american physiological society

JNP JOURNAL OF NEUROPHYSIOLOGY

J Neurophysiol 127: 571–585, 2022.
First published January 26, 2022; doi:10.1152/jn.00465.2021

RESEARCH ARTICLE

Sensory Processing

Oculomotor freezing indicates conscious detection free of decision bias

Alex L. White,[1] James C. Moreland,[2] and Martin Rolfs[3]

[1]*Department of Neuroscience & Behavior, Barnard College, Columbia University, New York, New York;* [2]*Department of Psychology, University of Washington, Seattle, Washington; and* [3]*Department of Psychology, Humboldt-Universität zu Berlin, Berlin, Germany*

Abstract

The appearance of a salient stimulus rapidly and automatically inhibits saccadic eye movements. Curiously, this "oculomotor freezing" response is triggered only by stimuli that the observer reports seeing. It remains unknown, however, whether oculomotor freezing is linked to the observer's sensory experience or their decision that a stimulus was present. To dissociate between these possibilities, we manipulated decision criterion via monetary payoffs and stimulus probability in a detection task. These manipulations greatly shifted observers' decision criteria but did not affect the degree to which microsaccades were inhibited by stimulus presence. Moreover, the link between oculomotor freezing and explicit reports of stimulus presence was stronger when the criterion was conservative rather than liberal. We conclude that the sensory threshold for oculomotor freezing is independent of decision bias. Provided that conscious experience is also unaffected by such bias, oculomotor freezing is an implicit indicator of sensory awareness.

NEW & NOTEWORTHY Sometimes a visual stimulus reaches awareness, and sometimes it does not. To understand why, we need objective, bias-free measures of awareness. We discovered that a reflexive freezing of small eye movements indicates when an observer detects a stimulus. Furthermore, when we biased observers' decisions to report seeing the stimulus, the oculomotor response was unaltered. This suggests that the threshold for conscious perception is independent of the decision criterion and is revealed by oculomotor freezing.

microsaccades; oculomotor freezing; perceptual awareness; perceptual decision-making

INTRODUCTION

You can often gain insight into another person's mind by observing how they move their eyes and what they choose to look at. But even when they attempt to keep their gaze still, tiny involuntary eye movements reveal aspects of their mental state. Interspersed among slower types of fixational eye movements, involuntary microsaccades rapidly shift the gaze direction by small amounts (1, 2). Microsaccades are in many ways similar to large saccadic eye movements (3–5), and their frequency and timing are affected by other cognitive and motor processes. For instance, microsaccade rates decrease in anticipation of sensory events (6–9) and before voluntary eye and hand movements (10, 11).

A particularly striking oculomotor phenomenon is oculomotor freezing (12): saccadic eye movements are momentarily and automatically inhibited by the appearance of new stimuli (5, 13–15). Specifically, the onset of a stimulus, be it auditory, tactile, or visual, causes a transient decrease in the spontaneous microsaccade rate that lasts from roughly 100 to 400 ms poststimulus, which is followed by a brief rebound above baseline (5, 8, 13, 14, 16, 17).

We recently found that oculomotor freezing is triggered only by stimuli that the observer detects (as measured by explicit report), revealing a possible link to visual awareness (12). In those experiments, we presented brief grating stimuli (Gabor patches) on half the trials and asked the observers to report stimulus presence or absence. We developed an algorithm to convert microsaccade rates into a measure of oculomotor sensitivity (o') that can be compared to perceptual sensitivity (d'). Contrast thresholds for the two sensitivity measures were indistinguishable [consistent with contemporaneous work by others (16, 17)]. Crucially, the same physical stimulus gave rise to full-fledged oculomotor freezing when

Correspondence: A. L. White (alwhite@barnard.edu).
Submitted 18 October 2021 / Revised 4 January 2022 / Accepted 21 January 2022

it was detected but caused no change in microsaccade rates when it was missed. Moreover, microsaccades were inhibited if observers reported having seen a stimulus even if none had appeared. Because of this correlation, a Bayesian algorithm could decode from observers' eye movement patterns whether they had detected a stimulus or not. This oculomotor link to perception may provide a new tool for studies of perception in incommunicative patients, children, or nonhuman animals and for "no-report" studies of consciousness (18).

The present study answers an important question left open by all previous studies: is oculomotor freezing triggered by observers' sensory experience or by their decision that a stimulus was present? Those two phenomena can be dissociated, and understanding which one lies at the origin of oculomotor freezing is vital to its interpretation and application. We consider two hypotheses to explain the established covariation between oculomotor responses and explicit perceptual detection (12, 16, 17, 19). Both assume a classical signal detection model: on each trial, the stimulus evokes an internal response that is compared against a criterion to decide whether to produce a response or not. Even when the physical stimulus and task demands are constant, the sensory response varies across trials, but the criterion is relatively stable. The two hypotheses concern whether the criterion for oculomotor freezing is the same as the criterion for explicit perceptual decisions.

1) Shared criterion: There is a single decision criterion that determines both explicit perceptual reports and oculomotor freezing. When the sensory response exceeds the criterion, it triggers both a "yes" decision and oculomotor freezing. The shared criterion can be strategically modified, to maximize expected rewards. From a physiological perspective, this is conceivable: manipulations of stimulus probability that shift decision bias also affect activity in the superior colliculus (20), which is causally involved in controlling microsaccades (21).

2) Distinct criteria: There are distinct criteria for triggering oculomotor freezing and for deciding that a stimulus was present. The observer can strategically change their perceptual decision criterion to maximize expected rewards as conditions change. In contrast, the oculomotor criterion is unaffected by such manipulations. Thus, the two criteria can diverge, breaking the link between explicit reports and oculomotor freezing. To explain our prior results (12), this hypothesis assumes that the participants reported exactly what they perceived and set their decision criterion very near the criterion for oculomotor freezing.

Several key studies on this topic appear to support the distinct-criteria hypothesis, although they are also consistent with the shared-criterion hypothesis. Two research teams have been able to predict perceptual contrast thresholds based on microsaccade patterns that were measured while the participant did not explicitly respond to the stimuli (16, 17, 19). These studies show that oculomotor freezing is not related to response execution. However, participants in those studies either had to silently count the stimuli (16) or prepare to respond on a random subset of trials (17, 19), so they were likely making covert decisions about each stimulus. Therefore, decision-making processes could have contributed to oculomotor freezing in those data, in line with the shared-criterion hypothesis.

We designed two experiments to discriminate between the two hypotheses defined above by manipulating observers' decision criterion in a detection task. The first experiment used weighted payoffs (real money won or lost on each trial), and the second varied the expected probability that a stimulus would appear on each trial. Such manipulations shift the theoretically optimal criterion to a point that corresponds to a particular likelihood ratio, β_{opt}, of target presence to absence, and have been shown to work empirically (22–24). Our question here is whether and how these bias manipulations affect the prevalence of oculomotor freezing. To answer it, we conduct two main analyses of microsaccade rates: the first separates trials according to the physical stimulus presence, and the second additionally separates trials according to the participants' reports of stimulus presence or absence. The shared-criterion hypothesis predicts an effect of bias condition in the first analysis but not the second; the distinct-criteria hypothesis predicts the opposite.

METHODS

Both experiments were preregistered (https://osf.io/ycjgr; https://osf.io/s9myc/). Data and analysis code are available at https://osf.io/zkcag/. The Ethics Committee of the German Society for Psychology (DGPs) approved the study.

Experiment 1

Participants.

We recruited a total of 16 observers from the Humboldt-Universität zu Berlin community with normal or corrected-to-normal vision. They participated in exchange for a payment that depended on performance (details below). Of the 14 observers who completed the study (see below), 6 were male, 8 were female, and their ages ranged from 19 to 34 yr (mean 26.3). All were naive as to the research aims and gave written informed consent.

The sample size was determined by a power analysis based on the data from White and Rolfs (12). In *experiment 3* of that study, we found an effect of orientation adaptation on microsaccade rates. That effect size was modest: the maximal difference at 350 ms after stimulus was 0.2 saccades/s. Averaging over the time window when the overall inhibitory effect of stimulus presence was significant, the mean effect was 0.13 saccades/s.

We made the conservative assumption that if there is an effect of payoff condition, it is 75% as large as the effect of orientation adaptation, at each individual time point. We conducted a power analysis to determine how many participants would be necessary to find such an effect with a power of 0.8. For each possible sample size (N) between 10 and 20, we simulated 100 experiments. For each experiment, we conducted a bootstrapping analysis: in each of 1,000 repetition, we drew N observers with replacement from the original data set in White and Rolfs (12). For each observer, we computed the difference in microsaccade rate between the unadapted and the adapted condition, at each poststimulus time point, multiplied by 0.75. We then averaged those differences across the resampled participants. Over 1,000 repetitions we built up a distribution of differences at each time

point, from which we could extract a *P* value. We applied the false discovery rate (FDR) correction to determine at which time points the difference was significant. For each simulated experiment, we considered the overall effect to be significant if the difference was significant in at least 10 individual time points. For each *N*, we defined power as the proportion of experiments with a significant effect. The minimal *N* to have a power >0.8 was 14 (estimated power = 0.87).

Two participants began the study but did not finish it and were not included in the analyses. One was unable to finish all the sessions, and another discontinued after three sessions with d' far above the acceptable range because of threshold estimation failure. Thus, the final sample included 14 participants.

Apparatus and stimuli.

Observers sat in a darkened room with their head on a chin rest, 270 cm from a projection screen that displayed stimuli with a gamma-linearized PROPixx projector (VPixx Technologies; 120 Hz, 1,920 × 1,080 pixel resolution). We recorded the gaze position of both eyes at 500 Hz with a head-mounted EyeLink II system (SR Research, Ottawa, ON, Canada). Stimuli were controlled and data collected with the Psychophysics and Eyelink Toolboxes (25–27). The grayscale display (1,920 × 1,080 pixels, 120 Hz refresh rate) had 8 bits of resolution in luminance. The background luminance was set to 35% of its maximum (18.15 cd/m^2).

The fixation mark was a 4 × 4-pixel black-and-white checkerboard pattern of width 0.09 degrees of visual angle (dva) at the center of the screen. In between trials, this mark was replaced by a circle (0.27 dva radius) of alternating black and white pixels. The target stimulus was a Gabor pattern: a 0.75 cycles/dva, vertically oriented sinusoidal grating windowed by a two-dimensional Gaussian (σ = 0.67 dva). Figure 1*A* shows the stimuli in an example trial.

Procedure.

Observers began each trial by fixating on the central mark. After 0.5–2.5 s, the target Gabor stimulus flashed for 8.3 ms at the screen center. The target's onset time had a roughly flat hazard rate: on each trial, the onset time was set to 0.5 s

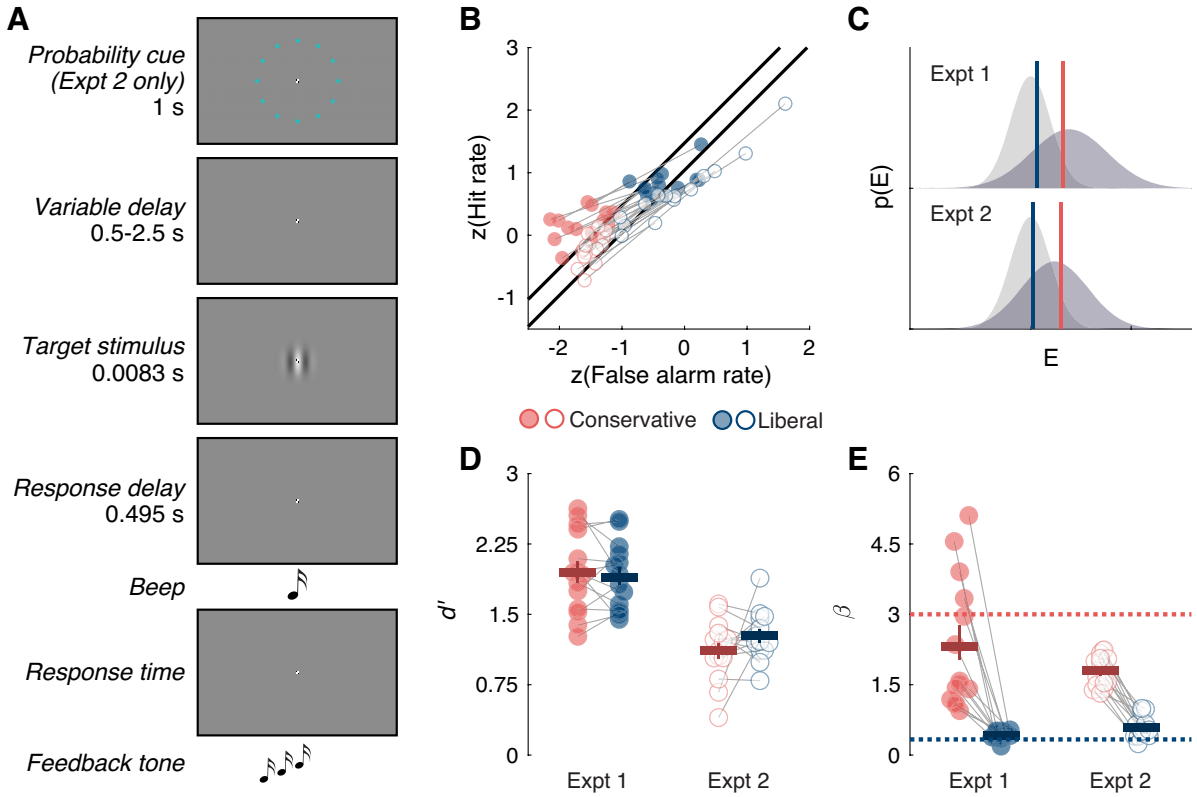

Figure 1. Trial sequence and effects of the bias manipulations on explicit perceptual reports. *A*: stimuli in the trial sequence, with time going from *top* to *bottom*. *B*: the receiver operating characteristic (ROC) showing individual *z*-transformed hit and false alarm rates. The 2 black lines with *slope 1* are the predictions of an equal-variance model for each experiment (*experiment 1* is the upper black line). The data have slopes < 1, suggesting that the distribution of sensory evidence has higher variance when the target is present rather than absent. *C*: signal detection models that account for the empirical hit and false alarm rates. These show probability distributions of sensory evidence *E* on target-absent trials (light gray) and target-present trials (dark gray). The standard deviations of the target-present distributions were derived from the average ROC slopes in *B*. The blue and red vertical lines are the mean empirical criteria (computed from false alarm rates) in the liberal and conservative conditions, respectively. *D*: individual participants' detection sensitivity d', assuming the standard deviations of sensory evidence as modeled in *B*. *Experiment 1* is in filled circles, *experiment 2* in open circles. Thin gray lines connect points from the same participant. The horizontal positions of individual data points are jittered to avoid total overlap, but points from the same participant have the same relative jitter. The horizontal lines represent the means, with error bars spanning the 68% bootstrapped confidence interval (approximately ±1SE). *E*: individual participants' decision bias β, again assuming unequal variance. Format as in *C*. Horizontal dotted lines are the optimal β for each condition (dark blue, liberal; light red, conservative).

plus a value drawn from an exponential distribution (mean = 0.65 s) clipped at 2 s. The target's phase on each trial was randomly set to either 0° or 180°. On 50% of the trials, the target had nonzero contrast (target-present trials). On the remaining trials, its contrast was set to 0, causing no change on the screen (target-absent trials). The fixation mark remained visible at the center of the Gabor (participants were required to maintain central fixation, as described in *Eye-tracking*). Four hundred ninety-five milliseconds after target offset, a beep (400 Hz, 50 ms, delivered through headphones) indicated that the trial was over.

The observer's task was to indicate whether the target was present or absent by pressing the up or down arrow, respectively, with the right hand. Response time was unlimited, but responses were not allowed before the beep. Tones delivered immediately after the response indicated whether the response was correct or incorrect and how many points were won or lost (details in *Payoff conditions*). After an intertrial interval (700 ms) containing only the circular mark at the screen center, the next trial began.

The first session began with practice and then two blocks of staircase trials to estimate the observer's contrast threshold. During the staircase blocks, the contrast was adjusted after each trial according to the single-interval adjustment matrix (SIAM) staircase procedure (28). The contrast adjustment depended on the stimulus and response: after a hit, −0.3 $\log_{10}$ units; miss, +0.3 $\log_{10}$ units; false alarm, +0.6 $\log_{10}$ units; correct rejection, no adjustment. The magnitudes of these steps were halved after the first and second staircase reversals. In each block, we interleaved two staircases, one starting at a relatively high level and the other at a low level of contrast. The block ended when both staircases underwent 10 additional reversals. The mean contrast of all but the first two reversal points provided the threshold estimate. We defined the observer's contrast threshold as the mean of four threshold estimates (2 from each of 2 blocks).

In the main experimental blocks (80 trials each), the target's contrast was set to the observer's estimated threshold. The mean stimulus contrast in included trials was 9% (ranging across individuals from 7% to 12%).

Payoff conditions.

Our main manipulation is to the reward structure for correct and incorrect responses on target-present and target-absent trials. On each trial the observer won or lost "points," which at the end of the experiment were converted to a monetary payment (1,600 points = €1). By varying payoffs, we aimed to manipulate the observer's detection criterion, that is, how much internal sensory evidence is required for the participant to report "target present" (22, 24). In the main experimental blocks, there were two payoff conditions: conservative and liberal. Additionally, a neutral condition was used in the initial staircase blocks to estimate contrast threshold. Following classic signal detection theory, we assumed that on each trial the observer bases their decision on a single value E, which is the amount of sensory evidence in favor of target presence. This model is built on Gaussian probability density distributions, which follow this equation:

$$f(x) = \frac{e^{-\frac{(x-\mu)^2}{2\sigma^2}}}{\sqrt{2\pi\sigma^2}} \quad (1)$$

The probability distribution of E on target-absent trials is $f_a(E)$, a Gaussian with $\mu = 0$ and $\sigma = 1$. The probability distribution of E on target-present trials is $f_p(E)$, a Gaussian with $\mu = d'$ and $\sigma = 1$. d' is the observer's sensitivity to the target. The observer's criterion can be expressed as c, the cutoff value of E needed to report presence. A related measure is the observer's bias, the likelihood ratio β:

$$\beta = \frac{f_p(c)}{f_a(c)} \quad (2)$$

After substituting the full Gaussian formulas for f_p and f_a, we can reduce the equation to

$$\beta = e^{cd' - \frac{d'^2}{2}} \quad (3)$$

The payoffs in each condition were set to achieve a desired optimal criterion β_{opt}: the value of β that maximizes the expected reward. The values of β_{opt} were 3 for the conservative condition, 1 for the neutral condition, and 1/3 for the liberal condition. We set the payoffs such that the optimal observer, with a d' of 1.5, would earn an average of 6.4 points per trial. Over 1,280 trials, that would yield a payment of €5.12 at our exchange rate of 1,600 points/€. By setting the target luminance contrast to detection threshold, we aimed to keep each observer's d' near 1.5. Given the average expected reward/trial (6.4 points) and the expected d', we computed the payoff matrix that would lead an ideal observer to set their criterion to the desired β_{opt}. Specifically, we computed the payoffs for target-present trials, R_p, and for target-absent trials, R_a. For each trial type j (j = p for target-present; j = a for target-absent), the reward for correct responses is R_j points and the reward for errors is $-R_j$ points.

On any given trial, there were four possible outcomes: hits or misses if a target was present or correct rejections or false alarms if there was no target. Given d' and β, we can compute the probabilities of each of those outcomes. Given R_p and R_a, we can then compute the expected reward V per trial:

$$V = P(\text{hit})R_p - P(\text{miss})R_p + P(\text{correct reject})R_a - P(\text{false alarm})R_a \quad (4)$$

Given that the prior probabilities of target presence and absence were both equal to 0.5, the optimal likelihood ratio criterion is the ratio of payoffs:

$$\beta_{opt} = \frac{R_a}{R_p} \quad (5)$$

Therefore, greater payoffs on target-absent trials should induce a conservative (higher) criterion, whereas greater payoffs on target-present trials should induce a liberal (lower) criterion. In our conservative condition (β_{opt} = 3), payoffs on target-absent trials should be three times payoffs on target-present trials. The inverse is true in the liberal condition. Working backward from the equations above, and given our desired d' and expected reward per trial (V), we computed the payoff matrix shown in Table 1.

The payoff on each trial was indicated by a feedback tone immediately after the response. These tones were composed

Table 1. *Payoff matrix*

Condition	Hit (R_p)	Miss ($-R_p$)	Correct Reject (R_a)	False Alarm ($-R_a$)
Conservative	4.9	−4.9	14.8	−14.8
Liberal	14.8	−14.8	4.9	−4.9
Neutral	11.7	−11.7	11.7	−11.7

For each condition, this table lists the number of points that can be won (positive values) or lost (negative values) for each type of response. The neutral condition was only used in the initial staircase blocks. R_a, payoff for target-absent trials; R_p, payoff for target-present trials.

of one, two, or three beeps, depending on the absolute value of the payoff (as shown in Table 1, there were 3 possible magnitudes). When there were multiple beeps, their pitches ascended in a major scale for correct responses or descended in a minor scale for incorrect responses. Each beep was separated by 20 ms of silence. In the liberal condition, for example, hits won 14.8 points and were followed by three ascending beeps, whereas false alarms cost 4.9 points and were followed by one low-pitched beep. The three beeps used for correct tones were 75 ms of 440 Hz, 80 ms of 587 Hz, and 85 ms of 659 Hz. The three beeps used for incorrect feedback tones were 75 ms of 196 Hz, 80 ms of 155 Hz, and 85 ms of 131 Hz.

The total number of points won was displayed at the end of each block. Prior to each block, instructions regarding the payoff structure were displayed on the screen. These instructions consisted of a 2 × 2 table showing the number of points that could be won or lost for reporting "Yes" or "No" depending on whether a target was present or absent. The values in this table were the same as in the corresponding condition's row in Table 1, rounded to the nearest integer. A single sentence was written above the table: in the conservative condition, "Rewards and penalties are greatest when the target is absent."; in the liberal condition, "Rewards and penalties are greatest when the target is present."

Importantly, the words "liberal" or "conservative" were never said to the participants, nor did experimenters tell them what the optimal strategy was for any given condition. However, in the first training session, the participant read a longer document of instructions that said, "In the main part of the experiment, we will vary the number of points you can win or lose depending on presence of the target and the response you make. There are two types of blocks that differ in the relative rewards and penalties on trials when the target was really present or absent. To win the most money, you should adjust how sure you need to be to say 'yes' or 'no,' depending on the points available for each type of response in the current block." When introducing the conservative condition, the instructions said: "You will win three times as many points when the target is absent and you say no than when a target is present and you say yes… and lose three times as many points when the target is absent and you say yes than when a target is present and you say no." Complementary instructions followed for the liberal condition. Observers were also instructed that they could win points and earn money during the staircase blocks as well as the main blocks.

In the first session, we informed observers that they would be paid a base hourly rate of €7/h, plus a bonus equal to the total number of points they accumulated during the trials divided by 1,600. The maximum bonus they could earn in any given hour-long session was €4. The mean bonus paid per participant for two main experimental sessions was €4.66 (range €4.13 to €5.30).

Each participant completed a total of eight blocks of each condition (for a total of 640 trials/condition). The first session began with practice, the staircase to estimate threshold, and, if time permitted, some main experimental blocks. In each subsequent session (~1 h each), the typical observer completed eight blocks: the first four of one payoff condition and the next four of the other condition. In each session, observers thus did an equal number of blocks of the two payoff conditions. The order of conditions alternated across sessions, and a random half of the observers started with the liberal condition.

Completing all 16 blocks required a total of three sessions for the typical participant (including the first staircase session). At the start of the second and third sessions, a practice block established whether the prior session's contrast threshold was still appropriate; in some cases, it was necessary to reset the contrast level for that session to keep d' near 1.5. If the overall d' in a full session (~8 blocks) was above 2.0 or below 1.0, we excluded those blocks from analysis and reran them in an extra session. This occurred when our threshold estimate was significantly inaccurate. A total of three sessions from three participants were excluded and rerun in that fashion. The reason to exclude them is that our analyses of interest depend on the target stimulus being at threshold visibility. Importantly, we always excluded and reran the same number of blocks of each payoff condition.

Eye-tracking.

At the start of each block, we performed a nine-point calibration within a central square region, 21 dva wide. Every 28 trials, we performed a standard drift correction by having the observer press a key while fixating a dot at the screen's center. If either eye's gaze position deviated >2 dva from the fixation mark between the start of a trial start and the beep, that trial was immediately terminated and repeated at the end of the block. We also detected fixation breaks offline by defining, for each trial, the fixation position as the median gaze coordinates during the first 100 ms of the trial and fixation breaks as deviations >2 dva from that. Trials with offline-detected fixation breaks were excluded from the analysis, but that only excluded an average of one trial per participant (maximum 3).

Experiment 2

Participants.

We recruited a total of 20 observers from the Humboldt-Universität zu Berlin community. All had normal or corrected-to-normal vision, participated in exchange for payment, and gave written informed consent. Of the 14 observers who completed the study and were included in the analysis (see below), 4 were male, 10 were female, and their ages ranged from 20 to 37 yr (mean 25.4).

We used the same number of participants as in *experiment 1*, but with twice as many trials per condition. The reason is that this experiment contained a condition in which the

target was half as likely to appear (and we needed to separately analyze trials with and without targets). Six participants were not included in the analysis because they discontinued participation before completing the study (in 2 cases because their d' was out of range in 1 or more completed sessions and they declined to repeat them). Thus, the final sample included 14 observers.

Procedure.

All stimuli and methods in *experiment 2* were the same as in *experiment 1*, except as noted here. Observers began each trial by fixating on the central mark. Then a probability cue appeared for 1 s. The target probability cues were formed of 12 dots (each 0.2 dva in diameter) arranged in a ring around fixation (radius 3 dva). See Fig. 1*A* for an example. The dots on each trial were all of the same color, either cyan or magenta. For half the observers, a cyan cue indicated low target probability and magenta indicated high target probability. For the other half of observers, the colors were reversed. Then, after a variable delay of 0.5–2.5 s, the target Gabor stimulus flashed for 1 frame (8.3 ms) at the screen center, and the trial proceeded as in *experiment 1*. The mean stimulus contrast in included trials was 6% (ranging across individuals from 5% to 9%).

Feedback and rewards.

The feedback and reward structures were matched to the "neutral" condition in *experiment 1* (used in the staircase blocks). The participants won 11.7 points on correct trials (hits or correct rejections) and lost 11.7 points on incorrect trials (misses or false alarms). The feedback tones were two ascending beeps or two descending beeps.

Probability conditions.

Our main manipulation was the probability of a target being present on each trial (P_T). In "low-probability" trials P_T = 0.25, and on "high probability" trials P_T = 0.75. Those trials were randomly intermixed, because if they were in separate blocks there could be hysteresis effects due to different amounts of stimulation in each block. The cyan or magenta precue indicated the target probability condition at the start of each trial.

Given the average expected reward/trial (6.4 points) and the expected d' (1.5), we computed the target probabilities that would lead an ideal observer to set their criterion to the desired β_{opt}. Using the expected reward on each trial (*Eq. 4*), we can compute β_{opt} from the ratio of payoffs, scaled by the ratio of the probability of no target and the probability of a target:

$$\beta_{opt} = \frac{R_a}{R_p}\frac{(1-P_T)}{P_T} \quad (6)$$

See Swets et al. (24) for an equivalent derivation. In *experiment 2*, $R_a = R_p$ = 11.7 points. Therefore,

$$\beta_{opt} = \frac{(1-P_T)}{P_T} \quad (7)$$

In the low-probability condition, P_T = 0.25 and β_{opt} = 3, the same as in the conservative payoff condition of *experiment 1*. In the high-probability condition, P_T = 0.75 and β_{opt} = 1/3, the same as in the liberal payoff condition of *experiment 1*. We therefore label the low-probability condition as the conservative condition and the high-probability condition as the liberal condition.

At the start of the experiment, the observer was instructed to pay attention to the colored probability cues and was told their exact meaning. We did not tell the observers how to use the cues, but we did tell them, "If you pay attention to the colored dots and adjust your responses accordingly, you could gain roughly 20% more money than if you ignore them!". Prior to each block, we displayed a reminder about what the probability cues mean.

Each participant completed a total of 32 blocks of the experiment (80 trials per block, for a total of 2,560 trials, 1,280 per condition). Completing all 32 blocks required a total of five or six sessions for the typical participant (including the first staircase session). The mean bonus paid per participant for all the main experimental sessions was €10.67 (range €7.61 to €13.64).

As in *experiment 1*, we excluded and reran sessions with d' above 2.0 or below 1.0. That occurred for a total of five sessions, one per each of five observers. On average, <0.1% of trials were excluded for offline fixation breaks (max 0.3%).

Analyses

Perceptual data analysis.

We excluded trials with reaction times >4 standard deviations (SDs) above the observer's median. Across participants, this criterion excluded an average of 1% of trials in *experiment 1* (maximum 1.6%) and an average of 0.7% in *experiment 2* (maximum 1.4%). We then computed perceptual sensitivity in each condition using the observer's hit rate (HR, the proportion of "yes" responses on target-present trials) and false alarm rate (FR, the proportion of "yes" responses on target-absent trials):

$$d' = z(\mathrm{HR}) - z(\mathrm{FR}) \quad (8)$$

where z is the inverse of the normal cumulative distribution function. To avoid undefined d' values, HR and FR were not allowed to fall below $1/(2N)$ or to exceed $[1 - 1/(2N)]$, where N is the number of target-present or -absent trials. For example, if the hit rate was 1, we assumed that, had we run twice as many trials, there would have been 1 miss. We also report the observer's criterion

$$c = z(1 - \mathrm{FR}) \quad (9)$$

From that, we compute the bias β, the likelihood ratio, using *Eq. 3* defined above.

To evaluate the effect of payoff condition on these perceptual measures (d' and β), we used bootstrapping to estimate 95% confidence intervals (CIs) between pairs of conditions. A difference is deemed significant if the 95% CI excludes 0 (a 2-tailed test).

Microsaccade detection.

The trial exclusion criteria applied in the perceptual data analysis (see above) also applied to the eye movement analysis. Our analysis of eye movement traces followed the procedure reported in our previous paper (12). We first transformed the raw gaze positions into velocities (dva/s) and smoothed them by averaging over neighboring pairs of two samples. Then, we identified microsaccadic events as shifts

in gaze position with two-dimensional (2-D) velocities that exceed, for at least three samples, an ellipse with horizontal and vertical radii equal to 5 times the horizontal and vertical median-based standard deviations, respectively (29). However, for six observers in *experiment 1* and three in *experiment 2*, the fixed threshold of 5 SDs yields very few microsaccades, so we lowered the threshold to 4.

Monocular microsaccadic events <10 ms apart were merged. We defined binocular microsaccades as those with at least one sample of overlap between the two eyes and, again, merged binocular microsaccades <10 ms apart. We defined microsaccade onset as the time at which the first of the two eye velocities exceeded the threshold and offset as the time point just before the last eye's velocity dropped below threshold. Other parameters (e.g., amplitude) were averaged over the two eyes. We included in the analysis only binocular microsaccades with durations $\geq$ 6 ms, amplitudes $\leq$ 1 dva, and peak velocities $\leq$ 250 dva/s. The "main sequence" plots, showing microsaccade amplitude versus peak velocity, are shown in Supplemental Fig. S1 (all Supplemental Figures are available at https://doi.org/10.17605/OSF.IO/T9BY7). (Note: the pupil size did not measurably change throughout the time period 0–500 ms after stimulus onset, in all conditions. Thus, microsaccade detection was unlikely to have been influenced by sudden changes in pupil size).

Microsaccade rate analysis.

We then determined the time-varying microsaccade rate for each experimental condition with a smoothing procedure. First, we counted the number of microsaccades detected at each millisecond τ relative to target onset, across all trials in each condition. Then, for each time point t, we computed a weighted sum of microsaccades in the local interval, using a "causal" kernel:

$$\omega(\tau) = \alpha^2 \tau e^{-\alpha\tau} \quad (10)$$

ω describes the weight given to microsaccades τ ms before time point t. We shifted the filter by $1/\alpha$ ms to avoid a temporal bias and give the most weight to microsaccades at point t (5, 30). The parameter α was set to 1/25. The smoothed rate r (t) is the weighted sum of microsaccades divided by the total number of trials in the sample and converted into units of saccades per second by multiplying by 1,000. Microsaccade rates were computed from −350 to +500 ms relative to target onset.

To estimate the statistical significance of changes in microsaccade rates at each time point, we bootstrapped them by simulating 1,000 repetitions of the experiment (31). On each repetition, we resampled with replacement from the set of observers, then took the mean between conditions. That gave us distributions of differences at each time point. The two-tailed bootstrapped P value is defined as twice the proportion of differences that fell below 0. When evaluating differences at many time points, we applied the false discovery rate correction (32). Two conditions are deemed significantly different if the 95% confidence interval of differences does not include 0 (corrected $P < 0.05$). (Note: this bootstrapping procedure differs from what we preregistered in that it is simpler and focuses on variability across observers rather than variability across trials within each observer, thus being a nonparametric analog of a t test). We used this bootstrapping to identify the time window when the oculomotor freezing effect is significant.

To directly compare changes in microsaccade rate to perceptual sensitivity, we computed an analogous estimate of oculomotor sensitivity (12). At each millisecond, the lack of a microsaccade following a stimulus is a "hit," and the lack of a microsaccade following no stimulus is a "false alarm." From the resulting oculomotor hit rate (HR) and false alarm rates (FRs), we can compute oculomotor d'_o at each time point t relative to stimulus onset ($0 \leq t \leq 500$):

$$d'_o(t) = z[\mathrm{HR}(t)] - z[\mathrm{FR}(t)] \quad (11)$$

Like perceptual d', this measure requires correction if HR or FR reaches extreme values. This can happen if no microsaccade were detected during a period around t as wide as the base B of the filter (~200 ms). Therefore, both rates were not allowed to fall below $1/(2NB)$ or to exceed $[1 - 1/(2NB)]$, where N is the number of target-present or -absent trials, respectively. That is, we assume that had we run twice as many trials, we would have found at least 1 microsaccade (a "miss") in the 200-ms time window surrounding any given time point. Nonetheless, because microsaccades occur only about once or twice every second, both HR and FR at individual (millisecond) time points were high (above 0.999). But because HR rose even higher than FR after stimulus presentation, we found positive values of d'_o.

To extract a single oculomotor sensitivity measure from an entire rate time course for a given condition, we defined a value o', the maximum of the cumulative sum of d'_o values across time (within 200–550 ms after stimulus). o' is unaffected by rate rebounds following inhibition, which result in negative d'_o. Pairwise differences in o' (across payoff conditions) were tested with bootstrapping, similar to perceptual d' as described above.

In addition to the preregistered analyses reported thus far, we conducted two exploratory analyses. First, to simplify the comparison of microsaccade rates across conditions (without relying on hundreds of noisy tests at many individual time points), we computed the microsaccade rates integrated across two time windows: for the baseline microsaccade rate on target-absent trials, we used the time window 0 to 500 ms. For target-present trials, we used the time window within which the microsaccade rate on target-present trials was significantly lower than the rate on target-absent trials, according to the bootstrapping procedure described above, for both bias conditions (bootstrapped FDR-corrected $P < 0.05$). This is the time window of significant oculomotor freezing (see RESULTS).

Second, compared with our previous studies, we found that baseline microsaccade rates were lower on average and more variable across observers, which complicates comparing rates by taking simple differences (liberal-conservative) between conditions. We therefore computed modulation indexes that are more robust to variation across observers in overall microsaccade rates: $(A - B)/(A + B)$, where A and B refer to a measure in specific conditions (e.g., microsaccade rate on conservative and liberal trials or report-present and report-absent trials). This index ranges from −1 to 1, where positive values indicate higher microsaccades rates in A compared with B and negative values indicate the opposite. In some cases, we did statistical tests [bootstrapping and

Bayes factors (BFs), as described below] to assess whether the mean modulation index is significantly different from 0. In other cases, we test whether the modulation indexes differ across conditions (e.g., whether the effect of perceptual report on mean microsaccade rates differs between the liberal and conservative conditions).

Finally, we supplement our pairwise tests with Bayes factors (BFs), which quantify strength of evidence. In this context, a BF is the ratio of the probability of the data under the alternate hypothesis (that 2 conditions differ) relative to the probability of the data under the null hypothesis (that there is no difference) (33, 34). As an example, a BF of 10 indicates that the data are 10 times more likely under the alternate hypothesis than the null hypothesis. Typically, BFs between 1 and 3 are regarded as weak evidence for the alternate hypothesis, BFs between 3 and 10 as substantial evidence, and BFs between 10 and 100 as strong evidence (35). Conversely, BFs between 1/3 and 1/10 are considered substantial evidence for the null hypothesis, etc. We computed BFs for pairwise *t* tests and two-way repeated-measures ANOVAs using the bayesFactor toolbox by Bart Krekelberg (https://doi.org/10.5281/zenodo.4394422).

RESULTS

Explicit Perceptual Reports: Bias Manipulations Affect Decision Criteria but Not Sensitivity

On each trial, observers reported the presence or absence of a brief Gabor stimulus with a luminance contrast that had been set to their individual detection threshold (Fig. 1A). The time of the target's onset was unpredictable, but the end of each trial was indicated by a beep 500 ms after the time of (potential) target appearance. The observers' goal was to win "points" that were converted to bonus monetary payments. Correct responses (hits and correct rejections) gained points, and incorrect responses (misses and false alarms) lost points.

In *experiment 1*, we introduced asymmetric monetary payoffs to manipulate decision bias. In the liberal condition, rewards were three times greater for hits than for correct rejections, and penalties were three times greater for misses than for false alarms. This reward structure places the optimal criterion at the level of sensory evidence that is three times as likely to be observed when the target is absent than when it is present. Thus, the optimal likelihood ratio β_{opt} is 1/3. In the conservative condition, rewards were three times greater for correct rejections than for hits, and penalties were three times greater for false alarms than for misses. That makes β_{opt} = 3. The reward structure varied across blocks of trials and was known to the participant in advance. Feedback at the end of each trial indicated the reward magnitude.

In *experiment 2*, we manipulated the probability that a target would appear and informed observers of that probability on each trial. In the liberal condition, there was a 75% chance that a target would appear (3 times likelier to be present than absent), which lowered the optimal criterion such that β_{opt} = 1/3. In the conservative condition, there was a 25% chance that a target would appear, raising the optimal criterion such that β_{opt} = 3 (as in *experiment 1*). These trial types were randomly intermingled within blocks, but a cue in the form of colored dots presented at the start of each trial informed the participant of the target probability. Payoffs on target-present and target-absent trials were of equal magnitude.

In both experiments, the bias manipulation strongly affected explicit perceptual reports of target presence. The mean hit and false alarm rates, their mean differences between bias conditions, along with the 95% confidence interval (CI) of those differences and Bayes factors, are listed in Table 2. Hit rates and false alarm rates were much lower in the conservative than the liberal condition, indicating that participants were less willing to report seeing a target when the potential payoffs were greater on target-absent trials (*experiment 1*) and when target presence was unlikely (*experiment 2*). The mean response time (RT) for all target-present conditions was at least 930 ms after stimulus onset (430 ms after the beep prompting response), well after the period in which we assess microsaccades, and not affected by bias condition. RTs on target-absent trials tended to be slower, especially in the liberal condition. RTs are plotted in Supplemental Fig. S2.

To interpret these psychophysical data, we adopt the classic signal detection model: the participant reports target presence if the magnitude of sensory evidence *E* exceeds a criterion level *c*. The variances of *E* on target-absent and target-present trials are often unequal and can be estimated with a receiver operating characteristic (ROC) graph (24).

Table 2. *Explicit reports in each condition of each experiment*

	Conservative	Liberal	Diff	Diff 95% CI	BF
Hit rate					
Expt 1	0.57 (0.02)	0.79 (0.01)	0.22 (0.03)	[0.18 0.27]	22,618
Expt 2	0.41 (0.02)	0.73 (0.04)	0.32 (0.05)	[0.22 0.40]	822
False alarm rate					
Expt 1	0.06 (0.01)	0.38 (0.04)	0.31 (0.04)	[0.24 0.40]	5,933
Expt 2	0.07 (0.01)	0.45 (0.07)	0.38 (0.07)	[0.27 0.52]	378
d'					
Expt 1	1.95 (0.12)	1.90 (0.10)	−0.05 (0.08)	[−0.22 0.09]	0.32
Expt 2	1.12 (0.09)	1.27 (0.07)	0.15 (0.11)	[−0.02 0.40]	0.56
β					
Expt 1	2.32 (0.38)	0.42 (0.02)	−1.89 (0.39)	[−2.72 −1.29]	114
Expt 2	1.80 (0.09)	0.59 (0.06)	−1.21 (0.11)	[−1.39 −0.99]	252,608

Conservative and liberal columns list the across-subject mean values with SE in parentheses. Diff is the average (and SE) difference liberal − conservative. The 95% bootstrapped confidence interval (CI) of the difference is also shown. When a CI excludes 0, we conclude that there is a significant effect of the bias condition. BF, Bayes factor for the comparison of liberal vs. conservative conditions. *d'* and β, sensitivity and bias measures assuming unequal variance of sensory evidence on target-present and target-absent trials (see text).

The ROC graph in Fig. 1*B* plots false alarm rates versus hit rates, each *z*-transformed through the inverse normal cumulative distribution function. For each participant, one line connects their points for the liberal (blue) and conservative (red) conditions. If the distributions of sensory evidence have equal variance, then these lines should have slopes equal to 1 (illustrated with thick diagonal black lines in Fig. 1*B*). The empirical slopes are consistently shallower: in *experiment 1* the mean slope was 0.53 (95% CI = [0.46 0.60]), and in *experiment 2* it was 0.60 (95% CI = [0.50 0.69]). Assuming that the target-absent distributions have standard deviations (SDs) equal to 1, the SDs of the target-present distributions are equal to the inverse of the ROC slopes: 1.90 in *experiment 1* and 1.66 in *experiment 2*. These best-fitting signal detection models are shown in Fig. 1*C*, with the mean criteria (computed directly from false alarm rates) as vertical blue and red lines.

Using these estimated variances, we computed d', a measure of sensitivity (Fig. 1*D*), and β, a measure of bias (Fig. 1*E*). d' is the distance between the mean E (sensory evidence) on target-present trials and the mean E on target-absent trials. β is the likelihood ratio of target presence to target absence when $E = c$. Using the formulas for β and d' (*Eqs. 3* and *8*) that typically assume equal variance, we substituted the best-fitting SDs into the probability and cumulative density functions. Statistics for both measures are reported in Table 2. d' did not significantly differ between the liberal and conservative conditions (CIs include 0), but β was significantly higher in the conservative condition, for all participants. The dotted lines in Fig. 1*E* are the optimal β_{opt} in each condition. Most participants did not shift their criteria quite far enough to reach the optimal levels (36). For the estimates of d' and β that (incorrectly) assume equal variance on target-present and -absent trials, see Supplemental Fig. S3.

In sum, both bias manipulations had large effects on decision criteria for explicit judgments, whereas sensitivity remained unaffected.

Microsaccade Rates Contingent on Physical Target Presence: Bias Manipulations Do Not Affect Oculomotor Freezing

Figure 2, *A* and *B*, show the mean microsaccade rates plotted as a function of time relative to target onset. The target, when present, was flashed at *time point 0*. In both experiments we observed oculomotor freezing on target-present trials (solid lines): the microsaccade rate begins to drop roughly 130–150 ms after stimulus onset and then returns to baseline 300–400 ms later. The key question is whether microsaccade rates differ between the liberal and conservative bias conditions. The distinct-criteria hypothesis predicts no difference. The shared-criterion hypothesis, which posits that oculomotor freezing is linked to explicit report

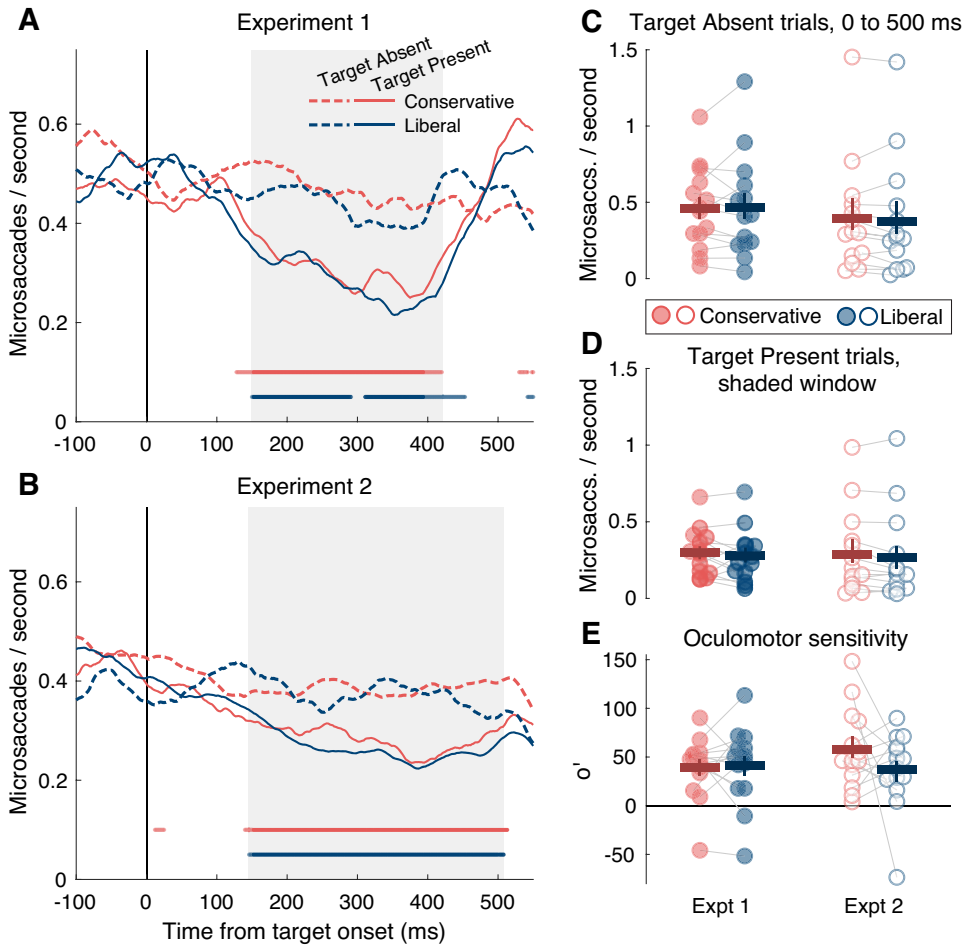

Figure 2. Bias manipulations do not affect overall microsaccade rates on target-present and target-absent trials. *A* and *B*: mean microsaccade rates as a function of time relative to target onset in *experiments 1* (*A*) and *2* (*B*), for target-absent trials (dashed lines) and target-present trials (solid lines). The horizontal lines at *bottom* of each plot indicate time points when the rate on target-present trials is significantly different from the rate on target-absent trials (corrected $P < 0.05$). The gray region of the background indicates the time window when the rate was significantly reduced on target-present trials in both conditions. *C*: mean microsaccade rates on target-absent trials in the time window between 0 and 500 ms. Format as in Fig. 1*B*. *D*: mean microsaccade rates on target-present trials in the time windows with significant inhibition in both conditions (shaded portions in *A* and *B*). There are no significant effects of bias condition. *E*: oculomotor sensitivity (o'), a measure of the difference in microsaccade rates between target-present and target-absent trials over the entire interval 0 to 500 ms. There are no significant effects of bias condition.

decisions, predicts a greater drop in microsaccade rates on target-present trials of the liberal condition, in which the participant reports "present" more often. The data do not support the shared-criterion hypothesis. Although the mean rate in the liberal condition (blue line) dips slightly lower than in the conservative condition (red line), that effect is small and not consistent across participants.

To simplify this analysis and maximize power, we integrated microsaccades over two key time windows: 0 to 500 ms for target-absent trials and the window of significant oculomotor freezing for target-present trials (shaded windows in Fig. 2, *A* and *B*; see METHODS). In *experiment 1* the window of significant freezing was from 149 ms to 421 ms, and in *experiment 2* it was from 145 ms to 509 ms. As shown in Fig. 2, *C* and *D*, there were no reliable effects of bias condition on the mean microsaccade rates in these time windows. We evaluated the effects both as mean differences (L − C, where L is the rate on liberal trials and C on conservative trials) and as modulation indexes [(L − C)/(L + C)] that adjust for individual differences in overall microsaccade rate. The mean microsaccade rates in each condition are listed in Table 3, along with statistics on the modulation indexes. With one exception, none of those effects was significant: 95% CIs include 0, and Bayes factors (BFs) support the null hypothesis at least 2:1 (BFs < 0.5). The one exception is for target-absent trials in *experiment 2*: when the effect is expressed as a modulation index, the baseline microsaccade rate was slightly but significantly lower on liberal than conservative trials (BF = 1.35). The mean difference (L − C) was only −0.02 saccades/s (95% CI = [−0.04 0.03]; BF = 0.37).

To combine across experiments, we entered these data into linear mixed-effects models (LMEs), with fixed effects for condition, experiment, and their interaction as well as random effects for participant. We fit one such model for the target-absent trials and another for the target-present trials. The effect of condition was negligible in both analyses (0.006 and 0.02 saccades/s, respectively) and not significant (for target-absent trials, $P = 0.75$, BF = 0.16; for target-present trials, $P = 0.14$, BF = 0.45). There were no effects of experiment or interaction between experiment and condition (all $P > 0.5$).

We also computed oculomotor sensitivity (o') as a measure of the strength of oculomotor freezing (12) (Fig. 2*E*), comparable to d'. In both experiments, o' did not differ significantly between bias conditions: 95% CIs were far from excluding 0, and Bayes factors supported the null hypothesis (see Table 3). An LME combining across experiments found no effect of condition ($P = 0.32$, 95% CI = [−4.7 14.12]; BF = 0.32) and no main effect of experiment or interaction (both $P > 0.2$).

Altogether, the microsaccade rates in this first analysis are consistent with the distinct-criteria hypothesis: oculomotor freezing is independent of bias manipulations that affect explicit perceptual reports. Next, we sorted the data further by the participant's report on each trial. Based on our prior study (12), we predicted more oculomotor freezing on trials when the participant reports seeing a stimulus than when they do not, but the magnitude of that effect may depend on the bias condition.

Microsaccade Rates Contingent on Explicit Perceptual Reports: Oculomotor Freezing Is Stronger in Conservative than Liberal Bias Conditions

When we analyze trials separately according to whether the participant reported target presence or absence, the shared-criterion hypothesis predicts no effect of bias condition. The observer's ultimate decision is the same on liberal hit trials as on conservative hit trials, so the prevalence of oculomotor freezing should be the same. In contrast, the distinct-criteria hypothesis predicts an effect of bias condition: when considering only trials in which the observer reports target presence (hits and false alarms), microsaccade rates should be lower in the conservative condition than the liberal condition. This is because in the conservative condition the sensory evidence must be stronger for the participant to report presence, and therefore it is also likely to trigger oculomotor freezing. In the liberal condition, some explicit reports of target presence are guesses with low sensory evidence, which will not exceed the criterion for oculomotor freezing, so microsaccade rates should be higher.

Figure 3*A* plots the mean microsaccade rates on target-absent trials, separated by bias condition and the participant's explicit report of whether a target was present or not (correct reject trials in dark lines, false alarm trials in bright lines). In a prior study (12), we found that microsaccade rates were lower on false alarm than correct reject trials, consistent with the notion that a spurious sensory signal triggered both an explicit false alarm and oculomotor freezing. The distinct-criteria hypothesis predicts that effect (the relative inhibition of microsaccades on false alarm trials) should be weakened in the liberal condition, when many false alarms

Table 3. *Effects of bias condition on microsaccade rates in key time windows of target-absent trials and target-present trial and the effect on oculomotor sensitivity o′*

	Conservative	Liberal	Modulation Index	Index 95% CI	BF
Target-absent microsaccade rate					
Expt 1	0.46 (0.07)	0.46 (0.09)	−0.04 (0.04)	[−0.11 0.03]	0.41
Expt 2	0.39 (0.10)	0.38 (0.10)	−0.09 (0.04)	[−0.17 −0.02]	1.35
Target-present microsaccade rate					
Expt 1	0.29 (0.04)	0.27 (0.05)	−0.07 (0.06)	[−0.17 0.07]	0.45
Expt 2	0.28 (0.07)	0.26 (0.08)	−0.06 (0.05)	[−0.16 0.04]	0.47
Oculomotor sensitivity (o')					
Expt 1	39.7 (8.4)	41.51 (10.5)	−0.01 (0.12)	[−0.27 0.20]	0.27
Expt 2	57.6 (10.9)	36.91 (10.5)	−0.07 (0.14)	[−0.36 0.20]	0.30

Effects of bias condition on microsaccade rates in key time windows of target-absent trials and target-present trials and the effect on oculomotor sensitivity (o'). For each measure, we report the mean (and SD) in each bias condition, followed by statistics for the effect of bias condition expressed as a modulation index. BF, the Bayes factor comparing the modulation index to 0; CI, confidence interval.

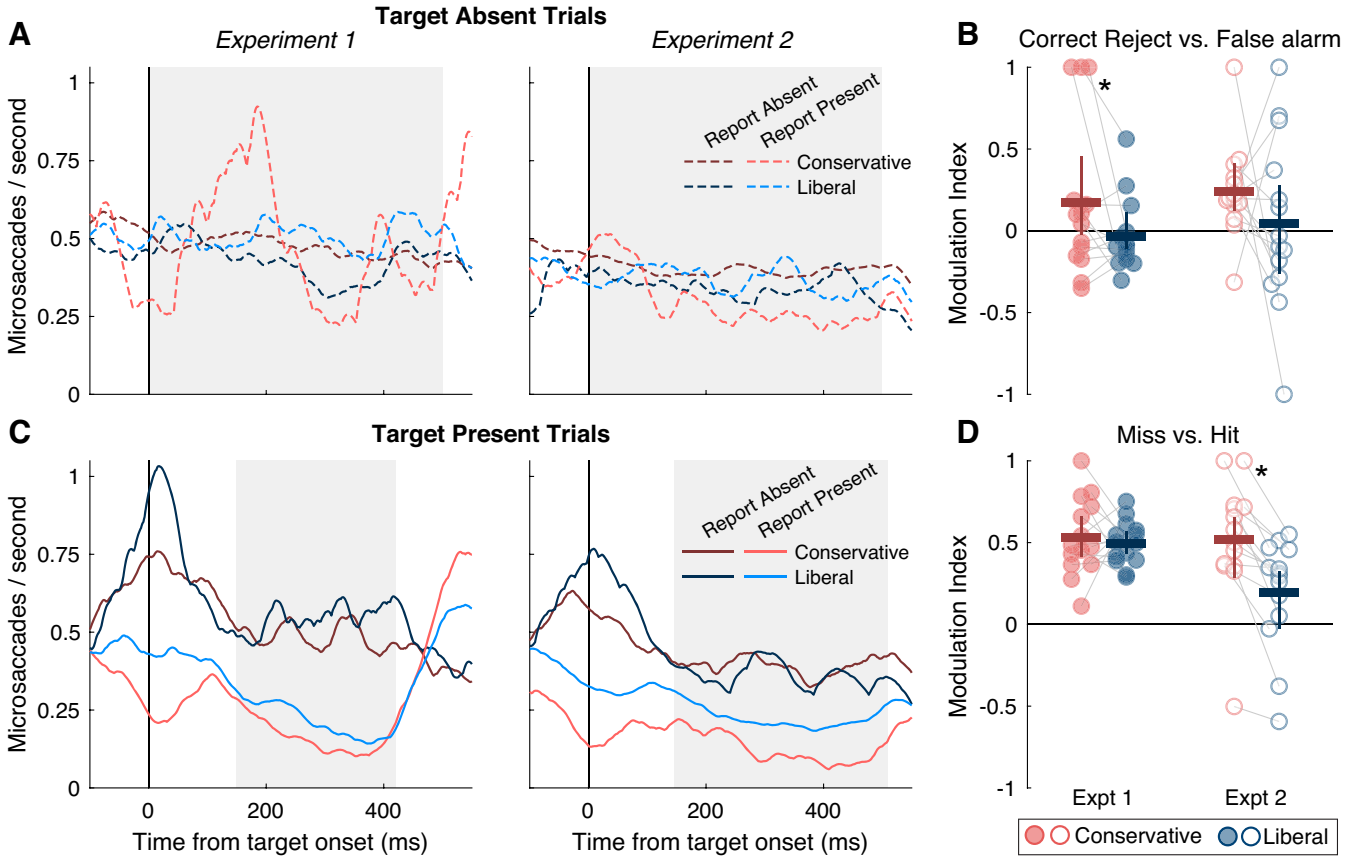

Figure 3. Microsaccade rate signatures as a function of bias condition and explicit report outcome. *A*: mean rates as a function of time on target-absent trials, separated by bias condition and by whether the participant reported target absent (correct reject trials, dark lines) or target present (false alarm trials, bright lines). Note there are very few false alarm trials in the conservative condition (bright red lines). The gray shading indicates the time window of 0–500 ms after stimulus onset (up until the beep) within which microsaccade rates were integrated for the analysis in *B*. *B*: the mean modulation indexes comparing microsaccade rates on correct reject trials and false alarm trials, integrated over 0 to 500 ms. Format as in Fig. 2, *C–E*, except that the error bars are bootstrapped 95% confidence intervals, to illustrate whether the mean modulation index in each condition deviates significantly from 0. The overall effect of perceptual report is significant and marginally higher on conservative than liberal trials. The black asterisk indicates that the mean modulation index is significantly greater in the conservative condition than in the liberal condition. *C*: mean microsaccade rates on target-present trials, separated by bias condition and by whether the participant reported target absent (miss trials, dark lines) or reported target present (hit trials, bright lines). The shaded areas are the intervals with significant stimulus-induced inhibition, the same as in Fig. 2. *D*: the mean indexes comparing microsaccade rates on miss and hit trials, integrated over the intervals with significant stimulus-induced inhibition (shaded in *C*). The 95% error bars exclude 0 in all conditions except the liberal condition of *experiment 2*. The black asterisk indicates that the modulation index is significantly larger in the conservative than the liberal condition.

are guesses without a sensory signal strong enough to inhibit microsaccades.

To test these predictions, we integrated microsaccade rates over the 0 to 500 ms time window (shaded region in Fig. 3*A*) and then computed the effect of explicit report as a modulation index: (CR − FR)/(CR + FR), where CR is the microsaccade rate on correct reject trials and FR is the microsaccade rate on false alarm trials. The occurrence of oculomotor freezing on false alarm trials predicts a positive index. In addition, the distinct-criteria hypothesis predicts a larger index in the conservative compared with the liberal condition. The mean indexes are plotted in Fig. 3*B* and listed in Table 4 with 95% CIs and BFs. Only in the conservative condition of *experiment 2* was the effect of report significant (as shown in Fig. 3*B*, the error bar does not overlap with 0; BF = 7.5). According to a linear mixed-effects model that combined experiments, there was a small but significant difference between microsaccade rates on correct reject and false alarm trials (mean index = 0.11, CI = [0.004 0.209], $P = 0.04$, BF = 1.4), a marginal effect of bias condition (index 0.2 larger in the conservative condition, CI = [0.002 0.41], $P = 0.053$; BF = 1.28), and no effect of experiment or interaction (both $P > 0.4$, BF < 0.25). All told, the data in Fig. 3*B* are consistent with our previous finding that false alarms are associated with inhibition of microsaccades and are consistent with the distinct-criteria hypothesis. However, this analysis is limited because of the small number of false alarm trials in the conservative condition (on average across participants, only 20 trials in *experiment 1* and 70 in *experiment 2*). The target-present trials provide supporting evidence.

Figure 3*C* plots microsaccade rates on target-present trials. These traces diverge around stimulus onset (i.e., 0 ms) because of the reductive effect of microsaccades on perceptual sensitivity (1, 12, 37, 38): a microsaccade that occurs close in time to the stimulus onset can make the participant miss the stimulus. Indeed, on roughly 2% of all trials a

Table 4. *Effects of perceptual report on microsaccade rates, expressed as modulation indexes*

		Correct Reject vs. False Alarm			Miss vs. Hit		
	Condition	Modulation Index	Index 95% CI	BF	Modulation Index	Index 95% CI	BF
Expt. 1	Conservative	0.17 (0.13)	[−0.039 0.437]	0.58	0.53 (0.06)	[0.429 0.657]	1.4×10^4
	Liberal	−0.03 (0.06)	[−0.121 0.116]	0.31	0.50 (0.04)	[0.433 0.566]	2.6×10^6
Expt. 2	Conservative	0.24 (0.08)	[0.128 0.403]	7.47	0.52 (0.10)	[0.248 0.671]	197.6
	Liberal	0.04 (0.14)	[−0.226 0.319]	0.27	0.20 (0.09)	[−0.008 0.333]	1.68

Modulation index values are means with SE across participants in parentheses. The 95% confidence intervals (CIs) are derived from bootstrapping. BF, Bayes factor.

microsaccade occurred within 25 ms of target onset time, and on those trials participants were far less likely to detect the stimulus (Supplemental Fig. S4c). Thus, miss trials are associated with a peak in the microsaccade rate near *time 0*. That peak is especially large in the liberal condition, when misses are less frequent and require a definite lack of target evidence. Conversely, hits are associated with fewer microsaccades near the time of stimulus onset, and thus there is a dip in microsaccade rate on hit trials. That dip is larger in the conservative condition, when hits require high certainty and would otherwise be turned to misses by microsaccades. To confirm that the drop in microsaccade rates on hit trials ~150–400 ms after stimulus is not an artifact of the divergent dips and peaks observed around 0 ms because of saccadic suppression of perception, we reanalyzed all our data by excluding the ~2% of trials with microsaccades that began <25 ms before or after *time 0*. All the results concerning oculomotor freezing in the key time windows (~150–400 ms after target onset) were confirmed in this analysis. For details, see Supplemental Fig. S4.

Our present research question focuses on the later time period, starting roughly 150 ms after stimulus, when stimulus detection is associated with inhibition of microsaccades. We tested whether that effect of perceptual report (misses vs. hits) is equal in the two bias conditions. The distinct-criteria hypothesis predicts greater inhibition on hit trials of the conservative condition, because conservative hits are "purer" (i.e., they contain fewer lucky guesses) and require a strong sensory signal that is also likely to trigger oculomotor freezing.

Indeed, the microsaccade rate dips lower on hit trials of the conservative condition (Fig. 3*C*, light red lines) than of the liberal condition (light blue lines). To summarize these effects, we integrated microsaccade rates over the time window with significant inhibition (shaded windows in Fig. 2, *A* and *B*, and Fig. 3*B*). For each bias condition we then computed the effect of explicit detection as a modulation index: $(M - H)/(M + H)$, where M is the microsaccade rate on miss trials and H is the rate on hit trials. The effect of explicit detection was significant (95% CI of the index excludes 0) in all conditions except the liberal condition of *experiment 2* (see Table 4). According to a linear mixed-effects model, that modulation index was significantly larger in the conservative than the liberal condition (by 0.18 on average; CI = [0.08 0.14], $P = 0.0004$; BF = 34.0). This is strong evidence that correct reports of target presence in the conservative condition are associated with stronger inhibition of microsaccades than in the liberal condition. The effect of bias condition was also larger in *experiment 2* than in *experiment 1* (interaction between condition and experiment, $P = 0.004$; BF = 6.32).

These data consistently support the distinct-criteria hypothesis: oculomotor freezing is triggered when a sensory signal crosses a threshold that is independent of the participant's decision bias. The sensory signal is more likely to have crossed that oculomotor threshold on hit trials of the conservative condition, when the criterion for explicit report is higher, than on hit trials of the liberal condition. Thus, when the participant was induced to adopt a more conservative decision bias, explicit detection of the stimulus was associated with more robust oculomotor freezing.

DISCUSSION

Detecting potentially relevant stimuli in the environment is a fundamental task of perceptual systems. Our data suggest that although sensory input is continuous and noisy, the brain switches into a qualitatively different state when there is sufficient evidence that a target is present (39). Passing this threshold gives rise to a conscious percept and an involuntary pause of saccadic eye movements (that is, oculomotor freezing). A pause in microsaccades can be considered the oculomotor system's "report" that it detected a stimulus. The participant's decision to respond voluntarily to the stimulus, for instance, by pressing a button, depends on the conscious percept as well as potential rewards and expectations.

Visual stimulus detection therefore has three consequences that are of interest to the present investigation: a conscious percept, a decision to report stimulus presence, and oculomotor freezing. It is crucial that we understand how those three consequences relate in terms of neural and cognitive mechanisms. Although perceptual decisions and oculomotor freezing can be measured directly, the conscious percept cannot be. But if oculomotor freezing is a proxy for conscious perception (as we argue below), researchers would be equipped with a "no-report" paradigm to investigate the neural correlates of consciousness (18) without interference from explicit cognitive tasks.

In five independent experiments across this study and a previous one (12), we consistently found that explicit reports and oculomotor freezing covary: the eyes only freeze in response to stimuli that the person reports seeing. To explain that covariation, in this study we manipulated the likelihood that participants reported stimulus presence. When rewards and penalties were greater on target-present than target-absent trials (*experiment 1*) or when the target probability was known to be high (*experiment 2*), participants adopted a liberal decision criterion, reporting target presence much more often than in the opposite (conservative) conditions (Fig. 1).

In contrast, the magnitude of the drop in microsaccade rates just after stimulus onset showed little to no effect of our bias manipulations (Fig. 2). We need not rely only on that null result, however, because we also found effects of the bias condition when splitting the trials according to the explicit report (Fig. 3). The difference in microsaccade rates between hit and miss trials, which indexes the link between explicit reports and oculomotor freezing, was larger in the conservative than the liberal condition. Our interpretation is that when participants make conservative decisions they only report sensations that are strong enough to also trigger oculomotor freezing. In contrast, when participants make liberal decisions they often make strategic guesses that a stimulus was present, even when the sensory signal was weak and oculomotor freezing was not triggered.

We therefore reject the shared-criterion hypothesis and support the distinct-criteria hypothesis (described in INTRODUCTION). The criteria in question specify the magnitude of sensory evidence required to trigger a response. There is one criterion for explicitly reporting stimulus presence, and it can be shifted to maximize rewards. There is also a distinct criterion for inhibiting eye movements, which is not affected by shifts of decision criterion.

A key feature of our theory is that oculomotor freezing is all or none, not graded. In a prior study (12), we varied the visibility of a target grating by varying its luminance contrast or by adapting the observer to the same or different orientation. Considering all target-present trials, the degree of oculomotor freezing scaled with explicit d'. However, when considering only hit trials, oculomotor freezing was equivalent across all contrast levels and adaptation states. Intense stimuli had no effect on eye movements if the observer missed them, and faint stimuli were accompanied by full-fledged inhibition provided they were detected. We found similar patterns in the new data reported above, providing consistent support that oculomotor freezing is a discrete all-or-none reflex that occurs if and only if a stimulus is consciously detected.

It is noteworthy that the results supporting the distinct-criteria hypothesis were stronger in *experiment 2*, which manipulated stimulus expectations, than in *experiment 1*, which manipulated rewards (see Fig. 3, *C* and *D*). This is reminiscent of other findings that reward manipulations have weaker effects on perceptual decisions than probability manipulations do (23, 40, 41). In our case there are several possible explanations: first, there were greater individual differences in explicit report criteria in *experiment 1* (Fig. 1*E*), perhaps due to variable valuations of the rewards. Such individual differences may have added noise to the microsaccade data as well. Second, overall d' levels were higher in *experiment 1* than in *experiment 2* (Fig. 1*D*). The bias manipulations are likely to have greater effects when the target is difficult to detect. Third, expectations of stimulus presence may have different effects on visual processing than rewards do. For instance, under the predictive coding framework (42), unexpected stimuli evoke "prediction errors." If oculomotor freezing is a response to prediction errors, then in *experiment 2* there should be a larger drop in microsaccades in target-present trials of the conservative (low probability) condition than of the liberal (high probability) condition. That did happen but, importantly, only on hit trials when the participant reported seeing the stimulus (Fig. 3*C*). In contrast, the reward manipulation in *experiment 1* did not change prediction errors, which may be why it yielded a smaller difference between conservative and liberal conditions.

Regardless, neither experiment reported above supports the shared-criterion hypothesis, and the analyses of all the target-present trials together unequivocally support the distinct-criteria hypothesis. We conclude that the sensory criterion for inhibiting saccades is not the same as the criterion for deciding to report stimulus presence.

It is somewhat more difficult to draw definitive conclusions about how conscious perception relates to oculomotor freezing and explicit reports, although we argue that some conclusions are more likely than others. The question is whether the sensory criterion for conscious perception is identical to either of the other two criteria. There are three possibilities: *1*) The perceptual criterion is identical to the oculomotor freezing criterion. It is therefore not affected by bias manipulations. This explains why we have consistently found that oculomotor freezing only occurs on trials when the participant sees the stimulus. *2*) The perceptual criterion is identical to the decision criterion. This means that it is also affected by bias manipulations. In this case, participants really would have consciously perceived the target more often in the liberal conditions than in the conservative conditions, whereas oculomotor freezing remained unaffected by these criterion manipulations. *3*) The perceptual criterion is independent of both the decision criterion and the oculomotor criterion. It is not affected by bias manipulations and usually aligns with the oculomotor criterion. We have yet to find a way to dissociate the perceptual criterion from the oculomotor criterion.

All three of these possibilities are consistent with our data and with the distinct-criteria hypothesis described in INTRODUCTION. We confidently favor the first possibility: oculomotor freezing and conscious perception are coupled because they share a sensory threshold. The third possibility involves three independent sensory criteria. Although that is possible, we favor the more parsimonious explanations (*possibilities 1* and *2*) in the absence of any supporting evidence for a more complex one (*possibility 3*). To arbitrate between *possibility 1* and *possibility 2*, we must interpret why the bias manipulations affect how often participants report target presence. Do those manipulations affect the decision criterion at a postperceptual stage (*possibility 1*) or the threshold for conscious perception (*possibility 2*)?

The research addressing this question directly has yielded mixed results. There is some neurophysiological evidence that expectations, as manipulated by the probability cues in *experiment 2*, can affect sensory processing (42, 43). One theory is that expecting a stimulus evokes a "template" in neural populations that prefer the expected features (44, 45). In contrast, one functional MRI (fMRI) study concluded that payoff and probability manipulations recruit frontal and parietal brain regions involved in decision-making to shift the starting point of evidence accumulation, similar to a criterion shift (23). The existing behavioral evidence is also ambiguous. One study argued that expectation improves detection by elevating the baseline of "signal-selective units" (46). Another found that probability cues presented after the

stimulus had effects similar to cues presented before the stimulus, in favor of a postperceptual criterion shift (47). It remains a matter of discussion, therefore, whether expectations affect conscious perception or decision processes (48–51).

We are left with our own interpretations of those prior studies and our fallible notions about what is plausible when surveying the data. Based on these considerations, we advance the hypothesis that the bias manipulations operate at the postperceptual decision stage: In the liberal conditions of both our experiments, observers reported "yes" more often because doing so maximized rewards, not because they actually saw the target more often. That is why the difference in microsaccade rates between hit and correct reject trials is weaker in the liberal than the conservative condition. The implication is that oculomotor freezing provides an implicit index of conscious perception that is free of bias.

One qualification relevant to the design of future studies is that the protocol we used here yielded a low baseline rate of microsaccades. Participants were instructed to maintain fixation while awaiting a faint and brief target, and therefore kept their eyes as still as they could. That makes oculomotor freezing difficult to detect without a large number of trials. Modifications to the protocol that encourage more saccades will increase the practical advantages of using oculomotor freezing to implicitly assess conscious perception.

To conclude, our results are consistent with the hypothesis that oculomotor freezing is unaffected by biases in perceptual decision-making and shares a sensory threshold with conscious detection. Alternate hypotheses that allow the threshold for conscious detection to diverge from the threshold for oculomotor freezing are more complicated. They must either postulate an additional free parameter, for a total of three sensory thresholds/criteria, or assume that the decision criterion is also the threshold for perception and thus bias manipulations truly affect perception. In our opinion, there is insufficient evidence for such bias effects on perception. Therefore, we argue that oculomotor freezing provides a valuable tool to measure conscious perception without requiring explicit reports, free of the influence of decision bias.

This study is part of a larger research effort to understand the branching sensory pathways that support oculomotor control and visual perception. In some cases, eye movements appear to be driven by sensory signals that differ from what is consciously perceived [reviewed by Spering and Carrasco (52)]. In other cases, including the experiments reported in this article, perceptual sensitivity and oculomotor sensitivity are linked. That makes oculomotor measures quite useful for measuring perception, as has been shown for optokinetic nystagmus (53) and smooth pursuit (54). We look forward to future discoveries in this field that are both theoretically enlightening and applicable to challenges in the laboratory and in the clinic.

DATA AVAILABILITY

The data from each individual trial in both experiments, along with analysis code that generates the figures and statistics in this article, are available for download here: https://osf.io/zkcag/.

SUPPLEMENTAL DATA

Supplemental Figs. S1–S4: https://doi.org/10.17605/OSF.IO/T9BY7.

ACKNOWLEDGMENTS

We are grateful to Michael Grubb, Jan-Nikolas Klanke, and Dobromir Rahnev for comments on this manuscript, to Jan-Nikolas Klanke for data collection, and to Richard Schweitzer for technical assistance.

GRANTS

This work was funded by The Humboldt University Talent Travel Award (A.L.W., J.C.M.), Deutsche Forschungsgemeinschaft (DFG) Grants RO3579/8-1 and RO3579/12-1 (M.R.), and National Institutes of Health Grant K99 EY029366 (A.L.W.).

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

A.L.W., J.C.M., and M.R. conceived and designed research; M.R. performed experiments; A.L.W. analyzed data; A.L.W. and M.R. interpreted results of experiments; A.L.W. prepared figures; A.L.W. drafted manuscript; A.L.W. and M.R. edited and revised manuscript; A.L.W., J.C.M., and M.R. approved final version of manuscript.

ENDNOTE

At the request of the authors, readers are herein alerted to the fact that additional materials related to this manuscript may be found at https://osf.io/zkcag/. These materials are not a part of this manuscript and have not undergone peer review by the American Physiological Society (APS). APS and the journal editors take no responsibility for these materials, for the website address, or for any links to or from it.

REFERENCES

1. **Rolfs M.** Microsaccades: small steps on a long way. *Vision Res* 49: 2415–2441, 2009. doi:10.1016/j.visres.2009.08.010.
2. **Rucci M**, **Poletti M.** Control and functions of fixational eye movements. *Annu Rev Vis Sci* 1: 499–518, 2015. doi:10.1146/annurev-vision-082114-035742.
3. **Hafed ZM.** Mechanisms for generating and compensating for the smallest possible saccades. *Eur J Neurosci* 33: 2101–2113, 2011. doi:10.1111/j.1460-9568.2011.07694.x.
4. **Otero-Millan J**, **Macknik SL**, **Langston RE**, **Martinez-Conde S.** An oculomotor continuum from exploration to fixation. *Proc Natl Acad Sci USA* 110: 6175–6180, 2013. doi:10.1073/pnas.1222715110.
5. **Rolfs M**, **Kliegl R**, **Engbert R.** Toward a model of microsaccade generation: the case of microsaccadic inhibition. *J Vis* 8: 5.1–23, 2008. doi:10.1167/8.11.5.
6. **Abeles D**, **Amit R**, **Tal-Perry N**, **Carrasco M**, **Yuval-Greenberg S.** Oculomotor inhibition precedes temporally expected auditory targets. *Nat Commun* 11: 3524, 2020. doi:10.1038/s41467-020-17158-9.
7. **Amit R**, **Abeles D**, **Carrasco M**, **Yuval-Greenberg S.** Oculomotor inhibition reflects temporal expectations. *Neuroimage* 184: 279–292, 2019. doi:10.1016/j.neuroimage.2018.09.026.
8. **Badde S**, **Myers CF**, **Yuval-Greenberg S**, **Carrasco M.** Oculomotor freezing reflects tactile temporal expectation and aids tactile perception. *Nat Commun* 11: 3341, 2020. doi:10.1038/s41467-020-17160-1.
9. **Denison RN**, **Yuval-Greenberg S**, **Carrasco M.** Directing voluntary temporal attention increases fixational stability. *J Neurosci* 39: 353–363, 2019. doi:10.1523/JNEUROSCI.1926-18.2018.

10. **Betta E**, **Turatto M.** Are you ready? I can tell by looking at your microsaccades. *Neuroreport* 17: 1001–1004, 2006. doi:10.1097/01.wnr.0000223392.82198.6d.
11. **Rolfs M**, **Laubrock J**, **Kliegl R.** Shortening and prolongation of saccade latencies following microsaccades. *Exp Brain Res* 169: 369–376, 2006. doi:10.1007/s00221-005-0148-1.
12. **White AL**, **Rolfs M.** Oculomotor inhibition covaries with conscious detection. *J Neurophysiol* 116: 1507–1521, 2016 [Erratum in *J Neurophysiol* 118: 648, 2017]. doi:10.1152/jn.00268.2016.
13. **Engbert R**, **Kliegl R.** Microsaccades uncover the orientation of covert attention. *Vision Res* 43: 1035–1045, 2003. doi:10.1016/S0042-6989(03)00084-1.
14. **Hafed ZM**, **Ignashchenkova A.** On the dissociation between microsaccade rate and direction after peripheral cues: microsaccadic inhibition revisited. *J Neurosci* 33: 16220–16235, 2013. doi:10.1523/JNEUROSCI.2240-13.2013.
15. **Reingold EM**, **Stampe DM.** Saccadic inhibition in voluntary and reflexive saccades. *J Cogn Neurosci* 14: 371–388, 2002. doi:10.1162/089892902317361903.
16. **Bonneh YS**, **Adini Y**, **Polat U.** Contrast sensitivity revealed by microsaccades. *J Vis* 15: 11, 2015. doi:10.1167/15.9.11.
17. **Scholes C**, **McGraw PV**, **Nyström M**, **Roach NW.** Fixational eye movements predict visual sensitivity. *Proc Biol Sci* 282: 20151568, 2015. doi:10.1098/rspb.2015.1568.
18. **Tsuchiya N**, **Wilke M**, **Frässle S**, **Lamme VA.** No-report paradigms: extracting the true neural correlates of consciousness. *Trends Cogn Sci* 19: 757–770, 2015. doi:10.1016/j.tics.2015.10.002.
19. **Denniss J**, **Scholes C**, **McGraw PV**, **Nam SH**, **Roach NW.** Estimation of contrast sensitivity from fixational eye movements. *Invest Ophthalmol Vis Sci* 59: 5408–5416, 2018. doi:10.1167/iovs.18-24674.
20. **Crapse TB**, **Lau H**, **Basso MA.** A role for the superior colliculus in decision criteria. *Neuron* 97: 181–194.e6, 2018. doi:10.1016/j.neuron.2017.12.006.
21. **Hafed ZM**, **Goffart L**, **Krauzlis RJ.** A neural mechanism for microsaccade generation in the primate superior colliculus. *Science* 323: 940–943, 2009. doi:10.1126/science.1166112.
22. **Macmillan NA**, **Creelman CD.** *Detection Theory: A User's Guide*. Mahwah, NJ: Lawrence Erlbaum Associates, 2005.
23. **Mulder MJ**, **Wagenmakers EJ**, **Ratcliff R**, **Boekel W**, **Forstmann BU.** Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *J Neurosci* 32: 2335–2343, 2012. doi:10.1523/JNEUROSCI.4156-11.2012.
24. **Swets JA**, **Tanner WP Jr**, **Birdsall TG.** Decision processes in perception. *Psychol Rev* 68: 301–340, 1961. doi:10.1037/h0040547.
25. **Brainard DH.** The psychophysics toolbox. *Spat Vis* 10: 433–436, 1997. doi:10.1163/156856897X00357.
26. **Cornelissen FW**, **Peters EM**, **Palmer J.** The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behav Res Methods Instrum Comput* 34: 613–617, 2002.
27. **Pelli DG.** The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10: 437–442, 1997. doi:10.1163/156856897X00366.
28. **Kaernbach C.** A single-interval adjustment-matrix (SIAM) procedure for unbiased adaptive testing. *J Acoust Soc Am* 88: 2645–2655, 1990. doi:10.1121/1.399985.
29. **Engbert R**, **Mergenthaler K.** Microsaccades are triggered by low retinal image slip. *Proc Natl Acad Sci USA* 103: 7192–7197, 2006. doi:10.1073/pnas.0509557103.
30. **Widmann A**, **Engbert R**, **Schröger E.** Microsaccadic responses indicate fast categorization of sounds: a novel approach to study auditory cognition. *J Neurosci* 34: 11152–11158, 2014. doi:10.1523/JNEUROSCI.1568-14.2014.
31. **Efron B**, **Tibshirani R.** *An Introduction to the Bootstrap*. New York: Chapman and Hall, 1993.
32. **Benjamini Y**, **Hochberg Y.** Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Ser B* 57: 289–300, 1995. doi:10.1111/j.2517-6161.1995.tb02031.x.
33. **Rouder JN**, **Speckman PL**, **Sun D**, **Morey RD**, **Iverson G.** Bayesian t tests for accepting and rejecting the null hypothesis. *Psychon Bull Rev* 16: 225–237, 2009. doi:10.3758/PBR.16.2.225.
34. **Rouder JN**, **Morey RD**, **Speckman PL**, **Province JM.** Default Bayes factors for ANOVA designs. *J Math Psychol* 56: 356–374, 2012. doi:10.1016/j.jmp.2012.08.001.
35. **Kass RE**, **Raftery AE.** Bayes factors. *J Am Stat Assoc* 90: 773–795, 1995. doi:10.1080/01621459.1995.10476572.
36. **Kubovy M.** A possible basis for conservatism in signal detection and probabilistic categorization tasks. *Percept Psychophys* 22: 277–281, 1977. doi:10.3758/BF03199690.
37. **Scholes C**, **McGraw PV**, **Roach NW.** Selective modulation of visual sensitivity during fixation. *J Neurophysiol* 119: 2059–2067, 2018. doi:10.1152/jn.00819.2017.
38. **Zuber BL**, **Crider A**, **Stark L.** Saccadic suppression associated with microsaccades. *Q Prog Rep* 74: 1964, 1964.
39. **Mashour GA**, **Roelfsema P**, **Changeux JP**, **Dehaene S.** Conscious processing and the global neuronal workspace hypothesis. *Neuron* 105: 776–798, 2020. doi:10.1016/j.neuron.2020.01.026.
40. **Leite FP**, **Ratcliff R.** What cognitive processes drive response biases? A diffusion model analysis. *Judgm Decis Mak* 6: 651–687, 2011.
41. **Simen P**, **Contreras D**, **Buck C**, **Hu P**, **Holmes P**, **Cohen JD.** Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J Exp Psychol Hum Percept Perform* 35: 1865–1897, 2009. doi:10.1037/a0016926.
42. **de Lange FP**, **Heilbron M**, **Kok P.** How do expectations shape perception? Perceptual consequences of expectation. *Trends Cogn Sci* 22: 764–779, 2018. doi:10.1016/j.tics.2018.06.002.
43. **Pajani A**, **Kok P**, **Kouider S**, **de Lange FP.** Spontaneous activity patterns in primary visual cortex predispose to visual hallucinations. *J Neurosci* 35: 12947–12953, 2015. doi:10.1523/JNEUROSCI.1520-15.2015.
44. **Kok P**, **Failing MF**, **de Lange FP.** Prior expectations evoke stimulus templates in the primary visual cortex. *J Cogn Neurosci* 26: 1546–1554, 2014. doi:10.1162/jocn_a_00562.
45. **Kok P**, **Mostert P**, **de Lange FP.** Prior expectations induce prestimulus sensory templates. *Proc Natl Acad Sci USA* 114: 10473–10478, 2017. doi:10.1073/pnas.1705652114.
46. **Wyart V**, **Nobre AC**, **Summerfield C.** Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proc Natl Acad Sci USA* 109: 3593–3598, 2012 [Erratum in *Proc Natl Acad Sci USA* 109: 6354, 2012]. doi:10.1073/pnas.1120118109.
47. **Bang JW**, **Rahnev D.** Stimulus expectation alters decision criterion but not sensory signal in perceptual decision making. *Sci Rep* 7: 17072, 2017. doi:10.1038/s41598-017-16885-2.
48. **Press C**, **Kok P**, **Yon D.** The perceptual prediction paradox. *Trends Cogn Sci* 24: 13–24, 2020. doi:10.1016/j.tics.2019.11.003.
49. **Rungratsameetaweemana N**, **Itthipuripat S**, **Salazar A**, **Serences JT.** Expectations do not alter early sensory processing during perceptual decision-making. *J Neurosci* 38: 5632–5648, 2018. doi:10.1523/JNEUROSCI.3638-17.2018.
50. **Rungratsameetaweemana N**, **Serences JT.** Dissociating the impact of attention and expectation on early sensory processing. *Curr Opin Psychol* 29: 181–186, 2019. doi:10.1016/j.copsyc.2019.03.014.
51. **Summerfield C**, **Egner T.** Feature-based attention and feature-based expectation. *Trends Cogn Sci* 20: 401–404, 2016. doi:10.1016/j.tics.2016.03.008.
52. **Spering M**, **Carrasco M.** Acting without seeing: eye movements reveal visual processing without awareness. *Trends Neurosci* 38: 247–258, 2015. doi:10.1016/j.tins.2015.02.002.
53. **Dakin SC**, **Turnbull PR.** Similar contrast sensitivity functions measured using psychophysics and optokinetic nystagmus. *Sci Rep* 6: 34514, 2016. doi:10.1038/srep34514.
54. **Mooney SW**, **Alam NM**, **Hill NJ**, **Prusky GT.** Gradiate: a radial sweep approach to measuring detailed contrast sensitivity functions from eye movements. *J Vis* 20: 17, 2020. doi:10.1167/jov.20.13.17.