Humans strategically shift decision bias by flexibly adjusting sensory evidence accumulation in visual cortex

Niels A. Kloosterman[*1,2,3,5], Jan Willem de Gee[2,4], Markus Werkle-Bergner[5], Ulman Lindenberger[1,5], Douglas D. Garrett[1,5+], Johannes Jacobus Fahrenfort[2,6+]

[1] Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

[2] Department of Psychology, University of Amsterdam, The Netherlands;

[3] Center for Brain and Cognition, Institute for Interdisciplinary Studies, University of Amsterdam, The Netherlands;

[4] Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, Germany;

[5] Center for Lifespan Psychology, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

[6] Department of Experimental and Applied Psychology, Vrije Universiteit, van der Boechorststraat 1, 1081 BT Amsterdam, The Netherlands

[+]Shared senior author

*Correspondence: kloosterman@mpib-berlin.mpg.de

Lead Contact Information

Niels A. Kloosterman, Ph.D.

Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Lentzeallee 94, 14195, Berlin, Germany

Phone: +49 30 82406 424

E-mail: kloosterman@mpib-berlin.mpg.de

Abstract

Decision bias is traditionally conceptualized as a flexible internal reference against which sensory evidence is compared. Here, in contrast, we show that experimental manipulation of decision bias adjusts the rate of evidence accumulation in visual cortex towards one of the choice alternatives. Participants performed a visual detection task during EEG recordings. We experimentally manipulated participants' response criterion using different stimulus-response reward contingencies, inducing liberal and conservative decision biases in different conditions. Drift diffusion modeling of choice behavior revealed that an experimentally induced liberal decision bias specifically biased the rate of sensory evidence accumulation towards 'yes' choices. In visual cortex, the liberal bias manipulation suppressed prestimulus 8—12 Hz (alpha) power, which in turn boosted cortical stimulus-related activity in the 59—100 Hz (gamma) range. Together, these findings show that observers can intentionally control cortical excitability to strategically bias evidence accumulation towards the decision bound that maximizes reward within a given ecological context.

Introduction

Perceptual decisions arise not only from the evaluation of sensory evidence, but are often biased towards one or the other response alternative by environmental factors, for example as a result of task instructions and/or stimulus-response reward contingencies (White & Poldrack, 2014). The ability to willfully control decision bias enables the behavioral flexibility required to survive in an ever-changing and uncertain environment. But despite its central and important role in decision making, the neural mechanisms underlying decision bias are not fully understood.

The traditional account of decision bias comes from signal detection theory (SDT) (Green & Swets, 1966). In SDT, decision bias is quantified by estimating the relative position of a decision point or 'criterion' in between sensory evidence distributions for noise and signal (see Figure 1A). In this framework, a more liberal decision bias arises by moving the criterion closer towards the noise distribution (see green arrow in Figure 1A). Although SDT has been very successful at quantifying decision bias, it has not done much to elucidate the mechanism behind it. One reason for this lack of insight may be that SDT does not have a temporal component to track how decisions are reached over time (Fetsch, Kiani, & Shadlen, 2014).

As an alternative to SDT, the drift diffusion model (DDM) conceptualizes perceptual decision making as the accumulation of noisy sensory evidence over time into an internal decision variable (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Gold & Shadlen, 2007; Ratcliff & McKoon, 2008). A decision in this model is made when the decision variable crosses one of two decision bounds corresponding to the choice alternatives. Within this framework, a strategic decision bias imposed by the environment can be modelled in two different ways: either by moving the starting point of evidence accumulation closer to one of the boundaries (see green arrow in Figure 1B), or by biasing the rate of the evidence accumulation process itself towards one of the boundaries (see green arrow in Figure 1C). In both the SDT and DDM frameworks, decision bias shifts have little effect on the sensitivity of the observer when distinguishing signal from noise; they predominantly affect the relative response ratios (and in the case of DDM the speed with which one or the other decision bound is reached). There has been some evidence to suggest that decision bias induced through shifting the response criterion is best characterized by a drift bias in the DDM (Urai, de Gee, & Donner, 2018; White & Poldrack, 2014). However, the drift bias parameter has as yet not been related to a well-described cortical mechanism.

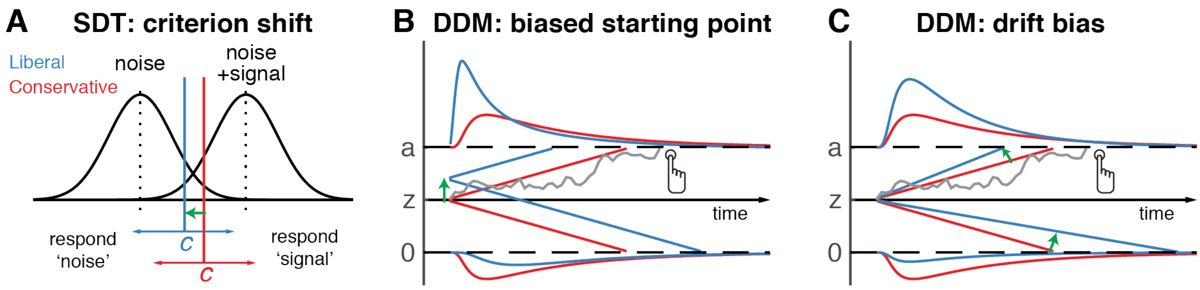

Figure 1 | Theoretical accounts of decision bias. A. The SDT account of decision bias. In this framework, signal and noise+signal distributions are plotted as a function of the strength of internal sensory evidence. Here, the decision point (or criterion) that determines whether to indicate signal presence or absence is plotted as a vertical criterion line *c*, reflecting the degree of decision bias. *c* can be shifted left- or rightwards to respectively model a more liberal or conservative bias (green arrow indicates a shift to liberal). In drift diffusion models (DDMs, panels **B.** and **C.**), decisions are modelled in terms of a dynamic process of sensory evidence accumulation. When sensory input is presented, evidence starts to accumulate (drift) over time departing from starting point *Z*. The rate at which evidence accumulates is called the drift rate, and a response is given when it either crosses decision boundary *a* (signal presence) or decision boundary *0* (no signal). DDMs are fitted to distributions of reaction times obtained over multiple trials. In panels B. and C., reaction time distributions for signal-present responses are plotted at the top and reaction time distributions for no-signal responses are plotted in mirror image at the bottom. DDMs can model bias in two different ways. In panel B., bias is modelled in terms of the DDM starting point *Z*, which is moved closer or further away from the decision bounds *a* and *0*. In panel C., decision bias is modelled in terms of drift bias, where the rate of evidence accumulation for signal and noise move upwards or downwards in tandem (green arrows indicate a shift to liberal). The predicted reaction time distributions under the models in B. and C. are plotted separately for a liberal and conservative bias above and below the graphs. Panels B. and C. are modified and reproduced with permission from Urai, de Gee, & Donner (2018).

Likewise, there have been a number of reports about a correlational relationship between cortical activity and decision bias. For example, spontaneous trial-to-trial variations in prestimulus oscillatory activity in the 8—12 Hz (alpha) band have been shown to correlate with decision bias and confidence (Iemi, Chaumon, Crouzet, & Busch, 2017; Limbach & Corballis, 2016; Samaha, Iemi, & Postle, 2017). Relatedly, alpha oscillations have been proposed to be involved in the gating of task-relevant sensory information (Jensen & Mazaheri, 2010), possibly encoded in high-frequency (gamma) oscillations in visual cortex (Ni et al., 2016; Popov, Kastner, & Jensen, 2017). Although these reports suggest a link between alpha suppression and decision bias, they do not uncover whether pre-stimulus alpha plays an instrumental role in decision bias and how exactly this might be achieved. For example, it is unknown whether an experimentally induced shift in decision bias is implemented in the brain by willfully adjusting pre-stimulus alpha in sensory areas.

Here, we explicitly investigate these potential mechanisms by employing a task paradigm in which shifts in decision bias were experimentally induced within participants through instruction and asymmetries in stimulus-response reward contingencies during a visual detection task. By applying drift diffusion modeling to the participants' choice behavior, we show that strategically adjusting decision bias specifically affects the rate of sensory evidence accumulation towards one of the two decision bounds. Further, we demonstrate that this drift bias is achieved by flexibly up- and down-regulating prestimulus alpha to control the response gain of stimulus-related gamma activity in visual cortex. Critically, we also show that gamma activity accurately predicts the strength of the evidence accumulation bias within subjects, providing a direct link between the proposed mechanism and decision bias. Together, these findings identify the neural mechanism by which intentional control of cortical excitability is applied to strategically bias perceptual decisions in order to maximize reward in a given context.

Results

Liberal decision bias manipulation shifts sensory evidence accumulation

In three EEG recording sessions, human participants (N = 16) viewed a continuous stream of horizontal, vertical and diagonal line textures alternating at a rate of 25 textures/second. The participants' task was to detect an orientation-defined square presented in the center of the screen and report it via a button press (Figure 2A). Trials consisted of a fixed-order sequence of textures (total sequence duration 1 second) embedded in the continuous stream. A square appeared in the fifth texture of a trial in 75% of the presentations (target trials), while in 25% a homogenous diagonal texture appeared in the fifth position (nontarget trials). Although the onset of trials within the continuous stream of textures was not explicitly cued, the similar distribution of reaction times in target and nontarget trials suggests that participants employed the temporal structure of the task even when no target appeared (Figure S1A). Consistent significant EEG power modulations after trial onset even for non-target trials further confirm that subjects registered trial onsets even without an explicit cue, plausibly using the onset of a fixed order texture sequence as an implicit cue (Figure S1B).

In alternating nine-minute blocks of trials, we actively biased participants' perceptual decisions by instructing them either to report as many targets as possible ("Detect as many targets as possible!"; liberal condition), or to only report high-certainty targets ("Press only if you are really certain!"; conservative condition). Participants were free to respond at any time during a block whenever they detected a target. We provided auditory feedback following missed targets (misses) in the liberal condition and falsely detected targets (false alarms) in the conservative condition and applied monetary penalties for these errors (Figure 2A; see Methods for details).

Participants reliably adopted the intended criterion shift (see Figure 2B showing that both the hit rate and the false alarm rate went down in tandem as a consequence of a more conservative criterion). The difference between hit rate and false alarm rate was not significant between conservative and liberal ($p = 0.81$, right bars in Figure 2B). However, detection performance (sensitivity) computed using standard SDT *d'* (reflecting the distance between the noise and signal distributions in Figure 1A) was slightly higher during conservative (liberal: $d' = 2.0$ (s.d. 0.90), versus conservative: $d' = 2.31$ (s.d. 0.82), $p = 0.0002$, two-sided permutation test, 10,000 permutations, see Figure 2C, left bars)(Green & Swets, 1966). We also computed the standard SDT criterion measure *c* reflecting bias in the decision process (see the blue and red vertical lines in Figure 1A), which uncovered a strong experimentally induced criterion shift (liberal: $c = -0.13$ (s.d. 0.4), versus conservative: $c = 0.73$ (s.d. 0.36), $p = 0.0001$, permutation test, see Figure 2C, right bars).

Because the SDT framework is static, we decided to further investigate how bias affected various components of the dynamic decision process by fitting different drift diffusion models (DDMs) to the behavioral data (Figure 1B, C) (Ratcliff & McKoon, 2008). DDMs postulate that perceptual decisions are reached by accumulating noisy sensory evidence towards one of two decision boundaries representing the choice alternatives. Crossing one of these boundaries can either trigger an explicit behavioural report to indicate the decision (for 'yes' responses in our experiment), or remain implicit (i.e. without active response, for 'no' decisions in our experiment) (Ratcliff, Huang-Pollock, & McKoon, 2016). We tested two different DDMs that can potentially account for decision bias: one in which the starting point of evidence accumulation moves closer to one of the decision boundaries ('starting point model', Figure 1B) (Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012), and one in which the evidence accumulation process (called the drift) itself is biased towards one of the boundaries (de Gee et al., 2017) ('drift bias model', see Figure 1C, referred to as drift criterion by Rattclif and McKoon (2008)). In the two respective models, we freed either the drift bias parameter (db, see Figure 2D) for the two conditions while keeping starting point (z) fixed across conditions (for the drift bias model), or vice versa (for the starting point model). The drift bias parameter is determined by estimating the contribution of an evidence-independent constant added to the drift (Figure 2B). These alternative models make different predictions about the shape of the RT distributions in combination with the response ratios: a shift in starting point produces large changes in both the leading edge and tail of the distribution, whereas a shift in drift bias produces large changes only in the tail (Ratcliff & McKoon, 2008; Urai et al., 2018), also see the RT distributions above and below the evidence accumulation graphs in Figure 1B and 1C.

We fitted both the starting point and drift bias models to each participant's RT distribution for 'yes' choices and the total number of implicit 'no' choices. In both models, all of the non-bias related parameters (drift rate v, boundary separation a and non-decision time u+w, see Figure 2D) were also allowed to vary by condition. We compared goodness of fit of the models to assess which model best explained the data. We found that the starting point model provided a worse fit to the data, as indicated by higher Bayesian Information Criterion (BIC) estimates than for the drift bias model (Figure 2E, see Methods for details). Specifically, for 14 out of the 16 participants the drift bias model provided better fits than the starting point model, for ten of which delta BIC was greater than six, indicating strong evidence in favor of the drift bias model. Finally, we compared these models to a model in which both drift bias and starting point were fixed across the conditions, while still allowing the non-bias-related parameters to vary per condition. This model provided the lowest goodness of fit (delta BIC greater than six for both models for all participants). See Figure S3 for model fits of the drift bias model for each participant.

Given the superior performance of the drift bias model, we further characterized decision making under the bias manipulation using parameter estimates from this model. Drift rate, reflecting the participants' ability to discriminate targets and non-targets, was somewhat higher in the conservative compared to the liberal condition (liberal: $v = 2.39$ (s.d. 1.07), versus conservative: $v = 3.06$ (s.d. 1.16), $p = 0.0001$, permutation test, Figure 2F, left bars). An almost perfect correlation across participants between DDM drift rate and SDT *d'* provided strong evidence that the drift rate parameter captures perceptual sensitivity (liberal, $r = 0.97$; conservative, $r = 0.95$, all p-values < 0.005, see Figure S2A).

Regarding the DDM bias parameters, the condition-fixed starting point parameter in the drift bias model was smaller than half the boundary separation (i.e. closer to the 'no' boundary: $z = 0.24$, $p < 0.0001$, tested against 0.5), indicating an overall conservative starting point across conditions (Figure S2D). Strikingly, however, whereas drift bias was on average not different from zero in the conservative condition (db = –0.04, p = 0.90), drift bias was strongly positive in the liberal condition (db = 2.08, p = 0.0001; liberal vs conservative: p = 0.0005; Figure 2F, right bars). The overall conservative starting point combined with a condition-specific neutral drift bias explained the conservative decision bias (as quantified by SDT criterion) in the conservative condition (Figure 2C). Likewise, in the liberal condition the overall conservative starting point combined with a condition-specific positive drift bias (pushing the drift towards the 'yes' boundary) explained the neutral bias observed with SDT (criterion around zero for liberal, see Figure 2C).

Converging with these modelling results, drift bias was strongly anti-correlated across participants with both SDT criterion (liberal, r = –0.83; conservative, r = –0.79, see Figure S2B) and reaction times (liberal, r = –0.66; conservative, r = –0.76, see Figure S2C). The strong correlations between DDM drift rate and SDT *d'* on the one hand, and DDM drift bias and SDT criterion on the other, provide converging evidence that the SDT and DDM frameworks captured similar underlying mechanisms, while the DDM additionally captured the dynamic nature of perceptual decision making by linking the decision bias manipulation to the evidence accumulation process itself.

Finally, the bias manipulation also affected two other parameters in the drift bias model that were not directly related to sensory evidence accumulation: boundary separation was slightly but reliably higher during liberal compared to conservative ($p < 0.0001$), and non-decision time (comprising time needed for sensory encoding and motor response execution) was shorter during liberal ($p < 0.0001$)(supplementary Figure S2D). In conclusion, a drift diffusion model of choice behavior implementing a bias in sensory evidence accumulation best explained how participants adjusted to the manipulations of decision bias. In the next sections, we used spectral analysis of the concurrent EEG recordings to identify a plausible neural mechanism that implements biased sensory evidence accumulation.

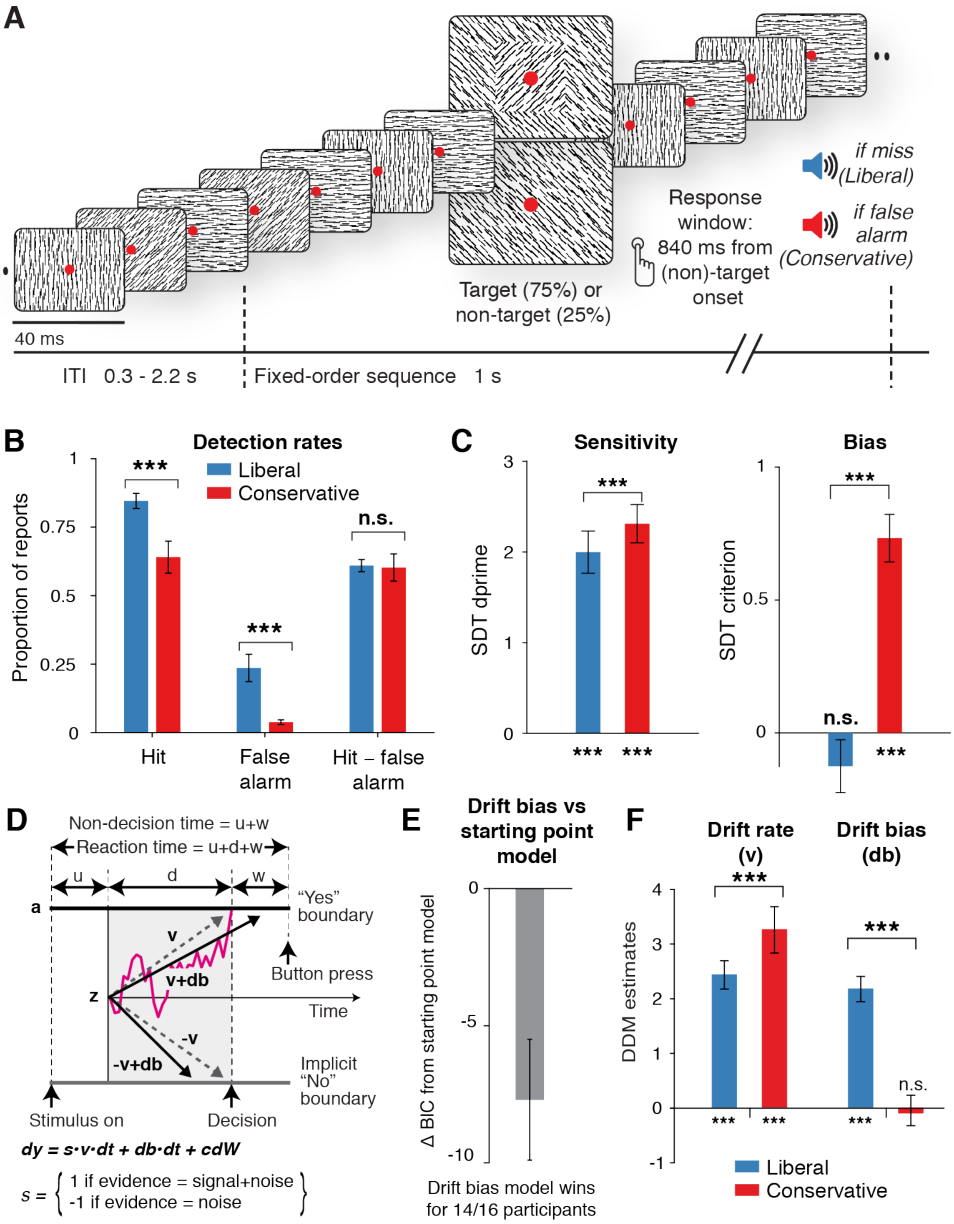

Figure 2 | Biased sensory evidence accumulation underlies experimentally induced liberal decision bias A. Schematic of the visual stimulus and task design. Participants viewed a continuous stream of full-screen diagonally, horizontally and vertically oriented textures at a presentation rate of 40 ms (25 Hz). After random inter-trial intervals (range 0.3—2.2 s), a fixed-order sequence (duration 1 s) was presented, embedded in the stream. The fifth texture in each sequence either consisted of a single diagonal orientation (nontarget), or contained an orthogonal orientation-defined square (target of either 45° or 135° orientation). Participants decided whether they had just seen a target, reporting detected targets by button press within 840 ms after target onset. Liberal and conservative conditions were administered in alternating nine-minute blocks by penalizing either misses or false alarms, respectively, using aversive tones and monetary deductions. Depicted square and fixation dot sizes are not to scale. **B.** Average detection rates (hits and false alarms) during both conditions. **C.** SDT parameters for sensitivity and criterion. **D**. Schematic and simplified equation of drift diffusion model accounting for reaction time (RT) distributions for explicit 'yes' and implicit 'no' decisions. Decision bias in this model can be implemented by either shifting the starting point of the evidence accumulation (*Z*), or by adding an evidence-independent constant ('drift bias', db) to the drift rate. See text and Figure 1 for details. Notation: dy, change in decision variable y per unit time dt; v·dt, mean drift (multiplied with 1 for signal + noise (target) trials, and -1 for noise-only (non-target) trials); db·dt, drift bias; and cdW, Gaussian white noise (mean = 0, variance = c2·dt). **E.** The difference in Bayesian Information Criterion (BIC) goodness of fit estimates for the drift bias and the starting point models, A lower delta BIC value indicates a better fit**,** showing superiority of the drift bias model to account for the observed results. **F.** Estimated model parameters for drift rate and drift bias in the drift bias model. Error bars, SEM across 16 participants. ***$p < 0.001$; n.s., not significant.

Task-relevant textures induce stimulus-related responses in visual cortex

Sensory evidence accumulation in the visual detection task presumably relies on stimulus-related signals processed in visual cortex. Such stimulus-related signals are typically reflected in cortical population activity exhibiting a rhythmic temporal structure (Buzsáki & Draguhn, 2004). Specifically, bottom-up processing of visual information has previously been linked to increased high-frequency (> 40 Hz, i.e. gamma) electrophysiological activity over visual cortex (Bastos et al., 2015; Michalareas et al., 2016; Popov et al., 2017; van Kerkoerle et al., 2014). Figure 3A shows time-frequency representations of EEG power modulations over posterior cortex for the low and high frequency bands, normalized with respect to the prestimulus baseline period.

We observed a total of four distinct stimulus-related power modulations after trial onset: two in the high-frequency range (> 36 Hz, Figure 3A, top panel) and two in the low frequency range (< 36 Hz, Figure 3A, bottom panel). First, a spatially focal modulation in a narrow frequency range around 25 Hz reflecting the steady state visual evoked potential (SSVEP) arising from entrainment by the visual stimulation frequency of our experimental paradigm (Figure 3B, lower panel). A second modulation from 42—58 Hz (Figure 3B, top panel) comprised the first harmonic of the SSVEP, as can be seen from their similar topographic distributions (Figure 3B, compare top and lower panel).

Third, we observed a 59—100 Hz gamma power modulation (Figure 3C, top panel), after carefully controlling for high-frequency EEG artifacts due to small fixational eye movements (microsaccades) by removing microsaccade-related activity from the data (Hassler, Trujillo-Barreto, & Gruber, 2011; Hipp & Siegel, 2013; Yuval-Greenberg, Tomer, Keren, Nelken, & Deouell, 2008), and by suppressing non-neural EEG activity through scalp current density transformation (Melloni, Schwiedrzik, Wibral, Rodriguez, & Singer, 2009; Perrin, Pernier, Bertrand, & Echallier, 1989) (see Methods for details). Importantly, the topography of the observed gamma modulation was confined to posterior electrodes (electrodes highlighted in Figures 3B and 3C, top panels), in line with the role of gamma in stimulus-related processing in visual cortex (Ni et al., 2016). Finally, we observed suppression of low-frequency beta (11—22 Hz) activity in posterior cortex, which typically occurs in parallel with enhanced stimulus-related gamma activity (Donner & Siegel, 2011; Kloosterman et al., 2015; Meindertsma, Kloosterman, Nolte, Engel, & Donner, 2017; Werkle-Bergner et al., 2014)(Figure 3A and 3C, lower panels). In the next section, we used the topographies of the high-frequency post-stimulus effects in visual cortex (Figures 3B and 3C, top panels) to identify a prestimulus neural mechanism that could explain the observed biased evidence accumulation resulting from the experimental decision bias manipulation.

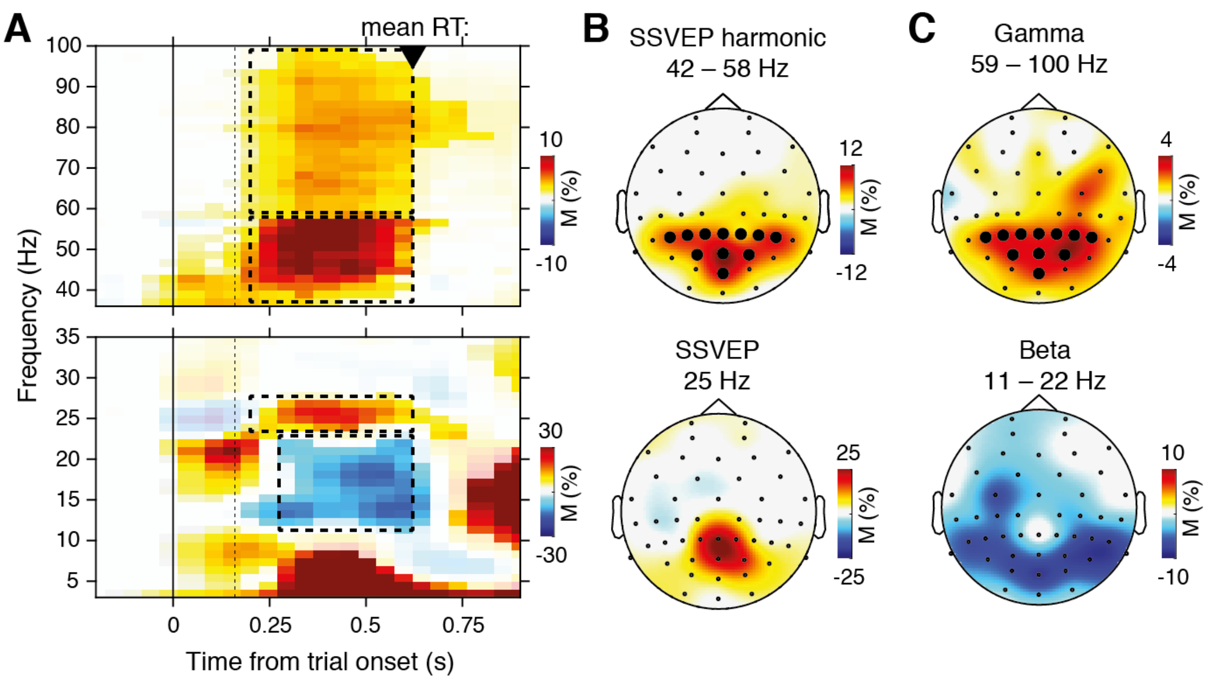

Figure 3 I Task-relevant textures induce stimulus-related responses in visual cortex. A. Time-frequency representations of high- (top) and low-frequency (bottom) EEG power modulations with respect to the prestimulus period (-0.4 - 0 s), pooled over the two conditions. Saturated colors indicate clusters of significant modulation, cluster threshold $p < 0.05$, two-sided permutation test across participants, cluster-corrected; N = 15). Solid and dotted vertical lines respectively indicate the onset of the trial and the target stimulus. M, power modulation. **B.** Scalp maps showing topography of the SSVEP power modulation around 25 Hz (top) and its harmonic from 42 – 58 Hz (bottom), from 0.2 – 0.6 s after trial onset. **C.** 59-100 Hz gamma power modulation from 0.2 – 0.6 s (top) and concurrent low frequency ('beta') power suppression from 11 – 22 Hz; see dashed outlines on time-frequency representations in A. Thick dots indicate electrodes used for the time-frequency representations in A, and which were selected for further analysis.

Adopting a liberal decision bias suppresses prestimulus alpha power

As a first step, we examined prestimulus power between 0.8 and 0.2 s before trial onset, using the same electrodes that showed the strongest post-stimulus effects (Figure 4A). This uncovered a highly specific modulation in the alpha range, which we confirmed to be strongest over the same cortical region that showed strong modulation in the gamma range (Figure 4B, white dots indicate electrodes showing stimulus-related gamma modulation). Indeed, when expressing spectral power during the liberal condition as the percentage signal change from the conservative condition, we observed a statistically significant cluster of suppressed frequencies precisely in the 8—12 Hz frequency range ($p < 0.05$, cluster-corrected for multiple comparisons) (Figure 4C), which again showed a posterior topography (Figure 4D). This shows that an experimentally induced liberal decision bias suppresses prestimulus alpha power, suggesting that alpha modulations are a hallmark of strategic bias adjustment rather than a mere correlate of spontaneous shifts in decision bias. Importantly, this finding implies that humans are able to actively control prestimulus alpha power in visual cortex, plausibly acting to bias sensory evidence accumulation towards the response alternative that maximizes rewards.

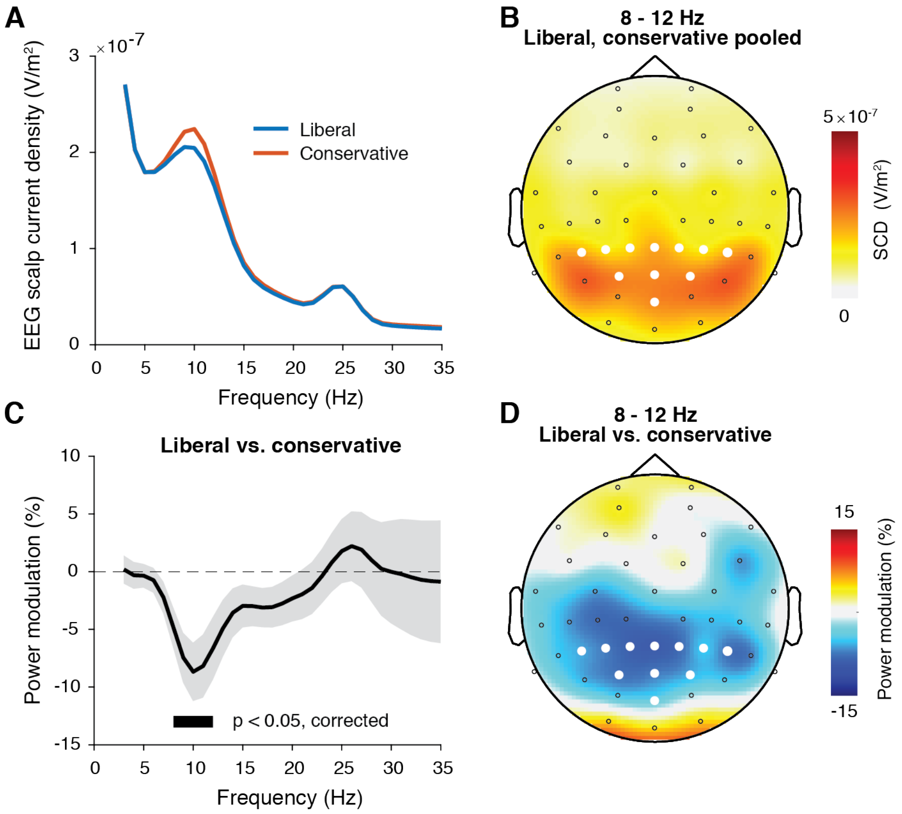

Figure 4 I Adopting a liberal decision bias suppresses prestimulus alpha power. A. Low-frequency power spectra of prestimulus neural activity for both conditions based on the electrodes that show large post-stimulus power modulations in Figure 3B and C (top panels). **B.** Scalp map of raw prestimulus EEG alpha power (8—12 Hz neural activity between 0.8 and 0.2 s before sequence onset), pooled over conditions. White symbols indicate visual cortical electrodes used for the power spectra in A. and C. **C.** Liberal versus conservative power spectrum. Black horizontal bar indicates statistically significant frequency range ($p < 0.05$, cluster-corrected for multiple comparisons, two-sided). Error bars, SEM across participants (N = 15). **D.** The corresponding scalp map of power modulation in the liberal condition, expressed as percent signal change from the conservative condition.

Alpha suppression enhances the gain of cortical gamma responses

How could suppression of prestimulus alpha activity bias the process of sensory evidence accumulation? One possibility is that alpha suppression influences evidence accumulation by modulating the susceptibility of visual cortex to sensory stimulation, a phenomenon dubbed 'neural excitability' (Iemi et al., 2017; Jensen & Mazaheri, 2010). We explored this possibility using a theoretical framework put forward by Rajagovindan and Ding (2011). This framework assumes that the relationship between total synaptic input activity that a neuronal ensemble receives and the total output activity it produces is characterized by a sigmoidal function (Figure 5A); a notion that is biologically plausible (Destexhe, Rudolph, Fellous, & Sejnowski, 2001; Freeman, 1979). Within this framework, both sensory input (i.e. as a result of sensory stimulation) and ongoing fluctuations in endogenous neural activity (i.e. levels of neural excitability) contribute to the synaptic input into visual cortex. The isolated effect of sensory input on the total output (i.e. the gain of the output response as caused by an input stimulus; see marked interval in Figure 5A), can then be expressed as the first order derivative (the slope) of the sigmoid in Figure 5A. In our experiment, stimulus-related input activity can be assumed to be more or less constant across trials since the same stimulus sequence was shown in each trial (see Figure 2A). Thus, modulations in the stimulus-related output gain generated in visual cortex are largely determined by the brain's excitability state. This can be seen in Figure 5B, where the stimulus-related output gain (the first order derivative, or slope from Figure 5A) is plotted as a function of neural excitability, yielding an inverted-U shaped function.

Figure 5B then shows the *effective range* in which the impact of neural excitability on the stimulus-related output response is largest, while its impact during low and high excitability is lower. When heightened excitability in the liberal condition is observed, this framework predicts enhanced output activity in visual cortex when compared to the conservative condition (Figure 5B), in particular when excitability differences between conditions occur in its effective range (Rajagovindan & Ding, 2011) (i.e. steeper slope of the solid blue curve compared to the red curve in Figure 5B).

We tested this model in our data by following the method put forward by Rajagovindan and Ding (2011), in which we operationalized neural excitability as pre-stimulus alpha (Jensen & Mazaheri, 2010), and stimulus-related output gain as post-stimulus gamma (Ni et al., 2016). We exploited the large number of trials per participant per condition across the multiple sessions in our study (range 543 to 1391 trials) by sorting each participant's trials into ten excitability bins based on equal-sized ranges of descending (log-transformed) prestimulus alpha power (indicating increasing excitability), separately for the conservative and liberal conditions. We subsequently computed and averaged the (log-transformed) gamma power across the trials within each excitability bin. Following Rajagovindan and Ding (2011), we removed individual differences in overall gamma power magnitude by subtracting the lowest binned gamma observation in the conservative condition from all observations. Finally, we plotted normalized gamma power as a function of excitability, separately for liberal and conservative (Figure 5C, see Methods for details).

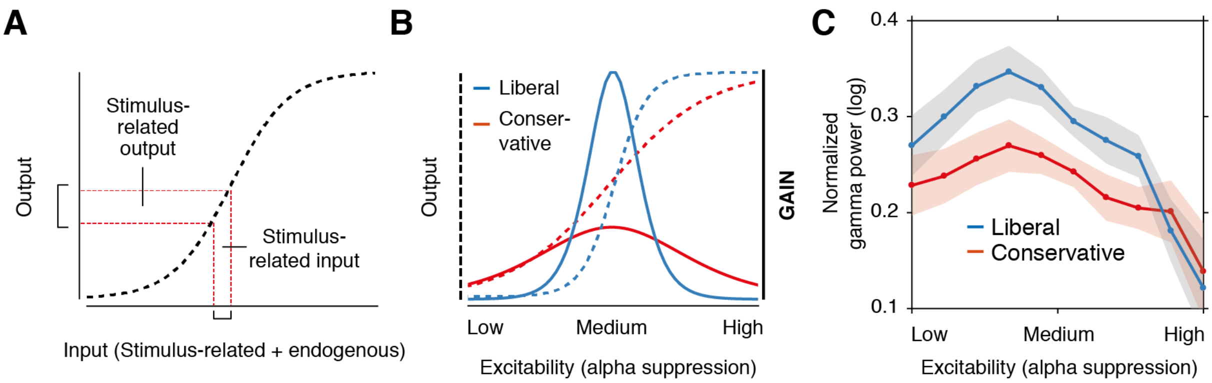

Figure 5 I Neural excitability boosts visual cortical responses by enhancing gain. A. Theoretical response gain model, which describes the transformation of input (both stimulus-related and endogenous) to the total output activity in visual cortex as a sigmoidal function. **B.** Model predictions. Stimulus-related output responses (solid lines) are formalized as the first derivative of the sigmoidal functions (dotted lines), resulting in inverse-U shaped response gain functions. The model predicts that a liberal decision bias increases the steepness of the sigmoidal function (right) compared to a conservative bias (left), resulting in stronger stimulus-related responses due to higher gain (Rajagovindan & Ding, 2011). **C.** Corresponding empirical test. Log-transformed gamma activity (normalized within participants by subtracting the minimum gamma power during the conservative condition from all bins) plotted as a function of neural excitability. Error bars, within-subject SEM across participants (N = 14).

The resulting plot indeed closely follows an inverted-U shaped relationship between excitability and stimulus-related gamma activity for both conditions, with particularly low gamma responses for the highest excitability bins (Figure 5C). Critically, average gamma power was higher in the liberal than in the conservative condition, except during the highest excitability bins (Figure 5C, rightmost two data points). Indeed, the flanks of the inverted-U curve for the liberal condition were steeper for the liberal condition, suggesting increased response gain. A three-way repeated measures ANOVA with factors condition (conservative, liberal), brain activity type (prestimulus alpha, poststimulus gamma) and bin level (1–10) revealed a significant three-way interaction ($F(9,117) = 2.96$, $p = 0.003$, partial $\eta^2 = 0.19$, Greenhouse-

Geisser corrected p = 0.046). Importantly, the marginally significant quadratic contrast ($F(1,13) = 3.47$, $p = 0.085$, partial $\eta^2 = 0.21$) fitted this interaction almost as well as a linear contrast ($F(1,13) = 4.69$, $p = 0.049$, partial $\eta^2 = 0.265$). This three-way quadratic interaction effect indeed suggests a more steeply U-shaped curve for gamma responses in the liberal condition, in line with enhanced gain. Taken together, these findings indicate that increased excitability during the liberal condition boosted input-related activity, which in turn indiscriminately biased sensory evidence accumulation towards 'yes' responses. In the next section, we confirm a direct link between drift bias and cortical stimulus response gain as measured through gamma.

Visual cortical gamma activity predicts strength of evidence accumulation bias

The findings presented so far suggest that behaviorally, a liberal decision bias shifts evidence accumulation towards 'yes' responses (drift bias in the DDM), while neurally it results in an increase of prestimulus cortical excitability concomitant with post-stimulus response gain expressed in gamma modulation in visual cortex. In a final analysis, we asked whether increases in gamma activity are directly related to a stronger drift bias. We predicted such a direct correspondence during the liberal condition, in which both drift bias and gamma activity were increased (see Figures 2F and 5C), but not during the conservative condition, in which drift bias was around zero and gamma was weaker than during liberal.

To test these predictions, we again applied the drift bias DDM to the behavioral data, but now freed the drift bias parameter not only for the two conditions, but also for the ten alpha suppression bins used to show the inverted-U-shaped relationship between excitability and stimulus-related gamma (see Figure 5C). We normalized the bin-resolved drift bias and gamma scalar values by z-scoring within each participant to remove individual differences in their ranges and averaged across participants within each alpha (excitability) bin. Finally, we directly tested the correspondence between drift bias and gamma using a within-subject group regression. Gamma activity indeed accurately predicted drift bias in the liberal condition ($R^2(9) = 0.77$, $p = 0.0008$, Figure 6 left panel). In contrast, drift bias was not well predicted by the corresponding gamma activity in the conservative condition ($R^2(9) = 0.0001$, $p = 0.98$,

Figure 6 right panel), which is perhaps unsurprising given the fact that drift bias was around zero in the conservative condition (see Figure 2F). Accordingly, predictive power was significantly greater in the liberal than in the conservative condition ($R^2_{LIB} - R^2_{CONS} = 0.77$, $p = 0.01$). The increase in gamma power in liberal versus conservative also predicted the increase in drift bias across the conditions ($R^2(9) = 0.56$, $p = 0.0126$, Figure S4), further suggesting that the experimental bias manipulation indeed enhanced gamma activity across excitability bins. We obtained qualitatively similar results without averaging across participants, but instead correlating across bins of all participants together using either ten or five bins per participant, suggesting these effects were not driven by single participants (data not shown). Taken together, these results show that enhanced post-stimulus gamma activity during the liberal condition underlies the evidence accumulation bias reflected in the drift bias parameter of the drift diffusion model.

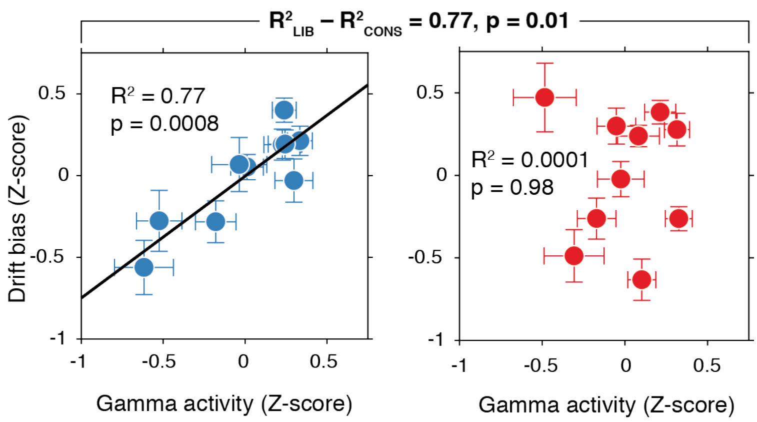

Figure 6 I Visual cortical gamma activity predicts strength of evidence accumulation bias. Linear regression of drift bias on gamma activity, separately for the two conditions. Gamma activity accurately predicts drift bias in the liberal (left), but not the conservative condition (right). Each dot represents an excitability bin and is obtained after averaging across participants (N = 14, see Methods for details). Error bars, SEM across participants.

Discussion

Humans possess a remarkable ability to strategically shift decision biases in order to flexibly adapt to the environment and maximize rewards. Traditionally, bias has been conceptualized in SDT as a criterion threshold that can be shifted towards or away from a noise or signal distribution. To date, however, the neural underpinnings of such bias shifts have remained elusive. Here, in contrast, we use a DDM drift bias model to demonstrate that an experimentally induced bias shift affects the process of sensory evidence accumulation itself, rather than shifting a threshold entity as SDT implies. Moreover, we reveal the neural signature of drift bias, by showing that a liberal decision bias increases alpha suppression (neural excitability) of visual cortex.

Although previous studies have shown correlations between suppression of prestimulus alpha (8—12 Hz) power and a liberal decision bias (Iemi et al., 2017; Limbach & Corballis, 2016), these studies have not established the effect of experimentally induced bias shifts. In the current study, by experimentally manipulating decision bias we show for the first time that prestimulus alpha plays an instrumental, and not merely a correlational role in decision bias. Further, we show that alpha suppression in turn boosts stimulus-related gamma activity through increased cortical response gain. Critically, gamma activity accurately predicted the strength of the drift bias parameter in the DDM drift bias model. Together, these findings show for the first time that humans are able to actively implement decision biases by flexibly adapting neural excitability to strategically shift sensory evidence accumulation towards one of two decision bounds.

Based on our results, we propose that decision biases are implemented by flexibly adjusting neural excitability in visual cortex. Figure 7 summarizes this proposed mechanism graphically by visualizing a hypothetical transition in neural excitability following an experimentally induced liberal decision bias, as reflected in visual cortical alpha suppression (left panel). This increased excitability translates into stronger gamma-band responses following stimulus onset (right panel, top). This increased gamma gain finally biases evidence accumulation towards the 'yes' decision boundary during a liberal state, resulting in more 'yes' responses, whereas 'no' responses are decimated (blue RT distributions; right panel, bottom). Our experimental manipulation of decision bias in different blocks of trials suggests that decision makers are able to control this biased evidence accumulation mechanism willfully by adjusting excitability, as reflected in alpha.

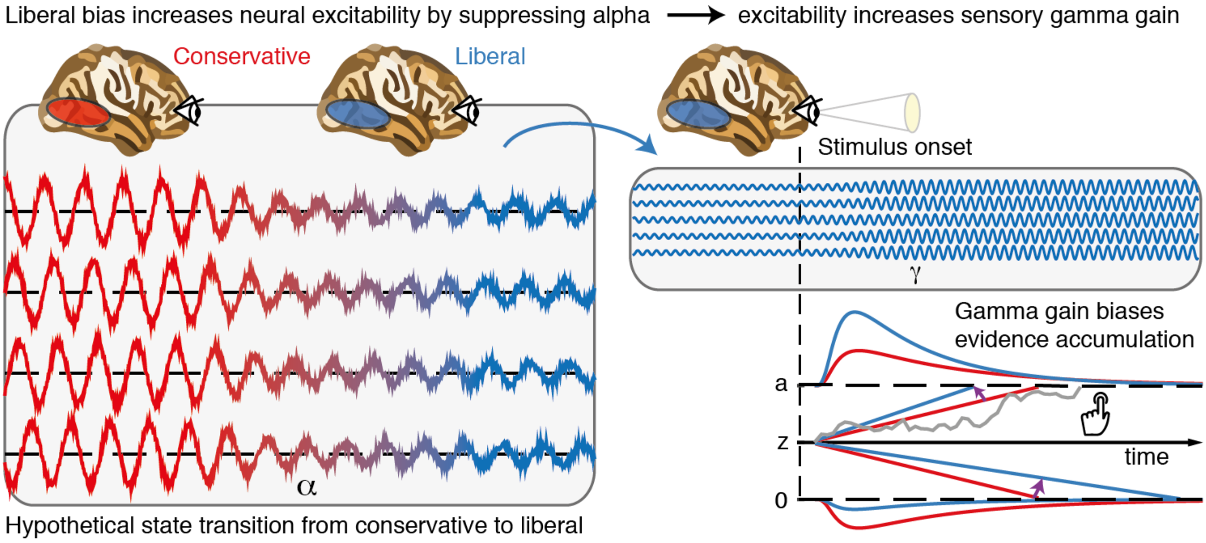

Figure 7 I Illustrative graphical depiction of the excitability state transition from conservative to liberal, and subsequent biased evidence accumulation under a liberal bias. The left panel shows the transition from a conservative to a liberal stimulus block. The experimental induction of a liberal decision bias causes alpha suppression in visual cortex, which increases neural excitability. The right top panel shows increased gamma-gain for incoming sensory evidence under conditions of high excitability. The right bottom panel shows how increased gamma-gain causes a bias in the drift rate, resulting in more 'target present' responses than in the conservative state.

A neural mechanism that could underlie bias-related alpha suppression may be under control of the catecholaminergic neuromodulatory systems, consisting of the noradrenaline-releasing locus coereleus (LC) and dopamine systems (Aston-Jones & Cohen, 2005). These systems are able to modulate the level of arousal and neural gain, and show tight links with pupil responses (de Gee et al., 2017; de Gee, Knapen, & Donner, 2014; Joshi, Li, Kalwani, & Gold, 2015; McGinley, David, & McCormick, 2015). Accordingly, prestimulus alpha power suppression has also recently been linked to pupil dilation (Meindertsma et al., 2017). From this perspective, our results reconcile previous studies showing relationships between a liberal bias, suppression of spontaneous alpha power and increased pupil size. Consistent with this, a recent monkey study observed increased neural activity during a liberal bias in the superior colliculus (Crapse, Lau, & Basso, 2018), a mid-brain structure tightly interconnected with the LC (Joshi et al., 2015). Taken together, a more liberal within-person bias (following experimental instruction) might activate neuromodulatory systems that subsequently increase cortical excitability and enhance sensory responses for both stimulus and 'noise' signals in visual cortex, thereby increasing a person's propensity for 'yes' responses (Iemi et al., 2017).

Rather than a link between alpha activity and decision bias, several previous studies have reported a link between alpha and task performance, particularly in the phase of alpha oscillations (Busch, Dubois, & VanRullen, 2009; Mathewson, Gratton, Fabiani, Beck, & Ro, 2009). Our findings can be reconciled with those by considering that detection sensitivity in many previous studies was often quantified in terms of raw stimulus detection rates, which do not dissociate objective sensitivity from response bias (see Figure 2B) (Green & Swets, 1966). Indeed, our findings are in line with recently reported links between decision bias and spontaneous fluctuations in excitability (Iemi et al., 2017; Iemi & Busch, 2017; Limbach & Corballis, 2016), suggesting an active role of neural excitability in decision bias.

Relatedly, a concern regarding our findings could be that the observed change in cortical excitability reflects a change in detection sensitivity (drift rate) rather than an intentional bias shift. This is unlikely because that would predict effects opposite to those we observed. We found increased excitability in the liberal condition compared to the conservative condition. If this were related to improved detection performance, one would predict higher sensitivity in the liberal condition, while in fact we found higher sensitivity in the conservative condition (compare drift rate to drift bias in both conditions in Fig. 2C). This finding convincingly ties cortical excitability in our paradigm to a strategically applied bias shift, as opposed to a change in detection sensitivity. Convergently, other studies also report a link between prestimulus low-frequency EEG activity and subjective perception, but not objective task performance (Benwell et al., 2017; Iemi & Busch, 2017).

Summarizing, our results show that stimulus-related responses are boosted during a liberal decision bias due to increased cortical excitability, in line with recent work linking alpha power suppression to response gain (Peterson & Voytek, 2017). Future studies can now establish whether this same mechanism is at play in other subjective aspects of decision-making, such as confidence and meta-cognition (Fleming, Putten, & Daw, 2018; Samaha et al., 2017) as well as in a dynamically changing environment (Norton, Fleming, Daw, & Landy, 2017). Explicit manipulation of cortical response gain during a bias manipulation by pharmacological manipulation of the noradrenergic LC-NE system (Servan-Schreiber, Printz, & Cohen, 1990) or by enhancing occipital alpha power using transcranial stimulation (Zaehle, Rach, & Herrmann, 2010) would further establish the underlying mechanisms involved in decision bias. In the end, although one may be unaware, every decision we make is influenced by biases that operate on the noisy evidence accumulation process towards one of the decision bounds. Understanding how these biases affect our decisions is key to becoming aware of these biases (Pleskac, Cesario, & Johnson, 2017), allowing us to control or invoke them adaptively. Pinpointing the neural mechanisms underlying bias in an elementary perceptual task (as used here) may paves the way for understanding how more abstract and high-level decisions are modulated by decision bias (Tversky & Kahneman, 1974).

Acknowledgments

This work was supported by the Max Planck Society. In addition, DDG and NAK are supported by an Emmy Nöther Programme grant from the German Research Foundation (to DDG) and by the Max Planck UCL Centre for Computational Psychiatry and Ageing Research. MWB is supported by a grant from the German Research Foundation (DFG; WE4296/5-1), as well as the Jacobs Foundation via an Early Career Research Fellowship.

Author Contributions

NAK and JJF designed research, NAK performed research, NAK, JWdG and JJF analyzed data, MWB and DDG provided theoretical background, NAK, JJF, UL, MWB, DDG, JWdG wrote the paper, NAK, JJF, UL, MWB, DDG, JWdG edited and commented on the manuscript.

Declaration of Interests

The authors declare no competing interests.

References

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*(1), 403–450. http://doi.org/10.1146/annurev.neuro.28.061604.135709

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., et al. (2015). Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron*, *85*(2), 390–401. http://doi.org/10.1016/j.neuron.2014.12.018

Benwell, C. S. Y., Tagliabue, C. F., Veniero, D., Cecere, R., Savazzi, S., & Thut, G. (2017). Pre-stimulus EEG power predicts conscious awareness but not objective visual performance. *eNeuro*, *4*(6), ENEURO.0182–17.2017. http://doi.org/10.1523/ENEURO.0182-17.2017

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. http://doi.org/10.1037/0033-295X.113.4.700

Busch, N. A., Dubois, J., & VanRullen, R. (2009). The Phase of Ongoing EEG Oscillations Predicts Visual Perception. *Journal of Neuroscience*, *29*(24), 7869–7876. http://doi.org/10.1523/JNEUROSCI.0113-09.2009

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science (New York, NY)*, *304*(5679), 1926–1929. http://doi.org/10.1126/science.1099745

Crapse, T. B., Lau, H., & Basso, M. A. (2018). A Role for the Superior Colliculus in Decision Criteria. *Neuron*, *97*(1), 181–194.e6. http://doi.org/10.1016/j.neuron.2017.12.006

de Gee, J. W., Colizoli, O., Kloosterman, N. A., Knapen, T., Nieuwenhuis, S., & Donner, T. H. (2017). Dynamic modulation of decision biases by brainstem arousal systems. *eLife*, *6*, 309. http://doi.org/10.7554/eLife.23232

de Gee, J. W., Knapen, T., & Donner, T. H. (2014). Decision-related pupil dilation reflects upcoming choice and individual bias. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(5), E618–25. http://doi.org/10.1073/pnas.1317557111

Destexhe, A., Rudolph, M., Fellous, J. M., & Sejnowski, T. J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, *107*(1), 13–24. http://doi.org/10.1016/S0306-4522(01)00344-X

Donner, T. H., & Siegel, M. (2011). A framework for local cortical oscillation patterns. *Trends in Cognitive Sciences*, *15*(5), 191–199. http://doi.org/10.1016/j.tics.2011.03.007

Efron, B., & Tibshirani, R. (1998). The problem of regions. *The Annals of Statistics*, *26*(5), 1687–1718. http://doi.org/10.1214/aos/1024691353

Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2007). Masking disrupts reentrant processing in human visual cortex. *Journal of Cognitive Neuroscience*, *19*(9), 1488–1497.

http://doi.org/10.1162/jocn.2007.19.9.1488&url_ctx_fmt=info:ofi/fmt:kev:mtx:ctx&rft_val_fmt=info:ofi/fmt:kev:mtx:journal&rft.atitle=Masking

Fetsch, C. R., Kiani, R., & Shadlen, M. N. (2014). Predicting the Accuracy of a Decision: A Neural Mechanism of Confidence. *Cold Spring Harbor Symposia on Quantitative Biology*, *79*, 185–197. http://doi.org/10.1101/sqb.2014.79.024893

Fleming, S. M., Putten, E. J., & Daw, N. D. (2018). Neural mediators of changes of mind about perceptual decisions. *Nature Neuroscience*, *21*(4), 617. http://doi.org/10.1038/s41593-018-0104-6

Freeman, W. J. (1979). Nonlinear gain mediating cortical stimulus-response relations. *Biological Cybernetics*, *33*(4), 237–247. http://doi.org/10.1007/BF00337412

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. http://doi.org/10.1146/annurev.neuro.29.051605.113038

Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics. *Society*, *1*, 521.

Hassler, U., Trujillo-Barreto, N., & Gruber, T. (2011). Induced gamma band responses in human EEG after the control of miniature saccadic artifacts. *NeuroImage*, *57*(4), 1411–1421. http://doi.org/10.1016/j.neuroimage.2011.05.062

Hipp, J. F., & Siegel, M. (2013). Dissociating neuronal gamma-band activity from cranial and ocular muscle activity in EEG. *Frontiers in Human Neuroscience*, *7*, 338. http://doi.org/10.3389/fnhum.2013.00338

Iemi, L., & Busch, N. A. (2017). Moment-to-moment fluctuations in neuronal excitability bias subjective perception rather than decision-making. *bioRxiv*, 151324. http://doi.org/10.1101/151324

Iemi, L., Chaumon, M., Crouzet, S. M., & Busch, N. A. (2017). Spontaneous Neural Oscillations Bias Perception by Modulating Baseline Excitability. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *37*(4), 807–819. http://doi.org/10.1523/JNEUROSCI.1432-16.2017

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Frontiers in Human Neuroscience*, *4*, 186. http://doi.org/10.3389/fnhum.2010.00186

Joshi, S., Li, Y., Kalwani, R. M., & Gold, J. I. (2015). Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron*, *0*(0), 221–234. http://doi.org/10.1016/j.neuron.2015.11.028

Kloosterman, N. A., Meindertsma, T., Hillebrand, A., van Dijk, B. W., Lamme, V. A. F., & Donner, T. H. (2015). Top-down modulation in human visual cortex predicts the stability of a perceptual illusion. *Journal of Neurophysiology*, *113*(4), 1063–1076. http://doi.org/10.1152/jn.00338.2014

Limbach, K., & Corballis, P. M. (2016). Prestimulus alpha power influences response criterion in a detection task. *Psychophysiology*, *53*(8), 1154–1164. http://doi.org/10.1111/psyp.12666

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–190. http://doi.org/10.1016/j.jneumeth.2007.03.024

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To See or Not to See: Prestimulus α Phase Predicts Visual Awareness. *Journal of*

Neuroscience, *29*(9), 2725–2732. http://doi.org/10.1523/JNEUROSCI.3963-08.2009

McGinley, M. J., David, S. V., & McCormick, D. A. (2015). Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. *Neuron*, *87*(1), 179–192. http://doi.org/10.1016/j.neuron.2015.05.038

Meindertsma, T., Kloosterman, N. A., Nolte, G., Engel, A. K., & Donner, T. H. (2017). Multiple Transient Signals in Human Visual Cortex Associated with an Elementary Decision. *Journal of Neuroscience*, *37*(23), 5744–5757. http://doi.org/10.1523/JNEUROSCI.3835-16.2017

Melloni, L., Schwiedrzik, C. M., Wibral, M., Rodriguez, E., & Singer, W. (2009). Response to: Yuval-Greenberg et al., “Transient Induced Gamma-Band Response in EEG as a Manifestation of Miniature Saccades.” Neuron 58, 429-441. *Neuron*, *62*(1), 8–10– author reply 10–12. http://doi.org/10.1016/j.neuron.2009.04.002

Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J.-M., Kennedy, H., & Fries, P. (2016). Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. *Neuron*, *89*(2), 384–397. http://doi.org/10.1016/j.neuron.2015.12.018

Mitra, P. P., & Pesaran, B. (1999). Analysis of Dynamic Brain Imaging Data. *Biophysical Journal*, *76*(2), 691–708. http://doi.org/10.1016/S0006-3495(99)77236-X

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *32*(7), 2335–2343. http://doi.org/10.1523/JNEUROSCI.4156-11.2012

Neath, A. A., & Cavanaugh, J. E. (2012). The Bayesian information criterion: background, derivation, and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, *4*(2), 199–203. http://doi.org/10.1002/wics.199

Ni, J., Wunderle, T., Lewis, C. M., Desimone, R., Diester, I., & Fries, P. (2016). Gamma-Rhythmic Gain Modulation. *Neuron*, *92*(1), 240–251. http://doi.org/10.1016/j.neuron.2016.09.003

Norton, E. H., Fleming, S. M., Daw, N. D., & Landy, M. S. (2017). Suboptimal Criterion Learning in Static and Dynamic Environments. *PLoS Computational Biology*, *13*(1), e1005304–28. http://doi.org/10.1371/journal.pcbi.1005304

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, *2011*(1), 1–9. http://doi.org/10.1155/2011/156869

Perrin, F., Pernier, J., Bertrand, O., & Echallier, J. F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and Clinical Neurophysiology*, *72*(2), 184–187. http://doi.org/10.1016/0013-4694(89)90180-6

Peterson, E. J., & Voytek, B. (2017). Alpha oscillations control cortical gain by modulating excitatory-inhibitory background activity. *Biorxiv.org*
. http://doi.org/https://doi.org/10.1101/185074

Pleskac, T. J., Cesario, J., & Johnson, D. J. (2017). How race affects evidence accumulation during the decision to shoot. *Psychonomic Bulletin & Review*, *18*(2), 1–30. http://doi.org/10.3758/s13423-017-1369-6

Popov, T., Kastner, S., & Jensen, O. (2017). FEF-Controlled Alpha Delay Activity Precedes Stimulus-Induced Gamma-Band Activity in Visual Cortex. *Journal of Neuroscience*, *37*(15), 4117–4127. http://doi.org/10.1523/JNEUROSCI.3015-16.2017

Rajagovindan, & Ding, M. (2011). From prestimulus alpha oscillation to visual-evoked response: an inverted-U function and its attentional modulation. *Journal of Cognitive Neuroscience*, *23*(6), 1379–1394. http://doi.org/10.1162/jocn.2010.21478

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. http://doi.org/10.1162/neco.2008.12-06-420

Ratcliff, R., Huang-Pollock, C., & McKoon, G. (2016, August 15). Modeling Individual Differences in the Go/No-Go Task With a Diffusion Model. http://doi.org/http://dx.doi.org/10.1037/dec0000065

Samaha, J., Iemi, L., & Postle, B. R. (2017). Prestimulus alpha-band power biases visual discrimination confidence, but not accuracy. *Consciousness and Cognition*. http://doi.org/10.1016/j.concog.2017.02.005

Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior. *Science (New York, NY)*, *249*(4971), 892–895.

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science (New York, NY)*, *185*(4157), 1124–1131. http://doi.org/10.1126/science.185.4157.1124

Urai, A. E., de Gee, J. W., & Donner, T. H. (2018). Choice history biases subsequent evidence accumulation. *bioRxiv*, 251595. http://doi.org/10.1101/251595

van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., & Roelfsema, P. R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(40), 14332–14341. http://doi.org/10.1073/pnas.1402773111

Werkle-Bergner, M., Grandy, T. H., Chicherio, C., Schmiedek, F., Lovden, M., & Lindenberger, U. (2014). Coordinated within-Trial Dynamics of Low-Frequency Neural Rhythms Controls Evidence Accumulation. *Journal of Neuroscience*, *34*(25), 8519–8528. http://doi.org/10.1523/JNEUROSCI.3801-13.2014

White, C. N., & Poldrack, R. A. (2014). Decomposing bias in different types of simple decisions. *Journal of Experimental Psychology Learning, Memory, and Cognition*, *40*(2), 385–398. http://doi.org/10.1037/a0034851

Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Frontiers in Neuroinformatics*, *7*. http://doi.org/10.3389/fninf.2013.00014

Yuval-Greenberg, S., Tomer, O., Keren, A. S., Nelken, I., & Deouell, L. Y. (2008). Transient Induced Gamma-Band Response in EEG as a Manifestation of Miniature Saccades. *Neuron*, *58*(3), 429–441. http://doi.org/10.1016/j.neuron.2008.03.027

Zaehle, T., Rach, S., & Herrmann, C. S. (2010). Transcranial Alternating Current Stimulation Enhances Individual Alpha Activity in Human EEG. *PLoS ONE*, *5*(11), e13766. http://doi.org/10.1371/journal.pone.0013766

Materials and Methods

Participants Sixteen participants (eight female, mean age 24.1 years, ± 1.64) took part in the experiment, either for financial compensation or in partial fulfillment of first year course requirements. Each participant completed three experimental sessions on different days, each session lasting ca. 2 hours, including preparation and breaks. One participant completed only two sessions, yielding a total number of sessions across subjects of 47. Due to technical issues, for one session only data for the liberal condition was available. One participant was an author. All participants were included in the signal-detection-theoretical and drift diffusion modeling analyses (Figure 2). One participant was excluded from the stimulus-related and the prestimulus alpha analysis (Figures 3 and 4) due to excessive noise (EEG power spectrum opposite of 1/frequency). One further participant was excluded from the alpha-versus-gamma power modulation (Figure 5) and gamma-versus-drift bias analyses (Figure 6) because the liberal-conservative difference in gamma power in this participant was > 3 standard deviations away from the other participants. All participants had normal or corrected-to-normal vision and were right handed. Participants provided written informed consent before the start of the experiment. All procedures were approved by the ethics committee of the University of Amsterdam.

Stimuli Stimuli consisted of a continuous semi-random rapid serial visual presentation (rsvp) of full screen texture patterns. The texture patterns consisted of line elements approx. 0.07° thick and 0.4° long in visual angle. Each texture in the rsvp was presented for 40 ms (i.e. stimulation frequency 25 Hz), and was oriented in one of four possible directions: 0°, 45°, 90° or 135°. Participants were instructed to fixate a red dot in the center of the screen. At random inter trial intervals (ITI's) sampled from a uniform distribution (ITI range 0.3–2.2 s), the rsvp contained a fixed sequence of 25 texture patterns, which in total lasted one second. This fixed sequence consisted of four stimuli preceding a (non-)target stimulus (orientations of 45°, 90°, 0°, 90°

respectively) and twenty stimuli following the (non)-target (orientations of 0°, 90°, 0°, 90°, 0°, 45°, 0°, 135°, 90°, 45°, 0°, 135°, 0°, 45°, 90°, 45°, 90°, 135°, 0°, 135° respectively) (Figure 2). The fifth texture pattern within the sequence (occurring from 0.16 s after sequence onset) was either a target or a nontarget stimulus. Nontargets consisted of either a 45° or a 135° homogenous texture, whereas targets contained a central orientation-defined square of 2.42° visual angle, thereby consisting of both a 45° and a 135° texture. 50% of all targets consisted of a 45° square and 50% of a 135° square. Of all trials, 75% contained a target and 25% a nontarget. Target and nontarget trials were presented in random order. To avoid specific influences on target stimulus visibility due to presentation of similarly or orthogonally oriented texture patterns temporally close in the cascade, no 45° and 135° oriented stimuli were presented directly before or after presentation of the target stimulus. All stimuli had an isoluminance of 72.2 cd/m^2. Stimuli were created using MATLAB (The Mathworks, Inc., Natick, MA, USA) and presented using Presentation (Neurobehavioral systems, Inc., Albany, CA, USA).

Experimental design The participants' task was to detect targets and actively report them by pressing a button using their preferred hand. Targets occasionally went unreported, presumably due to constant forward and backward masking by the continuous cascade of stimuli and unpredictability of target timing (Fahrenfort, Scholte, & Lamme, 2007). The onset of the fixed order of texture patterns preceding and following (non-)target stimuli was neither signaled nor apparent.

At the beginning of the experiment, participants were informed they could earn a total bonus of EUR 30, on top of their regular pay or course credit. In two separate conditions within each session of testing, we encouraged participants to use either a conservative or a liberal bias for reporting targets using both aversive sounds as well as reducing their bonus after errors. In the conservative condition, participants were instructed to only press the button when they were relatively sure they had seen the target. The instruction on screen before block onset read as follows: "Try to detect as many targets as possible. Only press when you are relatively sure you just saw a target." To maximize effectiveness of this instruction, participants were told the bonus would be diminished by ten cents after a false alarm. During the experiment, a loud aversive sound was played after a false alarm to inform the participant about an error.

During the liberal condition, participants were instructed to miss as few targets as possible. The instruction on screen before block onset read as follows: “Try to detect as many targets as possible. If you sometimes press when there was nothing this is not so bad”. In this condition, the loud aversive sound was played twice in close succession whenever they failed to report a target, and three cents were subsequently deducted from their bonus. The difference in auditory feedback between both conditions was included to inform the participant about the type of error (miss or false alarm), in order to facilitate the desired bias in both conditions. After every block, the participant’s score (number of missed targets in the liberal condition and number of false alarms in the conservative condition) was displayed on the screen, as well as the remainder of the bonus. After completing the last session of the experiment, every participant was paid the full bonus as required by the ethical committee.

During a block, participants continuously monitored the screen and were free to respond by button press whenever they thought they saw a target. Each block contained 240 trials, of which 180 target and 60 nontarget trials. Participants performed six blocks per session. The task instruction was presented on the screen before the block started. The condition of the first block of a session was counterbalanced across participants. Prior to EEG recording in the first session, participants performed a 10-minute practice run of both conditions, in which visual feedback directly after a miss (liberal condition) or false alarm (conservative) informed participants about their mistake, allowing them to adjust their decision bias accordingly.

Behavioral analysis We calculated participants criterion *c* (Green & Swets, 1966) across the trials in each condition as follows:

$$c = -\frac{1}{2}\left[Z(Hit\text{-}rate) + Z(FA\text{-}rate)\right]$$

where Z(...) is the inverse standard normal distribution. Furthermore, we calculated objective sensitivity measure d’ using:

$$d' = Z(Hit\text{-}rate) - Z(FA\text{-}rate)$$

as well as by subtracting hit and false alarm rates. Reaction times (RT's) were measured as the period between target onset and button press.

Drift diffusion modeling of choice behavior We fitted the drift diffusion model to our behavioural data for each subject individually, and separately for the liberal and conservative conditions. We fitted the model using a *G* square method based on quantile RT's (RT cutoff, 200 ms, for details, see Ratcliff et al. (2016)), using a modified version of the HDDM 0.6.0 package (Wiecki, Sofer, & Frank, 2013) (code will be made available). The RT distributions for 'yes' responses were represented by the 0.1, 0.3, 0.5, 0.7 and 0.9 quantiles, and, along with the associated response proportions, contributed to *G* square. In addition, a single bin containing the number of 'no' responses contributed to *G* square. Fitting the model to RT distributions for 'yes' and 'no' choices (termed 'stimulus coding' in Wiecki et al. (2013)), as opposed to the more common fits of correct and incorrect choice RT's (termed 'accuracy coding' in Wiecki et al. (2013)), allowed us to estimate parameters that could have induced biases in subjects' behavior.

Parameter recovery simulations showed that letting both the the starting point of the accumulation process and drift bias (an evidence-independent constant added to the drift toward one or the other bound) free to vary with experimental conditions is problematic for data with no explicit "no" responses (data not shown). Thus, to test whether shifts in drift bias or starting point underlied bias we fitted three separate models. In the first model ('fixed model'), we allowed only the following parameters to vary between the liberal and conservative condition: (i) the mean drift rate across trials; (ii) the separation between both decision bounds (i.e., response caution); and (iii) the non-decision time (sum of the latencies for sensory encoding and motor execution of the choice). Additionaly, the bias parameters starting point and drift bias were fixed for the experimental conditions. The second model ('starting point model') was the same as the fixed model, except that we let the starting point of the accumulation process vary with experimental condition, whereas the drift bias was kept fixed for both conditions. The third model ('drift bias model') was the same as the fixed model, except that we let the drift bias vary with experimental condition, while the starting point was kept fixed for both conditions. We used Bayesian Information Criterion (BIC) to select the model which provided the best fit to the data (Neath & Cavanaugh, 2012). The BIC compares models based on their maximized log-likelihood value, while penalizing for the number of parameters.

EEG recording Continuous EEG data were recorded at 256 Hz using a 48-channel BioSemi Active-Two system (Biosemi, Amsterdam, the Netherlands), connected to a standard EEG cap according to the international 10-20 system. Electrooculography (EOG) was recorded using two electrodes at the outer canthi of the left and right eyes and two electrodes placed above and below the right eye. Horizontal and vertical EOG electrodes were referenced against each other, two for horizontal and two for vertical eye movements (blinks). We used the Fieldtrip toolbox (Oostenveld, Fries, Maris, & Schoffelen, 2011) and custom software in MATLAB R2016b (The Mathworks Inc., Natick, MA, USA) to process the data (see below). Data were re-referenced to the average voltage of two electrodes attached to the earlobes.

Trial extraction and preprocessing We extracted trials of variable duration from 1 s before target sequence onset until 1.25 after button press for trials that included a button press (hits and false alarms), and until 1.25 s after stimulus onset for trials without a button press (misses and correct rejects). The following constraints were used to classify (non-)targets as detected (hits and false alarms), while avoiding the occurrence of button presses in close succession to target reports and button presses occurring outside of trials: 1) A trial was marked as detected if a response occurred within 0.84 s after target onset; 2) when the onset of the next target stimulus sequence started before trial end, the trial was terminated at the next trial's onset; 3) when a button press occurred in the 1.5 s before trial onset, the trial was extracted from 1.5 s after this button press; 4) when a button press occurred between 0.5 s before until 0.2 s after sequence onset, the trial was discarded. See Kloosterman et al. (2015) and Meindertsma et al. (2017) for similar trial extraction procedures. After trial extraction, channel time courses were linearly detrended and the mean of every channel was removed per trial.

Artifact rejection Trials containing muscle artifacts were rejected from further analysis using a standard semi-automatic preprocessing method in Fieldtrip. This procedure consists of bandpass-filtering the trials of a condition block in the 110–125 Hz frequency range, which typically contains most of the muscle artifact activity, followed by a Z-transformation. Trials exceeding a threshold Z-score were removed completely from analysis. We used as the threshold the absolute value of the minimum Z-score within the block, + 1. To remove eye blink artifacts from the time courses, the EEG data from a complete session were transformed using independent component analysis (ICA), and components (typically one or two of the 48) due to blinks was removed from the data. In addition, to remove microsaccade-related artifacts we included two virtual channels in the ICA based on channels Fp1 and Fp2, which included transient spike potentials as identified using the algorithm from Hassler et al. (2011). The two components loading high on these virtual electrodes (typically with a frontal topography) were also removed. Blinks and eye movements were then semi-automatically detected from the horizontal and vertical EOG (frequency range 1–15 Hz; z-value cut-off 4 for vertical; 6 for horizontal) and trials containing eye artefacts within 0.1 s around target onset were discarded. This step was done to remove trials in which the target was not seen because the eyes were closed. Finally, trials exceeding a threshold voltage range of 200 μV were discarded. To attenuate volume conduction effects and suppress any remaining microsaccade-related activity, the scalp current density (SCD) was computed using the second-order derivative (the surface Laplacian) of the EEG potential distribution (Perrin et al., 1989).

Spectral analysis of EEG power We used a sliding window Fourier transform ((Mitra & Pesaran, 1999); step size, 50 ms; window length, 400 ms; frequency resolution, 2.5 Hz) to calculate time-frequency representations (spectrograms) of the EEG power for each electrode and each trial. We used a single Hann taper for the frequency range of 3–35 Hz (spectral smoothing, 4.5 Hz, bin size, 1 Hz) and the multitaper technique for the 36 – 100 Hz frequency range (spectral smoothing, 8 Hz; bin size, 2 Hz; five tapers). See Kloosterman et al. (2015) and Meindertsma et al. (2017) for similar settings.

Spectrograms were aligned to the onset of the stimulus sequence containing the (non)target. Power modulations (denoted as M in Figure 3) during the trials were quantified as the percentage of power change at a given time point and frequency bin, relative to a baseline power value for each frequency bin. We used as a baseline the mean EEG power in the interval 0.4 to 0 s before trial onset. If this interval was not completely present in the trial due to preceding events (see Trial extraction), this period was shortened accordingly. We subtracted the trial-specific baseline value from each sample in the time course per frequency bin and divided by the mean baseline power across all trials within a session. For the analysis of raw prestimulus power modulations no baseline correction was applied. We focused our analysis of EEG power modulations around target onsets on those electrodes that processed the visual stimulus. To this end, we averaged the power modulations or raw power across eleven occipito-parietal electrodes that showed stimulus-induced responses in the gamma-band range (59–100 Hz). See Kloosterman et al. (2015) and Meindertsma et al. (2017) for a similar procedure.

Condition-related EEG power modulation To test at which frequencies raw EEG power differed for the liberal and conservative conditions, we averaged power modulation from 0.8 s up to 0.2 s (i.e. up to half the window size used for spectral analysis, to avoid contamination of post- with pre-stimulus activity (Iemi et al., 2017)) from trial onset. Then, we expressed the power at each frequency in units of percent signal change with respect to the conservative condition and statistically tested whether this signal differed from zero (Figure 4C) (see Statistical comparisons).

Response gain model test To test the prediction of increased gain during liberal of the gain model, we first averaged activity in the 8–12 Hz range from 0.8 to 0.2 s before trial onset (staying half our window size from trial onset, to avoid mixing pre- and post-stimulus activity, also see Iemi et al. (2017)) and took the log transform, yielding a single scalar alpha power value per trial expressing neural excitability. If this interval was not completely present in the trial due to preceding events (see Trial extraction), this period was shortened accordingly. Trials in which the scalar was > 3 standard deviations away from the participant's mean were excluded. We then sorted all single trials for each participant in ascending order of excitability and assigned them to ten equally-spaced bins ranging from the lowest to the highest excitability scalars of that participant. Adjacent bin ranges overlapped for 50% to stabilize estimates. Then we averaged the corresponding log-transformed gamma modulation of these trials (consisting of the average power within 59–100 Hz 0.2 to 0.6 s after trial onset) and normalized each participants response by subtracting the minimum gamma power during the conservative condition from all bins. Finally, we averaged across participants and plotted the excitability bin number against the normalized gamma power for each condition. See Rajagovindan and Ding (2011) for a similar procedure. To statistically test the gain prediction, we employed a three-way repeated measures ANOVA (see Statistical comparisons). For plotting purposes (Figure 5C), we computed within-subject error bars by removing within each participant the mean across conditions from the estimates.

Correlation between gamma power and drift bias To link DDM drift bias and cortical gamma power, we re-fitted the DDM drift bias model while freeing the drift bias parameter within each condition for the ten neural excitability bins as determined by prestimulus alpha suppression (see section Response gain model test), while freeing the other parameters (drift rate, boundary separation, non-decision time) for each condition and fixing starting point across conditions. Then, we normalized the obtained scalars for gamma power and drift bias separately within participants using a Z-transformation, and averaged across participants. Finally, we used within-subject group regression of the two measures across the ten bins for both conditions separately. In a control analysis, we conducted this regression after taking the liberal – conservative difference for each excitability bin before regressing, and obtained convergent results (Figure S4).

Statistical comparisons We used two-sided permutation tests (10,000 permutations) (Efron & Tibshirani, 1998) to test the significance of behavioral effects and the model fits. To quantify power modulations after (non-)target onset, we tested the overall power modulation for significant deviations from zero. For these tests, we used a cluster-based permutation procedure to correct for multiple comparisons (Maris & Oostenveld, 2007). For time-frequency representations of power modulation, this procedure was conducted across all time-frequency bins. For frequency spectra, this procedure was performed across all frequency bins. To test whether there was evidence for increased gain in the liberal compared to the conservative condition, we conducted a three-way repeated measures ANOVA (condition (conservative, liberal) x brain activity type (prestimulus alpha, poststimulus gamma power) x bin level (1–10)) using SPSS 23 (IBM, Inc.), inspecting linear and quadratic contrasts. As sphericity was violated in this model ($p = 0.0001$), we report both the uncorrected and Greenhouse-Geisser-corrected p-values. We used Pearson correlation to test the link between gamma power and drift bias. We tested the difference in correlation between the liberal and conservative conditions using the Fisher r-to-Z transformation and obtained the corresponding two-tailed p-value.

Supplemental Information

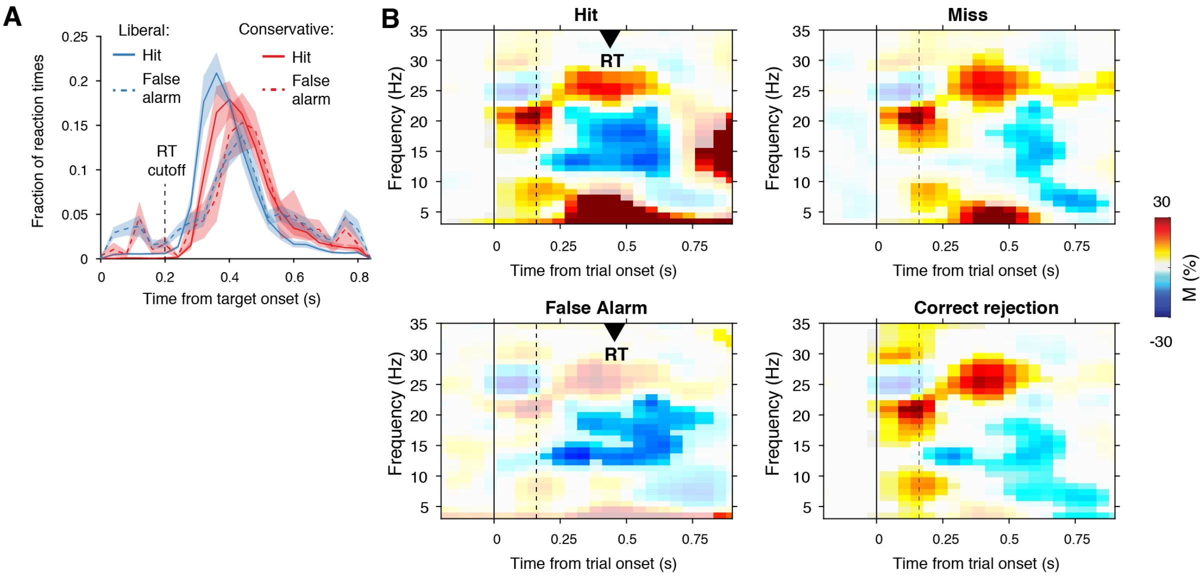

Figure S1 l Behavioral and neurophysiological evidence that participants were sensitive to the implicit task structure. A. Participant-average RT distributions for hits and false alarms in both conditions. The presence of similar RT distributions for false alarms and hits indicates that participants were sensitive to trial onset despite the fact that trial onsets were only implicitly signalled. Error bars, SEM. **B.** Time-frequency representations of low-frequency EEG power modulations with respect to the prestimulus period (–0.4 – 0 s), pooled across the two conditions. Significant low-frequency modulation occurred even for nontarget trials without overt response (correct rejections), indicating that participants detected the onset of a trial even when neither a target was presented nor a response was given. Saturated colors indicate clusters of significant modulation, cluster threshold $p < 0.05$, two-sided permutation test across participants, cluster-corrected; N = 15). Solid and dotted vertical lines respectively indicate the onset of the trial and the target stimulus. M, power modulation.

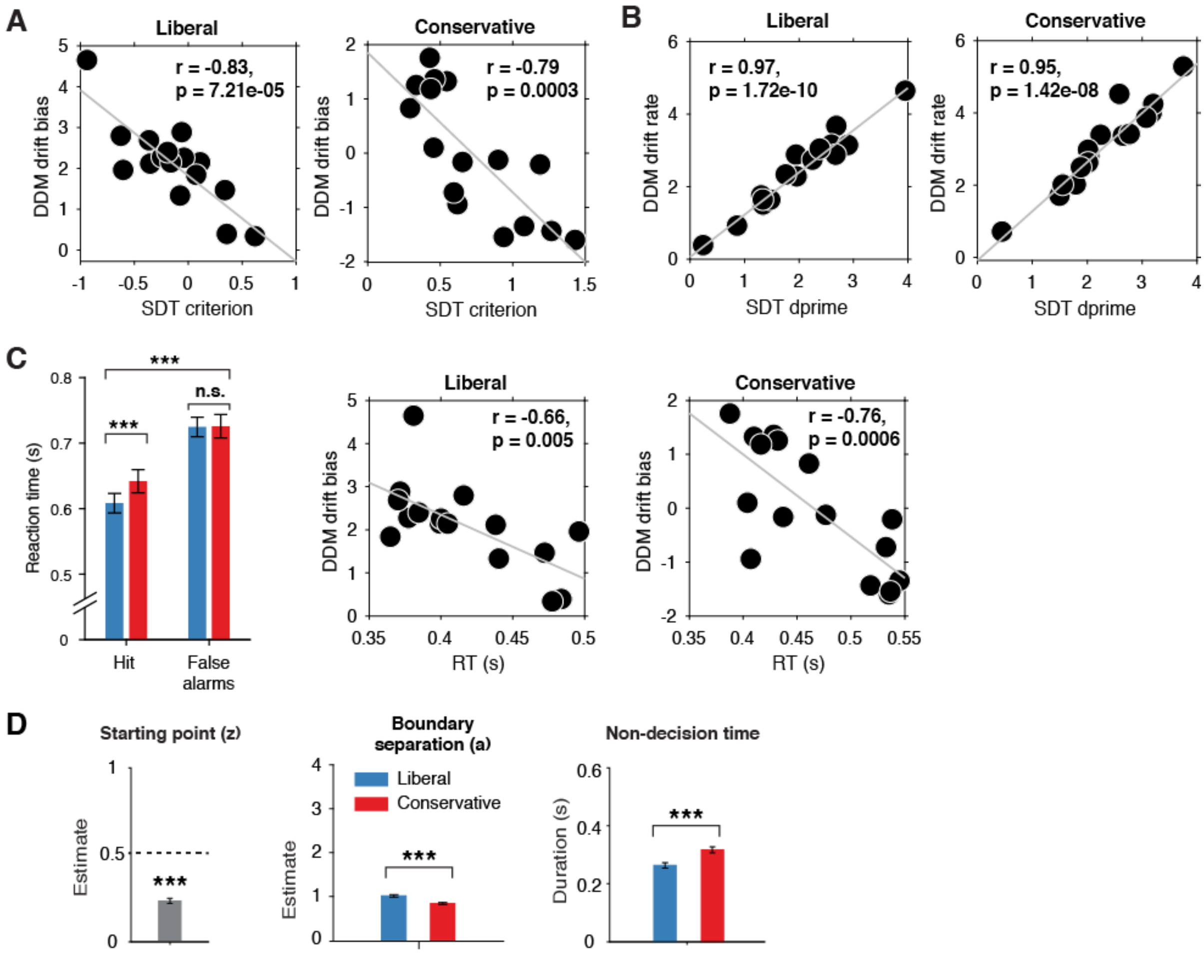

Figure S2 I Signal-detection-theoretic (SDT) behavioral measures during both conditions correspond closely to drift diffusion modeling (DDM) parameters. A. Across-participant Pearson correlation between criterion and DDM drift bias for the two conditions. The correlation is negative due to a lower criterion reflecting a stronger liberal bias. Each dot represents a participant. **B.** As A. but for correlation between dprime and drift rate. **C.** Left panel, mean reaction times (RT) for hits and false alarms for the two conditions. Middle and right panels, As A. but for correlation between RT and drift bias. **D.** Parameter estimates in the drift bias DDM not related to evidence accumulation (drift). ****p* < 0.001; n.s., not significant.

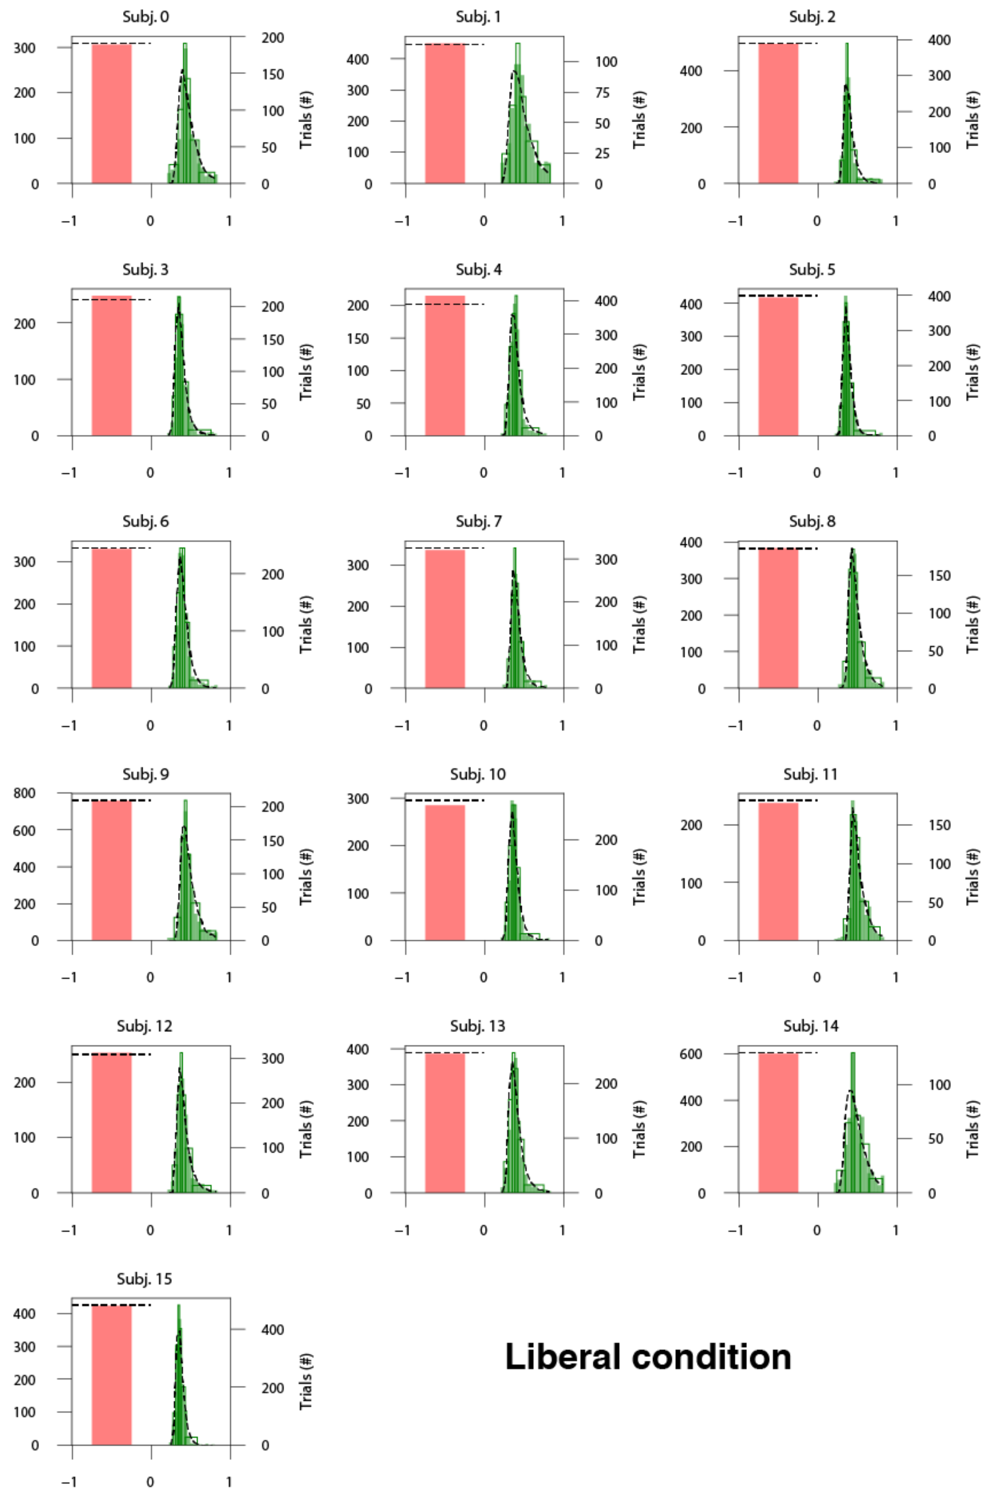
Subj. 0
Subj. 1
Subj. 2
Subj. 3
Subj. 4
Subj. 5
Subj. 6
Subj. 7
Subj. 8
Subj. 9
Subj. 10
Subj. 11
Subj. 12
Subj. 13
Subj. 14
Subj. 15
Trials (#)
Liberal condition

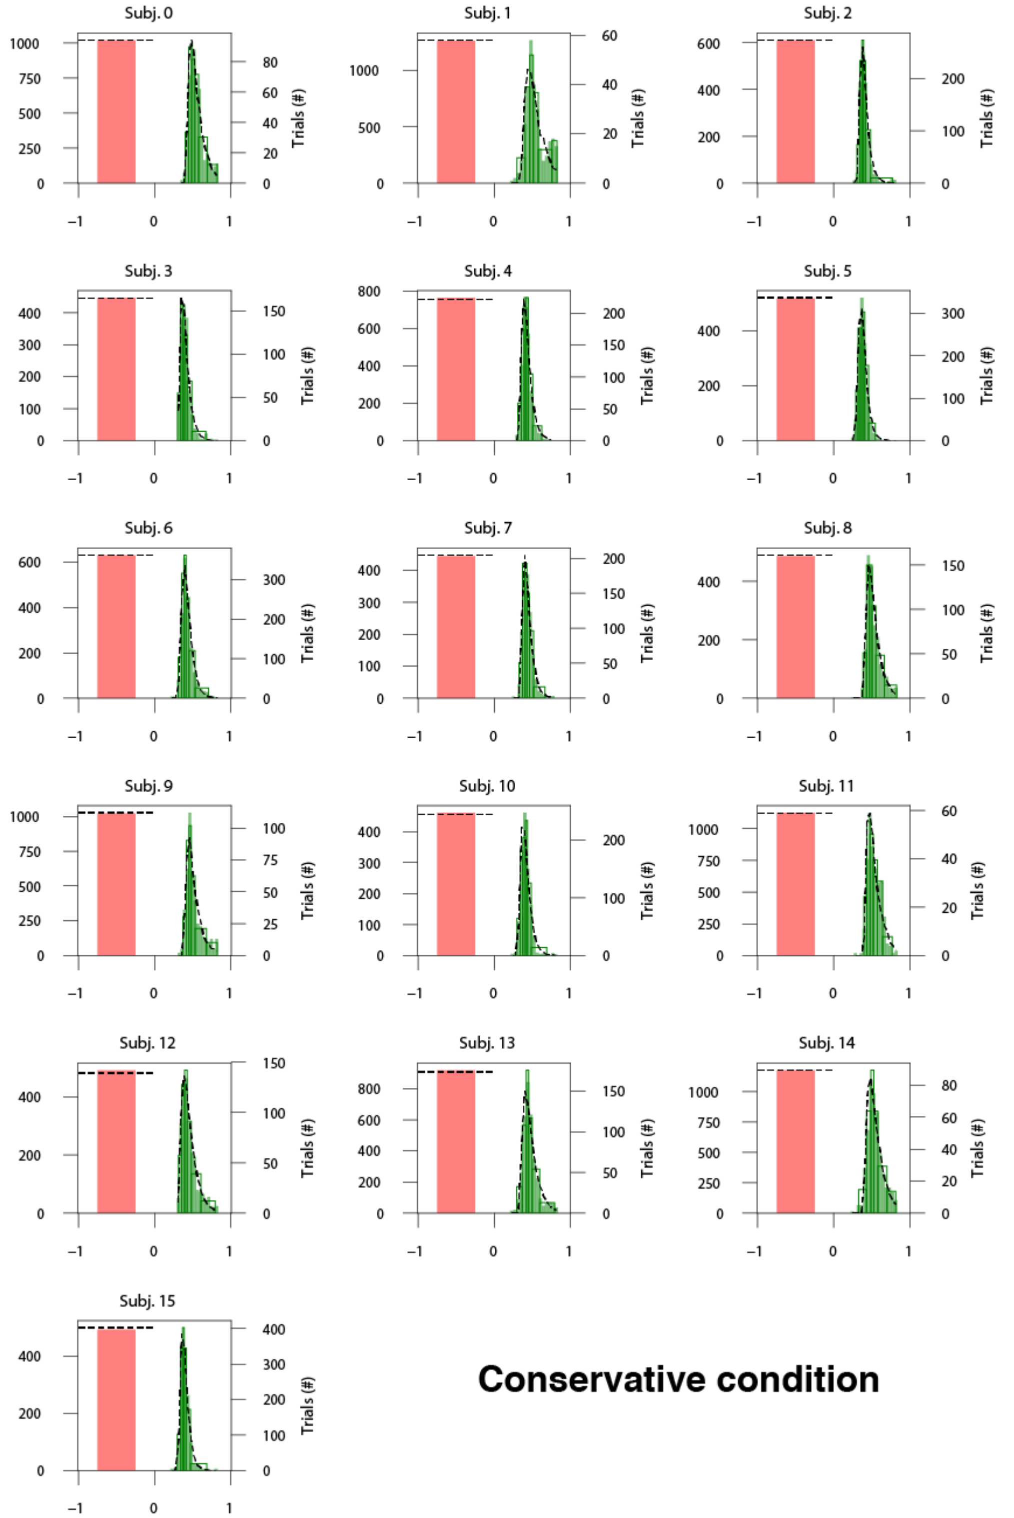

Figure S3 | Single-participant drift diffusion model fits for the drift bias model for both conditions Pink bars, number of "No" trials; Green bars, RT quantiles for "Yes" trials; dotted lines, model fits for the drift bias model.

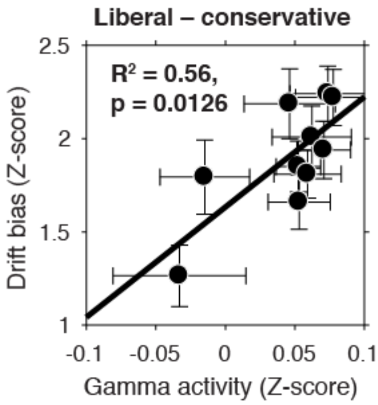

Figure S4 I Liberal – conservative gamma activity predicts corresponding drift bias increase, showing that the experimental bias manipulation enhanced gamma activity. A. Linear regression of drift bias on gamma power across excitability bins for the liberal – conservative contrast. Gamma and drift bias values were computed within participant within ten alpha suppression bins reflecting neural excitability, then Z-scored, and finally the conservative – liberal difference across bins was taken.