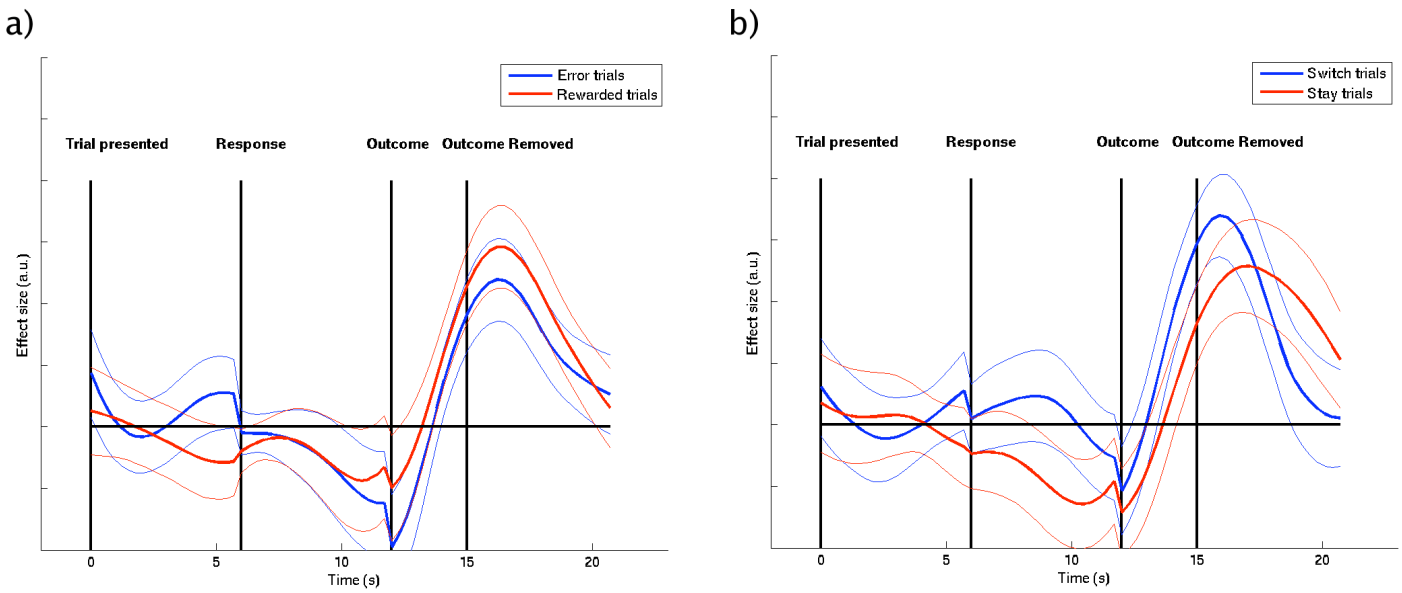

Figure S1: Plots show the effect of volatility at each timepoint when the data is divided into a) error and rewarded trials and b) switch and stay trials. The effect of volatility remains even if only error trials, or only rewarded trials, or only switch trials, or only stay trials are included in the analysis.

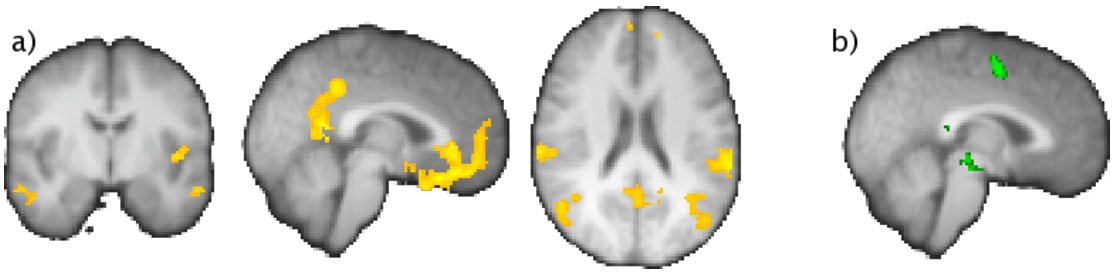

Figure S2:

Neural correlates of the value of the decision, during the DECIDE phase (a), and the prediction error during the MONITOR phase. Slices are taken at: coronal, y=-6, sagittal, x=-2, axial, z=22. a) clusters are shown for a cluster forming threshold of Z=2.8, (p<0.05 cluster corrected). b) No activations were present at corrected levels. The figure shows effects of prediction error in the presma, and Ventral tegmental Area for Z>2.3, p<0.01 uncorrected. These regions overlap with regions coding the predicted probability of an outcome during the INTERVAL phase. See supplementary and main text.

Supplementary Table 1

Model	#params (predictor)	#params (selector)	Prediction success	Log likelihood
Bayesian	0	36	87% (76%)	-1432
Bayesian jump	0	36	87% (76%)	-1435
RL 1	18	36	85% (69%)	-1516
RL2	36	72	86% (73%)	-1478

Learning the value of information in an uncertain world.

Timothy E.J. Behrens, Mark W. Woolrich, Mark E. Walton, Matthew F.S. Rushworth

Supplementary information.

Description of the Bayesian learner

Here we describe a Bayesian approach to tracking outcome probabilities given a past history of stochastic trial outcomes in an environment where these probabilities may change. At trial i a new outcome y_i is generated. In the context of this experiment, y_i is a reward either on blue or on green, where y_i is on blue with probability r_i, and on green with probability $(1- r_i)$. The goal of the subject is to estimate r_{i+1} from all previous data $y_{[0-i]}$ in order to choose optimally at the following trial. However, in a time-varying environment, r may change from trial-to-trial. It is beneficial to the subject to understand how r is changing. Effectively, it will allow them to know how best to discount information with passing trials. For example, in an environment where r is known to be stable, a subject should consider all historical trials with equal weight, in order to obtain the best prediction of r at the following trial. However, in an environment where r may change, recent trials are most likely to be representative of the current situation, so an outcome's predictive power is diminished as it moves further into the past.

The optimality of this approach relies on the assumption that the problem is solved in a Markovian fashion. That is, when an outcome is observed, the new outcome probability depends only on this observed outcome and on the previous outcome probability, but not on the full history of previous outcome probabilities. A Markovian approach is compelling since it can parsimoniously capture many of the

statistical regularities of the environment without explicitly considering states of the remote past. In other words, the environment can be well represented without needing to store the entire history of estimated reward probabilities and outcomes. Under this Markovian assumption, a diagram of the probability- tracking problem can be seen in Main figure 1, in the form of a generative model. In the Markovian setting, the changeability of r is represented by $p(r_{i+1} | r_i)$, i.e. the probability that r will move from the value r_i to the value r_{i+1} over the course of a single trial. This is a probability distribution on the next reward rate r_{i+1}, which must lie between 0 and 1. The distribution is centred on r_i, encoding the prior expectation that the reward rate at the next trial, r_{i+1}, depends upon the current reward rate, r_i . We represent this distribution as a *beta distribution*:

$$p(r_{i+1} | r_i , v) \sim \beta (r_i , V) \qquad (1)$$

where r_i defines the mean of the distribution and V=exp(v) defines the width of the distribution (See appendix for a definition of the reparameterised beta distribution as used here). Note that the exponential is taken so that integrals on the width parameter are performed in log space. A large value of v leads to a wide distribution, implying r may be expected to change greatly; a small value of v leads to a narrow distribution, implying r may be expected to change very little. Hence, we refer to v as the volatility as it controls the extent to which the observed outcomes of decisions can update the estimated reward probability between trials. If the subject treats the environment exactly as described above, they will make the assumption that the environment is always susceptible to the same rate of change, and they will therefore be unable to profit from periods of *local* stability to form more accurate estimates of the reward rate. If, in fact, the environment may go through stable phases and volatile phases, the

subject must understand this changing volatility in order to perform optimally. Mathematically, we track the changing volatility, v_i, in the same manner as the reward rate, r_i , by again assuming a Markovian progression. The changeability of v is represented as $p(v_{i+1} | v_i)$. Here, the distribution does not need to be constrained between 0 and 1, so the form of the transitional distribution can be taken to be Gaussian:

$$p(v_{i+1} | v_i , k) \sim N(v_i , K) \qquad (2)$$

where K=exp(k) controls the rate of change of volatility. Again, the exponent is taken so that integrals are performed in log space; the natural space for a variance parameter. A large value for k leads to a wide transitional distribution, and would describe an environment that moves quickly between stable and volatile periods; a small value for k leads to a narrow distribution and would describe an environment whose volatility could only change slowly. In other words, the parameter k represents the distrust in the constancy of the environment's volatility.

An alternative forward model.

The forward model described above is optimal for tracking drifting reward rates. However, an alternative model, is to assume that the reward rate remains constant for a period of time, and then jumps to a new unknown value. This second scenario is, in fact, the case in our experiment, but subjects cannot know this going into the task. Here we will lay out such a model, and later we show that the predictions of the Bayesian learner are not sensitive to this modeling choice.

For this model, we must introduce a latent variable S_i, which takes the value 1 when a jump in reward rate occurs, and 0 when the reward rate remains stable. The conditional probability distribution on r_i can now be written:

$$p(r_i \mid r_{i-1}, S_i) = \begin{cases} \delta(r_i - r_{i-1}) & S_i = 0 \\ U(0,1) & S_i = 1 \end{cases}$$

The volatility in this model is taken as the log probability of witnessing a reward rate jump at each trial. Hence:

$$p(S_i) = V_i$$

where, V_i=exp(v_i). We can now marginalise over the latent variable S_i, to leave:

$$p(r_i \mid r_{i-1}, V_i) = (1 - V_i)\partial(r_i - r_{i-1}) + V_i U(0,1).$$

And the conditional distribution on v_i is exactly as in the original model above.

$p(v_{i+1} \mid v_i, k) \sim N(v_i, K)$

Note that the only difference between the two models is therefore the exact form on the conditional distribution on v_i, and that the graphical description of the problem (main figure 1) is identical for the two models.

Inferring r from past observations.

During the experiment, the Bayesian learner's task is to use previous data (decision outcomes) to establish a belief about the statistics of the reward environment to use in the next trial. These unknown statistics are the outcome probability, r, the volatility, v, and the rate of change of volatility, k. It can be shown that this is done in the most efficient unbiased way by using Bayesian learning [1].

Bayesian learning is implemented using Bayes rule:

$$p(\text{unknowns} \mid \text{data}) \propto p(\text{data} \mid \text{unknowns})p(\text{unknowns})$$

where, in this experiment, the data is the history of decision outcomes from all trials; the unknowns are the statistics of the reward environment r, v, and k; and the likelihood, p(data | unknowns), is the generative model represented graphically in main figure 1b. This equation can be used to implement the Bayesian learner. It

provides the means to use the evidence in the data to update p(unknowns), the *prior* probability distributions of (or belief in) the unknown parameters *prior* to seeing the data at each trial, in order to obtain p(unknowns | data), the *posterior* probability distributions of (or belief in) the unknown parameters after seeing the outcome of the trial.

Bayes rule can be used to write down the *joint* probability distribution at trial i of all the parameters representing the statistics of the reward environment through the course of the experiment:

$$p(r_{\leq i}, v_{\leq i}, k \mid y_{\leq i}) \propto p(k)p(r_1)p(v_1)\prod_{j=1}^{i} p(y_j \mid r_j)p(r_j \mid r_{j-1}, v_j)p(v_j \mid v_{j-1}, k)$$

where $y_{\leq i}$ is all data (decision outcomes) up to and including trial i, $r_{\leq I}$ and $v_{\leq I}$ are the entire history of reward rates and volatilities at each trial and k is the distrust in the constancy of volatility (see main figure 1b). y_i,r_i and v_i represent the data, reward rate and volatility at a particular trial, *i*. Note that this makes use of the Markovian properties of r_i and v_i :

$$p(r_j \mid r_{\leq j}, v_{\leq j}) = p(r_j \mid r_{j-1}, v_j)$$
$$p(v_j \mid v_{\leq j}, k) = p(v_j \mid v_{j-1}, k)$$

Marginalizing (integrating) over the history of r and v previous to trial i gives the posterior probability distribution for the parameters r,v, and k at trial i as:

$$p(r_i, v_i, k \mid y_{\leq i}) \propto p(k)\int\cdots\int p(r_1)p(v_1)\prod_{j=1}^{i}[p(y_j \mid r_j)p(r_j \mid r_{j-1}, v_j)p(v_j \mid v_{j-1}, k)]dr_{\leq i-1}dv_{\leq i-1} \qquad (eq\ \ A)$$

A similar expression can be written down for the posterior probability for (belief in) the parameters r,v, and k at trial i+1:

$$p(r_{i+1}, v_{i+1}, k \mid y_{\leq i+1}) \propto p(k)\int\cdots\int p(r_1)p(v_1)\prod_{j=1}^{i+1}[p(y_j \mid r_j)p(r_j \mid r_{j-1}, v_j)p(v_j \mid v_{j-1}, k)]dr_{\leq i}dv_{\leq i}$$

Importantly, this can then be much simplified by substituting in the expression for the

posterior probability for (belief in) the parameters r,v, and k at the previous trial i (given in equation A):

$$p(r_{i+1},v_{i+1},k \mid y_{\leq i+1}) \propto p(y_{i+1} \mid r_{i+1}) \int \left[\int p(r_i,v_i,k \mid y_{\leq i}) p(v_{i+1} \mid v_i,k) dv_i \right] p(r_{i+1} \mid r_i,v_{i+1}) dr_i \qquad (eq \;\; B)$$

This equation provides a way of updating our belief in the parameters r, v, and k on a trial-by-trial basis without needing to store the full history of decision outcomes or the full history of statistics of the reward environment. This can be seen, as no data points previous to y_{i+1} and no parameters previous to trial i exist in the equation. The entire history of previous rewards is summarized in the joint probability distribution $p(r_i,v_i,k)$. This distribution, the current belief in the parameters r, v, and k, is all the information that is carried to trial i+1.

We perform the integrals in equation B using numerical integration on a 5-dimensional grid (the five dimensions represent r_{i+1},v_{i+1},r_i,v_i and k, with integration over two of them, r_i and v_i). This leaves a 3-dimensional grid representing the joint distribution $p(r_{i+1},v_{i+1},k)$, which is stored between trials. At the beginning of the experiment the joint distribution is set to be uniform, reflecting the fact that the subjects have no information about any of the parameters before the experiment begins.

In order to make a decision at trial i+1, the Bayesian learner requires an estimate of the reward rate r_{i+1}. This estimate is the mean of the probability distribution on r_{i+1}. The final computation for the Bayesian learner is therefore to compute the *marginal* distribution $p(r_{i+1})$ from the joint probability distribution $p(r_{i+1},v_{i+1},k)$. This is obtained by marginalising (integrating over) v_{i+1} and k:

$$p(r_{i+1}) = \int \int p(r_{i+1},v_{i+1},k) dv_{i+1} dk$$

The current estimate of the reward rate is then given by the mean of this distribution:

$$\hat{r}_{i+1} = E(r_{i+1}) = \int r_{i+1} p(r_{i+1}) dr_{i+1}$$

Again, the integrals in this and in the previous equation are performed numerically, this time across a 3-dimensional grid (the three dimensions are r_{i+1},v_{i+1} and k).

Experiment 1: Human Behaviour

Model performance:

In order to further test the performance of the model, we compared subject performance to predictions made by both the Bayesian learner (with both forward models) and by classical reinforcement-learning models. We used two different RL models for the comparison. The first (RL1) used one learning rate for each subject. The second (RL2) used separate learning rates in each phase (volatile and stable) in each subject. The RL1 model has 18 parameters (1 per subject), and the RL2 model has 36 parameters (2 per subject) that are fit to the subject data. In contrast, the Bayesian learner has no parameters that can be fit to the subject data. If the Bayesian model can perform as well as the reinforcement-learning models, despite having no free parameters to fit to the subject data, it is strong evidence for the Bayesian rationale for setting the learning rate.

On top of each of these "predictor" models, we used the "selector" model described above, and again fit model parameters using Bayesian estimation. We then made predictions of subject decisions from each model, assuming that the model would choose the action with the higher value at each trial. We compared these predictions against subject decisions in two ways (Table 1). First, we compared against subject decisions in all trials. However, due to the nature of the task, most of the decisions made by the subjects are trivial, and predicted by any reasonable model. We therefore further tested the models by examining only "difficult" trials. These

trials are defined as those where the expected Pascalian values of the two options are separated by less than 5 points (results in parentheses in table 1).

Despite the difference in number of parameters, the three models perform comparably over all trials, with the Bayesian model marginally out-performing the reinforcement learning models. Examining the data from the difficult trials, where chance performance is 50%, RL1 predicts subject decisions at 19% better than chance, RL2 predicts at 23% better than chance, and the Bayesian model predicts at 26% better than chance. Comparing the log likelihood of the three models with the fitted parameters reveals behavioural data to be best fit by the Bayesian model, followed by the RL2 model, again confirming that the data is best predicted by models in which the learning rate can change.

Comparing the two possible forward models for the Bayesian learner revealed a very similar level of prediction of the subject performance. This is unsurprising as, of all 580 trials in the two schedules, the two Bayesian forward models made different predictions of subject choice in only 9 trials.

Experiment II: Learning rule related activity in the ACC

FMRI experiment timing.

Subjects performed the same basic task as in the behavioural experiment described in the main text. Each trial was divided into 3 phases: **DECIDE (4-8seconds jittered):** The phase when subjects made the decision. The subjects could see the available options, but could not respond until a question mark appeared on the screen. **INTERVAL (4-8 seconds jittered):** After making the decision, but before the correct answer was revealed. **MONITOR (3ss):** Subjects observed the correct outcome of the trial in the centre of the screen. If the subject had guessed correctly at that trial, the prize bar moved forward by the distance associated with that option.

There was an inter-trial interval **(3-7 seconds jittered).** Total trial time was 20s on average. We had 120 trials, giving an average scan time of just over 40 minutes.

FMRI data

FMRI data were acquired in 20 subjects on a 3T Siemens TRIO scanner. Data were excluded from one subject due to rapid head motion, and from one subject as mean activation levels over the entire brain were 5 standard deviations from the population mean. The remaining 18 subjects were included in the analysis.

FMRI data were acquired with a voxel resolution of 3x3x3 mm^3, TR=3s, TE=30ms, Flip angle=87^o. The slice angle was set to 15^o and a local z-shim was applied around the orbitofrontal cortex to minimize signal dropout in this region[2], which had been implicated in other aspects of decision-making in previous studies. The number of volumes acquired depended on the behaviour of the subject. The mean number was 830, giving a total experiment time of around 42 minutes.

Field Maps were acquired using a dual echo 2D gradient echo sequence with echos at 5.19 and 7.65 ms, and repetition time of 444ms. Data were acquired on a 64x64x40 grid, with a voxel resolution of 3mm isotropic.

T1-weighted structural images were acquired for subject alignment using an MPRAGE sequence with the following parameters: Voxel resolution 1x1x1 mm^3 on a 176x192x192 grid, Echo time(TE)= 4.53 ms, Inversion time(TI)= 900 ms, Repitition time (TR)= 2200 ms.

FMRI analysis

FMRI analysis was carried out using FMRIB's Software library (FSL[3]).

Single subject processing.

Preprocessing

Data were preprocessed using FSL default options: motion correction was applied using rigid body registration to the central volume[4]; Gaussian spatial smoothing was applied with a full width half maximum of 5mm; brain matter was segmented from non-brain using a mesh deformation approach[5]; high pass temporal filtering was applied using a Gaussian-weighted running lines filter, with a 3dB cut-off of 100s.

Model estimation

A general linear model was fit in pre-whitened data space (to account for autocorrelation in the FMRI residuals)[6]. The following regressors were included in the model (see main text): **DECIDE** – times when the subjects were making the decisions, before they make a response; **INTERVAL** – times between making a response and the true outcome being displayed; **MONITOR** – 3 second period after the outcome has been displayed on the screen. **DECIDExVOLATILITY** – decision phase modulated by the volatility estimate at each trial; **INTERVALxVOLATILITY**– interval phase modulated by the volatility estimate at each trial; **MONITORxVOLATILITY**– monitor phase modulated by the volatility estimate at each trial; **ERROR** – monitor phase on trials when the subject chose the incorrect option. Note, for all volatility interactions, the most likely, or *modal* value was used from the posterior distribution on volatility at each trial.

These regressors were convolved with the FSL default haemodynamic response function (Gamma function, delay=6s, standard deviation =3s), and filtered by the same high pass filter as the data.

Field Inhomogeneity correction.

EPI distortions due to magnetic field inhomogeneity were corrected using the acquired field-maps. Field-maps were unwrapped using a common-phase region expansion approach[7], and the parameter and variance estimates from the model were spatially shifted according to the local field.

Group data processing.

Subjects were aligned to the MNI152 template using affine registration[8]. A general linear model was fit to estimate the group mean effect of the regressors described above[9].

Inference: Main effects.

For the main effects of **DECIDE** and **MONITOR**, which were expected to have widespread activations throughout the brain, inference was carried out using Gaussian random field theory and cluster-based thresholding[10], with a cluster-forming threshold of Z=2.8 and a family-wise false positive rate of p=0.05. Clusters surviving thresholding can be found in table 1.

Inference: Interaction effects.

The interaction between monitoring and volatility were expected to induce more focal activations than the main effects described above, and therefore not to be suitable for cluster-based thresholding. They were also expected to lie within the region of the anterior cingulate cortex (ACC). For this reason, inference for this interaction was carried out using Gaussian random field-based voxel-wise thresholding, with a small volume correction in the region of the ACC (3800 voxels). This gave a Z-threshold of 3.5. A region of 42 voxels in the Anterior Cingulate Cortex survived this threshold (max Z=4.23, at MNI x=-6,y=26,z=34 mm). Note that, as predicted no region in the brain outside the Anterior Cingulate Cortex would have survived at this threshold level. For purposes of display, main figure 3b is rendered with a Z-threshold of 3.1 (163 voxels in the ACC survived at this threshold level).

Region of interest analyses.

We performed several region of interest analyses on data take from the region of activation in the ACC. These analyses were performed in order to examine the nature of the BOLD signal fluctuations, and to exclude the possibility that these fluctuations could be accounted for by several potential confounding regressors.

We took BOLD data in each subject from a mask back-projected from the group ACC activation in the MONITORxVOLATILITY regressor. We separated each subject's timeseries into each trial, and resampled each trial to a duration of 20s, such that the decision was presented at 0s, the response was allowed at 6s and the outcome was presented from 12s-15s. The resampling resolution was 100ms. We then performed a GLM across trials at every trial timepoint in each subject independently. Lastly, we calculated the group average effect sizes at each timepoint, and their standard errors. The graphs in main figures 4 and 5, and supplementary figure 1 therefore show a timeseries of effect sizes throughout the trial, for the regressor of interest (volatility).

Supplementary figure 1 shows the effect of volatility on the BOLD signal separated out for a) error and rewarded trials and b) switch and stay trials. It can be seen that the effect of volatility is present independent of whether the trial in question was rewarded, and independent of whether the subject had switched response on the trial in question. Notice that in figure 1 and main figures 4,5 the effect of volatility is limited to the period when the subject is observing the outcome of the trial. This is the crucial period when the subject must use their estimated volatility, and therefore uncertainty, in updating their estimate of the future reward likelihood.

Other signals of interest generated by the model.

We performed a second analysis looking for the effects of a different form of uncertainty (the probability that a trial would be rewarded) to contrast with the effect of volatility we had already found. The crucial time for this form of uncertainty is the INTERVAL phase when the subject is awaiting the outcome of the trial. This new analysis followed the same procedure described above.

While our original experiment was carefully designed to control for any possible confounding effect of volatility, we were not able to control for possible confounding effects between other parameters in our model. In this analysis we therefore included other regressors that may confound with reward probability. We also included a prediction error regressor (and potential confounds), reasoning that brain regions that code for the predicted probability of a reward when the subject is waiting for an outcome, would be ideally placed to code for the error on this prediction when the outcome was observed.

We therefore used the following regressors:

Three regressors coding the phases of the experiment DECIDE,INTERVAL,MONITOR; three regressors coding their interaction with volatility; the estimated probability of a reward on the chosen option during the DECIDE and INTERVAL phases; the value (probability x outcome size) of the chosen option during the DECIDE and INTERVAL phases; the prediction error during the MONITOR phase; the actual reward obtained during the MONITOR phase.

We first note that, as predicted by the ROI analysis presented in the main manuscript, the effect of volatility was unchanged by the inclusion of the extra regressors.

The only new regressors to have significant activations either at voxel-wise or cluster corrected levels were:

1) The "value" regressors during the DECIDE phase showed significant clusters in: The ventro-medial prefrontal cortex (max Z=4.74), the posterior cingulate cortex (max Z=4.43), the left posterior insular (max Z=4.8), the supramarginal gyrus, bilaterally (max Z=4.2), the middle temporal gyrus bilaterally (max Z=3.95) and the superior occipital lobe bilaterally (max Z=3.8). These activations can be seen in figure S2.

2) The "value" regressor during the prediction phase shows significant effects in various visual and attention related areas in the occipital and parietal lobes. We do not detail these individually here.

3) The error trial regressor shows a significant effect at the SMA/preSMA boundary (max Z=3.6, MNI x=-4,y=0,z=52).

However, because of the firm hypotheses about the coding of reward probability in Fiorillo et al. 2003, we looked for the reward-prediction signal during the INTERVAL phase in the dopaminergic VTA at a lower threshold ($p<0.01$ uncorrected, $Z>2.3$) , and found a significant effect (main figure 6).

We further reasoned that the same regions that coded the reward prediction signal during the INTERVAL phase, would be ideally positioned to code for the prediction error signal during the MONITOR phase. This was the case at the same threshold, (figure S2).

The ACC region that codes for the volatility-related uncertainty showed no signal relating to the contrasting form of uncertainty: the probability of an outcome during the prediction phase (Z=0.6).

Appendix

Beta distribution

The standard form of the beta distribution is parameterised as follows:

$$\beta(x;\varsigma,\eta) = \frac{\Gamma(\xi+\eta)(1-x)^{\eta-1}x^{\varsigma-1}}{\Gamma(\xi)\Gamma(\eta)} \qquad 0 \le x \le 1$$

This function is valid for x between 0 and 1, and has a mean at $\mu = \frac{\xi}{\xi+\eta}$. The sharpness of the distribution is defined by the effective number of samples drawn from a binomial distribution, $I = \xi + \eta$. This can be easily interpreted in terms of estimating the probability of throwing a head from a number of coin tosses. The mean is the number of heads divided by the number of tosses, the sharpness (or certainty) of the distribution is the number of tosses.

We reparameterise this distribution in terms of the expected reward rate and the volatility.

$$\hat{r} = \mu = \frac{\zeta}{\zeta + \eta}$$

$$v = -\log(I) = -\log(\zeta + \eta)$$

Activated clusters for main effects of DECIDE and MONITOR, x,y,z are the MNI152 space locations of local maxima in the z-statistic.

DECIDE					
Cluster size (Voxels)	Brain region	x (mm)	y (mm)	z (mm)	z-score
14822					
	Striate and extrastriate visual cortex	38 -44 28 26	-88 -72 -96 -90	-12 -12 -2 -8	6.09 5.93 5.84 5.84
	Fusiform gyrus	34 -44	-44 -62	-20 -18	5.79 4.9
	Cerebellum	-36	-50	-34	6.29
3814					
	Superior temporal gyrus	60 66	-40 -46	20 16	5.11 4.77

	Insular Cortex	46	12	-4	5.1
	Frontal Operculum	42 38 42	18 20 18	-4 -14 -8	5.07 4.97 4.96
3250					
	Anterior Cingulate Cortex	6 4 2 -8	20 16 -8 -8	26 30 48 52	5.71 5.2 4.94 4.87
	Pre-SMA SMA	2 4 0	6 16 -8	60 56 68	5.17 5.14 3.97
2235	Insular cortex	-44 -38 -34 -32	-4 -2 18 14	2 6 6 10	4.87 4.76 4.74 4.67
	Frontal Opercular Cortex	-46 -28	8 24	0 0	4.77 4.8
	Globus pallidus	-20 18	2 0	-2 -2	4.18 3.9
2209	Parietal operculum	-54 -48 -56	-24 -26 -22	14 12 20	4.88 4.59 4.31

	Intraparietal sulcus Post-central Sulcus	-30 -34 -38	-56 -54 -34	52 54 42	4.47 4.44 4.36
1173	STN	8 -8	-22 -18	-4 -4	5.51 5.07
	hippocampus	24	-34	-6	4.7
	Centromedian thalamus	10 -12	-20 -20	2 0	4.64 3.9
940	Dorsolateral Prefrontal Cortex	34 30 36 32 28	28 30 46 48 54	26 24 34 22 12	4.92 4.79 4.69 4.4 4.14
455	Dorsolateral prefrontal cortex	-34 -32	46 56	24 20	4.83 4.53
284	Posterior cingulate cortex	6	-24	26	4.19
140	Parietal operculum	48	-24	16	4.06

MONITOR					
Cluster size (Voxels)	Brain region	x (mm)	y (mm)	z (mm)	z-score
34609	Parahippocampal Gyrus	20	-48	-14	6.13
	Striate and extrastriate visual and cortex	46 -4 18 -20	-70 -72 -74 -72	0 6 -10 -10	6.09 6.09 5.94 5.92
	Intraparietal sulcus	36 42 6 8 -36	-52 -52 -66 -46 -54	54 40 48 44 54	5.43 4.95 5.12 4.11 4.87
	Superior temporal sulcus	50	-42	12	5.11
	Ventral striatum	-10 16	6 6	-8 -4	4.75 3.4
	Midbrain	10	-20	-10	3.95
	Mediodorsal thalamus	12	-22	12	4.2
	Hippocampus	22 -20	-26 -22	-10 -10	5.45 5.45
	Premotor cortex	44 52	-4 6	50 16	5.31 5.01

	Dorsolateral prefrontal	52 50	14 36	34 20	5.25 4.45
1746	Anterior cingulate cortex	-4 6 4	34 36 32	6 10 24	4.94 4.74 4.48
	Pre-SMA/superior frontal gyrus	6	34	44	4.80
548	Frontal operculum	-28 -32	16 14	0 -10	4.9 4.8
	Insular cortex	-32	14	-10	4.8
417	Parietal operculum	-58	-14	28	4.45
	Ventral premotor	-54	2	42	3.59
259	Posterior cingulate cortex	8	-36	24	4.41
356	Superior temporal sulcus	-56	-50	8	3.53

References

1. Cox, R. T. Probability, frequency, and reasonable expectaion. Am. Jour. Phys 14, 1-13 (1946).
2. Deichmann, R., Gottfried, J. A., Hutton, C. & Turner, R. Optimized EPI for fMRI studies of the orbitofrontal cortex. Neuroimage 19, 430-41 (2003).

3. Smith, S. M. et al. Advances in functional and structural MR image analysis and implementation as FSL. Neuroimage 23 Suppl 1, S208-19 (2004).
4. Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage 17, 825-41 (2002).
5. Smith, S. M. Fast robust automated brain extraction. Hum Brain Mapp 17, 143-55 (2002).
6. Woolrich, M. W., Ripley, B. D., Brady, M. & Smith, S. M. Temporal autocorrelation in univariate linear modeling of FMRI data. Neuroimage 14, 1370-86 (2001).
7. Jenkinson, M. Fast, automated, N-dimensional phase-unwrapping algorithm. Magn Reson Med 49, 193-7 (2003).
8. Jenkinson, M. & Smith, S. A global optimisation method for robust affine registration of brain images. Med Image Anal 5, 143-56 (2001).
9. Woolrich, M. W., Behrens, T. E., Beckmann, C. F., Jenkinson, M. & Smith, S. M. Multilevel linear modelling for FMRI group analysis using Bayesian inference. Neuroimage 21, 1732-47 (2004).
10. Worsley, K. J., Evans, A. C., Marrett, S. & Neelin, P. A three-dimensional statistical analysis for CBF activation studies in human brain. J Cereb Blood Flow Metab 12, 900-18 (1992).
11. Yu, A. J. & Dayan, P. Uncertainty, neuromodulation, and attention. Neuron 46, 681-92 (2005).