

Journal Pre-proof

Bayesian modeling of temporal expectations in the human brain

Antonino Visalli, Mariagrazia Capizzi, Ettore Ambrosini, Ilaria Mazzonetto, Antonino Vallesi



PII: S1053-8119(19)30688-3

DOI: <https://doi.org/10.1016/j.neuroimage.2019.116097>

Reference: YNIMG 116097

To appear in: *NeuroImage*

Received Date: 1 May 2019

Revised Date: 16 July 2019

Accepted Date: 11 August 2019

Please cite this article as: Visalli, A., Capizzi, M., Ambrosini, E., Mazzonetto, I., Vallesi, A., Bayesian modeling of temporal expectations in the human brain, *NeuroImage* (2019), doi: <https://doi.org/10.1016/j.neuroimage.2019.116097>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2019 Published by Elsevier Inc.

Bayesian modeling of temporal expectations in the human brain

Antonino Visalli ^{a,b,*}, Mariagrazia Capizzi ^a, Ettore Ambrosini ^{b,c}, Ilaria Mazzonetto ^{a,d,1},

Antonino Vallesi ^{c,e}

^a *Department of Neuroscience, University of Padova, 35128 Padova, Italy*

^b *Department of General Psychology, University of Padova, 35131 Padova, Italy*

^c *Department of Neuroscience & Padova Neuroscience Center, University of Padova, 35131 Padova, Italy*

^d *Department of Information Engineering, University of Padova, 35131 Padova, Italy*

^e *Brain Imaging and Neural Dynamics Research Group, IRCCS San Camillo Hospital, 30126 Venice, Italy*

* Address for correspondence:

Antonino Visalli

Department of General Psychology, University of Padova

Via Venezia 8, 35131 Padova, Italy

Phone number: (+39) 049 8277424

Email: antonino.visalli@unipd.it

antonino.visalli.av@gmail.com

Note: ¹Dr. Ilaria Mazzonetto has moved to the Department of Developmental Psychology and Socialisation, University of Padova, Via Venezia 8, 35131 Padova, Italy.

Abstract

The brain predicts the timing of forthcoming events to optimize processes in response to them. Temporal predictions are driven by both our prior expectations on the likely timing of stimulus occurrence and the information conveyed by the passage of time. Specifically, such predictions can be described in terms of the hazard function, that is, the conditional probability that an event will occur, given it has not yet occurred. Events violating expectations cause surprise and often induce updating of prior expectations. While it is well-known that the brain is able to track the temporal hazard of event occurrence, the question of how prior temporal expectations are updated is still unsettled. Here we combined a Bayesian computational approach with brain imaging to map updating of temporal expectations in the human brain. Moreover, since updating is usually highly correlated with surprise, participants performed a task that allowed partial differentiation between the two processes. Results showed that updating and surprise differently modulated activity in areas belonging to two critical networks for cognitive control, the fronto-parietal (FPN) and the cingulo-opercular network (CON). Overall, these data provide a first computational characterization of the neural correlates associated with updating and surprise related to temporal expectation.

Keywords: Bayesian brain, fMRI, surprise, temporal prediction, updating

1 Introduction

¹The ability to generate accurate predictions about the timing of forthcoming events is essential to temporally optimize cognitive processes ranging from perception to action selection. Temporal predictions can be formally described in terms of the hazard function, that is, the conditional probability that an event will occur given it has not yet occurred (Janssen and Shadlen, 2005; Nobre and van Ede, 2018). Accordingly, temporal predictions depend on both prior expectations about the likely timing of events and the information conveyed by the elapse of time. To illustrate this, consider a common breakfast scenario in which you put your toast in the toaster. While waiting for the toast, you have some expectations about “when” the toast is more likely to jump out and, critically, such expectations grow over time. Previous reaction time (RT) experiments employing a foreperiod (FP; i.e., preparatory time interval) between warning signal and target provided compelling evidence that the brain shows an anticipatory activity that tracks the temporal hazard of target occurrence (Buetti et al., 2010; Herbst et al., 2018). However, how the brain forms and revises prior temporal expectations still remains an unsettled issue. Here, we took a computational approach in an fMRI experiment to address this fundamental question. Specifically, the Bayesian brain framework was applied to quantitatively describe belief updating about FP distributions.

According to the Bayesian brain hypothesis (Kersten et al., 2004; Knill and Pouget, 2004; Friston, 2005; Doya et al., 2007), the brain weighs current evidence (*likelihood*) on the basis of expectations about the environment (*prior* beliefs) and updates such beliefs into

¹ *Abbreviations:* aINS, anterior insula; CON, cingulo-opercular network; cPPI, correlational psychophysiological interaction; dACC, dorsal anterior cingulate cortex; D_{KL} , Kullback-Leibler divergence; EPI, echo-planar image; FWE, family-wise error; FD, frame-wise displacement; FDR, false discovery rate; FFG, fusiform gyrus; FP, foreperiod; FPN, fronto-parietal network; GLM, general linear model; IFG, inferior frontal gyrus; IPS, intra-parietal sulcus; I_s , Shannon information; ITI, inter-trial interval; LMM, linear mixed model; PCC, posterior cingulate cortex; PCun, precuneus; PM, parametric modulator; PPC, posterior parietal cortex; pre-SMA, pre-supplementary motor area; RT, reaction time; SD, standard deviation; TE, echo time; TR, repetition time; VOI, volume of interest.

posterior ones. Given an agent's beliefs, those events fulfilling our prior expectations can be predicted to optimize behavior. Conversely, those events violating our expectations are surprising, which commonly leads to behavioral costs and to an update of the internal model in order to improve future predictions. It is important to note, however, that a surprising observation does not always give rise to an update of predictive models (Itti and Baldi, 2009). Indeed, although surprise and updating are likely to co-occur (i.e., they are correlated), they reflect distinct cognitive processes that have started to be differentiated at the neural level only by a few recent studies (O'Reilly et al., 2013; Schwartenbeck et al., 2016; Kobayashi and Hsu, 2017).

To the best of our knowledge, no previous research has investigated in the same study the neural mechanisms associated with both updating of temporal expectations and the effect of temporally unexpected, surprising, events. To this end, in the present study we modeled our FP task after a previous spatial paradigm by O'Reilly and colleagues (2013), who investigated the neural correlates of updating and surprise about spatial locations. Briefly, the authors employed a task manipulation that allowed having a subset of surprising events not eliciting updating. Although updating always occurs in the presence of surprising events, their task was successful in identifying distinct brain areas associated with updating and surprise in the spatial domain.

In sum, the aim of the present study was to investigate the brain correlates of Bayesian updating about temporal expectations and to distinguish them from surprise correlates not involved in updating. To this end, we implemented an ideal Bayesian observer, which enabled us to capture participants' beliefs in terms of probability distributions and to model belief updating using Bayes' rule. The measures of surprise and updating obtained through

this model were then used as parametric explanatory variables of both behavioural and fMRI data to address our research questions.

2 Materials and Methods

2.1 Participants

The study included an initial sample of 26 participants, who were reimbursed 25 euros for their time. Data from two participants were discarded because of excessive head movements (see details on the MRI preprocessing section). Additionally, one participant was excluded due to falling asleep (11% of no responses) and another one due to low compliance with task instructions (the participant reported a change of strategy during the session that led to a lot of anticipations and to an overall accuracy of 49%). Therefore, the final sample comprised 22 participants² [10 females; mean age: 26 years (standard deviation, SD = 4 years), range: 20-34 years]. All of them were right-handed, as assessed with the Edinburgh Handedness Inventory (Oldfield, 1971) with an average score of 89.1 (SD = 11.7, range: 60-100), reported no history of neurological or psychiatric disorders, had normal color vision and normal or corrected-to-normal visual acuity (MRI-compatible glasses were used when appropriate). The procedures involved in this study were approved

² Based on a very recent study (Ostwald et al., 2019) presenting power functions for voxel- and cluster-level inference in the RFT-corrected multiple testing scenario, we can provide an estimation of the sensitivity to find significant cluster-level effects given our sample. Assuming an FWE-corrected significance level of $\alpha = .05$ and a minimal power of $1 - \beta = .80$, our sample size was large enough to detect a significant result with a medium-small effect size of $d = .35$ assuming a partial alternative hypothesis parameter of $\lambda = .1$ (or with an effect size of $d = .26$ assuming a partial alternative hypothesis parameter of $\lambda = .3$). Concerning behavioral analyses, it should be noted that the number of observations we had in our study (obs = 6.903) is well above that recommended for mixed-effects models in reaction time experiments (Brysbaert & Stevens, 2018), thus assuring that our study was properly powered.

by the Bioethical Committee of the Azienda Ospedaliera di Padova. Participants gave their written informed consent before the experiment, in accordance with the Declaration of Helsinki.

2.2 Task and procedure

As mentioned in the Introduction, we modeled our temporal preparation task after the spatial one developed by O'Reilly and colleagues (2013). The task was implemented in MATLAB (The MathWorks, Inc., Natick, Massachusetts, United States) using the PSYCHOPHYSICS TOOLBOX 3 (Brainard and Vision, 1997; Pelli, 1997; Kleiner et al., 2007). Each trial began with the presentation of an uninformative warning signal that consisted of a black fixation cross. The warning signal was displayed centrally against a gray background and remained on the screen for the whole FP duration for that trial. As soon as the FP elapsed, the warning signal was replaced by the target, which consisted of a colored circle (see below) with a diameter equal to the length of the cross arms, centrally presented for 1500 ms (Fig. 1A). Participants were required to respond to the onset of the target as quickly as possible by pressing a button of an MRI-compatible response box with their index finger. Half of the participants used their right hand and the other half their left hand. Within each task condition, the inter-trial interval (ITI; a blank screen) pseudo-randomly varied in order to approximate a Poisson distribution having a lambda of 2, which was shifted of 2 s to the right (i.e., in the range from 2 to 12 s).

In 80% of the trials, the FP duration was drawn from a Gaussian distribution (normal FP) with a mean and a standard deviation that remained fixed during a block, but that abruptly changed across blocks. Blocks were not temporally separated, such that the first trial of a new block followed directly the last trial of the previous block. In the remaining 20% of

trials, the FP duration was drawn from a uniform distribution (uniform FP) in the 200-3000 ms range. Consequently, the generative probability density function over FP duration (Fig. 1B-C) for each block was:

$$p(FP) = .80 p(FP|FP \sim \mathcal{N}(\mu, \sigma)) + .20 p(FP|FP \sim \mathcal{U}(200 \text{ ms}, 3000 \text{ ms})) \quad (1)$$

Importantly, participants were instructed to use the color of the target in order to distinguish the beginning of a new block from uniform trials. More specifically, in normal trials each target could be filled with one of four colors (vermillion, reddish purple, bluish green, and blue). A given color (i.e., blue) was kept constant over a block of trials and changed only when a new block started. Unlike normal trials, in uniform trials the target was always presented in white. In such a case only, participants were to respond to the target but with the explicit instruction that the temporal information associated with it (either earlier or later than the actual FP distribution in that block) was useless to anticipate the next target occurrence. In sharp contrast, a change in color signaled the beginning of a new block and prompted participants to use the temporal information from that trial to update their expectation and predict subsequent target onsets. Summarizing then, the color manipulation created three types of trials: “predictable” trials in which target onset in a given block could be easily predicted using the information from previous trials, “update” trials (signaled by a change in target color) in which the unpredicted FP led to strongly update the internal model in order to anticipate the next target onset, and “uniform” trials that did not require any updating of FP distribution despite their break of current temporal expectation.

The experiment was composed of 33 blocks with a total number of trials equal to 350. The length of a block was in the range of 7-13 trials (mean = 9.82, SD = 1.24). For each new block, the mean of the Gaussian distribution from which the normal FPs were extracted was at least 3 SDs away from the previous block in order to ensure a noticeable distribution change between blocks. Moreover, 16 blocks had an FP mean lower than the previous block and 16 blocks a mean higher than the previous one. Accordingly, half of the update trials had an FP duration shorter than the mean of the previous Gaussian distribution, while the other half had a longer FP duration. The same manipulation was applied to uniform trials such that they were divided into shorter and longer ones with respect to the mean of the current Gaussian distribution.

Overall, the 32 Gaussian distributions (not considering the first block) were derived from an orthogonal combination of 7 mean values (500, 800, 1100, 1400, 1700, 2000 and 2300 ms) and 4 SD values (20, 40, 60 and 80 ms). In total, there were 6 fMRI runs. Participants were informed that each new run started using the same FP distribution as that used on the previous block. At the beginning of a new run, the number of trials before the new block was in the range of 4-6 (mean = 5.2, SD = 0.8).

Before the fMRI session, participants practiced the task outside the MRI bore. They performed a shorter version of the experimental task comprising four blocks. In the first two blocks, participants were presented with normal FPs only in order to familiarize themselves with updating after a color change. In the subsequent blocks, we introduced uniform FPs and carefully explained the difference between them and normal ones. Participants were explicitly told that there was no relation between a given color and a given duration and to use the color just to distinguish uniform from update trials.




Figure 1 | Paradigm and schematic depiction of the experimental approach. (A) Each trial started with the presentation of an uninformative warning signal (fixation cross) that remained on the screen for a given foreperiod (FP) duration (as shown in panel B). The participant's task was to respond to the onset of a target (colored circle) that appeared after the FP elapsed. **(B)** Plot of FP duration over 30 trials. In most of the trials (i.e., 80%), the onset of the colored targets was predictable, because they occurred after an FP drawn from a Gaussian distribution (panel C) with a mean and standard deviation that remained fixed during each block of trials. Blocks were not temporally separated such that a change in the target color signaled the beginning of a different FP distribution than that used in the preceding block. On 20% of the trials, the FP duration was drawn from a uniform distribution in the 200-3000 ms range. The targets for these trials were unpredictable and were distinguished from the other trial types by their white color. Importantly, the white targets did not instantiate a new block but future target onsets followed the original FP distribution. **(C)** Plot of FP probability distributions within three blocks of trials from which the corresponding FP durations depicted in panel B were drawn.

2.3 Computational approach

2.3.1 Ideal Bayesian observer.

The model was an ideal Bayesian observer adapted from the one specifically designed by O'Reilly and colleagues (2013) for these kinds of learning tasks. As shown in Equation 1, FPs were drawn from a distribution that was a mixture of a Gaussian and a uniform distribution (i.e., generative model). The ideal Bayesian observer sequentially estimated the parameters of the Gaussian distribution that led to each FP duration on the basis of current and previous observations (i.e., inversion of the generative model). It is important to note that we are not assuming here that participants' updating is optimal or exclusively based on the information elaborated by the model (Waskom et al., 2017). Rather, the use of an ideal Bayesian observer represents "just a description of optimal behavior" (Friston, 2012), which allows formalizing hypotheses and making quantitative predictions about cognitive processes (Frank, 2013; see also Griffiths et al., 2012, for a thorough discussion on the use of Bayesian models).

The present model aims to iteratively infer the parameters μ and σ of the Gaussian distribution underlying normal FPs. After each new observation, it estimates the posterior probability for each possible pair of parameters μ and σ (i.e., the posterior probability over parameter space). The area of the employed parameter space was 300×15 , that is, the combination of all the means from 10 to 3000 ms and SD from 10 to 150 ms in steps of 10 ms.

Updating was differently computed according to the trial type. After *predictable* trials, updating was computed using Bayes' rule as:

$$p(FP \sim \mathcal{N}(\mu, \sigma) | FP_{1:n}, \text{predictable}) \propto p(FP_n | FP \sim \mathcal{N}(\mu, \sigma)) \times p(FP \sim \mathcal{N}(\mu, \sigma) | FP_{1:n-1}). \quad (2)$$

Thus, the posterior at trial n was proportional to the likelihood of the experienced FP duration (FP_n) times the prior at trial n , which is equivalent to the posterior at trial $n-1$. After *update* trials, the change of color explicitly signaled the start of a new distribution and, as a consequence, previous observations were no longer useful in estimating the posterior probability. For this reason, the prior in update trials was blanked with a uniform distribution, and updating was computed following Bayes' rule as:

$$p(FP \sim \mathcal{N}(\mu, \sigma) | FP_{1:n}, \text{update}) \propto p(FP_n | FP \sim \mathcal{N}(\mu, \sigma)) \times p(FP \sim \mathcal{U}(\mu, \sigma)). \quad (3)$$

According to the task instructions given to participants, no updating occurred after *uniform* trials³, such that the posterior probability over parameter space at trial n was derived from the prior without modifications:

$$p(FP \sim \mathcal{N}(\mu, \sigma) | FP_{1:n}, \text{uniform}) = p(FP \sim \mathcal{N}(\mu, \sigma) | FP_{1:n-1}). \quad (4)$$

The model then translated the estimates of the parameters μ and σ into probability density functions over time. Specifically, the prior over time for a subsequent trial $n + 1$ was derived from the posterior over parameter space on trial n as follows:

³ As already mentioned, the assumption that no updating occurred after uniform trials is “just a description of optimal behavior”. Accordingly, it is hard to completely exclude the possibility that participants might have implicitly used the information of uniform trials to update their beliefs. If this were the case, participants should have shown some behavioral costs on *predictable* trials following a *uniform* trial, because such trials are indeed highly predictable only refraining from updating on uniform trials. In contrast to this hypothesis, we did not find significant evidence of updating on uniform trials (Supplementary Fig. S1).

$$\begin{aligned}
p(FP_{n+1}|FP_{1:n}) = & p(predictable_{nb+1}) \sum_{\mu_{n+1}, \sigma_{n+1}} (FP_{n+1}|FP_{n+1} \sim \mathcal{N}(\mu_{n+1}, \sigma_{n+1})) \\
& \times p(FP_{n+1} \sim \mathcal{N}(\mu_{n+1}, \sigma_{n+1})|FP_{1:n}) + p(uniform_{nb+1} \cup update_{nb+1}) \\
& p(FP_{n+1}|\mathcal{U}(10\text{ ms}, 3000\text{ ms}),
\end{aligned} \tag{5}$$

where $p(predictable_{nb+1})$ and $p(uniform_{nb+1} + update_{nb+1})$ represent the probability of incurring, respectively, in a predictable or in a uniform/update trial at the next trial of the current block (nb indicates the trial number within a block, which differs from n that indicates trial number referred to the whole task). For simplicity, the combined probability to have an uniform or an update trial was set to the true proportion of those trial types at a given trial within a block, smoothed using a moving average in order to have a monotonic increase in the probability of having an update trial (e.g., the probability of having an update/uniform trial on $nb+1=14$ was higher than on $nb+1=13$ and so on, thus, with the lowest probability at $tb+1=2$, that is after the color change). The probability of a predictable trial was equal to $1-p(uniform_{nb+1} \cup update_{nb+1})$. The output of the model is presented in Supplementary Figure S2A.

2.3.2 Model-based measures of updating and surprise.

Two measures from information theory were used to formally quantify updating of temporal expectations and the surprise of observing the target after a specific FP. These measures were computed with reference to our model as follows.

Updating. Following Itti and Baldi (2009), we quantified the updating of the internal predictive model as the Kullback-Leibler divergence (D_{KL} ; Supplementary Fig. S2B) between prior and posterior on trial n :

$$D_{KL}(FP_n) = \sum_{FP=10}^{3000} p(FP_n|FP_{1:n-1}) \log \frac{p(FP_n|FP_{1:n-1})}{p(FP_{n+1}|FP_{1:n})}, \quad (6)$$

where $p(FP_n|FP_{1:n-1})$ and $p(FP_{n+1}|FP_{1:n})$ represent, respectively, the prior and posterior probability of FP durations in the range from 10 to 3000 ms in steps of 10 ms.

Surprise. Since during the trial the prior probability of target onset changed as a function of the elapse of time (Janssen and Shadlen, 2005), we quantified surprise at trial n as the Shannon information (I_S ; Supplementary Fig. S2C) associated with the value of the hazard function at target onset:

$$I_S(FP_n) = -\log h(FP_n). \quad (7)$$

The value of the hazard function at target onset at trial n , $h(FP_n)$, was computed as:

$$h(FP_n) = \frac{f(FP_n)}{1 - F(FP_n)}, \quad (8)$$

where $f(FP_n)$ is the prior probability of target occurrence after the FP duration at trial n , $p(FP_n|FP_{1:n-1})$, $F(FP_n)$ is the cumulative probability $p(FP \leq FP_n)$, and $1 - F(FP_n)$, represents the probability that the target had not yet occurred after the FP at trial n .

The correlation between D_{KL} and I_S was $r = 0.34$ ($R^2 = .12$), the correlation between D_{KL} and h was $r = -0.26$ ($R^2 = .07$), and the correlation between I_S and h was $r = -0.79$ ($R^2 = .63$).

2.4 Behavioral data analysis.

Data from error trials (i.e., anticipated or missing responses) and post-error trials were excluded from analysis (mean excluded trials: 6%; SD: 5%). Reaction times (RTs) were log-transformed to mitigate the influence of non-normally distributed and skewed data. Following the procedure proposed by Baayen and Milin (2010), log-transformed RTs were analyzed by conducting Linear Mixed Models (LMM) using the *lme4* library (Bates et al., 2015) in R (R Core Team, 2015, R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>). We investigated the behavioral correlates of surprise and updating by using I_S and D_{KL} as regressors of interest. A full LMM was specified as follows: I_S and D_{KL} (and their interaction) as well as TRIAL, which represents the rank-order of a trial, and log-RT at the preceding trial (PRECEDING RT), were entered as fixed-effects predictors. The random structure included correlated by-participant (ID) random intercepts and slopes for TRIAL, PRECEDING RT, I_S and D_{KL} . All these continuous predictors were standardized using Z-score in order to facilitate model convergence. The variables TRIAL and PRECEDING RT were included to control for the temporal dependencies that generally need to be addressed when modeling RTs (Baayen and Milin, 2010). Specifically, TRIAL was included to capture possible effects of learning and fatigue, while PRECEDING RT was used to take into account the RT autocorrelation between subsequent trials. Using the function *step* from the *lmerTest* library (Kuznetsova et al., 2017), a stepwise variable selection was used to perform backwards elimination of non-significant random and fixed effects of the full LMM.

2.5 Neuroimaging analysis

2.5.1 MRI data acquisition.

Structural and functional images were acquired using a 3T Ingenia Philips whole body scanner (Philips Medical Systems, Best, The Netherlands) equipped with a 32-channel head-coil, at the Neuroradiology Unit of the University Hospital of Padova, Italy. Functional data were obtained using a whole head T2-weighted echo-planar image (EPI) sequences (repetition time, TR: 2000 ms; echo time, TE: 30 ms; 39 axial slices with ascending acquisition; voxel size: $3 \times 3 \times 3$ mm; flip angle, FA: 76° ; acquisition matrix: 84×84).

Excluding the four dummy scans for stabilization of the T1-saturation effect, the functional acquisitions consisted of 8 minutes of resting state activity, whose analysis is not presented here, followed by a total of 39.4 minutes of task-related activity. To correct fMRI images for spatial distortion, two spin echo EPI images with reversed phase encoding directions were acquired at the beginning of each of the six runs. These images were geometrically matched (same field of view and voxel size) with the functional images (Glasser et al., 2013). After the functional session, high resolution T1- and T2- weighted anatomical images were acquired (T1w: TR/TE: 8.10/3.72 ms; 180 sagittal slices; voxel size: $1 \times 1 \times 1$ mm; FA: 8° ; acquisition matrix: 256×256 ; T2w: TR/TE: 2500/249 ms; 180 sagittal slices; voxel size: $0.97 \times 0.97 \times 1$ mm; FA: 90° ; acquisition matrix: 256×256). In order to avoid head movement during scanning, small foam cushions and sponge pads were placed around the participant's head. Participants also wore earplugs to reduce acoustic noise.

2.5.2 MRI preprocessing.

First, spatial distortion of functional data was corrected using the susceptibility-induced off-resonance field estimated from the two oppositely phase-encoded spin echo EPI images as

implemented in the FSL (FMRIB Software Library, version 5.0.7; Smith et al., 2004; Jenkinson et al., 2012) toolbox “topup” (Andersson et al., 2003; Smith et al., 2004). This step improves the following coregistration step between fMRI and structural image (Glasser et al., 2013). Functional data were then slice-timing corrected using the middle slice as the reference frame, rigidly realigned to the first volume and spatially smoothed using a Gaussian kernel with a full-width at half-maximum (FWHM) of 5 mm using SPM12 (Statistical Parametric Mapping software; Wellcome Department of Cognitive Neurology, London, UK; <http://www.fil.ion.ucl.ac.uk/spm>). Participant’s head movements were quantified by means of framewise displacement (FD) index, which represents the sum of the absolute values of the derivatives of the translational and rotational realignment parameters. Before calculating FD, the estimated rotations were converted to displacements on the surface of a sphere of radius equal to 50 mm (Power et al., 2012). Participants with a mean FD above two SDs from the mean of all the participants (group mean = 0.05 mm, standard deviation = 0.02 mm) were excluded. The deformation field that mapped the individual functional data to standard Montreal Neurological Institute (MNI) template was derived combining several steps, all implemented in FSL. First, T1-weighted anatomical image was bias-field corrected and a non-linear transformation to MNI template was estimated (T1>MNI). To coregister the functional data to the T1-weighted anatomical image, a T2-weighted anatomical image was used as an intermediate target. Indeed, it has the same acquisition modality of fMRI data, but the same high-resolution with clear region boundary contours of T1-weighted anatomical image. Both T2- and T1- weighted structural images were skull-stripped and then a 6-parameter transformation from the former to the latter was computed (T2>T1). At the end, a 12-parameter affine transformation from the first volume of the functional data to the T2-weighted skull-stripped anatomical image was estimated (fMRI>T2) and combined

with the T2>T1 and T1>MNI transformations. The resulting transformation was then used to map the results obtained at the individual level in the functional space to the MNI template.

2.5.3 Whole-brain fMRI analysis.

Statistical analyses were carried out using SPM12. For each participant, first-level analysis was performed into the subject space (i.e., not normalized) using two general linear models (GLMs). For each GLM, the task was modeled with three regressors that corresponded to the main effect of the FP duration, the main effect of target onset, and either D_{KL} (GLM 1) or I_S (GLM 2) as parametric modulator (PM) of target onset. Specifically, the main effect of the FP was modeled as a boxcar starting from the onset of the warning signal and with duration equal to the FP length; the main effect of target onset was modeled as a delta function at the target onset and it was modulated by the model-based PM, D_{KL} or I_S . All these regressors were convolved with the hemodynamic response function. As in the behavioral analysis, the PM was standardized using Z-score and orthogonalized with regard to target onset.

Estimates of head movements were also included as six additional regressors of no interest. Volumes with FD exceeding a threshold of 0.5 were modeled out adding a single time point regressor for each volume of high movement (Siegel et al., 2013). Slow signal drifts were removed using a 128 s high-pass filter. For each participant and each GLM, a t -contrast was computed for each PM versus zero (i.e., baseline). At the group level, individual participants' contrast maps (i.e., the weighted sum of related beta, where weights are the values in the corresponding t -contrast) were normalized to MNI template as described in the previous section. Then, for each GLM, group-level maps were generated with random-effect models using normalized participants' contrast maps. Group statistics were assessed for cluster-wise significance using a cluster-defining threshold of $p < .001$ and a cluster significance

threshold of $p < .05$ corrected for family-wise error (FWE). A third GLM (GLM 3) with the value of the hazard function at the target onset (i.e., $h(FP_n)$) as PM was run in order to compare our results with previous findings on temporal hazard in the human brain (e.g., Buetti et al., 2010). Again, first-level contrast maps derived from the t -contrast of the PM against zero were normalized to MNI and employed in the subsequent group-level random-effect analysis. Finally, we ran an analysis with a first-level GLM including both D_{KL} and I_S as PMs (GLM 4). Since the task manipulation employed in our study allowed us to have surprising trials with no updating, I_S was orthogonalized with respect to D_{KL} (Supplementary Fig. S3B) in order to obtain unique variance associated with surprise. By contrast, updating is driven by the surprise of the events. Forcing D_{KL} to be orthogonal to I_S would then lead to a spurious anticorrelation between them. This can be appreciated in Supplementary Figure S3A, which shows that after orthogonalization, update trials had the lowest D_{KL} values. Therefore, running a model with D_{KL} orthogonal to I_S would not actually test the neural correlates of updating, but the opposite hypothesis, that is, which regions show enhanced BOLD signal with low levels of updating.

GLM 4 was run in the MNI space to facilitate the extraction of the volumes of interest (VOIs) used in the functional connectivity analysis described in the next section. First-level contrast maps derived from the t -contrast of each PM against zero were used to generate group-level maps with random-effect models.

2.5.4 Functional connectivity analysis.

Task-related functional connectivity analysis was included after inspection of the whole-brain results with the aim of further investigating the pattern of results (see section 3.2). This analysis was computed using the correlational psychophysiological interaction (cPPI)

toolbox (Fornito et al., 2012). In classical PPI analyses (Friston et al., 1997), the activity time course from a specific seed region is extracted and multiplied by a task regressor of interest to isolate task-specific modulations in the functional coupling between the seed region and other brain regions. This is a regression-based approach, in the sense that for each pair of time series the seed region activity is used as a predictor of the activity in the other regions. This implies that PPI is a form of effective connectivity (Friston et al., 1997), which makes it a suitable tool to test clear hypotheses about which region modulates activity in other regions (Fornito et al., 2012). Since we had no specific predictions, we hence employed the cPPI approach, which provides a measure of functional connectivity that does not require directional assumptions. Briefly, for any given pair of brain regions their time course is multiplied by the task-regressor to obtain two PPI terms. Then, the cPPI uses partial correlation between the two PPI terms in order to isolate covariations in task-related modulations of functional connectivity as distinct from intrinsic task-unrelated connectivity noise, and coactivation effects induced by the task. Accordingly, starting from a set of regions the cPPI analysis returns a functional connectivity matrix of pair-wise correlations due to task-specific modulations of neural activity. We estimated two cPPI correlation matrices, one for each of the two model-based regressors, D_{KL} and I_S . For each participant, coordinates of the volumes of interest (VOIs) were defined by generating 6-mm-radius spheres. These were centered on the nearest local maximum of the group maximum peaks observed for D_{KL} and I_S in the GLM 4. When the obtained two peaks for D_{KL} and I_S were very close (< 6 mm), a VOI centered at the midpoint between the two peaks was created instead of having two largely overlapping VOIs. The time course from each VOI in each participant was extracted using the *spm_regions* function in SPM and corrected for the same movement parameters as in the whole-brain analysis prior to deconvolution. Coordinates of

each VOI are presented in Supplementary Table S1. For each matrix, the correlation between pairs of VOIs was estimated after partialling out coactivation induced by the main effect of all the task regressors, as well as the main effect of neural activity in the remaining VOIs. Two-tailed one-sample t -tests were computed to test the significance of D_{KL} - and I_S -related partial correlations (p -values were adjusted for multiple-testing by controlling the false discovery rate, FDR, at the .05 level).

3 Results

3.1 Behavioral results.

Log-transformed RTs were analyzed by means of a linear mixed model (LMM) in which I_S and D_{KL} were used as explanatory variables along with the rank-order of a trial ($TRIAL$), and log-RT at the preceding trial ($PRECEDING\ RT$) to control for trial-by-trial dependencies (Baayen and Milin, 2010). Backward elimination of non-significant effects resulted in a model specified as the following *lme4*-notation formula:

$$\log(RT) \sim TRIAL + PRECEDING\ RT + I_S * D_{KL} + (TRIAL + PRECEDING\ RT + I_S \mid ID). \quad (9)$$

Visual inspection of the residuals showed that the model was a bit stressed. As suggested by Baayen and Milin (2010), trials with absolute standardized residuals higher than 2.5 SD were considered outliers and removed (2.4% of the trials). After outlier trials removal, the model was refitted achieving reasonable closeness to normality. The R^2 of the final model was .39. Table 1 shows the statistical results of the type III ANOVA. A significant interaction was

found between I_S and D_{KL} . Figure 2 shows that RTs increased with increasing surprise (I_S) and this effect was slightly augmented with high D_{KL} in a multiplicative way.

Table 1 | Analysis of variance with type-III sums of squares.

Fixed Effects	Sum Sq	Num. df	Den. df	<i>F</i>	<i>p</i>	β
TRIAL	0.39	1	21.0	10.0	.005	0.07
PRECEDING RT	0.66	1	21.0	17.0	< .001	0.10
I_S	3.70	1	22.1	94.8	< .001	0.25
D_{KL}	0.33	1	6819.8	8.6	.003	-0.05
$I_S : D_{KL}$	0.99	1	6823.5	25.5	< .001	0.09

Notes: *F*-statistics and associated *p*-values were calculated using Kenward-Roger's approximation of degrees of freedom (Halekoh and Højsgaard, 2014). Additionally, standardized regression coefficients (β) are shown.




Figure 2 | Interaction plot for log-transformed RTs. The plot shows the effect of surprise (I_s) for update trials (black line) and no-update (i.e., uniform and predictable; grey line) trials. Fitted log-RT values were extracted from the final model using the *fitted.merMod* function in the *lme4* library. Then, separately for each trial type (i.e., update and no-update), they were averaged by I_s values across participants (black dots and grey triangles represent average I_s values for update and no-update trials, respectively).

3.2 Neuroimaging results.

We investigated the neural correlates associated with updating and surprise across the whole brain by means of GLMs in which our information theoretical measures were used as parametric modulators of target onset events. Table 2 reports the results of GLM 1 and GLM 2, while Supplementary Table S2 reports results for GLM 3 (see section 2.5.3).

Updating significantly modulated activity in a set of lateral frontal and parietal regions, as well as in medial parietal regions and in a cluster around the left fusiform gyrus (FFG) and the left cerebellum (Fig. 3A). Lateral frontal and parietal regions included bilateral inferior

frontal gyri (IFGs) and posterior parietal cortices (PPCs) around the intra-parietal sulci (IPSs). Medial parietal activations included the posterior cingulate cortex (PCC) and the precuneus (PCun).

Surprise significantly modulated activity in the right IFG, bilateral PPCs around the IPSs, dorsal anterior cingulate cortex (dACC) including the pre-supplementary motor area (pre-SMA), bilateral anterior insula (aINS), and left FFG (Fig. 3B). Surprise adjusted for updating (GLM 4) modulated activity in the same regions with the exception of the right IFG, which was not significantly activated (see Supplementary material, Fig. S4 for a map of R^2 averaged across participants and Table S2, for the mean and the median R^2 of our VOIs).

Table 2 | Significant cluster activations in SPM analyses.

Anatomical region	MNI			Peak Z	Cluster level	
	x	y	z		p	Size
GLM1. Regions modulated by updating (D_{KL})						
L. Fusiform Gyrus	-42	-60	-12	4.52	.001	240
	-38	-56	-6	4.20		
	-36	-62	-34	4.09		
L. Posterior Parietal Cortex	-34	-64	44	4.37	< .001	801
	-48	-44	48	4.08		
	-44	-64	44	3.99		
R. Inf. Frontal Gyrus	50	20	28	4.33	< .001	315
	44	28	20	4.06		
	36	26	20	4.03		
R. Posterior Parietal Cortex	34	-62	34	4.29	< .001	371
	34	-66	46	4.14		
	26	-62	36	4.08		
L. Inf. Frontal Gyrus	-40	4	30	4.08	.016	142
Precuneus	-8	-68	46	3.93	.012	152
	4	-68	46	3.73		
	12	-68	44	3.58		
Post. Cingulate Cortex	-2	-34	26	3.74	.043	112
GLM 2. Regions modulated by surprise (I_S)						
L. Ant. Insula	-36	22	4	5.36	< .001	515

	-38	14	6	4.59		
	-32	26	-2	4.13		
L. Posterior Parietal Cortex	-34	-60	46	5.26	< .001	932
	-40	-42	46	4.42		
	-36	-42	38	4.08		
R. Ant. Insula	34	24	2	4.90	.003	206
Dorsal Ant. Cingulate Cortex	-4	10	48	4.61	< .001	342
	-8	-4	64	3.65		
	10	16	40	3.26		
R. Posterior Parietal Cortex	34	-68	46	4.50	< .001	356
	30	-62	38	4.03		
	30	-48	44	3.75		
L. Fusiform Gyrus	-38	-60	-12	4.49	.008	170
R. Inf. Frontal Gyrus	40	26	20	4.22	.003	201

Note: L and R. indicate left and right.

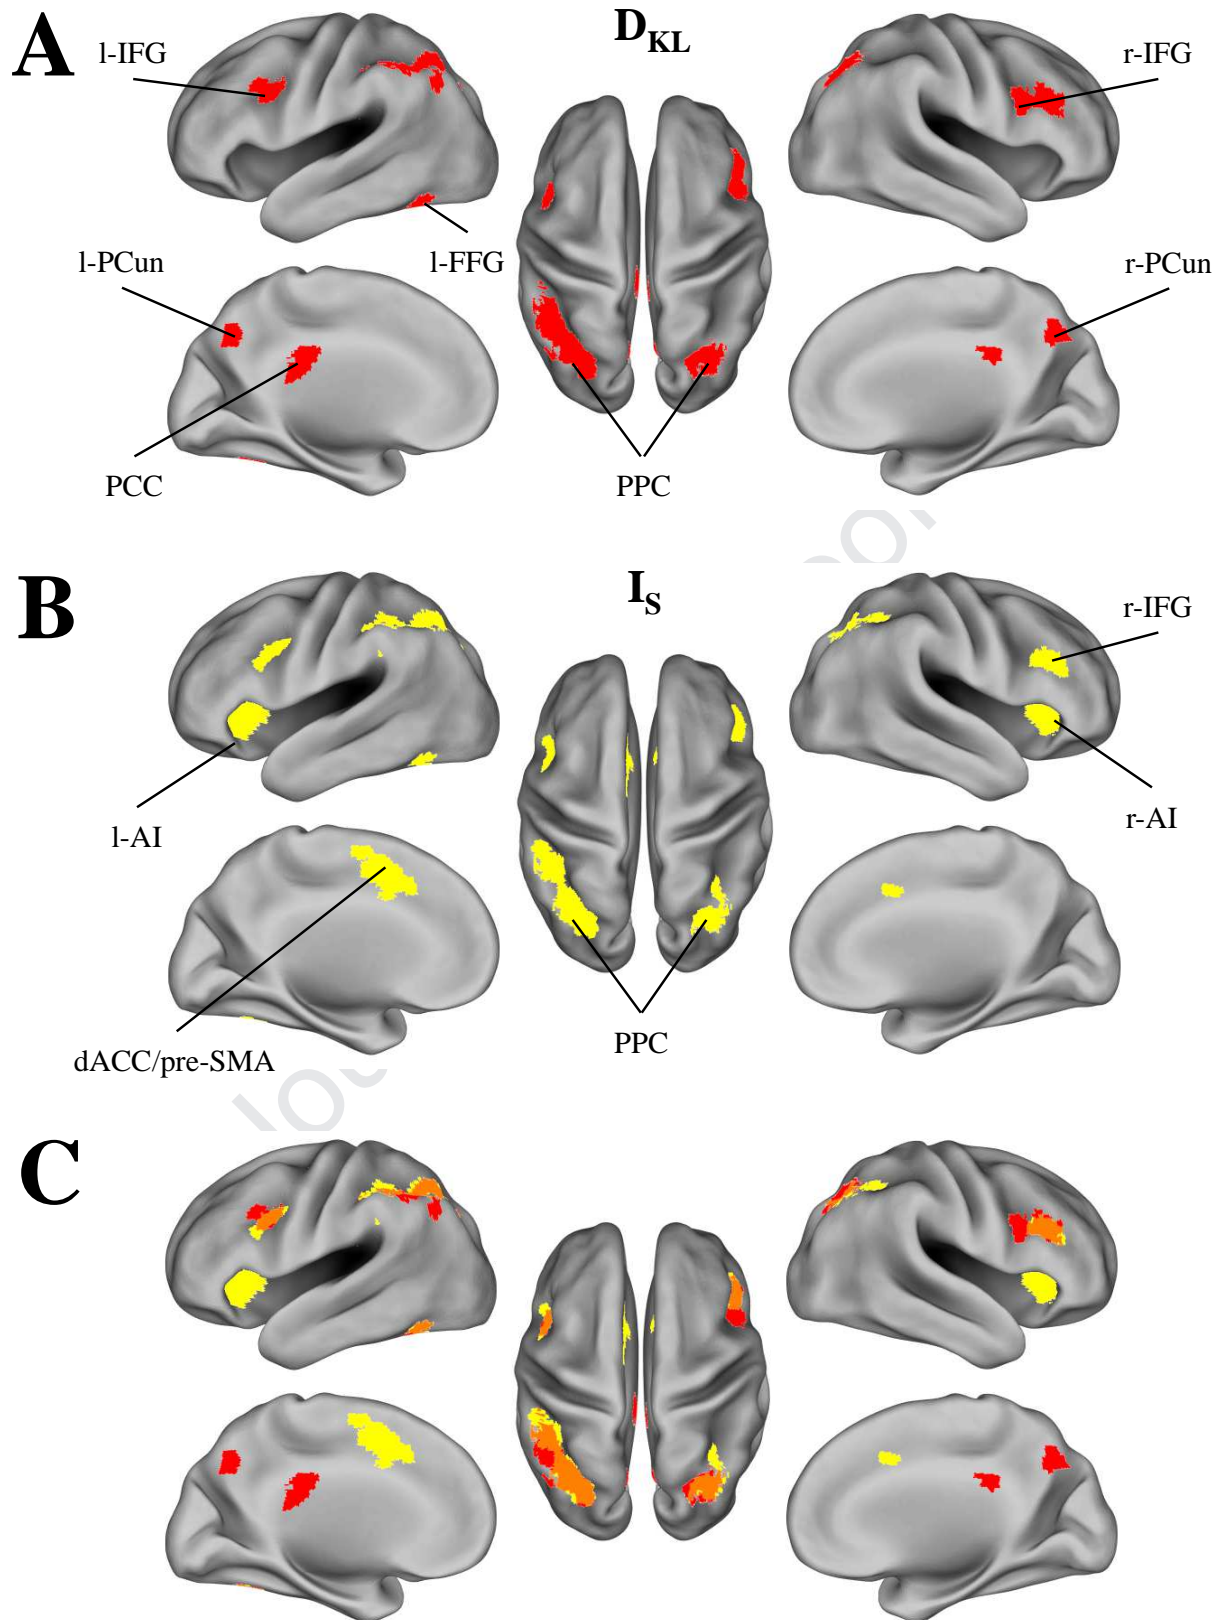


Figure 3 | Whole-brain fMRI analysis results. (A) GLM 1: Regions significantly modulated by updating (D_{KL}). **(B)** GLM2: Regions significantly modulated by surprise (I_S). **(C)** Overlapping of activation between D_{KL} and I_S . Abbreviations: IFG:

inferior frontal gyrus; PPC: posterior parietal cortex; PCun: precuneus; PCC: posterior cingulate cortex; FFG: fusiform gyrus; AI: anterior insula; dACC: dorsal anterior cingulate cortex; pre-SMA: pre-supplementary motor area.

With the exception of the FFG, all the clusters modulated by updating and surprise represent nodes of the two cognitive-control networks described in Dosenbach et al. (2008), namely, the fronto-parietal (FPN) and cingulo-opercular (CON) networks (FPN: bilateral PPC, PCun, PCC, and bilateral IFG; CON: bilateral AI and dACC/pre-SMA). Consequently, we employed a correlation psychophysiological interaction (cPPI) approach to explore how functional connectivity between these nodes was modulated by updating and surprise. The cPPI results are shown in Figure 4. Concerning updating, we observed a D_{KL} -related increase in functional connectivity in the FPN, mainly between posterior nodes (i.e., bilateral PPC, PCun, and PCC). Conversely, D_{KL} -related changes in functional connectivity within the CON (i.e., between INS and dACC/pre-SMA) as well as between-networks connectivity were very weak. Regarding the cPPI results on surprise, we observed a more distributed I_S -related increase in within- and between-networks functional connectivity. The results from the paired-sample t -tests (computed to compare cPPI correlations between D_{KL} and I_S) showed a slightly higher I_S -related functional connectivity, compared to D_{KL} , between the dACC/pre-SMA and nodes belonging to both CON (i.e., right INS) and FPN (i.e., left PPC nodes and PCun). In sum, we observed higher within-CON and between-network connectivity induced by I_S compared to D_{KL} .

To compare our results with previous literature on temporal preparation, we also performed a whole-brain analysis in which target onset events were parametrically modulated by the value of the hazard function at target onset (GLM 3). The results showed that the temporal hazard correlated with activity in bilateral lingual cortex, cuneus, and bilateral superior temporal gyrus (Supplementary Fig. S5). These findings corroborate

previous studies showing the involvement of sensory visual areas in shadowing temporal hazard (Vallesi et al., 2009; Buetti et al., 2010).




Figure 4 | cPPI analysis results. (A and B) The plots show task-related functional connectivity for updating (A) and surprise (B), respectively. The line width is (non-linearly) proportional to the t -value obtained from a two-tailed one-sample t -test against 0 (Bonferroni correction for 90 multiple comparisons). Non-significant correlations (p -value $> .05$) are not visible. Dashed and solid lines indicate significant correlations with p -values $> .001$ and $< .001$, respectively. Red nodes indicate regions that in the whole-brain univariate analyses were significantly modulated by D_{KL} in panel A and by I_S in panel B (cfr. Fig. 3), while nodes non-significantly modulated are depicted in orange. (C) The plot shows the results of the paired-sample t -test contrasting I_S against D_{KL} . Non-significant contrasts are not visible. The line width is (non-linearly) proportional to the obtained t -value.

4 Discussion

The present study tackled a fundamental aspect of temporal preparation often raised in the literature but never directly investigated so far, that is, the creation and revision of prior temporal expectations by the brain. To fill this gap, we combined a Bayesian computational approach with fMRI to investigate the neural correlates of temporal belief updating. Importantly, since (Bayes-optimal) belief updating takes place after events violating our prior expectations, we experimentally decoupled surprise from updating in order to tease

apart the neural processes involved in updating of temporal expectations from those facing surprising, poor temporally predicted, target onsets.

Consistently with O'Reilly and colleagues (2013), the behavioral results showed a strong effect of surprise on participants' RTs. Since RTs are a reliable index of information encoding and cognitive effort (Teichner and Krebs, 1974), this finding is in line with information-theoretic accounts of cognition, which assume that unexpected events are effortful to encode (see Zénon et al., 2019, for a recent overview). Furthermore, the significant effect of surprise on RTs indicates that those events that were surprising according to the ideal Bayesian observer were surprising to participants as well, confirming the ability of our model in providing an estimate of participants' expectations. We also found a behavioral cost due to updating that was larger as surprise increased. This finding extends current evidence in revealing that revising prior expectations might have a further cost on information processing that adds up to the cost due to stimulus encoding (i.e., surprise-related cost). Moreover, the fact that updating-related costs add up to surprise-related ones in a multiplicative way indicates that, compared to stimulus encoding, updating becomes harder the higher the violation of prior expectations is. Overall, the distinction between updating and surprise suggests that there could be an additional information cost in cognitive processing (Zénon et al., 2019) due to updating.

The fMRI results confirmed that our design differentiated between updating and surprise. Indeed, albeit there were some commonalities between the two, updating and surprise modulated distinct brain regions belonging to two well-known functional networks for cognitive control: the fronto-parietal network (FPN; Dosenbach et al., 2008) and the cingulo-opercular network (CON; Dosenbach et al., 2008; Menon, 2015).

Updating (D_{KL}) of temporal expectations was correlated mainly with FPN regions such as the posterior parietal cortex (PPC), precuneus (PCun), posterior cingulate cortex (PCC) and inferior frontal gyri (IFG). According to recent models of cognitive control (Dosenbach et al., 2008; Cocchi et al., 2013; Crittenden et al., 2016), the FPN is involved in the trial-by-trial adjustment of task-relevant information in order to implement top-down control. This property of the FPN is thus highly compatible with the idea of a Bayesian brain that updates its predictive models after each new observation is encountered.

A confirmation of the link between FPN and belief updating is also provided by several fMRI studies that have modeled belief updating about context information in domains other than time (Gläscher et al., 2010; Schwartenbeck et al., 2016; Kobayashi and Hsu, 2017; Waskom et al., 2017). As an example, Waskom et al. (2017) reported that the bilateral inferior frontal sulcus, bilateral intra-parietal sulcus (IPS), PCun and PCC responded to prediction error - derived from an ideal Bayesian observer – about the trial-relevant dimension, either color or motion direction, in a context-dependent perceptual discrimination task. Although prediction error is a measure more similar to I_S than D_{KL} , the fact that Waskom and colleagues' task did not distinguish between updating and surprise makes it unlikely that their findings reflected surprise alone, a claim further supported by the high anatomical concordance between their prediction error and our D_{KL} results. This point underscores the importance of decoupling updating and surprise in research aiming at understanding how the brain encodes context information. As a matter of fact, our FPN findings fit well with those reported in studies that have actually decomposed updating and surprise (Schwartenbeck et al., 2016; Kobayashi and Hsu, 2017). For instance, in Kobayashi and Hsu (2017) belief updating about the content of the urn in an Ellsberg three-color urn task modulated activity in bilateral middle frontal gyrus, IPS, and PCun. Summarizing then, it

seems plausible to hypothesize that the FPN regions we found to be modulated by updating of temporal expectations may play a domain-general role in belief updating. Specifically, the congruence between our temporal and other previous non-temporal studies points to the conclusion that updating is encoded by the FPN irrespective of the kind of information to be updated. However, there is also evidence that argues against this possibility (O'Reilly et al., 2013; Vossel et al., 2015). Indeed, O'Reilly et al. (2013) found that updating was mainly located to the ACC (an area also reported in Schwartenbeck et al., 2016). At first glance, the discrepancy between O'Reilly and colleagues' study and ours could seem quite counterintuitive, considering that our temporal paradigm was modeled after their spatial task. By contrast, we believe that such a difference suggests that belief updating in space and time could rely on different brain areas (Coull and Nobre, 1998). Future studies that manipulate updating of the two dimensions in the same task will provide direct evidence on this intriguing question.

Turning back to the FPN, it is interesting to consider a further fractionation of FPN nodes into laterally and medially located ones (see Cocchi et al., 2013, for an overview). Namely, the PCun and PCC are also involved in the "task-negative" default mode network (Fransson and Marrelec, 2008) or, as suggested recently, they can be viewed as a separate hub integrating information across different functional networks (Fornito et al., 2012; Leech et al., 2012). The distinction between lateral and medial FPN regions also emerged for our updating and surprise modulators in that the lateral nodes were modulated by both updating and surprise, whereas the medial nodes were activated by updating only. On the one hand, the finding that lateral regions were modulated to a considerable extent by surprise is broadly consistent with their proposed role in the phasic adjustment of top-down control in response to identified salient stimuli (Seeley et al., 2007; Dosenbach et al., 2008).

Furthermore, this result converges with that obtained in a recent magnetoencephalographic study that performed source reconstruction of the power modulation associated with surprise about event timing (Meindertsma et al., 2018). On the other hand, the distinctive modulation of medial regions by updating is consistent with converging evidence from non-human primates showing the involvement of these regions in environmental change detection (Hayden et al., 2008; Pearson et al., 2009; Hayden et al., 2010), thus corroborating the idea of a putative role of such areas in encoding the statistical properties of changing environments (Pearson et al., 2011).

The importance of trying to disentangle updating and surprise is further supported by the fact that surprise only, but not updating, modulated areas belonging to the CON. This finding strengthens the specificity of our updating results and also gives some hints about the neural correlates of surprise. Following the “dual-network” model (Dosenbach et al., 2008), the CON is in charge of maintaining relevant task-goals across trials, whereas according to other models, the AI along with the ACC forms a “salience network”, which is involved in the transient identification of relevant stimuli in order to guide behavior (Seeley et al., 2007). Specifically, Menon and Uddin (2010) proposed that the AI, which receives multimodal sensory inputs, detects salient stimuli and induces a response in the ACC, which in turn sends transient control signals to lateral FPN and medial parietal regions (Sridharan et al., 2008). This might explain the higher surprise-related connectivity - compared to updating - that we found between AI and dACC/pre-SMA, and between the dACC and parietal nodes of the FPN. However, these conclusions should be taken with caution, since our connectivity analysis was performed post-hoc to better interpret the whole-brain findings and because it did not show enhanced connectivity between CON and prefrontal FPN nodes. If any, we speculate that our surprise-related findings are more in line with

salience network models than with the sustained role of the CON surmised in the dual-network model.

A final result of our study concerns the areas responding to temporal hazard at target onset.

We found that regions located to the auditory and visual cortices were sensitive to the hazard function, a pattern in line with previous studies (Buetti et al., 2010; Vallesi et al., 2009). Interestingly, the finding of auditory activation by visual temporal expectations lends support to the existence of cross-modal associations in temporal preparation as the opposite pattern, that is, modulation of visual areas by auditory temporal expectations has been already shown elsewhere (Buetti and Macaluso, 2010).

To conclude, the present fMRI data showed that updating of internal models and surprise about the timing of relevant events rely on the functioning of two key cognitive control networks. In this regard, our study has the additional value of shedding new light into the differential role of the FPN and CON in higher-order cognition. More importantly, to our knowledge, it provides the first characterization of the brain processes involved in temporal belief updating. Considering that temporal expectations are a fundamental feature of cognitive brain functions, the present study made an important step toward a deeper understanding of the encoding processes involved in temporal belief updating.

Data and code availability statement: Data will be made available on request.

Funding: This work was supported by the European Research Council (n° 313692 to A.Va)

Conflict of Interest. The authors declare no competing financial interests.

References

- Andersson JL, Skare S, Ashburner J (2003) How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *Neuroimage* 20:870-888.
- Baayen RH, Milin P (2010) Analyzing reaction times. *Int J Psychol Res* 3:12-28.
- Baldi P, Itti L (2010) Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Netw* 23:649-666.
- Bates D, Mächler M, Bolker B, Walker S (2015) Fitting linear mixed-effects models using lme4. *J Stat Softw* 67:1-48.
- Brainard DH, Vision S (1997) The psychophysics toolbox. *Spatial vision* 10:433-436.
- Brysbaert M, Stevens M (2018) Power Analysis and Effect Size in Mixed Effects Models: A Tutorial. *J Cogn* 1(1) 9:1-20.
- Bueti D, Macaluso E (2010) Auditory temporal expectations modulate activity in visual cortex. *Neuroimage* 51:1168-1183.
- Bueti D, Bahrami B, Walsh V, Rees G (2010) Encoding of temporal probabilities in the human brain. *J Neurosci* 30:4343-4352.
- Cocchi L, Zalesky A, Fornito A, Mattingley JB (2013) Dynamic cooperation and competition between brain systems during cognitive control. *Trends Cogn Sci* 17:493-501.
- Coull JT, Nobre AC (1998) Where and when to pay attention: the neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *J Neurosci* 18:7426-7435.

- Crittenden BM, Mitchell DJ, Duncan J (2016) Task encoding across the multiple demand cortex is consistent with a frontoparietal and cingulo-opercular dual networks distinction. *J Neurosci* 36:6147-6155.
- Dosenbach NU, Fair DA, Cohen AL, Schlaggar BL, Petersen SE (2008) A dual-networks architecture of top-down control. *Trends Cogn Sci* 12:99-105.
- Doya K, Ishii S, Pouget A, Rao RP (2007) Bayesian brain: Probabilistic approaches to neural coding. MIT press.
- Fornito A, Harrison BJ, Zalesky A, Simons JS (2012) Competitive and cooperative dynamics of large-scale brain functional networks supporting recollection. *Proc Natl Acad Sci U S A* 109:12788-12793.
- Frank MC. 2013. Throwing out the bayesian baby with the optimal bathwater: Response to. *Cognition*. 128:417-423.
- Fransson P, Marrelec G (2008) The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis. *Neuroimage* 42:1178-1184.
- Friston K (2005) A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* 360:815-836.
- Friston K, Buechel C, Fink GR, Morris J, Rolls E, Dolan RJ (1997) Psychophysiological and Modulatory Interactions in Neuroimaging. *NeuroImage* 3:218-229.
- Glasser MF, Sotiropoulos SN, Wilson JA, Coalson TS, Fischl B, Andersson JL, Xu J, Jbabdi S, Webster M, Polimeni JR, Van Essen DC, Jenkinson M (2013) The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* 80:105-124.

- 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.
- Griffiths TL, Chater N, Norris D, Pouget A. 2012. How the Bayesians got their beliefs (and what those beliefs actually are): Comment on Bowers and Davis (2012).
- Halekoh U, Højsgaard S (2014) A Kenward-Roger Approximation and Parametric Bootstrap Methods for Tests in Linear Mixed Models—The R Package pbrtest. *J Stat Softw* 59:1-32.
- Hayden BY, Smith DV, Platt ML (2010) Cognitive control signals in posterior cingulate cortex. *Front Hum Neurosci* 4:223.
- Hayden BY, Nair AC, McCoy AN, Platt ML (2008) Posterior cingulate cortex mediates outcome-contingent allocation of behavior. *Neuron* 60:19-25.
- Herbst SK, Fiedler L, Obleser J (2018) Tracking Temporal Hazard in the Human Electroencephalogram Using a Forward Encoding Model. *eNeuro* 5: ENEURO-0017.
- Itti L, Baldi P (2009) Bayesian surprise attracts human attention. *Vision research* 49:1295-1306.
- Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8:234-241.
- Jenkinson M, Beckmann CF, Behrens TE, Woolrich MW, Smith SM (2012) FSL. *Neuroimage* 62:782-790.
- Kersten D, Mamassian P, Yuille A (2004) Object perception as Bayesian inference. *Annu Rev Psychol* 55:271-304.
- Kleiner M, Brainard D, Pelli D, Ingling A, Murray R, Broussard C (2007) What's new in psychtoolbox-3. *Perception* 36:1-16.

- Knill DC, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci* 27:712-719.
- Kobayashi K, Hsu M (2017) Neural Mechanisms of Updating under Reducible and Irreducible Uncertainty. *J Neurosci* 37:6972-6982.
- Kuznetsova A, Brockhoff PB, Christensen RHB (2017) lmerTest: tests in linear mixed effects models. *J Stat Softw* 82(13):1-26
- Leech R, Braga R, Sharp DJ (2012) Echoes of the brain within the posterior cingulate cortex. *J Neurosci* 32:215-222.
- Meindertsma T, Kloosterman NA, Engel AK, Wagenmakers EJ, Donner TH (2018) Surprise About Sensory Event Timing Drives Cortical Transients in the Beta Frequency Band. *J Neurosci* 38:7600-7610.
- Menon V (2015) Salience Network. In: *Brain Mapping* (Toga AW, ed), pp 597-611. Waltham: Academic Press.
- Menon V, Uddin LQ (2010) Saliency, switching, attention and control: a network model of insula function. *Brain Struct Funct* 214:655-667.
- Nobre AC, van Ede F (2018) Anticipated moments: temporal structure in attention. *Nat Rev Neurosci* 19:34-48.
- O'Reilly JX, Schuffelgen 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 U S A* 110:E3660-3669.
- Ostwald D, Schneider S, Bruckner R, Horvath L (2019) Power, positive predictive value, and sample size calculations for random field theory-based fMRI inference. *bioRxiv* 613331.

- Pearson JM, Hayden BY, Raghavachari S, Platt ML (2009) Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Curr Biol* 19:1532-1537.
- Pearson JM, Heilbronner SR, Barack DL, Hayden BY, Platt ML (2011) Posterior cingulate cortex: adapting behavior to a changing world. *Trends Cogn Sci* 15:143-151.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial vision* 10:437-442.
- Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE (2012) Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* 59:2142-2154.
- Schwartenbeck P, FitzGerald THB, Dolan R (2016) Neural signals encoding shifts in beliefs. *Neuroimage* 125:578-586.
- Seeley WW, Menon V, Schatzberg AF, Keller J, Glover GH, Kenna H, Reiss AL, Greicius MD (2007) Dissociable intrinsic connectivity networks for salience processing and executive control. *J Neurosci* 27:2349-2356.
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27:623-656.
- Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, Bannister PR, De Luca M, Drobnjak I, Flitney DE, Niazy RK, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady JM, Matthews PM (2004) Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23 Suppl 1:S208-219.
- Sridharan D, Levitin DJ, Menon V (2008) A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proc Natl Acad Sci U S A* 105:12569-12574.

- Teichner WH, Krebs MJ. 1974. Laws of visual choice reaction time. *Psychol Rev.* 81:75-98.
- Vallesi A, McIntosh AR, Shallice T, Stuss DT (2009) When time shapes behavior: fMRI evidence of brain correlates of temporal monitoring. *J Cogn Neurosci* 21:1116-1126.
- Vossel S, Mathys C, Stephan KE, Friston KJ (2015) Cortical Coupling Reflects Bayesian Belief Updating in the Deployment of Spatial Attention. *J Neurosci* 35:11532-11542.
- Waskom ML, Psychology Do, Frank MC, Psychology Do, Wagner AD, Psychology Do, Neurosciences Program SU, Stanford, CA 94305, USA (2017) Adaptive Engagement of Cognitive Control in Context-Dependent Decision Making. *Cerebral Cortex* 27:1270-1284.
- Zénon A, Solopchuk O, Pezzulo G. 2019. An information-theoretic perspective on the costs of cognition. *Neuropsychologia.* 123:5-18.