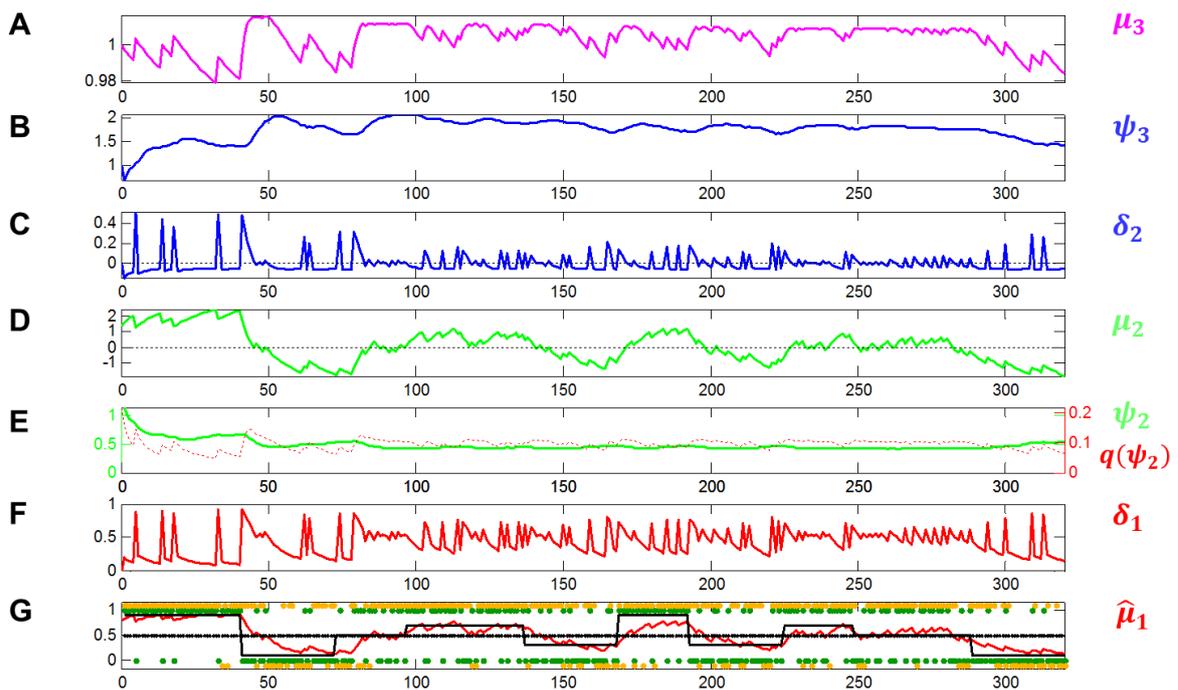


Figure S2: Computational trajectories from subject GL35, with parameter estimates $\vartheta = 2.55 \times 10^{-3}$ and $\kappa = 0.94$. For explanation of the different plots and variables, see legend to Figure S1. Related to Figure 1.

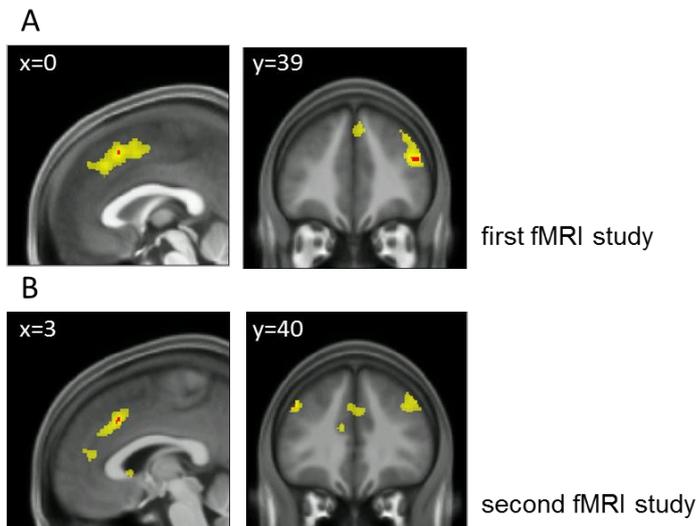


Figure S3: Regions whose BOLD activity correlated significantly with σ_2 . Red: $p < 0.05$, whole-brain FWE corrected at the peak level; yellow: $p < 0.05$, whole-brain FWE corrected at the cluster level under a voxel-level cut-off of $p < 0.001$. Related to Figure 1.

Supplemental Tables

Tables S1-S2: Modeling of behavioral data

BMS results	Behavioral study		fMRI study 1		fMRI study 2	
	PP	XP	PP	XP	PP	XP
HGF1	0.8435	1	0.7422	1	0.7166	1
HGF2	0.0259	0	0.0200	0	-	-
HGF3	0.0361	0	0.1404	0	0.1304	0
Sutton	0.0685	0	0.0710	0	0.0761	0
RW	0.0260	0	0.0264	0	0.0769	0

Table S1: The posterior probabilities (PP) and exceedance probabilities (XP) of all five models. This table is related to Figure 1.

Model	Behavioral study		fMRI study 1		fMRI study 2	
	Mean	SD	Mean	SD	Mean	SD
HGF1						
ϑ	0.0019	0.0017	0.0022	0.0016	0.0019	0.0013
κ	2.835	0.710	2.963	0.702	2.678	0.550
ω	-4	0	-4	0	-4	0
ζ	2.141	1.132	1.680	0.741	2.176	0.930

Table S2: Average maximum a posteriori estimates of the free parameters in the winning model (HGF₁). Related to Figure 1.

Tables S3-S7: fMRI results

A						B					
fMRI study 1	Hemi-sphere	x	y	z	t score	fMRI study 2	Hemi-sphere	x	y	z	t score
ϵ_2 : negative correlation						ϵ_2 : negative correlation					
Ventromedial prefrontal cortex	L	-4	58	0	8.14	Ventromedial prefrontal cortex	L	-6	57	-3	7.70
Mid/ posterior cingulate cortex	L	-4	-31	42	7.95	Posterior cingulate cortex	L	-3	-48	27	4.95
Auditory cortex	R	39	6	9	7.26	Auditory cortex	R	38	4	10	4.93
Orbitofrontal cortex	R	32	32	-12	6.87	Orbitofrontal cortex	R	24	30	-14	5.13

Table S3: Montreal Neurological Institute (MNI) coordinates and t-values of maxima of deactivations by ϵ_2 which were significant ($p < 0.05$, FWE whole-brain corrected) in both fMRI studies. Related to Figure 2.

A						B					
fMRI study 1	Hemi-sphere	x	y	z	t score	fMRI study 2	Hemi-sphere	x	y	z	t score
σ_2 : positive correlation						σ_2 : positive correlation					
Dorsolateral prefrontal cortex	R	45	39	21	5.19	Dorsolateral prefrontal cortex	L	-44	33	36	4.88
Anterior cingulate cortex/ SMA	R	0	20	51	4.95	Anterior cingulate cortex/ SMA	R	3	23	39	4.98

Table S4: Montreal Neurological Institute (MNI) coordinates and t-values of maxima of activations by σ_2 which were significant ($p < 0.05$, FWE whole-brain corrected) in both fMRI studies. Related to Figure 1.

A						B					
fMRI study 1	Hemi-sphere	x	y	z	t score	fMRI study 2	Hemi-sphere	x	y	z	t score
σ_3 : positive correlation						σ_3 : positive correlation					
Lingual gyrus	L	-6	-79	-18	5.88	Lingual gyrus	L	-6	-78	-18	6.27
Lingual gyrus	L	-10	-72	-12	5.44	Lingual gyrus	R	0	-76	-11	6.67
Cerebral cortex	L	-39	-64	-23	5.71	Cerebral cortex	L	-33	-63	-21	5.54

Table S5: Montreal Neurological Institute (MNI) coordinates and t-values of maxima of activations by σ_3 which were significant ($p < 0.05$, FWE whole-brain corrected) in both fMRI studies. Related to Figure 1.

A						B					
fMRI study 1	Hemi-sphere	x	y	z	t score	fMRI study 2	Hemi-sphere	x	y	z	t score
μ_3 : negative correlation						μ_3 : negative correlation					
Calcarine sulcus	L	-8	-85	-3	8.77	Calcarine sulcus	L	-2	70	10	6.32
Lingual gyrus	R	26	-78	-12	6.45	Lingual gyrus	R	22	-76	-15	5.11

Table S6: Montreal Neurological Institute (MNI) coordinates and t-values of maxima of deactivations by μ_3 which were significant ($p < 0.05$, FWE whole-brain corrected) in both fMRI studies. Related to Figure 1.

A						B					
fMRI study 1	Hemi-sphere	x	y	z	t score	fMRI study 2	Hemi-sphere	x	y	z	t score
ϵ_{ch} : positive correlation						ϵ_{ch} : positive correlation					
Nucleus accumbens	R	12	9	-12	9.96	Nucleus accumbens	R	12	6	-12	7.71
Nucleus accumbens	L	-14	6	-14	9.66	Nucleus accumbens	L	-10	8	-12	7.15
Ventromedial prefrontal cortex	L	-4	62	12	9.23	Ventromedial prefrontal cortex	L	-4	60	10	8.62
Posterior cingulate cortex	L	-2	-42	33	9.07	Posterior cingulate cortex	L	-2	-48	30	6.88
Orbitofrontal cortex	L	-33	35	-12	8.11	Orbitofrontal cortex	L	-38	29	-11	5.41
Inferior parietal cortex	L	-45	-66	-46	7.20	Inferior parietal cortex	L	-46	-63	45	5.73
Prefrontal cortex	R	15	44	43	7.16	Prefrontal cortex	R	16	42	48	7.57
Orbitofrontal cortex	R	36	34	-12	6.64	Orbitofrontal cortex	R	39	33	-9	7.95
Putamen	R	30	-16	3	6.37	Putamen	R	30	-19	0	5.52
Inferior temporal gyrus	L	-60	-48	-5	6.17	Inferior temporal gyrus	L	-56	-48	-5	5.32
Auditory cortex	R	38	2	13	6.16	Auditory cortex	R	34	4	10	5.59
Inferior temporal sulcus	L	-58	-16	-12	5.17	Inferior temporal sulcus	L	-58	-10	-15	4.94
Anterior cingulate cortex	R	6	32	12	4.96	Anterior cingulate cortex	L	-6	39	1	6.24

Table S7: Montreal Neurological Institute (MNI) coordinates and t-values of maxima of activations by ϵ_{ch} which were significant ($p < 0.05$, FWE whole-brain corrected) in both fMRI studies. Related to Figure 5.

Supplemental Material for Experimental Procedures

A. Precision-weighted prediction errors in the three-level HGF for categorical outcomes

In this study, we used the three-level HGF described by Mathys et al. (2011) based on the HGF Toolbox v2.1 (which can be obtained freely as part of the open source software TAPAS; www.translationalneuromodeling.org/tapas). As described in the main text, the first level of this model represents a sequence of environmental states x_1 (here: whether a face or house was presented), the second level represents the cue-outcome contingency x_2 (i.e. the conditional probability, in logit space, of the visual stimulus given the auditory cue), and the third level the log-volatility of the environment x_3 . Each of these hidden states is assumed to evolve as a Gaussian random walk, such that its variance depends on the state at the next higher level (see Figure 1C in the main paper):

$$p(x_1 | x_2) = s(x_1)^{x_2} (1 - s(x_2))^{1-x_2} = \text{Bernoulli}(x_1; s(x_2)) \quad (\text{A.1})$$

$$p(x_2^{(k)} | x_2^{(k-1)}, x_3^{(k)}) = \mathcal{N}(x_2^{(k)}; x_2^{(k-1)}, \exp(\kappa x_3^{(k)} + \omega)) \quad (\text{A.2})$$

$$p(x_3^{(k)} | x_3^{(k-1)}, \mathcal{G}) = \mathcal{N}(x_3^{(k)}; x_3^{(k-1)}, \mathcal{G}) \quad (\text{A.3})$$

where k is a trial index and s is a sigmoid function

$$s(x) = \frac{1}{1 + \exp(-x)} \quad (\text{A.4})$$

In the above equations, \mathcal{G} determines the speed of learning about the log-volatility of the environment; κ determines how strongly the second and third levels are coupled and thus how much the estimated environmental volatility affects the learning rate at the second level; ω is a constant component of the step size at the second level.

Under a variational approximation to ideal hierarchical Bayesian learning according to the above equations, analytical update equations can be derived that share a general form: At any level i of the hierarchy, the update of the belief on trial k (i.e., posterior mean $\mu_i^{(k)}$ of the state x_i) is proportional to the precision-weighted prediction error $\varepsilon_i^{(k)}$. This weighted prediction error (PE) is the product of the prediction error $\delta_{i-1}^{(k)}$ from the level below and a precision ratio $\psi_i^{(k)}$:

$$\begin{aligned}\mu_i^{(k+1)} - \mu_i^{(k)} &\propto \psi_i^{(k)} \delta_{i-1}^{(k)} \\ &= \varepsilon_i^{(k)}\end{aligned}\tag{A.5}$$

$$\psi_i^{(k)} = \frac{\hat{\pi}_{i-1}^{(k)}}{\pi_i^{(k)}}\tag{A.6}$$

(A note on notation: we use the superscript $\hat{}$ to refer to “prediction”. That is, $\hat{\mu}_1^{(k)}$ is the prediction on trial k before experiencing the trial outcome, and $\hat{\pi}_i^{(k)}$ is the precision of this prediction.)

The numerator of the precision ratio ψ_i (i.e., $\hat{\pi}_{i-1}^{(k)}$) represents the precision of the prediction about the input from the level below; the higher this precision, the more meaningful a deviation from this prediction is and therefore the higher the impact of the PE $\delta_{i-1}^{(k)}$ on belief updating. Conversely, the denominator ($\pi_i^{(k)}$) encodes the precision of the belief at the current level. The higher this precision (i.e., the more confident the agent is about its present belief) the smaller the impact of the PE on belief updating.

(Note that in the HGF for categorical trial outcomes, the sigmoidal mapping from the second to the first level in the HGF for categorical trial outcomes causes a slight deviation of ψ_2 from the general form given in Eq. A.6; see Mathys et al. 2011 for a detailed derivation).

In the context of our three-level HGF, two precision-weighted PEs occur, ε_2 and ε_3 , which are the main targets of our fMRI analyses and which we describe in some detail in the next sections. We focus on explaining the nature of the PEs; for the exact definition of the precision weights at different levels, see Mathys et al. (2011).

Precision-weighted PE at the second level

At the second level of the HGF, the precision-weighted PE about visual stimulus outcome (termed ε_2) serves to update the belief about x_2 (the cue-outcome contingency in logit space):

$$\begin{aligned}\mu_2^{(k+1)} - \mu_2^{(k)} &= \psi_2^{(k)} \delta_1^{(k)} \\ &= \varepsilon_2^{(k)}\end{aligned}\tag{A.7}$$

Here, the PE is the difference between the actual and the predicted outcome on trial k :

$$\delta_1^{(k)} = \mu_1^{(k)} - \hat{\mu}_1^{(k)}\tag{A.8}$$

where the prediction $\hat{\mu}_1^{(k)}$ about visual stimulus outcome results from a sigmoidal transformation of the previous belief about the cue-outcome contingency $\mu_2^{(k-1)}$:

$$\hat{\mu}_1^{(k)} = s\left(\mu_2^{(k-1)}\right) \quad (\text{A.9})$$

As detailed in Mathys et al. (2011), at the second level of the HGF the precision weight ψ_2 also corresponds to the variance σ_2 and can be seen as the equivalence of a dynamic learning rate in RL models (cf. Preusschoff & Bossaerts 2007). This is intuitively sensible because $\sigma_2^{(k)}$ represents the degree of our uncertainty about $x_2^{(k)}$. In other words, the less confident an agent is about its current estimate of $x_2^{(k)}$, the higher it should weigh the information conveyed by a prediction error.

Precision-weighted PE at the third level

At the third level, the precision-weighted PE ε_3 is proportional to the update of the belief about x_3 (environmental log-volatility):

$$\begin{aligned} \mu_3^{(k+1)} - \mu_3^{(k)} &\propto \psi_3^{(k)} \delta_2^{(k)} \\ &= \frac{\hat{\pi}_2^{(k)}}{\pi_3^{(k)}} \delta_2^{(k)} \\ &= \varepsilon_3^{(k)} \end{aligned} \quad (\text{A.10})$$

Here, the PE concerns the cue-outcome contingency, or more precisely, a variance ratio of its estimates (in logit space) after and before observing the sensory input, respectively:

$$\delta_2^{(k)} = \frac{\sigma_2^{(k)} + \left(\mu_2^{(k)} - \mu_2^{(k-1)}\right)^2}{\sigma_2^{(k-1)} + e^{K\mu_3^{(k-1)} + \omega}} - 1 \quad (\text{A.11})$$

A detailed explanation of the nature of this PE can be found in Mathys et al. (2011), pp. 8-9, which we quote here in part (where $u^{(k)}$ refers to the sensory input or outcome on trial k):

“... $\delta_2^{(k)}$ represents prediction error. It is positive if the updates at the second level (of μ_2 and σ_2) in response to input $u^{(k)}$ indicate that the agent was underestimating x_3 . Conversely, it is negative if the agent was overestimating x_3 . This can be seen by noting that the uncertainty about x_2 has two sources: *informational*, i.e. the lack of knowledge about x_2 (represented by σ_2), and *environmental*, i.e. the volatility of the environment (represented by $e^{K\mu_3 + \omega}$). Before

receiving input $u^{(k)}$ the total uncertainty is $\hat{\sigma}_2^{(k)} = \sigma_2^{(k-1)} + e^{\kappa\mu_3^{(k-1)} + \omega}$. After receiving the input, the updated total uncertainty is $\sigma_2^{(k)} + (\mu_2^{(k)} - \mu_2^{(k-1)})^2$... If the total uncertainty is greater after seeing $u^{(k)}$, the fraction in $\delta_2^{(k)}$ is greater than one, and μ_3 increases. Conversely, if seeing $u^{(k)}$ reduces total uncertainty, μ_3 decreases. (Since x_3 is on a logarithmic scale with respect to x_2 , the ratio and not the difference of quantities referring to x_2 is relevant for the prediction error in x_3 .)"

In other words, $\delta_2^{(k)}$ represents a PE about the certainty of the estimate of $x_2^{(k)}$, the cue-outcome contingency. This renders δ_2 conceptually similar to the construct of „expected uncertainty“ (EU) introduced by Yu & Dayan (2002, 2005). In this previous work, EU was implicitly defined (in slightly different ways across papers) as a higher-level PE in that it represents the difference between a conditional probability (degree of cue validity) and certainty (see Appendix to Yu and Dayan, 2005). However, while the two constructs are related, they are not identical: the model employed by Yu & Dayan (2005) is a hidden Markov model which models the probabilistic transition of a contextual state and ensuing discrete jumps in cue validity (when context changes). In contrast, the state representing cue validity in our model, $x_2^{(k)}$, evolves as a continuous quantity in logit space.

B. Outcome prediction errors and choice prediction errors

This section provides definitions of two different types of prediction errors (PE), both of which are related to the predictions about sensory outcome (visual stimulus category), i.e. predictions about \mathbf{x}_1 at the first level of our HGF model. In brief, the first PE, which we refer to as *outcome prediction error*, corresponds to $\boldsymbol{\varepsilon}_2$ as defined in Eq. A.7. As explained above, this PE is the difference between the actual visual stimulus outcome and its a priori probability (i.e., before trial outcome observation) according to the model. As we will explain in more detail below, this PE is always positive and bounded between 0 and 1. By contrast, the second PE, which we refer to as *choice prediction error*, is the difference between the correctness of the subject's choice and the a priori probability of this choice being correct (according to the model). This PE is positive when the subject's choice was correct and negative when his choice was wrong.

In the following, we will provide a detailed definition of these two different PEs, with a particular focus on the coding employed by our analyses.

Basic definitions

A *trial* consists of a sequence of three events:

1. An auditory *cue* is presented to the subject. This is either a *high tone* ($\boldsymbol{\jmath} = \uparrow$) or a *low tone* ($\boldsymbol{\jmath} = \downarrow$)
2. The subject predicts the visual stimulus category [*face* (f) or *house* (h)] that will follow the cue.
3. The visual *stimulus* is presented.

Outcome prediction error

The *outcome prediction error*, which drives learning at the second level of our HGF model, is the difference between the actual visual stimulus outcome s and its estimated probability before observation (cf. Eq. 25 in Mathys et al. 2011):

$$\boldsymbol{\varepsilon} \stackrel{\text{def}}{=} p(s|s) - p(s|\boldsymbol{\jmath}) = 1 - p(s|\boldsymbol{\jmath}) \quad (\text{B.1})$$

In our associative learning paradigm with categorical outcomes (face or house stimuli), the subject knows (by explicit instruction) that the following cue-stimulus contingency relations are in force:

$$p(f|\mathcal{J}=\uparrow) = 1 - p(h|\mathcal{J}=\uparrow) = p(h|\mathcal{J}=\downarrow) = 1 - p(f|\mathcal{J}=\downarrow) \stackrel{\text{def}}{=} \hat{\mu}_1 \quad (\text{B.2})$$

This enables us to write the subjects' prediction (the probability of seeing a particular stimulus, given the cue) in terms of $\hat{\mu}_1$:

$$p(f|\mathcal{J}) = \begin{cases} \hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases}, \quad p(h|\mathcal{J}) = \begin{cases} 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ \hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases} \quad (\text{B.3})$$

Then, the corresponding outcome prediction errors are:

$$\varepsilon(f|\mathcal{J}) = \begin{cases} 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ \hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases}, \quad \varepsilon(h|\mathcal{J}) = \begin{cases} \hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases} \quad (\text{B.4})$$

These prediction errors are always positive and bounded between 0 and 1, which is intuitively sensible for learning about categorical entities.

Under the above perspective, it would be natural to model two separate belief trajectories and obtain two prediction error trajectories (separately for face and for house outcomes). However, we can obtain equivalent results by modeling a single trajectory, provided we represent the categorical trial outcomes in contingency space using the following coding (where the assignment of 1 and 0 is arbitrary):

$$(f|\mathcal{J}=\uparrow) = 1, (f|\mathcal{J}=\downarrow) = 0, (h|\mathcal{J}=\downarrow) = 1, (h|\mathcal{J}=\uparrow) = 0 \quad (\text{B.5})$$

Based on this (arbitrary) coding and the definition of $\hat{\mu}_1$ above, we can compute the outcome prediction error in contingency space, ε_{cs} :

$$\varepsilon_{cs} | f = \begin{cases} 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ -\hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases} \quad (\text{B.6})$$

and

$$\varepsilon_{cs} | h = \begin{cases} -\hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases}. \quad (\text{B.7})$$

Clearly, the sign of this prediction error depends on the arbitrarily chosen coding. However, taking the absolute value of ε_{cs} gives unsigned prediction errors ε_u , such that on trials with face (f) outcomes

$$\varepsilon_u | f = \begin{cases} 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ \hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases}, \quad (\text{B.8})$$

and on trials with house (h) outcomes

$$\varepsilon_u | h = \begin{cases} \hat{\mu}_1 & \text{if } \mathcal{J}=\uparrow \\ 1 - \hat{\mu}_1 & \text{if } \mathcal{J}=\downarrow \end{cases}. \quad (\text{B.9})$$

Note that this renders Eqs. B.8 and B.9 identical to the expression in Eq. B.4. In other words, the unsigned (absolute) prediction error in contingency space (Eqs. B.8 and B.9) is identical to the prediction error when computed for each outcome category separately (Eq. B.4). For fMRI analyses, Eqs. B.8 and B.9 can be combined with trial-wise precision weights $\psi_2^{(k)}$ (see Eq. A.7) to yield a single trajectory for the precision-weighted prediction error about the visual outcome; this corresponds to $\varepsilon_2^{(k)}$ in Eq. A.7.

Finally, it is worth pointing out that the unsigned prediction errors in Eqs. B.8 and B.9 correspond to Bayesian surprise (cf. den Ouden et al., 2009; Ide et al. 2013). Therefore, the activations by ε_2 which we report in the paper are related but not identical to Bayesian surprise (because of the additional factor $\psi_2^{(k)}$).

Choice prediction error

Using the definitions of Eq. B.2 above, we can also derive the choice prediction error, i.e. the difference between the correctness of the subject's choice (1 if choice was correct, 0 otherwise) and the *a priori* probability of the subject's choice being correct (according to the running estimate of the Bayesian model).

On trials with outcome face (f) and correct choice f this means:

$$\varepsilon_{ch}|f = \begin{cases} 1 - \hat{\mu}_1 & \text{if } \mathcal{J} = \uparrow \\ 1 - (1 - \hat{\mu}_1) & \text{if } \mathcal{J} = \downarrow \end{cases} \quad (\text{B.10})$$

On trials with outcome f and incorrect choice house (h) this means:

$$\varepsilon_{ch}|h = \begin{cases} 0 - (1 - \hat{\mu}_1) & \text{if } \mathcal{J} = \uparrow \\ 0 - \hat{\mu}_1 & \text{if } \mathcal{J} = \downarrow \end{cases} \quad (\text{B.11})$$

On trials with outcome h and correct choice h this means:

$$\varepsilon_{ch}|h = \begin{cases} 1 - (1 - \hat{\mu}_1) & \text{if } \mathcal{J} = \uparrow \\ 1 - \hat{\mu}_1 & \text{if } \mathcal{J} = \downarrow \end{cases} \quad (\text{B.12})$$

On trials with outcome h and incorrect choice f this means:

$$\varepsilon_{ch}|f = \begin{cases} 0 - \hat{\mu}_1 & \text{if } \mathcal{J} = \uparrow \\ 0 - (1 - \hat{\mu}_1) & \text{if } \mathcal{J} = \downarrow \end{cases} \quad (\text{B.13})$$

This results in positive choice prediction errors whenever the subject's choice was correct, and in negative choice prediction errors whenever the choice was wrong.

C. Bayesian model selection (BMS) results

We used random effects BMS (Stephan et al. 2009) to examine alternative explanations for subjects' learning. Our model space included the following models (see also main text):

- **HGF₁**: a three-level HGF (representing cue-target associations at its second level and log-volatility at its third level), with two free parameters, ϑ and κ (ω was fixed to -4). The associated response model linked the predicted probability of a visual target given the auditory cue to behavioral responses (trial-wise predictions of visual stimulus category) by means of a sigmoid function with free parameter ζ .
- **HGF₂**: identical in structure to HGF₁, except that this model assumed that participants attempted to learn and predict trial-wise rewards (instead of visual stimuli).
- **HGF₃**: a reduced variant of HGF₁ where we eliminated the third level from the hierarchy by fixing both the log-volatility and the value of κ to zero.
- **RW**: a standard RL model based on the Rescorla-Wagner formalism, with a single free parameter representing the (fixed) learning rate.
- **Sutton**: a more flexible RL model with dynamic learning rate (algorithm K1 as described by Sutton 1992).

Importantly, BMS takes into account differences in model complexity (including the number of free parameters) and selects the model with the best trade-off between fit and complexity, thus shielding against overfitting (for details, see Penny et al. 2004; Stephan et al. 2009).

Comparison of all five models by BMS provided strong evidence for the superiority of HGF₁. Table S1 shows the posterior probabilities (PP) and exceedance probabilities (XP) of all five models (the exceedance probability of a model m is the probability that the posterior probability of m is higher than that of any other model considered; cf. Stephan et al. 2009). Please note that HGF₂ (which modeled predictions about trial-wise rewards) was not applied to data from fMRI study 2 because rewards were entirely absent in this study.

Table S2 shows the average maximum a posteriori estimates of the free parameters in the winning model (HGF₁) and their standard deviation (SD) across subjects.

D. Examples of Computational Trajectories in Two Individual Subjects

In this section, we plot the trajectories of the key computational variables used in our fMRI analyses for two individual subjects. These subjects were chosen because they provide examples of different individual learning styles, as reflected by different values of the parameter κ and ϑ : the first subject shows considerably higher values of both parameters than the second subject ($\kappa = 3.77$ vs. 0.94 , and $\vartheta = 4.69 \times 10^{-3}$ vs. 2.55×10^{-3}).

The plots of both subjects illustrate the general features of belief updating in the HGF (for example, that the updates of μ_i are proportional to the product of the PE from the next lower level, δ_{i-1} , and the precision weight ψ_i). Contrasting these two subjects, several individual differences are clearly visible (compare Figures S1 and S2):

- (i) Due to the higher estimates of κ and ϑ in the first compared to the second subject, the evolution of log-volatility μ_3 spans a wider range in the first subject.
- (ii) Due to high κ in the first subject, the updates of log-volatility μ_3 are dominated by PEs δ_2 from the second level; in contrast, the low κ in the second subject means that updates of μ_3 are relatively more influenced by the precision weighting ψ_3 .
- (iii) The high κ in the first subject means that the step size of the Gaussian random walk that μ_2 performs is strongly influenced by the log-volatility estimates μ_3 at the level above. It also means that the precision weight (or learning rate) ψ_2 at the second level is high and keeps changing throughout the experiment; in contrast, for the second subject with his low κ , this learning rate is smaller and changes less. These differences are also visible at the first level of the model (panel G) where, in the first subject, the updates of the predictions $\hat{\mu}_1$ are generally spikier, and this spikiness differs across time (e.g., compare the first and last block of the experiment).

One particular part of the plots, panel E, deserves an extra explanation. This panel shows two plots: ψ_2 (in green) is the precision weight which modulates the impact of outcome prediction error δ_1 on μ_2 , the conditional probability of the visual outcome (given the auditory cue).

Since μ_2 is in logit space (cf. Equation A.1), ψ_2 necessarily exhibits very large changes, even for small updates in the predicted visual outcome, whenever μ_1 is close to either 0 or 1. This can be seen, for example, during the first block of trials where the true cue-outcome contingency is 0.9 and where ψ_2 keeps rising during error-free periods even though μ_2 only shows minute changes. To alleviate the reader from having to take this sigmoid transformation into account when interpreting the plots and to make the function of ψ_2 as a dynamic learning rate more easily visible, the following transform is useful (which is motivated by considering that learning rate is essentially “rate of change”):

$$\begin{aligned}
 x_1 &= s(x_2) \\
 dx_1 &= \frac{dx_1}{dx_2} dx_2 \\
 &= s(x_2)(1-s(x_2))dx_2 \\
 \Rightarrow q(x_1) &:= s(x_2)(1-s(x_2))\psi_2
 \end{aligned}
 \tag{D.1}$$

where s is a sigmoid function as defined in Equation A.4.

Even though x_1 does not evolve as a Gaussian random walk by itself but is simply a sigmoidal transform of x_2 (see Figure 1 in the main text), Equation D.1 effectively transforms the learning rate at the second level into a learning rate at the first level. This transformed learning rate is $q(\psi_2)$, the red broken line in panel E, which declines during periods of correct predictions and shows marked jumps when prediction errors occur.

E. Additional fMRI results not reported in the main text

fMRI activations reflecting negative correlations with precision-weighted outcome prediction error ε_2

In the main text and Figure 2, we report brain regions whose activity, as measured by fMRI, showed a *positive* correlation with trial-by-trial precision-weighted outcome prediction error (ε_2 , see Eq. A.7 above). Here, we additionally report those brain regions which exhibited significant *negative* correlations with trial-by-trial ε_2 . These deactivations reflect brain responses which are inversely proportional to precision-weighted PEs about visual stimuli. This type of response may represent a form of “intrinsic reward” (of having made a correct prediction), but this is, of course, only a speculation. In any case, however, these activations are not related to the actual choice made by the subject. Instead, prediction errors related to choice are reported in the main text of this study.

In both fMRI studies, deactivations by ε_2 were found in the ventromedial prefrontal cortex (VMPFC) and auditory cortex ($p < 0.05$, whole-brain FWE corrected), as well as in the basal forebrain ($p < 0.05$, corrected for the anatomical ROI mask); see Table S3. Notably, these activations cannot reflect processing of trial-wise monetary rewards as these were orthogonal to visual outcome (in the first fMRI study) and were absent in the second fMRI study. In the first fMRI study, we also tested the orthogonal contrast between “high reward” and “low reward” trials explicitly and did not find significant effects in any of the regions mentioned above.

Activations reflecting uncertainty

The focus of the present manuscript is on prediction errors at different levels of the hierarchy embodied by our Bayesian model. However, the structure of our model also allows for testing which brain areas encode uncertainty at different hierarchical levels. Here, we report these secondary results in all brevity.

At the second level of our model, the uncertainty σ_2 can be understood as a dynamic learning rate which determines how the updating of the estimated cue-outcome contingency proceeds under the influence of PEs about visual stimuli and log-volatility estimates, respectively (see Mathys et al., 2011). We found that in both fMRI studies, σ_2 correlated positively with activity in right DLPFC and ACC/SMA (Figure S3 and Table S4). Interestingly, this activation was situated in a similar location as the activations due to volatility reported by Behrens et al. (2007).

However, strictly speaking, we do not consider our findings on σ_2 a proper replication across studies, since they did not spatially overlap under our requirement of a significant conjunction of peak-level corrected activations. Large spatial overlap was found, however, under whole-brain FWE correction at the cluster-level, as shown in Figure S3. The ROI analyses of dopaminergic and cholinergic regions did not yield significant results for σ_2 .

At the third level, the uncertainty σ_3 (to which the dynamic learning rate by which log-volatility estimates are updated is proportional) correlated with activity in the left cerebellum and left lingual gyrus ($p < 0.05$, whole-brain FWE corrected); see Table S5. Additionally, in our first fMRI study, the ROI analyses demonstrated a significant activation in cholinergic regions ($p < 0.05$, FWE corrected for our anatomical mask), both in the brain stem (PPT/LDT) as well as in the basal forebrain. However, this finding could not be replicated in the second fMRI study.

Activations reflecting log volatility

Finally, we examined the influence of trial-wise estimates of log-volatility, μ_3 . Whole-brain analyses demonstrated a negative correlation of μ_3 with activity in visual cortex, bilaterally in the calcarine sulcus and lingual gyrus ($p < 0.05$ FWE whole-brain corrected); see Table S6. The ROI analysis of dopaminergic and cholinergic regions did not yield significant results.

References

- Ide, J.S., Shenoy, P., Yu, A.J., Chiang-shan, R.L. (2013) Bayesian Prediction and Evaluation in the Anterior Cingulate Cortex. *J Neurosci* 33:2039-2047.
- Mathys, C., Daunizeau, J., Friston, K.J., and Stephan, K.E. (2011). A Bayesian foundation for individual learning under uncertainty. *Front Hum Neurosci* 5, 39
- Penny, W., Stephan, K., Mechelli, A., and Friston, K. (2004). Comparing dynamic causal models. *NeuroImage* 22, 1157-1172
- Preuschoff, K., and Bossaerts, P. (2007). Adding prediction risk to the theory of reward learning. *Ann N Y Acad Sci* 1104, 135-146.
- Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., and Friston, K.J. (2009). Bayesian model selection for group studies. *NeuroImage* 46, 1004-1017
- Sutton, R.S. (1992) Gain Adaptation Beats Least Squares? *Proceedings of the Seventh Yale Workshop on Adaptive and Learning Systems*, pp. 161–166, 1992.
- Yu, A.J., and Dayan, P. (2002). Acetylcholine in cortical inference. *Neural Networks* 15, 719-730.
- Yu, A.J., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron* 46, 681-692.