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

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

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

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

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

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

[5] 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

Corresponding author 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 an internal reference against which sensory evidence is compared. Here, we show that individuals are able to strategically shift this internal reference depending on current task demands by changing the rate of sensory evidence accumulation in visual cortex. Participants performed a target detection task during EEG recordings. We experimentally manipulated participants' decision criterion for reporting target-present using different stimulus-response reward contingencies, inducing liberal and conservative biases in different conditions. Drift diffusion modeling revealed that a strategic liberal bias shift specifically biased sensory evidence accumulation towards target-present choices. In visual cortex, the liberal bias suppressed pre-stimulus 8—12 Hz (alpha) power, which in turn mediated output activity of visual cortex, as expressed in 59—100 Hz (gamma) power. 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 another choice alternative by environmental factors, perhaps as a result of task instructions and/or stimulus-response reward contingencies (White & Poldrack, 2014). The ability to willfully control decision bias could potentially enable the behavioral flexibility required to survive in an ever-changing and uncertain environment. But despite its 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, how exactly bias affects decision making and how it is reflected in neural activity remains unknown.

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. After one of the bounds is reached, the corresponding decision can subsequently either be actively reported, for example by means of a button press indicating a detected signal, or it could remain without behavioral report when no signal is detected (Ratcliff, Huang-Pollock, & McKoon, 2016). 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 by shifting the 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 neural mechanism.

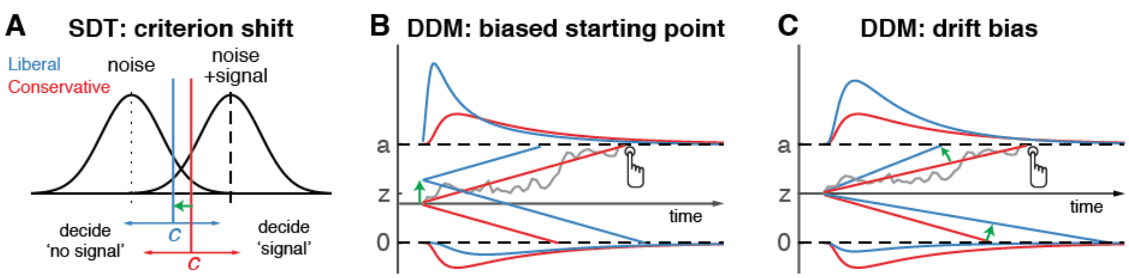

Figure 1 | Theoretical accounts of decision bias. A. Signal-detection-theoretic account of decision bias. Signal and noise+signal distributions are plotted as a function of the strength of internal sensory evidence. The decision point (or criterion) that determines whether to decide signal presence or absence is plotted as a vertical criterion line *c*, reflecting the degree of decision bias. *c* can be shifted left- or rightwards to denote a more liberal or conservative bias, respectively (green arrow indicates a shift towards more liberal). **B, C:** Drift diffusion model (DDM) account of decision bias, in which decisions are modelled in terms of a set of parameters that describe a dynamic process of sensory evidence accumulation towards one of two decision bounds. When sensory input is presented, evidence starts to accumulate (drift) over time after initialization at the starting point *z*. A decision is made when the accumulated evidence either crosses decision boundary *a* (signal presence) or decision boundary *0* (no signal). After a boundary is reached, the corresponding decision can be either actively reported by a button press (e.g. for signal-present decisions), or remain implicit, without a response (for signal-absent decisions). The DDM can capture decision bias through a shift of the starting point of the evidence accumulation process (panel B) or through a shift in bias in the rate of evidence accumulation towards the different choices (panel C). These mechanisms are dissociable through their differential effect on the shape of the reaction time (RT) distributions, as indicated by the curves above and below the graphs for target-present and target-absent decisions, respectively. Panels B. and C. are modified and reproduced with permission from Urai, de Gee, & Donner (2018).

Regarding the neural underpinnings of decision bias, there have been a number of reports about a correlational relationship between cortical population activity measured with EEG and decision bias. For example, spontaneous trial-to-trial variations in pre-stimulus 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). Alpha oscillations, in turn, 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 links between pre-stimulus alpha suppression, sensory information gating 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 (a) instruction and (b) asymmetries in stimulus-response reward contingencies during a visual target 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 pre-stimulus alpha as well as the output activity of visual cortex, as reflected in gamma power modulation. Critically, we show that gamma activity accurately predicts the strength of the evidence accumulation bias within participants, 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

Manipulation of decision bias affects 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 embedded in the continuous stream (total sequence duration 1 second). 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 a trial within the continuous stream of textures was not explicitly cued, the similar distribution of reaction times in target and nontarget trials suggests that participants used the temporal structure of the task even when no target appeared (Figure 2—figure supplement 1A). Consistent and significant EEG power modulations after trial onset (even for nontarget trials) further confirm that subjects registered trial onsets in the absence of an explicit cue, plausibly using the onset of a fixed order texture sequence as an implicit cue (Figure 2—figure supplement 1B).

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. A trial was considered a target present response when a button press occurred before the fixed-order sequence ended (i.e. within 0.84 s after onset of the fifth texture containing the (non)target, see Figure 2A). We provided auditory feedback and applied monetary penalties following missed targets in the liberal condition and following false alarms in the conservative condition (Figure 2A; see Methods for details). The median number of trials for each SDT category across participants was 1206 hits, 65 false alarms, 186 misses and 355 correct rejections in the liberal condition, and 980 hits, 12 false alarms, 419 misses and 492 correct rejections in the conservative condition.

Participants reliably adopted the intended decision bias shift across the two conditions, as shown by both the hit rate and the false alarm rate going down in tandem as a consequence of a more conservative bias (Figure 2B). The difference between hit rate and false alarm rate was not significantly modulated by the experimental bias manipulations ($p = 0.81$, two-sided permutation test, 10,000 permutations, see Figure 2B). However, target detection performance computed using standard SDT *d'* (perceptual sensitivity, reflecting the distance between the noise and signal distributions in Figure 1A)(Green & Swets, 1966) was slightly higher during conservative (liberal: $d' = 2.0$ (s.d. 0.90), versus conservative: $d' = 2.31$ (s.d. 0.82), $p = 0.0002$, see Figure 2C, left bars). We quantified decision bias using the standard SDT criterion measure *c*, in which positive and negative values reflect conservative and liberal biases, respectively (see the blue and red vertical lines in Figure 1A). This uncovered a strong experimentally induced bias shift from the conservative to the liberal condition (liberal: c = – 0.13 (s.d. 0.4), versus conservative: c = 0.73 (s.d. 0.36), p = 0.0001, see Figure 2C), as well as a conservative average bias across the two conditions (c = 0.3 (s.d. 0.31), p = 0.0013).

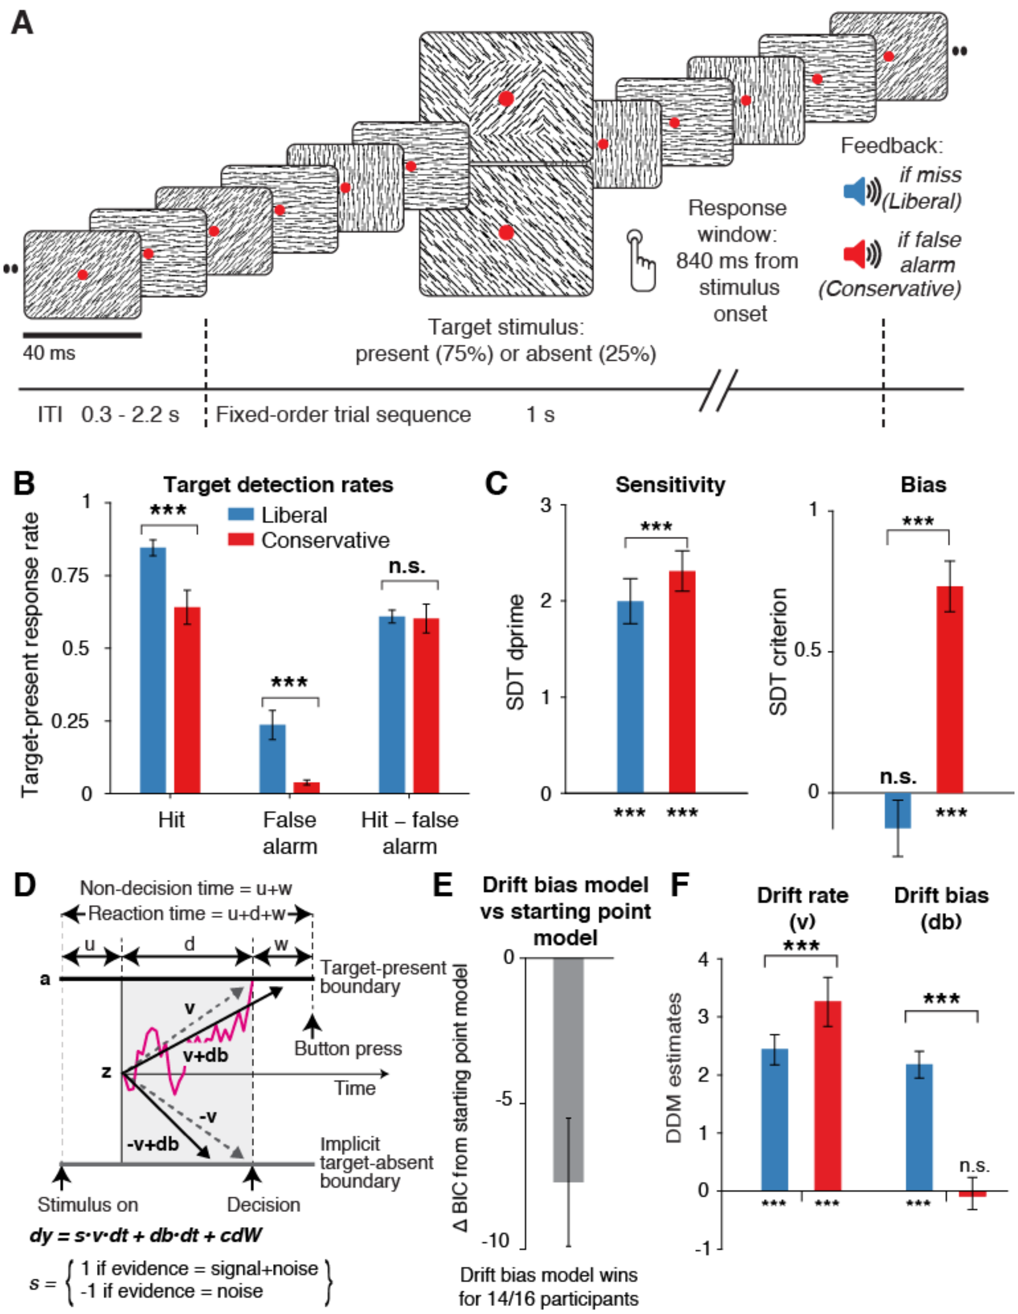

Figure 2 | Strategic decision bias shift towards liberal biases evidence accumulation. A. Schematic of the visual stimulus and task design. Participants viewed a continuous stream of full-screen diagonally, horizontally and vertically oriented textures at a presentation rate of 40 ms (25 Hz). After random 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 (target absent), or contained an orthogonal orientation-defined square (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. Miss rate is equal to 1 – hit rate since both are computed on stimulus present trials, and correct-rejection rate as 1 – false alarm rate since both are computed on stimulus absent trials, together yielding the four SDT stimulus-response categories **C.** SDT parameters for sensitivity and criterion. **D**. Schematic and simplified equation of drift diffusion model accounting for reaction time distributions for actively reported target-present and implicit target-absent decisions. Decision bias in this model can be implemented by either shifting the starting point of the evidence accumulation process (*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 (nontarget) trials); db·dt, drift bias; and cdW, Gaussian white noise (mean = 0, variance = c2·dt). **E.** 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. Panel D. is modified and reproduced with permission from (de Gee et al., 2017).

The following source data and figure supplements are available for Figure 2:

Source data 1. This csv table contains the data for Figure 2 panels B, C, E and F.

Figure supplement 1. Behavioral and neurophysiological evidence that participants were sensitive to the implicit task structure.

Figure supplement 2. Signal-detection-theoretic behavioral measures during both conditions correspond closely to drift diffusion modeling parameters.

Figure supplement 3. Single-participant drift diffusion model fits for the drift bias model for both conditions.

Because the SDT framework is static, we further investigated 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). The DDM postulates 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 behavioral report to indicate the decision (for target-present responses in our experiment), or remain implicit (i.e. without active response, for target-absent decisions in our experiment). The DDM captures the dynamic decision process by estimating parameters reflecting the rate of evidence accumulation (drift rate), the separation between the boundaries, as well as the time needed for stimulus encoding and response execution (non-decision time)(Ratcliff & McKoon, 2008). The DDM is able to estimate these parameters based on the shape of the RT distributions for actively reported (target-present) decisions along with the total number of trials in which no response occurred (i.e. implicit target-absent decisions) (Ratcliff et al., 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 drift rate 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)). The drift bias parameter is determined by estimating the contribution of an evidence-independent constant added to the drift (Figure 2D). 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). Permitting only one parameter at a time to vary freely between conditions allowed us to directly compare the models without having to penalize either model for the number of free parameters. These alternative models make different predictions about the shape of the RT distributions in combination with the response ratios: a shift in starting point results in more target-present choices particularly for short RTs, whereas a shift in drift bias grows over time, resulting in more target-present choices also for longer RTs (de Gee et al., 2017; Ratcliff & McKoon, 2008; Urai et al., 2018). The RT distributions above and below the evidence accumulation graphs in Figure 1B and 1C illustrate these different effects. 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 found that the starting point model provided a worse fit to the data than the drift bias model (starting point model, Bayesian Information Criterion (BIC) = 10287; drift bias model, BIC = 10279, Figure 2E, see Methods for details). Specifically, for 14 out of the 16 participants, the drift bias model provided a better fit than the starting point model, for 10 of which delta BIC > 6, 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 > 6 for both models for all participants). See Figure 2—figure supplement 3 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 nontargets, 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). Almost perfect correlations across participants in both conditions between DDM drift rate and SDT *d'* provided strong evidence that the drift rate parameter captures perceptual sensitivity (liberal, $r = 0.97$, $p = 1.7e^{-10}$; conservative, $r = 0.95$, $p = 1.4e^{-8}$, see Figure 2—figure supplement 2A).

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 target-absent boundary ($z = 0.24$ (s.d. 0.06), $p < 0.0001$, tested against 0.5)), indicating an overall conservative starting point across conditions (Figure 2—figure supplement 2D), in line with the overall positive SDT criterion (see Figure 2C, right panel). Strikingly, however, whereas the drift bias parameter was on average not different from zero in the conservative condition (db = –0.04 (s.d. 1.17), $p = 0.90$), drift bias was strongly positive in the liberal condition (db = 2.08 (s.d. 1.0), $p = 0.0001$; 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 target-present boundary) explained the neutral bias observed with SDT criterion (*c* around zero for liberal, see Figure 2C).

Convergent with these modelling results, drift bias was strongly anti-correlated across participants with both SDT criterion (liberal, $r = -0.83$; conservative, $r = -0.79$) and reaction times (liberal, $r = -0.66$; conservative, $r = -0.76$, all p-values < 0.005, see Figure 2—figure supplement 2B and 2C). The strong correlations between drift rate and *d'* on the one hand, and drift bias and *c* on the other, provide converging evidence that the SDT and DDM frameworks capture similar underlying mechanisms, while the DDM additionally captures 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 the liberal compared to the conservative condition ($p < 0.0001$), and non-decision time (comprising time needed for sensory encoding and motor response execution) was shorter during liberal ($p < 0.0001$) (Figure 2—figure supplement 2D). In conclusion, a drift diffusion model of choice behavior implementing a bias in sensory evidence accumulation best explained how participants adjusted to the decision bias manipulations. 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.

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

Sensory evidence accumulation in a visual target detection task presumably relies on stimulus-induced signals processed in visual cortex. Such stimulus-induced 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 condition-specific pre-stimulus baseline period. We observed a total of four distinct stimulus-induced 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, we found 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).

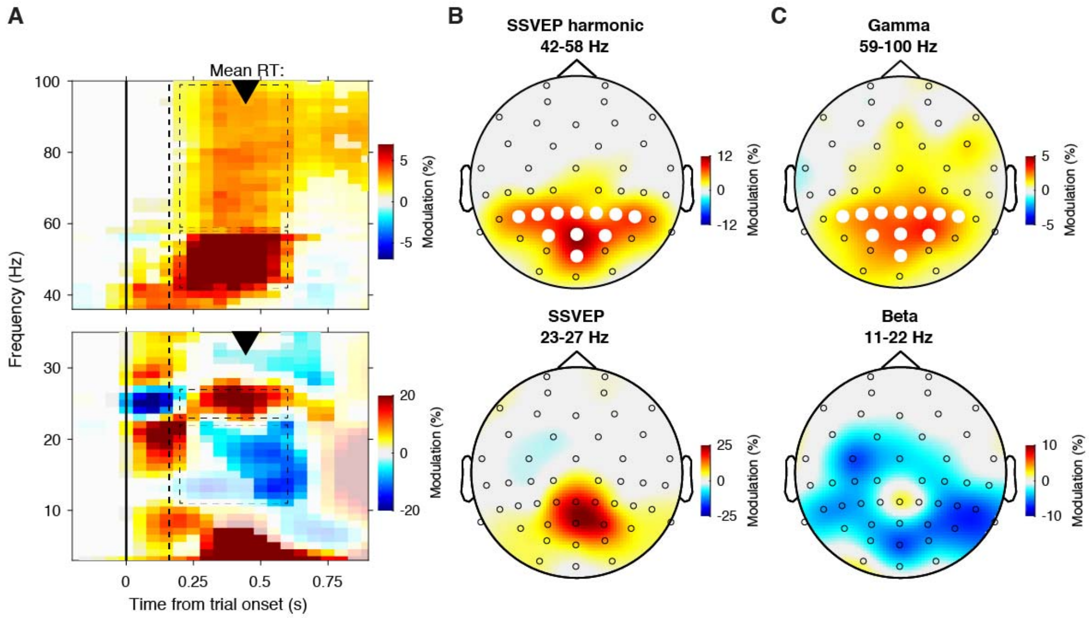

Figure 3 | Task-relevant textures induce stimulus-induced responses in visual cortex. A. Time-frequency representations of high- (top) and low-frequency (bottom) EEG power modulations with respect to the condition-specific pre-stimulus period (–0.4 to 0 s). Saturated colors indicate clusters of significant 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. **B.** Scalp maps showing topography of the steady-state visual evoked potential (SSVEP) power modulation around 25 Hz (bottom) and its harmonic from 42 – 58 Hz (top), 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 (bottom); see dashed outlines on time-frequency representations in A. White dots indicate electrodes used for the time-frequency representations in A, and which were selected for further analysis.

The following source data is available for Figure 3:

Source data 1. MATLAB .mat file containing the data used in panel A.

Figure 3 – Figure supplement 1. Liberal – conservative EEG power modulation contrast across space, time and frequency.

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 (SCD) 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 bottom-up 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-induced 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). Taken together, we observed several different stimulus-induced power modulations in posterior cortex. In the next section, we used the topographies of the high-frequency poststimulus effects in visual cortex (Figures 3B and 3C, top panels) to identify a pre-stimulus neural mechanism that could explain the observed biased evidence accumulation resulting from the experimental decision bias manipulation.

Adopting a liberal decision bias suppresses pre-stimulus alpha power

Next, we tested whether our bias manipulation affected the amplitude of pre-stimulus 8–12 Hz (alpha) oscillations in visual cortex. To this end, we examined the raw, low-frequency spectral power in the pre-stimulus interval in which a link between spontaneous alpha fluctuations and decision bias has previously been reported (0.8 to 0.2 s before trial onset) (Iemi et al., 2017). We focused this analysis on cortical regions processing visual information by selecting the electrode pooling that showed stimulus-induced high-frequency gamma power modulation (see Figures 3B and 3C). Spectral power averaged across the two conditions indeed uncovered a highly specific modulation around 10 Hz, which we confirmed to be strongest in the same electrodes that showed strong modulation in the gamma range (Figure 4A, white dots indicate electrodes showing stimulus-induced gamma modulation). Crucially, the liberal – conservative difference between conditions revealed a statistically significant cluster of suppressed frequencies precisely in the 8—12 Hz frequency range ($p < 0.05$, cluster-corrected for multiple comparisons), which again showed a posterior topography (Figure 4B). This small but highly consistent shift in the range in which alpha occurs during the liberal compared to the conservative condition is depicted in Figure 4C. Taken together, these findings show that a strategic liberal bias shift suppresses pre-stimulus 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 pre-stimulus alpha power in visual cortex, plausibly acting to bias sensory evidence accumulation towards the response alternative that maximizes rewards.

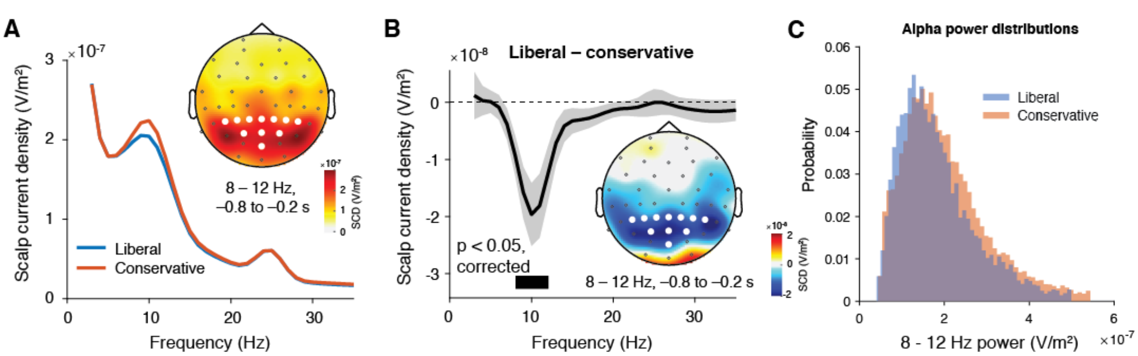

Figure 4 | Adopting a liberal decision bias suppresses pre-stimulus alpha power. A. Low-frequency raw power spectra of pre-stimulus neural activity for both conditions based on the electrodes that show large poststimulus power modulations in Figures 3B and 3C (top panels). Inset, scalp map of raw pre-stimulus EEG alpha power (8 — 12 Hz neural activity between 0.8 and 0.2 s before trial onset), pooled over conditions. White symbols indicate visual cortical electrodes used for the raw power spectra in A. and B. **B.** Liberal – conservative raw 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). Inset, corresponding liberal – conservative scalp map of the pre-stimulus raw power difference between conditions. SCD, scalp current density. **C.** Probability density distributions of single trial alpha power values for both conditions, averaged across participants.

The following source data is available for Figure 4:

Source data 1. MATLAB .mat file containing the data used in panel B.

Pre-stimulus alpha power mediates cortical gamma responses

How could suppression of pre-stimulus 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 response gain model coined by Rajagovindan and Ding (2011). This model postulates that the relationship between the total synaptic input activity that a neuronal ensemble receives and the total output activity it produces is characterized by a sigmoidal function (red line in Figure 5A) – a notion that is biologically plausible (Destexhe, Rudolph, Fellous, & Sejnowski, 2001; Freeman, 1979). In this model, sensory input (i.e. due to sensory stimulation) and ongoing fluctuations in endogenously generated (i.e. not sensory-related) neural activity together comprise the synaptic input into visual cortex. In our experiment, the sensory input into visual cortex can be assumed to be identical across trials, because the same sensory stimulus was presented in each trial (see Figure 2A). The endogenous input, in contrast, varies from trial to trial reflecting fluctuations in top-down cognitive processes such as attention, and is assumed to be reflected in alpha power. Given the combined constant sensory and variable endogenous input in each trial (see horizontal axis in Figure 5A), the strength of the output responses of visual cortex are largely determined by the trial-to-trial variation caused by endogenous activity (see vertical axis in Figure 5A). Furthermore, the sigmoidal shape of the input-output function results in an effective range in the center of the function's input side which yields the strongest stimulus-induced output responses since the sigmoid curve there is steepest. Mathematically, the effect of endogenous input on stimulus-induced output responses (see marked interval in Figure 5A) can be expressed as the first order derivative or slope of the sigmoid in Figure 5A, which is referred to as the response gain by Rajagovindan and Ding (2011). This derivative is plotted in Figure 5B (red line) across levels of pre-stimulus alpha power, predicting an inverted-U shaped relationship between alpha and response gain in visual cortex.

Regarding our experimental conditions, the model not only predicts that the suppression of pre-stimulus alpha observed in the liberal condition reflects a shift in the operational range of alpha (see Figure 4C), but also that it increases the maximum output of visual cortex (a shift from the red to the blue line in figure 5A). Thus, as the operational range of alpha shifts leftwards from conservative to liberal, the upper asymptote in Figure 5A moves upwards such that the total maximum output activity increases. This in turn affects the inverted-U-shaped relationship between alpha and gain in visual cortex (blue line in Figure 5B), leading to a steeper response curve in the liberal condition resembling a Gaussian (bell-shaped) function.

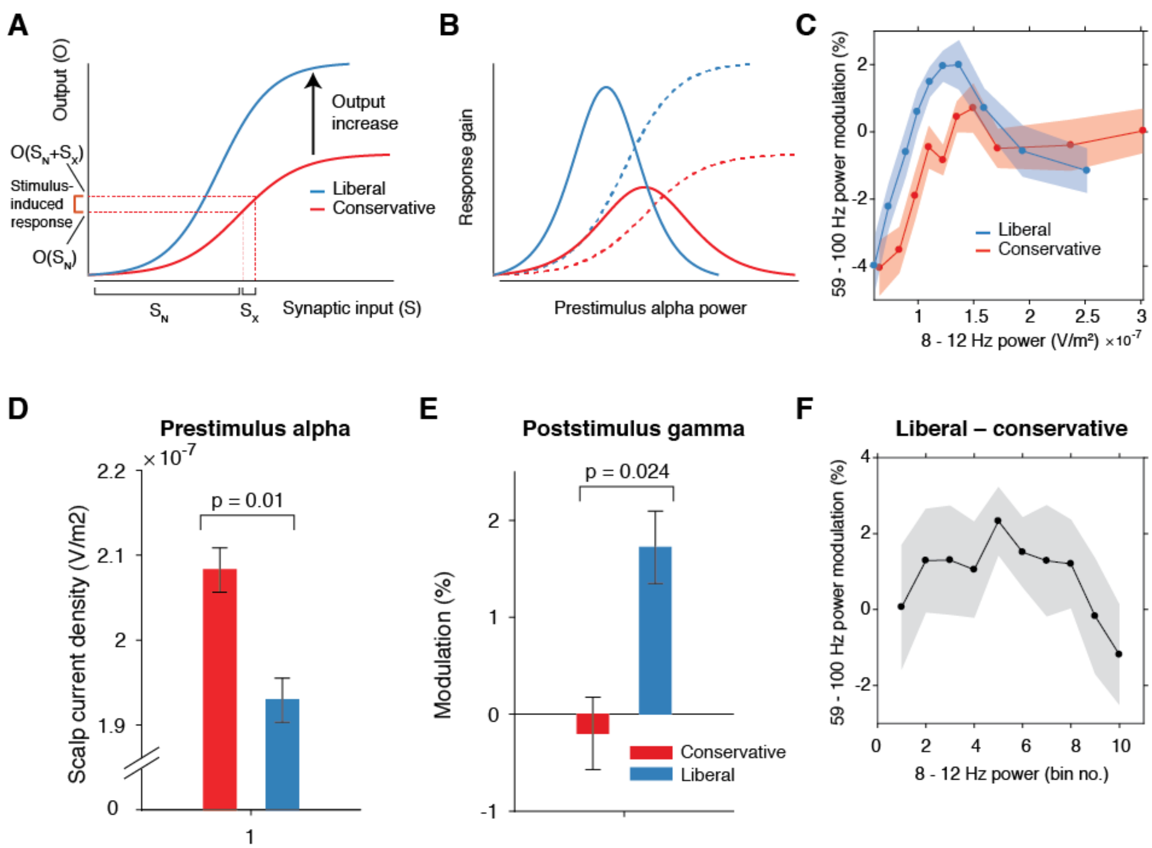

Figure 5 | Pre-stimulus alpha power mediates cortical gamma responses. A. Theoretical response gain model describing the transformation of stimulus-induced and endogenous input activity (denoted by S_x and S_N respectively) to the total output activity (denoted by $O(S_x + S_N)$) in visual cortex by a sigmoidal function. Different operational alpha ranges are associated with input-output functions with different slopes due to corresponding changes in the total output. **B.** Alpha-mediated output responses (solid lines) are formalized as the first derivative (slope) of the sigmoidal functions (dotted lines), resulting in inverted-U (Gaussian) shaped relationships between alpha and gamma, involving stronger response gain in the liberal than in the conservative condition **C.** Corresponding empirical data showing gamma modulation (same percent signal change units as in Figure 3) as a function of alpha bin. The location on the x-axis of each alpha bin was taken as the median alpha of the trials assigned to each bin and averaged across subjects. **D-F.** Model prediction tests. **D.** Raw pre-stimulus alpha power for both conditions, averaged across subjects. **E.** Post-stimulus gamma power modulation for both conditions averaged across the two middle alpha bins (5 and 6) in panel C. **F.** Liberal – conservative difference between the response gain curves shown in panel C, centered on alpha bin. Error bars, within-subject SEM across participants (N = 14).

The following source data is available for Figure 5:

Source data 1. SPSS .sav file containing the data used in panels C, E, and F.

To investigate sensory response gain across different alpha levels in our data, we used the post-stimulus gamma activity (see Figure 3) as a proxy for alpha-mediated output gain in visual cortex (Bastos et al., 2015; Michalareas et al., 2016; Ni et al., 2016; Popov et al., 2017; van Kerkoerle et al., 2014). We exploited the large number of trials per participant per condition (range 543 to 1391 trials) by sorting each participant's trials into ten equal-sized bins ranging from weak to strong alpha, separately for the two conditions. We then calculated the average gamma power modulation within each alpha bin and finally plotted the participant-averaged gamma across alpha bins in Figure 5C (see Methods for details). This indeed revealed an inverted-U shaped relationship between alpha and gamma, with a steeper curve for the liberal condition.

To assess the model's ability to explain the data, we statistically tested three predictions derived from the model. First, the model predicts overall lower average pre-stimulus alpha power for liberal than for conservative due to the shift in the operational range of alpha. This was confirmed in Figure 5D ($p = 0.01$, permutation test, see also Figures 4B and 4C). Second, the model predicts a stronger gamma response for liberal than for conservative around the peak of the gain curve (the center of the effective alpha range, see Figure 5B), which we indeed observed ($p = 0.024$, permutation test on the average of the middle two alpha bins)(Figure 5E).

Finally, the model predicts that the difference between the gain curves (when they are aligned over their effective ranges on the x-axis using alpha bin number, as shown in Figure 5 – figure supplement 1A, also resembles a Gaussian curve (Figure 5 – figure supplement 1B). Consistent with this prediction, we observed an interaction effect between condition (liberal, conservative) and bin number (1-10) using a standard Gaussian contrast in a 2-way repeated measures ANOVA ($F(1,13) = 4.6$, $p = 0.051$, partial $\eta^2 = 0.26$). Figure 5F illustrates this finding by showing the difference between the two curves in Figure 5C as a function of alpha bin number (see Figure 5 – figure supplement 1C for the curves of both conditions as a function of alpha bin number). Subsequent separate tests for each condition indeed confirmed a significant U-shaped relationship between alpha and gamma in the liberal condition with a large effect size ($F(1,13) = 7.7$, $p = 0.016$, partial $\eta^2 = 0.37$), but no significant effect in the conservative condition with only a small effect size ($F(1,13) = 1.7$, $p = 0.22$, partial $\eta^2 = 0.12$, one-way repeated measures ANOVA's with factor alpha bin, Gaussian contrast).

Taken together, these findings suggest that the alpha suppression observed in the liberal compared to the conservative condition boosted stimulus-induced activity in the liberal condition, which in turn might have indiscriminately biased sensory evidence accumulation towards the target-present decision boundary. In the next section, we investigate a direct link between drift bias and stimulus-induced activity 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 target-present responses (drift bias in the DDM), while neurally it suppresses pre-stimulus alpha while enhancing poststimulus gamma responses. In a final analysis, we asked whether alpha-binned gamma modulation is directly related to a stronger drift bias. To this end, we again applied the drift bias DDM to the behavioral data of each participant, but now freed the drift bias parameter not only for the two conditions, but also for the ten alpha bins for which we calculated gamma modulation (see Figure 5C). We directly tested the correspondence between DDM drift bias and gamma modulation using repeated measures correlation (Bakdash and Marusich, (2017), which takes all repeated observations across participants into account while controlling for non-independence of observations collected within each participant (see Methods for details). Gamma modulation was indeed correlated with drift bias in both conditions (liberal, $r(125) = 0.49$, $p = 5e-09$; conservative, $r(125) = 0.38$, $p = 9e-06$). We tested the robustness of these correlations by excluding the data points that contributed most to the correlations (as determined with Cook's distance) and obtained qualitatively similar results, indicating these correlations were not driven by outliers (Figure 6, see Methods for details). As a final control, we also performed this analysis for the SSVEP (23-27 Hz) power modulation (see Figure 3B, bottom) and found a similar inverted-U shaped relationship between alpha and the SSVEP for both conditions (Figure 6 – figure supplement 1A), but no correlation with drift bias (Figure 6 – figure supplement 1B). This suggests that the SSVEP is similarly coupled to alpha as the stimulus-induced gamma, but is unaffected by the experimental conditions and not predictive of decision bias shifts. Taken together, these results suggest that gamma modulation underlies biased sensory evidence accumulation.

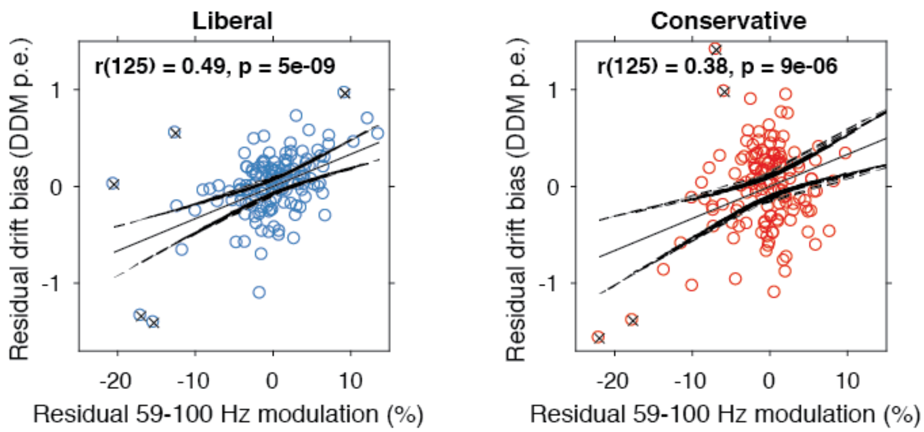

Figure 6 | Alpha-binned gamma modulation correlates with evidence accumulation bias. Repeated measure correlation between gamma modulation and drift bias for the two conditions. Each circle represents a participant's gamma modulation within one alpha bin. Drift bias and gamma modulation scalars were residualized by removing the average within each participant and condition, thereby removing the specific range in which the participants values operated. Crosses indicate data points that were most influential for the correlation, identified using Cook's distance. Correlations remained qualitatively unchanged when these data points were excluded (liberal, r(120) = 0.43, p = 8e-07; conservative, r(121) = 0.32, p = 0.0003) Error bars, 95% confidence intervals after averaging across participants.

The following source data and figure supplements are available for Figure 6:

Source data 1. MATLAB .mat file containing the data used.

Figure supplement 1. Alpha-binned post-stimulus SSVEP modulation.

Discussion

Traditionally, bias has been conceptualized in SDT as a criterion threshold that is positioned at an arbitrary location between noise and signal-embedded-in-noise distributions of sensory evidence strengths. The ability to strategically shift decision bias in order to flexibly adapt to stimulus-response reward contingencies in the environment presumably increases chances of survival, but to date such strategic bias shifts as well as their neural underpinnings have not been demonstrated. Here, we used a DDM drift bias model to show 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, and enhancing gamma activity by increasing response gain.

Although previous studies have shown correlations between suppression of pre-stimulus alpha (8—12 Hz) power and a liberal decision bias during spontaneous fluctuations in alpha activity (Iemi et al., 2017; Limbach & Corballis, 2016), these studies have not established the effect of experimentally induced bias shifts within person. In the current study, by experimentally manipulating stimulus-response reward contingencies we show for the first time that pre-stimulus alpha can be actively modulated by a participant to achieve changes in decision bias. Further, we show that alpha suppression in turn modulates gamma activity, in part by increasing the gain of cortical responses. Critically, gamma activity accurately predicted the strength of the drift bias parameter in the DDM drift bias model, thereby linking our behavioral and neural findings directly. 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 a strategic liberal bias shift, as reflected in visual cortical alpha suppression (left panel). This increased excitability translates into stronger gamma-band responses following stimulus onset (right panel, top). These increased gamma responses finally bias evidence accumulation towards the target-present decision boundary during a liberal state, resulting in more target-present responses, whereas target-absent 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.

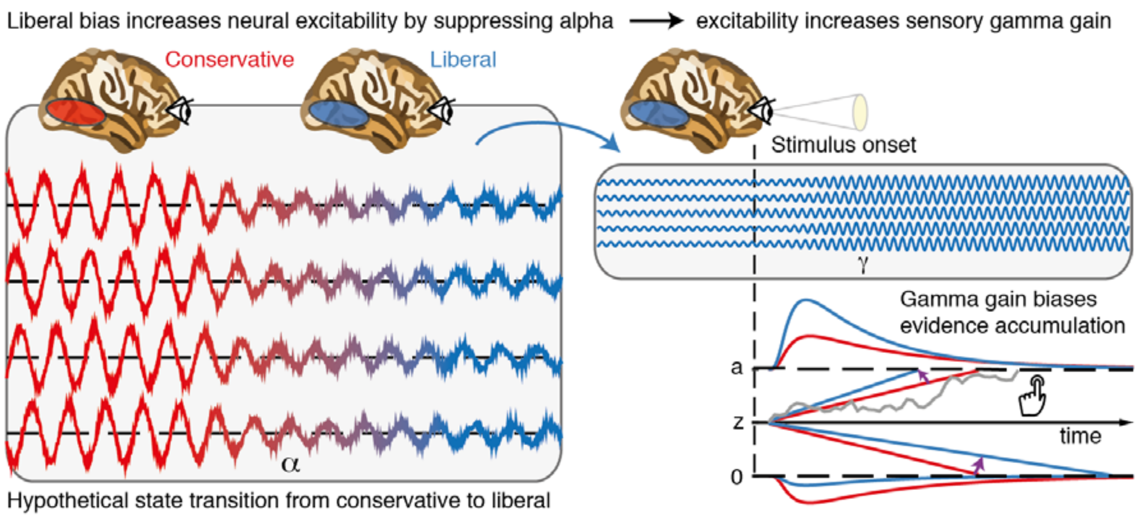

Figure 7 | Illustrative graphical depiction of the excitability state transition from conservative to liberal, and subsequent biased evidence accumulation under a liberal bias. The left panel shows the transition from a conservative to a liberal condition block. The experimental induction of a liberal 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 coeruleus (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, pre-stimulus 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 and/or reward) 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 target-present responses (Iemi et al., 2017).

We note that although the gain model is consistent with our data as well as the data on which the model was conceived (see Rajagovindan & Ding, 2011), we do not provide a plausible mechanism that could bring about the steepening in the U-curved function observed in Figures 5C and 5F. Although Rajagovindan and Ding report a simulation in their paper suggesting that increased excitability could indeed cause increased gain, this shift could in principle either be caused by the alpha suppression itself, by the same signal that causes alpha suppression, or it could originate from an additional top-down signal from frontal brain regions. To investigate this latter possibility, we performed a control analysis contrasting the conditions simultaneously across space, time and frequency to test whether any frontal brain region shows differences between conditions (see Figure 3 – figure supplement 1). We did not find any such regions, even when using a less stringent test by omitting the required correction for multiple comparisons. Thus, how exactly the gain enhancement comes about remains an open question that should be addressed in future research.

Whereas we report a unique link between alpha-mediated gamma modulation and decision bias through the gain model, several previous studies have reported a link between alpha and objective performance instead of bias, 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.

Further, one could ask whether the observed change in cortical excitability may reflect a change in target 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 we rather 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 decision bias, as opposed to detection sensitivity. Convergently, other studies also report a link between pre-stimulus low-frequency EEG activity and subjective perception, but not objective task performance (Benwell et al., 2017; Iemi & Busch, 2017).

In summary, our results suggest that stimulus-induced 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 brain stimulation (Zaehle, Rach, & Herrmann, 2010) could further establish the underlying neural mechanisms involved in decision bias.

In the end, although one may be unaware, every decision we make is influenced by biases that operate on one's noisy evidence accumulation process. Understanding how these biases affect our decisions is crucial to enable us to control or invoke them adaptively (Pleskac, Cesario, & Johnson, 2017). Pinpointing the neural mechanisms underlying bias in the current elementary perceptual task may pave the way for understanding how more abstract and high-level decisions are modulated by decision bias (Tversky & Kahneman, 1974).

Acknowledgments

The authors thank Timothy J. Pleskac for discussion.

Declaration of Interests

The authors declare no competing interests.

Data and code sharing

The data analyzed in this study are publicly available on Figshare (Kloosterman et al., 2018). Analysis scripts are publicly available on Github (https://github.com/nkloost1/critEEG).

Source data and figure supplements

The following source data and figure supplements are included in this article:

Figure 2 – Figure supplements 1, 2 and 3.

Figure 2 – Source data 1. (source_Figure2.csv)

Figure 3 – Source data 1. (source_Figure3.mat.zip)

Figure 3 – Figure supplement 1

Figure 4 – Source data 1. (source_Figure4.mat.zip)

Figure 5 – Source data 1. (source_Figure5.sav.zip)

Figure 6 – Source data 1. (source_Figure6.mat.zip)

Figure 6 – Figure supplement 1.

Figure supplement legends

Figure 2—figure supplement 1 | 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 signaled. Error bars, SEM. **B.** Time-frequency representations of low-frequency EEG power modulations with respect to the pre-stimulus 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 2—figure supplement 2 | Signal-detection-theoretic (SDT) behavioral measures during both conditions correspond closely to drift diffusion modeling (DDM) parameters. A. Across-participant Pearson correlation between d’ and drift rate for the two conditions. Each dot represents a participant. **B.** As A. but for correlation between criterion and DDM drift bias. The correlation is negative due to a lower criterion reflecting a stronger liberal bias. **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 for hits and drift bias. **D.** Parameter estimates in the drift bias DDM not related to evidence accumulation (drift rate). ***$p < 0.001$; n.s., not significant.

Figure 2—figure supplement 3 | Single-participant drift diffusion model fits for the drift bias model for both conditions. Pink bars, number of implicit target-absent choices; Green bars, RT distribution quantiles for target-present choices; dotted lines, model fits for the drift bias model.

Figure 3 – figure supplement 1 | Liberal – conservative contrast of EEG power modulations across space, time and frequency. The two conditions were contrasted across space-time-frequency bins using paired t-tests performed at each bin. Single bins were subsequently thresholded at $p < 0.05$ and clusters of contiguous bins were determined. Cluster significance was assessed using a cluster-based permutation procedure (1000 permutations). For visualization purposes, we integrated (using the matlab trapz function) power modulation in the time-frequency representations (TFR’s, left panels) across the highlighted electrodes in the topographies (right panels). For the topographies, modulation was integrated across the saturated time-frequency bins in the TFR’s. Saturated colors indicate bins that are part of clusters. p-values above the topographies indicate

cluster significance level tested across participants; N = 14). **A.** We found one significant cluster of positive sign (p = 0.005) located in the most occipital electrodes (I1, I2, and Iz, extending towards parietal regions) spanning the complete high-frequency range (> 20 Hz), reflecting enhanced broad-band gamma activity in the liberal condition. Note that our selected electrode pooling (see Figure 2) did not include these electrodes. **B.** Further, we observed one marginally significant negative cluster (p = 0.11) comprising the pre-stimulus alpha-suppression in the liberal condition (as reported in the manuscript) that was connected across time with alpha-beta band activity over motor cortex around the time of report (~0.5 s). Note that all participants responded with their right hand, yielding stronger left-lateralized motor-related activity. **C.** Finally, we observed a transient positive cluster around 10 Hz from 0.4 s post-trial onset with a spatial topography similar to the cluster in A, which was not significant (p = 0.35). This cluster possibly reflects either a stronger event-related potential, or stronger transient enhancement of theta oscillations (4–8 Hz) in the liberal condition around the time of the response. Topographies in all panels appear quite similar due to the strong modulation of the cluster depicted in panel A. However, the cortical locations of clusters in each panel are indicated by the thick black dots that indicate electrodes that are part of the cluster. Taken together, we observe no strong evidence for a frontal cluster that could potentially underlie the steepened inverse-U shape during the liberal condition observed in visual cortex.

Figure 5 – figure supplement 1 | Gain model predictions and corresponding empirical data plotted as a function of pre-stimulus alpha bin number. A. Model predictions for both conditions. The gain curve for the liberal condition is steeper than for the conservative condition. Binning trials based on alpha within each condition directly maps the peaks of the gain curves onto each other. **B.** Model prediction for liberal – conservative as a function of alpha bin number. The difference gain curve between the two conditions is again an inverted-U shaped function. **C.** Corresponding empirical data. The plot is identical to Figure 5C, except that the bin number is plotted instead of the actual alpha power for each condition.

Figure 6 – figure supplement 1 | Alpha-binned post-stimulus SSVEP modulation. A. Inverted-U shaped relationship between alpha and SSVEP modulation, computed as the percent signal change 23 – 27 Hz power modulation with respect to the pre-stimulus baseline. **B.** Correlations between SSVEP modulation and drift bias for both conditions. These non-significant correlations are overall weaker than for gamma (see Figure 6).

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

Bakdash, J. Z., & Marusich, L. R. (2017). Repeated Measures Correlation. *Frontiers in Psychology*, *8*, 491. http://doi.org/10.3389/fpsyg.2017.00456

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

Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). The spatiotemporal profile of cortical processing leading up to visual perception. *Journal of Vision*, *8*(1), 12–12. http://doi.org/10.1167/8.1.12

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), 617–624. 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

Kiani, R., Hanks, T. D., & Shadlen, M. N. (2008). Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment. *Journal of Neuroscience*, *28*(12), 3017–3029. http://doi.org/10.1523/JNEUROSCI.4761-07.2008

Kloosterman, N. A., de Gee, J. W., Werkle-Bergner, M., Lindenberger, U., Garrett, D. D., & Fahrenfort, J. J. (2018). Data from: Humans strategically shift decision bias by flexibly adjusting sensory evidence accumulation in visual cortex. http://doi.org/https://doi.org/10.6084/m9.figshare.6142940

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

Lamme, V. A. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *Journal of Neuroscience*, *15*(2), 1605–1615. http://doi.org/10.1523/JNEUROSCI.15-02-01605.1995

Lamme, V. A. F., Zipser, K., & Spekreijse, H. (2006). Masking Interrupts Figure-Ground Signals in V1. *Dx.Doi.org*, *14*(7), 1044–1053. http://doi.org/10.1162/089892902320474490

Limbach, K., & Corballis, P. M. (2016). Pre-stimulus 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: Pre-stimulus 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 pre-stimulus 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. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. http://doi.org/10.1037//0033-295X.85.2.59

Ratcliff, R. (2006). Modeling response signal and response time data. *Cognitive Psychology*, *53*(3), 195–237. http://doi.org/10.1016/j.cogpsych.2005.10.002

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). Pre-stimulus 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.

Supèr, H., Spekreijse, H., letters, V. L. N., 2003. (2003). Figure–ground activity in primary visual cortex (V1) of the monkey matches the speed of behavioral response. *Elsevier*

, *344*(2), 75–78. http://doi.org/10.1016/S0304-3940(03)00360-4

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 females, mean age 24.1 years, ± 1.64) took part in the experiment, either for financial compensation (EUR 10, - per hour) or in partial fulfillment of first year psychology 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 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.

Regarding sample size, our experiment consisted of 16 biological replications (participants) and either two (one participant) or three (fifteen participants) technical replications (i.e. experimental sessions). The sample size was determined based on two criteria: 1) obtaining large amounts of data per participant (thousands of trials), which is necessary to perform robust drift diffusion modelling of choice behavior and obtain reliable EEG spectral power estimates for each of the ten bins of trials that were created within participants, and 2) obtaining data from a sufficient number of participants to leverage across-subject variability in correlational analyses. Thus, we emphasized obtaining many data points per participant relative to obtaining many participants, while still preserving the ability to perform correlations across participants.

All participants were included in the signal-detection-theoretical and drift diffusion modeling analyses (Figure 2). One participant was excluded from the pre-stimulus alpha analysis (Figures 3 and 4) due to excessive noise (EEG power spectrum opposite of 1/frequency). One further participant was excluded from the gamma analyses (Figures 4, 5 and 6) because the liberal-conservative difference in gamma power in this participant was > 3 standard deviations away from the other participants.

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) (see Figure 2A). 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 and actively report targets by pressing a button using their right 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 of EUR 10, - per hour 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.

Participants performed six blocks per session lasting ca. nine minutes each. 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. 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. There were short breaks between blocks, in which participants indicated when they were ready to begin the next block.

Behavioral analysis We calculated each participant's 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 hit-rate is the proportion target-present responses of all target-present trials, false alarm (FA)-rate is the proportion target-present responses of all target-absent trials, and 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 (RTs) were measured as the duration between target onset and button press.

Drift diffusion modeling of choice behavior In order to be detected, the 40 ms-duration figure-ground targets used in our study undergo a process in visual cortex called figure-ground segregation. This process has been well characterized in man and monkey (Fahrenfort, Scholte, & Lamme, 2008; Lamme, 1995; Lamme, Zipser, & Spekreijse, 2006; Supèr, Spekreijse, letters, 2003, 2003), and results from recurrent processing to extract the surface region in visual cortex. Figure-ground segregation is known to extend far beyond the mere presentation time of the stimulus, thus providing a plausible neural basis for the evidence accumulation process. Further, a central assumption of the drift diffusion model is that the process of evidence accumulation is gradual, independent of whether sensory input is momentary. Indeed, the DDM was initially developed to explain reaction time distributions during memory retrieval, in which evidence accumulation must occur through retrieval of a memory trace within the brain, in the complete absence of external stimulus at the time of the decision (Ratcliff, 1978). Our observed RT distributions show the typical features that occur across many different types of decision and memory tasks, which the DDM is so well able to capture, including a sharp leading edge and a long tail of the distributions (see Figure 2-supplement 3). The success of the DDM in fitting these data is consistent with previous work (e.g. Ratcliff (2006)) and might reflect the fact that observers modulate the underlying components of the decision process also when they do not control the stimulus duration (Kiani, Hanks, & Shadlen, 2008).

We fitted the drift diffusion model to our behavioral 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). The RT distributions for target-present 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 target-absent responses contributed to *G* square. Fitting the model to RT distributions for target-present and target-absent 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 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 condition is problematic for data with no explicit target-absent responses (data not shown). Thus, to test whether shifts in drift bias or starting point underlie 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). Additionally, 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.

Distinguishing DDM drift bias and drift rate In our task, only target-present responses were coupled to a behavioral response (button-press), so we could measure reaction times only for these responses, whereas reaction times for target-absent responses remained implicit. Thus, in our fitting procedure, the RT distributions for target-present 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 target-absent responses contributed to G square. It has been shown that such a diffusion model with an implicit (no response) boundary can be fit to data with almost the same accuracy as fitting the two-choice model to two-choice data (Ratcliff et al., 2016). In a diffusion model with an implicit (no response) boundary, both an increase in drift rate and drift criterion would predict faster target-present responses. However, the key distinction is that an increase in drift additionally predicts more correct responses (for both target-present and target-absent responses), and an increase in drift criterion shifts the relative fraction of target-present and target-absent responses (decision bias). Because a single bin containing the number of target-absent responses contributed to G square, our fitting procedure can distinguish between decision bias versus drift rate.

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 due to blinks (typically one or two) were 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 saccadic artefact detection algorithm from Hassler et al. (2011). This yielded a total number of channels submitted to ICA of 48 + 2 = 50. 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 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 (Figure 3). We used as a baseline the mean EEG power in the interval 0.4 to 0 s before trial onset, computed separately for each condition. If this interval was not completely present in the trial due to preceding events (see Trial extraction), this period was shortened accordingly. We normalized the data by subtracting the baseline from each time-frequency bin and dividing this difference by the baseline (x 100 %). For the analysis of raw pre-stimulus power modulations, no baseline correction was applied on the raw scalp current density values. 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 raw EEG power change To test at which frequencies raw EEG power differed for the liberal and conservative conditions, we averaged raw power from 0.8 s up to 0.2 s before trial onset (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)). Then, we took the liberal – conservative difference at each frequency bin and statistically tested whether and at which frequency bins this signal differed from zero (Figure 4C) (see Statistical comparisons).

Response gain model test To test the predictions 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 poststimulus activity, also see Iemi et al. (2017)), yielding a single scalar alpha power value per trial. 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-trial alpha values for each participant and condition in ascending order and assigned them to ten bins of equal size, ranging from weakest to strongest alpha. Adjacent bin ranges overlapped for 50% to stabilize estimates. Then we averaged the corresponding gamma modulation of the trials belonging to each bin (consisting of the average power modulation within 59–100 Hz 0.2 to 0.6 s after trial onset, see Figure 3). Finally, we averaged across participants and plotted the median alpha value per bin averaged across participants against gamma modulation. See Rajagovindan and Ding (2011) for a similar procedure. To statistically test for the existence of inverted U-shaped relationships between alpha and gamma, we performed a one-way repeated measures ANOVA on gamma modulation with factor alpha bin (10 bins) to each condition separately and a two-way repeated measures ANOVA with factors bin and condition for testing the liberal–conservative difference (Figure 5F). Given the model prediction of a Gaussian-shaped relationship between alpha and gamma, we constructed a Gaussian contrast using the normal Gaussian shape with unit standard deviation (contrast values: -1000, -991, -825, 295, 2521, 2521, 295, -825, -991, -1000, values were chosen to sum to zero). For plotting purposes (Figure 5C-F), we computed within-subject error bars by removing within each participant the mean across conditions from the estimates.

Correlation between gamma modulation 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 both for each condition as well as for the ten alpha bins, while freeing the other parameters (drift rate, boundary separation, non-decision time) for each condition and fixing starting point across conditions. We then used repeated measures correlation to test whether stronger gamma was associated with stronger bias. Repeated measures correlation determines the common within-individual association for paired measures assessed on two or more occasions for multiple individuals by controlling for the specific range in which individuals' measurements operate, and correcting the correlation degrees of freedom for non-independence of repeated measurements obtained from each individual. Specifically, the correlation degrees of freedom were 14 participants $\times$ 10 observations – Number of participants – 1 = 140 – 14 – 1 = 125. Repeated measures correlation tends to have much greater statistical power than conventional correlation across individuals because neither averaging nor aggregation is necessary for an intra-individual research question. Please see Bakdash and Marusich (2017) for more information. We assessed the impact of single observations on the correlations by excluding observations exceeding five times the average Cook's distance of all values within each condition (five observations for liberal and four for conservative) and recomputing the correlations.

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. Permutation tests yield $p = 0$ if the observed value falls outside the range of the null distribution. In these cases, $p < 0.0001$ is reported in the manuscript. The standard deviation (s.d.) is reported as a measure of spread along with all participant-averaged results reported in the text. 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 the existence of inverted-U shaped relationships between gamma and alpha bins, we conducted repeated measures ANOVA's and Gaussian shaped contrasts (see section Response gain model test for details) using SPSS 23 (IBM, Inc.). We used Pearson correlation to test the link between parameter estimates of the DDM and SDT frameworks and repeated measures correlation to test the link between gamma power and drift bias (see previous section).