

Disruption of the Right Temporoparietal Junction Impairs Probabilistic Belief Updating

Paola Mengotti,¹ Pascasie L. Dombert,¹ Gereon R. Fink,^{1,2} and Simone Vossel^{1,3}

¹Cognitive Neuroscience, Institute of Neuroscience and Medicine (INM-3), Research Centre Juelich, 52425 Juelich, Germany, ²Department of Neurology, University Hospital Cologne, 50937 Cologne, Germany, and ³Department of Psychology, University of Cologne, 50923 Cologne, Germany

Generating and updating probabilistic models of the environment is a fundamental *modus operandi* of the human brain. Although crucial for various cognitive functions, the neural mechanisms of these inference processes remain to be elucidated. Here, we show the causal involvement of the right temporoparietal junction (rTPJ) in updating probabilistic beliefs and we provide new insights into the chronometry of the process by combining online transcranial magnetic stimulation (TMS) with computational modeling of behavioral responses. Female and male participants performed a modified location-cueing paradigm, where false information about the percentage of cue validity (%CV) was provided in half of the experimental blocks to prompt updating of prior expectations. Online double-pulse TMS over rTPJ 300 ms (but not 50 ms) after target appearance selectively decreased participants' updating of false prior beliefs concerning %CV, reflected in a decreased learning rate of a Rescorla–Wagner model. Online TMS over rTPJ also impacted on participants' explicit beliefs, causing them to overestimate %CV. These results confirm the involvement of rTPJ in updating of probabilistic beliefs, thereby advancing our understanding of this area's function during cognitive processing.

Key words: attention; cue validity; learning rate; probabilistic inference; transcranial magnetic stimulation

Significance Statement

Contemporary views propose that the brain maintains probabilistic models of the world to minimize surprise about sensory inputs. Here, we provide evidence that the right temporoparietal junction (rTPJ) is causally involved in this process. Because neuroimaging has suggested that rTPJ is implicated in divergent cognitive domains, the demonstration of an involvement in updating internal models provides a novel unifying explanation for these findings. We used computational modeling to characterize how participants change their beliefs after new observations. By interfering with rTPJ activity through online transcranial magnetic stimulation, we showed that participants were less able to update prior beliefs with TMS delivered at 300 ms after target onset.

Introduction

The right temporoparietal junction (rTPJ) responds to task-relevant stimuli, particularly when unexpected. Therefore, in the domain of attention, the main theoretical framework considers rTPJ as part of the ventral attentional system, with the specific role of reorienting attention toward upcoming stimuli, acting as a kind of “circuit-breaker” over top-down attention sustained by

the dorsal system (Corbetta and Shulman, 2002; Corbetta et al., 2008).

However, the role of rTPJ is particularly controversial and debated, as its activation is not restricted to the attentional domain. rTPJ has been shown to be involved in theory of mind and empathy (Decety and Lamm, 2007; Mitchell, 2008; Graziano and Kastner, 2011), as well as body representation and body awareness (Blanke et al., 2005; Tsakiris et al., 2008). The “contextual updating” hypothesis (Geng and Vossel, 2013) offers a unifying framework to interpret such broad involvement of the rTPJ across diverse cognitive domains. Contextual updating is particularly important when unexpected stimuli occur; therefore, this hypothesis proposes that the rTPJ is responsible for updating internal models of the current behavioral context to the aim of optimizing appropriate responses.

To this end, a precise characterization of the timing of the rTPJ response after target appearance is central: if rTPJ acts as a “circuit-breaker” over the dorsal system, its activation should

Received Nov. 30, 2016; revised April 10, 2017; accepted April 12, 2017.

Author contributions: P.M., G.R.F., and S.V. designed research; P.M. and P.L.D. performed research; P.M., P.L.D., and S.V. analyzed data; P.M., P.L.D., G.R.F., and S.V. wrote the paper.

This work was supported by Federal Ministry of Education and Research BMBF 01GQ1401 to S.V. We thank our colleagues from the Institute of Neuroscience and Medicine (INM-3) for valuable support and discussions.

The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Paola Mengotti, Cognitive Neuroscience, Institute of Neuroscience and Medicine (INM-3), Research Centre Juelich, Leo-Brandt-Str. 5, 52425 Juelich, Germany. E-mail: p.mengotti@fz-juelich.de.

DOI:10.1523/JNEUROSCI.3683-16.2017

Copyright © 2017 the authors 0270-6474/17/375419-10\$15.00/0

occur in an early time-window of target processing. However, if rTPJ is involved in contextual updating, its response should occur in a later time-window, as updating should occur after the initiation of attentional reorienting. More specifically, because ERP studies have linked contextual updating to the P300 component (Polich, 2007; Kolossa et al., 2015), we expected an involvement of rTPJ from 300 ms onwards.

Online transcranial magnetic stimulation (TMS) is a technique that allows to address this question because of both its temporal precision and its causal role in interfering with the activity of a circumscribed brain region. Online TMS applied to the angular gyrus (i.e., a region adjacent to rTPJ) revealed that the processing of invalidly cued targets could be impaired in two distinct time windows, from 90 to 120 ms and from 210 to 240 ms after target onset (Chambers et al., 2004).

Here, we used online double-pulse TMS to directly test the timing and causal role of the rTPJ response for contextual updating after the trialwise observation of cue-target outcomes. Participants performed a location-cueing paradigm (Posner, 1980), which required updating of prior expectations concerning the validity of the spatial cue for predicting the target's location. In spatial cueing tasks, behavioral responses (Vossel et al., 2014) and rTPJ activity (Vossel et al., 2015; Dombert et al., 2016) are highly sensitive to changes in the probability with which the cue predicts the target location, and this inference process has been described using computational learning models. In the present study, trial-by-trial learning of cue-target contingencies, and hence belief updating, was assessed by estimating the learning rate of a Rescorla–Wagner model (Rescorla and Wagner, 1972) on the basis of response speed (RS; inverse reaction time [RT]). In each experimental block, prior probabilistic beliefs were manipulated providing either true or false explicit information about the percentage of cue validity (%CV). We hypothesized that higher learning rates should be observed in blocks with false prior information. Moreover, by interfering with neural activity of the rTPJ with TMS applied 300–500 ms after target onset, we expected to impair contextual updating in false %CV blocks, without impairing the overall performance or reorienting of spatial attention toward invalidly cued targets. More specifically, we predicted that late TMS over rTPJ would decrease the Rescorla–Wagner learning rate, especially in the false %CV blocks, where contextual updating is required. If rTPJ is also involved in early circuit breaking and attentional reorienting, TMS in an earlier time window (50–250 ms after target appearance) should modulate the RS difference between invalid and valid trials, regardless of the current probabilistic context.

Materials and Methods

Experimental design and statistical analysis. Twenty-one healthy volunteers participated in the study. They had normal or corrected-to-normal vision, no past or present neurological or psychiatric condition, and they were naive to the purpose of the experiment. All participants but one were right-handed, as assessed with the Edinburgh Handedness Inventory (Oldfield, 1971). One participant was ambidextrous with a laterality quotient of 40 (i.e., on the boundary to right-handedness). They all gave written informed consent. The study had been approved by the ethics committee of the German Psychological Society and was performed in accordance with the Code of Ethics of the World Medical Association (Declaration of Helsinki).

After data acquisition, 3 participants had to be excluded from further analysis: 2 due to bad task performance (>2 SDs below the mean accuracy of all participants) and 1 for a technical problem with the recording of the manual responses. Therefore, the final sample comprised 18 par-

ticipants (11 females; age range: 18–33 years; mean age: 24.4 years; 17 right-handed).

We used a location-cueing paradigm with central cueing (Posner, 1980) (Fig. 1). Stimuli were presented on a 15.6 inch monitor (spatial resolution 1024 × 768 pixels, refresh rate 60 Hz) with a viewing distance of 90 cm. A central diamond was presented on a gray background for the total duration of the task, serving as a fixation point. On each trial, a triangle appeared for 400 ms behind the fixation diamond creating an arrowhead pointing to the left or to the right to indicate the hemifield in which the target would appear, therefore acting as spatial cue. After an 800 ms stimulus onset asynchrony, two diamonds appeared for 350 ms on the left or right side of the central diamond (6.6° eccentric in each visual field). The target diamond was missing a fragment in its upper or lower part. Participants were instructed to press a button with the index or middle finger of their right hand (i.e., ipsilateral to the TMS stimulated hemisphere) to indicate whether the upper or lower fragment of the target was missing. The response mapping was counterbalanced across participants. The intertrial interval was 2000 ms (Fig. 1a).

In each session, two runs of eight experimental blocks were presented to the participants. %CV was manipulated between blocks but was kept constant within each block. %CV within each block amounted to ~90% (87.5%), ~70% (71%), ~30% (29%), or ~10% (12.5%), respectively. In the 30% and 10% CV blocks, the cue was counter-predictive, as the majority of trials was invalid. As the %CV ranged from highly predictive (e.g., ~90%) to highly counter-predictive (e.g., ~10%) in the present task, valid and invalid trials are, on average, not necessarily associated with higher and lower RSs, respectively, as in a traditional Posner task. From the structure of the task, no overall mean differences should be detected between valid and invalid trials when looking at all trials, and if so, these differences can be attributed to a general tendency to follow the cue even when it is not predictive (Eimer, 1997). The %CV was explicitly reported to the participants at the beginning of each block. However, in half of the blocks (i.e., 4 blocks), the provided information was false, thereby creating blocks in which the prior expectation of the %CV was incorrect. In these false blocks, the expected %CV was inverted with respect to the true %CV (e.g., the expected %CV was ~30%, but the actual %CV was ~70%).

Participants were aware that, in some blocks, false %CV information would be provided, but they were not informed about the number of blocks or about the difference between the false and true %CV. Therefore, they were instructed to use the cue according to how much they “trust” it and to estimate the true %CV. Indeed, participants learned the actual cue-target contingencies, as reflected in the modulation of the validity effect (RT difference between invalid and valid trials) by %CV in the separate blocks (Fig. 1b). At the end of each block, participants were asked to report their explicit belief about the true %CV, using an interval scale ranging from 10% to 90% (in 10% steps). Each block comprised 48 trials resulting in a total of 768 trials for each session (384 trials per run). Breaks were possible at the end of each block, and a longer break was provided between the two runs. RTs and accuracy were recorded.

Each participant completed a short practice session at the beginning of each experimental session, divided into two short runs. The first run consisted of 48 trials with true 80% CV. The second run comprised three blocks of 48 trials each: with two true blocks (true 80% CV) and one false block (false 80% [true 20%] CV).

Each participant was presented with the same sequence of trials and the same sequence of blocks within each run. This is a standard procedure in computational studies of learning processes that require inference on conditional probabilities in time series. A single experimental session lasted ~1 h 30 min to 2 h.

T1-weighted structural MRI scans for each participant were used to perform neuronavigated TMS over the rTPJ (MNI coordinate: $x = 56$, $y = -44$, $z = 12$; Fig. 1c). The coordinate was chosen from the fMRI study of Vossel et al. (2015), as local maximum for the parametric modulation effect for probability-dependent attention. The coordinate was back-normalized from the MNI template to the single subject's brain, using the inverse normalization parameters from the normalization of each participant's brain to the MNI template. The anatomical localization of the stimulation site, located in the posterior part of the superior

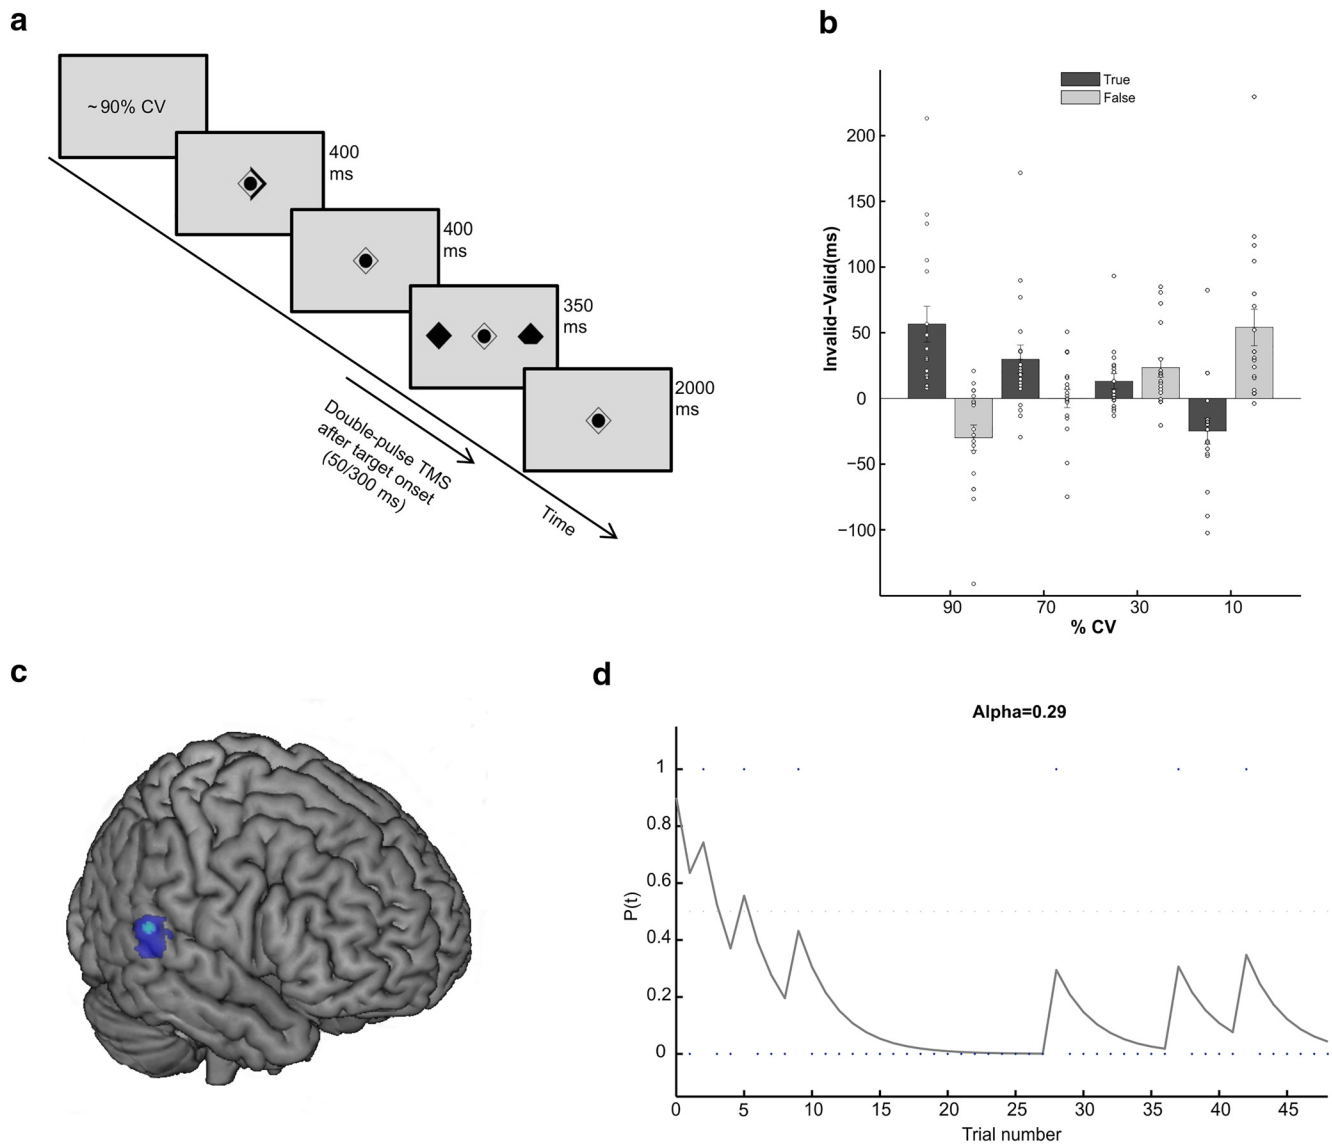


Figure 1. Experimental paradigm and illustration of the modeling approach. **a**, Timeline of the experimental paradigm for a validly cued trial. At the beginning of each block, the %CV (either true or false) was shown. This value was used as prior before the observation of the first trial in the modeling approach. On each trial, participants indicated whether the upper or lower corner of the target was missing. Central fixation was maintained throughout the experiment. Double-pulse online TMS (with 100 ms interstimulus interval) was delivered in every trial, either at 300 ms or at 50 ms after target appearance, in different sessions. Both active TMS and sham stimulation were performed, in two different runs, in each session. **b**, Validity effects (mean \pm SEM) for each true and false %CV block in the pooled sham condition. Participants were able to use the cue and to estimate the true cue-target contingencies, as shown by the linear pattern of the data. White circles indicate single-subjects data. **c**, The MNI coordinate of the rTPJ where TMS was applied is shown in light blue ($x = 56, y = -44, z = 12$) in a standard brain, together with the fMRI cluster (blue) for the parametric modulation effect for probability-dependent attention found by Vossel et al. (2015). **d**, Example of one participant's trial-by-trial changes in the probability estimate that the cue will be valid ($P(t)$) as derived from the learning model in a block with a false 90% CV prior. Dots indicate valid trials (outcome = 1) and invalid trials (outcome = 0) in this experimental block.

temporal gyrus, was then checked visually in the single subjects' brains and adjusted in case the normalization located the coordinate in an anatomically different region ($n = 9$ subjects; mean adjustment = 4 ± 1.4 mm).

TMS-coil placement was guided and continuously monitored throughout the whole experimental session via a neuronavigation system (Brainsight2, Rogue Research). A Magstim SuperRapid stimulator (Magstim) equipped with a 70 mm figure-of-eight coil was used to apply double-pulse stimulation over the rTPJ. Two TMS pulses were delivered with an interstimulus interval of 100 ms after 50 ms or after 300 ms from target onset.

The stimulation intensity was adjusted for each participant by incrementally increasing it, starting from 30% of the stimulator output in increments of 3% up to a maximum of 50%. If participants reported discomfort or muscle twitches at a given intensity, the stimulation was decreased in steps of 1% to reach the maximum stimulation intensity

possible without creating discomfort (for a similar procedure, see Azanón et al., 2010). The mean stimulation intensity was 38% ($\pm 6\%$, SD). This intensity was kept constant between the two sessions of the same participant, unless discomfort or muscle twitching was reported during the preparation for the second session. With these procedures, participants reported no discomfort or adverse effects during TMS in the task. Results from a pilot study performed on 6 participants showed that this procedure led to intensities between 70% and 87% of the resting motor threshold ($78 \pm 6\%$, mean \pm SD).

During the double-pulse TMS stimulation, the coil was held tangential to the scalp, with the handle pointing backward. Participants were provided with earplugs to reduce the noise produced by the stimulating coil. In the sham condition, the coil was held perpendicularly to the surface of the scalp over the same target position, to mimic the noise and the mechanical vibration of TMS without actually stimulating brain tissue. During the whole experimental session, an additional ineffective TMS coil

was placed over the left hemisphere in the analogous location of the stimulating coil (tilted perpendicularly to the surface of the scalp). This was done to prevent that the sound of the TMS stimulation served as an exogenous cue, as location-cueing paradigms can be sensible to lateralized auditory stimulation (Chica et al., 2011). Therefore, the ineffective coil on the left side was locked to the timing of the right-side stimulation and set at the same intensity of the actual stimulation.

In a within-subject design, participants completed four runs of the experimental task, distributed over 2 d. The experiment consisted of two sessions with two different timings of the TMS and sham stimulation. The sessions were performed on 2 different days (intersession interval: 1–8 d) at the same daytime. The run order and the association with the TMS protocol were counterbalanced across participants. In each of the two sessions, one run was performed with active TMS and one run with sham stimulation. Only one timing was tested per session (either 300 or 50 ms after target onset, both for TMS and sham) and the second timing was tested in the second session.

RS (1/RT) was assessed for each trial. RSs were used because, in contrast to RTs, they are more normally distributed (Carpenter and Williams, 1995; Brodersen et al., 2008). Anticipations (RT < 100 ms), misses, and incorrect responses were excluded from the analyses of RSs.

A Rescorla–Wagner model was applied to trialwise RSs in the different blocks, to quantify belief updating concerning %CV in true and false blocks. The choice of a Rescorla–Wagner model, instead of the hierarchical Bayesian model used in previous studies (Vossel et al., 2014, 2015), was motivated by the different structure of the task and the substantially smaller number of trials entering the modeling. In particular, we used short separate blocks (due to the technical and practical constraints imposed by the use of online TMS), which were modeled independently instead of one long continuous sequence with hidden changes in %CV as in some of the previous work. Here, we used the explicit prior information concerning %CV provided to the subjects before each block as the prior mean, but we also allowed the model to adjust this starting value to the observed data. Although the Rescorla–Wagner model is a heuristic model, the learning rate has been found to be significantly correlated with the parameter governing updating of cue validity in a hierarchical Bayesian model (Vossel et al., 2014). In both types of models, updating is determined by the weighting of prediction errors (the discrepancy between observed and predicted outcomes) by a learning rate. Hierarchical Bayesian models have been proven particularly useful in volatile environments where cue-target contingencies change within a sequence of several hundreds of trials, and where continuous changes in beliefs about the volatility of the cue validity could be additionally captured. However, this is not required in our short blocks with constant %CV. Moreover, it should be noted that we modeled each single block separately, so that any changes in learning across blocks, as well as the hypothesized difference in the learning rates for true and false blocks in the sham condition, are accounted for by our modeling approach.

According to the Rescorla–Wagner model, the update of the belief that the cue will be valid in a given trial is the product of a learning rate α and the prediction error $\delta^{(t)}$ (i.e., the difference between the observed and the predicted outcome in the respective trial t).

The new revised prediction after observation of the trial t , $P^{(t)}$, is then given by the sum of the prediction from the previous trial and the product of learning rate and prediction error as follows:

$$P^{(t)} = P^{(t-1)} + \alpha \delta^{(t)}$$

Thus, the learning rate α determines the extent to which prediction errors may change the subject's belief from trial to trial. At the same time, algebraic manipulation of the above formula reveals that the learning rate α dictates the extent to which past events influence the subject's prediction because it determines the steepness of the exponential decay of the influence of preceding trials (Rushworth and Behrens, 2008).

Single-trial RSs were used to estimate the learning rate α in each experimental block (Fig. 1*d*). The observation or response model, which was used to map from the subject's belief about cue validity to observed responses (RSs), assumed a linear relationship between $RS^{(t)}$ and the prediction before the observation of the outcome of the trial $P^{(t-1)}$ (for a similar procedure, see Vossel et al., 2014) as follows:

$$RS^{(t)} = \begin{cases} \zeta_{1_valid} + \zeta_2 P^{(t-1)} & \text{for valid trials} \\ \zeta_{1_invalid} + \zeta_2 (1 - P^{(t-1)}) & \text{for invalid trials} \end{cases}$$

ζ_{1_valid} , $\zeta_{1_invalid}$, and ζ_2 are additional subject-specific parameters that are estimated from the data. Whereas ζ_{1_valid} and $\zeta_{1_invalid}$ determine the constants of the linear equation (i.e., the overall levels of RSs), ζ_2 parameterizes the slope of the affine function (i.e., the strength of the increase in RS with increased estimated cue validity $P^{(t-1)}$). The learning rate α , the starting values $P^{(0)}$, and the three parameters from the observation model were estimated from trialwise RSs using variational Bayes as implemented in the HGF toolbox (<http://www.translationalneuromodeling.org/tapas/>) running on MATLAB (R2014a, The MathWorks). Variational Bayes optimizes the (negative) free-energy F as a lower bound on the log evidence, such that maximizing F minimizes the Kullback–Leibler divergence between exact and approximate posterior distributions or, equivalently, the surprise about the inputs encountered (for details, see Friston et al., 2007).

To test for the appropriateness of the Rescorla–Wagner model, we formally compared our model with an alternative model with no learning ($\alpha = 0$; i.e., only RS for valid and invalid trials is estimated).

In a second analysis, we tested the possibility that, in the blocks with false priors, participants directly switched to the inverse probability after a couple of observations rather than gradually learned the probability. Here, we formally compared the Rescorla–Wagner model with two alternative models. In these two alternative models, switches from the prior to the inverse probability would occur after the first quarter (12) or after half (24) of the 48 trials in one block.

For both control analyses, the log-model evidences resulting from the variational Bayesian model inversion of the competing models were compared using random-effects Bayesian Model Selection (BMS) (Stephan et al., 2009) for all blocks in the first comparison and for false blocks only in the second comparison, and separately for the sham, TMS300, and TMS50 conditions. Protected exceedance probabilities and the Dirichlet parameter estimates (representing the effective number of subjects in which a given model generated the observed data) of the competing models are reported. The protected exceedance probability (Rigoux et al., 2014) uses the Bayesian omnibus risk to compute a Bayesian model average of the exceedance probability, that is, the estimation of the likelihood of a particular model being the best compared with any other model, given the data. Importantly, random-effects BMS treats the model itself as being selected probabilistically by each subject in the population, enabling group-level inference while accounting for interindividual differences (e.g., the optimal model can vary across subjects).

The learning rate α was averaged separately for the blocks with true and false prior information concerning %CV and tested for normality (skewness and kurtosis within ± 2) (George and Mallery, 2010). For the group-level analyses, we first compared the two sham conditions to investigate putative differences caused by the sound of the TMS at two different timings. Therefore, a within-subject ANOVA with the factors timing (300 ms, 50 ms) and prior (true, false) was performed. No main effect of timing or significant interaction with timing was present. Therefore, the two sham conditions were pooled together in the analysis presented in Results, so that the learning rates were analyzed with a within-subject ANOVA with the factors stimulation (sham, TMS300, TMS50) and prior (true, false). All group-level analyses were performed with SPSS (SPSS Statistics for Windows, version 22.0, IBM). Results from these analyses are reported at a significance level of $p < 0.05$ after Greenhouse–Geisser correction. Paired-sample or one-sample t tests (with Bonferroni correction) were computed to interpret interaction effects. In addition, Bayesian paired-sample or one-sample t tests (nondirectional, Cauchy prior = 0.707) were computed in JASP (version 0.8.1.1) and Bayes factors are reported (B_{10} for all comparisons and B_{01} in those with $B_{10} < 3$). The Bayes factor B_{10} reflects the evidence for H1 (i.e., the data from the two conditions are different) compared with H0 (i.e., the data from the two conditions are not different). B_{01} reflects the evidence in favor of the alternative hypothesis H0. Bayes factors > 3 , > 10 , and > 30 indicate moderate, strong, or very strong evidence for a difference, respectively. Bayes factors between 1 and 3 indicate inconclusive evidence.

Because two conditions showed mild violations from normality, we also used nonparametric Wilcoxon tests as control analyses. The results of the ANOVA were confirmed by the nonparametric tests; therefore, only the former will be presented.

In addition to the analysis with the pooled sham conditions, we ran an analysis in which the α values of each sham condition were subtracted from the respective TMS condition (TMS300-SHAM300 and TMS50-SHAM50). These differences were analyzed with a within-subject ANOVA with the factors stimulation (TMS300-SHAM300, TMS50-SHAM50) and prior (true, false). One condition showed a mild violation from normality, so we added nonparametric Wilcoxon tests as control analyses. The results of the ANOVA were confirmed by the nonparametric tests; therefore, only the former will be presented.

As a further control analysis, we ran a within-subject ANOVA with the factors stimulation (sham, TMS300, TMS50), prior (true, false), and block (high difference, 90%, 10%; low difference, 70%, 30%) on α values. α values in true and false 90% and 10% blocks and α values in true and false 70% and 30% blocks were averaged together.

Mean scores of explicit evaluation of %CV, as given by the participants at the end of each block, were analyzed with a within-subject ANOVA with the factors stimulation (sham, TMS300, TMS50) and prior (true, false). Two conditions showed mild violations from normality, so we added nonparametric Wilcoxon tests as control analyses. The results of the ANOVA were confirmed by the nonparametric tests; therefore, only the former will be presented. An additional stimulation (TMS300-SHAM300, TMS50-SHAM50) \times prior (true, false) ANOVA was performed using the differences between the TMS and sham condition in each timing.

Mean RSs and accuracy scores expressed in percentage of correct responses were used in two additional within-subject ANOVAs with the factors stimulation (sham, TMS300, TMS50), prior (true, false), and validity (valid, invalid). There were no deviations from normality for the RS data. An additional stimulation (TMS300-SHAM300, TMS50-SHAM50) \times prior (true, false) ANOVA was performed using the differences between the TMS and sham condition in each timing.

Eye movement data recording. Eye movements were monitored with an Eye-Link 1000 (SR Research) eye-tracking system with a sampling rate of 500 Hz. At the start of the experiment, a calibration and validation of the eye-tracker were performed (validation error $<1^\circ$ of visual angle). Analysis of the data was performed using MATLAB (R2014a, The MathWorks). The critical period analyzed for gaze deviations from the center was the time window between the presentation of the cue and the target display (cue-target period). Saccades were identified as gaze deviations from fixation $>1.5^\circ$ visual angle in the cue-target period, and they were determined and expressed as a percentage score over the total number of trials. Differences in percentage of saccades were analyzed with a within-subject ANOVA with factors stimulation (sham, TMS300, TMS50), prior (true, false), and validity (valid, invalid). One participant had to be excluded from this analysis due to bad quality of the signal; therefore, eye-movement data from 17 of the 18 participants were analyzed.

Participants maintained fixation on average in 98.4% ($\pm 0.5\%$, SEM) of the trials. Analysis of the eye-movement data with a within-subjects ANOVA with factors stimulation (sham, TMS300, TMS50), prior (true, false), and validity (valid, invalid) did not reveal any significant effect of stimulation or any interaction with stimulation.

Results

TMS300-induced reduction of belief updating (learning rate α)

Figure 2*a* shows the learning rate α in the sham, TMS300, and TMS50 conditions for true and false blocks. The within-subject ANOVA with the factors stimulation (sham, TMS300, TMS50) and prior (true, false) on the learning rate α revealed a main effect of prior ($F_{(1,17)} = 39.6, p = 0.000008$) with higher learning rates in false blocks, but also a significant stimulation \times prior interaction ($F_{(1.88,32)} = 4.15, p = 0.027$; Fig. 2*a*). We ran paired *t* tests to interpret the interaction (nine comparisons, Bonferroni-corrected

threshold $p < 0.0056$). We found a significant difference between α in blocks with true and false prior both for sham ($t_{(17)} = -6.4, p = 0.000006, B_{10} = 3354$) and for TMS50 ($t_{(17)} = 4.7, p = 0.0002, B_{10} = 152$), with higher learning in false blocks, but no significant difference between false and true blocks for TMS300 ($t_{(17)} = -0.76, p = 0.46, B_{10} = 0.31, B_{01} = 3.2$). Moreover, α in false blocks in TMS300 was significantly different from sham ($t_{(17)} = -3.8, p = 0.0014, B_{10} = 27$), with reduced learning for TMS300. In addition, we found a trend toward significance between α in the TMS300 and TMS50 in false blocks ($t_{(17)} = -2.67, p = 0.016, B_{10} = 3.5$), again with slower learning for TMS300.

The analysis on the differences between α in each TMS and sham condition using a within-subject ANOVA with the factors stimulation (TMS300-SHAM300, TMS50-SHAM50) and prior (true, false) revealed a significant stimulation \times prior interaction ($F_{(1,17)} = 4.8, p = 0.043$; Fig. 2*b*). *Post hoc* paired *t* tests (four comparisons, Bonferroni-corrected threshold $p < 0.0125$) showed a significant difference between α in blocks with true and false prior for TMS300-SHAM300 ($t_{(17)} = -2.9, p = 0.01, B_{10} = 5.3$), with a relative increase in α with TMS compared with sham in blocks with true prior and a relative decrease in α with TMS compared with sham in blocks with false prior. We also tested for the deviance of these differences from zero with one-sample *t* tests. Here, the only effect that was different from zero was in the TMS300-SHAM300 in blocks with false prior ($t_{(17)} = -2.2, p = 0.04, B_{10} = 1.7, B_{01} = 0.6$), with a TMS-induced decrease in α ; however, this effect did not pass the threshold for Bonferroni correction (two comparisons, threshold $p < 0.025$).

Although the critical test for our hypothesis was the modulation of the learning rate α as described above, we wanted to find signatures of this TMS effect in a model-free analysis. To this aim, we ran single-subject linear regressions on RS with the number of repetitions of the same (most frequent) trial type (valid trials in the predictive blocks and invalid trials in the counter-predictive blocks, respectively), to test whether the acceleration of RS with repeating trials was present in all stimulation conditions. If participants show fast learning, RS should increase with an increasing number of repetitions of the same trial type, and the acceleration should be captured by the β parameter (slope) of the regression. The within-subject ANOVA with the factors stimulation (sham, TMS300, TMS50) and prior (true, false) revealed a trend toward significance for the stimulation \times prior interaction ($F_{(1.5,25.5)} = 3, p = 0.08$; Fig. 2*c*). We ran paired *t* tests to interpret the trend found in the interaction (nine comparisons, Bonferroni-corrected threshold $p < 0.0056$). We found a significant difference between β in blocks with true and false prior for TMS300 ($t_{(17)} = 3.5, p = 0.0027, B_{10} = 16$), with higher acceleration for the blocks with true prior. The same comparisons for sham ($t_{(17)} = -0.23, p = 0.82, B_{10} = 0.2, B_{01} = 4$) and for TMS50 ($t_{(17)} = 0.07, p = 0.94, B_{10} = 0.2, B_{01} = 4.1$) were not significant. In addition, we found a trend toward significance between β in the TMS300 and sham in false blocks ($t_{(17)} = 2.7, p = 0.015, B_{10} = 3.8$), with lower acceleration for TMS300, and between β in the TMS300 and TMS50 in false blocks ($t_{(17)} = 2.16, p = 0.046, B_{10} = 1.5, B_{01} = 0.6$), with lower acceleration for TMS300.

In what follows, we present additional evidence for the validity of our modeling approach. The mean adjustment (averaged over all CV blocks) from the prior value (e.g., from 90% CV) in the perceptual model was -0.02% ($\pm 3.04\%$ SD) for sham, 2.04% ($\pm 6.7\%$) for TMS300, and -1.21% ($\pm 5.61\%$) for TMS50.

Table 1 shows the values of the remaining parameters of the response model, which determine the constant of the linear equa-

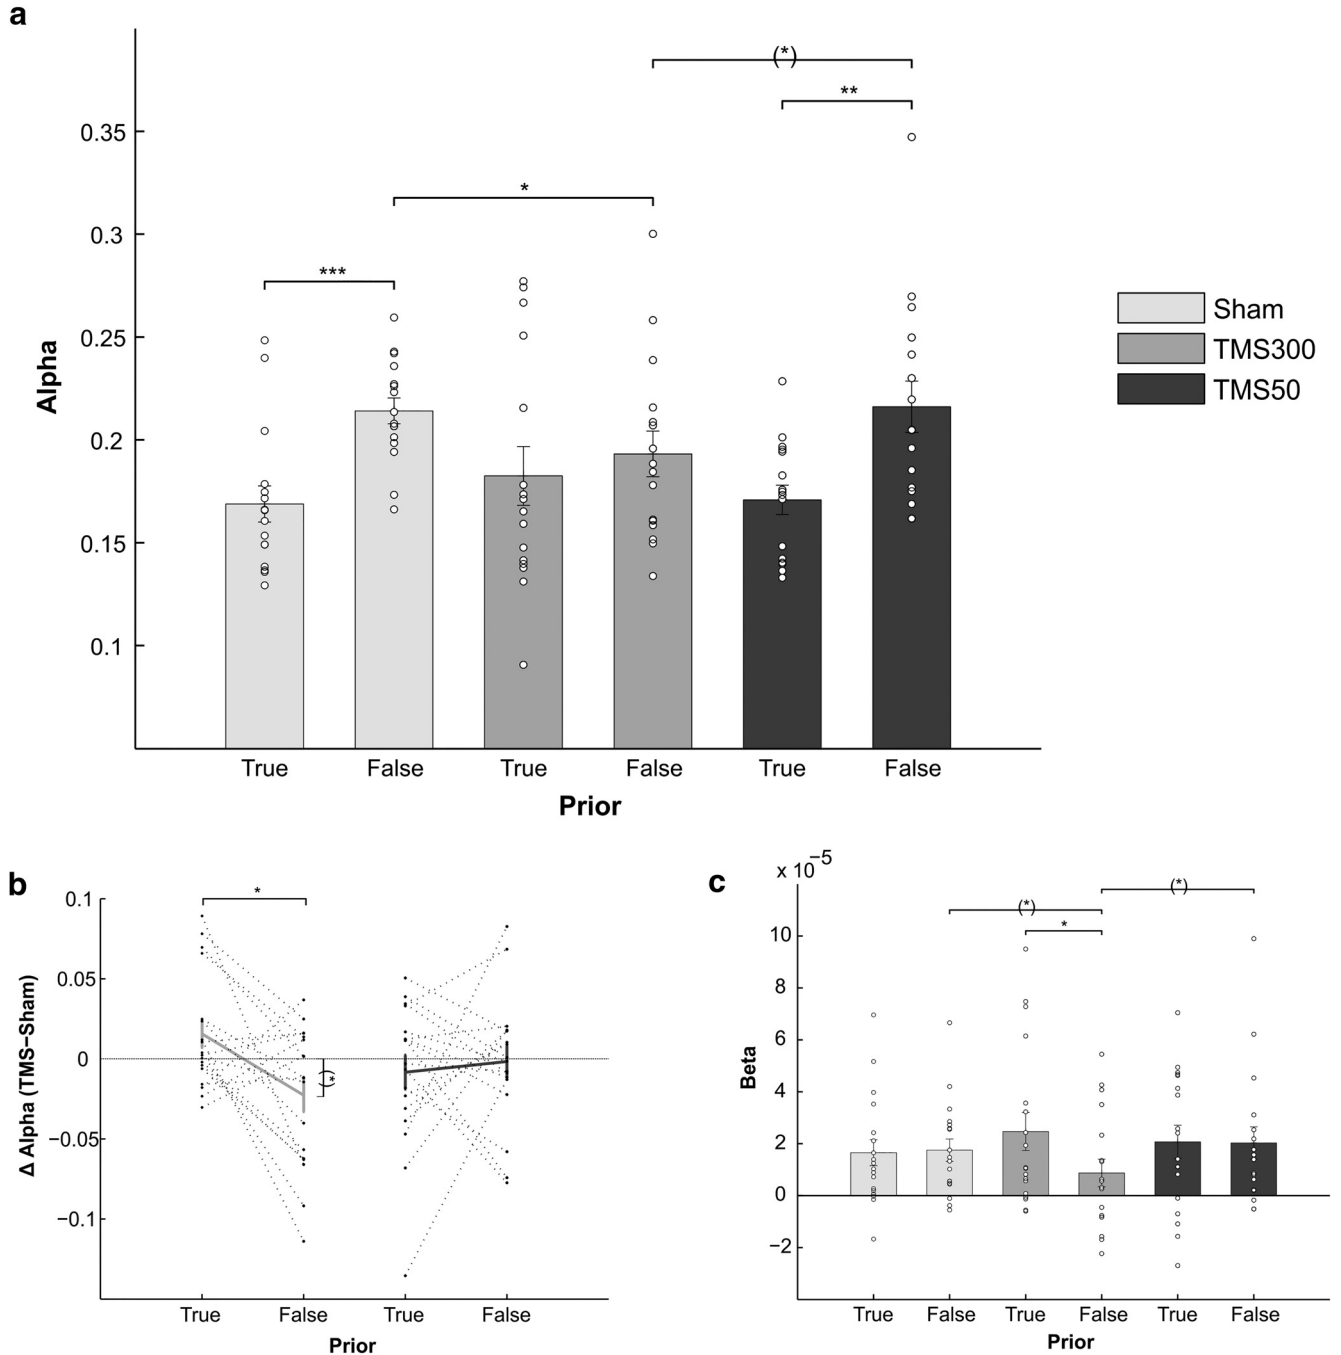


Figure 2. Belief updating results. **a**, Learning rate α for all stimulation (TMS300, TMS50, sham) and prior (true, false) conditions. Error bars indicate SEM. White circles indicate single-subjects data. **b**, Mean deviations from baseline (TMS300-SHAM300 and TMS50-SHAM50), indicating time-specific effects. Dotted lines indicate single-subjects data. **c**, Mean β values indicating the acceleration of RSs for increasing number of consecutive repetitions of the same type of trials for true and false prior in all stimulation conditions. White circles indicate single-subjects data. *** $p < 0.0001$. ** $p < 0.001$. * $p < 0.05$. (*) $p < 0.05$, not passing the Bonferroni-corrected threshold for multiple comparisons.

Table 1. Mean (\pm SEM) values for the parameters of the response model

	True prior blocks			False prior blocks		
	Sham	TMS300	TMS50	Sham	TMS300	TMS50
ζ_{1_valid}	1.8E-03 (\pm 9.23E-05)	1.76E-03 (\pm 9.84E-05)	1.83E-03 (\pm 9.93E-05)	1.79E-03 (\pm 9.3E-05)	1.83E-03 (\pm 9.57E-05)	1.81E-03 (\pm 9.92E-05)
$\zeta_{1_invalid}$	1.74E-03 (\pm 9.76E-05)	1.75E-03 (\pm 9.55E-05)	1.75E-03 (\pm 9.56E-05)	1.77E-03 (\pm 9.44E-05)	1.78E-03 (\pm 9.39E-05)	1.74E-03 (\pm 9.68E-05)
ζ_2	4.83E-04 (\pm 1.59E-05)	5.26E-04 (\pm 2.35E-05)	4.7E-04 (\pm 2.51E-05)	4.33E-04 (\pm 1.6E-05)	4.36E-04 (\pm 2.57E-05)	4.7E-04 (\pm 2.51E-05)

tion ζ_1 (i.e., the overall levels of response speeds in valid and invalid trials) and the slope of the affine function ζ_2 (i.e., the strength of the increase in response speed with increased precision-dependent attention).

On the basis of these response model parameters and the subject- and session-specific values of α and $P^{(0)}$, we calculated observed and predicted responses in the different experimental blocks and stimulation sessions. Figure 3 shows that the observed

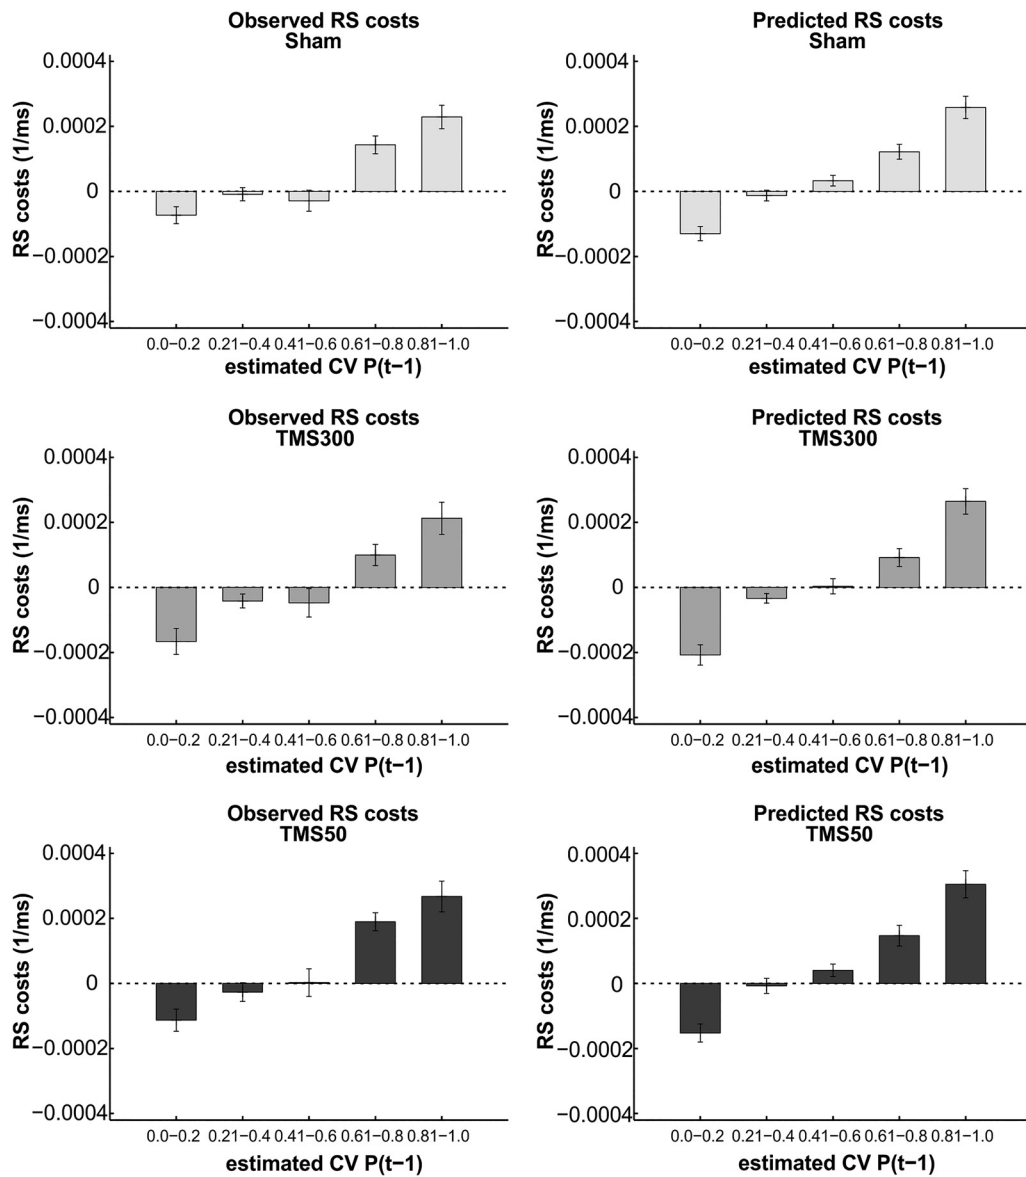


Figure 3. Observed and predicted pattern of RS costs from the learning model. RS costs were calculated by subtracting RSs of invalid trials from valid trials and are shown in relation to the participants' estimated probability that the cue will be valid in an upcoming trial $P^{(t-1)}$, binned in intervals of 0.2. As expected, RS costs increased with increased subjective %CV, and the pattern of observed RS costs matched the RS costs predicted by the model on the basis of individual values for α , $P^{(0)}$, ζ_{1_valid} , $\zeta_{1_invalid}$ and ζ_2 in the different blocks and sessions. Error bars indicate SEM.

Table 2. Results of the BMS selection^a

Model	Sham		TMS300		TMS50	
	PXP	Dirichlet parameter	PXP	Dirichlet parameter	PXP	Dirichlet parameter
All blocks						
Learning model	0.98	17.4	0.99	17.6	1	18
No learning model	0.02	2.6	0.01	2.4	<0.001	2
Only false blocks						
Learning model	0.91	13.2	0.86	12.5	0.99	15.2
Switch-half model	0.02	1.3	0.02	1.2	<0.001	1.2
Switch-quarter model	0.07	6.5	0.12	7.2	0.01	4.6

^aPXP, Protected exceedance probability.

RS pattern closely matched the RS pattern predicted by the learning model in all three stimulation conditions, corroborating the validity of our modeling approach.

Moreover, we compared our model with a null model with no learning ($\alpha = 0$). Random-effects BMS showed that the Rescorla-

Wagner learning model had higher model evidence than the no learning model in all three stimulation conditions (Table 2).

In an additional analysis, we tested for the possibility that our probability manipulation of using inverse probabilities in false blocks was too obvious and that the subjects would merely switch to the inverse probability after having observed a couple of trials. Therefore, we formally compared the log-model evidence of the learning model with two alternative models, which predicted that participants in false blocks would directly switch after the first quarter or after half of the trials from the expected %CV (i.e., the prior information provided at the beginning of the blocks) to the inverse %CV. Random-effects BMS showed that the learning model had higher model evidence in all three stimulation conditions (Table 2).

We also tested, in a further control analysis, whether the difference in α in false blocks between different stimulation conditions depended on the high or low violation of the expectation concerning %CV. In blocks with 90% or 10% CV, the violation of

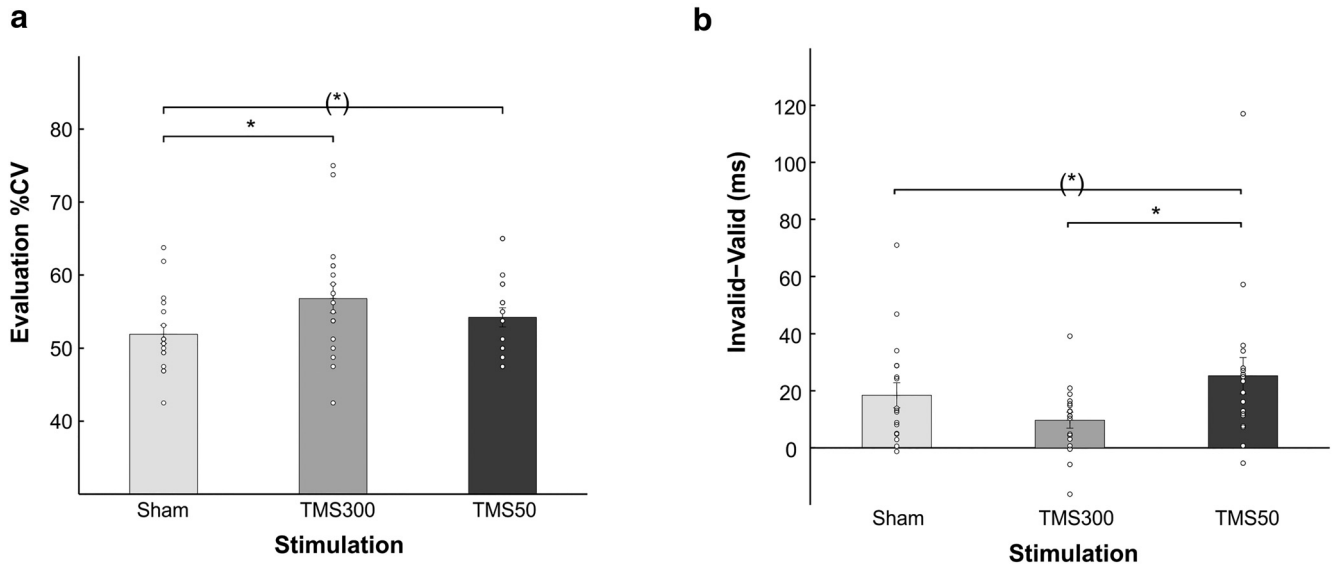


Figure 4. Mean explicit evaluation of %CV and mean (probability-independent) validity effects. **a**, Explicit evaluation of the %CV for the three stimulation conditions. Error bars indicate SEM. **b**, Validity effects (i.e., the tendency to rely on the cue) in the three stimulation conditions. Error bars indicate SEM. * $p < 0.05$. (*) $p < 0.05$, not passing the Bonferroni-corrected threshold for multiple comparisons. White circles indicate single-subjects data.

the expectations was stronger than in 70% and 30% blocks. We therefore ran a within-subject ANOVA with the factors stimulation (sham, TMS300, TMS50), prior (true, false), and block (high difference, 90%, 10%; low difference, 70%, 30%) on the α parameter. In addition to the previously shown effects of prior ($F_{(1,17)} = 39.6, p = 0.000008$) and stimulation \times prior interaction ($F_{(1.88,32)} = 4.15, p = 0.027$), we found a main effect of block ($F_{(1,17)} = 13.5, p = 0.002$), with higher learning rates for high difference blocks. However, no significant interaction between block and stimulation or prior was found, suggesting that a higher or lower level of violation of expectation did not impact on our stimulation results.

TMS-induced explicit overestimation of cue validity

The within-subject ANOVA with the factors stimulation (sham, TMS300, TMS50) and prior (true, false) on the explicit evaluation of the %CV revealed a main effect of prior ($F_{(1,17)} = 17.3, p = 0.0007$), with higher estimates of %CV in blocks with false prior ($57 \pm 1.6\%$ CV) than in blocks with true prior ($52 \pm 1.2\%$ CV), and a main effect of stimulation ($F_{(1.7,29)} = 5.9, p = 0.009$; Fig. 4a). Pairwise comparisons (three comparisons, Bonferroni-corrected threshold $p < 0.017$) showed a significant difference between TMS300 and sham ($t_{(17)} = 3.1, p = 0.007, B_{10} = 7.5$), with higher explicit evaluation of %CV with TMS at 300 ms. In addition, we found a trend toward significance between TMS50 and sham ($t_{(17)} = 2.1, p = 0.048, B_{10} = 1.5, B_{01} = 0.7$), with a tendency toward higher explicit evaluation of %CV with TMS at 50 ms.

The within-subject ANOVA on the TMS-sham differences with the factors stimulation (TMS300-SHAM300, TMS50-SHAM50) and prior (true, false) revealed a main effect of prior ($F_{(1,17)} = 6.9, p = 0.018$), with higher explicit evaluation of %CV for blocks with false prior, and a trend toward significance for the main effect of stimulation ($F_{(1,17)} = 3.3, p = 0.089$), with a tendency toward increased estimates of %CV with TMS300-SHAM300 compared with TMS50-SHAM50. We also tested for the deviance of the TMS effects from zero for the main effect of stimulation with one-sample t tests (two comparisons, Bonferroni-corrected threshold $p < 0.025$), which revealed that the only condition that

was different from baseline was in the TMS300-SHAM300 ($t_{(17)} = 2.9, p = 0.009, B_{10} = 5.6$), with an increase in the explicit estimate of %CV compared with baseline.

TMS-induced increase in average validity effects

We also analyzed mean condition-specific RSs to test for a TMS effect on attentional reorienting (i.e., the validity effect: RS difference between valid and invalid trials).

The within-subject ANOVA on mean RSs with the factors stimulation (sham, TMS300, TMS50), prior (true, false), and validity (valid, invalid) revealed a main effect of validity ($F_{(1,17)} = 32.8, p = 0.00025$), with lower RSs for invalid trials, and a significant stimulation \times validity interaction ($F_{(1.81,30.77)} = 5.25, p = 0.0035$). No main effect of stimulation was found, suggesting that there were no unspecific effects of TMS on RSs.

We ran paired t tests to interpret the effect found in the interaction (nine comparisons, Bonferroni-corrected threshold $p < 0.0056$). We found a significant difference between valid and invalid trials in all three stimulation conditions ($p < 0.002$). No other comparison was significant ($p > 0.3$). Therefore, the interaction effect must lie in the size of RS difference between valid and invalid trials (i.e., the validity effect) in the three stimulation conditions. Accordingly, we calculated validity effects, subtracting RSs of valid and invalid trials. Validity effects in the three stimulation conditions were then compared with paired t tests (three comparisons, Bonferroni-corrected threshold $p < 0.017$; Fig. 4b). A significant difference was found between TMS300 and TMS50 ($t_{(17)} = 3.55, p = 0.002, B_{10} = 17$), with higher validity effects for the TMS50 condition, suggesting an effect of TMS interference on validity effects when TMS was applied at 50 ms. In addition, we found a trend toward significance between validity effects in the TMS50 and sham ($t_{(17)} = -2.56, p = 0.02, B_{10} = 2.9, B_{01} = 0.3$). In the present task, as half of the trials were presented with a counter-predictive cue ($\sim 30\%$ and $\sim 10\%$ blocks), validity effects theoretically should not be detected. Therefore, positive validity effects should be interpreted as resulting from a general tendency to rely on the cue, even when it is not predictive.

The analysis on the differences between TMS and sham with a within-subject ANOVA with the factors stimulation (TMS300-SHAM300, TMS50-SHAM50), prior (true, false), and validity (valid, invalid) revealed only a significant prior \times validity interaction ($F_{(1,17)} = 5, p = 0.039$) but no effect of, or interaction with, stimulation.

TMS effects on accuracy

Overall, the average accuracy amounted to 94% (± 0.8 SEM). The within-subject ANOVA on accuracy scores with the factors stimulation (sham, TMS300, TMS50), prior (true, false), and validity (valid, invalid) revealed a main effect of validity ($F_{(1,17)} = 32, p = 0.00003$), with higher accuracy for valid ($95 \pm 0.6\%$) than for invalid ($92 \pm 0.9\%$) trials. However, no significant effect of stimulation or interaction with this factor was found, suggesting that the TMS did not affect accuracy (neither after 300 ms after 50 ms).

Discussion

In the present study, we applied online TMS to probe the causal role of rTPJ in belief updating and to characterize the chronometry of this process. As hypothesized, our data revealed that disrupting rTPJ activity by applying TMS 300 ms after target appearance selectively decreased participants' updating of false prior beliefs concerning %CV without affecting reorienting of attention. Additional analyses showed that observed response speed indeed followed the predictions of the employed learning model, and that models with no learning or models according to which the subject merely switched to the inverse probability level in false prior blocks after a couple of trials had lower model evidence than the Rescorla–Wagner model. The effect of TMS interference over rTPJ also had an impact on the participants' explicit beliefs concerning cue validity, although with less clear time specificity. Moreover, TMS interference did not cause any unspecific effects on RSs or accuracy.

Our results are consistent with previous neuroimaging evidence (Vossel et al., 2015), showing that rTPJ activity is modulated by violations of expectations concerning upcoming external events. However, given the correlative nature of fMRI, the present work, for the first time, provides evidence for a causal role of rTPJ in probabilistic belief updating. Hence, the present findings are consistent with, but extend, prior work. Together, neuroimaging and TMS evidence offer a unifying interpretation of previously observed rTPJ activations across multiple cognitive domains, from theory of mind to body awareness (Blanke et al., 2005; Decety and Lamm, 2007; Mitchell, 2008; Tsakiris et al., 2008). Indeed, all the tasks used in these studies share the need for updating prior information with new upcoming information provided by other individuals, external sensory stimuli, or internal body signals. Therefore, our result, even though embedded in a spatial attention task, might not be restricted to the attentional domain only.

These novel insights about the function and chronometry of rTPJ activity can also be linked to the ERP literature according to which contextual updating after unexpected events is associated with the P300 component (Polich, 2007; Kolossa et al., 2015). It has been speculated that the neural origin of the P300 may be in frontal and temporoparietal brain regions (Polich, 2007; Kolossa et al., 2015). Our present data provide direct evidence for the involvement of rTPJ in belief updating after 300 ms from target onset, thereby supporting the hypothesis of a connection between P300 component and contextual updating.

When looking more generally on average RSs in all valid and invalid trials, we also found an effect when interfering with rTPJ

activity, causing an increase of validity effects. This effect was more pronounced with TMS 50 ms from target onset (compared with the late TMS stimulation at 300 ms), but the analysis with TMS-sham differences for the two timings did not reveal any specific differences between TMS stimulation at 50 and at 300 ms, making it difficult to draw any definitive conclusion from this observation.

Likewise, the effects of TMS on the subject's explicit evaluation of cue validity were not clear-cut. Here, more research is warranted to systematically explore the relationship between parameters based upon RS and explicit ratings of probabilities.

Our present results show that the function of rTPJ is not limited to reorienting and proposes an additional novel role in updating prior beliefs in a later phase of stimulus processing. This interpretation also helps to understand why patients with right-hemispheric parietal damage are not just impaired in reorienting attention toward invalidly cued targets (e.g., Friedrich et al., 1998) but also do not show speeded responses after repeated presentations of a stimulus (Shaqiri and Anderson, 2012) and have difficulties to improve their task performance through feedback (Seki et al., 1999). In addition, the suggested new role of rTPJ based upon our current TMS and previous neuroimaging data (Vossel et al., 2015) may explain the broad involvement of this region in diverse cognitive domains, which require reconciling predicted and observed input and adjusting internal models accordingly.

As far as the specificity of rTPJ for contextual updating is concerned, we cannot exclude that additional regions are involved in the updating of prior expectations. Nevertheless, the main aim of the study, inspired by the consistent involvement of rTPJ in modulating beliefs in spatial cueing tasks found in fMRI studies (Vossel et al., 2015; Dombert et al., 2016), concerned the causal involvement of this region in contextual updating together with a characterization of the timing of the processing. The results strongly support the hypothesis of a causal role of rTPJ in contextual updating.

References

- Azañón E, Longo MR, Soto-Faraco S, Haggard P (2010) The posterior parietal cortex remaps touch into external space. *Curr Biol* 20:1304–1309. [CrossRef Medline](#)
- Blanke O, Mohr C, Michel CM, Pascual-Leone A, Brugger P, Seeck M, Landis T, Thut G (2005) Linking out-of-body experience and self-processing to mental own-body imagery at the temporoparietal junction. *J Neurosci* 25:550–557. [CrossRef Medline](#)
- Brodersen KH, Penny WD, Harrison LM, Daunizeau J, Ruff CC, Duzel E, Friston KJ, Stephan KE (2008) Integrated Bayesian models of learning and decision making for saccadic eye movements. *Neural Netw* 21:1247–1260. [CrossRef Medline](#)
- Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62. [CrossRef Medline](#)
- Chambers CD, Payne JM, Stokes MG, Mattingley JB (2004) Fast and slow parietal pathways mediate spatial attention. *Nat Neurosci* 7:217–218. [CrossRef Medline](#)
- Chica AB, Bartolomeo P, Valero-Cabré A (2011) Dorsal and ventral parietal contributions to spatial orienting in the human brain. *J Neurosci* 31:8143–8149. [CrossRef Medline](#)
- Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3:201–215. [CrossRef Medline](#)
- Corbetta M, Patel G, Shulman GL (2008) The reorienting system of the human brain: from environment to theory of mind. *Neuron* 58:306–324. [CrossRef Medline](#)
- Decety J, Lamm C (2007) The role of the right temporoparietal junction in social interaction: how low-level computational processes contribute to meta-cognition. *Neuroscientist* 13:580–593. [CrossRef Medline](#)
- Dombert PL, Kuhns A, Mengotti P, Fink GR, Vossel S (2016) Functional mechanisms of probabilistic inference in feature- and space-based attentional systems. *Neuroimage* 142:553–564. [CrossRef Medline](#)

- Eimer M (1997) Uninformative symbolic cues may bias visual-spatial attention: behavioral and electrophysiological evidence. *Biol Psychol* 46:67–71. [CrossRef Medline](#)
- Friedrich FJ, Egly R, Rafal RD, Beck D (1998) Spatial attention deficits in humans: a comparison of superior parietal and temporal-parietal junction lesions. *Neuropsychology* 12:193–207. [CrossRef Medline](#)
- Friston K, Mattout J, Trujillo-Barreto N, Ashburner J, Penny W (2007) Variational free energy and the Laplace approximation. *Neuroimage* 34:220–234. [CrossRef Medline](#)
- Geng JJ, Vossel S (2013) Re-evaluating the role of TPJ in attentional control: contextual updating? *Neurosci Biobehav Rev* 37:2608–2620. [CrossRef Medline](#)
- George D, Mallery P (2010) *SPSS for Windows step by step: a simple guide and reference*. Boston: Pearson.
- Graziano MS, Kastner S (2011) Human consciousness and its relationship to social neuroscience: a novel hypothesis. *Cogn Neurosci* 2:98–113. [CrossRef Medline](#)
- Kolossa A, Kopp B, Fingscheidt T (2015) A computational analysis of the neural bases of Bayesian inference. *Neuroimage* 106:222–237. [CrossRef Medline](#)
- Mitchell JP (2008) Activity in right temporo-parietal junction is not selective for theory-of-mind. *Cereb Cortex* 18:262–271. [CrossRef Medline](#)
- Oldfield RC (1971) The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia* 9:97–113. [CrossRef Medline](#)
- Polich J (2007) Updating P300: an integrative theory of P3a and P3b. *Clin Neurophysiol* 118:2128–2148. [CrossRef Medline](#)
- Posner MI (1980) Orienting of attention. *Q J Exp Psychol* 32:3–25. [CrossRef Medline](#)
- Rescorla RA, Wagner ARA (1972) A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement. In: *Classical conditioning, Vol II: Current research and theory* (Black AH, Prokasy WF, eds), pp 64–99. New York: Appleton-Century-Crofts.
- Rigoux L, Stephan KE, Friston KJ, Daunizeau J (2014) Bayesian model selection for group studies—revisited. *Neuroimage* 84:971–985. [CrossRef Medline](#)
- Rushworth MF, Behrens TE (2008) Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat Neurosci* 11:389–397. [CrossRef Medline](#)
- Seki K, Ishiai S, Koyama Y, Sato S (1999) Unassociated responses to two related task demands: a negative factor for improvement of unilateral spatial neglect. *Neuropsychologia* 37:75–82. [CrossRef Medline](#)
- Shaqiri A, Anderson B (2012) Spatial probability cuing and right hemisphere damage. *Brain Cogn* 80:352–360. [CrossRef Medline](#)
- Stephan KE, Penny WD, Daunizeau J, Moran RJ, Friston KJ (2009) Bayesian model selection for group studies. *Neuroimage* 46:1004–1017. [CrossRef Medline](#)
- Tsakiris M, Costantini M, Haggard P (2008) The role of the right temporo-parietal junction in maintaining a coherent sense of one's body. *Neuropsychologia* 46:3014–3018. [CrossRef Medline](#)
- Vossel S, Mathys C, Daunizeau J, Bauer M, Driver J, Friston KJ, Stephan KE (2014) Spatial attention, precision, and bayesian inference: a study of saccadic response speed. *Cereb Cortex* 24:1436–1450. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)