

Neural Mechanisms for Integrating Prior Knowledge and Likelihood in Value-Based Probabilistic Inference

Chih-Chung Ting,¹ Chia-Chen Yu,³ Laurence T. Maloney,^{4,5,6} and Shih-Wei Wu^{1,2,6}

¹Institute of Neuroscience and ²Brain Research Center, National Yang-Ming University, Taipei, 112 Taiwan, ³School of Medicine, Taipei Medical University, Taipei, 110 Taiwan, and ⁴Department of Psychology, ⁵Center for Neural Science, and ⁶Institute for the Interdisciplinary Study of Decision Making, New York University, New York, New York 10003

In Bayesian decision theory, knowledge about the probabilities of possible outcomes is captured by a prior distribution and a likelihood function. The prior reflects past knowledge and the likelihood summarizes current sensory information. The two combined (integrated) form a posterior distribution that allows estimation of the probability of different possible outcomes. In this study, we investigated the neural mechanisms underlying Bayesian integration using a novel lottery decision task in which both prior knowledge and likelihood information about reward probability were systematically manipulated on a trial-by-trial basis. Consistent with Bayesian integration, as sample size increased, subjects tended to weigh likelihood information more compared with prior information. Using fMRI in humans, we found that the medial prefrontal cortex (mPFC) correlated with the mean of the posterior distribution, a statistic that reflects the integration of prior knowledge and likelihood of reward probability. Subsequent analysis revealed that both prior and likelihood information were represented in mPFC and that the neural representations of prior and likelihood in mPFC reflected changes in the behaviorally estimated weights assigned to these different sources of information in response to changes in the environment. Together, these results establish the role of mPFC in prior-likelihood integration and highlight its involvement in representing and integrating these distinct sources of information.

Key words: Bayesian decision theory; Bayesian integration; decision making; judgment under uncertainty; medial prefrontal cortex

Introduction

In making decisions, we balance the desirability of possible outcomes against the feasibility of getting any one of them (Gilboa, 2010). Decision can be modeled as a choice among lotteries where a lottery is a list of mutually exclusive outcomes each paired with a probability of occurrence: the lottery A (0.5, \$100; 0.5, \$0) offers a 50:50 chance of \$100 or \$0. The lottery B (1, \$40) offers the certainty of \$40. The decision maker must choose A or B.

In most everyday tasks, we do not have complete information about probabilities and must estimate them to make decisions. There are two potential sources of information that we can rely on to perform such tasks: one source is our past experience (prior knowledge) and the other is current sensory information (likelihood). How humans weigh and integrate these two sources of information has been extensively studied in psychology (Kahneman et al., 1982), experimental and behavioral economics (Cam-

erer, 1987), finance (Kluger and Wyatt, 2004), and perception and action (Knill and Richards, 1996; Körding and Wolpert, 2006).

Previous studies on reward-based learning revealed how mid-brain dopamine systems and associated frontostriatal circuits form expectations about future outcomes through experience (Schultz et al., 1997; O'Doherty et al., 2003; Fiorillo et al., 2003; Daw et al., 2006). Recently, decision-making studies have begun to reveal how prior expectations, established through the presentation of cues before a decision, affects choice and neural computations in frontoparietal and sensory systems (Forstmann et al., 2010; Li et al., 2011; Rahnev et al., 2011; Mulder et al., 2012; Kok et al., 2013). Vilares et al. (2012) investigated the neural representations of uncertainty associated with prior and likelihood information and found separate systems for representing these sources of uncertainty, with prior uncertainty represented in the orbito-frontal cortex and likelihood uncertainty in the occipital cortex. In a study that examined the encoding of event probability, d'Acremont et al. (2013) found separate frontoparietal networks for representing event frequency and posterior event probability.

Despite these recent advances, the neurobiological basis of the integration of prior and likelihood information, particularly in the domain of reward probability, remains largely unknown. It is not clear whether the systems that represent prior and likelihood information about reward probability are distinct from the systems that integrate them. It is little known whether the neural systems involved in prior-likelihood integration overlap with

Received July 31, 2014; revised Nov. 27, 2014; accepted Dec. 9, 2014.

Author contributions: L.T.M. and S.-W.W. designed research; C.-C.T., C.-C.Y., and S.-W.W. performed research; C.-C.T., C.-C.Y., L.T.M., and S.-W.W. analyzed data; C.-C.T., L.T.M., and S.-W.W. wrote the paper.

This work was supported by the National Science Council in Taiwan (Grants NSC 99-2410-H-010-013-MY2 and NSC 101-2628-H-010-001-MY4 to S.-W.W.) and the National Institutes of Health (Grant NEI019889 to L.T.M.). We acknowledge magnetic resonance imaging support from National Yang-Ming University, Taiwan, which is in part supported by the Ministry of Education plan for the top University.

The authors declare no competing financial interests.

Correspondence should be addressed to Shih-Wei Wu, PhD, Institute of Neuroscience, National Yang-Ming University, No. 155, Sec. 2, Linong Street, Taipei, 112 Taiwan. E-mail: swwu@ym.edu.tw.

DOI:10.1523/JNEUROSCI.3161-14.2015

Copyright © 2015 the authors 0270-6474/15/351792-14\$15.00/0

those involved in the computation of subjective value, particularly the medial prefrontal cortex (mPFC) and the ventral striatum, during value-based decision making (Platt and Huettel, 2008; Kable and Glimcher, 2009; Padoa-Schioppa, 2011; Bartra et al., 2013; Clithero and Rangel, 2013). To address these questions, we designed a novel lottery decision task that allowed us to examine the neural representations of prior and likelihood and the neural mechanisms involved in integrating them. We developed an ideal decision maker model to evaluate how subjects weigh these sources of information, how close they are to ideal performance, and how the weighting might change as we systematically manipulated both prior and likelihood information.

Materials and Methods

Overview of Bayesian integration

Bayesian decision theory provides a normative framework for modeling how people make decisions. Suppose that a subject is presented with a visual stimulus that signals a probability θ of obtaining a monetary reward $\$x$ or nothing. The stimulus and the monetary value form a lottery ($\theta, \$x$). The subject does not know the true value of θ and therefore has to estimate it so that s/he can make proper decisions about ($\theta, \$x$). What should the subjects do to estimate θ ?

In the Bayesian framework, knowledge about θ has two sources: a prior distribution $\pi(\theta)$ on θ based on past experience and a likelihood function $L(\theta)$ based on sensory data about the true value of θ that is currently available to the decision maker. The product of prior and likelihood—after normalization—is the posterior distribution, obtained as follows:

$$\pi'(\theta) = C_{\pi'} \pi(\theta) L(\theta) \tag{1}$$

where $C_{\pi'}$ is a normalization constant chosen so that the area under the posterior $\pi'(\theta)$ is 1. According to Bayesian decision theory (Berger, 1985; O'Hagan, Forster and Kendall, 2004; Zhang and Maloney, 2012), the posterior distribution $\pi'(\theta)$ captures all the information that the subject has about θ .

Prior distribution. We chose priors that were beta distributions (O'Hagan et al., 2004)

$$\pi(\theta) = C_{\pi} \theta^{s_0} (1 - \theta)^{f_0} \tag{2}$$

where C_{π} is the normalization constant and s_0, f_0 are the two free parameters. Here, we use a nonstandard parameterization of the beta distribution for convenience. The typical parameters α, β are related to ours by $\alpha = s_0 + 1, \beta = f_0 + 1$.

In Figure 1A, we illustrate the two prior distributions on θ that we used in the experiment. The blue curve is the probability density function for one hypothetical prior $\pi(\theta)$ for θ . With this prior, θ is less likely to be >0.5 and more likely to be in the range of 0.1 to 0.4. In contrast, the prior $\pi(\theta)$ in red gives more weight to values of $\theta >0.5$. One distribution (blue) had a mean of 0.3 ($s_0 = 2, f_0 = 6$) and the other (red) had a mean of 0.8 ($s_0 = 7, f_0 = 1$). In our experiment, each prior distribution was paired with a visual stimulus (a white-on-black icon, shown above each distribution in Fig. 1A). We refer to the stimulus as an “icon” or “symbol” interchangeably throughout the text. Subjects first had to learn these two prior distributions (see Session 1 in “Procedure”) and their associated icons before engaging in a task (see Session 2 in “Procedure”) that required the integration of prior and likelihood information.

Likelihood function. In the integration task (see Session 2 in “Procedure”), on each trial, subjects were presented with both prior and likelihood information about θ . In Figure 1B, we give an example on how likelihood information was generated. Suppose that, on a given trial, θ is sampled from the prior distribution in blue and $\theta = 2/3$. Likelihood information here is a sample from a binomial distribution with parameters (θ, n) where n is the sample size. In Figure 1B, we show three samples that differ in sample size. Each sample consists of red and white dots that respectively represent winning a reward or nothing.

The likelihood function for each sample is a beta distribution (O'Hagan et al., 2004) calculated as follows:

$$L(\theta) = C_L \theta^s (1 - \theta)^f \tag{3}$$

where s is the number of successes (winning a reward of $\$x$) in the sample, f , the number of failures (winning nothing), and C_L a normalization constant chosen so that the area under the likelihood function is 1. The sample size is $n = s + f$.

Posterior distribution. The posterior distribution is the product of Equations 2 and 3 and is also a beta distribution, as follows:

$$\pi'(\theta) = C_{\pi'} \theta^{s+s_0} (1 - \theta)^{f+f_0} \tag{4}$$

where $C_{\pi'}$ is a normalization constant chosen so that the area under the likelihood function is 1. The beta distribution is conjugate to the binomial (O'Hagan et al., 2004) leading to the simple relation between prior, likelihood and posterior that we find. Figure 1C provides a graphical example of a prior distribution (blue curve in the left figure), three alternative likelihood functions (gray curves in the left figure), and the resulting posterior functions (red curves on the right figure). Three possible samples based on sampling from (θ, n) are shown as examples in Figure 1B for how sampled data might look like with different sample sizes $n = [3, 15, 75]$ used in the experiment. The frequency of reward (s/n) in each sample is the proportion of red dots. In Figure 1C, we plot the likelihood function (gray curve) for each of the sample shown in Figure 1B.

For convenience, in this example, we set the frequency of reward to be $2/3$ in all 3 samples. The peak of each likelihood function is therefore at $2/3$ and is the same across these 3 samples differing in sample size. However, the variances of the likelihood functions are different. As a result, when an ideal decision maker combines likelihood (gray curve) with the prior (blue curve) as in Equation 4 to arrive at the posterior distribution on θ (red curve on the right figure in Fig. 1C), s/he gets different posterior distributions on θ .

For any given relative frequency of reward in the sample, the posterior distribution is closer to the prior distribution when the sample size is smaller, and is closer to the likelihood function when the sample size becomes bigger. Put differently, when the sample size is smaller, the ideal subject would weight prior information relatively more compared with likelihood information. However, as sample size increases, likelihood information is weighted relatively more.

Finally, the expected value of the lottery ($\theta, \$x$) with respect to the posterior is:

$$EV = \int \theta x \pi'(\theta) d\theta = x \int \theta \pi'(\theta) d\theta = x \mu_{\pi'} \tag{5}$$

where $\mu_{\pi'}$ is the mean of the posterior distribution. The mean of the posterior beta distribution is $\mu_{\pi'} = (s + s_0 + 1)/(s + s_0 + f + f_0 + 2)$, which we take to be the estimate of θ derived from the posterior. A common alternative to the mean is the mode of the posterior, the maximum a posteriori (MAP) estimator. For the beta distributions considered, the mean and MAP are very close and use of either would lead to the same conclusions here.

Procedure

We designed the experiment to find out how people combine information about prior and likelihood to estimate reward probability and to determine the neural systems involved in performing such computations. There were two sessions in the experiment. The first session was conducted in a behavioral testing room. The second session was conducted inside an MRI scanner and was the main session of interest. The tasks were programmed using the Psychophysics Toolbox in MATLAB (Brainard, 1997; Pelli, 1997).

Session 1 (behavioral): learning prior distributions. The goal of the session was to familiarize the subject with the reward probabilities associated with the two different priors, blue and red, in Figure 1A. There were two different visual stimuli shown as white-on-black icons in Figure 1A, each placed above the peak of the prior it represented.

The blue prior led, on average, to a 30% chance of reward, whereas the red prior had an 80% average chance of reward. Before the session, the subjects did not know the prior probability distributions associated with

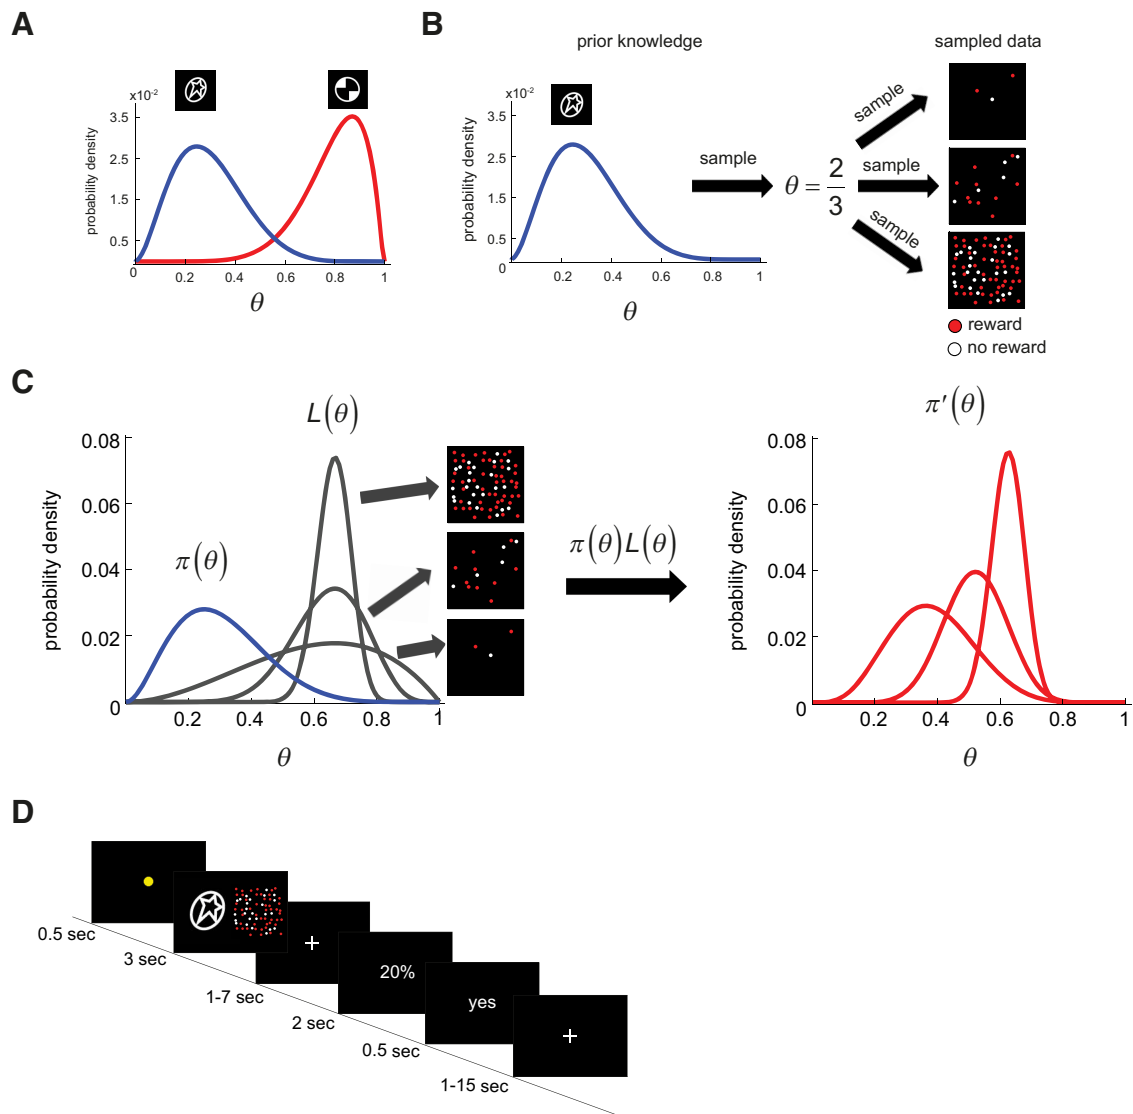


Figure 1. Overview of Bayesian integration and experimental design. **A**, Prior distributions. Two probability density functions on a parameter θ that takes on the value between 0 and 1. In this study, θ was the reward probability associated with a lottery option, termed the “symbol lottery.” There were two density functions, each represented by a white-on-black icon shown above the peak of each distribution. These density functions served as prior distributions on θ . Subjects had to first learn each density function in the behavioral session (Session 1). **B**, Generating current sensory information (likelihood). In the integration task (Session 2), on each trial, a sample was drawn from the prior distribution corresponding to the icon displayed to determine the reward probability of the symbol lottery (θ). In this example, we assume that $\theta = 2/3$ for convenience but, of course, it may take on any value between 0 and 1. The subjects did not know this value and had to estimate it. One source of information was the prior experience s/he had with the icon in the previous behavioral session. The other source was current sensory information about the true value of θ , which was generated by sampling with replacement based on θ . The sample was presented as randomly distributed dots (red and white dots that indicated reward and no reward respectively). There were three possible sample sizes (three dots, 15 dots, 75 dots) implemented in this study. **C**, Bayesian integration of prior and likelihood information. To optimally combine prior and likelihood information, an ideal Bayesian decision maker multiplies the prior distribution, $\pi(\theta)$, which represents prior beliefs, and the likelihood function, $L(\theta)$, which represents current sensory information about θ , resulting in the posterior distribution $\pi'(\theta)$. **D**, Trial sequence in the fMRI session (Session 2). Information about the symbol lottery was presented for 3 s and was followed by a variable 1–7 s interstimulus interval. After the interstimulus interval, the choice screen appeared. Information about the reward probability associated with the alternative lottery was revealed explicitly in numeric form and subjects had up to 2 s to make a choice between the symbol lottery and the alternative lottery with a button press.

the two icons. They had to learn each probability distribution through feedback in this session. The session consisted of eight blocks of 30 trials each. Within a block, only one of the two stimuli was presented. Each stimulus was allocated with equal number of blocks with the order randomized differently for each subject.

On each trial, the subjects were first presented with an icon designating a prior distribution and a probability of reward θ was drawn from it. Only the icon was presented; no other information about the prior probability distribution was presented. Therefore, on each trial, the stimulus represented a lottery ($\theta, \$x$) where $\$x$ was the reward. Notably, θ would vary from trial to trial because it was sampled from a prior distribution. The gain ($\$x$) associated with the lottery was fixed at 10 points (100 points = 1 National Taiwan Dollar (NTD); 1 USD = 30 NTD). When the visual

icon was presented, the subject’s task was to estimate the probability of reward in the current trial. The subject was instructed to enter his/her estimate from 0 to 100 in steps of 1 with key presses. For example, if the subject wished to enter 56, s/he pressed key 5 and then key 6 before hitting the return key to solidify his/her answer.

Critically, to motivate the subjects to learn the probability distributions, we implemented an incentive compatible procedure. On each trial, subjects received an extra monetary gain or loss depending on how close his/her estimate in the current trial was to the reward probability sampled from the distribution in that trial. Before the session, subjects received detailed instruction on this design: s/he would receive 100 points if the difference was within 5%, 50 points if the difference was between 5% and 15%, 0 if the difference was >15% but smaller than 30%, and would lose

50 points if the difference was >30%. Note that, to emphasize learning of the prior distribution, we purposefully made the consequence (extra gain or loss) of the accuracy of probability estimate much bigger than the size of the gain the subjects would receive based on the current lottery. After entering an estimate, the subjects received feedback on the actual reward probability θ of the current trial, the amount of bonus based on deviation of his/her probability estimate from θ , and whether they won a reward or not based on executing the lottery (θ , $\$x$).

Session 2 (fMRI): Integrating prior and likelihood. The goals of the session were to investigate how subjects integrated prior knowledge and likelihood and to determine the neural mechanisms underlying integration using fMRI.

The trial sequence is illustrated in Figure 1D. On each trial, the subjects were asked to choose between two lotteries, referred to as the “symbol lottery” and the “alternative lottery.” The two lotteries differed only in reward probability. Reward magnitude was the same and was fixed throughout the experiment. During each trial, information about the symbol lottery was presented first. The subjects were presented with one of the two visual icons from the previous session, signifying which before use. As in Session 1, the reward probability θ associated with the symbol lottery of the current trial was drawn from the prior distribution associated with the icon presented in that trial. Subjects did not know θ because it was not explicitly revealed to him/her.

What was different from Session 1 is that, in this session, the subjects were given an additional piece of information about θ . The piece of this new information came from sampling with replacement from the current sampled reward probability—the sample served as an additional source of information about the reward probability on the current trial.

As described in “Overview of Bayesian integration,” the sample is drawn from a binomial distribution with two parameters (θ , n) where θ is the reward probability associated with the symbol lottery on the current trial and n is the sample size. Figure 1B shows 3 samples of different sizes implemented in the experiment ($n = [3, 15, 75]$). In the examples shown in Figure 1B, the reward probability θ sampled from the prior distribution is 2/3. The colored dots represented the outcomes sampled from θ . The sampled results (red and white dots) were presented along with the icon designating the prior for 3 s. The location of the prior icon and sampled results (left or right) was randomized across trials.

Up to this point, the subjects had obtained information about only one of the lottery options. Information about the remaining option (termed the alternative lottery) was revealed after a variable interstimulus interval (1–7 s, uniform distribution) after the presentation of the prior icon and the sampling results. The reward probability associated with the alternative lottery (θ_{alt}) was explicitly specified in numeric form. No integration was needed to estimate θ_{alt} . On each trial, θ_{alt} was drawn randomly from the set [0.01, 0.2, 0.4, 0.6, 0.8, 0.99]. At this point, the subjects were instructed to choose between the symbol lottery and the alternative lottery within 2 s. No feedback was given to the subjects after their choice so that subjects could not update knowledge about the prior distributions through feedback.

Subjects were asked to indicate his/her preference level for the symbol lottery with a four-point scale (strongly yes, yes, no, strongly no). “Strong yes” or “yes” indicated that subjects chose the symbol lottery, whereas “strong no” or “no” indicated that s/he chose the alternative lottery. To exclude motor-related confounds that was not the interest of this study, the button mapping was balanced (left to right; right to left) across subjects. In the left-to-right mapping, the left middle finger, left index finger, right index finger, right middle finger indicated “strong yes,” “yes,” “no,” and “strong no,” respectively. The reverse was true for the right-to-left mapping. After the subjects made a response, a brief feedback (0.5 s) on the indicated preference level was given to the subjects to confirm their choice. There were 5 blocks of trials, each having 30 trials. We implemented a 2 (prior) by 3 (likelihood sample size) factorial design. Each combination was presented on 25 trials.

Subjects. Thirty-two subjects participated in the experiment (16 males; mean age, 25.4 years; age range, 19–33) and completed two sessions in 2 d. All participants had no history of psychiatric or neurological disorders. The study was approved by the Institutional Review Board at the Taipei Veterans General Hospital. Before the experiment, all subjects

gave written consent to participate in the study. Subjects were paid 600 NTD for their participation and monetary bonus (average: 332 NTD) obtained throughout the experiment. All subjects were paid after they completed the experiment.

Behavioral analysis 1: logistic regression analysis on choice and model comparison. The goal of this analysis was to compare how well different models described subjects’ choices. In a logistic regression analysis, the difference in reward probability between symbol lottery (θ_{sym}) and the alternative lottery (θ_{alt}), $\theta_{sym} - \theta_{alt}$, was implemented as the regressor. A choice of the symbol lottery was coded as 1, otherwise 0.

There were three models, the prior model, the likelihood model, and the posterior model. In the prior model, the value of θ_{sym} was the mean of the prior distribution on reward probability. In the likelihood model, the value of θ_{sym} was the mean of the likelihood function on reward probability. In the posterior model, the value of θ_{sym} was the mean of the posterior distribution on reward probability. For each subject, we estimated each model separately and computed its Bayesian information criterion (BIC). We then performed three pairwise model comparisons (posterior compared with prior, posterior compared with likelihood, likelihood compared with prior) based on BIC.

Let model A and model B denote the pair of models being compared. For each subject, if the BIC of model A is smaller than that of model B, which indicates that model A is better than B, assign a value of 1 to that subject; otherwise, assign 0 to the subject. As a result, the comparison of BIC between the two models becomes a binomial random variable (1 if model A is better, 0 if model B is better). We then performed a sign test to test the null hypothesis that the two models do equally well in describing the data. That is, the probability that A is better is equal to the probability that B is better: $H_0: p(\text{A is better}) = p(\text{B is better}) = 0.5$.

Behavioral analysis 2: logistic regression analysis on choice and the use of prior and likelihood information. Here, we performed a logistic regression analysis on choice to understand how subjects used information about prior and likelihood. A choice of the symbol lottery was coded as 1, whereas a choice of the alternative lottery was coded as 0. The analysis looked at all trials and implemented the following regressors: (R1) the mean of the prior distribution (μ_π), (R2) the mean of the likelihood function (μ_L), the frequency of reward in the current sensory information indicated by the fraction of red dots), (R3) the standard deviation (SD) of the likelihood function (σ_L), (R4) the interaction between R1 and R3 ($\mu_\pi \times \sigma_L$), (R5) the interaction between R2 and R3 ($\mu_L \times \sigma_L$), and (R6) the reward probability of the alternative lottery (θ_{alt}). We estimated this model for each subject separately.

Behavioral analysis 3: ideal decision maker analysis. The goal of this analysis was to evaluate how close subjects were to an ideal decision maker that optimally integrates prior and likelihood information. The critical comparison we made between actual performance and ideal performance was on the weight w_π assigned to prior information μ_π relative to the weight w_L assigned to likelihood information μ_L . We computed an index $RW = w_\pi/w_L$ and referred to it as the relative weight (RW).

We compared the actual RW with the ideal RW separately for each sample-size condition. If subjects combined prior and likelihood as predicted by the Bayesian model, then, as sample size increases, RW should decrease.

To estimate RW for each subject and each sample-size condition separately, we performed a logistic regression with the following regressors: (R1) the mean of the prior distribution (μ_π), (R2) the mean of likelihood function (μ_L), the frequency of reward in the current sensory information indicated by the fraction of red dots), (R3) the SD of the likelihood function (σ_L), and (R4) the reward probability of the alternative lottery (θ_{alt}). RW was estimated by the β estimate of R1 divided by the β estimate of R2.

To estimate the RW of an ideal subject, we simulated 10,000 experiments performed by an ideal decision maker. In each experiment, the ideal decision maker faced trials similar to those faced by the subjects. Based on Bayesian integration, we simulated what the ideal decision maker would choose. On each trial, we computed the posterior distribution on reward probability of the symbol lottery based on prior and likelihood information (Eq. 4). We then sampled from the posterior distribution to determine the ideal subject’s estimate on the reward prob-

ability of the symbol lottery $\hat{\theta}_{sym}$ in the current trial. By treating θ_{sym} as a random variable, we effectively added noise to the probabilistic inference process and the posterior distribution. The ideal subject would choose the symbol lottery if $\hat{\theta}_{sym} > \theta_{alt}$. Otherwise, the alternative lottery would be chosen. We also repeated all analyses using the mean of the posterior distribution as the subject's estimate instead of a sample from the posterior. This change in simulation did not affect our conclusions.

Second, based on simulated choice, we then estimated the *RW* using the same logistic regression model described above. We estimated *RW* separately for each sample-size condition. Finally, we computed the mean of the simulated *RW*, also for each sample-size condition separately, across the 10,000 simulated experiments and used it to define the *RW* of an ideal subject.

fMRI data acquisition. Imaging data were collected with a 3T MRI whole-body scanner (Siemens) equipped with a 12-channel head array coil. T2*-weighted functional images were collected using an EPI sequence (TR = 2 s, TE = 30 ms, 33 oblique slices acquired in ascending interleaved order, $3.4 \times 3.4 \times 3.4$ mm isotropic voxel, 64×64 matrix in a 220 mm field of view, flip angle 90°). Each subject completed five runs in the scanning experimental session. There were 30 trials in each run. Each run consisted of 270 images. T1-weighted anatomical images were collected after the functional scans using an MPRAGE sequence (TR = 2.53 s, TE = 3.03 ms, flip angle = 7° , 192 sagittal slices, $1 \times 1 \times 1$ mm isotropic voxel, 224×256 matrix in a 256 mm field of view).

fMRI preprocessing. The following preprocessing steps were applied using FMRIB's Software Library (FSL) (Smith et al., 2004). First, motion correction was applied using MCFLIRT to remove the effect of head motion during each run (Jenkinson et al., 2002). Second, we applied spatial smoothing using a Gaussian kernel of FWHM 8 mm. Third, a high-pass temporal filtering was applied using Gaussian-weighted least square straight line fitting with $\sigma = 50$ s. Fourth, registration was performed using a two-step procedure. First, the unsmoothed EPI image that was the midpoint of the scan was used to estimate the transformation matrix (seven-parameter affine transformations) from EPI images to the subject's high-resolution T1-weighted structural image with nonbrain structures removed using FSL's BET (Brain Extraction Tool). Second, we estimated the transformation matrix (12-parameter affine transformation) from the high-resolution T1-weighted structural image with non-brain structures removed to the standard MNI template brain.

General linear models of BOLD response

We estimated the following three general linear models (GLMs) of the blood oxygen-level dependent (BOLD) signals using FSL's FEAT module (fMRI Expert Analysis Tool). Time series were prewhitened with local autocorrelation correction (Woolrich et al., 2001). The GLM analysis was then performed in three steps (Beckmann et al., 2003). First, first-level FEAT analyses were performed for each run of each subject. Second, a second-level fixed-effect (FE) analysis was performed for each subject that combined the first-level FEAT results from different runs using the summary statistics approach (Beckmann et al., 2003). That is, the parameter estimate (β) of each contrast specified in the first-level analysis was treated as data for the FE analysis. Finally, a third-level mixed-effect analysis using FSL's FLAME module (FMRIB's Local Analysis of Mixed Effects) was performed across subjects by taking the FE results from the previous level and treating subjects as a random effect (Woolrich et al., 2004).

All reported whole-brain results were corrected for multiple comparisons (familywise error rate $p < 0.05$) using Gaussian random field theory. This was accomplished by first defining clusters of activation based on a z statistic (the cluster-forming threshold). We then estimated a familywise error-corrected p -value of each cluster based on its size using Gaussian random field theory (Worsley et al., 1992; Forman et al., 1995).

We denote the mean of the posterior distribution on reward probability associated with the symbol lottery as μ_{π} , the SD of the posterior distribution as σ_{π} , the mean of the prior distribution on reward probability as $\mu_{\pi'}$, the frequency of reward indicated by current sensory information (fraction of red dots) that would correspond to the mean of the likelihood function as μ_L , the SD of the likelihood function as σ_L , the

reward probability of the chosen lottery as θ_C , and the reward probability of the nonchosen lottery as θ_{NC} .

GLM-1. The purpose of this GLM was to identify regions in which activity is correlated with information about the posterior distribution on reward probability associated with the symbol lottery when it was presented. The model had the following regressors: (R1) an indicator function for the presentation of the symbol lottery with a length of 3 s when sample size $n = 3$, (R2) R1 multiplied by $\mu_{\pi'}$ when sample size $n = 3$, (R3) R1 multiplied by $\sigma_{\pi'}$ when sample size $n = 3$, (R4) an indicator function for the presentation of the symbol lottery with a length of 3 s when sample size $n = 15$, (R5) R4 multiplied by $\mu_{\pi'}$ when sample size $n = 15$, (R6) R4 multiplied by $\sigma_{\pi'}$ when sample size $n = 15$, (R7) an indicator function for the presentation of the symbol lottery with a length of 3 s when sample size $n = 75$, (R8) R7 multiplied by $\mu_{\pi'}$ when sample size $n = 75$, (R9) R7 multiplied by $\sigma_{\pi'}$ when sample size $n = 75$, (R10) an indicator function for the choice period with a duration of subject's reaction time, (R11) R10 multiplied by the rating the subjects gave to the symbol lottery, (R12) R10 multiplied by $\theta_C - \theta_{NC}$.

GLM-2. The purpose of this GLM was to identify regions that correlate with information about the prior distribution and the likelihood function of reward probability associated with the symbol lottery at its presentation. The model had the following regressors: (R1) an indicator function for the presentation of the symbol lottery with a length of 3 s, (R2) R1 multiplied by μ_{π} , (R3) R1 multiplied by μ_L orthogonalized with respect to μ_{π} , (R4) R1 multiplied by σ_L , (R5) R1 multiplied by $|\mu_{\pi} - \mu_L|$, (R6) R1 multiplied by the size of the sample, (R7) an indicator function for the choice period with a duration of subject's reaction time, (R8) R7 multiplied by the rating the subjects gave to the symbol lottery, (R9) R7 multiplied by $\theta_C - \theta_{NC}$.

Because μ_L is sampled from the reward probability that was obtained by sampling from the prior distribution, μ_L and μ_{π} are expected to be correlated with the correlation increasing as sample size increases. The correlation across all trials was 0.78 (averaged across subjects). This high correlation was primarily driven by trials where sample size was 15 (0.8927) and 75 (0.8370). When sample size was 3, the mean correlation was 0.02. To avoid multicollinearity issues, we chose to orthogonalize μ_L with respect to μ_{π} . Therefore, signals with variance that is captured by the shared variance component between μ_{π} and μ_L will be attributed to the μ_{π} regressor, whereas signals that are unique to μ_L (deviation of likelihood from prior) will be attributed to the μ_L regressor.

GLM-3. This GLM was identical to *GLM-2* except that we added two additional regressors, one for the interaction between μ_L and σ_L ($\mu_L \times \sigma_L$), and one for the interaction between μ_{π} and σ_L ($\mu_{\pi} \times \sigma_L$). To avoid multicollinearity issues due to their correlation with σ_L , these regressors were orthogonalized with respect to σ_L . We also analyzed the data without applying orthogonalization and did not find that orthogonalization significantly altered the results.

All regressors in the GLMs were convolved with a canonical gamma hemodynamic response function. We implemented temporal derivatives of each regressor in each model as the regressors of no interest. This implementation often serves as an alternative to slice-timing correction and is useful to model the nonlinear neural and vascular effects on the timing shift of BOLD response (Henson et al., 1999; Calhoun et al., 2004). We also included the estimated motion parameters as regressors of no interest.

Construction of independent, unbiased regions of interest

Following Esterman et al. (2010) and Litt et al. (2011), we created an independent, unbiased region of interest (ROI) for each subject separately. For each subject, we first performed a mixed-effect analysis on the contrast of interest using all other subjects' data. As a result, we obtained a statistical parametric map (SPM) for the contrast of interest. The SPM was created for each subject separately and was independent of each subject's data. We then performed the standard cluster-based thresholding to correct for multiple comparisons using Gaussian random field theory. The voxel with the maximum z statistic within the cluster of interest was identified and a sphere mask centered at the peak voxel was then created (radius = 6 mm).

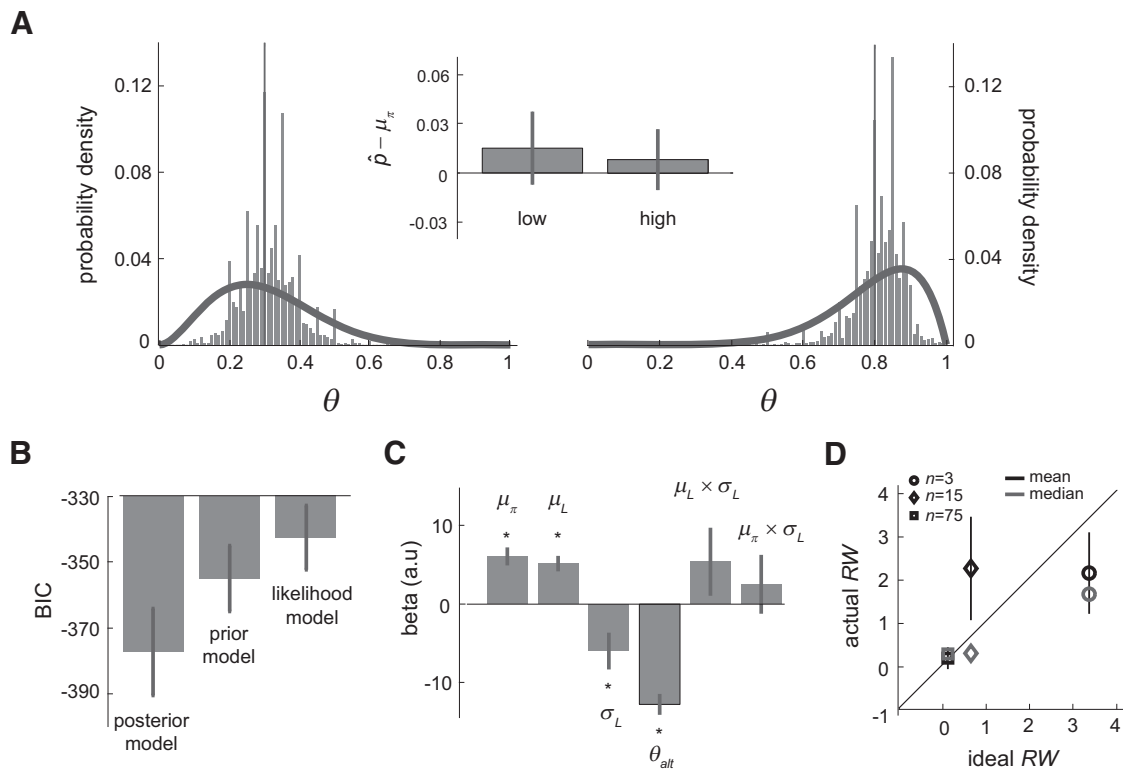


Figure 2. Behavioral results. **A**, Prior-learning session (Session 1, behavioral). Subjects learned two prior distributions on θ through sampling from each and receiving feedback on the sampled outcomes. In addition, subjects provided trial-by-trial estimates of θ , which are plotted as histogram here. The curves were the prior distributions and the vertical line indicated the mean of the distributions. The inset figure shows that when we compared subjects' probability estimate with the mean of the distribution, $\hat{p} - \mu_{\pi}$, there was no significant difference in both prior distributions ($p < 0.05$). **B–D**, Integration session (Session 2, fMRI). **B**, Model comparison: BIC. We estimated three logistic regression models on choice, which we referred to as the posterior model, the prior model, and the likelihood model, and plotted the mean of the BIC averaged across subjects for each model. **C**, Logistic regression analysis on the use of prior and likelihood information. Using choice data, we estimated a logistic regression model that consisted of the prior information (μ_{π}), the likelihood information (μ_L), likelihood uncertainty (σ_L), the alternative lottery (θ_{alt}), interaction between μ_L and σ_L ($\mu_L \times \sigma_L$), and interaction between μ_{π} and σ_L ($\mu_{\pi} \times \sigma_L$). The estimated β coefficients are plotted here. $*p < 0.05$, one-sample t test. **D**, Ideal decision maker analysis. We plot the mean RW across subjects (black) and the median RW (gray) against the ideal RW. Different symbols represent different sample-size conditions. Error bars represent ± 1 SEM.

Time-series analysis

We performed three time-series analyses on the independent ROIs that we identified using a method similar to that of Boorman et al. (2013). For each subject, we extracted the BOLD time series from an ROI. For each analysis, a generalized linear regression was performed on the BOLD time-series data.

Analysis 1. We modeled the 14 s after the onset of symbol-lottery presentation. At each time point after (and including) the onset of the presentation, we implemented an indicator regressor that took the value of 1 and a parametric regressor that took the value of μ_{π} . The regression thus gave us a time series of β coefficient of μ_{π} . At each time point, we then transformed the β coefficient to a z -statistic and plotted it in Figure 3B as a function of sample size (color coded).

Analysis 2. We modeled the 14 s after the onset of symbol-lottery presentation. At each time point after (and including) the onset of the presentation, we implemented an indicator regressor that took the value of 1, and two parametric regressors: one regressor took the value of μ_{π} , whereas the other took the value of μ_L orthogonalized with respect to μ_{π} . For μ_{π} (prior) and μ_L (likelihood) separately, we computed the z -statistic of the estimated β coefficient for each time point and plotted them in Figure 4B.

Analysis 3. The analysis was similar to Analysis 2 except that we estimated μ_{π} (prior) and μ_L (likelihood) separately for each sample-size condition and plotted them in Figure 5A–C.

Multiple-comparison correction. In the time-series analysis, at each time point, we estimated the β coefficient of a regressor of interest (e.g., μ_{π}) and computed its z -statistic. This effectively made each time point an independent test, requiring correction for multiple comparisons. We determined the corrected z threshold based on the following two considerations. First, we were only interested in testing the time points that

reflected BOLD signals in response to symbol-lottery presentation. Second, we allowed for the possibility that certain computations might take place during the delay period (the fixation period after symbol-lottery presentation and before the presentation of the alternative lottery). These considerations led us to test the time points between 4 and 12 s after stimulus onset based on the hemodynamic delay (4–6 s after stimulus onset) and the total duration of the symbol lottery and the delay period (7 s on average). Given that the TR was 2 s, there were a total of 5 time points and thus 5 tests to be corrected for. Therefore, we used the z value that corresponds to the p -value of $0.01 = 0.05/5$, which is 2.33, as the threshold. The dashed line in the relevant figures represents the corrected z threshold.

Results

Behavioral results: learning prior distributions

Subjects provided trial-by-trial estimates of the probability of reward associated with the two icons, each representing a prior probability density function on reward probability. Their estimates are plotted as histograms in Figure 2A. For both distributions, the mean estimate across subjects did not differ significantly from the mean of the distribution ($p > 0.05$, one-sample t test), indicating that subjects acquired accurate knowledge about the probability of reward associated with these two different stimuli.

Behavioral results: integrating prior and likelihood

To address how subjects used information about prior and likelihood in the fMRI session (see “Session 2: Integrating prior and

Table 1. Model comparison based on choice behavior

	n	n_A	$p(x \geq n_A H_0 \text{ is true})$
A: Posterior model	32	26	0.001
B: Prior model			
A: Posterior model	32	28	0.001
B: Likelihood model			
A: Prior model	32	22	0.025
B: Likelihood model			

Three model comparisons were made. For each comparison, we tested which model (A or B) is better based on comparing the BIC of each individual subject between the two models. The results were coded as 1 when model A was better and 0 when model B is better. A sign test was then carried out. n , Number of subjects; n_A , number of subjects in whom model A was better than model B; the binomial random variable; H_0 : $p(A \text{ is better}) = p(B \text{ is better}) = 0.5$.

likelihood”), we performed three sets of complementary behavioral analyses using subjects’ choice data. We summarize the results below.

Analysis 1: logistic regression on choice and model comparisons

We evaluated three models and compared how well these models described choice data. The three models were the posterior model, the prior model, and the likelihood model (see Materials and Methods for details). For each model, we summed the BIC over all subjects and plotted it in Figure 2B. Here, we observe that the posterior model had the smallest BIC value, indicating that this model described choice data better than the other two models at the aggregate level. We then performed three pairwise model comparisons each using a sign test based on the BIC (see “Behavioral analysis 1: logistic regression analysis on choice and model comparisons” for details). The results are summarized in Table 1. We found that the posterior model described the choice data better than either the likelihood model ($p < 0.001$) or the prior model ($p < 0.001$). This is the first evidence in this dataset indicating that, instead of using either prior or likelihood information alone, subjects used the posterior information about the symbol lottery, the product of prior and likelihood information, when making decisions. In addition, we also found that the prior model was better than the likelihood model ($p = 0.025$).

Analysis 2: logistic regression on choice and the use of prior and likelihood information

Once we obtained the initial results supporting Bayesian integration, we further examined the use of prior and likelihood information using logistic regression analysis. We modeled the choice data by including the regressors on (1) prior (μ_π , the mean of the prior distribution on reward probability associated with the symbol lottery), (2) the likelihood (μ_L , the mean of the likelihood function on reward probability associated with the symbol lottery), (3) the likelihood uncertainty (σ_L , the SD of the likelihood function on reward probability associated with the symbol lottery), (4) the interaction between μ_π and σ_L ($\mu_\pi \times \sigma_L$), (5) the interaction between μ_L and σ_L ($\mu_L \times \sigma_L$), and (6) the reward probability of the alternative lottery θ_{alt} . We estimated the model for each subject separately.

We removed one subject’s data because his/her parameter estimates were >3 SDs away from the group mean. Based on the remaining 31 subjects, we then computed the mean of the parameter estimates (β). See Figure 2C for the β plots and Table 2 for details on the estimation results. We found that the β estimate of μ_π and μ_L were both positive and significantly different from 0 ($p < 0.0001$). This suggests that both prior mean and likelihood mean had a positive impact on the subjects choosing the symbol lottery. The β of σ_L was negative, suggesting that as likelihood information became more uncertain, subjects tended to

Table 2. Logistic regression analysis on choice behavior

	beta	SEM	p -value
μ_π	6.08	1.15	<0.0001
μ_L	5.17	0.99	<0.0001
σ_L	−5.98	2.34	0.0008
θ_{alt}	−12.78	1.33	<0.0001
$\mu_L \times \sigma_L$	5.40	4.34	0.11
$\mu_\pi \times \sigma_L$	2.52	3.75	0.25

Each row represents the mean beta estimate across subjects, SEM, and the p -value (testing significance from 0 using 1-sample t test) of each regressor implemented in the regression model.

choose the symbol lottery less often. The β of θ_{alt} was also negative, suggesting that, as the reward probability of the alternative lottery increased, subjects tended to choose the symbol lottery less often.

The interaction terms, in particular their sign, could potentially be used to test whether subjects performed Bayesian integration. If they do, then subjects should assign more weight to likelihood information when likelihood uncertainty becomes smaller (a negative effect of $\mu_L \times \sigma_L$) but should weight prior information more when likelihood uncertainty becomes greater (a positive effect of $\mu_\pi \times \sigma_L$). However, the β weights of both interaction terms were positive but neither was significantly different from 0.

We were surprised that the β weights of both interaction terms were not significantly different from 0. To explore further, we simulated the ideal decision maker’s choices for the actual trials that had been presented to each subject. See “Analysis 3: ideal decision-maker analysis” in Materials and Methods for how we simulated choice of an ideal decision maker.

Once we obtained the simulated choice data for all 32 subjects, we then estimated the same logistic regression model described above for each subject, obtaining 32 sets of parameter estimates. We then computed the mean and SE of the parameter estimates just as we had done in analyzing the actual data. We repeated the above procedure 10,000 times.

We found that the estimated β weight of $\mu_L \times \sigma_L$ was not significantly different from 0 at the 0.05 level in $\sim 97\%$ of the 10,000 simulated experiments. For $\mu_\pi \times \sigma_L$, the estimated β weight was not significantly different from 0 in $\sim 50\%$ of the 10,000 experiments.

We also found that, of 10,000 simulated experiments, ~ 2500 times the β of $\mu_L \times \sigma_L$ was positive and the β of $\mu_\pi \times \sigma_L$ was negative, 3500 times the β of $\mu_L \times \sigma_L$ was negative and the β of $\mu_\pi \times \sigma_L$ was positive, 3700 times both terms were positive, and 300 times both were negative.

That is, even if all of our 32 subjects were ideal Bayesian decision-makers, we would expect the pattern of results we found in actual data (the β weights of both interaction terms being positive) about 37% of the time. We would expect the correct pattern of results (the signs of the β weights of the interaction terms in agreement with those of their population values) only 35% of the time. About half of the time, we would expect neither interaction β weight to be significantly different from 0, the observed outcome. We conjecture that one reason for the observed lack of power is that the correlation between these two interaction terms was on average 0.86 in the experiment each subject ran (150 trials) and hence in our simulations. Including them in the same model would introduce a multicollinearity problem and lead to difficulty in estimating the sign of the β weight of either interaction term. In summary, we cannot reject the Bayesian integration hypothesis.

Given that learning of the prior occurred before the fMRI session and that the subjects were not given further information that would

allow him/her to update it, one wonders whether representation of the prior became noisier over time. To address this question, we analyzed the behavioral data by separating the first half of the fMRI session and the second half of the session. We found no significant difference between the first and second half of the session in the β coefficient of prior information in terms of its mean (paired t test, $t_{df=31} = 0.4582$, $p = 0.6288$) and variance (F-ratio test, $F_{df_1=31, df_2=31} = 0.6357$, $p = 0.2129$). This indicated that there was no significant change in the representation and use of prior information over time during the fMRI session.

Analysis 3: ideal decision maker analysis

In the current experiment, as we manipulated the sample size of the likelihood information, σ_L was systematically changed. Bayesian integration predicts that as σ_L increases, subjects should rely more on prior information and less on likelihood information. Hence, the weight assigned to prior relative to likelihood information, termed the relative weight (RW), should increase. Conversely, when sample size increases, RW should decrease, indicating that subjects should rely more on likelihood information. In this analysis, we seek to address two questions. First, how might RW change in response to changes in sample size? How would those changes compare with the changes an ideal decision maker would make?

For each subject, we performed one logistic regression analysis on choice behavior for each sample-size condition. We estimated RW as the ratio of β estimate of prior (μ_π) to that of the likelihood (μ_L). For each sample-size condition, an $RW > 2.5$ SDs away from the mean (across subjects) was identified as an outlier and was removed from subsequent analysis. We detected one outlier in each condition.

After outlier removal, for each condition, we computed the mean of RW of the remaining 31 subjects and plotted it against the ideal RW in Figure 2D (black). The mean RW was 2.14 when the sample size was small ($n = 3$). When sample size was medium ($n = 15$), the mean RW was 2.25, which was indistinguishable to the mean RW when $n = 3$. When sample size was large ($n = 75$), the mean RW was 0.18. This indicated that when sample size was small, subjects typically assigned more weight to prior than to likelihood information ($RW > 1$). As sample size became bigger, the subjects started to assign more weight to likelihood than to prior ($RW < 1$). The median of subjects' RW was also plotted against ideal for each condition (color coded in gray). The median RW was approximately the same as the mean RW when $n = 3$ and $n = 75$, but was different when $n = 15$.

Such difference was likely due to the asymmetry in the distribution of RW . Because the distribution was likely to be non-Gaussian, we used nonparametric bootstrap methods (Efron and Tibshirani, 1993) to estimate the confidence interval of actual mean RW under each sample-size condition. That is, for each sample-size condition, we first resampled with replacement from the actual RW obtained from the subjects ($n = 31$ in all 3 conditions) and computed the mean of the RW based on the resampled dataset. We then repeated this procedure for 10,000 times. As a result, we obtained a distribution of the mean actual RW , which we used to construct the 95% confidence interval. If the ideal RW were inside the confidence interval, we would conclude that subjects' mean RW was indistinguishable from the ideal RW . Otherwise, we would conclude that subjects significantly deviated from ideal. For $n = 3$, the 95% confidence interval CI of the actual mean RW was [0.37 4.02] (the ideal $RW = 3.45$). For $n = 15$, the 95% CI of the actual mean RW was [0.24 4.73] (the ideal $RW = 0.66$). For $n = 75$, the 95% CI of the actual mean RW was [−0.31

0.66] (the ideal $RW = 0.12$). In all three conditions, the ideal RW fell within the CI of actual RW . Therefore, we concluded that the actual RW was statistically indistinguishable from the ideal RW predicted by Bayesian integration.

Summary of behavioral results

In three separate analyses, we found that subjects integrated information about prior and likelihood in a fashion that was statistically indistinguishable from the ideal Bayesian integration. The first analysis demonstrated that a model based on using the posterior information about reward probability (influenced by both prior and likelihood information) better described the choice data than either the prior or the likelihood information alone. The second analysis further revealed that the subjects considered the mean of prior distribution the mean and the SD of the likelihood information when making decisions. Finally, the third analysis quantitatively revealed that the relative weight subjects assigned to prior and likelihood changed in response to changes in sample size and that the relative weights were statistically indistinguishable from an ideal decision maker that optimally integrates prior and likelihood information.

Neural correlates of the posterior distribution

In this study, the posterior distribution is the posterior probability density function on the probability of reward associated with the symbol lottery. The posterior distribution reflects the integration of prior knowledge and current sensory information in the form of the likelihood function. We first looked at regions in the brain for which activity correlated with information about the posterior distribution. In particular, we looked for regions that correlated with the mean of the distribution.

We found that the mPFC correlated with the mean of the posterior distribution (μ_π ; *GLM-1*; see Figure 3A). The maximum z -statistic = 3.95 at voxel coordinates $[x, y, z] = [4, 48, 2]$ (MNI space). The results were corrected for multiple comparisons ($z > 3$, familywise error-corrected $p < 0.05$ using Gaussian random field theory). At this threshold, the mPFC was the only region that correlated with μ_π . To understand how mPFC represents μ_π under different sample-size conditions, we constructed independent ROIs in mPFC based on the μ_π contrast and estimated the time series of effect size associated with μ_π under different sample-size conditions (see Materials and Methods: "Analysis 1, Time-series analysis"). The results are plotted in Figure 3B. The inset shows a summary of the mPFC ROI used in this analysis. Because of the method used (see Materials and Methods: "Construction of independent, unbiased region of interest"), each subject had a different mPFC mask that varied slightly in location. To get a general summary of the location of the mPFC ROIs, we added the ROI used for each subject before binarizing the summed mask. The mPFC ROI shown in yellow thus depicts the voxels that were from at least one subject's ROI. The black thick line in the figure indicated the time points we were interested in (4–12 s after stimulus onset).

Looking at the time-series results, we first observed similar dynamics across all three conditions. The effect size of μ_π , indicated by its z -statistic, was smallest at 4 s and gradually increased before reaching its peak. In both $n = 3$ and $n = 15$ conditions, the effect size of μ_π peaked at 6–8 s after the onset of the symbol lottery and was significant ($p < 0.05$, corrected for the number of time points tested). Conversely, when sample size was the largest ($n = 75$), no time point between 4 and 12 s after stimulus onset was significantly correlated with μ_π . Although the result was not significant, it still carried some useful information. The peak was

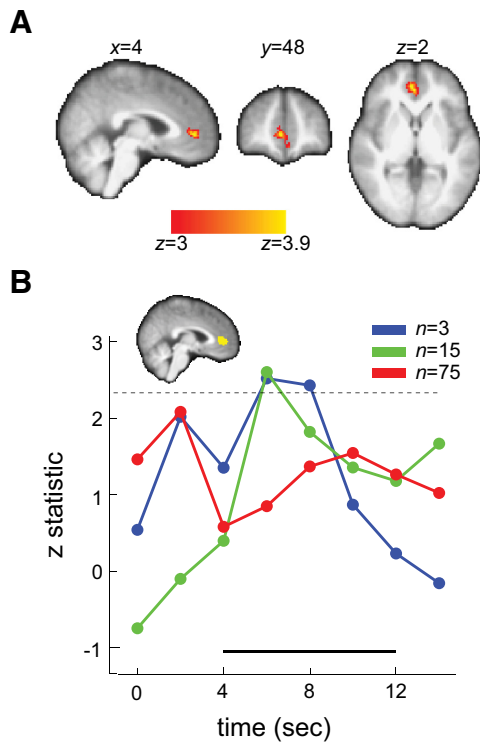


Figure 3. Neural representation of posterior information in the mPFC. **A**, Whole-brain results. Activity in mPFC was positively correlated with the mean of the posterior distribution on reward probability associated with the symbol lottery when prior and likelihood information were presented ($z > 3$, familywise error-corrected $p < 0.05$ using Gaussian random field theory). **B**, Independent ROI analysis. We constructed independent, unbiased ROIs in mPFC that significantly tracked the posterior information. We then extracted the BOLD time series from these independent, unbiased ROIs, estimated the effect size of the posterior information, transformed it to z statistic, and plotted it as a function of time. Time point 0 indicates onset of symbol lottery presentation. The thick black line at the bottom indicates the time points of interest. The dashed horizontal line indicated the z -threshold corresponding to $p = 0.05$ corrected for the number of time points tested.

~10 s after the stimulus onset. Compared with peaks in the other two conditions, this indicated that the posterior computation in the large sample-size trials might take more time to perform, possibly reflecting the visual processing associated with more dots shown on the screen.

The posterior distribution is the result of the integration of prior knowledge and likelihood information (Eq. 4). The analysis just described showed regions that correlate with information about the posterior distribution, the result of such integration. In the following analysis, we looked for regions that correlate with information about prior distribution and likelihood function.

Neural correlates of the prior distribution and likelihood function

Using a GLM that separately modeled prior distribution and likelihood function (GLM-2), we found that regions in the mPFC and the dorsal anterior cingulate cortex (dACC) correlated with subjects' prior knowledge about probability of reward when information about the symbol lottery was revealed. Activity in these regions correlated with the mean of the prior distribution on reward probability (μ_π) ($z > 3$, familywise error-corrected $p < 0.05$ using Gaussian random field theory). At this threshold, the cluster including mPFC (maximum z -statistic = 3.68, $[x, y, z] = [10, 44, -4]$) and dACC (maximum z -statistic = 4.03, $[x, y, z] = [-10, 44, 20]$) was the only region that correlated with μ_π (Fig. 4A). For the likelihood function, no region was significantly

correlated with μ_π during the period in which the symbol lottery was presented at the 0.05 level after correcting for multiple comparisons using Gaussian random field theory with a cluster-forming threshold $z > 2.3$.

Given that mPFC represented both the posterior and prior information, we next focused on this region with more detailed ROI analyses. Two ROI analyses were performed and both were independent and unbiased. In the first analysis, we looked at the representation of prior and likelihood information across all trials (Fig. 4B). In the second analysis, we investigated how prior and likelihood representations in mPFC might change under different sample-size conditions (Fig. 5). The mPFC ROIs used and summarized on the inset in Figures 4B and 5 were the same as in Figure 3B.

In the first analysis (see Materials and Methods: "Analysis 2, Time-series analysis"), we found significant representations of both prior and likelihood information in mPFC ($p < 0.05$, corrected for the number of time points tested; Fig. 4B). Significant prior representation was consistent with the results from the whole-brain analysis. Significant likelihood representation, however, was not consistent with whole-brain results. We found that, in contrast to an earlier plateau of prior information coding (6–8 s after stimulus onset), the representation of likelihood information had a later peak (~10 s after stimulus onset). Recall that information about the symbol lottery was presented for 3 s before entering the delay period (Fig. 1D). Given a typical hemodynamic lag of 6 s, this result indicated that likelihood coding in mPFC might not take place when information about the symbol lottery was revealed. Instead, it took place after that, during the delay period. This could explain why we did not find significant representation of likelihood in the whole-brain analysis because the GLM attempted to capture likelihood representation during symbol-lottery presentation.

Another distinct pattern that we observed from the ROI analysis is the ordering in information coding in mPFC. Despite the fact that we revealed information about prior (the icon) and likelihood (the dots) simultaneously, BOLD signals in mPFC indicated that prior information was coded first (as it had an earlier peak) and was only later by the coding of likelihood information. This temporal ordering may simply be due to the experimental design. First, likelihood information was presented in the form of dots, whereas prior information was represented by a symbol icon. It is plausible that the amount of time required to process these two sources of information was simply different. Second, familiarity might also contribute. Subjects should be relatively more familiar with the symbol icons because of the prior learning session that took place before the fMRI session.

In the second independent ROI analysis (see Materials and Methods: "Analysis 3, Time-series analysis"), we analyzed prior and likelihood coding for each sample-size condition separately. The results are shown in Figure 5. First, we noticed very similar patterns for $n = 3$ and $n = 15$. In both conditions, prior coding dominated the mPFC signals 6–8 s after stimulus onset. This is consistent with the previous analyses showing that mPFC represented prior information when symbol lottery was presented. Conversely, likelihood representation was not significant between 4 and 12 s after stimulus onset. Even though the results were not significant, we observed that, in both cases, likelihood coding peaked at 10 s after stimulus onset.

In contrast to significant prior coding but nonsignificant likelihood coding, we observed the opposite pattern when sample size was large ($n = 75$). Likelihood representation rose and peaked at 12 s after stimulus onset. At this point, it was marginally

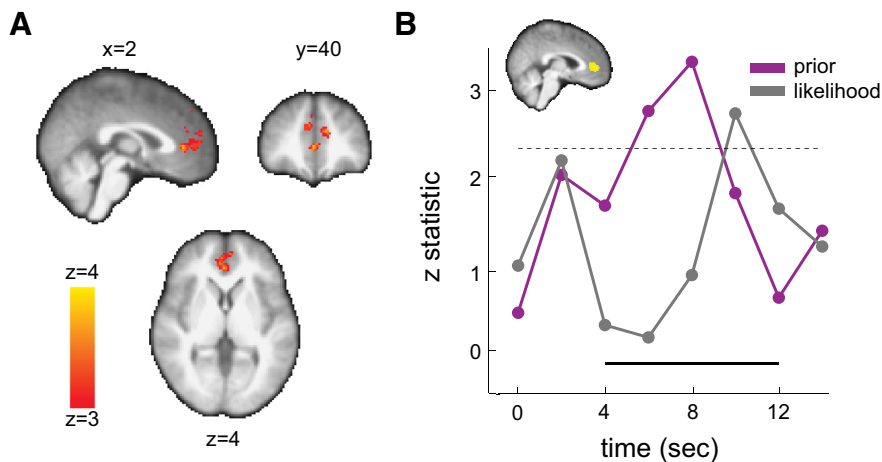


Figure 4. Neural representation of prior and likelihood information in the mPFC. **A**, Whole-brain results. Activity in mPFC was positively correlated with the mean of the prior distribution on reward probability associated with the symbol lottery when prior and likelihood information were revealed ($z > 3$, familywise error-corrected $p < 0.05$ using Gaussian random field theory). **B**, Independent ROI analysis. Using the independent ROIs identified as in Figure 3B and an analysis identical to that in Figure 3B, we plot the z-statistic time series of effect size for the prior information (purple) and the likelihood information (gray). Prior information here is the mean of the prior distribution on reward probability associated with the symbol lottery (μ_π), whereas likelihood information is the mean of the likelihood function (μ_L). The thick black line at the bottom indicates the time points of interest. The dashed horizontal line indicated the z-threshold corresponding to $p = 0.05$ corrected for the number of time points tested.

significant ($p = 0.014$). Conversely, prior representation changed little from 4–12 s after stimulus onset and was not significant at all the time points tested during this period.

These results together were qualitatively consistent with the behavioral data. In the logistic regression analysis on choice we performed, we found that subjects weighted prior information more than likelihood information when $n = 3$ and $n = 15$. Recall that RW is the ratio of the estimated β coefficient of prior information (μ_π) to that of likelihood (μ_L). We found that the mean RW across subjects was identical between the two conditions ($RW = 2.25$ when $n = 3$, $RW = 2.14$ when $n = 15$). This is consistent with the neural findings that prior information coding dominated over likelihood coding in mPFC activity in both of these conditions (Fig. 5A,B). In contrast, when sample size was large ($n = 75$), RW inferred from choice behavior suggested that subjects weighed likelihood information more than prior information ($RW = 0.18$). This is consistent with the results showing marginally significant likelihood coding, but nonsignificant prior coding in mPFC (Fig. 5C).

Based on these observations, we next attempted to define RW quantitatively at the neural level based on the results of the time-series analysis. To quantitatively estimate the neural RW , we did the following exploratory analysis. For each sample-size condition, we identified the maximum z-statistic during the time points of interest (4–12 s after stimulus onset) separately for μ_π and μ_L . The maximum z-statistic reflected the highest correlation with the variables of interest during this period. Let z_π^{\max} denote the maximum z-statistic of μ_π and z_L^{\max} denote the maximum z-statistic for μ_L . We defined the neural measure of RW by z_π^{\max}/z_L^{\max} .

We plot the behaviorally estimated RW against the neural RW in Figure 5D. The behavioral RW associated with each sample-size condition plotted here was identical to that shown in Figure 2D (the “actual RW ” plotted on the y-axis). Based on the 95% confidence interval of the behavioral RW (see Results: “Analysis 3: Ideal decision maker analysis”), the behavioral RW was statistically indistinguishable from the neural RW .

Neural correlates of likelihood uncertainty
We defined likelihood uncertainty by the SD of the likelihood function σ_L , which reflects how “trustworthy” information about the likelihood of reward is in the data. In the behavioral analysis, we found that likelihood uncertainty correlated with choice: subjects were less likely to choose the symbol lottery when σ_L became greater. Whole-brain GLM ($GLM-3$) indicated that BOLD response in a distributed network of regions including the occipital cortex and the striatum negatively correlated with σ_L (Fig. 6, in yellow) ($z > 2.6$, familywise error-corrected $p < 0.05$ using Gaussian random field theory). See Table 3 for a complete list of regions that were negatively correlated with σ_L .

This outcome indicated that activity in these regions increased when likelihood uncertainty decreased. There was no region that positively correlated with likelihood uncertainty. The sign of correlation was consistent the results of logistic regression on choice, which showed that the likelihood uncertainty negatively correlated with the probability of choosing the symbol lottery.

One potential concern for the fMRI results on likelihood uncertainty is its correlation with sample size, which is the number of dots shown on the screen. Given that likelihood uncertainty negatively correlates with sample size, activity correlated with likelihood uncertainty can be driven simply by sample size. To address this concern, we included a separate regressor for sample size in the GLM (see Materials and Methods: “ $GLM-2$ ” and “ $GLM-3$ ” for regressor descriptions). We found that, even when sample size was added as a parametric regressor, the regions reported above that correlated with likelihood uncertainty were still activated. This indicated that activity in these regions correlated with likelihood uncertainty and cannot simply be explained by the number of dots presented in the trial.

For the interaction between σ_L and μ_L ($\mu_L \times \sigma_L$, $GLM-3$), we found that the occipital cortex significantly and positively correlated with it (Fig. 6, in blue) ($z > 2.6$, familywise error-corrected $p < 0.05$ using Gaussian random field theory). No negative correlation was found at the whole-brain level. For the interaction between σ_L and μ_π ($\mu_\pi \times \sigma_L$, $GLM-3$), we found that BOLD response in the ventral striatum positively correlated with it (Fig. 6, in green; $z > 2.6$, familywise error-corrected $p < 0.05$ using Gaussian random field theory).

Although we did not find behaviorally that these two interaction terms were significant (see Results: “Analysis 2: Logistic regression on choice and the use of prior and likelihood information”), the significant correlation between BOLD response in these regions and the interaction terms provided evidence for the neural systems, in addition to mPFC, that may be involved in the integration process.

Discussion

We investigated how the brain represents prior knowledge and likelihood based on current sensory information and how these two sources of information are integrated. To study prior-likelihood integration, we designed a novel lottery decision task.

We systematically manipulated prior knowledge in the form of a probability density function on reward probability; we systematically manipulated current sensory information in the form of a likelihood function by varying the size of the sample of the sensory data on reward probability. We compared subjects' choice behavior to the predictions of Bayesian decision theory.

Integration of prior and likelihood information

We found that subjects flexibly adjust the weights assigned to prior and likelihood information. When likelihood information was less reliable because it was based on a smaller sample size, subjects tended to weight prior information more heavily than likelihood and vice versa. The weights subjects assigned to these two sources of information were statistically indistinguishable from the predictions based on Bayesian decision theory (Fig. 2D).

The relative weight given to prior and likelihood information had received a great deal of attention in various disciplines. This is a key issue in judgment under uncertainty and the subject of extensive research in psychology (Kahneman et al., 1982; Gigerenzer and Hoffrage, 1995; Zhang and Maloney, 2012), in experimental economics and finance (Camerer, 1987; Friedman, 1998; Kluger and Wyatt, 2004; Slembeck and Tyran, 2004; Kluger and Friedman, 2010), and in human perception and action (Geisler, 1989; Knill and Richards, 1996; Körding and Wolpert, 2006; Trommershäuser et al., 2008; Zhang and Maloney, 2012; Pouget et al., 2013). One key issue in this literature is to understand how humans select weights for these different sources of information. The second issue concerns how actual performance might deviate from that predicted by an ideal observer or decision maker that optimally integrates prior and likelihood information. For example, it has been shown that, in many cognitive judgment tasks, subjects are suboptimal in that they tend to assign more weights to current information (likelihood) than they should, that is, the base-rate fallacy (Kahneman et al., 1982).

Other studies, such as in multisensory integration and motor decision making, revealed that humans integrate prior and likelihood information nearly optimally (Ernst and Banks, 2002; Körding and Wolpert, 2004; Tassinari et al., 2006). What we added to the existing literature is that, in the domain of reward probability, human subjects flexibly adjust the weights assigned to likelihood and prior information and their adjustments are consistent with the predictions based on Bayesian decision theory.

Neural representations of prior, likelihood, and posterior information

There is a growing interest in decision neuroscience in studying the effect of prior expectation on behavioral performance and the underlying neural mechanisms, ranging from perceptual deci-

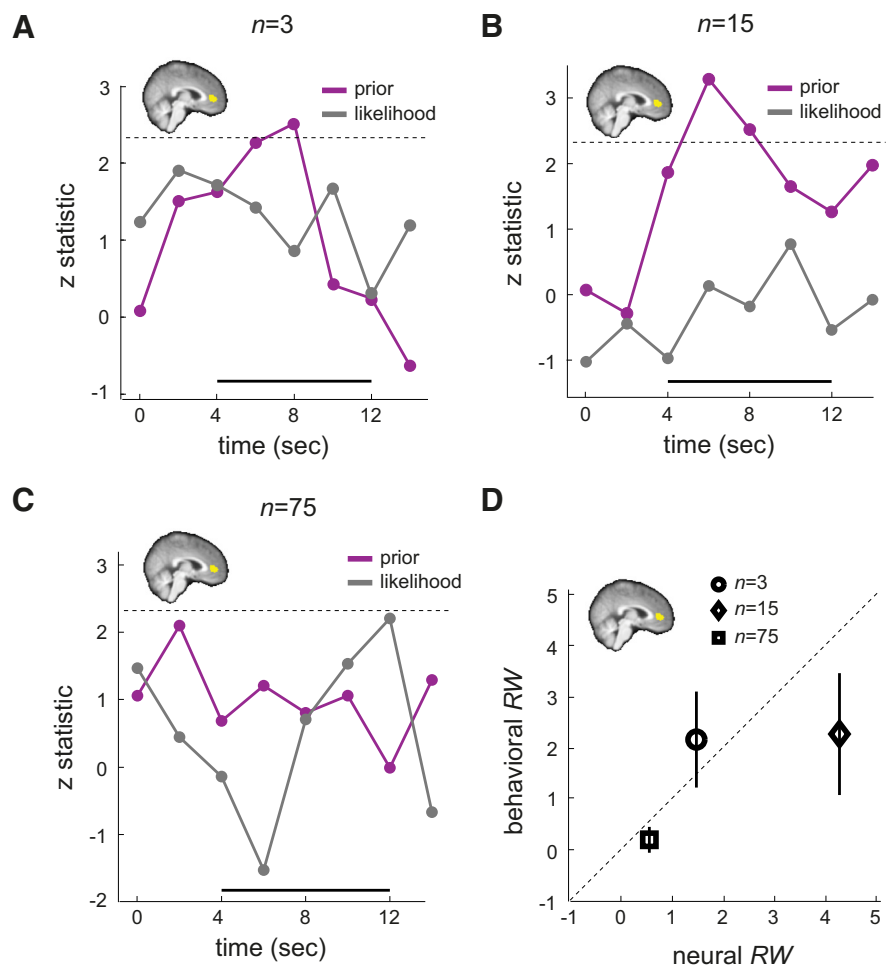


Figure 5. Independent ROI analysis on mPFC. For the independent, unbiased ROIs in mPFC that were found to track posterior information, we estimated the time-series of effect size for prior information (purple) and likelihood (gray) separately for each sample-size condition. The y-axis represented the z-statistic of the effect size. **A**, When sample size was small ($n = 3$). **B**, When sample size was medium ($n = 15$). **C**, When sample size was large ($n = 75$). The thick black line at the bottom indicates the time points of interest. The dashed horizontal line indicated the z-threshold corresponding to $p = 0.05$ corrected for the number of time points tested. **D**, Behaviorally estimated relative weight (behavioral RW) is plotted against the neural measure of RW (neural RW) obtained from mPFC. Error bars represent ± 1 SEM. Different symbols represent different sample-size conditions.

sion making (Forstmann et al., 2010; Rahnev et al., 2011; Mulder et al., 2012; Kok et al., 2013), reinforcement learning (Li et al., 2011), belief updating in financial decision making in the presence of social information (Huber et al., 2014), and in trust games (Fouragnan et al., 2013). Prior expectations in most studies were manipulated by varying the reliability of cues that predicts the upcoming stimulus (Forstmann et al., 2010; Li et al., 2011; Rahnev et al., 2011; Mulder et al., 2012; Kok et al., 2013) or simply the presence or absence of prior information (Fouragnan et al., 2013).

In perceptual decision making, the reliability of prior information is represented in a frontoparietal network (Rahnev et al., 2011; Mulder et al., 2012) that modulates the effective connectivity between sensory systems and the dorsolateral prefrontal cortex (DLPFC) (Rahnev et al., 2011) and affects the sensory representations in the visual cortex in a way that is consistent with prior-induced perceptual bias (Kok et al., 2013). In a financial decision task, Huber et al., (2014) found that frontoparietal regions including the inferior parietal cortex and the DLPFC update beliefs about option value in the presence of social information. To summarize, a growing body of work has begun to reveal the neural systems involved in representing prior expectations and the compu-

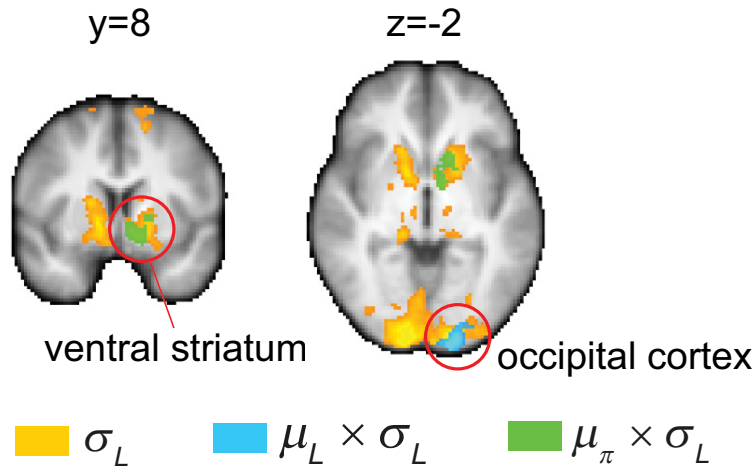


Figure 6. Neural correlates of likelihood uncertainty. Whole-brain results ($z > 2.6$, familywise error-corrected $p < 0.05$ using Gaussian random field theory) showing regions that significantly correlated with likelihood uncertainty (σ_L , in yellow), the interaction between likelihood uncertainty and the mean of the likelihood function (μ_L , frequency of reward indicated in the current sensory information), $\mu_L \times \sigma_L$ (in blue), and the interaction between σ_L and the mean of the prior distribution (μ_π), $\mu_\pi \times \sigma_L$ (in green). The occipital cortex correlated with both σ_L and $\mu_L \times \sigma_L$. The ventral striatum correlated with both σ_L and $\mu_\pi \times \sigma_L$.

Table 3. Regions in which the BOLD signal was negatively correlated with the SD of the likelihood function (likelihood uncertainty)

Cluster	Hemisphere	Cluster size	z-max	z-max (x, y, z)
Occipital cortex	L	6550	5.67	(−14, −98, 0)
Striatum	L	3639	4.39	(−8, 4, −8)
Inferior parietal lobule	R	1395	4.28	(54, −24, 38)
Inferior parietal lobule	L	917	3.87	(−46, −32, 34)
Precuneous	L	897	3.61	(−12, −66, 54)
Frontal pole	L	583	3.69	(−32, 50, 8)
Posterior cingulate cortex		450	3.76	(0, −24, 32)
Superior frontal gyrus	L	371	3.71	(−28, −4, 62)

Clusters that survived cluster-based correction ($z > 2.6$, familywise error-corrected $p < 0.05$ using Gaussian random field theory). The z-max column represents the MNI coordinates of the maximum z-statistic.

tational mechanisms for how prior information affects decision computations.

Our study adds to this literature by revealing the neural mechanisms underlying the integration of prior and likelihood information in the computation of reward probability. First, we found that the mPFC represented the posterior information (the mean of the posterior distribution on reward probability; Fig. 3). Because such information is a result of the integration of prior and likelihood information, this finding indicated that mPFC is a candidate region for performing such integration. Second, in a separate analysis, we found that mPFC represented both prior and likelihood information (Fig. 4). Third, we found that there was a close match between subjects' behavior and mPFC signals while likelihood and prior information were being evaluated (Fig. 5). The neural measure of relative weight (the weight assigned to prior information relative to the likelihood information) was statistically indistinguishable from the behavioral measure of relative weight (Fig. 5D).

The mPFC and the orbitofrontal cortex had also been shown to be involved in valuation during value-based decision making (Wallis and Miller, 2003; Padoa-Schioppa and Assad, 2006; Kable and Glimcher, 2007; Plassmann et al., 2007; Tom et al., 2007; Chib et al., 2009; for reviews, see Kable and Glimcher, 2009; Padoa-Schioppa, 2011; Wallis and Kennerley, 2010; Levy and Glimcher, 2011; Wu et al., 2011; for meta analysis, see Clithero and Rangel, 2013).

Despite the notion that valuation and the integration of prior and current information are intimately related, it remained an open question as to whether these two computations are performed by the same neural system. What we showed is that, at least in the domain of reward probability, the previously identified canonical valuation network is also engaged in integrating prior knowledge and current sensory information to estimate reward probability.

In a recent study, d'Acremont et al. (2013) investigated the neural mechanisms for the integration of prior and likelihood information in the computation of event probability. In two different tasks, subjects had to estimate the probability of occurrence associated with two complementary events to obtain monetary rewards. Both prior and likelihood information were systematically manipulated.

They found that a distributed frontoparietal network including the mPFC encoding likelihood information (event frequency) independent of prior knowledge, whereas a separate and distinct network including the inferior frontal gyrus represented the Bayesian posterior probability (the integration of prior and likelihood information).

Our results are consistent with d'Acremont et al. (2013) concerning likelihood coding in mPFC. Although the present study highlighted the role of mPFC in integrating prior and likelihood information, the d'Acremont et al. (2013) study revealed separate systems for encoding likelihood and posterior information. These differences in findings likely reflect differences in the variable of interest and in task design.

First, in the d'Acremont et al. (2013) study, the variable of interest was event probability, whereas we were interested in reward probability. Second, although subjects had to accumulate samples sequentially to compute likelihood and thus the posterior information in d'Acremont et al. (2013), likelihood information was not revealed sequentially in the present study. This could in principle explain why a distributed frontoparietal network, which had been shown to be critical to evidence accumulation (Heekeren et al., 2004; Heekeren et al., 2006; Kahnt et al., 2011; Liu and Pleskac, 2011; Filimon et al., 2013; Hebart et al., 2014), was involved in encoding and updating event frequency in the d'Acremont et al. (2013) study but not in our study. Finally, there were differences in how subjects acquired prior information in the two studies. Subjects in the present study established prior knowledge through an extensive learning session before the fMRI session, whereas prior information in d'Acremont et al. (2013) was explicitly revealed in graphical form on a trial-by-trial basis. Together, these differences could subsequently affect the neural systems involved in the integration of prior and likelihood information. Future investigations with an equivalent task design are therefore needed to further interpret the differences in findings.

Neural representations of likelihood uncertainty

Our neural results on likelihood uncertainty were in part consistent with a recent study (Vilares et al., 2012) that also found evidence for the representation of likelihood uncertainty in the occipital cortex. What we found is that, in addition to the occipital cortex, the precuneus, inferior parietal cortex, frontal pole,

and striatum are also involved in representing likelihood uncertainty. Furthermore, we also found that the occipital cortex correlated with the interaction between likelihood uncertainty and information about the frequency of reward in the likelihood function and that the ventral striatum correlated with the interaction between prior knowledge and likelihood uncertainty. Together, these results suggest that, in addition to mPFC, these regions may also participate in the integration process. Future studies are needed to better understand the exact computation roles of these regions on prior-likelihood integration. For example, it is possible that the involvement of the occipital cortex was driven by how we presented the likelihood information. It remains to be seen whether the same regions would encode these variables when one replaces the current way of presentation (showing colored dots) with other forms (e.g., showing numbers).

References

- Bartra O, McGuire JT, Kable JW (2013) The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *Neuroimage* 76:412–427. [CrossRef Medline](#)
- Beckmann CF, Jenkinson M, Smith SM (2003) General multilevel linear modeling for group analysis in fMRI. *Neuroimage* 20:1052–1063. [CrossRef Medline](#)
- Berger JO (1985) *Statistical decision theory and Bayesian analysis*. New York: Springer.
- Boorman ED, O'Doherty JP, Adolphs R, Rangel A (2013) The behavioral and neural mechanisms underlying the tracking of expertise. *Neuron* 80:1558–1571. [CrossRef Medline](#)
- Brainard DH (1997) The psychophysics toolbox. *Spatial Vision* 10:433–436. [CrossRef Medline](#)
- Calhoun VD, Stevens MC, Pearlson GD, Kiehl KA (2004) fMRI analysis with the general linear model: removal of latency-induced amplitude bias by incorporation of hemodynamic derivative terms. *Neuroimage* 22:252–257. [CrossRef Medline](#)
- Camerer C (1987) Do biases in probability judgment matter in markets? Experimental evidence. *American Economic Review* 77:981–997.
- Chib VS, Rangel A, Shimojo S, O'Doherty JP (2009) Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. *J Neurosci* 29:12315–12320. [CrossRef Medline](#)
- Clithero JA, Rangel A (2013) Informatic parcellation of the network involved in the computation of subjective value. *Soc Cogn Affect Neurosci* 9:1289–1302. [CrossRef Medline](#)
- d'Acremont M, Schultz W, Bossaerts P (2013) The human brain encodes event frequencies while forming subjective beliefs. *J Neurosci* 33:10887–10897. [CrossRef Medline](#)
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ (2006) Cortical substrates for exploratory decisions in humans. *Nature* 441:876–879. [CrossRef Medline](#)
- Efron B, Tibshirani RJ (1993) *An introduction to the bootstrap*. New York: Chapman and Hall.
- Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415:429–433. [CrossRef Medline](#)
- Esterman M, Tamber-Rosenau BJ, Chiu YC, Yantis S (2010) Avoiding non-independence in fMRI data analysis: leave one subject out. *Neuroimage* 50:572–576. [CrossRef Medline](#)
- Filimon F, Philastides MG, Nelson JD, Kloosterman NA, Heekeren HR (2013) How embodied is perceptual decision making? Evidence for separate processing of perceptual and motor decisions. *J Neurosci* 33:2121–2136. [CrossRef Medline](#)
- Fiorillo CD, Tobler PN, Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299:1898–1902. [CrossRef Medline](#)
- Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC (1995) Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of cluster-size threshold. *Magn Reson Med* 33:636–647. [CrossRef Medline](#)
- Forstmann BU, Brown S, Dutilh G, Neumann J, Wagenmakers EJ (2010) The neural substrate of prior information in perceptual decision making: a model-based analysis. *Front Hum Neurosci* 4:40. [CrossRef Medline](#)
- Fouragnan E, Chierchia G, Greiner S, Neveu R, Avesani P, Coricelli G (2013) Reputational priors magnify striatal responses to violation of trust. *J Neurosci* 33:3602–3611. [CrossRef Medline](#)
- Friedman D (1998) Monty hall's three doors: Construction and deconstruction of a choice anomaly. *American Economic Review* 88:933–946.
- Geisler WS (1989) Sequential ideal-observer analysis of visual discriminations. *Psychol Rev* 96:267–314. [CrossRef](#)
- Gigerenzer G, Hoffrage U (1995) How to improve Bayesian reasoning without instruction: Frequency formats. *Psychol Rev* 102:684–704. [CrossRef](#)
- Gilboa I (2010) *Rational choice*. Cambridge, MA: MIT.
- Hebart MN, Schriever Y, Donner TH, Haynes J-D (2014) The relationship between perceptual decision variables and confidence in the human brain. *Cereb Cortex*. In press.
- Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG (2004) A general mechanism for perceptual decision-making in the human brain. *Nature* 431:859–862. [CrossRef Medline](#)
- Heekeren HR, Marrett S, Ruff DA, Bandettini PA, Ungerleider LG (2006) Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *Proc Natl Acad Sci U S A* 103:10023–10028. [CrossRef Medline](#)
- Henson RNA, Buechel C, Josephs O, Friston KJ (1999) The slice-timing problem in event-related fMRI. *Neuroimage* 9:S125.
- Huber RE, Klucharev V, Rieskamp J (2014) Neural correlates of informational cascades: brain mechanisms of social influence on belief updating. *Soc Cogn Affect Neurosci*. In press.
- Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17:825–841. [CrossRef Medline](#)
- Kable JW, Glimcher PW (2007) The neural correlates of subjective value during intertemporal choice. *Nat Neurosci* 10:1625–1633. [CrossRef Medline](#)
- Kable JW, Glimcher PW (2009) The neurobiology of decision: consensus and controversy. *Neuron* 63:733–745. [CrossRef Medline](#)
- Kahneman D, Slovic P, Tversky A (1982) *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University.
- Kahnt T, Grueschow M, Speck O, Haynes JD (2011) Perceptual learning and decision-making in human medial frontal cortex. *Neuron* 70:549–559. [CrossRef Medline](#)
- Kluger BD, Friedman D (2010) Financial engineering and rationality: Experimental evidence based on the Monty Hall problem. *J Behav Finance* 11:31–49. [CrossRef](#)
- Kluger BD, Wyatt SB (2004) Are judgment errors reflected in market prices and allocations? Experimental evidence based on the Monty Hall problem. *J Finance* 59:969–998. [CrossRef](#)
- Knill DC, Richards W (1996) *Perception as Bayesian inference*. Cambridge: Cambridge University.
- Kok P, Brouwer GJ, van Gerven MA, de Lange FP (2013) Prior expectations bias sensory representations in visual cortex. *J Neurosci* 33:16275–16284. [CrossRef Medline](#)
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247. [CrossRef Medline](#)
- Körding KP, Wolpert DM (2006) Bayesian decision theory in sensorimotor control. *Trends Cogn Sci* 10:319–326. [CrossRef Medline](#)
- Levy DJ, Glimcher PW (2011) Comparing apples and oranges: Using reward-specific and reward-general subjective value representation in the brain. *J Neurosci* 31:14693–14707. [CrossRef Medline](#)
- Li J, Delgado MR, Phelps EA (2011) How instructed knowledge modulates the neural systems of reward learning. *Proc Natl Acad Sci U S A* 108:55–60. [Medline](#)
- Litt A, Plassmann H, Shiv B, Rangel A (2011) Dissociating valuation and saliency signals during decision-making. *Cereb Cortex* 21:95–102. [Medline](#)
- Liu T, Pleskac TJ (2011) Neural correlates of evidence accumulation in a perceptual decision task. *J Neurophysiol* 106:2383–2398. [CrossRef Medline](#)
- Mulder MJ, Wagenmakers EJ, Ratcliff R, Boekel W, Forstmann BU (2012) Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *J Neurosci* 32:2335–2343. [CrossRef Medline](#)
- O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2002) Temporal difference models and reward-related learning in the human brain. *Neuron* 38:239–337. [Medline](#)

- O'Hagan A, Forster J, Kendall MG (2004) Bayesian inference. London: Arnold.
- Padoa-Schioppa C (2011) Neurobiology of economic choice: a good-based model. *Annu Rev Neurosci* 34:333–359. [CrossRef Medline](#)
- Padoa-Schioppa C, Assad JA (2006) Neurons in the orbitofrontal cortex encode economic value. *Nature* 441:223–226. [CrossRef Medline](#)
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision* 10:437–442. [CrossRef Medline](#)
- Plassmann H, O'Doherty J, Rangel A (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J Neurosci* 27:9984–9988. [CrossRef Medline](#)
- Platt ML, Huettel SA (2008) Risky business: the neuroeconomics of decision making under uncertainty. *Nat Neurosci* 11:398–403. [CrossRef Medline](#)
- Pouget A, Beck JM, Ma WJ, Latham PE (2013) Probabilistic brains: known and unknowns. *Nat Neurosci* 16:1170–1178. [CrossRef Medline](#)
- Rahnev D, Lau H, de Lange FP (2011) Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. *J Neurosci* 31:10741–10748. [CrossRef Medline](#)
- Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. *Science* 275:1593–1599. [CrossRef Medline](#)
- Slembeck T, Tyran JR (2004) Do intuitions promote rationality? An experimental study of the three-door anomaly. *J Econ Behav Organ* 54:337–350. [CrossRef](#)
- Smith SM, et al. (2004) Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23:S208–S219. [Medline](#)
- Tassinari H, Hudson TE, Landy MS (2006) Combining priors and noisy visual cues in a rapid pointing task. *J Neurosci* 26:10154–10163. [CrossRef Medline](#)
- Tom SM, Fox CR, Trepel C, Poldrack RA (2007) The neural basis of loss aversion in decision-making under risk. *Science* 315:515–518. [CrossRef Medline](#)
- Trommershäuser J, Maloney LT, Landy MS (2008) Decision making, movement planning and statistical decision theory. *Trends Cogn Sci* 12:291–297. [CrossRef Medline](#)
- Vilares I, Howard JD, Fernandes HL, Gottfried JA, Kording KP (2012) Differential representations of prior and likelihood uncertainty in the human brain. *Curr Biol* 22:1641–1648. [CrossRef Medline](#)
- Wallis JD, Kennerly SW (2010) Heterogeneous reward signals in prefrontal cortex. *Curr Opin Neurobiol* 20:191–198. [CrossRef Medline](#)
- Wallis JD, Miller EK (2003) Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci* 18:2069–2081. [CrossRef Medline](#)
- Woolrich MW, Ripley BD, Brady M, Smith SM (2001) Temporal autocorrelation in univariate linear modeling of fMRI data. *Neuroimage* 14:1370–1386. [CrossRef Medline](#)
- Woolrich MW, Behrens TE, Beckmann CF, Jenkinson M, Smith SM (2004) Multilevel linear modeling for FMRI group analysis using Bayesian inference. *Neuroimage* 21:1732–1747. [CrossRef Medline](#)
- Worsley KJ, Evans AC, Marrett S, Neelin P (1992) A three-dimensional statistical analysis for CBF activation studies in human brain. *J Cereb Blood Flow Metab* 12:900–918. [CrossRef Medline](#)
- Wu SW, Delgado MR, Maloney LT (2011) The neural correlates of subjective utility of monetary outcome and probability weight in economic and in motor decision under risk. *J Neurosci* 31:8822–8831. [CrossRef Medline](#)
- Zhang H, Maloney LT (2012) Ubiquitous log odds: A common representation of probability and frequency distortion in perception, action, and cognition. *Front Neurosci* 6.