

Neural Mechanisms of Updating under Reducible and Irreducible Uncertainty

 Kenji Kobayashi¹ and  Ming Hsu^{1,2}

¹Helen Wills Neuroscience Institute and ²Haas School of Business, University of California, Berkeley, Berkeley, California 94720

Adaptive decision making depends on an agent's ability to use environmental signals to reduce uncertainty. However, because of multiple types of uncertainty, agents must take into account not only the extent to which signals violate prior expectations but also whether uncertainty can be reduced in the first place. Here we studied how human brains of both sexes respond to signals under conditions of reducible and irreducible uncertainty. We show behaviorally that subjects' value updating was sensitive to the reducibility of uncertainty, and could be quantitatively characterized by a Bayesian model where agents ignore expectancy violations that do not update beliefs or values. Using fMRI, we found that neural processes underlying belief and value updating were separable from responses to expectancy violation, and that reducibility of uncertainty in value modulated connections from belief-updating regions to value-updating regions. Together, these results provide insights into how agents use knowledge about uncertainty to make better decisions while ignoring mere expectancy violation.

Key words: ambiguity; decision making; fMRI; risk; updating; valuation

Significance Statement

To make good decisions, a person must observe the environment carefully, and use these observations to reduce uncertainty about consequences of actions. Importantly, uncertainty should not be reduced purely based on how surprising the observations are, particularly because in some cases uncertainty is not reducible. Here we show that the human brain indeed reduces uncertainty adaptively by taking into account the nature of uncertainty and ignoring mere surprise. Behaviorally, we show that human subjects reduce uncertainty in a quasioptimal Bayesian manner. Using fMRI, we characterize brain regions that may be involved in uncertainty reduction, as well as the network they constitute, and dissociate them from brain regions that respond to mere surprise.

Introduction

Adaptive decision making in the real world depends critically on an agent's ability to constantly make use of environmental signals to reduce uncertainty. Agents should take into account not only physical properties of signals but also the nature of uncertainty, as not all kinds of uncertainty can be reduced in the same way. In particular, two types of uncertainty, risk and ambiguity, have received widespread emphasis in decision-making literature (Keynes, 1921; Knight, 1921; Ellsberg, 1961). While ambiguity can be reduced by signals that carry new information and supplement an agent's prior knowledge about the environment, risk cannot be reduced by any signals.

To illustrate the implication of these types of uncertainty for the adaptive use of signals, consider a variation on the classic example of gambler's fallacy (Tversky and Kahneman, 1974), a gamble that depends on tosses of two coins A and B. You know that coin A is fair (risk), but you do not know whether coin B is biased or fair (ambiguity). If you observe 10 consecutive "heads" of B, it would suggest that B is biased, but the same sequence of A does not provide any new information. Although these two signals are similar and both surprising, appropriate decision making requires agents to use them differently because of the different natures of uncertainty.

This, however, poses a challenge for theoretical frameworks, such as reinforcement learning (RL), which do not incorporate explicit notions of uncertainty. Under RL, values of actions are updated to the extent that an observed outcome violates prior expectancy, such that agents cannot ignore merely expectancy-violating outcomes under risk (Sutton and Barto, 1998). One solution to this, possible under normative Bayesian and model-based RL accounts, is to posit that agents construct internal models or beliefs about the environment, which may be sensitive to the nature of uncertainty (Behrens et al., 2007; Itti and Baldi,

Received Feb. 23, 2017; revised May 29, 2017; accepted June 7, 2017.

Author contributions: K.K. and M.H. designed research; K.K. performed research; K.K. analyzed data; K.K. and M.H. wrote the paper.

This research was supported by National Institutes of Health Grants R01 MH098023 and R01 DA043196 to M.H.

The authors declare no competing financial interests.

Correspondence should be addressed to Ming Hsu, 2220 Piedmont Avenue, Haas School of Business and Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley CA 94720. E-mail: mhsu@haas.berkeley.edu.

DOI:10.1523/JNEUROSCI.0535-17.2017

Copyright © 2017 the authors 0270-6474/17/376972-11\$15.00/0




Figure 1. Experimental paradigm. **a**, An exemplar urn. This urn contains two balls in yellow (the risky color) and two balls in either red or green (the ambiguous colors). Subjects do not know whether the urn contains two red balls, two green balls, or a red and a green ball. See Table 1 for urn contents in the actual experiments. **b**, Gambles. Subjects received \$10 if a resolution draw from the urn matched a predetermined winning color. The gamble is called ambiguous when the winning color is an ambiguous color, and risky otherwise. **c**, Belief updating and expectancy violation of observed draws. A ball is drawn from the urn, revealing its color, and is returned. Top, A draw of a red ball (ambiguous color) updates belief, because it demonstrates that at least one red ball is in the urn [$\Delta P(R) > 0$]. The draw also violates prior expectancy to a certain extent [$1 - P(R) > 0$]. Bottom, A draw of a yellow ball (risky color) does not update belief, because the urn is already known to contain two yellow balls [$\Delta P(Y) = 0$]. It is still associated with expectancy violation [$1 - P(Y) > 0$]. **d**, Value updating. Left, When the winning color is red, a red observed draw increases the expected value [$\$10 \cdot \Delta P > 0$]. Middle, When the winning color is green, a red draw decreases the expected value [$\$10 \cdot \Delta P < 0$]. Right, When the winning color is yellow, a red draw does not affect winning probability [$\$10 \cdot \Delta P(Y) = 0$]. **e**, Presence of belief updating, value updating, and expectancy violation. Expectancy violation does not guarantee belief updating or value updating.

2009; Nassar et al., 2010; Payzan-LeNestour and Bossaerts, 2011; Ma and Jazayeri, 2014). This allows agents to update beliefs based on signals under conditions of ambiguity, but not under risk, and use updated beliefs to update values.

There is growing evidence that the human brain constructs and makes use of beliefs about the external world (Behrens et al., 2007, 2008; Gläscher et al., 2010; Daw et al., 2011), and that it represents risk and ambiguity (Hsu et al., 2005; Huettel et al., 2006; Bach et al., 2011). However, we know little about how the reducibility of uncertainty is taken into account in neural processes involved in belief updating. In this study, we aim to go beyond the presence of beliefs or uncertainty representation and investigate how the reducibility of uncertainty affects processing of environmental signals for adaptive decision making.

We adapted the classic Ellsberg three-color urn problem. Subjects know the number of balls of one color in the urn, but not the number of balls of the other two colors. The monetary outcome is contingent on the color of a ball drawn from the urn upon the resolution of the gamble (Ellsberg, 1961). We introduced an environmental signal to this problem in the form of a draw, which was shown to subjects and returned to the urn before the resolution. From a Bayesian perspective, this draw may or may not reduce uncertainty in beliefs on urn contents and expected value, depending on its color and outcome contingency (Fig. 1), allowing trial-by-trial manipulation of the reducibility of uncertainty. Notably, since the draw's color is probabilistic, it violates expectancy to a certain extent, even under irreducible uncertainty (Itti and Baldi, 2009; O'Reilly et al., 2013).

We hypothesized that neural processes of uncertainty reduction would be dissociable from responses to expectancy violation. Specifically, we predicted that the lateral frontoparietal cortex and the medial prefrontal cortex (MPFC), which have been associated with updating in belief and value respectively (Gläscher et al., 2010; Daw et al., 2011; O'Reilly et al., 2013), would respond to uncertainty reduction as opposed to mere expectancy violation. We further tested whether the reducibility of uncertainty would modulate interactions among regions involved in updating.

Materials and Methods

We conducted two experiments, one a behavioral experiment and one involving fMRI, using almost identical paradigms. The behavioral experiment was conducted first to examine whether subjects were sensitive to the reducibility of uncertainty in value updating, and to what extent their behavior could be characterized by our Bayesian model. The fMRI experiment was conducted to look for neural correlates of belief updating, value updating, and expectancy violation. Below, the task paradigm is illustrated first, followed by the procedures of each experiment.

Task paradigm

Subjects were presented with a gamble, involving a number of balls in an urn. They knew the exact number of balls of one color (hereafter the risky color), but they did not know the exact number of balls of the other two colors (the ambiguous colors). For example, an urn contains four balls, two balls in yellow and two in either red or green; it could contain two red balls, one red ball and one green ball, or two green balls (Fig. 1a). The monetary outcome of the gamble was determined by a resolution draw from the urn; subjects could win \$10 if the ball drawn matches a predetermined winning color, and nothing otherwise (Fig. 1b). We called a

Table 1. The urn contents used in the experiments, and the quantitative measurement of belief updating, value updating, and expectancy violation (derived under the Bayesian model with binomial prior probability distribution over urn contents)

Urn contents		Ambiguous-color draw				Risky-color draw			
Balls in ambiguous colors	Balls in risky color	Belief updating	Value updating		Expectancy violation	Belief updating	Value updating		Expectancy violation
			Ambiguous gamble	Risky gamble			Ambiguous gamble	Risky gamble	
1	1	0.250	±2.50	0	0.750	0	0	0	0.500
1	2	0.167	±1.67	0	0.833	0	0	0	0.333
2	1	0.167	±1.67	0	0.667	0	0	0	0.667
1	3	0.125	±1.25	0	0.875	0	0	0	0.250
2	2	0.125	±1.25	0	0.750	0	0	0	0.500
3	1	0.125	±1.25	0	0.625	0	0	0	0.750

gamble ambiguous when its winning color was one of the ambiguous colors, and risky when it was the risky color.

We introduced environmental signals to this gamble in the form of a draw. Prior to the resolution of the gamble, a ball was drawn from the urn, so that the subject knew its color, and then returned to the urn. We postulated that this observed draw first updates the prior belief about the urn content (Fig. 1c, belief updating) and then the value of the gamble (Fig. 1d, value). Note that the draw's color was well specified and there was no perceptual ambiguity.

This paradigm specifies and manipulates the reducibility of uncertainty in beliefs and value as follows. First, because the subject does not know the composition within ambiguous-color balls in the urn, a draw of an ambiguous-color ball, but not a draw of a risky-color ball, should update the subject's belief. In our exemplar urn, a red draw updates belief because it demonstrates that the urn holds at least one red ball, increasing probability of a future draw in red [$\Delta P(R) > 0$; Fig. 1c]. On the other hand, a yellow draw does not carry any information, because it is already known that the urn contains two yellow balls [$\Delta P(Y) = 0$].

Second, value should be updated as a consequence of belief updating only in ambiguous gambles, but not in risky gambles. This is because the chance to win \$10 is not perfectly specified when the winning color is an ambiguous color. In our exemplar urn, a red draw increases the probability of a red resolution draw and decreases that of a green draw (Fig. 1d). As a consequence, the value of the gamble increases if the winning color is red [$\$10 \cdot \Delta P(R) > 0$] and decreases if the winning color is green [$\$10 \cdot \Delta P(G) < 0$]. On the other hand, if the winning color is yellow, a red draw does not update its value, because probability of a yellow draw is unaffected [$\$10 \cdot \Delta P(Y) = 0$].

Therefore, if subjects rationally combine the prior knowledge about uncertainty with the color of the observed draw, they would update belief only after ambiguous-color draws, and update value only in ambiguous gambles (Fig. 1e). Such sensitivity would not be observed if updating is primarily driven by expectancy violation; since the draw's color is unpredictable, any draw in any gamble is associated with some level of expectancy violation, measured as $1 - P(\text{draw})$ [since $P(\text{draw}) < 1$, $1 - P(\text{draw}) > 0$ for any draw; Figure 1c]. To decouple updating from expectancy violation more clearly, we manipulated the urn composition across trials (Table 1). For instance, increasing yellow balls in our exemplar urn would increase expectancy violation of a red draw, but decrease the magnitudes of belief updating and value updating the red draw causes. The manipulation of the urn composition enabled us to look for neural correlates of belief updating and value updating while statistically controlling for expectancy violation, and vice versa, in fMRI analysis.

Behavioral experiment

Subjects. Ten undergraduate students in University of California, Berkeley (six women) participated. They provided written informed consent. All procedures were approved by the University of California, Berkeley Committee for the Protection of Human Subjects. The experiment was conducted individually in a self-paced manner in isolated cubicles. The experiment program was written on Matlab (Mathworks, RRID: SCR_001622) and Psychtoolbox (Brainard, 1997; Pelli, 1997; RRID: SCR_002881) and run on a laptop.

Procedure. Subjects were presented with gambles, each of which consisted of a winning color and an urn containing a number of balls in red,

green, or yellow. One of the gambles was randomly selected and resolved at the end of the experiment; a ball was randomly drawn from the urn, and subjects received \$10 only if it matched the gamble's winning color (in addition to the baseline payment for task completion). Subjective values of these gambles, both predraw and postdraw, were elicited as willingness to sell (WTS), i.e., the amount of money subjects were willing to give up for the opportunity to gamble. A standard Becker–DeGroot–Marschak (BDM) bidding procedure was used (Becker and Brownson, 1964); the gamble's price was randomly determined at the end of the experiment (uniform distribution between \$0 and \$10), and subjects sold the gamble for its price only if it exceeded their WTS. In total, 18 gambles were presented in a randomized order [6 urn contents (Table 1) \times 3 winning colors].

Each trial started with the presentation of the urn content and the winning color. Subjects were informed of the number of balls of one color (the risky color) and the total number of balls of the other two colors (the ambiguous colors), but not of the exact composition within the latter. Each ball in the risky color was visually represented as a full circle, and each ball in the ambiguous colors as a pair of half circles. The winning color was shown above the urn contents. After subjects indicated predraw WTS, observed draws in red, green, and yellow were presented in a randomized order, after each of which, postdraw WTS was indicated. Thus, four WTSs in total were obtained in each gamble. Upon the resolution of the gamble, the experiment program randomly determined whether subjects observed the draw or not (50%), and which color the observed draw had (probability following the urn composition).

Data analysis. To examine value updating in a model-free manner, trial-wise difference between predraw and postdraw subjective values (WTSs) were calculated and categorized according to the normative prediction of its valence. To more quantitatively characterize subjective values, predictions from our quantitative Bayesian model (see below) were fitted to WTSs in mixed-effect modeling implemented on R software (RRID:SCR_001905) and lmer package, with subjects as a random effect. To test one-sided deviation, the fixed-effect constant term (intercept) was compared against zero.

Bayesian modeling. Our quantitative Bayesian model consists of two stages, belief and valuation. The belief stage concerns probability distribution on a future draw's color. Before the observed draw, subjects were informed of the total number of balls in the two ambiguous colors n_a and the number of balls in the risky color n_r , but the number of balls in each ambiguous color n_{a1} or n_{a2} was unknown. Thus, while the probability of a future draw in the risky color $P_{\text{pre}}(r)$ can be easily uniquely specified as $n_r/(n_r + n_a)$, the probabilities of a future draw in the ambiguous colors $P_{\text{pre}}(a1)$ and $P_{\text{pre}}(a2)$ cannot. To generate their unique point estimates, all possible urn contents were considered, weighted according to their probability, and averaged. Specifically, we assumed that prior probability over urn contents followed a binomial distribution, i.e., the number of balls in one ambiguous color n_{a1} followed $P_{\text{pre}}(n_{a1}) = \frac{1}{2^{n_a}} \binom{n_a}{n_{a1}}$. The probability of a future draw in one ambiguous color can then be obtained as: $P_{\text{pre}}(a1) = \sum_{n_{a1}=0}^{n_a} P_{\text{pre}}(n_{a1}) \cdot n_{a1}/(n_r + n_a)$.

After the observed draw, probability over urn contents was updated according to Bayes' rule. As we adopted binomial distribution in

$P_{\text{pre}}(n_{a1})$, the posterior probability over urn contents $P_{\text{post}}(n_{a1})$ also follows binomial distribution (proof omitted). Namely, if the observed draw is in $a1$, then $P_{\text{post}}(n_{a1}) = \frac{1}{2^{n_{a1}-1}} \binom{n_a-1}{n_{a1}-1}$; if the observed draw is in $a2$, then $P_{\text{post}}(n_{a1}) = \frac{1}{2^{n_{a1}-1}} \binom{n_a-1}{n_{a1}}$; if the observed draw is in r , no belief updating occurs [$P_{\text{post}}(n_{a1}) = P_{\text{pre}}(n_{a1})$].

In the valuation stage, predraw and postdraw values of gambles are calculated as expected outcome, i.e., $\$10 \times$ probability of a draw in the winning color.

It is straightforward to prove that this modeling is mathematically equivalent to a more heuristic account, which only considers “effective” urn content (proof omitted). In this account, each ambiguous ball is treated as a pair of half (0.5) balls in the ambiguous colors. When an ambiguous-color draw is observed, one of such pairs is replaced with a full ball in the draw’s color. Critically, agents taking this heuristic strategy are still sensitive to the empirically manipulated reducibility of uncertainty. Due to this mathematical equivalence, the current study does not directly test whether updating processes are fully Bayesian or not; rather, our interest lies in whether updating processes take into account the reducibility of uncertainty, which is an important feature of Bayesian theories.

fMRI experiment

Subjects. Twenty subjects (mean age, 21.7 years; 11 women) participated after being screened for standard MRI contraindications. They provided written informed consent. All procedures were approved by University of California, Berkeley Committee for the Protection of Human Subjects. One subject declined to participate after the task instructions but before the scanning, and two subjects were discarded from analysis due to unsatisfactory performance in auxiliary tasks (see below), resulting in data from 17 subjects analyzed. During scanning, the experiment program was run on Matlab and Psychtoolbox, with which subjects interacted via an MRI-compatible button box.

Main task procedure. During scanning, subjects observed gambles in a randomized order one of which was randomly selected and resolved at the end of the experiment. Thirty gambles were presented in each of the three echo-planar imaging (EPI) runs (90 in total), and the winning color was changed across runs (remained the same within each run). Six urn contents (Table 1) were presented 15 times each, 5 times for risky gambles and 10 for ambiguous gambles ($6 \times 15 = 90$). Frequency of the observed draw’s color across gambles approximately followed the urn composition.

Each trial started with the fixation cross in the winning color (2 s), followed by the urn content presentation (visual representation similar to the behavioral experiment). The urn content was presented for 5–12 s (after a variable delay of 4–6 s, the urn lid opened in a 0.5 s animation, on which subjects were asked to press a button within 5 s; upon their button press, the balls moved into the urn in another 0.5 s animation; this process was introduced to keep subjects alert). After a variable interval (3–6 s), a gray ball moved out of the urn in a 0.5 s animation, and revealed its color after another variable interval (1–3 s). After 3 s, the drawn ball was returned to the urn in a 0.5 s animation, followed by a variable intertrial interval (2.5–4.5 s).

Auxiliary tasks. To verify each subject’s engagement and understanding throughout the main observation task, we asked them to respond to three types of auxiliary tasks. They were presented immediately after randomly selected 27 gambles (nine for each task type). In the memory task, subjects were asked to choose the correct description of the previous gamble (the winning color, the urn content, and the observed draw) from two options. In the value-updating judgment task, subjects were asked to indicate whether it is (1) less likely, (2) equally likely, or (3) more likely to win the gamble after the draw. In the surprise rating task, subjects were asked to rate their surprise of the observed draw on a three-point scale. Since the surprise rating task was purely subjective, we used the memory task and the value-updating judgment task to test each subject’s engagement and understanding. Subjects each received up to $\$10$ based on their performance in these two tasks. Two subjects were excluded from the subsequent analyses because of their unsatisfactory performance (>2 wrong responses or >2 trials without responses within 10 s in either

task). The remaining subjects were able to classify the valence of value updating consistently with the Bayesian prediction, which was in line with the behavioral experiment’s results.

fMRI data acquisition. MR images were acquired by a 3T Siemens Trio scanner and a 12-channel head coil. Functional images were obtained using T2*-weighted gradient-echo EPI pulse sequence (TR = 2000 ms; TE = 30 ms; voxel size, $3 \times 3 \times 3$ mm; interslice gap, 0.3 mm; in-plane resolution, 64×64 ; 32 oblique axial slices). Slices were tilted by 30° from the anterior commissure–posterior commissure line to alleviate signal dropout from the orbitofrontal cortex (Weiskopf et al., 2006). T1-weighted structural images ($1 \times 1 \times 1$ mm) were also obtained using a magnetization-prepared rapid-acquisition gradient-echo pulse sequence.

Preprocessing. Preprocessing was conducted using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm/>, RRID:SCR_007037). Structural images were segmented into gray matter and white matter. Functional images were motion corrected (aligned to the mean image), slice-time corrected, coregistered with segmented structural images, normalized to the MNI EPI template, and smoothed with a Gaussian kernel of 8 mm FWHM.

Whole-brain univariate analysis. Whole-brain analysis used general linear modeling (GLM) of BOLD time series and group-level random-effect models on SPM8. To quantitatively relate BOLD signals to belief updating, value updating, and expectancy violation, these variables were parametrically defined under the Bayesian model: belief updating as the absolute difference between predraw and postdraw probability of ambiguous colors, value updating as the signed difference between predraw and postdraw expected value of gambles, and expectancy violation as $1 -$ predraw probability of the observed draw. See Table 1 for their variations as functions of urn contents. We noted that, although our behavioral results may imply probability underweighting under ambiguity, adopting such a model has minimal effects on parametric modulators of belief updating and value updating ($r = 0.97$ and 0.99 , respectively).

These variables were included in GLMs as parametric modulators of a regressor at the observed draws. To adjust for the correlation between belief updating and expectancy violation, we included both parametric modulators in the GLMs so that their coefficients only captured variance uniquely explained by each of them, not the shared variance (Mumford et al., 2015). Since SPM8 orthogonalizes the second parametric modulator against the first by default, we implemented two GLMs: one in which the first parametric modulator was expectancy violation and the second was belief updating (GLM 1), and the other one in which the order was reversed (GLM 2). GLM 1 was used to look for neural correlates of belief updating adjusted for expectancy violation, and GLM 2 was used to look for neural correlates of expectancy violation adjusted for belief updating. To illustrate how results could have been affected if the shared variance were not removed and adjusted for, we also reported coefficient estimates for unorthogonalized parametric modulators, i.e., expectancy violation in GLM 1 and belief updating in GLM 2. GLM 3 included value updating as a sole parametric modulator.

All GLMs also included regressors that modeled events of gamble presentation, button press, question presentation, and question response. All event-related regressors were convolved with the SPM’s double- γ canonical hemodynamic response function. Additionally, they included six movement parameters estimated in the motion-correction procedure, 128 s high-pass filtering, and the AR(1) model of serial autocorrelation. Coefficient estimates of the parametric modulators were then entered into group-level analysis. For clusters defined by voxel-level threshold $p < 0.001$ (uncorrected), cluster-level p values with whole-brain correction for family-wise error (FWE) were calculated using nonparametric permutation in SnPM13 package (Nichols and Holmes, 2002; Hayasaka and Nichols, 2003; Woo et al., 2014; RRID:SCR_002092).

Regions-of-interest analysis. Activations in regions of interest (ROIs) were examined using the Marsbar package (Brett et al., 2002; RRID:SCR_009605) in the following steps. First, ROIs were determined based on the group-level whole-brain analysis at the cluster-level threshold of $k > 20$. Second, subject-specific ROIs were defined from group-level activation maps based on the other 16 subjects’ data (leave-one-subject-out; Boorman et al., 2013; Hunt et al., 2014). We used clusters that

Table 2. Clusters associated with belief updating (adjusted for expectancy violation), value updating, and expectancy violation (adjusted for belief updating)

Region	# Voxels	Cluster-level p (whole-brain FWE corrected)	Peaks			$T_{(16)}$	Voxel-level p (uncorrected)	Mean (SD) of # voxels in leave- one-subject-out
			x	y	z			
Belief updating								
Right middle frontal gyrus/superior frontal sulcus	61	0.0568	27	2	55	5.77	0.000014	50 (29)
			27	14	46	4.30	0.000275	
Precuneus	96	0.0280	6	−61	31	5.61	0.000020	82 (26)
			−6	−61	40	4.81	0.000096	
			−6	−58	28	4.40	0.000224	
Right IPS	146	0.0128	30	−40	49	5.54	0.000022	117 (45)
			45	−31	46	5.06	0.000058	
			27	−52	46	4.64	0.000136	
Left IPS	51	0.0724	−24	−43	46	5.18	0.000046	40 (33)
			−42	−37	40	4.66	0.000131	
Left middle frontal gyrus/superior frontal sulcus	39	0.1000	−30	11	58	4.67	0.000128	40 (42)
Value updating								
Right VMPFC	39	0.0966	15	29	−14	6.36	0.000005	34 (17)
			15	38	−11	4.68	0.000125	
			24	26	−11	4.13	0.000393	
Left superior temporal gyrus	43	0.0868	−48	−7	−2	5.75	0.000015	
MPFC	210	0.0118	−15	53	7	5.66	0.000018	168 (72)
			15	53	4	4.77	0.000104	
			−12	41	4	4.66	0.000131	
Cingulate cortex	107	0.0302	−9	−22	43	5.55	0.000022	104 (88)
			−18	−22	40	5.16	0.000047	
			−6	−10	37	4.47	0.000193	
Subgenual area	31	0.1194	9	26	7	5.53	0.000023	26 (9)
Cingulate cortex	29	0.1284	−6	−1	52	4.40	0.000224	N/A
			3	2	46	4.20	0.000339	
Left inferior parietal lobule	21	0.1690	−48	−37	25	4.37	0.000238	38 (10)
Expectancy violation								
Right insula	45	0.0660	30	20	−8	6.22	0.000003	42 (12)
Left insula	27	0.1060	−30	26	−2	5.05	0.000059	23 (17)
			−42	20	−2	4.13	0.000393	

Clusters were formed at voxelwise threshold $p < 0.001$. All clusters > 20 voxels are reported here. VMPFC, ventromedial prefrontal cortex.

survived uncorrected voxel-level $p < 0.001$, $k \geq 10$, with the local maxima located within 16 mm from all-subject group-level peaks (one cingulate value-updating ROI was discarded because it was not robustly identified in some iterations at this threshold). Third, mean BOLD time series from each ROI was extracted from the hold-out subject's data, to which GLMs 1, 2, and 3 were fitted. The parametric modulators' coefficient estimates were normalized according to the baseline of time series, similarly to conventional calculation of percentage signal change (although our coefficients were derived from parametric modulators and thus could not be interpreted as percentage signal change per se). Fourth, estimates were entered into mixed-effect modeling with subjects as a random effect, conducted on R and lmer package.

Note that, although the circularity problem in ROI definition (Kriegeskorte et al., 2009) is slightly alleviated by the leave-one-subject-out procedure, it is not eliminated. This is because the sets of ROIs were determined based on data from all subjects, including the one held out. As a consequence, for belief-updating ROIs, coefficient estimates for belief updating (GLM 1) may have a slightly positive bias, and estimates for expectancy violation (GLM 2) may be negatively biased; vice versa for expectancy-violation ROIs (since value updating is orthogonal to belief updating and expectancy violation by design, there is no bias in any coefficient estimates from value-updating ROIs, as well as in coefficient estimates for value updating from any ROIs).

Also note that our statistical inference in ROI analysis did not compare coefficients of different parametric modulators (belief updating, value updating, and expectancy violation); it only concerned whether each coefficient was different from zero, not whether one coefficient was higher than others. The parametric modulators are not in the same unit, and their coefficient estimates are not directly comparable.

Dynamic causal modeling. Dynamic causal modeling (DCM) analysis was conducted using the DCM module in SPM8. DCM is a generative

model of BOLD time series from multiple ROIs, and includes three types of factors: direct regional input, stationary inter-regional connections, and, importantly for our purpose, temporal modulations in inter-regional connections (Friston et al., 2003). DCM is agnostic about whether connections are monosynaptic.

To test modulation of inter-regional connections based on the type of gambles (ambiguous or risky), we constructed and compared three families of DCMs; Family 1 allowed for modulations in connections from belief-updating ROIs to value-updating ROIs, Family 2 allowed for modulations in connections from expectancy-violation ROIs to value-updating ROIs, and Family 3 did not allow for any modulations. We adopted this familywise approach because, while we were interested in testing existence of modulations among the sets of ROIs, we were not interested in discriminating among contributions of specific ROIs. We aimed in particular to allow the possibility that ROIs exhibiting modulation are heterogeneous across individuals (e.g., handedness might affect laterality). Modulations were implemented as differential strength of connections during ambiguous and risky gambles (boxcar functions from gambles' presentations to the trials' termination).

We constructed DCMs with eight ROIs: four belief-updating ROIs (bilateral frontal and parietal cortex), two value-updating ROIs (MPFC and right ventromedial PFC [VMPFC]), and two expectancy-violation ROIs (bilateral anterior insula). Instantiating all possible sets of modulated connections from belief-updating ROIs to value-updating ROIs, Family 1 thus contained $2^8 - 1 = 255$ models. Similarly, instantiating all possible sets of modulated connections from two expectancy-violation ROIs to two value-updating ROIs, Family 2 contained $2^4 - 1 = 15$ models. Family 3 contained 1 model with no modulations. We chose these ROIs primarily because they were identifiable for every subject (see below). Due to large model space, it is computationally prohibitive to explore alternative network specifications (e.g., including additional ROIs). Although it is possible that the DCM results depend on this selec-




Figure 2. Behavioral results. **a**, Experiment procedure. First, the urn content and the winning color were presented, and pre-draw subjective value of gambles were elicited. Next, the observed draw was presented, followed by elicitation of post-draw value. Post-draw valuation was repeated for all of the three draw colors. The urn content and the winning color content of the urn and the color of the winning ball were manipulated across trials. **b**, Histogram of value updating. Trials were classified based on the predicted valence (Fig. 1e). Proportions of the observed valence agreed with the predictions (χ^2 test of independence, $p < 0.05$). **c**, Quantitative modeling of subjective values. Observation was compared against prediction of our Bayesian model. Error bars are SEMs, but not visible in some cases due to their small sizes. Left, Reported pre-draw values were successfully predicted in both ambiguous (top) and risky (bottom) gambles (log-likelihood ratio test, $p < 0.05$). Middle, Reported post-draw values were successfully predicted in both ambiguous (top) and risky (bottom) gambles, after both ambiguous-color draws (blue open dots) and risky-color draws (red closed dots; $p < 0.05$). Right, Nonzero updating was only predicted in ambiguous gambles after ambiguous-color draws (top, blue open dots), and linear prediction was successful ($p < 0.05$). Although ambiguous-color draws slightly decreased values of risky gambles (bottom, blue open dots; $p < 0.05$), its effect size was negligible compared with ambiguous gambles. Risky-color draws did not update values both in ambiguous and risky gambles (red closed dots, $p > 0.10$).

tion of ROIs, type-I error has not been affected because ROI selection was not conditional on DCM results.

ROIs were defined in a subject-specific manner. From subject-wise activation maps (GLMs 1–3), ROIs that survived uncorrected voxelwise threshold $p < 0.20$ with local maxima located within 16 mm from all-subject group-level peaks were selected (Smith et al., 2006). Three subjects were discarded from DCM analysis because some ROIs could not be identified (one subject for the left frontal belief-updating cluster, one for the MPFC value-updating cluster, and one for left insula expectancy-violation cluster). Next, we extracted the principal eigenvariate of BOLD time series from 4-mm-radius spheres centered on the local maxima.

In addition to modulated connections, all DCMs included the same regional inputs and stationary inter-regional connections. Regional inputs modeled the events of the observed draws. To explain away intraregional computational processes captured in univariate analysis, we included parametric modulators of both belief updating and expectancy violation to inputs to belief-updating and expectancy-violation ROIs (as in GLM 1 and 2, with which these ROIs were originally defined), and value updating to value-updating ROIs (as in GLM 3). Stationary inter-regional connections modeled bidirectional influence between ROIs from the different categories (belief updating, value updating, and expectancy violation), but not from the same category.

Group-level, random-effect, family-level inference was conducted to compare the three families of models. The family comparison procedure takes into account both goodness of fit and model complexity, aggregates performance across all models in each family, and calculates exceedance probability, i.e., probability in which each family was better than the other two families (Stephan et al., 2009; Penny et al., 2010).

Results

Behavioral sensitivity to the nature of uncertainty

First, we tested the extent to which value updating is sensitive to the nature of uncertainty at the behavioral level. Subjective values of gambles were elicited via a BDM bidding procedure before and after the draw (Fig. 2a). If subjects are sensitive to the nature of uncertainty as normatively predicted, they would update value only when they observe ambiguous-color draws in ambiguous gambles. On the other hand, if the updating process is driven by expectancy violation, value would be also updated by risky-color draws or in risky gambles. We tested these predictions by classifying trials according to whether positive, negative, or zero value

updating was normatively predicted given the draw color (Fig. 2b). Distribution of observed value updating indeed varied consistently with the normative prediction (χ^2 test of independence, $\chi^2_{(4)} = 1493.432$, $p < 10^{-10}$).

Quantitative modeling of updating

To quantitatively relate updating processes to BOLD responses in the following fMRI analysis, we next sought to provide a quantitative model of updating. We postulated that agents first construct and update belief about probability of a future draw, and then use it to determine the expected value of the gamble ($\$10 \times$ the probability of winning). We further assumed that, in belief, probability over all possible urn contents is considered and updated according to Bayes' rule (see Materials and Methods for modeling details; note that this formalization instantiates so-called second-order probability, one of the most widely used approaches to ambiguity in decision theory; Camerer and Weber, 1992; Klibanoff et al., 2005; Nau, 2006; Ergin and Gul, 2009; Seo, 2009). More specifically, prior probability over urn contents is assumed to follow a binomial distribution.

We found that this model explained subjective values of gambles well (Fig. 2c). Pre-draw and post-draw subjective values were consistent with model predictions in both ambiguous and risky gambles (log-likelihood ratio test, pre-draw ambiguous gambles, $p = 2.84 \times 10^{-5}$; pre-draw ambiguous gambles, $p = 7.07 \times 10^{-9}$; post-draw ambiguous gambles, $p = 1.12 \times 10^{-5}$; pre-draw ambiguous gambles, $p = 5.72 \times 10^{-8}$; Fig. 2c, left and middle). More importantly, our model successfully predicted significant value updating after ambiguous-color draws in ambiguous gambles ($p = 8.05 \times 10^{-5}$; Fig. 2c, top right), as well as negligible value updating in the other cases (+\$0.068 after risky-color draws in ambiguous gambles, $p = 0.300$; -\$0.228 after risky-color draws in risky gambles, $p = 0.182$; but -\$0.143 after ambiguous-color draws in risky gambles, $p = 0.004$; Fig. 2c, right). We also noted that subjects exhibited ambiguity aversion (Ellsberg, 1961); while neither overvaluation nor undervaluation was observed in risky gambles (-\$0.063 before draws, $p = 0.714$; -\$0.234 after draws, $p = 0.272$ respectively), ambiguous gambles exhibited small un-

dervaluation ($-\$0.252$ before draws, $p = 0.065$; $-\$0.292$ after ambiguous-color draws, $p = 0.001$; but $-\$0.184$ after risky-color draws, $p > 0.239$). Overall, even though our model did not aim to account for ambiguity aversion, it is a successful first-order approximation of updating.

Neural correlates of belief updating

After establishing each subject's behavioral sensitivity to the reducibility of uncertainty, we conducted an fMRI experiment to examine how the brain processes environmental signals according to uncertainty. Specifically, we looked for neural correlates of belief updating and value updating and tested whether they were dissociable from expectancy violation. During scanning, subjects observed a series of gambles (Fig. 3*a*); each trial started with presentation of the gamble's winning color, followed by the urn content and the observed draw. The gambles were resolved only after the scanning. This observational task was adopted to isolate processes related to updating as opposed to choices. To ensure that subjects paid enough attention to the task, we used auxiliary tasks to elicit subjective assessment of the directionality of value updating and memory on the urn contents (see Materials and Methods).

We first looked for brain regions where activation was correlated with belief updating. We quantified belief updating as the absolute difference between pre-draw and post-draw probability of ambiguous-color draws under our Bayesian model (Table 1). Even though our paradigm quantitatively dissociates belief updating from expectancy violation (measured as $1 - \text{prior probability of the draw}$), the activation map of belief updating overlapped that of expectancy violation (Fig. 3*b*). This is because the correlation between these two variables was still not negligible ($r = 0.70$; Fig. 3*c*, left). To adjust for the correlation, we included both trial-wise belief updating and expectancy violation as parametric modulators in a single GLM and looked for regions where a significant amount of variance could be explained uniquely by belief updating (Mumford et al., 2015). We found bilateral clusters in the posterior middle frontal gyrus and the superior frontal sulcus, bilateral clusters in the intraparietal sulcus (IPS), and a cluster in the precuneus (Fig. 3*d*; cluster-forming voxel-level threshold $p < 0.001$, uncorrected, and cluster-size threshold $k > 20$; see Table 2 for cluster-level p values corrected for whole-brain FWE). The clusters in the frontal cortex may corre-




Figure 3. fMRI results. *a*, Experiment procedure. The winning color was first presented in the fixation cross, followed by the urn content and then the observed draw. We analyzed BOLD signals time-locked to the observed draw. The content of the urn and the color of the drawn ball were manipulated across trials while the winning color was manipulated across the scan runs. *b*, Activation maps of belief updating and expectancy violation without adjustment for their correlation. They greatly overlap because of correlation between the variables. Shown are clusters that survived voxel-level threshold $p < 0.001$ (uncorrected) and cluster-size threshold $k > 20$. *c*, Correlation between belief updating, value updating, and expectancy violation, in our experiment settings (Table 1). Left, Even though belief updating does not coincide with expectancy violation, the correlation between belief updating and expectancy violation is not negligible ($r = 0.70$). Right, As value updating could be positive or negative, it is orthogonal to both belief updating (top) and expectancy violation (bottom) by design. *d*, Activation maps of belief updating adjusted for expectancy violation (top), value updating (middle), and expectancy violation adjusted for belief updating (bottom). Belief updating was correlated with activation in the premotor/FEF, IPS, and precuneus; value updating with the MPFC and the cingulate cortex; expectancy violation with the AI. See Table 2 for cluster-level whole-brain FWE-corrected p values. *e*, ROI analysis. Top, Coefficient estimates for belief updating. Belief updating was not associated with value-updating ROIs ($p > 0.10$) or expectancy-violation ROIs ($p > 0.10$). Middle, Coefficient estimates for value updating. Value updating was not associated with belief-updating ROIs ($p > 0.05$) or expectancy-violation ROIs ($p > 0.10$). Bottom, Coefficient estimates for expectancy violation. Expectancy violation was not associated with belief-updating ROIs or value-updating ROIs ($p > 0.10$). Note that the coefficient estimates for belief updating (top), value updating (middle), and expectancy violation (bottom) are not directly comparable due to differences in scale of regressors. Error bars: SEMs. $*p < 0.05$.

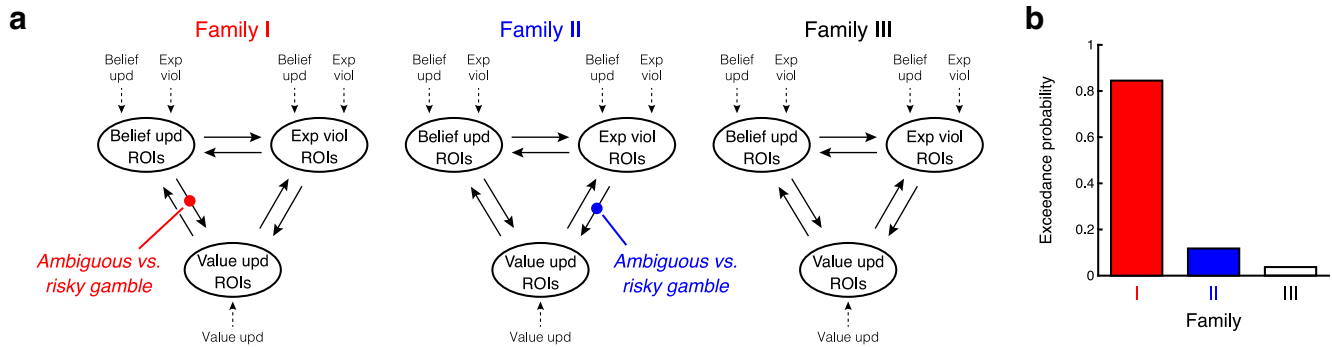


Figure 4. Dynamic causal modeling. **a**, Families of models. Dotted arrows indicate regional inputs, solid arrows indicate inter-regional connections, and dots on solid arrows indicate modulation according to the type of gambles. All models contained four belief-updating ROIs (bilateral frontal and parietal), two value-updating ROIs (MPFC and right ventromedial PFC [VMPFC]), and two expectancy-violation ROIs (bilateral insula). Families I and II instantiated all possible sets of modulation in connections from belief-updating ROIs or expectancy-violation ROIs to value-updating ROIs, respectively. Family III contained no modulation in connections. **b**, Exceedance probability of the families, i.e., the probability of each family performing better than the other two families. Family I was supported by data.

spond to premotor region or frontal eye field (FEF; Vernet et al., 2014).

Neural correlates of value updating

We then examined neural correlates of value updating. While belief can be updated both in ambiguous and risky gambles, value should be updated only in ambiguous gambles. Given this difference, we expected that neural correlates of value updating are anatomically distinct from belief updating. We looked for brain regions in which activation was correlated with value updating (quantified as the signed difference between predraw and postdraw expected values; Table 1; this parametric modulator is orthogonal to belief updating and expectancy violation by design; Fig. 3c, right). We found clusters in the VMPFC, the anterior and middle cingulate, and the left superior temporal gyrus (Fig. 3d; Table 2). These clusters did not overlap with the belief-updating clusters.

Neural correlates of expectancy violation

Next, we tested whether these updating regions responded to expectancy violation. Responses to expectancy violation, salience, or surprise has been long studied in cognitive neuroscience (Sokolov, 1963; Courchesne et al., 1975; Squires et al., 1975). However, those studies have been motivated by the assumption that surprising signals tend to be relevant for agents, and dissociation between expectancy violation and uncertainty reduction has not been well studied. We found that activation in bilateral AI was correlated with expectancy violation, even when the correlation with belief updating was adjusted for (Fig. 3d; Table 2). Importantly, these clusters did not overlap with the belief-updating clusters or the value-updating clusters. This localization of expectancy violation is consistent with previous reports that the AI responds to salient events in various domains (Corbetta et al., 2008; Singer et al., 2009; Menon and Uddin, 2010).

ROI analysis of separable neural correlates

To further illustrate the dissociation among neural correlates of belief updating, value updating, and expectancy violation, we conducted an ROI analysis, where ROIs were defined in a leave-one-subject-out fashion (see Materials and Methods; see Table 2 for ROI sizes in this analysis). We found that BOLD activation in belief-updating ROIs was correlated with belief updating as expected ($p < 10^{-4}$), but not with value updating or expectancy violation ($p = 0.096$ and 0.559 , respectively; Fig. 3e). Similarly,

activation in value-updating ROIs was correlated with value updating ($p < 10^{-4}$), but not with belief updating or expectancy violation ($p = 0.281$ and 0.874 , respectively). More critically, activation in expectancy-violation ROIs was correlated with expectancy violation ($p = 0.004$), but not with belief updating or value updating ($p = 0.953$ and 0.875 , respectively). These results show that neural processes of uncertainty reduction are anatomically dissociable from expectancy violation.

Interaction between updating regions

Last, we explored how these regions interact to drive appropriate value updating. Based on our results, we made two predictions about inter-regional interactions. First, we hypothesized that connections from belief-updating regions to value-updating regions would be modulated by the type of gamble. Since belief updating should contribute to value computation only in ambiguous gambles, inter-regional connections would be temporarily enhanced under ambiguous gambles (or temporarily weakened under risky gambles). Second, we hypothesized that connections from expectancy-violation regions to value-updating regions would not show such modulation as much, since expectancy violation does not drive value updating regardless of the type of gamble, both theoretically and behaviorally. To test these predictions, we conducted DCM analysis. It is appropriate for our purpose because it seeks to explain BOLD time series from more than 2 ROIs simultaneously and can include directional connections modulated by experimental manipulations (Friston et al., 2003).

We compared three scenarios: (1) belief-updating ROIs contribute to value updating, (2) expectancy-violation ROIs contribute to value updating, and (3) neither of them contributes to value updating (Fig. 4a; see Materials and Methods). To implement the first scenario, we constructed a family of models that instantiated every possible set of modulated connections from belief-updating ROIs to value-updating ROIs (Family I). Similarly, to implement the second scenario, another family of models instantiated every possible set of modulated connections from expectancy-violation ROIs to value-updating ROIs (Family II). The last scenario was implemented as a single model with no modulation (Family III). We included four belief-updating ROIs from the bilateral frontal (premotor/FEF) and parietal (IPS) regions, two value-updating ROIs in the MPFC and right VMPFC, and two expectancy-violation ROIs from the bilateral AI in DCM.

If neural processes of belief updating contribute to value updating in the way we hypothesized, our fMRI data should be best

explained by Family I. This prediction was supported by the results of a Bayesian model selection procedure (Stephan et al., 2009; Penny et al., 2010; Fig. 4*b*). Probability that Family I outperformed both Families II and III (“exceedance probability”) was >80%, while that of Family II was ~10% and Family III <5%. Among the models in Family I, the most supported model contained two modulated connections: from the right premotor/FEF to the MPFC and from the left IPS to the MPFC. These results suggest that uncertainty reducibility is reflected in interaction between belief-updating regions, rather than expectancy-violation regions, to value-updating regions.

Discussion

For adaptive uncertainty reduction, it is critical to understand whether current uncertainty is reducible, and by which signals it can be reduced. Specifically, we should not rely solely on signals’ expectancy violation (Itti and Baldi, 2009; O’Reilly et al., 2013). Distinction between uncertainty reduction and expectancy violation is not clear in some traditionally prevalent frameworks, such as RL and Pearce-Hall (Rescorla and Wagner, 1972; Pearce and Hall, 1980; Sutton and Barto, 1998; Pearce and Bouton, 2001; Roesch et al., 2012). In both, learning is driven by the degree to which prior expectancy about an event (e.g., the timing and amount of reward delivery) is violated. These theories do not explicitly state how agents can successfully ignore surprising, yet irrelevant, signals widespread in natural environments.

In this study, we showed that the human brain is sensitive to the nature of uncertainty by demonstrating dissociation between uncertainty reduction and expectancy violation at the behavioral and neural levels. This is relevant to two lines of decision-making studies. First, model-based RL and Bayesian theories postulate that agents construct and update beliefs about the environment. Such agents may possess representation of uncertainty, by which uncertainty reduction could be decoupled from expectancy violation (Behrens et al., 2007; Itti and Baldi, 2009; Nassar et al., 2010; Payzan-LeNestour and Bossaerts, 2011; O’Reilly et al., 2013; Ma and Jazayeri, 2014). Second, decision theory has long emphasized a distinction between reducible and irreducible uncertainty, often referred to as ambiguity and risk respectively (Keynes, 1921; Knight, 1921; Ellsberg, 1961; Camerer and Weber, 1992). Past studies have mainly investigated influence of these types of uncertainty on value-based choices and its neural basis (Hsu et al., 2005; Huettel et al., 2006; Bach et al., 2011). Our findings complement them and show that they are also important determinants of neural processing of environmental signals as well as signals’ behavioral consequences.

Behaviorally, we found evidence that value updating is sensitive to the nature of uncertainty. Value was updated only under reducible uncertainty, regardless of the extent to which expected signals were violated (Fig. 2*b*), and could be quantitatively characterized by a Bayesian model (Fig. 2*c*). Representation of uncertainty has been shown in Bayesian sensory and sensorimotor literature (Ernst and Banks, 2002; Körding and Wolpert, 2004; Pouget et al., 2013; Ma and Jazayeri, 2014), and our results suggest that such representation also plays a key role in the processes that link perception of signals to valuation. Note that, given the relatively simple nature of our task, it is possible that subjects used some heuristics instead of full-Bayesian computation (see Materials and Methods). Even if the updating processes may not necessarily be Bayesian, subjects appear capable of incorporating their knowledge about the reducibility of uncertainty into updating processes.

At the neural level, we found that processes associated with belief updating and expectancy violation were anatomically separable, which have been frequently confounded in past studies. Indeed, to our knowledge, there exist only two studies that explicitly decoupled these two (O’Reilly et al., 2013; Schwartenbeck et al., 2016). Despite differences in task design, these studies also found anatomical dissociation between belief updating and expectancy violation, and in particular, representation of belief updating in lateral frontoparietal regions. This also accords well with Gläscher et al. (2010), who used a Markov decision task to capture state prediction error (SPE) that updated beliefs on state-action-state transition probabilities. Even though SPE was formally equivalent to expectancy violation in their task, our results make it unlikely that their results reflected expectancy violation alone.

However, these studies differ in the precise localization of belief updating within frontoparietal regions, as well as recruitment of other regions. They also mapped expectancy violation onto different regions. These differences may stem from the differences in the details of the tasks; O’Reilly et al. (2013) used saccadic planning task, which did not include reward but visuomotor learning, while Schwartenbeck et al. (2016) used a task in which one of two cues predicted valence of monetary outcomes. Therefore, their mapping results may reflect visuomotor or reward-related processes to a certain extent. Given that few studies separated updating and expectancy violation, more studies are necessary to assess their mapping under various tasks.

One important characteristic of our task was the presence of “mere” expectancy violation, which might be important for sensory encoding but may have no consequences on subsequent behavior or valuation. Our localization of expectancy violation in the AI is consistent with previous reports that the AI responds to abrupt or rare stimuli (Singer et al., 2009; Menon and Uddin, 2010) and is involved in reorienting to surprising events (Sokolov, 1963; Corbetta et al., 2008). Our results raise the possibility that the AI is primarily involved in detecting and rapidly broadcasting expectancy violation through distinctly large bipolar neurons called Von Economo neurons (Allman et al., 2010; Evrard et al., 2012), which may be irrelevant for ongoing tasks but still important for agents’ survival.

We also dissociated neural correlates of value updating from expectancy violation; value updating was correlated with activation in the MPFC and the cingulate cortex, not the AI (Fig. 3). These regions have been long associated with value-related processing, such as valuation of choices or RL (Bartra et al., 2013). Our value updating is conceptually close to model-based reward prediction error (RPE), which has been associated with these regions (Behrens et al., 2008; Daw et al., 2011). Our results provide evidence that value-related computation in these regions is not driven solely by expectancy violation. Interestingly, we did not find value-updating representation in the striatum, even at a liberal threshold, which has also been associated with model-based RPE (Behrens et al., 2008; Gläscher et al., 2010; Daw et al., 2011). This could be because we did not provide reward feedback to prevent learning over trials; while the MPFC and the cingulate may be involved in reduction of value uncertainty regardless of feedback, the striatum might primarily respond to reward feedback. Alternatively, the striatum could be more involved in learning over time, possibly through corticostriatal loops (Balleine et al., 2007), than one-shot updating.

Although we found that belief and value updating were anatomically distinct, their computational processes should not be independent under a Bayesian framework. However, in the ex-

treme, association between belief updating and activations in premotor/FEF and the IPS could be an epiphenomenal reflection of more general cognitive processes, such as working memory (Mohr et al., 2006; Reinhart et al., 2012). Contrary to this possibility, we found using DCM that processes in premotor/FEF and the IPS contribute to processes in the MPFC via connections modulated by the reducibility of value uncertainty (Fig. 4). We speculate that such modulation might be biologically efficient; when value uncertainty is irreducible and the valuation system can safely ignore incoming signals, energies to maintain synaptic transmission from belief-updating regions can be temporarily saved. Additionally, we did not find evidence for modulation in connections from expectancy-violation regions (AI) to value-updating regions. Even though interpretation of negative results requires caution, this suggests that reducibility-based modulation is not a general phenomenon across the cortex, and further emphasizes the difference between functional roles of the AI and the lateral frontoparietal regions.

Even though this study is not sufficient to describe a detailed and holistic picture of information flow across regions, one possibility is that the modulation is caused by previously reported representations of risk and ambiguity (Hsu et al., 2005; Huettel et al., 2006; Bach et al., 2011). In addition to affecting value-based choices directly, these representations might also influence how valuation is updated by incoming signals. Thus, representations of risk and ambiguity might have wider implications for neural mechanism of decision making under uncertainty than previously thought. Their causal role could be more directly investigated by such methodologies as transcranial magnetic stimulation or neurofeedback.

It is worth noting that DCM does not quantify or assume monosynaptic connections. Although previous studies have found anatomical connections from premotor/FEF and the IPS to the cingulate motor area, connections to the more anterior portion of the MPFC, where our value-updating ROIs are located, have not been reported (Tomassini et al., 2007; Beckmann et al., 2009; Mars et al., 2011; Eradath et al., 2015; Neubert et al., 2015). We speculate that connections from premotor/FEF and the IPS to the MPFC could be mediated by anterior–posterior connections within cingulate cortex (Margulies et al., 2007) or by corticostriatal connections (Balleine et al., 2007; Di Martino et al., 2008).

Our findings suggest that the nature of uncertainty influences the way the human brain processes environmental signals. To understand neural mechanisms of uncertainty reduction more comprehensively, future studies can expand this study's approach in a number of directions. For instance, while our paradigm delivered one signal at a time, real-world scenarios typically involve many signals. In these cases, the knowledge about uncertainty reducibility may also be crucial for adaptive allocation of attentional resources (Gottlieb and Balan, 2010). Additionally, more studies are necessary to clarify the encoding scheme of belief updating. In particular, due to the relatively small urn sizes, the current study is unable to distinguish between different operationalizations of belief updating. Richer parameterization of urn sizes would be useful, albeit at a cost of substantially increased task complexity. Finally, the knowledge about the reducibility of uncertainty in the real world is often far from complete and accurate. It is an open question how agents estimate the reducibility of uncertainty, particularly in nonstationary environments (Yu and Dayan, 2005; Behrens et al., 2007; Payzan-LeNestour and Bossaerts, 2011).

Notes

Supplemental material for this article is available at http://neuroecon.berkeley.edu/public/papers/Kobayashi_Ambiguity_SOM.pdf. This material has not been peer reviewed.

References

- Allman JM, Tetreault NA, Hakeem AY, Manaye KF, Semendeferi K, Erwin JM, Park S, Goubert V, Hof PR (2010) The von Economo neurons in fronto-insular and anterior cingulate cortex in great apes and humans. *Brain Struct Funct* 214:495–517. [CrossRef Medline](#)
- Bach DR, Hulme O, Penny WD, Dolan RJ (2011) The known unknowns: neural representation of second-order uncertainty, and ambiguity. *J Neurosci* 31:4811–4820. [CrossRef Medline](#)
- Balleine BW, Delgado MR, Hikosaka O (2007) The role of the dorsal striatum in reward and decision-making. *J Neurosci* 27:8161–8165. [CrossRef Medline](#)
- 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](#)
- Becker SW, Brownson FO (1964) What price ambiguity? Or the role of ambiguity in decision-making. *J Pol Econ* 72:62–73. [CrossRef](#)
- Beckmann M, Johansen-Berg H, Rushworth MF (2009) Connectivity-based parcellation of human cingulate cortex and its relation to functional specialization. *J Neurosci* 29:1175–1190. [CrossRef Medline](#)
- Behrens TE, Woolrich MW, Walton ME, Rushworth MF (2007) Learning the value of information in an uncertain world. *Nat Neurosci* 10:1214–1221. [CrossRef Medline](#)
- Behrens TE, Hunt LT, Woolrich MW, Rushworth MF (2008) Associative learning of social value. *Nature* 456:245–249. [CrossRef Medline](#)
- Boorman ED, Rushworth MF, Behrens TE (2013) Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. *J Neurosci* 33:2242–2253. [CrossRef Medline](#)
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436. [CrossRef Medline](#)
- Brett M, Anton J-L, Valabregue R, Poline JB (2002) Region of interest analysis using an SPM toolbox. *Neuroimage* 16:2.
- Camerer C, Weber M (1992) Recent developments in modeling preferences: uncertainty and ambiguity. *J Risk Uncertainty* 5:325–370. [CrossRef](#)
- Corbetta M, Patel G, Shulman GL (2008) The reorienting system of the human brain: from environment to theory of mind. *Neuron* 58:306–324. [CrossRef Medline](#)
- Courchesne E, Hillyard SA, Galambos R (1975) Stimulus novelty, task relevance and the visual evoked potential in man. *Electroencephalogr Clin Neurophysiol* 39:131–143. [CrossRef Medline](#)
- Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ (2011) Model-based influences on humans' choices and striatal prediction errors. *Neuron* 69:1204–1215. [CrossRef Medline](#)
- Di Martino A, Scheres A, Margulies DS, Kelly AM, Uddin LQ, Shehzad Z, Biswal B, Walters JR, Castellanos FX, Milham MP (2008) Functional connectivity of human striatum: a resting state fMRI study. *Cereb Cortex* 18:2735–2747. [CrossRef Medline](#)
- Ellsberg D (1961) Risk, ambiguity, and the Savage axioms. *Q J Econ* 75:643–669. [CrossRef](#)
- Eradath MK, Abe H, Matsumoto M, Matsumoto K, Tanaka K, Ichinohe N (2015) Anatomical inputs to sulcal portions of areas 9m and 8Bm in the macaque monkey. *Front Neuroanat* 9:30. [CrossRef Medline](#)
- Ergin H, Gul F (2009) A theory of subjective compound lotteries. *J Econ Theory* 144:899–929. [CrossRef](#)
- Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415:429–433. [CrossRef Medline](#)
- Evrard HC, Forro T, Logothetis NK (2012) Von Economo neurons in the anterior insula of the macaque monkey. *Neuron* 74:482–489. [CrossRef Medline](#)
- Friston KJ, Harrison L, Penny W (2003) Dynamic causal modelling. *Neuroimage* 19:1273–1302. [CrossRef Medline](#)
- Gläscher J, Daw N, Dayan P, O'Doherty JP (2010) States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66:585–595. [CrossRef Medline](#)
- Gottlieb J, Balan P (2010) Attention as a decision in information space. *Trends Cogn Sci* 14:240–248. [CrossRef Medline](#)

- Hayasaka S, Nichols TE (2003) Validating cluster size inference: random field and permutation methods. *Neuroimage* 20:2343–2356. [CrossRef Medline](#)
- Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF (2005) Neural systems responding to degrees of uncertainty in human decision-making. *Science* 310:1680–1683. [CrossRef Medline](#)
- Huetzel SA, Stowe CJ, Gordon EM, Warner BT, Platt ML (2006) Neural signatures of economic preferences for risk and ambiguity. *Neuron* 49:765–775. [CrossRef Medline](#)
- Hunt LT, Dolan RJ, Behrens TE (2014) Hierarchical competitions subserving multi-attribute choice. *Nat Neurosci* 17:1613–1622. [CrossRef Medline](#)
- Itti L, Baldi P (2009) Bayesian surprise attracts human attention. *Vis Res* 49:1295–1306. [CrossRef Medline](#)
- Keynes JM (1921) *A treatise on money*. London: McMillan.
- Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. *Econometrica* 73:1849–1892. [CrossRef](#)
- Knight FH (1921) *Risk, uncertainty and profit*. Boston: Houghton Mifflin.
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247. [CrossRef Medline](#)
- Kriegeskorte N, Simmons WK, Bellgowan PS, Baker CI (2009) Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci* 12:535–540. [CrossRef Medline](#)
- Ma WJ, Jazayeri M (2014) Neural coding of uncertainty and probability. *Annu Rev Neurosci* 37:205–220. [CrossRef Medline](#)
- Margulies DS, Kelly AM, Uddin LQ, Biswal BB, Castellanos FX, Milham MP (2007) Mapping the functional connectivity of anterior cingulate cortex. *Neuroimage* 37:579–588. [CrossRef Medline](#)
- Mars RB, Jbabdi S, Sallet J, O'Reilly JX, Crosson PL, Olivier E, Noonan MP, Bergmann C, Mitchell AS, Baxter MG, Behrens TE, Johansen-Berg H, Tomassini V, Miller KL, Rushworth MF (2011) Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and macaque resting-state functional connectivity. *J Neurosci* 31:4087–4100. [CrossRef Medline](#)
- Menon V, Uddin LQ (2010) Saliency, switching, attention and control: a network model of insula function. *Brain Struct Funct* 214:655–667. [CrossRef Medline](#)
- Mohr HM, Goebel R, Linden DE (2006) Content- and task-specific dissociations of frontal activity during maintenance and manipulation in visual working memory. *J Neurosci* 26:4465–4471. [CrossRef Medline](#)
- Mumford JA, Poline JB, Poldrack RA (2015) Orthogonalization of regressors in fMRI models. *PLoS One* 10:e0126255. [CrossRef Medline](#)
- Nassar MR, Wilson RC, Heasly B, Gold JJ (2010) An approximately Bayesian Δ -rule model explains the dynamics of belief updating in a changing environment. *J Neurosci* 30:12366–12378. [CrossRef Medline](#)
- Nau RF (2006) Uncertainty aversion with second-order utilities and probabilities. *Management Sci* 52:136–145. [CrossRef](#)
- Neubert FX, Mars RB, Sallet J, Rushworth MF (2015) Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. *Proc Natl Acad Sci U S A* 112:E2695–E2704. [CrossRef Medline](#)
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp* 15:1–25. [CrossRef Medline](#)
- O'Reilly JX, Schüffelen U, Cuell SF, Behrens TE, Mars RB, Rushworth MF (2013) Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proc Natl Acad Sci* 110:E3660–E3669. [CrossRef Medline](#)
- Payzan-LeNestour E, Bossaerts P (2011) Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Comput Biol* 7:e1001048. [CrossRef Medline](#)
- Pearce JM, Bouton ME (2001) Theories of associative learning in animals. *Annu Rev Psychol* 52:111–139. [CrossRef Medline](#)
- Pearce JM, Hall G (1980) A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol Rev* 87:532–552. [CrossRef Medline](#)
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442. [CrossRef Medline](#)
- Penny WD, Stephan KE, Daunizeau J, Rosa MJ, Friston KJ, Schofield TM, Leff AP (2010) Comparing families of dynamic causal models. *PLoS Comput Biol* 6:e1000709. [CrossRef Medline](#)
- Pouget A, Beck JM, Ma WJ, Latham PE (2013) Probabilistic brains: knowns and unknowns. *Nat Neurosci* 16:1170–1178. [CrossRef Medline](#)
- Reinhart RM, Heitz RP, Purcell BA, Weigand PK, Schall JD, Woodman GF (2012) Homologous mechanisms of visuospatial working memory maintenance in macaque and human: properties and sources. *J Neurosci* 32:7711–7722. [CrossRef Medline](#)
- Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: *Classical conditioning II: current research and theory* (Black AH, Prokasy WF, eds), pp 64–99. New York: Appleton Century Crofts.
- Roesch MR, Esber GR, Li J, Daw ND, Schoenbaum G (2012) Surprise! Neural correlates of Pearce–Hall and Rescorla–Wagner coexist within the brain. *Eur J Neurosci* 35:1190–1200. [CrossRef Medline](#)
- Schwartenbeck P, FitzGerald TH, Dolan R (2016) Neural signals encoding shifts in beliefs. *Neuroimage* 125:578–586. [CrossRef Medline](#)
- Seo K (2009) Ambiguity and second-order belief. *Econometrica* 77:1575–1605. [CrossRef](#)
- Singer T, Critchley HD, Preuschoff K (2009) A common role of insula in feelings, empathy and uncertainty. *Trends Cogn Sci* 13:334–340. [CrossRef Medline](#)
- Smith AP, Stephan KE, Rugg MD, Dolan RJ (2006) Task and content modulate amygdala-hippocampal connectivity in emotional retrieval. *Neuron* 49:631–638. [CrossRef Medline](#)
- Sokolov EN (1963) Higher nervous functions: the orienting reflex. *Annu Rev Physiol* 25:545–580. [CrossRef Medline](#)
- Squires NK, Squires KC, Hillyard SA (1975) Two varieties of long-latency positive waves evoked by unpredictable auditory stimuli in man. *Electroencephalogr Clin Neurophysiol* 38:387–401. [CrossRef Medline](#)
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. *Neuroimage* 46:1004–1017. [CrossRef Medline](#)
- Sutton RS, Barto AG (1998) *Reinforcement learning: an introduction*. Cambridge, MA: MIT.
- Tomassini V, Jbabdi S, Klein JC, Behrens TE, Pozzilli C, Matthews PM, Rushworth MF, Johansen-Berg H (2007) Diffusion-weighted imaging tractography-based parcellation of the human lateral premotor cortex identifies dorsal and ventral subregions with anatomical and functional specializations. *J Neurosci* 27:10259–10269. [CrossRef Medline](#)
- Tversky A, Kahneman D (1974) Judgment under uncertainty: heuristics and biases. *Science* 185:1124–1131. [CrossRef Medline](#)
- Vernet M, Quentin R, Chanes L, Mitsumasu A, Valero-Cabré A (2014) Frontal eye field, where art thou? Anatomy, function, and non-invasive manipulation of frontal regions involved in eye movements and associated cognitive operations. *Front Integr Neurosci* 8:66. [CrossRef Medline](#)
- Weiskopf N, Hutton C, Josephs O, Deichmann R (2006) Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *Neuroimage* 33:493–504. [CrossRef Medline](#)
- Woo CW, Krishnan A, Wager TD (2014) Cluster-extent based thresholding in fMRI analyses: pitfalls and recommendations. *Neuroimage* 91:412–419. [CrossRef Medline](#)
- Yu AJ, Dayan P (2005) Uncertainty, neuromodulation, and attention. *Neuron* 46:681–692. [CrossRef Medline](#)