Criterion Setting Modulates Neural Excitability of Human Visual Cortex

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

[1] Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany
[2] Department of Psychology, University of Amsterdam, Amsterdam, The Netherlands; Amsterdam
[3] Center for Brain and Cognition, Institute for Interdisciplinary Studies, University of Amsterdam, Amsterdam, The Netherlands;
[4] Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, Germany;
[5] Center for Lifespan Psychology, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany
[6] Department of Experimental and Applied Psychology, Vrije Universiteit, van der Boechorststraat 1, 1081 BT Amsterdam, The Netherlands
* Shared senior author

Contact Information

Niels A. Kloosterman, Ph.D.

Max Planck UCL Centre for Computational Psychiatry and Ageing Research, Max Planck Institute for Human Development

Lentzeallee 94, 14195, Berlin, Germany

Phone: +49 30 82406 424

Summary

Biases, systematic tendencies toward one choice option, are hallmarks of decision-making under uncertainty. In perceptual decision-making, bias can be conceptualized as an internal reference to which incoming sensory evidence is compared. This reference, often called “criterion”, can be flexibly adjusted to match external asymmetries in the payoffs for certain outcomes. Yet, very little is known about how the human brain implements such strategic criterion shifts. Recent studies suggest that spontaneous fluctuations in neural excitability (indexed by suppression of prestimulus alpha-band (8-12 Hz) power in posterior cortex) may impact the criterion. Crucially however, it is currently unknown whether neural excitability and criterion can flexibly and intentionally be adjusted to meet external demands. Here, we experimentally induced criterion shifts in humans through through verbal instruction and assymetric reward contingencies and show for the first time that neural excitability is enhanced when humans adopt a liberal criterion compared to a more conservative criterion. Moreover, we show how increased excitability boosts subsequent stimulus-related visual cortical EEG activity in the gamma (59-100 Hz) range by enhancing sensory response gain. Drift diffusion modeling of choice behaviour further confirms that a liberal criterion is achieved by biasing the sensory evidence accumulation process towards ‘yes’ choices. Together, these findings show that humans are able to intentionally and flexibly adapt neural excitability to current

task demands, and that such changes in excitability implement criterion shifts by biasing sensory evidence accumulation.

Keywords

Perceptual decision-making, human, criterion, bias, decision threshold, drift-diffusion modeling, electroencephalography, neural oscillations, response gain

Highlights

- Human participants reported faint visual targets under task-induced liberal and conservative decision criteria.
- A liberal decision criterion suppressed prestimulus alpha-band EEG power and boosted sensory gamma responses in visual cortex.
- Participants achieved a liberal criterion by biasing sensory evidence accumulation towards 'yes' choices.
- Thus, criterion shifts (i.e. biases) are implemented by changing visual cortical neural excitability.

Results

Decisions are often reached not only through accumulation of sensory evidence, but can also be heavily biased by pre-existing attitudes and biases (Gross, 2017). These biases can psychologically be conceptualized as a flexible threshold of minimally required evidence for making a particular decision. The height of this threshold, dubbed "criterion" in signal detection theory (Green and Swets, 1966), is often imposed by the behavioral context, resulting in an optimal speed-accuracy tradeoff for the task at hand (Bogacz et al., 2010). Open questions in decision neuroscience are how this criterion is realized in the human brain, and how a liberal or conservative criterion affects the neural dynamics underlying decision making. Recently, a number of studies have revealed a link between a liberal criterion and spontaneous increases in "neural excitability", i.e. the brain's level of responsivity to external stimuli (Iemi et al., 2017; Limbach and Corballis, 2016). Specifically, these studies observed that perceptual decisions preceded by reduced rhythmic brain activity in the alpha (8-12 Hz) band were associated with an increased tendency to report stimulus presence, independent of the actual sensory input (Iemi et al., 2017; Limbach and Corballis, 2016; Samaha et al., 2017). It is unknown, however, (i) whether these correlations between neural excitability and response behavior reflect naturally occuring fluctuations (e.g., related to vigilance) or whether the act of criterion setting actively modulates neural excitability and (ii) how exactly the prestimulus excitability state affects post-stimulus neural processes involved in the decision process itself.

Here, we asked whether experimentally induced criterion shifts within persons would yield a corresponding modulation of neural excitability as expressed in prestimulus alpha activity. This would implicate neural excitability in the act of criterion setting. Further, previous work has linked suppression of prestimulus alpha to increased sensory responses, possibly by enhancing sensory response gain (Rajagovindan and Ding, 2011; Voytek, n.d.). This finding suggests that a liberal criterion might be implemented by enhancing cortical responses to task-relevant stimuli, such that the threshold for reporting the presence of a target is more often

reached. Finally, we applied drift diffusion modeling of our behavioural data to test the notion that criterion setting specifically affects the process of sensory evidence accumulation.

Experimental task manipulations induce robust criterion shifts

While their EEG was recorded, human participants viewed an ongoing stream of full-screen texture patterns presented at a frame rate of 40 ms. The orientation of the texture patterns varied randomly from frame to frame. The participants' task was to detect faint target stimuli embedded within this ongoing stream and report detected targets via a button press (Figure 1A). Targets were embedded within a fixed-order sequence of orientation textures, the start of which was implicitly signaled by a texture of diagonal orientation. In alternating nine-minute blocks of trials, we manipulated participants' criterion for reporting targets by instructing them either to report as many targets as possible ("Try to detect as many targets as possible!"; "liberal condition"), or to only report high-certainty targets ("Try to 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. To support the task instructions, we provided auditory feedback following missed targets (misses) in the liberal condition and falsely detected targets (false alarms) in the conservative condition, and applied monetary penalties for these errors (Figure 1A; see Methods for details).

Although the onset of targets was unpredictable and the transition into a fixed-order sequence was only implicitly signaled, reaction times (RT's) were clustered in time not only following target-present trials (hits), but also following target-absent trials (false alarms) (Figure 5C). This finding indicates that although trial onset was not explicitly cued, subjects were in fact sensitive to trial onset due to the diagonal orientation of the first texture in the fixed sequence (which was identical to one of the non-target stimuli). The EEG data support this notion by revealing significant stimulus-related power modulations even in trials in which a target was neither presented nor (falsely) detected (see next section and Figure S1).

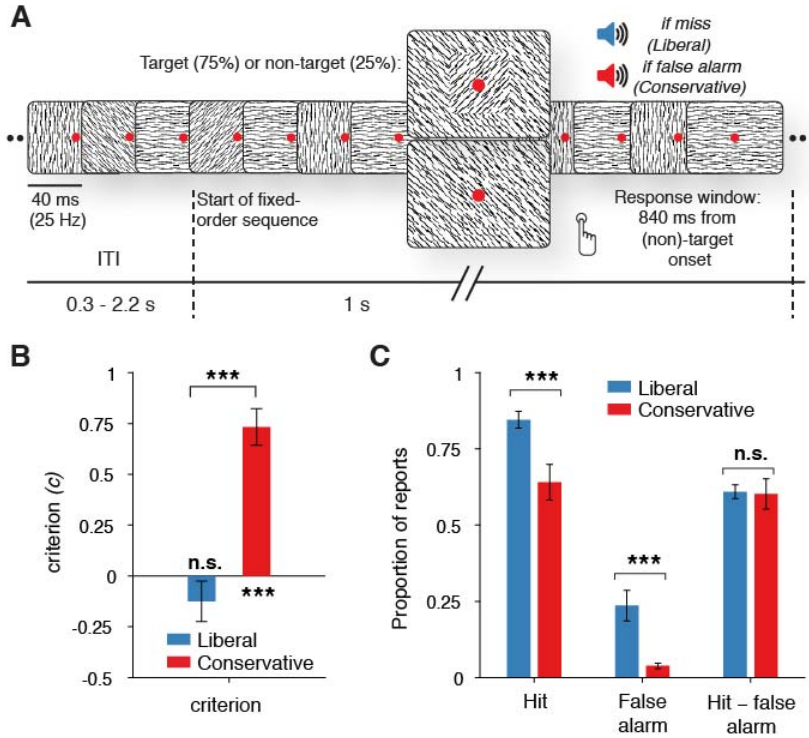

Figure 1 | Experimental design and behavioral results. A. Schematic of the visual stimulus and task design. Participants viewed an ongoing stream of full-screen diagonally, horizontally and vertically oriented textures at a frame rate of 40 ms (25 Hz). After random inter-trial intervals (range 0.3-2.2 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 (non-target), or contained an orthogonal orientation-defined square (target). Participants had to decide 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 both aversive tones and monetary deductions. Depicted square and fixation dot sizes are not to scale. **B.** Signal-detection-theoretic criterion during both conditions. **C.** Average proportion of hit and false alarm rates for each condition and average sensitivity quantified by the difference between hit and false alarm rates. Error bars, SEM across participants (N = 16). *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$; n.s., not significant.

Participants adjusted their criterion (*c*, quantified with signal detection theory (Green and Swets, 1966); see Methods) depending on the experimental condition. The average criterion across participants was significantly greater than zero in the conservative condition ($c = 0.73$, all participants > 0, $p < 0.0001$, two-sided permutation test, 10,000 permutations), and slightly below zero in the liberal condition ($c = -0.13$, 11 out of 16 participants < 0, $p = 0.22$). Critically, *c* was significantly lower in the liberal than in the conservative condition ($p < 0.0001$). Given that $c = 0$ indicates a neutral criterion (i.e. no bias), this finding indicates that participants indeed adopted more liberal and more conservative criteria during both experimental conditions, while overall maintaining a relatively conservative attitude, in line with previous studies (de Gee et al., 2017a; 2014).

In contrast to subjective criterion, objective target detection performance was relatively comparable between conditions. Specifically, the difference between hit and false alarm rates did not differ between the experimental conditions ($p = 0.8185$; Figure 1C). Observers were, however, less sensitive as quantified using signal-detection theoretic *d'* (Green and Swets, 1966); see Methods)) in the liberal condition: *d'* of 2.0 vs 2.31, $p = 0.0002$. Interestingly, reaction times were consistently shorter in the liberal condition too (mean RT liberal: 0.43 s (s.d. 0.03), conservative: 0.47 s (s.d. 0.05); $p = 0.0002$). Together, the decreased performance and faster responses in the liberal condition suggest that participants adapted their speed-accuracy tradeoff (SAT) depending on experimental condition to minimize losses, while detection ability remained similar (Bogacz et al., 2006). Indeed, whereas *c* was on average 132% (s.d. 77) lower in the liberal than in the conservative condition, *d'* decreased only 17% (s.d. 14)(*d'* vs. *c*: $p < 0.0001$), indicating that the experimental manipulations primarily affected the criterion. Taken together, the behavioral findings validate the experimentally induced criterion manipulation by demonstrating that participants adapted their criterion, thereby affecting their SAT while objective performance largely remained constant.

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

Understanding the neural mechanisms underlying criterion setting first entails a careful examination of the stimulus-related signals on which the criterion is thought to operate. Such stimulus-related signals are typically reflected in visual cortical population activity exhibiting rhythmic temporal structure (Buzsáki and Draguhn, 2004). Specifically, bottom-up processing of visual information has previously been linked to increased high-frequency (> 40 Hz) electrophysiological activity over visual cortex (Bastos et al., 2015; Michalareas et al., 2016; Popov et al., 2017; van Kerkoerle et al., 2014). Figure 2A shows time-frequency representations of the EEG activity recorded over visual cortex during the trial period. We observed a robust high-frequency (40-100 Hz, 'gamma') power increase immediately following presentation of the task-relevant texture from 0.2 to 0.6 s after stimulus sequence onset ($p < 0.05$, cluster-corrected for multiple comparisons, two-sided)(Figure 2A, upper panel). The topography of this power modulation was confined to posterior electrodes (Figure 2B, top), in line with a role in visual cortical processing. Thus, the appearance of task-relevant texture pattern in the continuous visual stream reliably induced a high-frequency stimulus-related response in visual cortex (Ni et al., 2016).

Concurrent with this enhanced gamma activity, we observed several additional power modulations related to visual processing. First, power increased in electrodes Pz and POz in a narrow frequency range around 25 Hz, reflecting the visual stimulation frequency of our experimental paradigm (Figure 2A, lower panel)(Müller et al., 1997). This visual evoked potential was similarly expressed around 50 Hz – the first harmonic of the stimulation frequency (Figure 2A, top panel). Finally, we observed suppression of low-frequency (11-22 Hz) activity in posterior cortex, which typically occurs in parallel with enhanced stimulus-related gamma activity (Donner and Siegel, 2011; Kloosterman et al., 2015; Meindertsma et al., 2017; Werkle-Bergner et al., 2014)(Figure 2A, lower panel and 2B, right). Importantly, these low-frequency power modulations were significant even for correct rejection trials, in which a target was neither presented nor (falsely) detected (Figure S1), indicating that participants were sensitive to the structure of the detection task,

despite the lack of clear trial onset cues. To investigate criterion-related low-frequency power modulations in visual cortex before trial onset (described in the next section), we selected eleven posterior electrodes exhibiting enhanced gamma (59-100 Hz) activity post-stimulus (highlighted electrodes in scalp maps in Figure 2B, top).

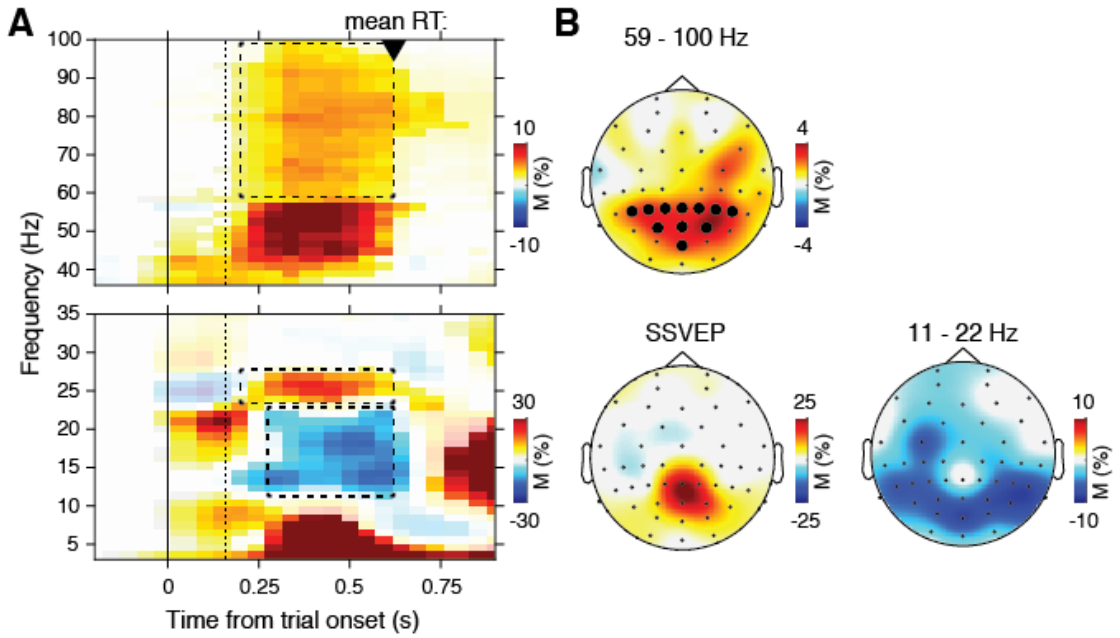

Figure 2 | Cortical responses to the visual stimulus. A. Time-frequency representations of high- (top) and low-frequency (bottom) EEG power modulations with respect to the prestimulus period (-0.4 - 0 s), pooled over the two conditions. Saturated colors indicate clusters of significant modulation, cluster threshold $p < 0.05$, two-sided permutation test across participants, cluster-corrected; N = 15). Solid and dotted vertical lines respectively indicate the onset of the trial and the target stimulus. M, power modulation. **B.** Scalp maps showing topography of stimulus-related modulations (0.2–0.6 s (gamma and SSVEP) or 0.25–0.6 s (low frequency suppression) after stimulus onset; see dashed outlines on time-frequency representations in A. Thick dots indicate electrodes used for the TFR's in A and which were selected for further analysis. SSVEP, steady state visual evoked potential.

Adopting a more liberal criterion increases neural excitability

Previous studies have shown that neural excitability, as reflected in suppression of prestimulus alpha (8-12 Hz) power, correlates with a decision-maker's propensity to categorize sensory input as a target (Iemi et al., 2017; Limbach and Corballis, 2016). This finding, however, does not elucidate the role of prestimulus alpha in criterion setting. Is the suppression of alpha activity merely a byproduct of a more liberal criterion? Or do controlled criterion shifts actively modulate prestimulus alpha power, suggesting a critical role of alpha in criterion setting? To address this issue, we examined the effect of experimentally induced criterion shifts on alpha power between 0.8 and 0.2 s before trial onset. The spatial topography of prestimulus alpha power pooled across the two conditions is plotted in Figure 3A. Alpha power was indeed strongest over our visual cortical electrode pooling of interest, suggesting involvement of this alpha activity in visual processing. Strikingly however, prestimulus alpha power was suppressed during the experimentally induced liberal criterion when compared to the experimentally induced conservative criterion (Figure

3B), suggesting within-person modulation of alpha power during criterion setting. Expressing spectral power during the liberal condition as the percentage signal change from the conservative condition revealed a statistically significant cluster of suppressed frequencies precisely in the 8-12 Hz range ($p < 0.05$, cluster-corrected for multiple comparisons)(Figure 3D). This alpha suppression was located in a posterior region of the cerebral cortex largely overlapping with our visual electrode pooling of interest (Figure 3C).

Taken together, these results show that experimentally induced criterion shifts are associated with a decrease of pre-stimulus alpha under a liberal criterion when compared to a conservative criterion. This finding suggests that alpha modulations are a hallmark of criterion setting, rather than being a spontaneously occuring haphazard influence on the criterion. Importantly, our finding suggests that humans are able to actively control the excitability of their brain in service of upcoming decisions.

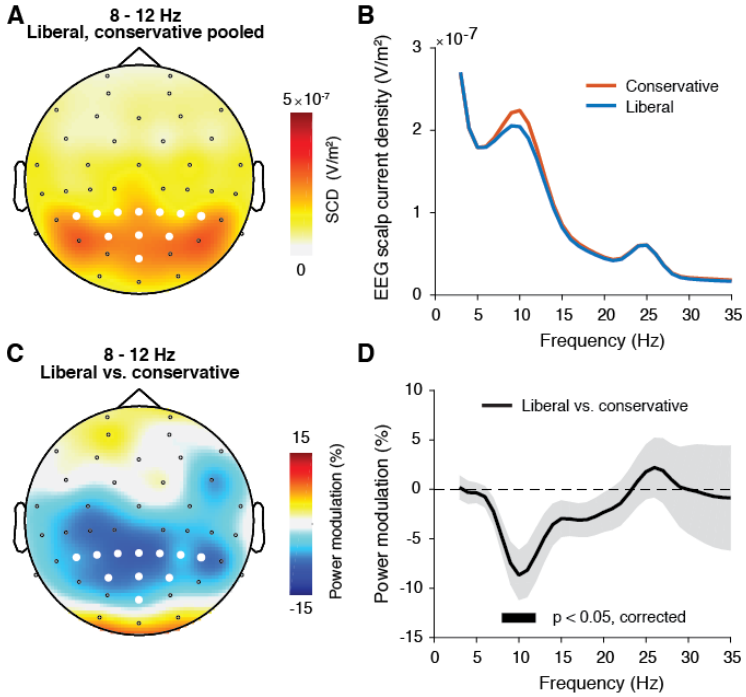

Figure 3 | Liberal decision criterion suppresses prestimulus alpha power. A. Scalp map of raw prestimulus EEG alpha power (8-12 Hz neural activity between 0.8 and 0.2 s before sequence onset), pooled over conditions. White symbols indicate visual cortical electrodes used for the power spectra in B. and D. **B.** Low-frequency power spectra of prestimulus neural activity for both conditions. **C.** Scalp map of power modulation in the liberal with respect to the conservative condition, expressed as percent signal change from the conservative condition. **D.** Corresponding liberal versus conservative power spectrum. Black horizontal bar indicates statistically significant frequency range ($p < 0.05$, cluster-corrected for multiple comparisons, two-sided). Error bars, SEM across participants (N = 15).

Enhanced neural excitability boosts sensory responses in visual cortex

How could increased neural excitability result in a more liberal response attitude? One possibility is that increased excitability enhances stimulus-related activity on both target and no-target trials, thereby increasing the likelihood that this activity exceeds the subjective threshold for detecting a target. We explored this possibility

using an existing theoretical framework that models the output activity of visual cortex as an s-shaped (sigmoidal) function of two factors: (1) the brain's excitability state in the current trial, here represented by prestimulus alpha power, and (2) stimulus-related neural activity, here represented by post-stimulus gamma power (Rajagovindan and Ding, 2011) (Figure 4A). Assuming that stimulus-induced input activity is roughly similar across trials, the magnitude of the stimulus-related response in visual cortex is captured by the first derivative (i.e. the slope) of the sigmoidal function, yielding an inverted-U shaped function of prestimulus alpha power (Figure 4B) (Destexhe et al., 2001; Freeman, 1979). Thus, according to the model, intermediate levels of neural excitability produce the strongest, and lower and higher excitability levels produce weaker, stimulus-related activity. Given observed increased excitability in the liberal condition, the model predicts enhanced stimulus processing during the liberal condition (Figure 4B), particularly at intermediate prestimulus alpha levels (Rajagovindan and Ding, 2011), through an increase in the response gain (i.e. steeper slope of the solid lines in Figure 4B).

To test our prediction, we followed the method put forward by Rajagovindan et al (2011). We exploited the large number of trials per participant per condition (range 543 to 1391 trials) by sorting each participant's trials within each condition into ten bins based on neural excitability, as reflected in log-transformed prestimulus alpha power ranging from strong (indicating low excitability) to weak (indicating high excitability)(Figure 4C). We then averaged across the trials within each excitability bin the log-transformed stimulus-related visual cortical gamma power, and normalized the binned average gamma power by subtracting the participant's mean gamma power in the conservative condition. Finally, we plotted the excitability bin number against the normalized gamma power, averaged across participants (see Methods for details and (Rajagovindan and Ding, 2011) for a similar procedure).

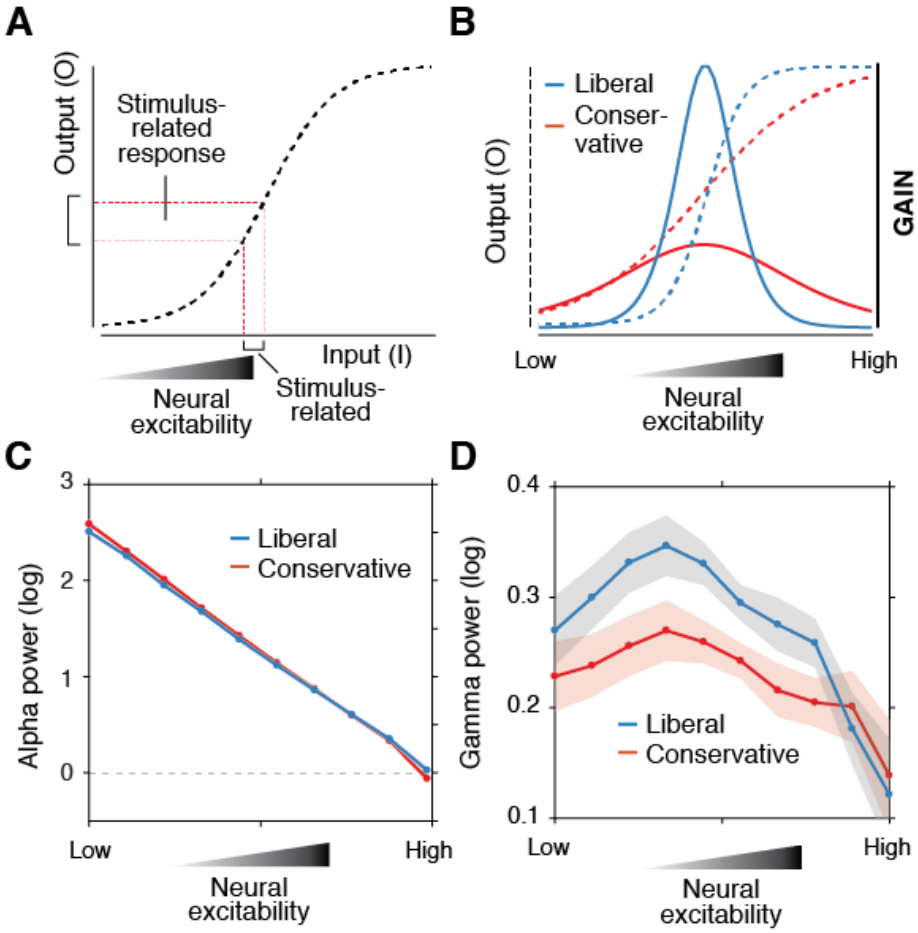

Figure 4 | Theoretical response gain model and empirical test. A. Response gain model, which describes the relationship stimulus-related input activity to output activity in visual cortex as a sigmoidal function, such that the current excitability state determines the output strength. **B.** Model predictions. Stimulus-related cortical output responses (solid lines) are formalized as the first derivative of the sigmoidal functions (dotted lines), resulting in an inverse-U shaped response gain function. The model predicts that a liberal criterion increases the steepness of the sigmoidal function (right) compared to a conservative criterion (left), resulting in stronger stimulus-related responses (Rajagovindan and Ding, 2011). **C.** Neural excitability reflected in single trial, log-transformed prestimulus alpha power sorted across ten bins from high (indicating low excitability) to low (indicating high excitability), separately for both conditions. **D.** Corresponding log-transformed gamma activity (normalized within participants by subtracting the minimum gamma power during the conservative condition) plotted as a function of neural excitability. Error bars, within-subject SEM across participants (N = 14).

The resulting plot closely resembles an inverted-U shaped relationship between excitability and stimulus-related gamma activity for both conditions, with particularly low gamma responses for high excitability trials (Figure 4D). Critically, average gamma activity was higher for the liberal than for the conservative condition, except during the highest excitability states (Figure 4D, rightmost two data points). Indeed, the curve for the liberal condition was steeper for the liberal condition, suggesting increased response gain. A three-way repeated measures ANOVA with factors condition (conservative, liberal), brain activity type (prestimulus alpha, poststimulus gamma) and bin level (1-10) revealed a significant three-way interaction between these factors ($F(9,117) = 2.96$, $p = 0.003$, partial $\eta^2 = 0.19$, Greenhouse-Geisser corrected $p = 0.046$). Importantly, the marginally significant quadratic

contrast ($F(1,13) = 3.47$, $p = 0.085$, partial $\eta^2 = 0.21$) fitted this interaction almost as well as a linear contrast ($F(1,13) = 4.69$, $p = 0.049$, partial $\eta^2 = 0.265$). This three-way quadratic interaction indeed suggests a more steeply U-shaped curve for gamma responses in the liberal condition, in line with enhanced gain. Taken together, these findings suggest that the increased excitability during the liberal condition boosted sensory stimulus processing. In turn, this boosted activity might have increased the participant's propensity to categorize both target and non-target stimuli as signal, resulting in a more liberal response attitude.

Evidence accumulation bias underlies experimentally induced liberal criterion setting

Finally, we used computational modeling of our behavioral data to investigate whether an experimentally induced liberal criterion specifically affects the process of sensory evidence accumulation, as suggested by our observed enhancement of stimulus-related activity. To this end, we fitted our reaction time data (Figure 5C) using the drift diffusion model, an established dynamic model of two-choice decision processes (Ratcliff and McKoon, 2008)(Figure 5A). This model postulates that decisions are reached by accumulation of noisy sensory evidence towards one of two decision boundaries, which can either be explicit (for 'yes' responses in our experiment), or implicit (i.e. without active response, for 'no decisions' in our experiment) (de Gee et al., 2017b; Ratcliff et al., 2016) . Within this model, a liberal decision bias can emerge in two different ways: either by moving the starting point of evidence accumulation closer to the 'yes' decision boundary ('starting point' parameter), or by driving the evidence accumulation process itself more towards the 'yes' boundary ('drift criterion' parameter, equivalent to a bias in evidence accumulation, implemented by adding an evidence-independent constant to the drift). Previous work has shown that an evidence accumulation bias drives a less conservative (i.e. more liberal attitude during pupil-linked arousal (de Gee et al., 2017a).

To test whether an evidence accumulation bias similarly underlies decision bias during an experimentally induced liberal criterion, we fitted the model letting various parameters free to vary per experimental condition (drift rate, boundary separation, non-decision time and drift criterion) while keeping starting point fixed. This model ("drift criterion model") fitted our data well (Figures 5B and S2). Indeed, whereas actual drift rate was significantly lower in the liberal than in the conservative condition ($p < 0.0001$, permutation test), drift rate was supported by a highly positive drift criterion parameter during the liberal condition, whereas this parameter was not significantly different from zero for the conservative condition ($p = 0.90$ for conservative; liberal vs, conservative: $p < 0.0001$; Figure 5D, top). Strikingly, the condition-induced shifts in SDT criterion and drift criterion were strongly correlated across participants ($r = -0.89$, $n = 16$, $p = 4.1e^{-6}$), indicating that these two measures reflect similar aspects of our data. In addition, starting point (fixed across conditions) was below 0.5, indicating that this parameter was overall closer to the "no" boundary, in line with an overall conservative response attitude ($p < 0.0001$) as also observed in SDT criterion (compare Figures 1B and 5D, bottom left). Finally, boundary separation also increased slightly but reliably during liberal ($p = 0.0001$) (Figure 5D, top left), whereas non-decision time decreased ($p = 0.0001$) (Figure 5D, bottom right).

As a first control of the goodness of fit of the condition-dependent effect on drift criterion we re-fitted the model, but now fixing drift criterion for both experimental conditions, while still allowing all other of the above (non-bias-related) parameters to vary freely. This “basic model” provided a worse fit to the data, as indicated by higher Bayesian Information Criterion (BIC) estimates than for the drift criterion model (Figure 5E)(see Methods for details). As a second control, we fitted the model again while fixing drift criterion for both experimental conditions, and instead allowing starting point as well as all other of the above parameters to vary (“starting point model”). This model also provided a worse fit to the data (Figure 5E). Specifically, for 12 out of 16 participants the drift criterion model provided better fits to behavior than the starting point model, and for 10 of which delta BIC was greater than 10 (indicating very strong evidence against the starting point model). Taken together, our modelling results suggest that participants achieved a liberal decision bias specifically by biasing the evidence accumulation process towards ‘yes’ decisions, not its starting point, while non-bias-related parameters were also affected.

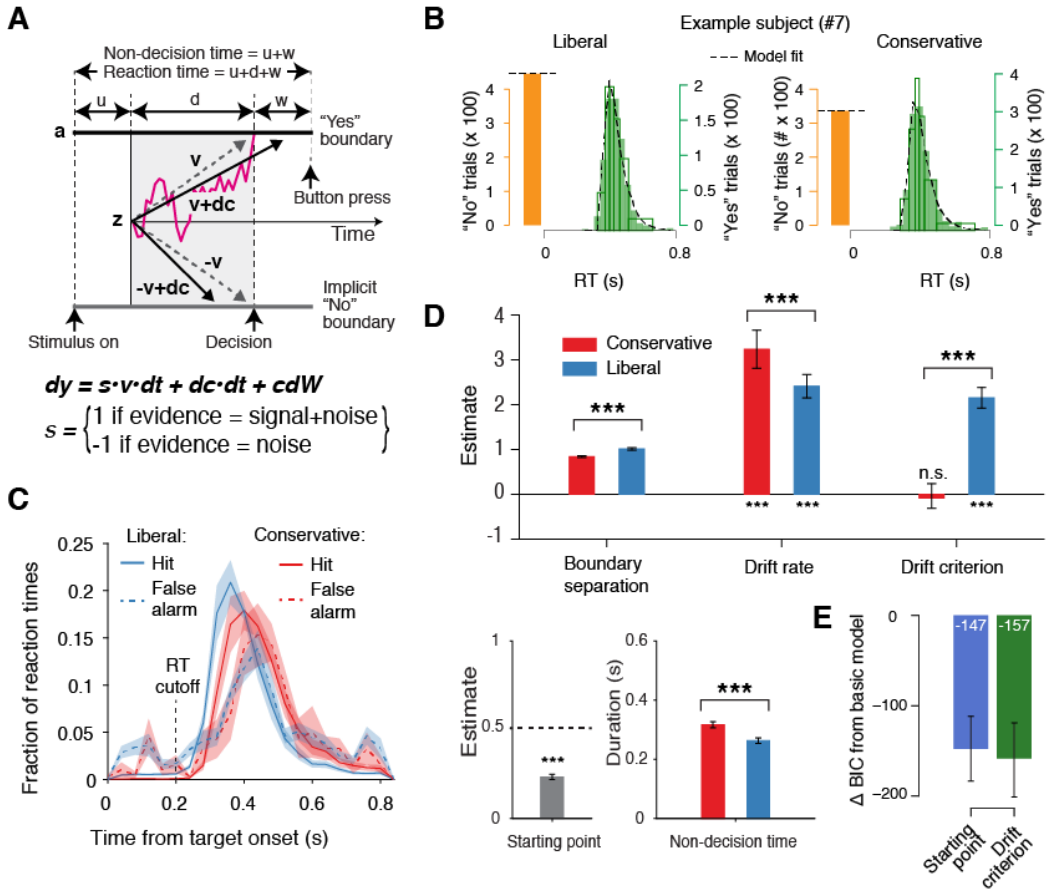

Figure 5 | Experimentally induced liberal criterion is realized through an evidence accumulation bias. A. Schematic and simplified equation of drift diffusion model accounting for RT distributions for explicit ‘yes’- and implicit ‘no’-choices (‘stimulus coding’; see Materials and methods). Notation: dy, change in decision variable y per unit time dt; v.dt, mean drift (multiplied with 1 for signal +noise trials, and 1 for noise trials); dc.dt, drift criterion (an evidence-independent constant added to the drift); and cdW, Gaussian white noise (mean = 0, variance = c2 dt). **B.** RT distributions of an example subject for ‘yes’- choices and the number of implicit ‘no’ choices, separately for the two conditions. See Figure S2 for all participant data. Green bars, observed RT quantiles. **C.** Subject-average RT distributions for hits and false alarms in both conditions. **D.** Estimated model parameters for the drift criterion model. Conventions as in Figure 1. N = 16 participants. **E.** BIC goodness of fit

estimates for the starting point and drift criterion models, expressed with respect to a basic model without bias parameters. A lower delta BIC value indicates a better fit.

Discussion

The subjective threshold for choosing a course of action plays a central role in every decision we make. To date, however, the neural underpinnings of decision criterion setting have remained elusive. Here, we demonstrate that instructed modulations of an observer's criterion resulted in robust changes in the excitability state of the brain, as reflected in prestimulus alpha activity over posterior cortex. Moving beyond mere description, we demonstrated that a more excitable prestimulus brain state, associated with a liberal response criterion, boosted stimulus-related high-frequency cortical responses over visual cortex by increasing response gain. Drift diffusion modeling of our reaction time data confirmed the notion that criterion shifts are achieved by biasing the evidence accumulation process. Our results suggest that criterion-induced enhanced neural processing equally affects both stimulus-related and internal 'noise' signals in visual cortex, thereby increasing both the number of correct (hits), but also incorrect (false alarms) behavioral reports (Iemi et al., 2017).

One neural mechanism that could underlie this enhanced processing may be under control of the catecholaminergic neuromodulatory systems, consisting of the noradrenaline and dopamine systems (Aston-Jones and Cohen, 2005). These systems are able to modulate the level of arousal and neural gain in the brain, and show tight links with pupil responses (de Gee et al., 2017a; 2014; Joshi et al., 2015; McGinley et al., 2015). In line with this notion, prestimulus alpha power suppression has also recently been linked to pupil dilation (Meindertsma et al., 2017). From this perspective, our results reconcile previous studies showing relationships between a liberal criterion and modulation of spontaneous alpha power as well as pupil size. Taken together, a more liberal within-person criterion (following experimental instruction) might activate neuromodulatory systems, which subsequently increase cortical excitability and enhance sensory responses.

One difference with previous studies investigating the link between neural excitability and decision criterion is that previous experimental designs contained active 'no' decisions (Iemi et al., 2017; Limbach and Corballis, 2016), whereas our participants only executed active 'yes' decisions and 'no' decisions were implicit. This design was chosen to create a naturalistic setting, in which participants were encouraged to detect targets whenever they appeared to them. Although our finding of increased neural excitability during a liberal criterion is consistent with previous studies, future studies could include active 'no' decisions to test whether a liberal response attitude increases excitability similarly for both response alternatives. A concern may be that participants did not understand the task structure in the continuous sequence of visual events. However, hits and false alarms showed normal reaction time distributions following both targets and non-targets, suggesting that participant were able to extract the implicit task structure. The tight drift diffusion model fits (Figure S2), as well as the low-frequency EEG power modulations (Figure S1) confirm this notion. Another concern might be that our results are due to differences between the conditions in detection performance and RT. While we cannot fully rule out this possibility, we speculate that a more liberal response attitude might in general be associated with an emphasis on speed versus accuracy. As previous criterion studies do not report reaction times, future studies could address this issue further.

A last issue might be that the high-frequency EEG power modulation we observed was due to microsaccade-related activity (Yuval-Greenberg et al., 2008). We deem this scenario unlikely for several reasons. First, our highly dynamic and strong visual stimulus was optimal to elicit high-frequency responses, as has been shown with MEG, which is less susceptible to microsaccades (Hoogenboom et al., 2006; Siegel et al., 2006). Second, we reduced the sensitivity of the EEG to microsaccades by using the earlobes instead of the nose as the reference (Yuval-Greenberg et al., 2008), by removing microsaccade-related activity from the data (Hassler et al., 2011; Joerg F Hipp, 2013), and by applying a current source density transformation to our data (Melloni et al., 2009)(see Methods for details). Finally, the onset of a task-relevant event is often associated with a drop in microsaccade rate (Bonneh et al., 2010; Rolfs, 2009), which should have resulted in a corresponding drop in high-frequency power; conversely, we find that high frequency power increases at target onset. Future work could further address this issue by measuring eye movements during the experimental tasks.

On a final note, our results seem to suggest that stimulus-related responses are boosted during the liberal condition due to increased cortical response gain, which is further supported by recent work linking alpha power suppression to gain (Voytek, n.d.) and another recent paper showing that excitability biases perceptual experiences, but not the decision-making strategy (Iemi and Busch, 2017). More conclusive evidence, however, could be provided by explicitly manipulating cortical response gain during a criterion manipulation, for instance by pharmacological manipulation of the noradrenergic LC-NE system (Servan-Schreiber et al., 1990), or by enhancing occipital alpha power using transcranial stimulation (Zaehle et al., 2010).

Although we are mostly unaware of it, every decision we make is tainted by implicit biases that work on the noisy evidence accumulation process towards one of the alternatives. Understanding how these counterproductive biases affect our decisions is key to becoming aware of these biases and successfully counteract them. Pinpointing the neural mechanisms underlying bias in an elementary perceptual task paves the way for studying how more abstract, ecologically relevant decisions are affected by bias (Tversky and Kahneman, 1974).

Author Contributions

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

Acknowledgments

MWB is supported by a grant from the German Research Foundation (DFG; WE4296/5-1), as well as the Jacobs Foundation via an Early Career Research Fellowship.

References

Aston-Jones, G., Cohen, J.D., 2005. An integrative theory of locus coeruleus-

norepinephrine function: adaptive gain and optimal performance. Annu Rev Neurosci 28, 403–450. doi:10.1146/annurev.neuro.28.061604.135709
Bastos, A.M., Vezoli, J., Bosman, C.A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J.R., De Weerd, P., Kennedy, H., Fries, P., 2015. Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. Neuron 85, 390–401. doi:10.1016/j.neuron.2014.12.018
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. Psychol Rev 113, 700–765. doi:10.1037/0033-295X.113.4.700
Bogacz, R., Wagenmakers, E.-J., Forstmann, B.U., Nieuwenhuis, S., 2010. The neural basis of the speed-accuracy tradeoff. Trends Neurosci 33, 10–16. doi:10.1016/j.tins.2009.09.002
Bonneh, Y.S., Donner, T.H., Sagi, D., Fried, M., Cooperman, A., Heeger, D.J., Arieli, A., 2010. Motion-induced blindness and microsaccades: Cause and effect. J Vis 10, 22–22. doi:10.1167/10.14.22
Buzsáki, G., Draguhn, A., 2004. Neuronal oscillations in cortical networks. Science 304, 1926–1929. doi:10.1126/science.1099745
de Gee, J.W., Colizoli, O., Kloosterman, N.A., Knapen, T., Nieuwenhuis, S., Donner, T.H., 2017a. Dynamic modulation of decision biases by brainstem arousal systems. eLife Sciences 6, 309. doi:10.7554/eLife.23232
de Gee, J.W., Knapen, T., Donner, T.H., 2014. Decision-related pupil dilation reflects upcoming choice and individual bias. Proc. Natl. Acad. Sci. U.S.A. 111, E618–25. doi:10.1073/pnas.1317557111
de Gee, J.W., Tsetsos, K., McCormick, D.A., McGinley, M.J., Donner, T.H., 2017b. Phasic pupil-linked arousal reduces decision biases in mice and men. Soc for Neurosci abstr 713.07 / UU49.
Destexhe, A., Rudolph, M., Fellous, J.M., Sejnowski, T.J., 2001. Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. Neuroscience 107, 13–24. doi:10.1016/S0306-4522(01)00344-X
Donner, T.H., Siegel, M., 2011. A framework for local cortical oscillation patterns. Trends Cogn Sci 15, 191–199. doi:10.1016/j.tics.2011.03.007
Efron, B., Tibshirani, R., 1998. The problem of regions. The Annals of Statistics 26, 1687–1718. doi:10.1214/aos/1024691353
Fahrenfort, J.J., Scholte, H.S., Lamme, V.A.F., 2007. Masking disrupts reentrant processing in human visual cortex. J Cogn Neurosci 19, 1488–1497. doi: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
Freeman, W.J., 1979. Nonlinear gain mediating cortical stimulus-response relations. Biol Cybern 33, 237–247. doi:10.1007/BF00337412
Green, D.M., Swets, J.A., 1966. Signal detection theory and psychophysics. Society 1, 521.
Gross, M., 2017. Can we change our biased minds? Curr Biol 27, R1089–R1091. doi:10.1016/j.cub.2017.10.013
Hassler, U., Trujillo-Barreto, N., Gruber, T., 2011. Induced gamma band responses in human EEG after the control of miniature saccadic artifacts. Neuroimage 57, 1411–1421. doi:10.1016/j.neuroimage.2011.05.062
Hoogenboom, N., Schoffelen, J.-M., Oostenveld, R., Parkes, L.M., Fries, P., 2006. Localizing human visual gamma-band activity in frequency, time and space. Neuroimage 29, 764–773. doi:10.1016/j.neuroimage.2005.08.043

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

Iemi, L., Chaumon, M., Crouzet, S.M., Busch, N.A., 2017. Spontaneous Neural Oscillations Bias Perception by Modulating Baseline Excitability. J Neurosci 37, 807–819. doi:10.1523/JNEUROSCI.1432-16.2017

Joerg F Hipp, M.S., 2013. Dissociating neuronal gamma-band activity from cranial and ocular muscle activity in EEG. Front Hum Neurosci 7, 338. doi:10.3389/fnhum.2013.00338

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, 221–234. doi:10.1016/j.neuron.2015.11.028

Kloosterman, N.A., Meindertsma, T., Hillebrand, A., van Dijk, B.W., Lamme, V.A.F., Donner, T.H., 2015. Top-down modulation in human visual cortex predicts the stability of a perceptual illusion. J Neurophysiol 113, 1063–1076. doi:10.1152/jn.00338.2014

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

Maris, E., Oostenveld, R., 2007. Nonparametric statistical testing of EEG-and MEG-data. J. Neurosci. Methods 164, 177–190. doi:10.1016/j.jneumeth.2007.03.024

McGinley, M.J., David, S.V., McCormick, D.A., 2015. Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. Neuron 87, 179–192. doi: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, 5744–5757. doi: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, 8–10– author reply 10–12. doi: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, 384–397. doi:10.1016/j.neuron.2015.12.018

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

Müller, M.M., Teder, W., Hillyard, S.A., 1997. Magnetoencephalographic recording of steadystate visual evoked cortical activity. Brain Topogr 9, 163–168. doi:10.1007/BF01190385

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

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

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, 1–9.

doi: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, 184–187. doi:10.1016/0013-4694(89)90180-6
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, 4117–4127. doi:10.1523/JNEUROSCI.3015-16.2017
Rajagovindan, R., Ding, M., 2011. From prestimulus alpha oscillation to visual-evoked response: an inverted-U function and its attentional modulation. J Cogn Neurosci 23, 1379–1394. doi:10.1162/jocn.2010.21478
Ratcliff, R., Huang-Pollock, C., McKoon, G., 2016. Modeling Individual Differences in the Go/No-Go Task With a Diffusion Model [WWW Document]. psycnet.apa.org. doi:http:// dx.doi.org/10.1037/dec0000065
Ratcliff, R., McKoon, G., 2008. The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. Neural Comput 20, 873–922. doi:10.1162/neco.2008.12-06-420
Rolfs, M., 2009. Microsaccades: small steps on a long way. Vis. Res. 49, 2415–2441. doi:10.1016/j.visres.2009.08.010
Samaha, J., Iemi, L., Postle, B.R., 2017. Prestimulus alpha-band power biases visual discrimination confidence, but not accuracy. Consciousness and Cognition. doi: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 249, 892–895.
Siegel, M., Donner, T.H., Oostenveld, R., Fries, P., Engel, A.K., 2006. High-Frequency Activity in Human Visual Cortex Is Modulated by Visual Motion Strength. Cerebral Cortex 17, 732–741. doi:10.1093/cercor/bhk025
Tversky, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. Science 185, 1124–1131. doi:10.1126/science.185.4157.1124
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. Proc. Natl. Acad. Sci. U.S.A. 111, 14332–14341. doi:10.1073/pnas.1402773111
Voytek, B., n.d. **Alpha oscillations control cortical gain by modulating excitatory- inhibitory background activity** [WWW Document]. google.com. URL https://www.google.com/url?hl=en-GB&q=http://www.biorxiv.org/content/biorxiv/early/2017/09/07/185074.full.pdf&source=gmail&ust=1504955978313000&usg=AFQjCNHq4s99IMfcV9UfIs5yiY_pdNeuIw (accessed 9.8.17).
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, 8519–8528. doi:10.1523/JNEUROSCI.3801-13.2014
Wiecki, T.V., Sofer, I., Frank, M.J., 2013. HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. Front. Neuroinform. 7. doi: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, 429–441. doi: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, e13766. doi:10.1371/journal.pone.0013766

Figure and Table Legends

[In the text for now]

Materials and Methods

Participants Sixteen participants (eight female, mean age 24.1 years, ± 1.64) took part in the experiment, either for financial compensation or in partial fulfillment of first year course requirements. Each participant completed three experimental sessions on different days, each session lasting ca. 1.5 hours, including preparation and breaks. One participant completed only two sessions, yielding a total number of 1-h measurements across subjects = 47. Due to technical issues, for one session only data for the liberal condition was available. One participant was an author. All participants were included in the behavioral and drift diffusion modeling analyses. One participant was excluded from the alpha-power analysis (Figure 3) due to excessive noise (EEG power spectrum opposite of 1/f). One further participant was excluded from the single-trial gamma power modulation analyses (Figure 4) because the liberal-conservative difference in gamma power in this participant was > 3 standard deviations away from the other participants. In summary, 16 participants were included in the analyses presented in Figures 1 and 5, 15 participants in Figures 2 and 3, and 14 participants were included in Figure 4. All participants had normal or corrected-to-normal vision and were right handed. Participants provided written informed consent before the start of the experiment. All procedures were approved by the ethics committee of the University of Amsterdam.

Stimuli Stimuli consisted of a continuous semi-random rapid serial visual presentation (rsvp) of full screen texture patterns. The texture patterns consisted of line elements approx. 0.07° thick and 0.4° long in visual angle. Each texture in the rsvp was presented for 40 ms (i.e. stimulation frequency 25 Hz), and was oriented in one of four possible directions: 0°, 45°, 90° or 135°. Participants were asked to fixate on a red dot in the center of the screen. At random inter trial intervals (ITI's) sampled from a uniform distribution (ITI range 0.3 - 2.2 s), the rsvp contained a fixed sequence of 25 texture patterns, which in total lasted one second. This fixed sequence consisted of four stimuli preceding a (non-)target stimulus (orientations of 45°, 90°, 0°, 90° respectively) and twenty stimuli following the (non)-target (orientations of 0°, 90°, 0°, 90°, 0°, 45°, 0°, 135°, 90°, 45°, 0°, 135°, 0°, 45°, 90°, 45°, 90°, 135°, 0°, 135° respectively) (Figure 1). The fifth texture pattern within the sequence (occurring from 0.16 s after sequence onset) was either a target or a non-target stimulus. Non-targets 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. Of all trials, 75% contained a target and 25% a non-target. Target and non-target trials were presented in random order. To avoid specific influences on target stimulus visibility due to presentation of similarly or orthogonally oriented texture patterns temporally close in the cascade, no 45° and 135° oriented stimuli were presented directly before or after presentation of the target stimulus. All stimuli had an isoluminance of 72.2 cd/m^2. Stimuli were

created using MATLAB (The Mathworks, Inc., Natick, MA, USA) and presented using Presentation (Neurobehavioral systems, Inc., Albany, CA, USA).

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

At the beginning of the experiment, participants were informed they could earn a total bonus of EUR 30, on top of their regular pay or course credit. In two separate conditions within each session of testing, we encouraged participants to use either a conservative or a liberal criterion 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 criteria 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 to eliminate differences in payment between participants.

During a block, participants continuously monitored the screen and were free to respond by button press whenever they thought they saw a target. Participants performed six blocks per session. Each block contained 240 trials, 180 target and 60 non-target 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 criterion accordingly.

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

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

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

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

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

EEG recording Continuous EEG data were recorded at 256 Hz using a 48-channel BioSemi Active-Two system (Biosemi, 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 et al., 2011) and custom software in MATLAB (version R2016b, The Mathworks) to process the data (see below). Data were 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 0.8 s after target offset; 2) when the onset of the the next target stimulus sequence started before trial end, the trial was terminated at next trials 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 a similar trial extraction procedure. After trial extraction, channel time courses were linearly detrended and the mean of every channel was removed per trial.

Artifact rejection Trials containing muscle artifacts were rejected from further analysis using a standard semi-automatic preprocessing method in Fieldtrip. This procedure consists of bandpass-filtering the trials of a condition block in the 110–125 Hz frequency range, which typically contains most of the muscle artifact activity, followed by a Z-transformation. Trials exceeding a threshold Z-score were removed completely from analysis. We used as the threshold the absolute value of the minimum Z-score within the block, + 1. To remove eye blink artifacts from the time courses, the EEG data from a complete session were transformed using independent component analysis (ICA), and components (typically one or two of the 48) due to blinks was removed from the data. In addition, to remove microsaccade-related artifacts we included two virtual channels based on channels Fp1 and Fp2 in the ICA, which included transient spike potentials as identified using the algorithm from (Hassler et al., 2011). The two components loading high on these virtual electrodes 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 and 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 frequency range of 36 – 100 Hz (spectral smoothing, 8 Hz; bin size, 2 Hz; five tapers). See (Kloosterman et al., 2015; Meindertsma et al., 2017) for similar settings.

Spectrograms were aligned to the onset of the stimulus train containing the (non-)target. Power modulations (denoted as M in Figure 2) during the trials were quantified as the percentage of power change at a given time point and frequency bin, relative to a baseline power value for each frequency bin. We used as a baseline the mean EEG power in the interval 0.4 to 0 s before trial onset. If this interval was not completely present in the trial due to preceding events (see Trial extraction), this period was shortened accordingly. We subtracted the trial-specific baseline value from each sample in the time course per frequency bin and divided by the mean baseline power across all trials. For the analysis of raw pre-stimulus power modulations no baseline correction was applied. We focused our analysis of EEG power modulations around target onsets on those electrodes that processed the visual stimulus. To this end, we averaged the power modulations or raw power across eleven occipito-parietal electrodes that showed stimulus-induced responses in the gamma-band range (59-100 Hz). See (Kloosterman et al., 2015) and (Meindertsma et al., 2017) for a similar procedure.

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

Empirical test of the response gain model To test the prediction of increased gain during liberal of the gain model, we first averaged activity in the 8-12 Hz range from 0.8 to 0.2 s before trial onset (staying half our window size from trial onset, to avoid mixing pre- and post-stimulus activity, also see (Iemi et al., 2017)) and took the log transform, yielding a single scalar value per trial expressing neural excitability. If this interval was not completely present in the trial due to preceding events (see Trial extraction), this period was shortened accordingly. Trials in which the scalar was > 3 standard deviations away from the participant's mean were excluded. We then sorted all single trials for each participant in ascending order of excitability and assigned them to ten equally-spaced bins ranging from the lowest to the highest excitability scalars present within that participant. Adjancent bin ranges overlapped for 50% to stabilize estimates (see (Rajagovindan and Ding, 2011) for a similar procedure). Then we averaged the corresponding log-transformed gamma

modulation of these trials (consisting of the average power within 59-100 Hz 0.2 to 0.6 s after trial onset) and normalized each participants response by subtracting the minimum gamma power during the conservative condition. Finally, we averaged across participants and plotted the excitability bin number against the normalized gamma power for each condition. To statistically test the gain prediction, we employed a three way repeated measures ANOVA (see Statistical comparisons). For plotting purposes (Figure 4D), we computed within-subject error bars by removing within each participant the mean across conditions from the estimates.

Drift diffusion modeling We fitted the drift diffusion model to our behavioural data, for each subject individualy, and separately for the liberal and conservative condition. 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 tailored version of the HDDM 0.6.0 package (Wiecki et al., 2013) (code available at Github). The RT distributions for go-choices 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 no-go response proportion contributed to *G* square. Fitting the model to RT distributions for go- and no-go choices (termed 'stimulus coding' in (Wiecki et al., 2013)), as opposed to the more common fits of correct and incorrect choice RT's (termed 'accuracy coding' in (Wiecki et al., 2013)), allowed us to estimate parameters that could have induced biases in subjects' behavior.

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

Statistical comparisons We used two-sided permutation tests (10,000 permutations) (Efron and Tibshirani, 1998) to test the significance of behavioral effects and the model fits (Figures 1, 5). 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 (Maris and Oostenveld, 2007) to correct for multiple comparisons. For time-frequency representations of power modulation (Figure 2), this procedure was conducted across all time-frequency bins. To test whether there was any evidence for increased

gain in the liberal compared to the conservative condition, we conducted a three-way repeated measures ANOVA (condition (conservative, liberal) x brain activity type (prestimulus alpha, poststimulus gamma power) x bin level (1-10)) using SPSS 23 (IBM, Inc.), inspecting linear and quadratic contrasts. As sphericity was violated in this model ($p < 0.0001$), we report both the uncorrected and Greenhouse-Geisser-corrected p-values.

Supplemental Information

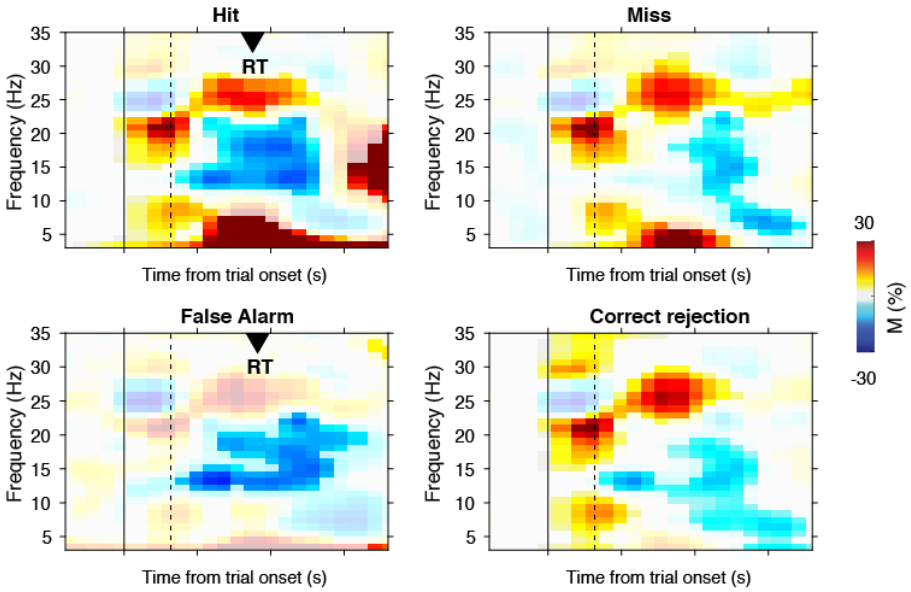

Figure S1 | Significant EEG power modulation following trial onset in all four signal detection trial categories. Time-frequency representations of low-frequency (bottom) EEG power modulations with respect to the prestimulus period (-0.4 - 0 s), pooled over the two conditions. Saturated colors indicate clusters of significant modulation, cluster threshold $p < 0.05$, two-sided permutation test across participants, cluster-corrected; N = 15). Solid and dotted vertical lines respectively indicate the onset of the trial and the target stimulus. M, power modulation.

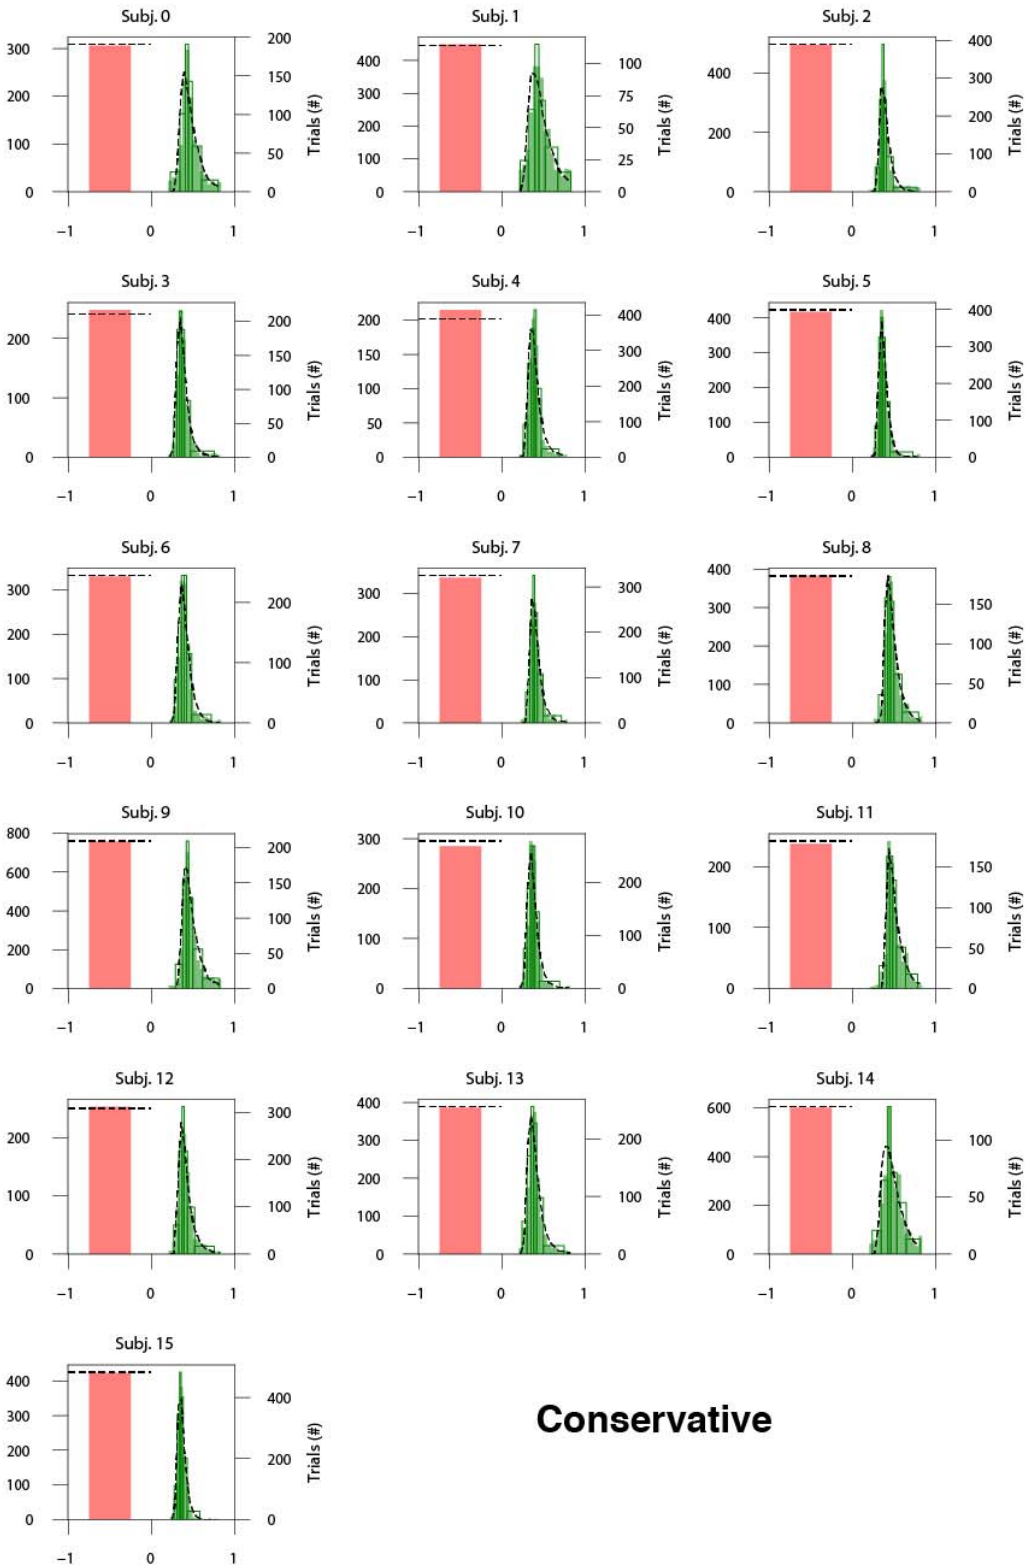

Conservative

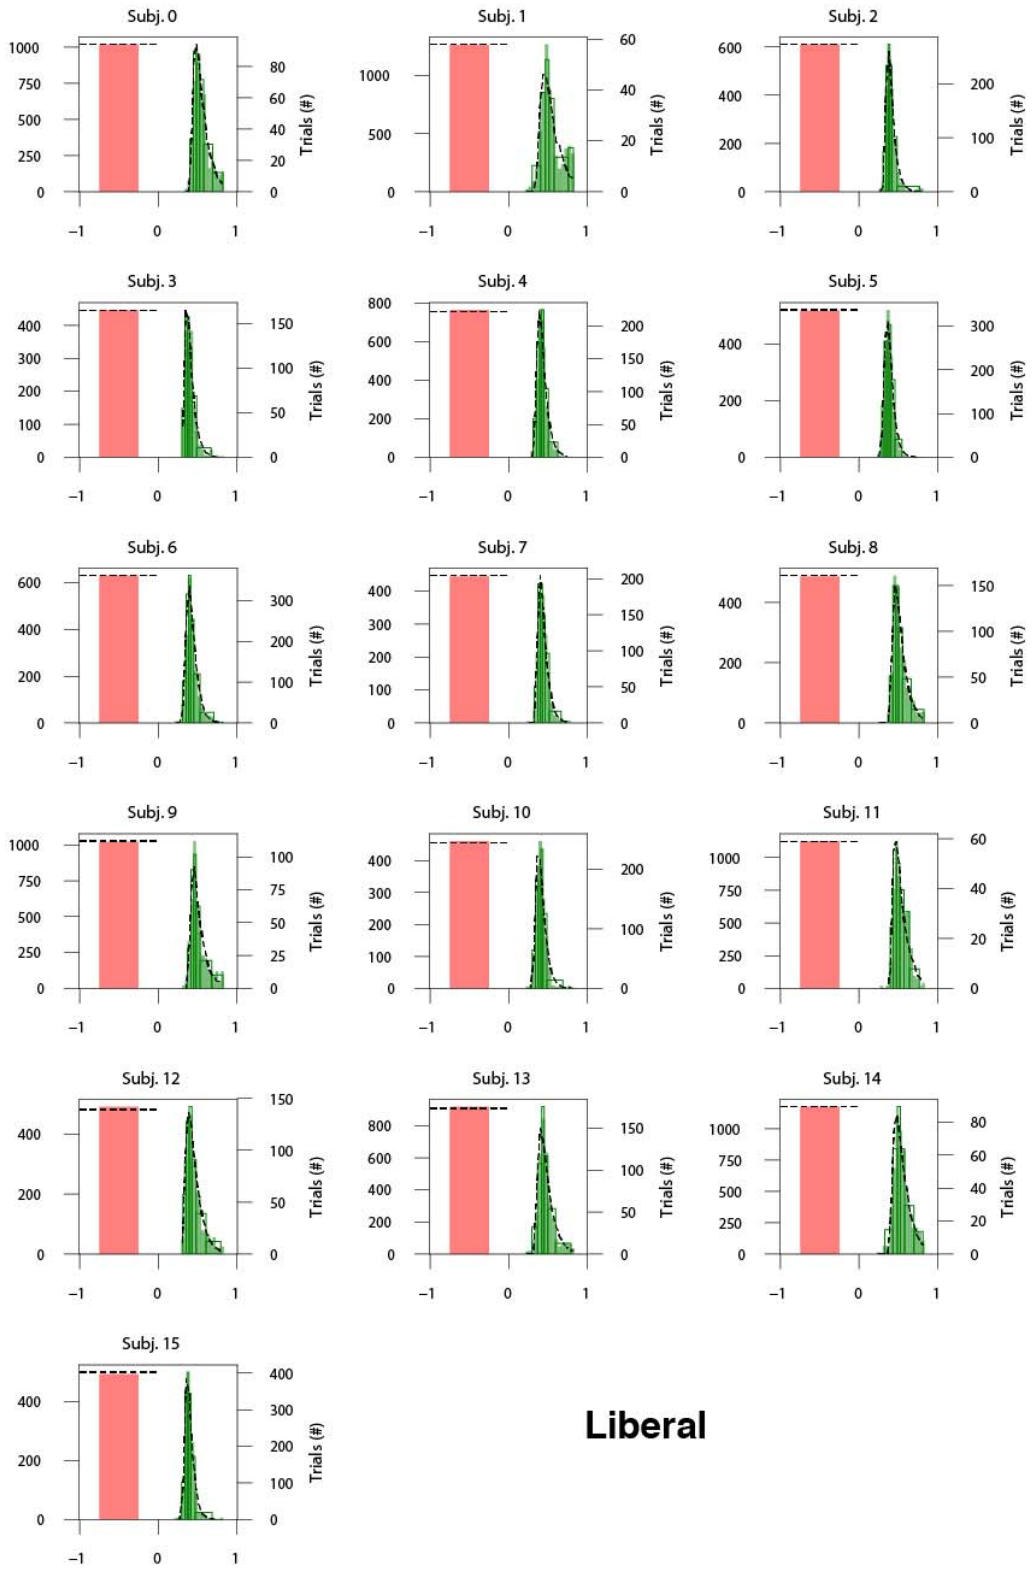

Figure S2 | Drift diffusion model fits for each subject and condition. Pink, number of No-go trials; Green, RT distribution for Go trials; dotted lines, model fits for the drift criterion model.