Title: Context-invariant neural dynamics underlying the encoding of Bayesian uncertainty, but not confidence

Abbreviated Title: Network dynamics of uncertainty and confidence

Vincenzo G. Fiore[1*], Xiaosi Gu[1,2,3*]

[1] Department of Psychiatry, Icahn School of Medicine at Mount Sinai, New York, NY 10029-5674, USA

[2] Nash Family Department of Neuroscience, Icahn School of Medicine at Mount Sinai, New York, NY 10029-5674, USA

[3] The Mental Illness Research, Education, and Clinical Center (MIRECC VISN 2) at the James J. Peter Veterans Affairs Medical Center, Bronx, NY

[*] Correspondence should be addressed to: Vincenzo G. Fiore, vincenzo.fiore@mssm.edu, Xiaosi Gu, xiaosi.gu@mssm.edu

Number of Pages: 22

Number of figures: 4

Number of words: 238 (abstract), 597 (introduction), 864 (discussion)

Conflict of Interest. The authors declare no financial and non-financial competing interest.

Acknowledgments. We thank Prof. Karl Friston for his comments and kind suggestions in shaping this manuscript. We thank Jennifer Jung for her help in setting up the code for both tasks (Psychtoolbox library in Matlab), Ann-Cathrin V. Guertler, Chandana C. Tatineni, and Ju-Chi Yu, for their help in collecting the data. XG is supported by NIDA R01DA043695 and the Mental Illness Research, Education, and Clinical Center (MIRECC VISN 2) at the James J. Peter Veterans Affairs Medical Center, Bronx, NY.

Abstract

Beliefs about action-outcomes contingencies are often updated in opaque environments where feedbacks might be inaccessible and agents might need to rely on other information for evidence accumulation. It remains unclear, however, whether and how the neural dynamics subserving confidence and uncertainty during belief updating might be context-dependent. Here, we applied a Bayesian model to estimate uncertainty and confidence in healthy humans (n=28) using two multi-option fMRI tasks, one with and one without feedbacks. We found that across both tasks, uncertainty was computed in the anterior insular, anterior cingulate, and dorsolateral prefrontal cortices, whereas confidence was encoded in anterior hippocampus, amygdala and medial prefrontal cortex. However, dynamic causal modelling (DCM) revealed a critical divergence between how effective connectivity in these networks was modulated by the available information. Specifically, there was directional influence from the anterior insula to other regions during uncertainty encoding, independent of outcome availability. Conversely, the network computing confidence was driven either by the anterior hippocampus when outcomes were not available, or by the medial prefrontal cortex and amygdala when feedbacks were immediately accessible. These findings indicate that confidence encoding might largely rely on evidence accumulation and therefore dynamically changes as a function of the available sensory information (i.e. symbolic sequences monitored by the hippocampus, and monetary feedbacks computed by amygdala and medial prefrontal cortex). In contrast, uncertainty could be triggered by any information that disputes existing beliefs (i.e. processed in the insula), independent of its content.

Significance Statement

Our choices are guided by our beliefs about action-outcome contingencies. In environments where only one action leads to a desired outcome, high estimated action-outcome probabilities result in confidence, whereas low probabilities distributed across multiple choices result in uncertainty. These estimations are continuously updated, sometimes based on feedbacks provided by the environment, but sometimes this update takes place in opaque environments where feedbacks are not readily available. Here, we show that uncertainty computations are driven by the anterior insula, independent of feedback availability. Conversely, confidence encoding dynamically adapts to the information available, as we found it was driven either by the anterior hippocampus, when feedback was absent, or by the medial prefrontal cortex and amygdala, otherwise.

Introduction

In classic decision-making theories, the presence of explicit outcomes or feedback after choice selections is considered a crucial source of information for behavioral adaptability and belief updating, e.g. by triggering prediction error signals (Rangel et al., 2008; Glimcher, 2011). However, real-life decision-making often occurs in opaque environments where outcomes can be sporadic or temporarily inaccessible, so that confidence or uncertainty in estimated action-outcome contingencies must rely on other sources of information (Pouget et al., 2013; Ma and Jazayeri, 2014; Meyniel et al., 2015a). Previous studies on value-based decision-making with explicit outcomes (Platt and Huettel, 2008; Tajima et al., 2016) have revealed multiple neural regions involved in encoding uncertainty or confidence in action-outcome estimations (White et al., 2014; Morriss et al., 2018). However, it is yet to be determined whether the neural dynamics underlying one's uncertainty or confidence are affected by the sources of available information.

We developed two multi-option decision-making tasks to allow non-binary decision-making (Churchland et al., 2008; Churchland and Ditterich, 2012) in the presence or absence of explicit stochastic outcomes in humans (**Fig. 1**). Both tasks were designed to allow a quantifiable manipulation of uncertainty and confidence on a trial-by-trial basis. The environments granted only one optimal choice per trial, and participants could estimate action-outcome contingencies on the basis of the evidence made available. Thus, given n available choices, decision-making resulted in nuanced transitions between trials characterized by high uncertainty (i.e. $\frac{1}{n}$ probability to produce the desired outcome, for all available choices) and those characterized by high confidence (i.e. one choice selection associated with the desired outcome with high probability). At the same time, we aimed at reducing second-order uncertainty (Allen et al., 2016; Spence et al., 2016), by relying on the

simplicity of task structures, with unambiguous stimuli used as evidence, and a low level of stochasticity, where outcomes were available. These features of the tasks were exploited in a Bayesian model, used to quantify uncertainty and confidence in both tasks.

In Bayesian terms, beliefs represent – as probability distributions – the estimated likelihood any action will produce a desired outcome (Kording and Wolpert, 2006; Payzan-LeNestour and Bossaerts, 2011; Fleming and Daw, 2017). As uncertainty normally increases in presence of conflicting information and confidence increases in the face of accumulating confirming evidence, it has been proposed that one derives from the inverse of the other (Ma and Jazayeri, 2014; Meyniel et al., 2015b; Meyniel et al., 2015a; Adler and Ma, 2018). The theory of Bayesian inference has found empirical support in studies demonstrating that probabilistic distributions are taken into account in choice selections (Orban and Wolpert, 2011; Sanders et al., 2016) and they are represented by the activity of neural populations (Ma et al., 2006; Pouget et al., 2013; Rich et al., 2015). Investigations into the neural computations of uncertainty and confidence have revealed a rich network of regions involved (Pouget et al., 2016; Meyniel and Dehaene, 2017; Morriss et al., 2018). In particular, neural regions monitoring conflict and errors, with a prominent role for insular (Seymour et al., 2004; Palminteri et al., 2012; Gu et al., 2015), and dorsal anterior cingulate (ACC) cortices (Behrens et al., 2007; Botvinick, 2007; Rushworth and Behrens, 2008), and areas such as the medial prefrontal cortex (mPFC) or the amygdala, implicated in the processing of value estimates, in presence of risk (Koechlin and Hyafil, 2007; Sarinopoulos et al., 2010; Peters et al., 2017). Based on this literature, and considering that our model-defined uncertainty was formalized as the inverse of our model-defined confidence for both task environments, we hypothesized that both uncertainty- and confidence-encoding networks would be found to change, depending on the context.

Materials and Methods

Experimental design and statistical analyses I: 3-options continuous choice tasks. In a first task, termed *bead task* (**Fig. 1A**) (modified from: Huq et al., 1988), a new visual stimulus (a red, blue or green *bead*) was presented at each trial, adding a congruent (e.g. a red bead after another red one) or incongruent (e.g. a blue bead following a green one) color stimulus in a sequence of beads. Participants had 2 seconds to decide from which of three jars displayed on the monitor the bead had been extracted. The jars were illustrated on screen as containing beads of three colors in a ratio of 80%-10%-10%. After each button press, the selected jar would be highlighted with a black rectangle, for the remaining time on the clock allowed for the choice selection, plus 0.5 seconds. The last five extracted beads were always present on screen (**Fig. 1A**), and the participants were allowed to make the first choice selection starting from the 5th bead extraction. Between trials, a grey square would appear to conceal the new extracted bead for a variable time of 2 to 3.5 seconds. Participants had access to immediate feedbacks during a training session, which would show that each correct guess would result in accumulating 100 points. However, no feedback was provided during the MRI task and the participants were made aware that the accumulated points would be disclosed only at the end of the task.

In the second task, termed *card task* (**Fig. 1B**) (modified from: O'Doherty et al., 2001), the participants had 2 seconds to select among three cards presented on the screen and characterized by three geometric figures (randomly selected among triangle, square, circle, star and diamond). Each card was assigned a different predominant value among the three possible outcomes of 100, 10 and 0 and after selecting a card, the screen would display a stochastic outcome, highlighted in green (variable duration: 1-2.5 seconds), with assigned

probability of 80%, for the predominant value, and 10% for each of the remaining outcomes. A second message on screen also signaled the amount won on white screen (fixed interval: 0.5 seconds), followed by a fixation cross (variable interval: 1-2.5 seconds), which would precede a new trial. The total amount of points accumulated, per block, was always displayed in the lower part of the screen.

Finally, the participants were instructed that the computer would randomly change the jar from which it extracted beads, or the card-value associations. The pace of these pseudo-random changes was determined in an interval of 5 to 7 trials for both tasks and it was independent of the performance of the participants. For both tasks, participants were compensated with $1 for every 500 points, selecting the points accumulated during one random block per task. Both tasks consisted in 3 blocks of 71 trials each, so the participants were told that the maximum amount of bonus they could gain consisted in about $15 dollars from each task. Three identical sequences -one per block- were used for all subjects, for both the bead colors and the card-value associations. Task order and block order were counterbalanced across subjects.

Experimental design and statistical analyses II: model and parametric modulators. We used two similar Bayesian models to estimate: 1) in the bead task, the probability that the latest color bead had been extracted from the selected jar; 2) in the card task, the probability that the selected card was associated with the 100 value, given the outcomes received from the previous choices. In both cases, we used a standard Bayes theorem in the context of Bayes updating:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

Where *P(A|B)*, indicates the conditional probability of the event A to occur, given the evidence B. *P(A)* indicates the distribution of probabilities before new evidence is provided (carried from a previous trial), and *P(B)* is derived from the subject-specific assumptions about the distribution of events in the task environment. We assumed these were normally distributed:

$$f(B|\sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(B)^2}{2\sigma^2}}$$

where the mean of the normal distribution is assumed to be zero ($\mu = 0$), the variable σ (standard deviation of the normal distribution) was fixed per each subject, to compute *P(B)*, using the law of total probability. Under these conditions, higher values of σ resulted in simulated behaviors characterized by slower updates. For the Bead task, a new colored bead evidence would update the probability distribution associated with the three jars, increasing the probability associated with the matched colored jar and equally decreasing the probability associated with the remaining ones. Similarly, for the card task, an outcome equivalent to 100 would be used as positive evidence to increase the probability to find the same outcome again by selecting the same card. Outcomes of 10 or 0 were used to decrease the probability to find the maximum reward by repeating the same choice, therefore equally increasing the probability associated with both non-selected cards.

We used a Monte Carlo method to estimate the value of the parameter σ, per each task, that would better match real choice selections expressed by the 28 participants. We then coupled the trial-by-trial distribution of probabilities resulting from these model-based estimations with the actual choice selection of each participant, having in turn the estimated probability *p* associated with each choice. This resulted in the trial-by-trial estimated choice-dependent probability to reach the desired outcome (i.e. the estimated probability to guess the correct extraction jar or the card associated with 100 points) and therefore was used as a parametric

modulator for the analysis of confidence. The inverse *1-p*, indicating the estimated probability to receive any outcome but the desired one, given the chosen action, was used as parametric modulator of uncertainty.

Participants. We recruited 28 healthy volunteers (16 females), age 24.8 ± 7.0. Participants taking any medication, or with a history of mental disorder or drug abuse were excluded from the study. One subject was excluded from all fMRI analysis involving the card task, due to excessive movement. The study was approved by the Institutional Review Board at the University of Texas, Dallas and University of Texas Southwestern Medical Center. Informed consent was obtained from all subjects and all participants were informed that they could withdraw from the study at any point.

fMRI data acquisition and preprocessing. Functional MRI data were acquired using a Philips 3-Tesla MR scanner at the Advanced Imaging Research Center at University of Texas Southwestern Medical Center. The anatomical scan sequence (multi-echo MPRAGE) was carried out with a resolution of 1 mm, Multi Parametric Maps. Functional images (EPI) were acquired with a resolution of 3.4 x 3.4 x 4mm, repetition time of 2000 milliseconds, echo time of 25 milliseconds, 38 axial slices, flip angle=90°, and a field of view of 240 mm. We used Statistical Parametric Mapping (SPM12, Wellcome Department of Imaging Neuroscience) for data preprocessing. The functional images were motion realigned to the first volume, coregistered to the participant's anatomical scan, MNI normalized, and spatially smoothed using an isotropic 8-mm full-width at half-maximum (FWHM) Gaussian kernel.

Model-based fMRI. For the GLM analysis, we used default SPM12 functions to observe BOLD signals associated with the onsets of choice selections (i.e. button presses). This is

consistent with our computational definition of confidence and uncertainty, as well as with existing literature (De Martino et al., 2013; Ma and Jazayeri, 2014; Meyniel et al., 2015a), as where these phenomena in association with a choice, i.e. confidence or uncertainty that the selected action with deliver the desired result. RTs were used as trial-by-trial covariates in both tasks, and parametric modulators (*p* for confidence and *1-p* for uncertainty) were used to identify the relationship between the -parametrically modulated- task events and the hemodynamic response. We concatenated the three blocks for each task. Whole brain activations were determined using a threshold of $P<.05$ corrected for familywise errors ($P_{FWE}<.05$; **Fig. 2A,B**).

Dynamic causal modelling (DCM). For the DCM estimation, we relied on SPM12 default functions. Our model-based GLM revealed two distinct networks subserving uncertainty (insula, ACC, dlPFC; **Fig. 2A,C,E**) and confidence (hippocampus, amygdala, mPFC; **Fig. 2B,D,F**). We extracted fMRI time series from individual ROIs, using their principal eigenvariates: we relied on anatomical maps manually generated for the amygdala and the anterior hippocampus, and we used spherical ROIs (8mm radius) for all cortical regions. These were centered on task-specific, group-averaged, peaks activity, at the following coordinates: anterior insula (bead: [-27, 23, -1] and [33, 23, 2]; card:[-30, 20, -7] and [30, 23, -4]), ACC (bead: [-6, 29, 41] and [6, 29, 41]; card:[-9, 29, 23] and [9, 29, 26]), dlPFC (bead: [-42, 23, 26] and [42, 29, 29]; card:[-36, 26, 32] and [36, 26, 29]), mPFC (bead: [-9, 56, -1] and [9, 56, -1]; card: [-9, 56, -4]). We also included the visual cortex as a network input region, with ROIs centred on the following coordinates: bead: [12 -88 32] and [-12 -94 -7]; card: [21 -91 11] and [-21 -91 11]. With the exception of the response to the signal of confidence recorded during the card task, which was limited to the ROIs in the left hemisphere, we found significant increase in BOLD response bilaterally, for all ROIs. Thus,

we used independent DCM analysis for both hemispheres for a cross validation of the results within tasks.

For the network architectures (the *A matrix*), we assumed reciprocal (extrinsic) connections among all four regions (the signal-dependent ROIs triplets and the input ROI of the visual cortex), as well as recurrent (self) connections. Then we established the targets of the modulatory inputs (the *B matrix*), represented by the signal of estimated confidence (*p*) or uncertainty (*1-p*). We assumed that these modulatory inputs would affect the connectivity between visual cortex and the main three active ROIs, bilaterally, for all models, but they would target the connectivity between the three main ROIs in one direction only. Finally, we assumed the sensory input would only directly drive the activity of the visual cortex (*C matrix*). These assumptions resulted in eight models, which were characterized by identical architecture and targets of sensory input, but differed for their *B matrix* or targets of modulatory inputs (combination of 2 possible directions between 3 pairs of nodes, **Fig. 3A-B**).

We used a Bayesian family-wise random effect approach to estimate the directed connectivity between each pair of nodes within the two network triplets, grouping the eight models each time into two families of four models, each family characterized by one common modulated directed connectivity. This method allowed comparing, for instance, the four models with an insula-to-ACC modulated connectivity (with varying connectivity between insula and dlPFC and between ACC and dlPFC, **Fig. 3A**) vs the four models characterized by an ACC-to-insula modulated connectivity (and varying connectivity for the remaining couples of nodes). The method allowed to assess the likelihood a modulator signal would affect the information flow in a specific direction, within a single pair of nodes, independent of the presence of a single

winning model that would have emerged in a single model comparison. Finally, we tested our results independently across hemispheres. The estimated exceedance probability was deemed inconclusive for the ACC-dlPFC connectivity, in both tasks, and for the connectivity between Amg and mPFC, limited to the bead task, due to lack of convergence across hemispheres and insufficient difference between the compared families.

Results

Parameter regression with Monte Carlo resulted in subject-specific values determining each participant's own pace of belief update (σ=.497±.095, σ=.321±.156, respectively for the bead and card task, **Fig. 1C**). These values resulted in optimized Bayesian models that provided a mean behavioral prediction accuracy of 85.75%±4.73% for the bead task and 75.93%±6.44% for the card task.

In the context of our study, we defined confidence as the estimated probability to succeed in securing the desired outcome (*p*), and uncertainty as the inverse probability (*1-p*), given a selected action (De Martino et al., 2013; Meyniel et al., 2015a; Pouget et al., 2016). Specifically, the model was used to estimate, on a trial-by-trial basis, each participant's confidence (*p*) that their selected jar was the extraction jar (bead task) or their selected card was associated with the maximum reward (card task). A subject by subject correlation analysis revealed this estimated confidence (*p*) was negatively correlated with the reaction times (RTs) in the bead task (r= -.43 ± .16; *P*=.01 ± .05), but not in the card task (r= -.01 ± .11; *P*=.33 ± .31; cf. **Fig. 1D** and **E**). Next, we included RTs values as covariates and we examined the neural activations associated with the estimated uncertainty (1-*p*) or confidence (*p)* (**Fig. 1D,E**). Consistent with previous literature (Payzan-LeNestour et al., 2013; Meyniel and Dehaene, 2017), we found that uncertainty was encoded in bilateral anterior insula,

bilateral ACC, and right dorsolateral prefrontal cortex (dlPFC), across tasks [**Fig. 2A;** $P<.05$, family-wise error (FWE) corrected]. Conversely, confidence was encoded in the mPFC, left anterior hippocampus (aHip) and left amygdala (Amg), across both tasks (**Fig. 2B**, $P_{FWE}<.05$). A contrast between the two tasks did not reveal any difference, after FWE correction. A less stringent threshold ($P<.005$, uncorrected) revealed increased activity in the right lateral frontopolar cortex, in association with uncertainty processing, when comparing Card vs Bead task BOLD activity, but this result would not survive FWE correction.

Next, we used DCM (Friston et al., 2003) to investigate the causal dependencies among uncertainty- and confidence-encoding regions. Family-wise model comparison estimated independently on both hemispheres revealed that, for both tasks, *uncertainty* modulated insula-to-ACC [exceedance probability for bead task: 94% (left hemisphere), 67% (right hemisphere), **Fig. 4A;** card task: 65% (left),75% (right), **Fig. 4C**] and insula-to-dlPFC connectivity [exceedance probability for bead task 75% (left),78% (right), **Fig. 4A**; card task: 69% (left),72% (right), **Fig. 4C**]. However, the directed influences associated with *confidence-encoding* depended on feedback availability. Confidence primarily modulated the connections from the anterior hippocampus to other regions in the bead task (aHip-to-mPFC: 94% (left),83% (right) and aHip-to-Amg: 97% (left),61% (right), **Fig. 4B**), when immediate outcome was absent. In contrast, when feedback was available, confidence primarily modulated mPFC-to-aHip, mPFC-to-Amg, and Amg-to-aHip connectivity (exceedance probability, left hemisphere: 95% and 81%, and 82%, **Fig. 4D**).

Discussion

Uncertainty and confidence are typically considered two sides of the same coin, where one increases as the other decreases. Our findings reveal a more complex and nuanced picture,

demonstrating that the relationship between uncertainty and confidence is not the same across all Marrian levels of analysis (Marr and Poggio, 1976). Specifically, despite the fact we defined uncertainty and confidence as computationally symmetric processes (i.e. model-estimated p and $1-p$), we found asymmetry and divergence in the implementation of these estimations in terms of neural network dynamics. In our tasks, both confidence- and uncertainty-encoding neural regions remain unchanged, i.e. they are independent from the task frame. Both tasks presented similar action-outcome volatility, implying that different sources of information (colored beads vs. card/money) were used by the same neural areas (hippocampus-mPFC-amygdala and insula-ACC-dlPFC) to estimate changes in probability distributions of expected outcomes, independent of feedback availability. Furthermore, the use of dynamic causal modelling revealed that uncertainty-related neural network dynamics were also not affected by the source of information available. Specifically, we found that the anterior insula drove the uncertainty-encoding network and exerted its directed influence over ACC and dlPFC, regardless of the presence of immediate monetary feedback. The insula-ACC-dlPFC circuit has been consistently implicated in various cognitive tasks (e.g. attention, response selection, cognitive control), leading to the hypothesis that they are part of a ubiquitous salience network (Seeley et al., 2007). This network also responds to any threat to homeostasis, as, for instance, in the case of information conflicting with established beliefs, which seems to be independent of the nature of such information (Gu and FitzGerald, 2014). In this sense, our finding of the insula driving uncertainty computations across tasks is consistent with its role in salience processing and belief updating (Gu and FitzGerald, 2014; Uddin, 2015), demonstrating that the brain monitors conflicting information in action-outcome contingencies, even in the absence of immediate outcomes.

In sharp contrast, we found that the directional influence among brain regions encoding confidence in action-outcome contingencies changed depending on the availability of explicit outcomes and the source of available information. As confidence encoding in the bead task relies on the sequence of colored beads as evidence, we speculate that this process engages first the hippocampus to track and accumulate this information (Strange et al., 2005; Zeidman and Maguire, 2016), and then mPFC and amygdala for the value estimations on the basis of this information (Gluth et al., 2015). Differently, in the card task, evidence accumulation is based on expected values estimated from the stochastic outcomes; thus, in this case, the mPFC and amygdala are the first to be engaged to monitor changes in action-outcome contingencies (Koechlin and Hyafil, 2007; De Martino et al., 2013).

This is an intriguing decoupling of the network implementation for confidence and uncertainty, despite the fact that the computational definitions of the two estimations are perfectly anticorrelated in both tasks. Importantly, these results were achieved relying on multi-alternative decision-making processes, which present several advantages. First, they avoid a critical downside of binary decision tasks, which are biased by a confirmation/exclusion criterion, as the exclusion of one option inherently implies the only alternative available is the correct choice to make. Second, despite the significant increase in computational demands associated with 3-options tasks (Churchland et al., 2008; Churchland and Ditterich, 2012; Itthipuripat et al., 2018; Tajima et al., 2019), the tasks used for this study rely on a simple set of rules, easy-to-compute stochastic associations (80% for the most prevalent bead in a jar and for the most likely outcome value associated to a card) and only three types of discrete evidence (i.e. the three feedback values or the three bead colors), reducing the effect of second-order uncertainty (Fleming and Dolan, 2012; Fleming and Daw, 2017). By avoiding these issues, our tasks allow a more precise manipulation of first-order

uncertainty and confidence, as the maximum uncertainty (determined by an even distribution of probabilities among all available options) increases with the number of possible choices ($1 - \frac{1}{n}$). This feature eventually results in a large number of trials available for the investigation of uncertainty in particular, strengthening the fMRI BOLD responses as well as the DCM analysis. On this matter, the key findings supporting our interpretation are consistent across hemispheres and tasks, after running four independent sets of DCM analyses (in the case of the uncertainty network) or three independent sets of DCM analyses (for the confidence network). Thus, even in those cases reporting lower values for the exceedance probabilities, the cross validations increased the robustness of the overall result in terms of estimated effective connectivity.

Finally, it can be argued that, in terms of active inference, this dissociation we found between Bayesian *uncertainty* and *confidence* in belief updating (Friston et al., 2017) can be accounted considering the former corresponds to *uncertainty about hidden states* of the world (that does not depend upon reward), while the latter corresponds to *uncertainty about policies* being pursued (that depends upon prior preferences). This dissociation also fits comfortably with the functional anatomy described above. Taken together, our findings reveal both the distinct neurocomputational mechanisms underlying uncertainty and confidence encoding, and the importance of outcomes in directing information flow between the neural regions computing confidence, but not uncertainty.

Limitations: simplicity of the model does not account for volatility e.g. compare with (Adams et al., 2016) for a model of Beads task

References:

Adams RA, Huys QJ, Roiser JP (2016) Computational Psychiatry: towards a mathematically informed understanding of mental illness. J Neurol Neurosurg Psychiatry 87:53-63.

Adler WT, Ma WJ (2018) Comparing Bayesian and non-Bayesian accounts of human confidence reports. PLoS Comput Biol 14:e1006572.

Allen M, Frank D, Schwarzkopf DS, Fardo F, Winston JS, Hauser TU, Rees G (2016) Unexpected arousal modulates the influence of sensory noise on confidence. Elife 5.

Behrens TE, Woolrich MW, Walton ME, Rushworth MF (2007) Learning the value of information in an uncertain world. Nat Neurosci 10:1214-1221.

Botvinick MM (2007) Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. Cognitive, affective & behavioral neuroscience 7:356-366.

Churchland AK, Ditterich J (2012) New advances in understanding decisions among multiple alternatives. Curr Opin Neurobiol 22:920-926.

Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. Nat Neurosci 11:693-702.

De Martino B, Fleming SM, Garrett N, Dolan RJ (2013) Confidence in value-based choice. Nat Neurosci 16:105-110.

Fleming SM, Dolan RJ (2012) The neural basis of metacognitive ability. Philos Trans R Soc Lond B Biol Sci 367:1338-1349.

Fleming SM, Daw ND (2017) Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. Psychol Rev 124:91-114.

Friston K, FitzGerald T, Rigoli F, Schwartenbeck P, Pezzulo G (2017) Active Inference: A Process Theory. Neural Comput 29:1-49.

Friston KJ, Harrison L, Penny W (2003) Dynamic causal modelling. Neuroimage 19:1273-1302.

Glimcher PW (2011) Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. Proc Natl Acad Sci U S A 108 Suppl 3:15647-15654.

Gluth S, Sommer T, Rieskamp J, Buchel C (2015) Effective Connectivity between Hippocampus and Ventromedial Prefrontal Cortex Controls Preferential Choices from Memory. Neuron 86:1078-1090.

Gu X, FitzGerald TH (2014) Interoceptive inference: homeostasis and decision-making. Trends Cogn Sci 18:269-270.

Gu X, Wang X, Hula A, Wang S, Xu S, Lohrenz TM, Knight RT, Gao Z, Dayan P, Montague PR (2015) Necessary, yet dissociable contributions of the insular and ventromedial prefrontal cortices to norm adaptation: computational and lesion evidence in humans. J Neurosci 35:467-473.

Huq SF, Garety PA, Hemsley DR (1988) Probabilistic judgements in deluded and non-deluded subjects. Q J Exp Psychol A 40:801-812.

Itthipuripat S, Cha K, Deering S, Salazar AM, Serences JT (2018) Having More Choices Changes How Human Observers Weight Stable Sensory Evidence. J Neurosci 38:8635-8649.

Koechlin E, Hyafil A (2007) Anterior prefrontal function and the limits of human decision-making. Science 318:594-598.

Kording KP, Wolpert DM (2006) Bayesian decision theory in sensorimotor control. Trends Cogn Sci 10:319-326.

Ma WJ, Jazayeri M (2014) Neural coding of uncertainty and probability. Annu Rev Neurosci 37:205-220.

Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. Nat Neurosci 9:1432-1438.

Marr D, Poggio T (1976) From understanding computation to understanding neural circuitry. Cambridge: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.

Meyniel F, Dehaene S (2017) Brain networks for confidence weighting and hierarchical inference during probabilistic learning. Proc Natl Acad Sci U S A 114:E3859-E3868.

427 Meyniel F, Sigman M, Mainen ZF (2015a) Confidence as Bayesian Probability: From Neural Origins to
428 Behavior. Neuron 88:78-92.
429 Meyniel F, Schlunegger D, Dehaene S (2015b) The Sense of Confidence during Probabilistic Learning:
430 A Normative Account. PLoS Comput Biol 11:e1004305.
431 Morriss J, Gell M, van Reekum CM (2018) The uncertain brain: A co-ordinate based meta-analysis of
432 the neural signatures supporting uncertainty during different contexts. Neurosci Biobehav
433 Rev.
434 O'Doherty J, Kringelbach ML, Rolls ET, Hornak J, Andrews C (2001) Abstract reward and punishment
435 representations in the human orbitofrontal cortex. Nat Neurosci 4:95-102.
436 Orban G, Wolpert DM (2011) Representations of uncertainty in sensorimotor control. Curr Opin
437 Neurobiol 21:629-635.
438 Palminteri S, Justo D, Jauffret C, Pavlicek B, Dauta A, Delmaire C, Czernecki V, Karachi C, Capelle L,
439 Durr A, Pessiglione M (2012) Critical roles for anterior insula and dorsal striatum in
440 punishment-based avoidance learning. Neuron 76:998-1009.
441 Payzan-LeNestour E, Bossaerts P (2011) Risk, unexpected uncertainty, and estimation uncertainty:
442 Bayesian learning in unstable settings. PLoS Comput Biol 7:e1001048.
443 Payzan-LeNestour E, Dunne S, Bossaerts P, O'Doherty JP (2013) The neural representation of
444 unexpected uncertainty during value-based decision making. Neuron 79:191-201.
445 Peters MAK, Thesen T, Ko YD, Maniscalco B, Carlson C, Davidson M, Doyle W, Kuzniecky R, Devinsky
446 O, Halgren E, Lau H (2017) Perceptual confidence neglects decision-incongruent evidence in
447 the brain. Nature human behaviour 1.
448 Platt ML, Huettel SA (2008) Risky business: the neuroeconomics of decision making under
449 uncertainty. Nat Neurosci 11:398-403.
450 Pouget A, Drugowitsch J, Kepecs A (2016) Confidence and certainty: distinct probabilistic quantities
451 for different goals. Nat Neurosci 19:366-374.
452 Pouget A, Beck JM, Ma WJ, Latham PE (2013) Probabilistic brains: knowns and unknowns. Nat
453 Neurosci 16:1170-1178.
454 Rangel A, Camerer C, Montague PR (2008) A framework for studying the neurobiology of value-
455 based decision making. Nat Rev Neurosci 9:545-556.
456 Rich D, Cazettes F, Wang Y, Pena JL, Fischer BJ (2015) Neural representation of probabilities for
457 Bayesian inference. Journal of computational neuroscience 38:315-323.
458 Rushworth MF, Behrens TE (2008) Choice, uncertainty and value in prefrontal and cingulate cortex.
459 Nat Neurosci 11:389-397.
460 Sanders JI, Hangya B, Kepecs A (2016) Signatures of a Statistical Computation in the Human Sense of
461 Confidence. Neuron 90:499-506.
462 Sarinopoulos I, Grupe DW, Mackiewicz KL, Herrington JD, Lor M, Steege EE, Nitschke JB (2010)
463 Uncertainty during anticipation modulates neural responses to aversion in human insula and
464 amygdala. Cereb Cortex 20:929-940.
465 Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD (2007)
466 Dissociable intrinsic connectivity networks for salience processing and executive control. J
467 Neurosci 27:2349-2356.
468 Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS
469 (2004) Temporal difference models describe higher-order learning in humans. Nature
470 429:664-667.
471 Spence ML, Dux PE, Arnold DH (2016) Computations underlying confidence in visual perception. J
472 Exp Psychol Hum Percept Perform 42:671-682.
473 Strange BA, Duggins A, Penny W, Dolan RJ, Friston KJ (2005) Information theory, novelty and
474 hippocampal responses: unpredicted or unpredictable? Neural Netw 18:225-230.
475 Tajima S, Drugowitsch J, Pouget A (2016) Optimal policy for value-based decision-making. Nat
476 Commun 7:12400.

477 Tajima S, Drugowitsch J, Patel N, Pouget A (2019) Optimal policy for multi-alternative decisions. Nat
478 Neurosci.
479 Uddin LQ (2015) Salience processing and insular cortical function and dysfunction. Nat Rev Neurosci
480 16:55-61.
481 White TP, Engen NH, Sorensen S, Overgaard M, Shergill SS (2014) Uncertainty and confidence from
482 the triple-network perspective: voxel-based meta-analyses. Brain and cognition 85:191-200.
483 Zeidman P, Maguire EA (2016) Anterior hippocampus: the anatomy of perception, imagination and
484 episodic memory. Nat Rev Neurosci 17:173-182.

485

Figure 1. Experimental paradigm, choice behavior, and model estimations for trial-by-trial confidence and uncertainty. (**A**) Bead task. Participants chose from which of three jars, containing predominantly red, blue or green beads (80%-10%-10% ratio), the latest bead was extracted from. The latest five beads extracted were always displayed on screen, and no feedback was provided after each choice selection. (**B**) Card task. Participants chose one card among three cards characterized by different geometric figures. Each card yielded 100, 10 or 0 points (immediate feedback) with a probability distribution of 80%-10%-10%. The extraction jar and the card-value associations were changed every 5 to 7 trials, in three pre-established pseudo-random sequences that were used for all participants, in a counterbalanced order. (**C**) Subject-specific values of the coefficient σ, used in the Bayesian models to estimate confidence and uncertainty in each participant on a trial-by-trial basis. Different σ values determine different rates of belief update: for instance, the higher the value, the more confirming evidence is required to increase confidence and decrease uncertainty. The color of each extracted new bead or the feedback provided after each card selection were used as discrete new evidence to update the estimated probability distribution or priors of each participant. (**D-E**) The actual choices of all the participants (upper row, expressed in percentages) match the model-estimated probabilities (p) assigned to each available option (lower row, expressed as mean and standard deviation). In grey, superimposed on the participant choice selection percentages, mean and standard deviation of the reaction times, per task.

Figure 2. Model-based BOLD responses encoding Bayesian confidence and uncertainty. (A-B) Model-based fMRI response recorded across tasks in association with the model-estimated signal of uncertainty (left column) or confidence (right column), used as parametric modulator and using reaction times as covariates (P_{FWE}<0.05). (C-F) Task-specific, model-

based, fMRI responses in association with the model-estimated signals of uncertainty and confidence ($P_{uncorrected}$<0.005, k=50).

Figure 3. Neural network architectures tested in the family-wise DCM. These neural architectures illustrate the eight models used to test the effective connectivity associated with the signal of uncertainty (A) or confidence (B). The same models have been tested in both tasks for a comparison of the results. The general architecture (*A matrix*) remains unchanged, whereas the targets of the modulatory inputs (*B matrix*) are different in each model. This differentiation results in the generation of the 8 models, then divided into two families of 4, to allow family-wise comparisons for each pair of nodes in a triplet (Ins-ACC-dlPFC or aHip-Amg-mPFC).

Figure 4. DCM results using family-wise comparison. Histograms report the family Exceedance probability for each tested pair-wise connection in the confidence- and uncertainty-encoding networks. The network diagrams illustrate a summary of all pair-wise analysis, with thick red arrow used to highlight the direction of causal effects and thin red lines used for those connections for which DCM analysis did not provide a conclusive result. In the bead task, the insula drives the uncertainty-encoding network (**A**) and the anterior Hippocampus drives the confidence-encoding network (**B**). In the card task, the network estimating uncertainty remains unchanged in comparison with the bead task, despite the presence of the immediate feedback (**C**). Conversely, the estimation of confidence is driven in this task by the medial prefrontal cortex and amygdala (**D**).

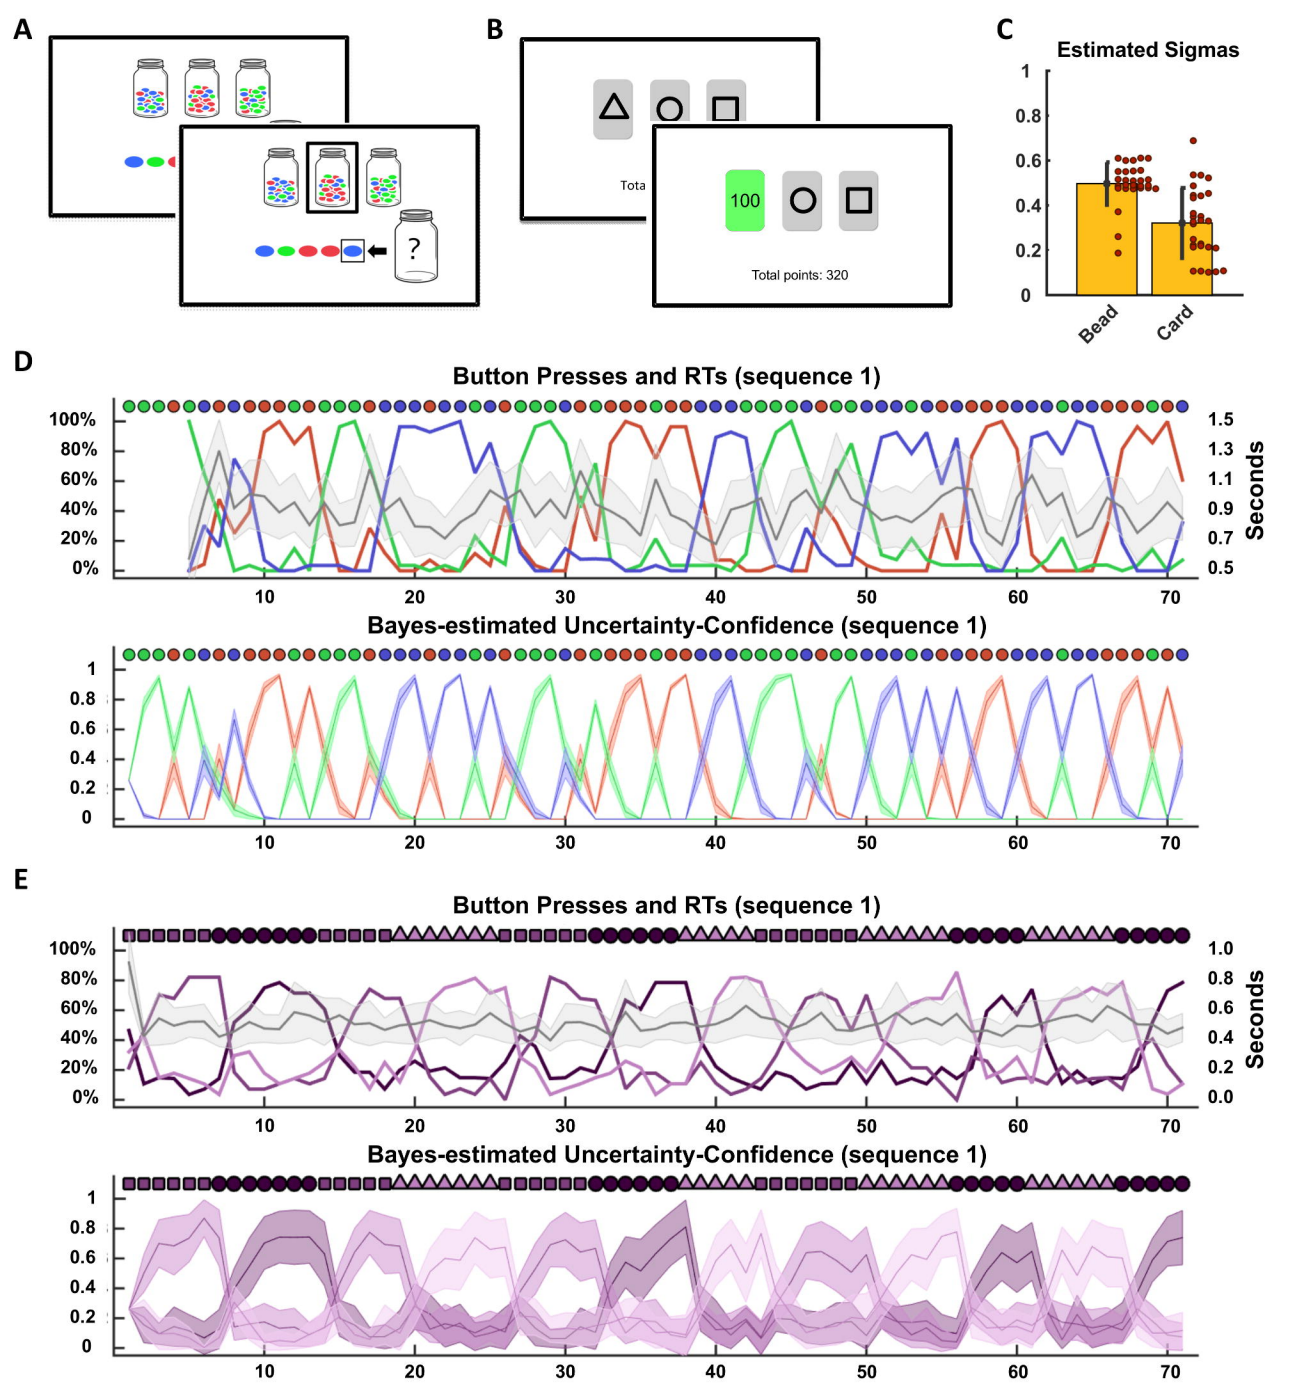
A
B
Tota
100
Total points: 320
C
Estimated Sigmas
Bead
Card
D
Button Presses and RTs (sequence 1)
Seconds
Bayes-estimated Uncertainty-Confidence (sequence 1)
E
Button Presses and RTs (sequence 1)
Seconds
Bayes-estimated Uncertainty-Confidence (sequence 1)

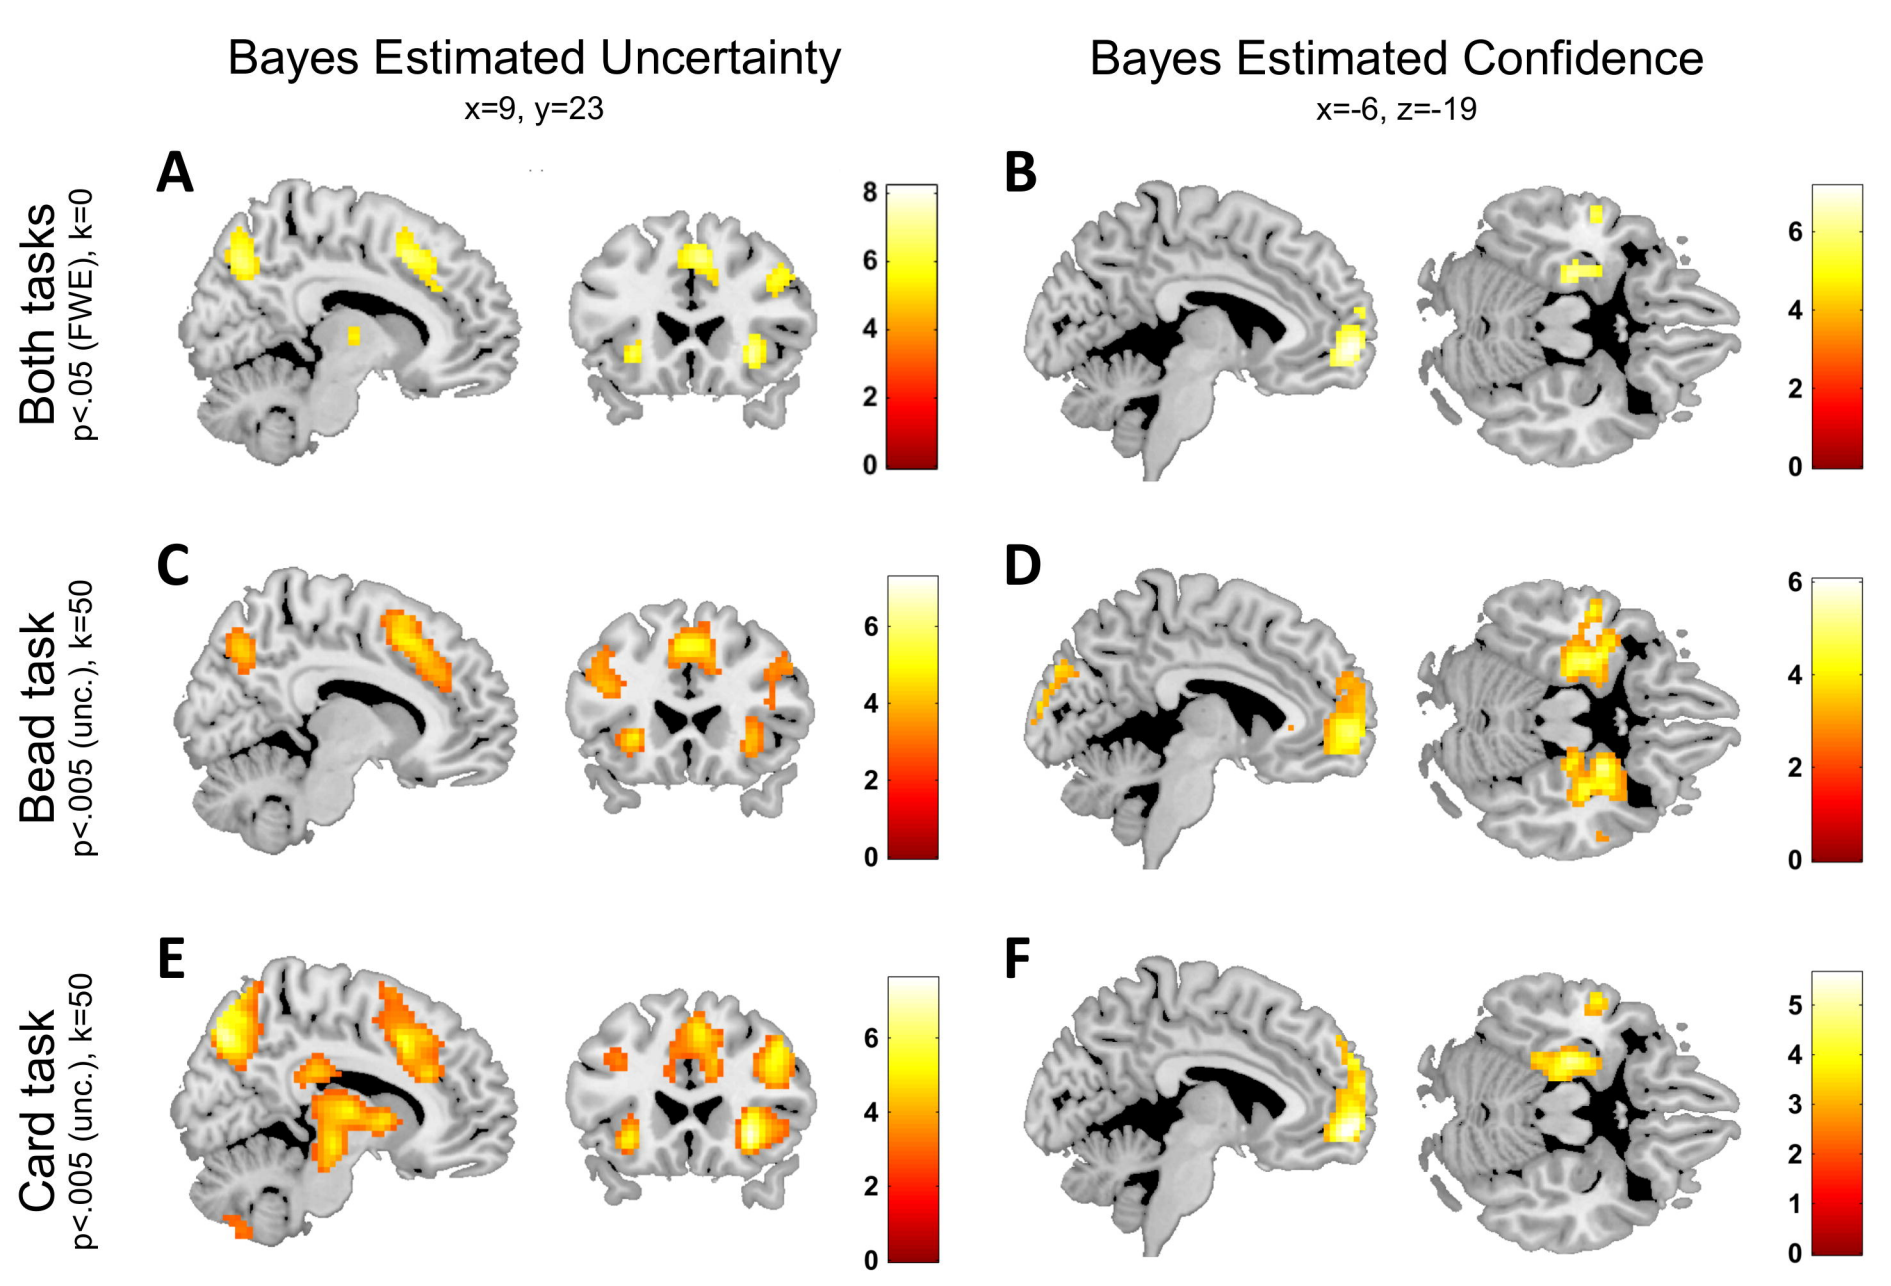
Bayes Estimated Uncertainty
x=9, y=23
Bayes Estimated Confidence
x=-6, z=-19
Both tasks
p<.05 (FWE), k=0
Bead task
p<.005 (unc.), k=50
Card task
p<.005 (unc.), k=50
A
B
C
D
E
F
8
6
4
2
0
5
3
1

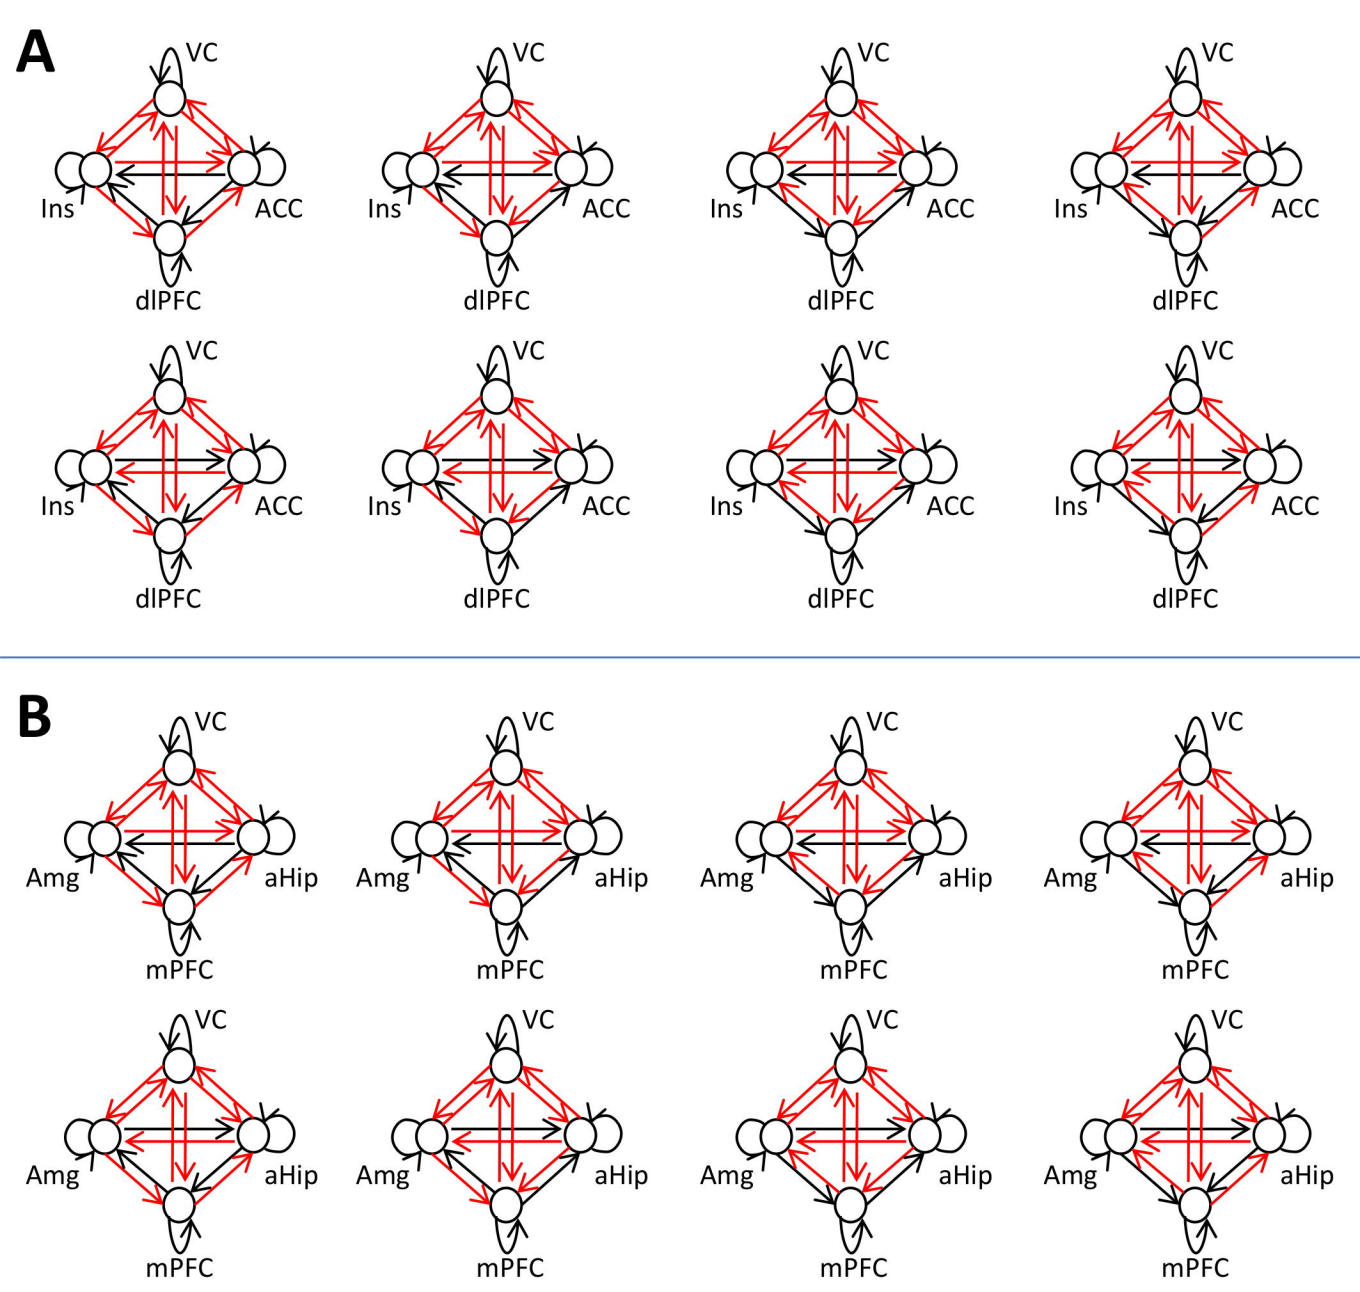
A
VC
Ins
ACC
dlPFC
B
VC
Amg
aHip
mPFC

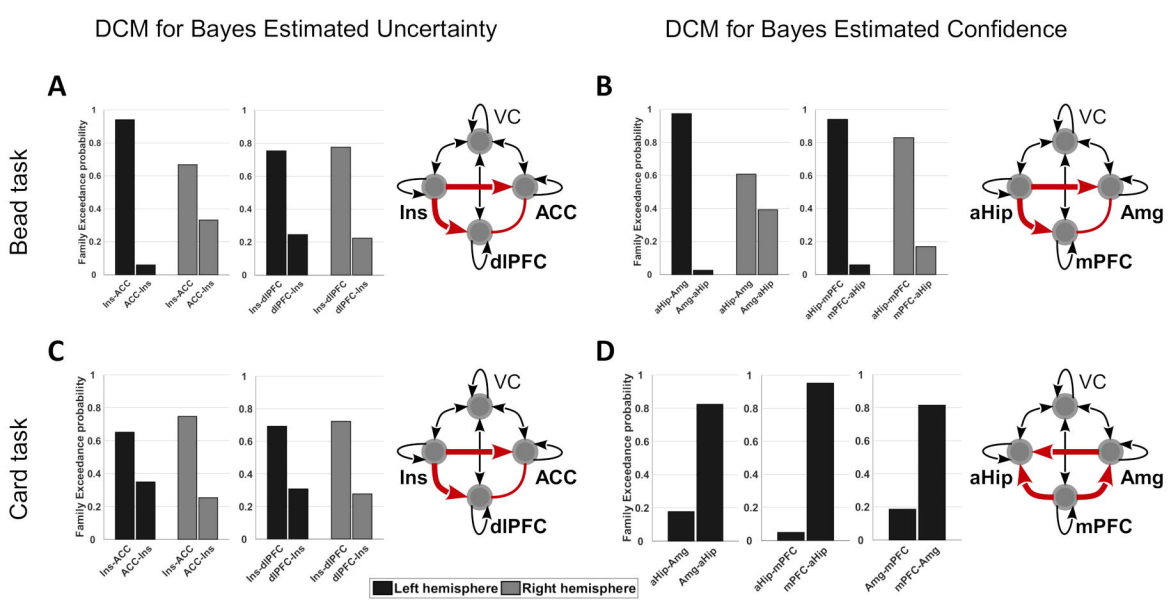
DCM for Bayes Estimated Uncertainty
DCM for Bayes Estimated Confidence
A
B
C
D
Bead task
Card task
Family Exceedance probability
Ins-ACC
ACC-Ins
Ins-dlPFC
dlPFC-Ins
aHip-Amg
Amg-aHip
aHip-mPFC
mPFC-aHip
Amg-mPFC
mPFC-Amg
VC
Ins
ACC
dlPFC
aHip
Amg
mPFC
Left hemisphere
Right hemisphere