Surprise About Sensory Event Timing Drives Cortical Transients in the Beta Frequency Band

T. Meindertsma[1,2,3], N.A. Kloosterman[2,3,4], A.K. Engel[1], E.J. Wagenmakers[2], T.H. Donner[1,2,3]

[1]Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, Hamburg, Germany; [2]Department of Psychology, [3]Amsterdam Brain and Cognition, University of Amsterdam, Amsterdam, Netherlands; [4]Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Max Planck Institute for Human Development, Berlin, Germany

Address for correspondence:
Thomas Meindertsma
Dept. of Psychology
University of Amsterdam
Nieuwe Achtergracht 129B, 1018 WS Amsterdam, The Netherlands
Email: t.meindertsma@uva.nl
t.donner@uke.de

Author contributions:
Conceptualization, T.M., N.A.K. , E.J.W. and T.H.D.; Investigation, T.M.; Formal analysis T.M.; Analytic tools, E.J.W. and N.A.K.; Writing - Original draft, T.M. and T.H.D; Writing - Review & Editing, T.M., N.A.K., A.K.E., E.J.W. and T.H.D.; Funding Acquisition, A.K.E. and T.H.D.; Supervision, T.H.D.

Number of pages: 20
Number of figures: 5
Number of tables/multimedia/3D models: 0
Number of words Abstract: 250
Number of words Introduction: 119
Number of words Discussion: 1311

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

Acknowledgements: This work was supported by the Netherlands Organization for Scientific Research (NWO, dossiernummer 406-14-016, to T.H.D. and T.M.); the Amsterdam Brain and Cognition priority program (ABC2014-01, to T.H.D.); the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 604102 (Human Brain Project) (to T.H.D. and A.K.E.); and the German Research Foundation (DFG): Heisenberg Professorship DO 1240/3-1 (to T.H.D.), the Collaborative Research Centers SFB 936 (Projects A2/A3, A7, to A.K.E., T.H.D.) and TRR 169 (Project B1 to A.K.E.). We thank all members of the Donner lab for helpful discussion.

Abstract

Understanding the statistical structure of the environment is crucial for adaptive behavior. Humans and non-human decision-makers seem to track such structure through a process of probabilistic inference, which enables predictions about behaviorally relevant events. Deviations from such predictions cause surprise, which in turn helps improve the inference. Surprise about the timing of behaviorally relevant sensory events drives phasic responses of neuromodulatory brainstem systems, which project to the cerebral cortex. Here, we developed a computational model-based magnetoencephalography (MEG) approach for mapping the resulting cortical transients across space, time, and frequency, in the human brain. We used a Bayesian updating model to estimate the predicted timing of the next stimulus change in a simple visual detection task. This model yielded quantitative trial-by-trial estimates of temporal surprise. The model-based surprise variable predicted trial-by-trial variations in reaction time more strongly than the externally observable interval timings alone. Trial-by-trial variations in surprise were negatively correlated with the power of cortical population activity measured with MEG. This surprise-related power suppression occurred transiently around the behavioral response, specifically in the beta frequency band. It peaked in left lateral prefrontal as well as in frontal midline regions, and its cortical distribution was distinct from the movement-related suppression of beta power in motor cortex. Our results indicate that surprise about sensory event timing transiently suppresses ongoing beta-band oscillations in association cortex. This transient suppression of frontal beta-band oscillations might reflect an active reset triggered by surprise, and is in line with the idea that beta-oscillations help maintain cognitive sets.

Significance statement

Agents continuously track the statistical structure of the environment, in order to make predictions about behaviorally relevant sensory events. Deviations from such predictions cause surprise, which in turn drives phasic responses of neuromodulatory brainstem systems that project to the cerebral cortex. We developed a computational model-based magnetoencephalography approach, which enabled us to map out transients changes in cortical population dynamics elicited by surprise about sensory event timing, across space, time, and frequency, in the human brain. The model-based estimates of surprise predicted behavior as well as a transient suppression of beta frequency-band oscillations in frontal cortical regions. Our results are in line with conceptual accounts that have linked neural oscillations in the beta-band to the maintenance of cognitive sets.

Introduction

Humans and other organisms continuously adapt their behavior to the statistical structure of their environment. This suggests that the brain is equipped with powerful machinery for statistical learning, which can interact with the neural processes driving goal-directed behavior. Of particular importance here is surprise (Dayan and Yu, 2006; O'Reilly et al., 2013), a violation of one's expectation about the next event, which might indicate a sudden change in the environmental structure, which might transiently boost central arousal state, increasing the organism's sensitivity and learning rate (Yu and Dayan, 2005; Nassar et al., 2012).

Expectation, uncertainty, and surprise are intricately related concepts. The precision of expectations scales with uncertainty, that is, the width of the distribution of events: high uncertainty precludes forming precise expectations. Violations of expectations cause surprise, the level of which depends on the difference between the expected and actually observed event (often termed prediction error). These intuitions can be readily formalized within the framework of Bayesian statistics and used to search for neurophysiological correlates (see Materials and Methods: *Bayesian model of surprise and uncertainty*).

One dimension of environmental statistics that has profound effects on behavior is the timing of behaviorally relevant sensory events (Gibbon et al., 1997; Nobre et al., 2007) Two lines of work have studied the neural basis of temporal expectation effects. One has shown that environments with rhythmic (i.e., precise) temporal structure entrain neural oscillations in the cerebral cortex, the phase of which then modulates sensory cortical responses, perception, and cognition (Lakatos et al., 2008; Schroeder and Lakatos, 2009; Rohenkohl and Nobre, 2011; Rohenkohl et al., 2012; Riecke et al., 2015; van Ede et al., 2017). Because in these periodic contexts, surprise is minimized (once the structure is learned expectations match observations), this work has not identified neural correlates of surprise.

Another line of work has instead studied neural responses of subcortical, neuromodulatory centers (specifically, dopaminergic midbrain) to sensory events (specifically, rewards). Because event timing here varied non-periodically from trial to trial as in many natural environments, this work could link phasic neuromodulatory responses to temporal surprise (Hollerman and Schultz, 1998; Fiorillo et al., 2008). It is likely that such surprise-driven phasic responses also occur in other neuromodulatory systems (e.g., the noradrenergic system; Dayan and Yu, 2006) with widespread projections to the cortical networks underlying goal-directed behavior. But little is known about the cortical responses to surprise about event timing.

Here, we present a computational approach for comprehensively mapping cortical transients encoding temporal surprise across space, time, and frequency. We developed a Bayesian learning model that used previous interval durations to estimate the subjects' belief about the temporal structure of the environment in a simple detection task. The model output enabled us to compute trial-to-trial measures of uncertainty and surprise. Correlating these computational quantities to brain-wide cortical dynamics measured with magnetoencephalography (MEG) pinpointed clusters in the time-frequency-space domain encoding surprise. This revealed widespread modulations of cortical dynamics in the beta band (around 20 Hz).

Materials & Methods

This paper reports a re-analysis of an MEG data set that has previously been used for a study into decision-related feedback signals in visual cortex (Meindertsma et al., 2017). Here, we focus on those aspects of the experimental design that are most

relevant for the issue addressed in the current paper: uncertainty and surprise about the timing of the experimental events specified below. We refer to our previous paper (Meindertsma et al., 2017) for a more detailed description of the visual stimulus and the behavioral task.

Participants

Thirty-one volunteers participated in the experiment. Two participants were excluded due to incomplete data and one participant did not complete the experiment due to poor quality of simultaneously acquired pupil data. Thus, 28 participants (17 female, age range 20 - 54 years, mean age 28.3, SD 9.2) were included in the analysis. All participants had normal or corrected-to-normal vision and no known history of neurological disorders. The experiment was conducted in accordance with the Declaration of Helsinki and approved by the local ethics committee of the Hamburg Medical Association. Each participant gave written informed consent.

Stimulus

MEG was measured while subjects viewed the intermittent presentation of a target (full contrast Gabor patch; diameter: 2°) and reported the on- and offset of the target (Figure 1A). The Gabor patch contained two cycles and flickered at 10 Hz. Target flicker was implemented by counter-phasing the sinusoid used to generate the Gabor patch. The target was located in either the lower left or lower right visual field quadrant (eccentricity: 5°, counterbalanced between subjects), surrounded by a rotating mask (17°x17° grid of black crosses), and superimposed on a gray background. The mask rotated at a speed of 160°/s. The target was separated from the mask by a gray "protection zone" subtending about 2° around the target (Bonneh et al., 2001). Subjects fixated on a fixation mark (red outline, white inside, 0.8° width and length) centered on the mask in the middle of the screen. Stimuli were presented using the Presentation Software (NeuroBehavioral Systems, Albany, CA, USA). Stimuli were back-projected on a transparent screen using a Sanyo PLC-XP51 projector with a resolution of 1024x768 pixels at 60 Hz. Subjects were seated 58 centimeters from the screen in a whole-head magnetoencephalography (MEG) scanner setup in a dimly lit room.

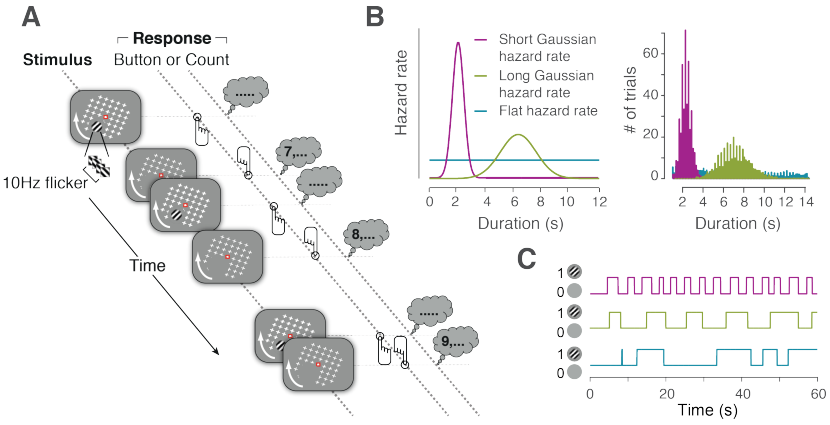

Figure 1: Behavioral task. A. Schematic depiction of the stimulus and task. A salient, flickering target (Gabor patch) temporarily appeared and disappeared on a rotating background. Subjects fixated on the red fixation mark and reported stimulus changes either by direct button press or silently counting the disappearances and reporting the total number at the end of the run. **B.** The interval duration between stimulus changes was randomly drawn from one of three distributions that corresponded to three hazard rates (left), resulting in distinct distributions of intervals (right, average histogram over

subjects). **C.** Example time courses of target presence (1 = present, 0 = absent) drawn from these distributions.

Behavioral task and experimental design

The subjects' task was to maintain stable fixation and detect the physical offsets and onsets of the target, the predictability of which fluctuated from trial to trial, and the mean predictability of which varied systematically across blocks. To this end, the interval durations between stimulus changes were sampled from three different distributions in the different blocks. These distributions were computed so as to produce three predetermined so-called hazard functions, which describe the probability that an event will occur at a particular time, given that it has not occurred yet. The hazard function formalizes the expectation of a change and affects human reaction times in simple detection tasks (Luce, 1986). The hazard function can be computed as follows:

$$\lambda_t = \frac{f_t}{1-F_t}, \qquad \text{Eq. 1}$$

where λ_t is the value of the hazard function at time point t, f_t is the value of distribution f on time point t, and F_t is the area under the curve of distribution f from $-\infty$ to time t.

We used the following procedure to construct three conditions, referred to as 'Short', 'Long', and 'Flat' below. We first selected three hazard functions that systematically differed in their level of predictability (Figure 1B, C). We then computed the actual distributions of intervals by re-arranging Eq. 1 as follows:

$$f_t = \lambda_t * (1 - F_t) , \qquad \text{Eq. 2}$$

The interval durations were then randomly selected from f. Specifically, the conditions were defined as follows:

Short: The hazard function was a narrow Gaussian distribution with a mean of 2 s and a standard deviation of 0.2 s. This resulted in nearly periodic and, thus, largely predictable intervals between events.

Long: This condition used the same hazard function as the previous condition, but with a larger mean and standard deviation (6 s and 0.6 s, respectively) thus rendering event timings less predictable (Fiorillo et al., 2008).

Flat: The hazard function was flat with a mean of 6s, yielding the least predictable interval durations. The resulting distribution of interval durations, f_t, therefore, approximated an exponential distribution; characterizing a memory-less process (i.e. the timing of the next event could not be predicted from previously encountered intervals, Feller, 1959).

Computational analysis with a Bayesian model (Fig. 2) described below confirmed that the sampled intervals from these three conditions gave rise to different mean levels of uncertainty and surprise (Fig. 2 D, E). The three experimental conditions were presented in separate blocks, which were divided in three-minute runs of continuous presentation.

Subjects were asked to report the stimulus changes either by button press or by silently counting the number of target offsets for later report, a manipulation that was critical for the analyses reported in our previous paper (Meindertsma et al., 2017). Those two conditions were randomly selected before each run under the constraint that both would occur equally often. The corresponding instructions were displayed on the screen before the run started. Subjects could only start the next run after they confirmed the instructions to the experimenter over the intercom.

Here, we focused on the condition entailing immediate behavioral report so as to study the impact of surprise on RTs and on response-related cortical dynamics. Subjects reported target offsets and onsets by pressing a button with the index finger and middle finger of their right hand.

All subjects completed a total of 6 runs of the Short condition, and 16 runs of the other two conditions. Additionally, subjects performed a motion-induced blindness task and a functional localizer task, which were not relevant for the current study, but are reported in our previous paper (Meindertsma et al., 2017). The order of blocks was counter-balanced across subjects.

Bayesian model of surprise and uncertainty

We developed a Bayesian updating model to quantify surprise and uncertainty about the timing of sensory events (i.e., the target on- and offsets). The model tracked the evolving predictive distribution of upcoming interval durations; more specifically, it computes the posterior predictive of unobserved interval durations, conditional on the observed data, throughout each block of the experiment. We assumed that subjects' tracked the temporal statistics of the task in a similar way, and we used the posterior predictive distribution as a proxy of the subjects' belief states (i.e., their prediction of the timing of the next stimulus change).

We assumed that the subjects used a model in which the observed intervals have been generated from a gamma distribution with parameters alpha (shape) and beta (scale). These parameters were given uninformative prior distributions (Lee and Wagenmakers, 2013), which were updated by the data to posterior distributions. Then we could obtain the expectations about to-be-observed intervals by generating posterior predictives (i.e., drawing an alpha-beta pair from the joint posterior distribution and then drawing a predicted interval from the associated gamma distribution; repeating this process many times yields a posterior predictive distribution for the to-be-observed intervals). We assumed that the subjects updated their belief state after each observation of a new interval duration. Likewise, the model was updated after every interval t by computing a new posterior predictive distribution, based on the durations of intervals $1{:}t$ and the prior.

We generated posterior predictive distributions over intervals using Gibbs sampling (a Markov chain Monte Carlo, or MCMC, algorithm (Andrieu et al., 2003) in the software JAGS (Plummer, 2003) and Matlab (version R2013a). We used two Markov chains with different starting points of 10,000 samples per chain with 1000 samples burn-in. We transformed the distribution of MCMC samples into a continuous probability density function by fitting a gamma function to the pooled distribution of both MCMC chains (Figure 2A,B):

$$f_t = gamma(\alpha_t, \beta_t \mid D_{1:t}), \qquad \text{Eq. 3}$$

where f_t was the probability density after observing interval t and α_t and β_t were the parameters of the gamma function and $D_{1:t}$ was the sampled distribution (i.e., the distribution of MCMC samples).

To be able to relate trial-to-trial uncertainty and surprise to behavior and the MEG data, we extracted two information theoretic metrics from the time-evolving posterior predictive distribution (i.e., belief) f_t.

Uncertainty: We quantified trial-to-trial uncertainty about the timing of the next interval $t{+}1$ as the entropy of the posterior probability distribution:

$$H_t = -\int_0^\infty (f_t(x) * \log f_t(x))\, dx, \qquad \text{Eq. 4}$$

where H_t was entropy after interval t, x were all possible instances of the probability function (i.e., interval durations). Entropy depended on the width of f_t, and thus uncertainty was higher when predictions of interval durations were less precise (Figure 2A,C). From here on, we will use the term entropy when referring to uncertainty, for the sake of mathematical precision.

Surprise: For every succeeding event *t+1*, we computed the surprise about the corresponding interval duration in terms of the Shannon information conveyed by the interval duration x_{t+1}, given the posterior predictive distribution (f_t) estimated from the previous interval (i.e., based on intervals *1:t*):

$$I_{t+1} = -\log f_t(x_{t+1}), \qquad \text{Eq. 5}$$

where I_{t+1} was the information gained by adding interval *t+1*, given f_t. Thus, surprise was defined as the negative log-likelihood of the next interval, given the intervals that had been presented so far, whereby the posterior distribution from the previous interval f_t was used as prior distribution f_{t+1} for the next interval in the updating process. We added one further transformation in the computation of surprise . The surprise measures defined in Eq. 5 quantified the surprise about the current event timing based on the prior distribution estimated from all previous interval durations, but disregarding the time elapsed on the current trial. It is unlikely that exactly this distribution translated into subjects' level of surprise: as time passed and no event occurred on a given trial, all interval durations shorter than the elapsed time became impossible. Subjects likely discounted these impossible intervals in their expectation of the timing of the upcoming event, which should have also affected their level of surprise. In other words, their internal representation of the prior distribution changed dynamically throughout each trial, as a function of elapsed time. To capture this process, we constructed a time-varying version of the prior distribution f_t, which was also conditioned on the elapsed time on trial t. This version was equal to f_t for elapsed time equal to 0 and then increasingly deviated from f_t as elapsed time grew. We approximated this time-varying distribution, denoted as f'_t in the following, by setting all probabilities in f_t up to the current time point to zero and renormalizing the remaining distribution to integrate to 1. We then computed surprise based on this new distribution f'_t using Eq. 5. The time-variant prior f'_t converged to 1 as time passed, and thus surprise approached zero for longer intervals.

Regressing computational variables against behavior

We used reaction time (RT) during Detection-button as behavioral readout of the impact of uncertainty and surprise. Accuracy approached ceiling for all subjects, due to the high saliency of the target. We computed and compared mean RTs per condition and stimulus event (target off- and onset). Furthermore, we assessed the Pearson correlation between log-transformed single-trial RTs and the trial-to-trial estimates of surprise or entropy. RT was log-transformed to normalize the (skewed) distributions of 'raw' RT before computing correlation coefficients. Differences from zero and differences between conditions were tested using permutation tests over subjects (two-sided, 10,000 permutations). We tested the difference in magnitude of the correlation of previous interval of log(RT) compared to surprise and entropy to log(RT) by testing the difference in absolute value across subjects using permutation tests (Figure 3C) and computed within subject 95% confidence intervals using Steiger's correlation test (Zou, 2007).

MEG data collection

Magnetoencephalography (MEG) data were acquired on a CTF 275 MEG system (VSM/CTF Systems, Port Coquitlam, British Columbia, Canada) with a sample rate of 1200 Hz. The location of the subjects' head was measured in real-time using three fiducial markers placed in the both ears and on the nasal bridge to control for excessive movement. Furthermore, electrooculogram (EOG) and electrocardiogram (ECG) were recorded to aid artifact rejection. All data were recorded in blocks of four runs of three minutes duration (or two runs at the end of a block), which corresponded to the runs of experimental conditions defined above.

MEG data analysis

Preprocessing. The data were analyzed in Matlab (version R2013a, The Mathworks, Natick, MA, USA) using the Fieldtrip (Oostenveld et al., 2011) toolbox and custom-made software.

Trial extraction. In runs involving subjects' reports, we extracted trials of variable duration, centered on subjects' button presses, from the 3 min runs of continuous stimulation. We call this method for trial extraction "response-locked". The following constraints were used to avoid mixing data segments from different percepts when averaging across trials: (i) The maximum trial duration ranged from −1.5 s to 1.5 s relative to report; (ii) when another report occurred within this interval, the trial was terminated 0.5 s from this report; (iii) when two reports succeeded one another within 0.5 s, no trial was defined; (iv) for the analysis of Detection-button runs, we included only those reports that were preceded by a physical change of the target stimulus within 0.2 to 1 s, thus discarding reports not related to stimulus changes. We used this method for the analyses related to surprise. In an alternative analysis of all Detection runs, trials were defined in the same way as described above, but now aligned to physical target on- and offsets ("stimulus-locked"). In the Detection-count conditions, no button responses were given during the run, so stimulus-locked trial extraction was the only option. We used this method for the analysis related to entropy (see Kloosterman et al., 2015b & Meindertsma et al., 2017 for a similar procedure).

Artifact rejection. All epochs that contained artifacts caused by environmental noise, eye-blinks, muscle activity or squid jumps were excluded from further analysis using standard automatic methods included in the Fieldtrip toolbox. Epochs that were marked as containing an artifact were discarded after every artifact detection step. For all artifact detection steps the artifact thresholds were set individually for all subjects. Both of these choices aimed at optimization of artifact exclusion. Line-noise was filtered out by subtracting the 50, 100, 150 and 200 Hz frequency components from the signal.

Time-frequency decomposition. We used a sliding window Fourier transform to compute the time-frequency representation for each sensor and each trial of the MEG data. The sliding window had a length of 200 ms and a step size of 50 ms, with one Hanning taper (frequency range 5-35 Hz, frequency resolution 2.5 Hz and bin size 1 Hz). The data was baseline corrected for every frequency bin and MEG sensor separately. The baseline was computed by averaging single-trial power over the baseline time window. The baseline time windows ranged from −1.25 to −0.75 s for response-locked and −1 to −0.5 s for stimulus-locked analyses, respectively. The time course of every frequency bin and sensor combination was first baseline corrected by subtracting the single-trial baseline and then normalized by dividing by the mean over the baselines of all trials within a condition (*Short*, *Long* or *Flat*).

Source reconstruction. We used an adaptive linear spatial filtering method called linear beamforming (Van Veen et al., 1997; Gross et al., 2001) to estimate single-trial modulations of MEG power at the source level. We computed a common

filter for a baseline time window (1 to 0.5 s before response), a 'transient' time window, and a frequency band of interest (0 to 0.5 s after response, 20 Hz +/- 4 Hz spectral smoothing, see dashed box in Figure 4A). The transient time window and frequency band of interest were selected based on cluster-based statistics at the sensor level (see next section). We used the measured head positions and individual single-shell volume conductor models, based on individual images from T1-weighted structural MRI. We computed the power values, in both baseline and transient time windows, for each trial and source grid point (i.e., voxel) as follows. First, we projected the sensor-level MEG power values from the time window of interest as well as from a baseline time window through the common spatial filter. Second, we converted the estimated power values during the time window of interest into units of power modulation, again by subtracting and dividing by the corresponding baseline power values.

Correlating single-trial computational variables to MEG power

We correlated the MEG power modulation to our measures of entropy and surprise, as derived using our Bayesian model (see *Bayesian model of surprise and uncertainty*) across trials.

Entropy: We correlated entropy to the MEG power modulation separately in every MEG sensor and frequency bin. This was done within subject and separately for the three hazard rate conditions. There are structural differences in entropy and surprise between these conditions (Figure 2D,E), thus pooling over these conditions might result in inflated correlations that reflect session differences instead of the true correlation between entropy and MEG power. We reasoned that entropy should affect baseline or tonic arousal, where high entropy should cause higher arousal. As our task was continuous, we considered the time window right before the stimulus change the best reflection of a baseline state. For this reason we averaged the MEG power over the time period right before a stimulus change (-0.5 to -0.25s with respect to the target disappearance or reappearance) before correlating to entropy.

The results were then averaged over the three conditions and transformed with the Fisher z transformation (Fisher, 1915):

$$z = 0.5 \cdot \ln\left(\frac{1+r}{1-r}\right) \qquad \text{Eq. 6}$$

We used two-tailed permutation tests with a cluster-based correction for multiple corrections to test the correlation coefficients against zero (Efron and Tibshirani, 1998; Maris and Oostenveld, 2007).

Surprise: Correlations between surprise and MEG power modulation were performed using the same method, with the following exceptions. First, we attuned the analysis in two ways to account for the correlation between surprise and RT (Figure 3). Because of this correlation, any post-stimulus correlations between surprise and MEG power modulation might reflect differences in the timing of the button press. We performed this analysis response-locked, because these RT differences are difficult to disentangle from genuine effects of surprise when the power modulations are time-locked to the stimulus change. Additionally, to account for confounding effects of RT and the duration of the previous interval, we also performed a partial correlation analysis between surprise and MEG power modulation with the interval duration preceding the stimulus change or RT as covariate. Second, for the correlation between surprise and MEG power modulation we did not average over a specific time window, but instead performed correlations separately for every time point, resulting in a 3-dimensional matrix of correlations (sensor * frequency bin * time point). Consequently, we also performed cluster-

based permutation statistics over these three dimensions. The correlations that survived cluster correction were visualized by integrating (i.e. computing the area under the curve) over sensors and frequency bins (for the time course), sensors and time points (for the frequency spectrum), frequency bins and time points (for the topography) or just over sensors for the time frequency representation (see Hipp et al., 2012 for a similar approach).

To assess the robustness of the emerging clusters we performed a cross-validation analysis using a leave-one-out procedure. To this end, we repeated the analyses on all possible iterations of N-1 subjects, each time using the resulting cluster as a mask to calculate the average correlation in the left-out subject, separately for target offset and onset trials. These values were tested against zero and against each other across subjects using permutation tests (10.000 permutations).

Trial-to-trial surprise, and to a lesser extend entropy, correlated to log(RT) (Figure 3B). We interpreted this as evidence that our surprise metric indeed captures a process that is behaviorally relevant to the subjects. From this perspective, we predicted that the surprise-related MEG cluster was related to RT as well. To test this hypothesis we computed the correlation between trial-to-trial power modulation averaged over the whole cluster and log(RT). The resulting correlations were tested against zero across subjects using a permutation test.

The transient modulations of MEG power estimated for each voxel in the source grid derived by means of source reconstruction (see *MEG data analysis: Source reconstruction*) was correlated to the trial-to-trial measure of surprise. This was done separately within each subject and the resulting correlations averaged over subjects after Fischer's z-transformation (Eq. 6). For comparison, we also computed the average modulations of MEG power in the same time window and frequency band. The resulting maps of correlation or average power modulation were nonlinearly aligned to a template brain (Montreal Neurological Institute) using the individual images from structural MRI.

Results

Subjects (N=28) performed a simple visual detection task reporting on- and offsets of a small, but salient target stimulus (Figure 1A). In different blocks, target events were administered using three different temporal conditions (Figure 1B,C) translating into different overall levels of uncertainty and surprise about the timing of target events (Figure 2D,E). In order to quantify these two computational variables not only across conditions, but also across individual trials, we developed a Bayesian belief-updating model. The model approximated subjects' evolving beliefs (i.e. the prior or posterior predictive distributions in Bayesian terms) about the temporal intervals between the sensory events, which were dynamically updated across trials and even within trials (for surprise, see *Materials and Methods*). From these time-evolving probability distributions, we extracted trial-by-trial measures of information-theoretic entropy (quantifying uncertainty) and surprise.

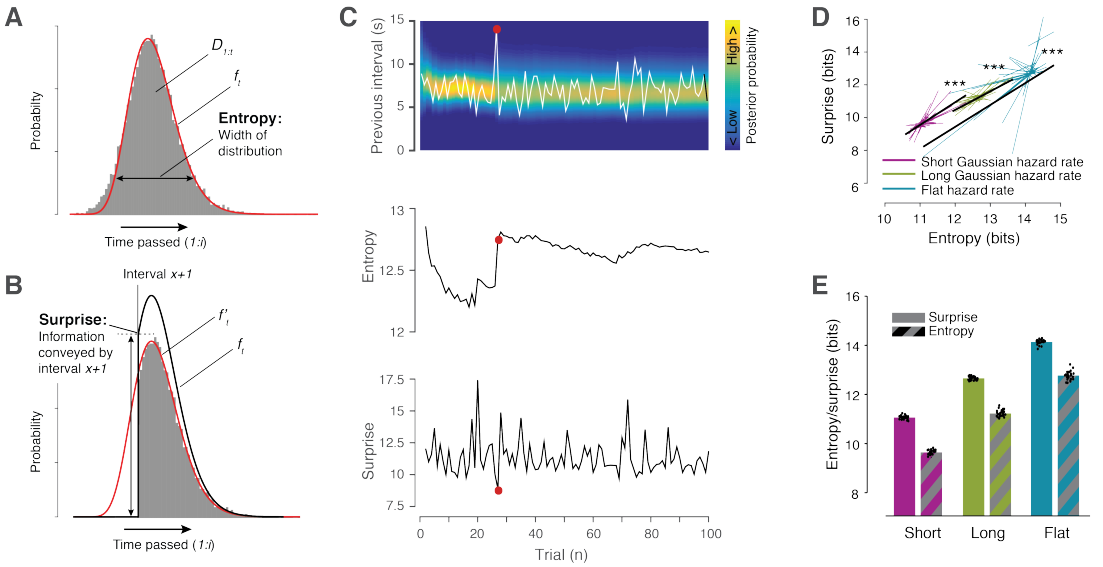

Figure 2: Bayesian updating model of belief about temporal structure A-C. The model estimated the posterior predictive distribution over timings of stimulus changes for each trial *t*. This distribution is denoted as f_t. The gray histogram shows the frequency distribution of intervals from all trials up to trial *t*, denoted as $D_{1:t}$. f_t was estimated by fitting a gamma probability density function (red line) to $D_{1:t}$; it was then used to extract two different information-theoretic computational variables for each trial: entropy and surprise. **A.** Entropy, a measure of the uncertainty about the timing of the interval duration from the current trial, computed from the complete distribution f_t using Eq. 4 (see main text). The wider the distribution, the higher entropy. **B.** Surprise, a measure of information provided by each new interval duration, was also computed from the posterior predictive distribution, but with one extra step (see main text): the part of the distribution up to the current interval duration was truncated, and the remainder of the distribution re-normalized to integrate to 1 (f'_t , black line). Surprise was defined based on this truncated function using Eq. 5 (see main text). **C.** Relationship between interval durations (white line in top panel, from the long Gaussian condition), posterior predictive distribution *f* (color coded in top panel), entropy (middle), and surprise (bottom). Red dot: exceptionally long interval (see duration in top panel). Surprise on this trial was low (bottom panel) because time dependent surprise decreased over time. After observing this interval entropy increased (middle panel) because the observed interval was longer than the expected duration, given previous intervals. **D.** Regression of surprise on entropy. Thin colored lines, regression lines of single subjects; black lines, group average regression. **E.** Trial-averaged surprise and entropy for the three experimental conditions defined in Fig. 1. Bars, group average; black dots, single subjects. *** p<0.001, permutation tests across subjects.

Estimates of entropy and surprise fluctuated across trials, especially in the early part of each block (Figure 2C). The trial averages of both measures within each block also varied lawfully between the different experimental conditions, scaling with the predictability of the stimulus changes (Figure 2D,E). Estimates were smallest for the Short condition, intermediate for the Long condition, and largest for the Flat condition. Further, variations in entropy and surprise were weakly correlated across trials (Figure 2D). This was expected, because both measures draw information from the same posterior predictive distribution. However, despite of this correlation the measures reflected functionally different concepts (uncertainty and surprise) and thus both merited further investigation.

Surprise predicts reaction time

The model-derived computational variables entropy and surprise predicted subjects' reaction time (RT) in the detection task (Figure 3). Mean RT scaled with the experimental conditions in the same way as surprise, with the fastest RTs for Short and slowest RTs for the Flat condition (Figure 3A, compare with Figure 2E). In

addition, the trial-to-trial variations of surprise within conditions (and in some conditions entropy) were correlated to trial-to-trial variations in (log-transformed) RT (Figure 3B). For surprise (bottom panel), this correlation was significant for all experimental conditions, with the exception of target offsets on the Long condition, as well as in the average across conditions. For entropy (top) this correlation was only present in the Flat condition, but not the other two. The overall weaker correlation for entropy than surprise likely reflected the overall smaller fluctuations of entropy as time progressed throughout the blocks (Figure 2C); those fluctuations were particularly small during the Short and Long conditions but stronger during Flat.

We are agnostic to the precise inference process that our subjects used to track the temporal structure of the task. Specifically, we do not claim that they implemented the exact computations prescribed by the model. But in line with a substantial body of evidence from psychology (Sutton and Barto, 1998; Gold and Shadlen, 2007; Glaze et al., 2015), we did postulate that subjects integrated, in some way, observations throughout each block. Our model implements this by integrating across the entire history of the observations (here: of interval durations) and updating internal representations accordingly. The surprise and entropy metrics correlated to the actual interval durations (white line in Figure 2C), but deviated progressively from them as time progressed throughout the block. The assumption that subjects also integrated over more than just the previous interval was supported by the analysis below.

The duration of the interval preceding a stimulus change also correlated to RT, with longer interval durations corresponding to faster responses (Figure 3C, left panel). However, surprise predicted RT better than previous interval duration, reflected by a significant difference between the magnitudes of correlations between surprise/previous interval and RT (Figure 3C, '*Difference surprise – prev. inter.*', permutation test). This indicates that the correlation between surprise and RT reflects subjects' sensitivity to the temporal structure of the task, rather than just their growing expectation of a stimulus change as the interval evolves. As for the direct correlations between surprise and RT (Figure 3B), these differences were significant for all but the target offsets in the Long condition (figure 3D). This difference was not observed for entropy (Figure 3C,D).

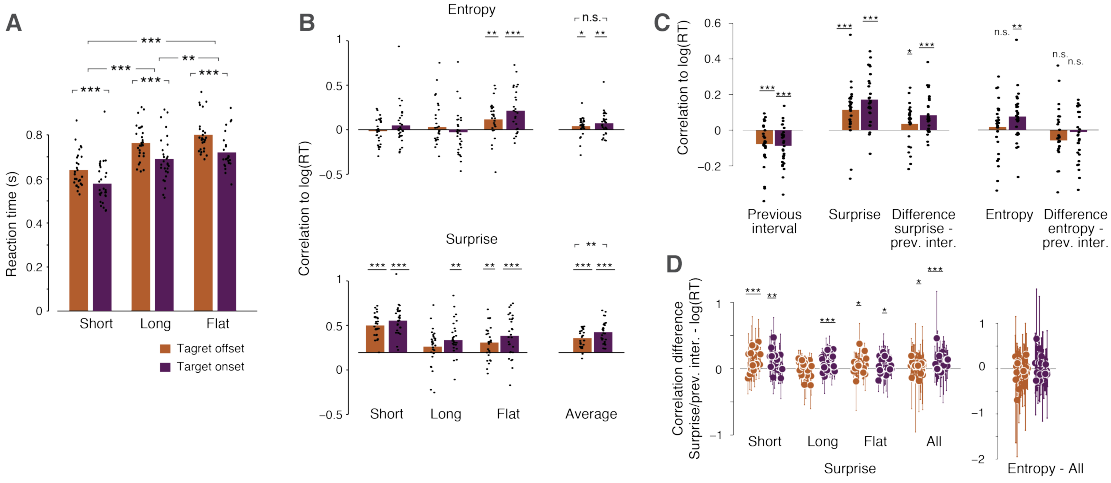

Figure 3: Link between computational variables and behavior. A. Average reaction time (RT) per interval distribution, separate for reports of disappearance and reappearance of the target. Bars show average over subjects; black dots depict average per subject. **B.** Correlation between log(RT) and entropy (top) and surprise (bottom) for each experimental condition. Bars, group average; black dots, individual subjects. **C.** Correlation between log(RT) (averaged across conditions) and the interval

duration preceding stimulus change, and the difference in correlation magnitude between surprise/previous interval and RT. Same for entropy. Bars show average over subjects; black dots depict average per subject. **D.** Difference between magnitude of correlation of surprise and log(RT) and previous interval duration and log(RT), per condition. Dots depict correlation difference per subject; error bars show 95% confidence intervals. *** p<0.001, ** p<0.01, * p<0.05; n.s., not significant; permutation tests.

Taken together, the results from Figure 3 indicate that the surprise variable derived from the Bayesian model captured variations in behavior. We next searched the whole-brain MEG data for a dynamical neurophysiological signature of this process. To this end, we focused on the trial-to-trial fluctuations of surprise within conditions, which were more pronounced than the differences in mean surprise between Short, Long, and Flat conditions (recorded in separate MEG runs).

Widespread cortical beta-band transient driven by surprise

We mapped out the cortical responses to trial-to-trial fluctuations in surprise by correlating the model-based surprise measures to modulations of MEG power, around the time of subjects' behavioral responses to sensory events. We did this in an exhaustive fashion across every time and frequency bin and MEG sensor and tested for clusters of significant correlations across these three dimensions, while applying cluster-based multiple comparison correction (*Materials and Methods*). This approach revealed negative correlations in the beta (~20 Hz) frequency range, indicating that higher surprise was associated with lower beta power. This negative correlation cluster started about 0.2 s before and peaked about 0.25 s after subjects' report of the stimulus change. This cluster exhibited several peaks over central, left frontal, and to a lesser extend left parietal cortex (Figure 4A,C).

The surprise-related cluster was robust and not driven by outliers, and the effect was not specific to the type of stimulus event (target on- or offset; Figure 5). To assess these possibilities, we computed the correlation between surprise and power in the cluster (figure 4A), separately for target offsets and onsets. We found robust negative correlations for both event types, although the correlations were somewhat stronger for target onsets (Figure 4D). Furthermore, we used a leave-one-out cross-validation procedure to test the robustness of the correlations on both target on- and offsets. Again, the correlation was significantly negative in both cases (Figure 4E).

As expected from previous work on modulations of MEG power around motor responses (Donner et al., 2009), the overall modulation of MEG power in the time-frequency window of the cluster (16-24 Hz, 0 – 0.5 s from response) peaked in bilateral motor cortex (Figure 4B). But the component of beta-power modulations that correlated with trial-by-trial surprise showed a different cortical distribution, with negative correlations that peaked in the central sulcus, extending from motor- to more frontal cortex, and in left frontal and parietal cortex (compare Figure 4B and 4C).

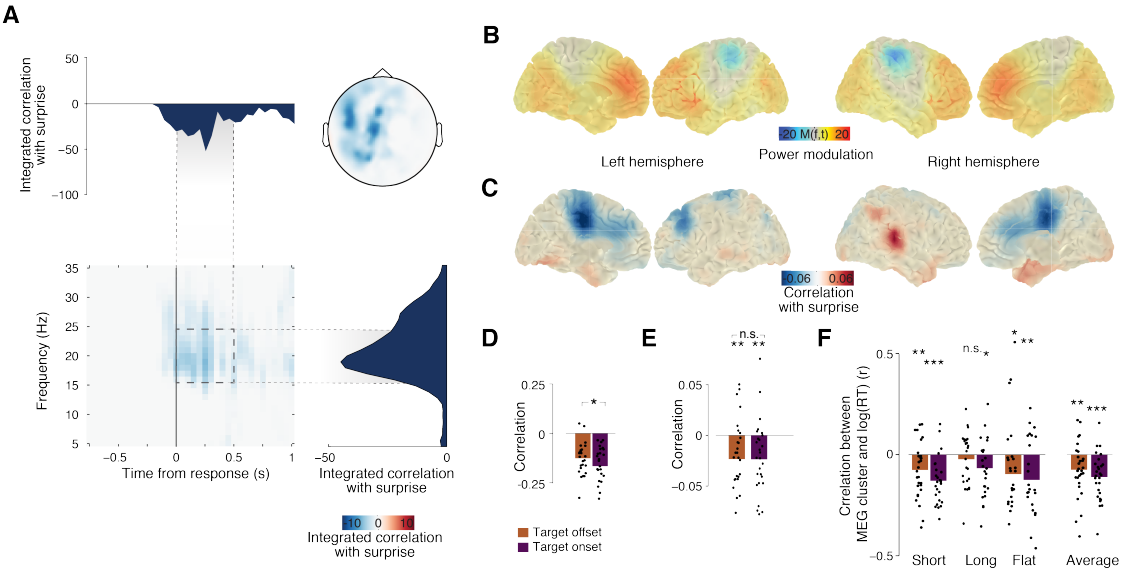

Figure 4: Widespread cortical beta-band transient driven by surprise. A. Exhaustive correlation between trial-to-trial measures of surprise and MEG power modulation in all sensors, time and frequency bins results in one cluster (cluster-based correction for multiple comparison, $p < 0.05$, two-sided) of negative correlation. Different panels show different dimensions of the cluster by integrating over the other dimensions; top left: time course, top right: spatial topography, bottom left: time-frequency representation; bottom right: frequency spectrum. **B.** Source reconstruction of the power modulation in the time window in which surprise-MEG correlation was strongest (dashed box in panel A). **C.** Source-reconstructed illustration of the correlation between transient modulation and trial-to-trial surprise depicted in panel B. These source maps are not statistically thresholded, but instead serve for comparing the correlation's spatial distribution with the transient power modulation in panel B. **D.** Comparison of correlation between surprise and power modulation between target offsets and onsets. Correlations are evaluated within the cluster from panel A. **E.** Leave-one-out cross-validation of the cluster found in panel A, separately for target offsets and onsets. Cluster-based permutation was performed on N-1 subjects and the average correlation in the resulting cluster was computed for the remaining subject (black dots); bars show averages over subjects. Correlation values were tested against 0 (permutation test; ** $p<0.01$). **F.** Correlation between MEG power in the cluster and log(RT) for separate distributions and average RT; *** $p<0.001$, ** $p<0.01$, * $p<0.05$; permutation tests.

The surprise-related cluster for target offsets exhibited a bimodal pattern in both the time and frequency domains: next to the peak around 20 Hz just after response, an additional peak was evident in the lowest frequency bin resolved (5 Hz) around 0.5 s before response. The topography showed peaks over parietal and occipital cortex and over left frontal cortex (Figure 5A). By contrast, the cluster for target onsets exhibited a single peak around 20 Hz just after response (Figure 5B). Taken together, our results suggest that perceptual surprise about both target on- and offsets elicited cortical transients in the beta-band. We consider them general dynamical correlates of temporal surprise monitoring. In addition, target offsets seem to have recruited additional processes expressed in the very low (<= 5 Hz) frequency range. This latter modulation might have been specific to the experimental context, which entailed monitoring illusory target disappearances in other runs, analyzed in previous reports (Kloosterman et al., 2015b; Meindertsma et al., 2017).

Finally, we asked whether the trial-to-trial fluctuations in beta-power modulations also predicted trial-to-trial variations in subjects' (log-transformed) RTs. Here, we used the Pearson correlation values (i.e., without regressing out RT; *Materials and Methods*). Just as surprise, beta-power in the cluster also robustly predicted RT (Figure 4F). These correlations were negative, as expected based on the negative correlation between surprise and beta-power (Figure 4A).

We excluded several potential confounding factors in separate correlation analyses. The negative correlation was present for (i) the 'raw' correlation between surprise and MEG power (data not shown), (ii) the partial correlation including the interval duration preceding the stimulus change as a covariate (data not shown), as well as (iii) the partial correlation using RT as a covariate. The latter is the more conservative measured and shown here.

No robust correlations between MEG baseline power and entropy

We did not find any evidence for a correlation of the raw baseline (-0.5 to 0 s with respect to stimulus change) MEG power with uncertainty, as measured in entropy. Correlations between entropy and MEG power spectra in the time window before stimulus change did not result in any significant (sensor-frequency) clusters that survived multiple-comparison correction (data not shown). It is likely that this lack of robust correlation reflected the continuous reduction in trial-to-trial variations of entropy over the course of each block (Figure 2C), which also translated into only minor effects on RT (Figure 3B).

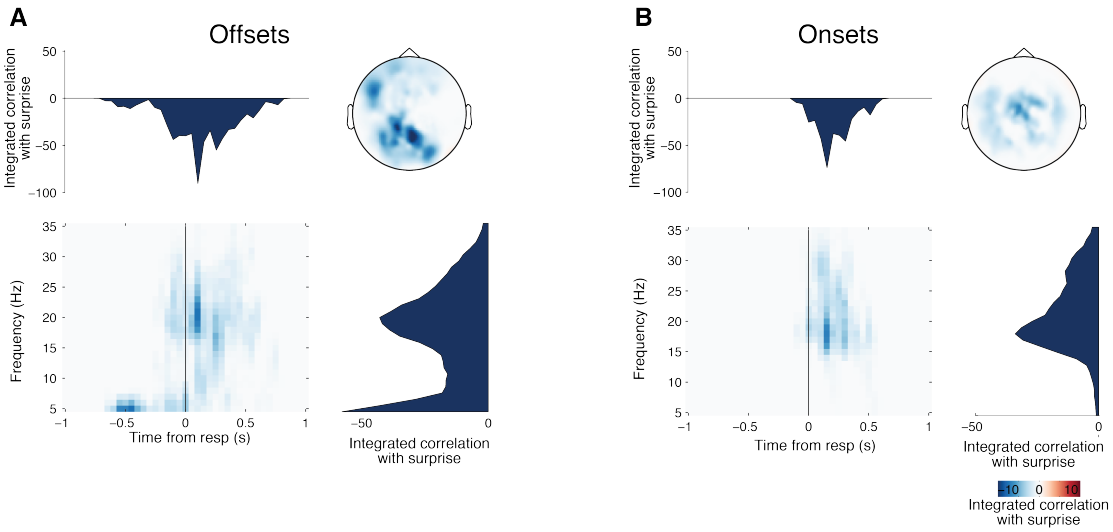

Figure 5: Separate correlations for offsets and onsets. Correlation analysis and cluster-based statistics performed separately for target offsets (**A**) and target onsets (**B**).

Discussion

In this study, we comprehensively mapped cortical transients elicited by surprise about the timing of sensory events. We used a Bayesian updating model to estimate trial-to-trial variations of surprise and correlated these to subjects' behavior as well as to neural dynamics, across the cortical surface. The model-derived surprise estimates robustly predicted across-trial and condition variations in RT. The surprise estimates also predicted transient suppressions of beta-band power in a widespread network comprising motor-, prefrontal and parietal cortical regions, predominantly in the left hemisphere. The model-derived surprise estimates were more closely related to both behavior and cortical dynamics than the mere trial-to-trial variations in externally observable interval timings.

The signatures of surprise we uncovered in the beta frequency band were quite similar around target on- and offset (Figure 5). This stands in sharp contrast to the opposite beta-band modulation during (illusory or veridical) target disappearances and reappearances, proposed to reflect a decision-related feedback signal to in visual cortex (Meindertsma et al., 2017). The beta-band

transients identified here likely reflected a distinct process that did not encode the content of the perceptual change, but rather the level of surprise about it.

We also measured average pupil responses during the three different conditions, the results of which were reported previously (Kloosterman et al., 2015a). This revealed differences in the transient pupil response amplitude for the three conditions that were in line with encoding surprise. In this study, we did not find such a relation between pupil size and the trial-to-trial metric of surprise. Although we here mapped the condition-wise distributions and trial-to-trial variations of interval timings with the same metric of surprise (Figure 2), it is possible that these two types of variations recruit separate neural computations. Future work using more widely spaced intervals to enable reliable trial-by-trial tracking of pupil response amplitudes is needed in order to better understand the role of pupil-linked arousal systems (de Gee et al., 2017) in the cortical correlates of the surprise computations which we have studied here.

Temporal expectation has been extensively studied in the context of temporal difference learning and the activity of the dopaminergic system of the midbrain (Hollerman and Schultz, 1998). Phasic neural responses in the striatum, as well as in dopaminergic nuclei, encode not only reward, but also the expected timing of reward arrival. The strength of these phasic neuronal responses inversely scales with condition in line with encoding the predictability of the reward timing, and it also predicted behavioral anticipation reward (i.e. licking behavior) in monkeys (Fiorillo et al., 2008). Our current study provides a critical complement to this previous work, by unraveling the cortex-wide dynamics elicited by surprising events. Our design did not involve rewards but rather neutral, yet behaviorally relevant sensory events. Still, it is likely that phasic neuromodulatory responses were nonetheless here driven by temporal expectation and surprise, possibly also in other systems such as the noradrenaline system (Dayan and Yu, 2006). Phasic neuromodulator release in cortex is a possible candidate source of the widespread transient modulations of beta-band activity observed here (Belitski et al., 2008; Donner and Siegel, 2011).

It is tempting to relate our results to conceptual accounts of the functional role of beta-band oscillations in the brain (Engel and Fries, 2010; Spitzer and Haegens, 2017). One account (Engel and Fries, 2010) holds that beta-band oscillations help maintain the current sensorimotor or cognitive state (termed the ‘status quo’). Another account (Spitzer and Haegens, 2017) holds that beta-band oscillations help activate the currently relevant task sets. In both frameworks, the need for maintaining the current status quo, or task set, is low when surprise (the violation of expectation, or probability of change in the environment) is high, in line with our observation of a suppression of beta-band oscillations under high surprise. Our current results point to the phasic release of neuromodulators, in particular dopamine, as a mechanistic source of the modulations of ongoing cortical beta-band oscillations, an idea not explicitly incorporated in either of these proposals so far.

Our current study provides a comprehensive picture of the cortical transients elicited by surprise, by systematically mapping these transients across the cortical surface and time-frequency plane. Previous work in humans has also studied neural correlates of model-derived measures of surprise, although this entailed surprise about stimulus identity and not timing. Electrophysiological work found surprise about cue identity to modulate the P3 component of the event-related potential as well as motor cortical excitability (Bestmann et al., 2008; Mars et al., 2008). Functional magnetic resonance imaging work linked surprise about the spatial location of stimuli to transients in posterior parietal cortex (O’Reilly et al., 2013). An EEG study dissociated oscillatory neural signatures of surprise and evidence

accumulation (Gould et al., 2012). This study also found surprise-related modulation of beta-band power primarily at frontal and parieto-occipital electrodes, but the underlying cortical distribution was not estimated. Future studies of surprise in other domains (e.g. about cue identity) should use a similar approach to assess if surprise-related cortical transients are domain-general or -specific.

Some previous work on the P3 component of the electroencephalogram (EEG) showed that this component is sensitive to the expected timing of events (Polich et al., 1994; Mertens and Polich, 1997; Polich, 2007). Although we used more continuous interval distributions that were computed from specific hazard functions, our task is similar to the ones used in these studies, and would thus likely induce a P3 response. We did not detect a clear P3-like component in the event-related fields (data not shown), in line with earlier studies (Schurger et al., 2015). It is possible that the surprise signature we observed in beta-band is a different reflection of the same widespread cortical process that also drives the P3. Our current signal is functionally most closely related to the P3b component, which is observed in response to the occurrence of attended but rare stimuli (Polich, 2007). Indeed, the P3b has been proposed to by driven by the phasic release of noradrenaline in cortex (Nieuwenhuis et al., 2005).

Another line of work has investigated the functional role of externally entrained low-frequency oscillations in temporal expectation. For fixed intervals, alpha phase in sensory cortices was found to be predictive of expected time of target arrival and lowered the threshold for sensory detection (Lakatos et al., 2008; Cravo et al., 2011, 2013; Rohenkohl and Nobre, 2011). Alpha oscillations might reflect rhythmic fluctuations in cortical excitability, entrained by rhythmic sensory input, which aids stimulus processing and perceptual performance (Schroeder and Lakatos, 2009). The high variability in interval durations (see Figure 1B,C inset) might explain the lack of alpha-band effects in our study. First, the range of possible durations was too broad to form predictions that fall within a specific phase of an alpha cycle. Second, even when oscillatory phase was modulated by temporal expectation in our task, the trial-to-trial variability would make it difficult to align trials and make these modulations visible.

While our current work presents an important first step towards unraveling the modulation of cortical dynamics by surprise, it is limited in that we only studied environments with constant statistical structure within each block. Once a posterior distribution has been learned, there remains no unexpected uncertainty, only expected uncertainty (Yu and Dayan, 2005). By contrast, the statistical structure of natural environments is often volatile. Richer experimental designs, that are volatile and include unmarked changes, allow for probing into richer, presumably hierarchical dynamics (Sugrue et al., 2004; Nassar et al., 2012; Meyniel et al., 2015). Our ongoing work aims to push beyond these limits by using richer environmental statistics that require more complex inference processes.

To conclude, we here uncovered a novel signature of temporal surprise that affected an elementary perceptual decision (target detection) and was characterized by a temporally focal, but spatially widespread, modulation of cortical population activity. This modulation might be instrumental in translating inferences about the behaviorally-relevant temporal structure into its consequences for action.

References

Andrieu C, de Freitas N, Doucet A, Jordan MI (2003) An Introduction to MCMC for Machine Learning. Mach Learn 50:5–43.

Belitski A, Gretton A, Magri C, Murayama Y, Montemurro MA, Logothetis NK, Panzeri S (2008) Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. J Neurosci 28:5696–5709.

789 Bestmann S, Harrison LM, Blankenburg F, Mars RB, Haggard P, Friston KJ,
790 Rothwell JC (2008) Influence of uncertainty and surprise on human
791 corticospinal excitability during preparation for action. Curr Biol 18:775–780.
792 Bonneh YS, Cooperman A, Sagi D (2001) Motion-induced blindness in normal
793 observers. Nature 411:798–801.
794 Cravo a. M, Rohenkohl G, Wyart V, Nobre a. C (2011) Endogenous modulation of
795 low frequency oscillations by temporal expectations. J Neurophysiol 106:2964–
796 2972.
797 Cravo AM, Rohenkohl G, Wyart V, Nobre AC (2013) Temporal expectation enhances
798 contrast sensitivity by phase entrainment of low-frequency oscillations in visual
799 cortex. J Neurosci 33:4002–4010.
800 Dayan P, Yu AJ (2006) Phasic norepinephrine: a neural interrupt signal for
801 unexpected events. Network 17:335–350.
802 de Gee JW, Colizoli O, Kloosterman NA, Knapen T, Nieuwenhuis S, Donner TH
803 (2017) Dynamic modulation of decision biases by brainstem arousal systems.
804 Elife 6:e23232.
805 Donner TH, Siegel M (2011) A framework for local cortical oscillation patterns.
806 Trends Cogn Sci 15:191–199.
807 Donner TH, Siegel M, Fries P, Engel AK (2009) Buildup of choice-predictive activity
808 in human motor cortex during perceptual decision making. Curr Biol 19:1581–
809 1585.
810 Efron B, Tibshirani RJ (1998) An Introduction to the Bootstrap. Boca Raton, FL:
811 Chapman & Hall/CRC Press.
812 Engel AK, Fries P (2010) Beta-band oscillations-signalling the status quo? Curr Opin
813 Neurobiol 20:156–165.
814 Feller W (1959) An Introduction to Probability Theory and Its Applications (Shewhart
815 WE, WIlks SS, eds)., Second. Wiley Publication in Statistics.
816 Fiorillo CD, Newsome WT, Schultz W (2008) The temporal precision of reward
817 prediction in dopamine neurons. Nat Neurosci 11:966–973.
818 Fisher RA (1915) Frequency Distribution of the Values of the Correlation Coefficient
819 in Samples from an Indefinitely Large Population. Biometrika 10:507.
820 Gibbon J, Malapani C, Dale CL, Gallistel C (1997) Toward a neurobiology of
821 temporal cognition: advances and challenges. Curr Opin Neurobiol 7:170–184.
822 Glaze CM, Kable JW, Gold JI (2015) Normative evidence accumulation in
823 unpredictable environments. Elife 4:1–27.
824 Gold JI, Shadlen MN (2007) The neural basis of decision making. Annu Rev
825 Neurosci 30:535–574.
826 Gould IC, Nobre AC, Wyart V, Rushworth MFS (2012) Effects of decision variables
827 and intraparietal stimulation on sensorimotor oscillatory activity in the human
828 brain. J Neurosci 32:13805–13818.
829 Gross J, Kujala J, Hamalainen M, Timmermann L, Schnitzler A, Salmelin R (2001)
830 Dynamic imaging of coherent sources: Studying neural interactions in the
831 human brain. Proc Natl Acad Sci U S A 98:694–699.
832 Hipp JF, Hawellek DJ, Corbetta M, Siegel M, Engel AK (2012) Large-scale cortical
833 correlation structure of spontaneous oscillatory activity. Nat Neurosci 15:884–
834 890.
835 Hollerman JR, Schultz W (1998) Dopamine neurons report an error in the temporal
836 prediction of reward during learning. Nat Neurosci 1:304–309.
837 Kloosterman N a., Meindertsma T, van Loon AM, Lamme V a. F, Bonneh YS, Donner
838 TH (2015a) Pupil size tracks perceptual content and surprise. Eur J Neurosci
839 41:1068–1078.
840 Kloosterman NA, Meindertsma T, Hillebrand A, Van Dijk BW, Lamme VAF, Donner
841 TH (2015b) Top-down modulation in human visual cortex predicts the stability

of a perceptual illusion. J Neurophysiol 113:1063–1076.
Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE (2008) Entrainment of Neuronal Oscillations as a Mechanism of Attentional Selection. Science (80-) 320:110–113.
Lee MD, Wagenmakers E-J (2013) Bayesian Cognitive Modeling: A Practical Course, 1st ed. New York: Cambridge University Press.
Luce RD (1986) Response Times: Their Role in Inferring Elementary Mental Organization, 1st ed. Oxford University Press, USA;
Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. J Neurosci Methods 164:177–190.
Mars RB, Debener S, Gladwin TE, Harrison LM, Haggard P, Rothwell JC, Bestmann S (2008) Trial-by-trial fluctuations in the event-related electroencephalogram reflect dynamic changes in the degree of surprise. J Neurosci 28:12539–12545.
Meindertsma T, Kloosterman NA, Nolte G, Engel AK, Donner TH (2017) Multiple Transient Signals in Human Visual Cortex Associated with an Elementary Decision. J Neurosci 37:5744–5757.
Mertens R, Polich J (1997) P300 from a single-stimulus paradigm: Passive versus active tasks and stimulus modality. Electroencephalogr Clin Neurophysiol - Evoked Potentials 104:488–497.
Meyniel F, Schlunegger D, Dehaene S (2015) The Sense of Confidence during Probabilistic Learning: A Normative Account. PLOS Comput Biol 11:e1004305.
Nassar MR, Rumsey KM, Wilson RC, Parikh K, Heasly B, Gold JI (2012) Rational regulation of learning dynamics by pupil-linked arousal systems. Nat Neurosci 15:1040–1046.
Nieuwenhuis S, Aston-Jones G, Cohen JD (2005) Decision making, the P3, and the locus coeruleus-norepinephrine system. Psychol Bull 131:510–532.
Nobre A, Correa a., Coull J (2007) The hazards of time. Curr Opin Neurobiol 17:465–470.
O'Reilly JX, Schüffelgen U, Cuell SF, Behrens TEJ, Mars RB, Rushworth MFS (2013) Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. Proc Natl Acad Sci U S A 110:E3660-9.
Oostenveld R, Fries P, Maris E, Schoffelen J-M (2011) FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Comput Intell Neurosci 2011:156869.
Plummer M (2003) JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. Proc 3rd Int Work Distrib Stat Comput (DSC 2003):20–22.
Polich J (2007) Updating P300: An integrative theory of P3a and P3b. Clin Neurophysiol 118:2128–2148.
Polich J, Eischen SE, Collins GE (1994) P300 from a single auditory stimulus. Electroencephalogr Clin Neurophysiol Evoked Potentials 92:253–261.
Riecke L, Sack AT, Schroeder CE (2015) Endogenous Delta/Theta Sound-Brain Phase Entrainment Accelerates the Buildup of Auditory Streaming. Curr Biol 25:3196–3201.
Rohenkohl G, Cravo AM, Wyart V, Nobre AC (2012) Temporal Expectation Improves the Quality of Sensory Information. J Neurosci 32:8424–8428.
Rohenkohl G, Nobre a. C (2011) Alpha Oscillations Related to Anticipatory Attention Follow Temporal Expectations. J Neurosci 31:14076–14084.
Schroeder CE, Lakatos P (2009) Low-frequency neuronal oscillations as instruments of sensory selection. Trends Neurosci 32:9–18.
Schurger A, Sarigiannidis I, Naccache L, Sitt JD, Dehaene S (2015) Cortical activity is more stable when sensory stimuli are consciously perceived. Proc Natl Acad Sci 112:E2083–E2092.

Spitzer B, Haegens S (2017) Beyond the Status Quo: A Role for Beta Oscillations in Endogenous Content (Re-) Activation. eneuro:ENEURO.0170-17.2017.

Sugrue LP, Corrado GS, Newsome WT (2004) Matching Behavior and the Representation of Value in the Parietal Cortex. Science (80-) 304:1782–1788.

Sutton RS, Barto AG (1998) Reinforcement Learning : An Introduction, First. Cambridge: MIT Press.

van Ede F, Niklaus M, Nobre AC (2017) Temporal Expectations Guide Dynamic Prioritization in Visual Working Memory through Attenuated α Oscillations. J Neurosci 37:437–445.

Van Veen BD, van Drongelen W, Yuchtman M, Suzuki A (1997) Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. IEEE Trans Biomed Eng 44:867–880.

Yu AJ, Dayan P (2005) Uncertainty, neuromodulation, and attention. Neuron 46:681–692.

Zou GY (2007) Toward Using Confidence Intervals to Compare Correlations. Psychol Methods 12:399–413.