Differential representations of prior and likelihood uncertainty in the human brain

Iris Vilares[a,b,c,1], **James D Howard**[d], **Hugo L Fernandes**[a,c,e], **Jay A Gottfried**[d,f], and **Konrad P Kording**[a,g,h]

[a]Department of Physical Medicine and Rehabilitation; Northwestern University and Rehabilitation Institute of Chicago; Chicago, Illinois, 60611; USA.

[b]International Neuroscience Doctoral Programme, Champalimaud Neuroscience Programme; Institutito Gulbenkian de Ciência; 2780-156 Oeiras; Portugal.

[c]Instituto de Tecnologia Química e Biológica; Universidade Nova de Lisboa; 2780-156 Oeiras; Portugal.

[d]Department of Neurology; Northwestern University Feinberg School of Medicine; Chicago, Illinois, 60611; USA.

[e]PhD Program in Computational Biology; Instituto Gulbenkian de Ciência; 2780-156 Oeiras; Portugal.

[f]Department of Psychology, Northwestern University, Weinberg College of Arts and Sciences, Evanston, Illinois, 60208; USA.

[g]Department of Physiology; Northwestern University; Chicago, Illinois, 60611; USA.

[h]Department of Applied Mathematics, Northwestern University, Chicago, Illinois, USA.

SUMMARY

Background—Uncertainty shapes our perception of the world and the decisions we make. Two aspects of uncertainty are commonly distinguished: uncertainty in previously acquired knowledge (prior) and uncertainty in current sensory information (likelihood). Previous studies have established that humans can take both types of uncertainty into account, often in a way predicted by Bayesian statistics. However, the neural representations underlying these parameters remain poorly understood.

Results—By varying prior and likelihood uncertainty in a decision-making task while performing neuroimaging in humans, we found that prior and likelihood uncertainty had quite distinct representations. While likelihood uncertainty activated brain regions along the early stages of the visuomotor pathway, representations of prior uncertainty were identified in specialized brain areas outside this pathway, including putamen, amygdala, insula, and orbitofrontal cortex. Furthermore, the magnitude of brain activity in the putamen predicted individuals' personal tendencies to rely more on either prior or current information.

Corresponding author: (iris-vilares@northwestern.edu; tel: (312) 238-6515)..

fMRI data acquisition, processing and analysis. See SI and **Sup. Tables** for details.

Conclusions—Our results suggest different pathways by which prior and likelihood uncertainty map onto the human brain, and provide a potential neural correlate for higher reliance on current or prior knowledge. Overall, these findings offer insights into the neural pathways that may allow humans to make decisions close to the optimal defined by a Bayesian statistical framework.

INTRODUCTION

Uncertainty is intrinsic to our world. For any given event, there is uncertainty in what our senses currently tell us – this is usually denoted likelihood uncertainty. There is also uncertainty in our preexisting knowledge of that event – this is known as prior uncertainty. For example, when judging the probability of rain, we combine current information obtained through our senses (are there clouds visible in the sky?) with previous acquired knowledge we possess about the chance of rain at our particular location (are we in Lisbon or London?). In such cases, the uncertainty associated with each piece of information determines how we should combine them. The combination of information gathered in the past (prior) with new information (likelihood) is critical for effective decision making[1] and can thus be seen as a central objective of the nervous system.

Bayesian statistics describe how prior and likelihood information can be optimally combined as a function of their respective uncertainties to give a posterior probability estimate. The uncertainty of this optimal estimate (posterior uncertainty, or general uncertainty) is generally smaller than the uncertainty associated with either prior or likelihood alone[2]. Several recent studies comparing Bayesian predictions to human behavior show that humans are close to optimal in a wide range of tasks, including estimation[3], learning[4] and movement[5-7]. The fact that behavior was close to the Bayesian optimal in these tasks indicates that human participants detect and use information about both prior and likelihood uncertainty. Nevertheless, in spite of a large body of behavioral research, it is still unclear how and where these types of uncertainty are represented in the brain.

Given that uncertainty is fundamental to behavior, there is an extensive modeling literature that hypothesizes how it could be represented. However, these different theoretical models do not tend to distinguish between the representation of priors and likelihoods. Also, they differ in their predictions of where uncertainty should be represented[8-12]. One set of theories assumes that uncertainty is a fundamental part of the way any pair of neurons exchanges information, and thus the representation of the uncertainty of a variable is always co-localized with the representation of the variable itself[8, 10]. A different set of theories assumes that there are specialized brain regions that encode and process uncertainty[9, 11-13]. Although these theories are not necessarily mutually exclusive, they offer different predictions, and none so far has received strong neurobiological support. It thus remains unknown whether uncertainty is represented along the sensorimotor pathway or within specialized brain areas outside this pathway, and whether different forms of uncertainty have different representations.

To try to find the neural correlates of uncertainty, several studies in monkeys have analyzed how uncertainty in a stimulus can change neural firing. For instance, in a classic visual discrimination task, monkeys view a cloud of randomly moving dots and need to identify their net direction[14, 15]. Varying the percentage of dots with a coherent motion demonstrated that the activity of neurons in the lateral intraparietal area (LIP) represents not only the direction of the stimulus but also the uncertainty associated with it[15]. Such studies can be interpreted as changing the likelihood. Other studies have changed the probability that even before seeing the stimulus a monkey would have to saccade to a given target, and found that this modulated neuronal activity[16]. These studies can be interpreted as changing the prior[14]. However, all of these studies are based on relatively simple

oculomotor tasks, with a focus on very specific brain areas. It remains poorly understood whether neural representations of uncertainty are also encoded elsewhere in the brain, and whether the findings in monkeys translate to other experimental tasks and settings in humans.

Interesting recent human studies in neuroeconomics have started to examine how uncertainty about reward is represented. These studies show that increased uncertainty about whether a reward is going to be obtained (risk) correlates with increased activations in the orbitofrontal cortex (OFC)[17, 18], midbrain[11], cingulate cortex[18, 19], and insula[9, 20, for a review, see 21]. However, in all these cases uncertainty was treated as one single parameter, and thus general uncertainty was perfectly correlated with prior or likelihood uncertainty, which were not distinguished as separate sources of uncertainty. Thus, although in humans there is an increasing knowledge of where uncertainty in reward is represented, it is still unclear where prior and likelihood uncertainty are, and if these areas coincide with the areas involved in uncertainty in reward.

Here we devised a sensory-motor decision-making task in which human subjects could use both prior and likelihood to estimate positions of hidden visual targets. Uncertainty was systematically varied on each trial in a two-by-two factorial design, such that two of the conditions were matched for performance accuracy but differed in the relative balance of prior and likelihood uncertainty. By combining event-related functional magnetic resonance imaging (fMRI) approaches with computational models of behavioral performance, we were able to characterize the neural representations of the two kinds of uncertainty while controlling for confounds related to expected reward. The behavioral results indicate that subjects are adept at using both kinds of uncertainty to optimize performance, in keeping with Bayesian predictions. The imaging results suggest that likelihood uncertainty is primarily represented in the early stages of the sensorimotor network, while prior uncertainty is represented in limbic and paralimbic decision-related areas outside of traditional sensorimotor pathways. Together these findings suggest fundamentally different representations by which prior and likelihood uncertainty in a decision-making task map onto the human brain.

RESULTS

In this study we wanted to know where prior and likelihood uncertainty are represented in the human brain. To this end, we developed a visual decision-making task in which subjects had to guess the position of a hidden target (a "coin") on a computer screen (see Fig. 1A and [6, 22]). Subjects were given noisy visual information about the target position in the form of a dot-cloud drawn from a Gaussian distribution centered at the true target position. To successfully estimate the position of the target, subjects could use both the likelihood, obtained from the displayed dots, and the prior, obtained from the distribution of previous target positions. Uncertainty in the likelihood varies with the dispersion of the displayed dots. This dispersion varied randomly from trial to trial, and thus could not be predicted beforehand. The average position of the target (the mean of the prior) was the middle of the screen, and subjects could estimate its uncertainty (prior uncertainty) from the distribution of target positions in previous trials (see Experimental Procedures and SI for details). Given that subjects had ample experience with the task from the behavioral experiment, they quickly acquired the prior. Successful estimates of the position were rewarded with points, which had motivational significance to the subjects (see Sup. Information -SI- and Fig. S1). The conditions comprised a two-by-two factorial design (Fig. 1B), with two levels of prior uncertainty (wide, more uncertain prior: **P**, and narrow, less uncertain prior: **p**) and two levels of likelihood uncertainty (wide likelihood: **L**, and narrow likelihood: **l**). Varying the

uncertainty in the prior and in the likelihood independently, together with fMRI imaging, allowed us to find where they are represented in the brain.

Prior and likelihood uncertainty affect behavior

We first wanted to know if variations of prior and likelihood uncertainty in our task influence the estimation behavior. If subjects ignore the prior information and rely only on the current sensory feedback (i.e., likelihood information) then the weight of sensory feedback (slopes of Fig. 2A) in the estimation should be one. On the other hand, if subjects rely only on their prior knowledge and ignore likelihood information, then the weight of the sensory feedback should be zero. Use of this metric demonstrated that scanned subjects relied on both prior and likelihood information ($0 < slope < 1$; Fig. 2B). Data obtained from all subjects (both those who were scanned and those who were not) showed the same effects (Fig. S2A). Furthermore, subjects relied more on the likelihood information as the prior uncertainty increased and as the likelihood uncertainty decreased (main effect of prior, $F_{1,43} = 207$, $p < 10^{-6}$; main effect of likelihood, $F_{1,43} = 35$, $p < 10^{-6}$, ANOVA repeated measures (r.m.)). Qualitatively, such behavior would be expected if they used a Bayesian strategy[23] (see also SI). Thus, in our experiment, subjects utilized knowledge of both prior and likelihood uncertainty for perceptual decision-making.

Both kinds of uncertainty may be expected to change the precision of subjects' estimates and thus their expected task performance. Not surprisingly, performance was better when each of the uncertainties was lower (Fig. 2C, main effect of prior, $F_{1,43} = 161$, $p < 10^{-6}$; main effect of likelihood, $F_{1,43} = 84$, $p < 10^{-6}$, ANOVA r.m.). Importantly, despite the fact that the slopes, and thus the relative weighting of the uncertain sources of information, differed significantly between the pL and Pl conditions (Fig. 2B, $p < 1 \times 10^{-4}$, W=0, Wilcoxon signed rank test), estimation performance between these two conditions was matched (Fig. 2C and Fig. S2C, $p > 0.05$, W=31, Wilcoxon signed rank test). Even testing estimation performance on each individual subject did not reveal any significant difference between these two conditions ($p > 0.05$ for each subject, corrected for multiple comparisons, comparing two proportions test). Therefore, as noted below, a comparison of these two conditions should provide a robust way to infer imaging-based differences between prior and likelihood uncertainty, while minimizing potential confounds due to differences in general uncertainty, performance, or reward expectation.

Prior and likelihood uncertainty have distinct neural representations

The uncertainty associated with the prior affects behavior (Fig. 2B). We thus wanted to ask if and where it affects BOLD activities (see SI for details). We found that areas more active with increased uncertainty in the prior include putamen, amygdala, insula, and OFC ($p < 0.05$ whole-brain corrected; family-wise-error (FWE) level; Fig. 3A). Cross-validated condition-specific activation profiles (Fig. 3B) demonstrate that the activation in each of these regions is specifically related with prior uncertainty, and not likelihood uncertainty. Please note that, because in our design there is no "zero prior uncertainty" condition no real baseline exists, and thus what is relevant in the activation profiles is the relative difference between conditions. There was a significant main effect of prior uncertainty (all regions depicted are significant at $p < 0.01$, $F > 9$; except left amygdala, with a $p = 0.059$, $F_{1,43} = 3.76$; ANOVA r.m.), with no significant main effects of likelihood ($F_{1,43} < 0.73$, $p > 0.05$; ANOVA r.m.). Even if we use a two-fold cross-validation, the same results still hold (see SI for details). We did not find any imaging evidence of an interaction between priors and likelihoods (for additional tests and controls, see SI and Fig. S3A). Areas more active with low prior uncertainty (i.e. more active with increased precision/confidence) encompassed the caudate nucleus, prefrontal cortex and areas adjacent to the anterior cingulate cortex (see Fig. S3B). Together,

these results suggest that wide regions of the brain, primarily outside of the traditional sensory-motor pathway, encode prior uncertainty.

Given that activity related to prior uncertainty encompassed a wide set of regions, we wanted to know if the uncertainty associated with the currently presented stimulus (the likelihood uncertainty) leads to activity in the same areas. We should expect that a neural representation exists, given that likelihood uncertainty also affected behavior (Fig. 2B). Areas more active with high vs. low likelihood uncertainty were localized to bilateral regions of superior occipital visual cortex (p<0.05 FWE corrected; see Fig. 4) and nowhere else in the brain at this threshold (see SI and Fig. S4A for controls). Taking advantage of the fact that the degree of likelihood uncertainty differed on every trial (where greater dot dispersion corresponded to higher visual variance and higher likelihood uncertainty) and that the between-trial variability in dot position produces some overlap between the low and high likelihood uncertainty groups (see SI and Fig. S4B-C), we also implemented a parametric fMRI model to test whether dot-dispersion variance modulated the same areas of occipital cortex on a trial-by-trial basis. This model revealed the same bilateral activations in the superior occipital visual cortex. These activations increased parametrically with the increase in likelihood uncertainty, had a higher level of significance and extended all the way to the calcarine sulcus and lingual gyrus (p<0.05 FWE corrected; Fig. 4C). Likelihood uncertainty in our task thus seems to affect BOLD signal mainly in areas corresponding to the early stages of the visuomotor pathway.

Correlations across subjects between prior-related behavior and activations

Thus far, we have analyzed how the two factors, prior and likelihood uncertainty, affect activities in the brain, and identified a small number of key regions. For those analyses we did not use observed behavior but rather the experimentally and mathematically defined uncertainties. However, if any of these brain regions are involved in the decision making process then we should expect them to be correlated with behavioral variance across subjects (see also Experimental Procedures and SI).

We hypothesized that subjects who were able to better detect the changes in prior uncertainty (i.e. those who changed more their behavior when the prior uncertainty changed) could also show higher differences in brain activity between the two prior conditions. This is indeed what was observed for both the right and left putamen (Fig. 5A-B, r=0.65 and r=0.61, p<0.02, Spearman correlation) and the OFC (Fig. 5C, r=0.66, p<0.01, Spearman correlation). The other reported areas (namely, amygdala and insula) did not show a significant correlation. These results suggest that, besides tracking prior uncertainty, putamen and OFC might be directly related with behavioral change.

We also wanted to look for a behavioral measure that could represent a specific "personality trait". The average slope of a subject (averaged over the four conditions) can provide such a measure, as a higher value indicates that subjects on average tend to rely more on current sensory information and less on prior knowledge, and vice versa. Indeed, as we have seen above, subjects varied behaviorally in how much, on average, they relied on prior information (see Fig. 2), and also varied in their mean fMRI activation response to uncertainty in prior-related brain areas (see Fig. 3). We found that this average slope was positively correlated with the degree of differential brain activation in the right putamen (Fig. 5D, r=0.66, p<0.01, Spearman correlation). Note that the effect in the putamen (in terms of difference between conditions) is actually the highest compared to the other brain areas (see Sup. Table S1). This finding raises the possibility that a general higher involvement of the putamen in the task, potentially by signaling greater prior uncertainty, may enhance learning from new sensory information, and therefore subjects in whom the putamen is more engaged tend to rely more on new than in prior information.

Online representation of Bayesian slopes

We wanted to ask how the brain computes prior versus likelihood information online in order to arrive at a sensorimotor decision. For that decision, at each trial the subject needs to decide how much to weigh the likelihood relative to the prior. If subjects behave as predicted by Bayesian statistics, then this slope can be estimated for each trial (see Experimental Procedures and SI for details). Testing for a correlation between this "instantaneous Bayesian slope" and BOLD activation, we found significant decision related activations in the prefrontal cortex (PFC), roughly in Brodmann area 9 (p<0.0001 unc., see Fig. 6). This finding suggests a possible role of prefrontal cortex in combining prior and likelihood information to estimate the target's position.

DISCUSSION

In this study we tested where prior and likelihood uncertainty are represented in the brain. We were particularly interested in knowing: 1) Are prior and likelihood uncertainty represented in the same set of areas? 2) Is uncertainty for visuomotor tasks represented in the traditional visuomotor pathway or in specialized areas? By combining a psychophysical paradigm with fMRI analysis in which two of the conditions were matched for performance and posterior uncertainty, we could disentangle the specific effects of prior and likelihood uncertainty. We found that greater prior uncertainty evoked increased brain activity in specialized brain areas that include the putamen, amygdala, OFC and parts of the insula. In contrast, greater likelihood uncertainty primarily affected neural activities in the occipital cortex, in areas that belong to the traditional visuomotor pathway. Prior and likelihood uncertainty were thus represented in largely non-overlapping areas, highlighting the importance of distinguishing between these kinds of uncertainty.

There are numerous reasons why the brain could use a strategy where prior uncertainty is processed in these specialized brain areas. Computationally, the forming of a prior results from integration over time, requiring long-term memory, and may be more difficult to implement in domain-specific sensory/motor brain areas. Anatomically, OFC, amygdala, putamen and insula receive inputs from sensory areas, which are needed to build a prior, and project to motor effector systems in the brainstem and cortex[12, 24-26] which are needed to use the prior. These connections may allow the optimization of behavior in the context of varying prior uncertainty and, if the prior is too uncertain, facilitate behavior change to gather relevant new information[26]. Additionally, these areas are connected or even directly involved with reward processing[11, 21, 24], and could thus more easily combine the need for more information with its potential value. Moreover, previous research has highlighted their involvement in uncertainty in reward[12, 17, 18, 20, 27] and, more specifically, in signaling ambiguity and the need to learn more about the world[13]. Thus, computational demands, anatomical connections, and previous research support the involvement of these areas in signaling prior uncertainty.

Likelihood uncertainty, in contrast, was associated with activations in areas that are part of the traditional visuomotor pathway. Our task is visual, and thus information about the task is also transmitted through the same route. Hence, these results concur with the hypothesis presented by theoretical models, such as probabilistic population codes and sampling theories[8, 10], that likelihood uncertainty is part of the inherent code by which neurons transmit information. According to these models, the activity of the same neurons transmits information (e.g. the position) along with uncertainty about this information. From a computational perspective, sensory information needs to be continuously used to calculate estimates of likelihood uncertainty[4] and hence sensory areas seem best suited for this ongoing update. Indeed, previous studies in human visual perception have also found that uncertainty affects brain activity in the corresponding sensory areas[28, 29], and, moreover,

they found that activity in visual areas was higher when the visual stimulus was more uncertain[28, 30, 31]. This occurred for random/ nonrandom dot motion[31], incoherent/ coherent shapes[30] and blurry/non-blurry images[29], indicating that the effects observed are not exclusive to the particular stimulus we used. Together, our results suggest that likelihood uncertainty about a visual stimulus may be processed along with the stimulus itself in the visual cortex.

For Bayesian decision-making, the brain needs not only to compute prior and likelihood uncertainty, but also to use them for appropriate weighting of both pieces of information. Activity in the brain areas where this occurs should then relate to these weights, which depend nonlinearly on both uncertainties. We found that an area in the PFC tracks the trial-by-trial weight on current vs. prior information. This area would then be a candidate area to receive information from both prior and likelihood uncertainty and calculate accordingly how much weight should be placed on new information. Indeed, the PFC is known for its role in planning and cognitive control[32] and, interestingly, has even been specifically associated with Bayesian decision-making[33]. Our results thus suggest how a network of brain areas may give rise to Bayesian instantiations of perception and behavior.

Although our study focused on the integration of priors with visual information, uncertainty may be represented differently for other sensory modalities or tasks. Our task dealt with a new, rapidly acquired prior over a series of trials, in which subjects, during each trial, had unlimited time to make a decision. Hence, it can be considered a cognitive task. Studies using different kinds of tasks, but that were also cognitive, have previously associated the prior-uncertainty areas identified here with decision-making and reward uncertainty[11, 18, 20, 34], making it more likely that our results would hold for other types of cognitive tasks as well. However, it is possible that other types of priors, such as the ones involved in early sensory perception (e.g. "light from above"), which might have been learned over many years (or even generations), have distinct representations[35]. As for likelihood uncertainty, if indeed it is transmitted concurrently with the sensory information itself, then a non-visual task should activate non-visual brain areas. Interesting future studies could help to answer these questions, for example by applying a computationally equivalent task in which the sensory feedback information given is not visual but, say, olfactory or auditory; or by implementing a faster/more unconscious sensory perception task.

The association of specific brain areas with prior uncertainty and potentially with subjects' individual tendency to rely on new sensory information (as is the case for putamen) may have implications for the understanding of learning disabilities and abnormal decision-making behavior. Learning, in a Bayesian sense, can be interpreted as the weight given to new evidence over prior beliefs[36]. If the prior is more certain then less learning from new information should occur, and vice versa[36]. Changes in the brain areas that represent prior uncertainty might then lead for example to an underestimation of prior uncertainty, potentially affecting learning from new information. Interestingly, it has been found that patients with damage to the amygdala[34], OFC[13], insula[37] and putamen[38] show considerable deficits in making decisions that involve uncertainty and learning from feedback, and that the existence of intact connections between these structures is essential for learning from new sensory information[39, 40]. If indeed one of the reasons why these deficits occur is based on an underestimation of prior uncertainty, then these could potentially be reversed by providing more certain current sensory information or by giving explicit information about prior uncertainty.

The finding of the areas involved in prior and likelihood uncertainty representation provides insight for an ongoing debate in the computational literature: is uncertainty part of the general code by which neurons exchange information and thus encoded in every neuron's

output[8, 10], or are there specialized areas that deal with the encoding of uncertainty[9, 11-13]? Our findings suggest that both of these hypotheses might be correct, but for different kinds of uncertainty: likelihood uncertainty seems intrinsically embodied in the stimulus encoding itself, as it is represented in sensorimotor areas, while prior uncertainty is encoded in specialized areas. Future models of brain function should take into account uncertainty in the prior and in the likelihood separately, since both their neural representations and their behavioral effects are distinct.

EXPERIMENTAL PROCEDURES

Subjects

Twenty-seven healthy subjects (12 women; age range 19-35 years; mean age=27 years) participated in the experiment. Of these, 17 participated in the fMRI experiment, from which data from 15 were used (9 women). All participants were right-handed, had normal or corrected-to-normal vision, were naive to the goals of the experiment, signed consent forms and were paid to participate. Subjects that only performed the behavioral part of the experiment were paid $20. Subjects that performed both the behavioral and the fMRI parts of the experiment were paid $70. All protocols were approved by the Northwestern University IRB.

Behavioral task

Subjects performed a decision-making task, which consisted of guessing the position of a hidden coin on a screen, in a task similar to one described in prior studies[6, 22]. Subjects were told the cover story of a coin being tossed into a pond and informed that their task was to guess where the coin had fallen. They could not see the coin, but they could see 5 blue dots that were the "splashes" produced by the coin falling in. They were told that the person who threw the coin aimed, albeit imperfectly, at the center of the screen (mean of prior). They were also told that, between blocks, the thrower changed, and the new one might be better or worse at throwing (i.e. they were indirectly informed that the variance of the prior changed). To estimate the coin position, subjects could use (although they were never explicitly told so) both the coin position's likelihood, obtained from the "splashes", and its prior (the distribution of previous coin locations). There was no temporal deadline.

Stimuli

The position of the coin was drawn from a Gaussian distribution, centered on the center of the screen with a standard deviation (std) that was either low (σ_p=2.5% of screen width) or high (σ_P=8.5% of screen width). This distribution was the *prior* of the experiment. Subjects were given the mean of the prior ("the coin throw is aimed at the screen center") but not its variance, which they could only estimate from the distribution of previous coin throws. The standard deviation of the prior was kept constant within blocks, but changed across blocks. On every trial, a cluster of five dots was shown on the screen. The x-position of each of these dots was drawn independently from a second Gaussian distribution in which the mean was the coin's horizontal location on that trial and the standard deviation was either low (σ_l=6%) or high (σ_L=15%). The distribution of these five dots defined the *likelihood*. The std of the likelihood was varied pseudo-randomly from trial to trial but counterbalanced across trials. We made the std of the likelihood vary pseudorandom from trial to trial so that subjects could not predict *a priori* the overall uncertainty that the trial would have. In total there were thus four conditions: low prior uncertainty and low likelihood uncertainty (pl); low prior uncertainty and high likelihood uncertainty (pL); high prior uncertainty and low likelihood uncertainty (Pl) and high prior uncertainty and high likelihood uncertainty (PL). See SI for more details.

Procedure Behavioral session

See SI for details.

Procedure fMRI session

From the 27 subjects recruited for the behavioral task, 17 were asked to return and perform the same task while undergoing fMRI (See SI for details).

Data Analysis

We used a Bayesian model to understand the behavior of the subjects in our task (see SI for details).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

REFERENCES

1. Cox RT. Probability, Frequency and Reasonable Expectation. American Journal of Physics. 1946; 14:1–13.

2. Glimcher, PW. Decisions, uncertainty, and the brain: The science of neuroeconomics. The MIT Press; 2004.

3. Miyazaki M, Nozaki D, Nakajima Y. Testing Bayesian models of human coincidence timing. Journal of Neurophysiology. 2005; 94:395–399. [PubMed: 15716368]

4. Burge J, Ernst MO, Banks MS. The statistical determinants of adaptation rate in human reaching. J Vision. 2008; 8:1–19.

5. Kording KP, Wolpert DM. Bayesian integration in sensorimotor learning. Nature. 2004; 427:244–247. [PubMed: 14724638]

6. Tassinari H, Hudson TE, Landy MS. Combining priors and noisy visual cues in a rapid pointing task. Journal of Neuroscience. 2006; 26:10154–10163. [PubMed: 17021171]

7. Stevenson IH, Fernandes HL, Vilares I, Wei KL, Kording KP. Bayesian Integration and Non-Linear Feedback Control in a Full-Body Motor Task. Plos Comput Biol. 2009; 5:1–9.

8. Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. Nature Neuroscience. 2006; 9:1432–1438.

9. Schultz W, Preuschoff K, Camerer C, Hsu M, Fiorillo CD, Tobler PN, Bossaerts P. Explicit neural signals reflecting reward uncertainty. Philosophical Transactions of the Royal Society B-Biological Sciences. 2008; 363:3801–3811.

10. Fiser J, Berkes P, Orban G, Lengyel M. Statistically optimal perception and learning: from behavior to neural representations. Trends in Cognitive Sciences. 2010; 14:119–130. [PubMed: 20153683]

11. Preuschoff K, Bossaerts P, Quartz SR. Neural differentiation of expected reward and risk in human subcortical structures. Neuron. 2006; 51:381–390. [PubMed: 16880132]

12. Singer T, Critchley HD, Preuschoff K. A common role of insula in feelings, empathy and uncertainty. Trends in Cognitive Sciences. 2009; 13:334–340. [PubMed: 19643659]

13. Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF. Neural systems responding to degrees of uncertainty in human decision-making. Science. 2005; 310:1680–1683. [PubMed: 16339445]

14. Gold JI, Shadlen MN. Neural computations that underlie decisions about sensory stimuli. Trends in Cognitive Sciences. 2001; 5:10–16. [PubMed: 11164731]

15. Shadlen MN, Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. Journal of Neurophysiology. 2001; 86:1916–1936. [PubMed: 11600651]

16. Basso MA, Wurtz RH. Modulation of neuronal activity by target uncertainty. Nature. 1997; 389:66–69. [PubMed: 9288967]

17. Tobler PN, O'Doherty JP, Dolan RJ, Schultz W. Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. Journal of Neurophysiology. 2007; 97:1621–1632. [PubMed: 17122317]

18. Critchley HD, Mathias CJ, Dolan RJ. Neural activity in the human brain relating to uncertainty and arousal during anticipation. Neuron. 2001; 29:537–545. [PubMed: 11239442]

19. Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS. Learning the value of information in an uncertain world. Nature Neuroscience. 2007; 10:1214–1221.

20. Huettel SA, Song AW, McCarthy G. Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices. Journal of Neuroscience. 2005; 25:3304–3311. [PubMed: 15800185]

21. Rangel A, Camerer C, Montague PR. A framework for studying the neurobiology of value-based decision making. Nature Reviews Neuroscience. 2008; 9:545–556.

22. Berniker M, Voss M, Kording K, Brezina V. Learning priors for Bayesian computations in the nervous system. PloS one. 2010; 5:e12686. [PubMed: 20844766]

23. Vilares I, Kording K. Bayesian models: the structure of the world, uncertainty, behavior, and the brain. Ann N Y Acad Sci. 2011; 1224:22–39. [PubMed: 21486294]

24. Rolls ET. The orbitofrontal cortex and reward. Cerebral Cortex. 2000; 10:284–294. [PubMed: 10731223]

25. Albin RL, Young AB, Penney JB. The Functional-Anatomy of Basal Ganglia Disorders. Trends in Neurosciences. 1989; 12:366–375. [PubMed: 2479133]

26. Whalen PJ. Fear, vigilance, and ambiguity: Initial neuroimaging studies of the human amygdala. Current Directions in Psychological Science. 1998; 7:177–188.

27. Kepecs A, Uchida N, Zariwala HA, Mainen ZF. Neural correlates, computation and behavioural impact of decision confidence. Nature. 2008; 455:227–U255. [PubMed: 18690210]

28. Beauchamp MS, Pasalar S, Ro T. Neural substrates of reliability-weighted visual-tactile multisensory integration. Front Syst Neurosci. 2010; 4:25. [PubMed: 20631844]

29. Helbig HB, Ernst MO, Ricciardi E, Pietrini P, Thielscher A, Mayer KM, Schultz J, Noppeney U. The neural mechanisms of reliability weighted integration of shape information from vision and touch. Neuroimage. 2012; 60:1063–1072. [PubMed: 22001262]

30. Murray SO, Kersten D, Olshausen BA, Schrater P, Woods DL. Shape perception reduces activity in human primary visual cortex. Proceedings of the National Academy of Sciences of the United States of America. 2002; 99:15164–15169. [PubMed: 12417754]

31. McKeefry DJ, Watson JDG, Frackowiak RSJ, Fong K, Zeki S. The activity in human areas V1/V2, V3, and V5 during the perception of coherent and incoherent motion. Neuroimage. 1997; 5:1–12. [PubMed: 9038280]

32. Miller EK. The prefrontal cortex and cognitive control. Nature Reviews Neuroscience. 2000; 1:59–65.

33. Hampton AN, Bossaerts P, O'Doherty JP. The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. Journal of Neuroscience. 2006; 26:8360–8367. [PubMed: 16899731]

34. Brand M, Grabenhorst F, Starcke K, Vandekerckhove MMP, Markowitsch HJ. Role of the amygdala in decisions under ambiguity and decisions under risk: Evidence from patients with Urbach-Wiethe disease. Neuropsychologia. 2007; 45:1305–1317. [PubMed: 17070876]

35. Fischer BJ, Pena JL. Owl's behavior and neural representation predicted by Bayesian inference. Nature Neuroscience. 2011; 14:1061–U1163.

36. Courville AC, Daw ND, Touretzky DS. Bayesian theories of conditioning in a changing world. Trends in Cognitive Sciences. 2006; 10:294–300. [PubMed: 16793323]

37. Clark L, Bechara A, Damasio H, Aitken MRF, Sahakian BJ, Robbins TW. Differential effects of insular and ventromedial prefrontal cortex lesions on risky decision-making. Brain. 2008; 131:1311–1322. [PubMed: 18390562]

38. Shohamy D, Myers CE, Grossman S, Sage J, Gluck MA, Poldrack RA. Cortico-striatal contributions to feedback-based learning: converging data from neuroimaging and neuropsychology. Brain. 2004; 127:851–859. [PubMed: 15013954]

39. Cohen MX, Elger CE, Weber B. Amygdala tractography predicts functional connectivity and learning during feedback-guided decision-making. Neuroimage. 2008; 39:1396–1407. [PubMed: 17997112]

40. Baxter MG, Parker A, Lindner CCC, Izquierdo AD, Murray EA. Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. Journal of Neuroscience. 2000; 20:4311–4319. [PubMed: 10818166]

HIGHLIGHTS

- Determines the neural correlates of two important Bayesian statistics parameters

- Likelihood uncertainty is represented in the associated sensory pathway (visual)

- In contrast, prior uncertainty is represented in putamen, amygdala, insula and OFC

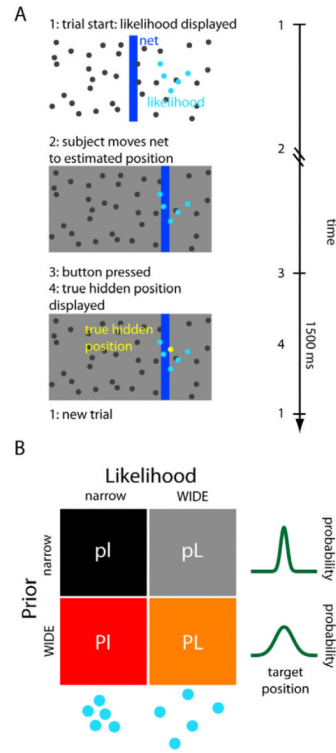- Activity in putamen correlates with subjects' tendency to rely on prior information

Figure 1. Experimental setup
(A) Illustration of the task. Subjects guess the position of a hidden target (the "coin", represented by the yellow dot) using a net (vertical blue bar) which they can displace horizontally. At the onset of each trial, subjects receive noisy information about the position of the hidden target in the form of a set of 5 blue dots (the likelihood). Subjects then move the net to the guessed position and press the mouse button to confirm their choice, after which the true target position is displayed. If the target is inside the net, the score increases by one point (see also SI and Fig. S1). A new trial then begins 1500 ms later. Left: illustration of the computer display that was presented to the subjects. Right: typical time course of a trial. (B) The 4 conditions of the experiment. The experiment consisted of a two-by-two factorial design, with two types of prior (p=narrow prior; P=wide prior) and two types of likelihood (l=narrow likelihood; L=wide likelihood). The wider conditions are the ones with more associated uncertainty.
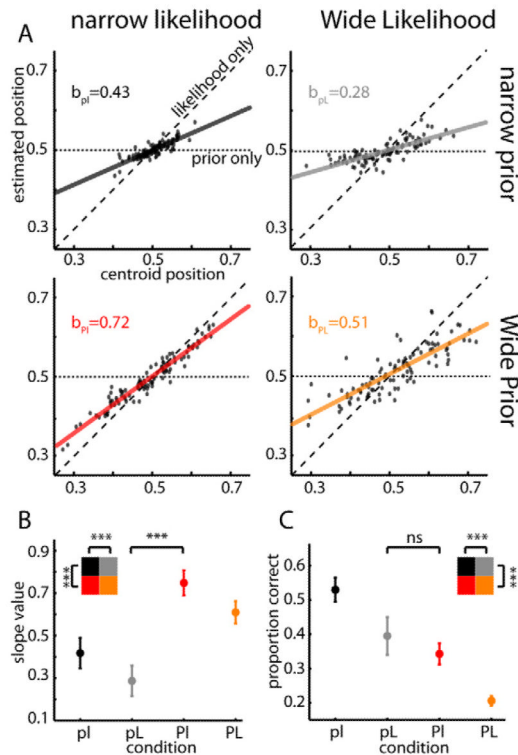
Figure 2. Behavioral results
(A) Estimates of the target position for one representative subject are shown as a function of the centroid of the displayed dots (likelihood). Displayed next to the graphs is the slope value of the linear regressions (solid line). The dashed line represents what the linear regression would look like if the subject only used likelihood information (slope=1). The horizontal dotted line represents a potential situation in which only prior information is taken into account (slope=0, localized at 0.5 which is the middle of the screen and the mean of the prior). (B) Average slope of the linear regression for the behavior of the 15 subjects during the scanning session, separated by condition. The slope quantifies the degree to which subjects rely on the current visual stimulus (likelihood) vs. the prior. The small blue rectangles represent the optimal Bayesian values (see Supplemental Experimental Procedures for details). (C) Average proportion of trials in which the subject accurately guessed the position of the target, separated by condition. Error bars in *(B)* and *(C)* represent 95% confidence intervals for the mean (s.e.m.). Inset shows that there is a significant effect of both prior and likelihood uncertainty. *n.s.*= non-significant, p>0.05; *** significant, p<0.001. See also Fig. S2.
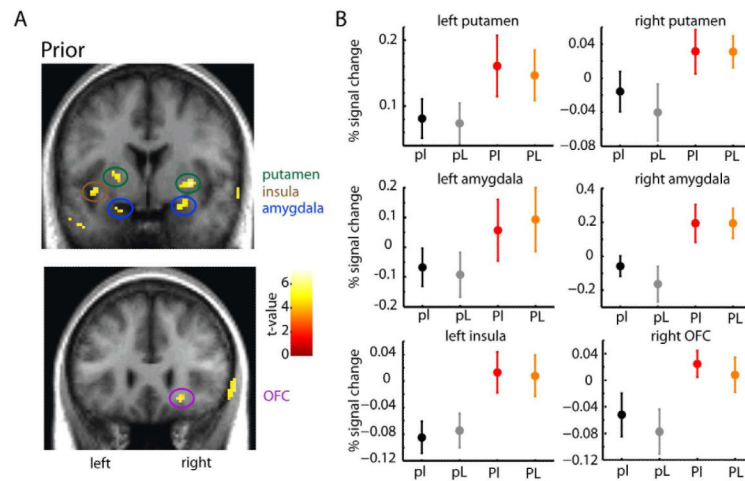
Figure 3. Brain regions more active during wide (more uncertain) prior conditions
(A) Stronger activations associated with high (vs. low) prior uncertainty were seen bilaterally in the putamen, amygdala, insula (top; y=0), and OFC (bottom; y=26). Functional activations are overlaid on coronal sections of the average of each subject's T1-weighted structural brain scan (display threshold at p<0.0001 unc., minimum 10 voxels; n=15). Activity in the right insula and left OFC appears at a less stringent p-value (p<0.001; not shown). In this and all subsequent figures, the right side of the brain corresponds to the right side of the image. (B) Percent signal change (PSC) by condition in the areas represented in *(A)*. Data were extracted from the peak (most significant) voxel in each cluster using a leave-one-subject-out cross-validation procedure (see Supplemental Experimental Procedures for details). Plots represent the subject-averaged parameter estimates converted to PSC for the four conditions (n=15). Error bars represent standard error of the mean (*s.e.m.*). See also Fig. S3.
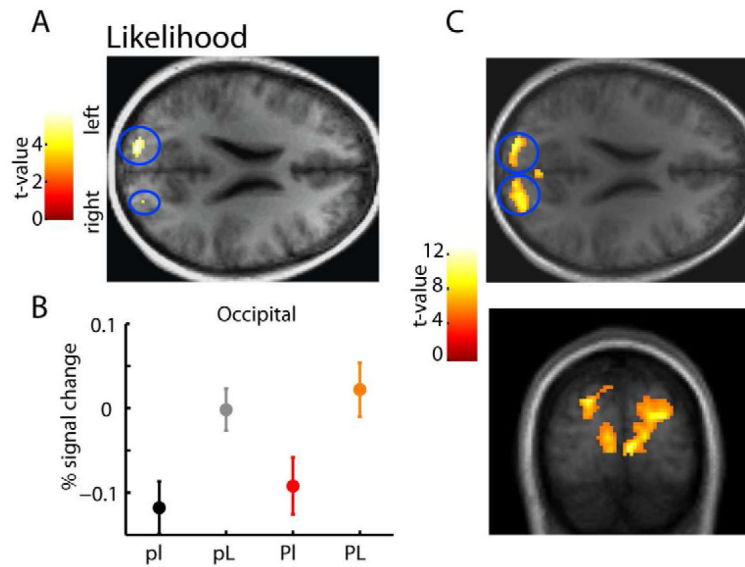
Figure 4. Brain regions more active during higher likelihood uncertainty conditions
(A) Stronger activations associated with high (vs. low) likelihood uncertainty were seen bilaterally in the superior occipital cortex (z=22). Functional activations are overlaid on the axial section of the subject-averaged scan (n=15). (B) Condition-specific percent signal change in the left occipital area shows a main effect of likelihood uncertainty significant at p<0.0001, with no significant main effect of prior (p>0.05). Data were extracted using a leave-one-subject-out cross-validation procedure. The same results hold if using a two-fold cross-validation procedure. The data plot represents the average parameter estimates (±s.e.m., cross-validated) converted to percentage signal change for the four conditions. (C) Brain regions parametrically correlated with higher likelihood uncertainty. Stronger activations associated with a parametric increase in likelihood uncertainty (standard deviation of the displayed dots at each trial) were seen bilaterally in the superior occipital cortex, extending down all the way to the calcarine sulcus and lingual gyrus (y=−88; z=22). Functional activations are overlaid on the axial (up) and coronal (down) sections of the subject-averaged scan (for all functional activations shown, display threshold is at p<0.0001 unc., minimum 10 voxels; n=15). See also Fig. S4.
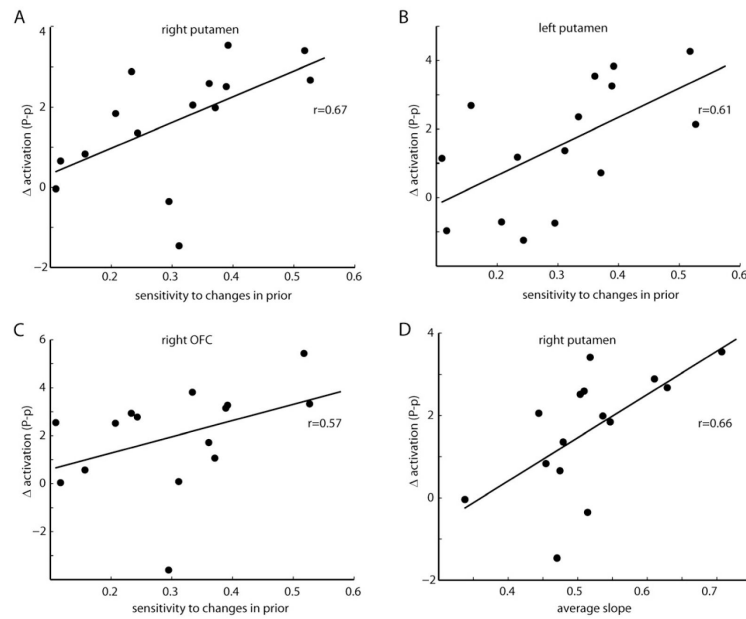
Figure 5. Relationship between prior uncertainty-related brain activity and behavioral measures
(A-C) Prior uncertainty differences in activation for individual participants at the right putamen (A; peak at 30, −2, −2), left putamen (B; peak at −22,0,4) and OFC (C; peak at 22,26,−12) are regressed against individuals' sensitivity to prior change. Presented on the graphics are the respective Spearman correlations. (D) Prior uncertainty differences in activation for individual participants at the right putamen (peak at 30, −2, −2) are regressed against the average weight given to likelihood vs prior information (the slope). Presented on the graphic is the respective Spearman correlation.
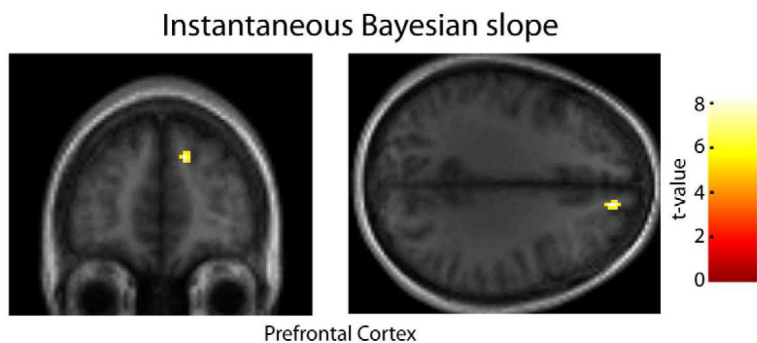
Instantaneous Bayesian slope

Prefrontal Cortex

Figure 6. Brain regions parametrically correlated with the instantaneous Bayesian slope
Stronger activations associated with a parametric increase in the instantaneous Bayesian slope (a measure that indicates how much the subject weights prior and likelihood information, at each trial) were seen in the right superior medial prefrontal cortex (y=48; z=32). Functional activations are overlaid on the coronal section (left) and axial (right) sections of the subject-averaged scan (display threshold at p<0.0001 uncorrected, minimum 10 voxels).