Journal of Experimental Psychology:
Learning, Memory, and Cognition

ISSN: 0278-7393

2022, Vol. 48, No. 2, 226–241
https://doi.org/10.1037/xlm0001055

How Do Expectations Change Behavior? Investigating the Contributions at Encoding Versus Decision-Making

Syaheed B. Jabar and Daryl Fougnie
Department of Psychology, New York University Abu Dhabi

Expectations about the environment play a large role in shaping behavior, but how does this occur? Do expectations change the way we perceive the world, or just our decisions based on unbiased perceptions? We investigated the relative contributions of priors to these 2 stages by manipulating *when* information about expected color was provided. We compared cases where the prior could affect encoding into perceptual/working memory representations (e.g., when provided prestimulus) against cases where it could not (e.g., when given at response after a delay). Although priors had a minor influence on encoding, the bulk of the effects were at decision-making. Furthermore, these effects appeared to be distinct. The effect on decision-making was Bayesian-like, with priors inducing bias while improving precision. In contrast, the same priors at encoding improved precision without causing changes in bias. Priors do not just affect encoding *or* decision-making, but appear to affect both, via distinct mechanisms.

Keywords: expectation, Bayesian priors, visual perception, decision-making, working memory

Supplemental materials: https://doi.org/10.1037/xlm0001055.supp

Imagine being asked to select, from a color palette, the color of the last stop sign that you saw. Odds are that you are going to be biased toward the prototypical red associated with such signs, even if the last exemplar observed happened to be atypical. Although biases toward our expectations are commonly observed in perceptual decision-making (Ma et al., 2006; Summerfield & de Lange, 2014; Summerfield & Egner, 2016; Wei & Stocker, 2015), the mechanisms that lead to such effects are unclear.

One possible mechanism is that expectations affect how information is perceptually encoded by changing the specificity or sensitivity of neural responses in visual cortex (de Lange et al., 2018; Summerfield & de Lange, 2014; Zhou et al., 2020). Indeed, repeated exposure to a stimulus has been shown to change neural firing patterns in visual cortical areas (Schoups et al., 2001), suggesting that expectations do alter perceptual encoding. Furthermore, this can happen in a fast timescale, even within a single session, because imbalances of stimuli (e.g., Gabors with differing orientation probabilities) can be shown to affect V1 activity and result in increased precision of perceptual estimations (Jabar et al., 2017). Computationally, having information flow through probabilistic population codes in the visual cortex is sufficient to mimic Bayesian decision-making (Beck et al., 2008) and other models of informed decision-making, for example applying efficient coding (Wei & Stocker, 2015) or attractor dynamics (Bitzer et al., 2015) also assume changes in population codes that are associated with perception or sensory representation. Expectations induced by explicit cues (rather than acquired though piecemeal experience) can also produce changes in visual cortex. For example, if informative auditory cues are provided before (visual) stimuli, it results in representational changes in the visual cortex (Kok et al., 2013; Kok et al., 2014; Kok et al., 2012). Magnetoencephalogram studies also show that auditory cues can modulate early visual processing (e.g., Shams et al., 2005). These studies therefore suggest that within-trial cues can also affect perceptual encoding.

An alternate hypothesis is that the effects of priors are on decision-making. For instance, judging the absence/presence of Gabor orientations, clearly a perceptual task, is nonetheless affected by nonperceptual information, such as reward (Summerfield & Koechlin, 2010). It could be that in perceptual tasks that probability distributions are encoded by the cortex, but are only collapsed onto estimates when decisions are needed (Ma et al., 2006). Bang and Rahnev (2017) attempted to study the effect of expectations on sensory processing versus decision-making using a task where participants had to make a choice about whether the average of a series of Gabor orientations were clockwise or counterclockwise from vertical. Participants were provided cues about the likely correct response (left or right) either before or after the stimulus, with the result that the postcue accounted for all the effects of the precue. Accordingly, they concluded that stimulus expectations affect decision criteria rather than sensory representations, which is in line

Syaheed B. Jabar https://orcid.org/0000-0001-7583-0744
Daryl Fougnie https://orcid.org/0000-0001-8245-883X
The data that supports this project has been deposited at https://osf.io/cnzmj/ (Jabar & Fougnie, 2020).
Correspondence concerning this article should be addressed to Syaheed B. Jabar, Department of Psychology, New York University Abu Dhabi, Saadiyat Island, P.O. Box 129188, Abu Dhabi, United Arab Emirates. Email: sbj2@nyu.edu

with other studies arguing that expectations do not alter early sensory processing/encoding (Rungratsameetaweemana et al., 2018; Rungratsameetaweemana & Serences, 2019).

How do we reconcile the hypothesis that priors affect encoding, with the hypothesis that priors affect decision-making[1]? One issue is that the Bang and Rahnev (2017) task involved explicitly cuing the response dimension, which only indirectly provided information about the likely stimuli. Under such conditions, it is perhaps not surprising that a cue during decision would have a strong effect. The present study adopts a similar pre- and postcue manipulation within a delayed perceptual matching task, but where cues would inform participants of the exact stimulus distribution. Furthermore, the influence of expectations can have dissociable effects on behavior: It can bias responses toward the expected value and/or take advantage of the additional information to make responses more precise (both are expected if participants are acting in a Bayesian manner; Geisler, 2011). The relative contribution of pre- and postcues on the precision and bias of responses has not been examined in previous studies. We therefore had participants use a continuous report, which allowed us to separate bias and precision of responses (see Figure 1; Zhang & Luck, 2008). We expected that color estimations should not only be biased toward the prior, but that estimation should also be more precise overall (Figure 1c). To preview our findings, we found dissociable influences of expectation at prestimulus and poststimulus stages. When prior information is available at response, estimates become more accurate and are biased toward the expected value, consistent with Bayesian-like inference. Unlike the Bang and Rahnev study, we also found a small additional benefit of priors during stimulus presentation (but this benefit was found to improve quality in a nonbiased manner). Our results suggest that expectations need not have a singular effect, and that they influence both encoding and decision-making in distinct ways.

Experiment 1: Effects of Response Priors

Is the introduction of a prior after perceptual encoding sufficient to produce changes in behavior? If the bulk of the effects of priors are on perceptual encoding, then priors given only at response should produce considerably smaller changes in behavioral performance compared to showing the same prior at prestimulus. Otherwise, if the bulk of the effects of priors are due to postperceptual/decisional process, then response priors should have an effect on performance with the addition of prestimulus priors contributing little to no additional performance benefits. We therefore compare the condition where participants are given no priors, versus only given priors at response, versus given the prior both prestimulus and at response.

Method

Participants

Based on a pilot study (within-subjects, $n = 11$), we found that a no-prior condition ($M = 17.3°$, $SD = 6.7°$) had significantly larger raw errors than response prior condition ($M = 15.5°$, $SD = 6.2°$) compared to a response prior condition, $t(11) = 2.72$, $p = .022$. Sampling from this data, the null hypothesis was rejected (alpha cutoff of $<.05$) at least 80% of the time when the sample size was 19 or higher. Therefore, twenty naïve participants (three male, 17 female, median age = 19.5) were recruited for Experiment 1.

Participants had normal or corrected-to-normal vision and declared normal color vision. Participants were recruited at New York University Abu Dhabi and took part either for course credits, or for a subsistence allowance of 50AED per hour. Written consent was obtained from each participant before the experiments. The experiments were approved by the New York University Abu Dhabi Institutional Review Board.

Apparatus and Stimuli

Stimuli were presented on a color-calibrated 24" BenQ XL2411 monitor (144 Hz refresh rate, 1,920 × 1,080 pixels), placed 57 cm away from the participant. All experiments reported here were coded in Matlab using the Psychophysics Toolbox (Brainard, 1997; Kleiner et al., 2007). The background of the screen was black (0.1cd/m^2) throughout the experiments, and room lighting was dim. A white (89cd/m^2) fixation cross was used, spanning 1.5° visual angle in both height and width. Cues used as priors were Gaussians (standard deviation of 30°) comprising of 360 color bands of tapering height (each band corresponds to an integer value of color orientation). We chose a prior width of 30° to maximize the effect of priors. Given an expected uncertainty of approximately 30° in memory, prior widths between 30° to 40° lead to the maximum bias toward the target. We chose the minimal end of this so that participants can more easily tell the differences in height (Gaussian approaches uniform as its standard deviation increases). This cue was also shown mirrored across the horizontal axis to emphasize the height differences more (e.g., instead of a Gaussian, it was a lens shape). Crucially, the height of the color bands accurately indicated the relative likelihood of that specific color (see, e.g., Figure 2). The cues were centrally located and were of maximal luminance (25cd/m^2) and height (1.5° visual angle) at the midpoint (which indicated the center of the prior). The visible horizontal span of the cue was approximately 8° visual angle. The targets and the response color wheel were located 8.3° away from fixation diagonally (displaced in the x- and y-axes by 5.9°). The response wheel occupied approximately 8° visual angle, while each target occupied approximately 2.2°. The mean color of each of the four targets on a trial were independently drawn from the same 30° prior in color space, even if the prior was not communicated to the participant (i.e., in the control condition). This was to ensure no differences in stimulus properties across conditions. To increase stimulus noise/uncertainty, for each individual target, pixels were randomly resampled from a Gaussian distribution of 30° around that target's mean color. This was done as the effect of priors are expected to increase as stimulus uncertainty increases (Geisler, 2011). The colors for the cue, the color wheel, and for the targets were isoluminant (approximately 25cd/m^2) and sampled from the MemToolbox (Suchow et al., 2013) color space.

[1] By "encoding," we do not distinguish between perceptual or memory encoding. The key distinction here is whether it is these encoding processes, or later post-encoding "decision-making" processes (which may include processes such as working memory retrieval) that are responsible for the effects of expectations or "priors".

Figure 1
Sample Error Distributions

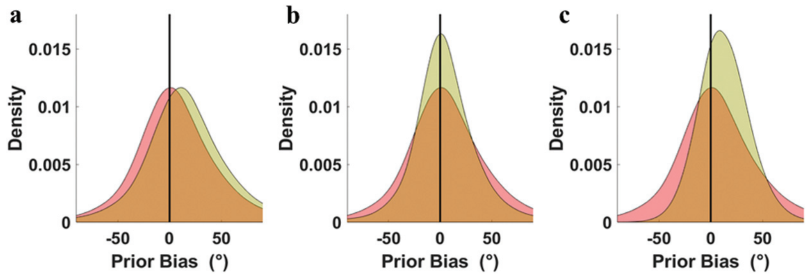

Note. Positive error (x-axis) indicates a bias towards the prior. Red (dark gray) indicates a sample control distribution, mustard (lighter gray) indicates the modified version of the control distribution. a: Error distribution predicted by a change in bias only. b: Error distribution predicted by a change in precision only. c: Error distribution predicted by an optimal Bayesian observer (see the online supplementary material). Note how the Bayesian process will result in both a change in bias as well as a change in distribution width. See the online article for the color version of this figure.

Procedure

Once written consent was obtained, participants were instructed that depending on the trial, the cues might either be given only at response, or both at response and prestimulus, or not at all. It was also made clear to the participants that the height of the cue corresponded to the likelihood of the target colors on that trial, and that while colors outside of the span of the visible cue were unlikely, they were nonetheless possible. Participants were instructed to be as precise as possible, and that the cash bonus was tied to their performance. Participants were given 20 practice trials to familiarize themselves with the task, and then were given three blocks of one hundred trials each, with breaks in-between. Each block corresponded to a different condition and block order was counterbalanced across participants.

Figure 2
General Paradigm

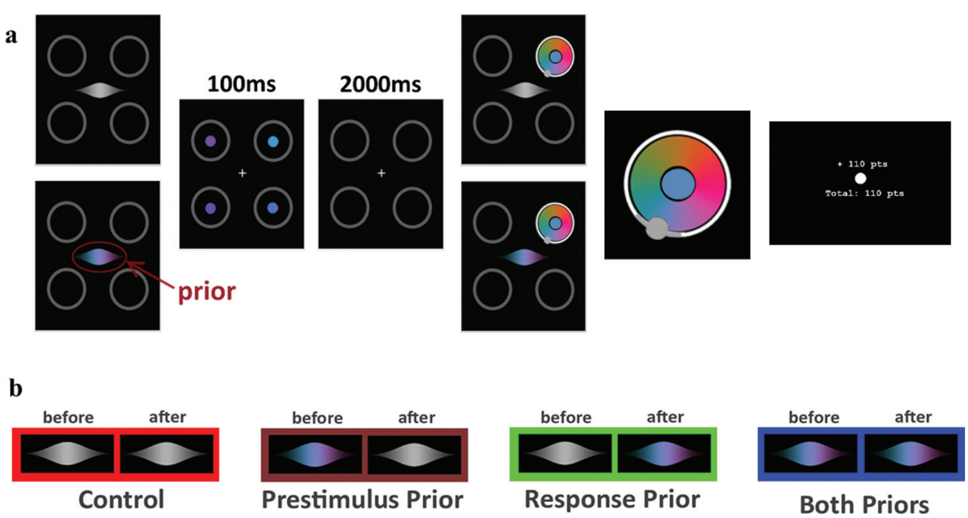

Note. a: Participants began with a central prior cue, the height of which is indicative of the likelihood of the color. In the case where there is no prior to be given, for example, in the control condition, the cue was represented in gray. The four targets appeared 500 ms after the cue was clicked, for a duration of 100 ms. This was followed by a delay of 2,000 ms. Thereafter, the cue appeared again, and one of the four circles (equiprobable) was filled with a color wheel, which indicated the to-be-reported location. Participants used a mouse to click on the color of the wheel to make their estimates. Participants then moved the mouse away from the point of click to extend a confidence arc (zoom-in is for illustration only, sizes of stimuli appear fixed on-screen). The participant then obtains points depending on the size of the confidence arc drawn. b: The possible cue combinations. Note that even in the control condition, while the prior might not be advertised, the four dots were still drawn from a 30° Gaussian distribution centered on some random color. Also note that the prestimulus prior condition only appears in Experiment 2 and 4.

Each trial began with a self-paced prior cue (the prestimulus prior). Participants were given as long as they desired to look at the prior, and they initiated the trials by clicking on the middle of the prior. This is followed by 500 ms (72 frames) of a black screen with a central fixation cross. The four stimuli then appeared simultaneously for approximately 100 ms (14 frames). This was followed by another blank screen with a fixation cross for 2,000 ms. A color wheel then appeared in one of the fours target locations randomly (spatial probability is equiprobable), and participants were to use the mouse cursor to click on what they thought was the color of the target that was in that cued location. The central part of the color wheel started off as white and then was updated to reflect the current choice that the participants were hovering over (see Figure 2a). Once the first click registered, participants were given the opportunity to draw a confidence arc over the region of the color they thought they saw. For example, if they were confident that they saw a particular shade of green, they could just draw a narrow arc. If they were unsure, they were told to make the arcs larger. Participants locked-in their responses with a second click. Arcs were always centered on the initial point of estimation (the first click). Participants were made to draw these arcs as this method could be used to gauge implicit confidence (Honig et al., 2020).

After the confidence estimate, participants were given visual feedback about the trial. Points were awarded on each trial. Maximum points per trial was 180, with points deducted based on the size of the confidence arc. For example, if the target were within the confidence arc, spanning 20°, the participant would have gotten 160 points for that trial. No points were awarded if the correct answer fell outside the arc. The optimal play for participants to obtain the maximum points is therefore to both be as precise as possible in making their estimate, as well as minimizing their confidence window when they were certain of what they encoded (Honig et al., 2020). Both the cumulative and points earned on-trial was shown on-screen, and participants were informed beforehand that the points earned scaled to performance and would be translated to bonus money that they could earn. Performance that averaged to 90 out of the maximum 180 points earned no bonus. Maximum bonus was 25 AED. This performance scaling was to maintain the motivation to be as accurate as possible.

Data Analysis

For data analysis, we first examined overall report error across conditions. However, error differences could arise from changes in bias or precision. Bias was determined, for each condition, by taking the magnitude of the errors of each trial and assigning a sign depending on whether the error was closer to (positive) or further away (negative) to the (central value of the) prior than it should be. The median of these signed errors was then taken. As with other perceptual matching tasks (e.g., Anderson, 2014), to estimate precision independent of bias the errors were adjusted by the overall condition bias (via a subtraction), and then the median absolute error was estimated. Simulations on the width and shifting of Gaussian distributions (e.g., those used to generate Figure 1) show that this type of analysis accurately returned the original Gaussian parameters. Model-free analyses were used as the main analysis since these measures makes fewer assumptions than model-based approaches.[2]

Importantly, we also conducted analyses using a mixture modeling approach and found no qualitative differences between the two approaches (see the online supplementary material). Furthermore, estimates of guess responses were small, arguing against the need of an analysis approach aimed to separate guess from nonguess response contributions.

Results

An alpha cutoff of $p < .05$ is used for the pairwise t-tests. Bayes Factor[3] (*BF*) analyses were also done for these tests (using the *BayesFactor* package; Morey et al., 2015) to determine how strongly the data favor null or alternative hypotheses.

Raw Error

Based on the raw errors (before calculating bias/precision, Figure 3b) there was a (nonsignificant) trend toward larger errors in the control condition (where no priors were given; $M = 20.2°$, $SD = 5.5°$) compared to the response prior condition (providing the prior at response only, $M = 18.9°$, $SD = 3.4°$), $t(19) = 1.18$, $p = .252$, $\eta^2 = .07$, $BF = 1.25$). Errors were significantly larger in the control than both priors condition (providing the prior prestimulus and at response, $M = 17.7°$, $SD = 4.6°$), $t(19) = 2.39$, $p = .027$, $\eta^2 = .23$, $BF = 5.92$). There was no significant difference in error between the Response Prior condition and the both priors condition, $t(19) = 1.67$, $p = .110$, $\eta^2 = .13$, $BF = .54$. At the very least, this suggests that the introduction of priors can result in a decrease in overall error compared to the control condition. However, the error data are ambiguous on whether adding a prestimulus prior had an effect.

Bias Toward the Prior

Simply looking at the raw errors does little to dissociate effects on bias from effects on precision. For example, Figure 3a, which depicts the error distribution, is clearly suggesting that the presence of the prior cues is causing a shift in bias toward the prior. Using the calculation outlined earlier, we found that there was a larger bias in the response prior condition ($M = 9.3°$, $SD = 6.1°$) than in the control condition ($M = 2.5°$, $SD = 4.6°$), $t(19) = 4.81$, $p < .001$, $\eta^2 = .56$ $BF = 158$). The both priors condition ($M = 9.9°$, $SD = 6.7°$) was also significantly more biased toward the prior than the control condition, $t(19) = 5.53$, $p < .001$, $\eta^2 = .63$, $BF = 248$. The response prior condition did not significantly differ from the both priors condition, $t(19) = .61$, $p = .549$, $\eta^2 = .02$, $BF = .27$.

This suggests, first, that priors given at response are sufficient to cause a bias. Second, having the same prior additionally be available to affect encoding does not contribute additional bias.

[2] For example, do we assume a fixed precision or variable precision model? (Fougnie et al., 2012; Shen & Ma, 2019; van den Berg et al., 2012).

[3] Bayes factors can be used to compare any two models to each other, for instance an alternative to a null model (Rouder et al., 2009). Any value above 1 reflects evidence for the model in the numerator, while any value below that reflects evidence for the model in the denominator. The commonly used guidelines of $BF > 3$ or $BF < .33$ was adopted in this study as the criteria of evidence for either model.

Figure 3
Experiment 1 Results

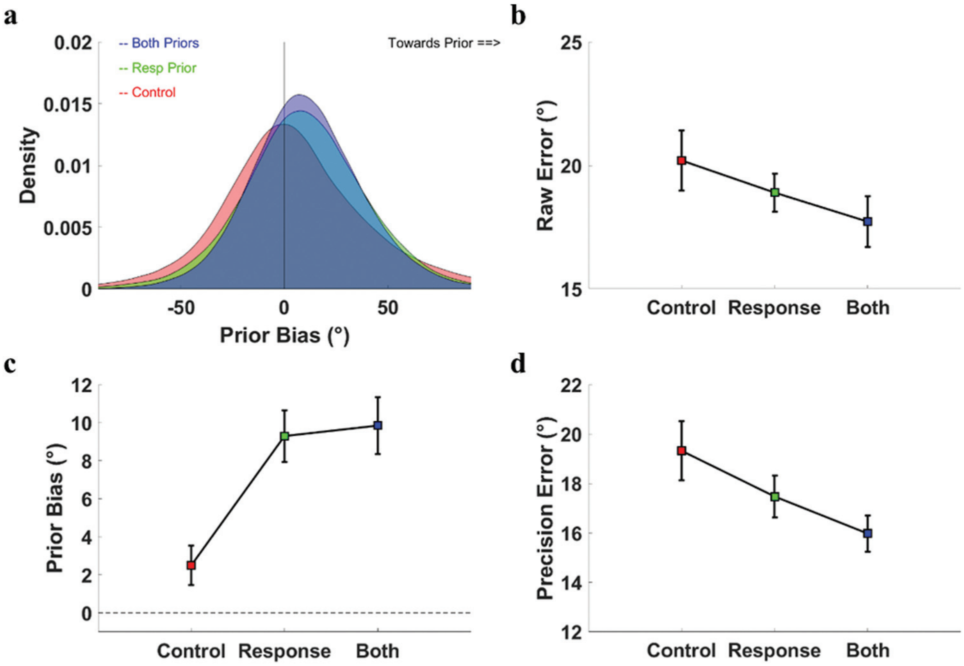

Note. a: Error distribution across participants. Each participant's errors per condition was put though a kernel density fit, and we calculate the mean density across participant at each possible integer value of error. Positive error (x-axis) indicates a bias towards the prior. b: Raw errors (°) for the three conditions. These errors were split into (b) bias and (c) precision errors. Red denotes the control condition, green denotes the response prior condition, and blue denotes the both priors condition. Error bars indicates one standard error, and the markers indicate the means across subjects.

Precision (Bias-Adjusted)

Because a Bayesian prior predicts not only an increase in bias (Figure 1c) but also precision, we also analyzed the spread of responses around the bias to determine if the prior influenced the amount of uncertainty. The control condition (M = 19.3°, SD = 5.3°) was not significantly more or less precise than the response prior condition (M = 17.5°, SD = 3.8°), $t(19) = 1.63$, $p = .119$, $\eta^2 = .13$, $BF = 2.17$. The control condition was significantly less precise than the both priors condition though (M = 16.0°, SD = 3.3°), $t(19) = 3.50$, $p = .002$, $\eta^2 = .41$, $BF = 75.2$. The response prior condition was marginally associated with less precision than the both priors condition, $t(19) = 1.95$, $p = .066$, $\eta^2 = .17$, $BF = 1.12$. In sum, although it is clear that priors can improve precision, it is unclear from this experiment what are the relative contributions of the prestimulus and response priors to this effect.

Confidence-Arc Sizes

Although the original intent of having the arc sizes drawn was to obtain a measurement of implicit confidence, we found no significant difference between the three conditions: control (M = 49.9°, SD = 8.9°), response prior (M = 48.2°, SD = 8.9°), both priors (M = 48.8°, SD = 9.7°), all $ps > .05$. This lack of arc size effects extends to the other experiments as well, and as such will not be further reported on. Either priors do not affect implicit confidence, or more likely, this measurement was too insensitive to capture differences in confidence levels.

Discussion

The introduction of priors results in a clear Bayesian-like effect (compare to Figure 1c to 3a), with estimates being both biased toward the prior. Of interest though is whether there is a benefit to introducing a prior prestimulus so that the prior can influence the encoding process, or whether showing the prior after perception, at response is sufficient to produce effects typical of priors. When errors were broken down into bias versus precision, we found no evidence that responses were more biased toward the prior. In fact, the Bayesian analysis found supporting evidence that adding the prestimulus prior did not create additional bias effects.

The results also suggest that priors could led to more precise responses. The evidence was inconclusive on whether this might be due to the prior at decision alone or whether the addition of priors at encoding boosted this effect. One possibility is that there is a small benefit on precision from encoding priors, and that this is difficult to observe relative to the benefits from decision priors. Indeed, a bootstrap simulation on our data suggested that to observe a significant additional prior effect (both priors vs. response prior, $p < .05$) at least 80% of the time on response precision would require data from approximately 80 participants. This question of priors at encoding is addressed later in the paper. First, we address potential confounds with Experiment 1's design.

Experiment 2 was designed to deal with the concern with that participants may not have focused on the prestimulus prior since it was redundant with the response prior. In Experiment 1, the prestimulus prior would *always* reappear during the response stage, and this may have led participants to discount or completely ignore the information. Hence, we tested a case where sometimes only the prestimulus prior was shown, and we randomized conditions within blocks to make it impossible to predict whether a prior would reappear at response, in order to encourage participants to use the prestimulus prior.

Experiment 2: Mixed Blocks and the Effects of the Prestimulus Prior Alone

To ensure that exactly half of the trials had a prestimulus prior, the Prestimulus Prior only trial type was introduced (for a total of four trial types) and conditions were randomized within blocks. This also had the benefit of allowing us to directly compare the effect of the prestimulus prior to that of the response prior. If priors do not affect perceptual encoding, then the effects of the prestimulus prior should be approximately equal to the effect of the response prior, on the assumption that previously shown priors would not be forgotten at the time of test.

Method

Twenty additional participants (14 male, six female, median age = 22) were recruited for Experiment 2. Apart from the introduction of the prestimulus prior condition and that the trials were mixed rather than blocked, all details of Experiment 2 were identical to Experiment 1. The *mu* of the priors were randomly selected every trial.

Results

A graphical depiction of the raw error, bias and precision results is given in Figure 4b–4d. Figure 4a shows the distribution of errors, with errors toward the prior coded as positive. Because of the large number of possible comparisons (six pairwise comparisons for each set of analyses) we made an a priori decision to only compare the conditions that bear directly on our hypotheses (e.g., whether the addition of prestimulus priors causes an effect over and above response priors). Those interested in additional comparisons can look at Table 1 where we list the statistics for a larger-N replication of this study (Experiment 4).

Raw Error

The control condition showed (M = 23.0°, SD = 8.0°) significantly larger errors than the prestimulus prior condition (M = 20.1°, SD = 4.2°), $t(19) = 2.33$, $p = .031$, $\eta^2 = .23$, $BF = 3.02$, and the response prior condition (M = 20.0°, SD = 4.1°), $t(19) = 2.66$, $p = .016$, $\eta^2 = .28$, $BF = 3.10$. Furthermore, the prestimulus prior

Figure 4
Experiment 2 Results

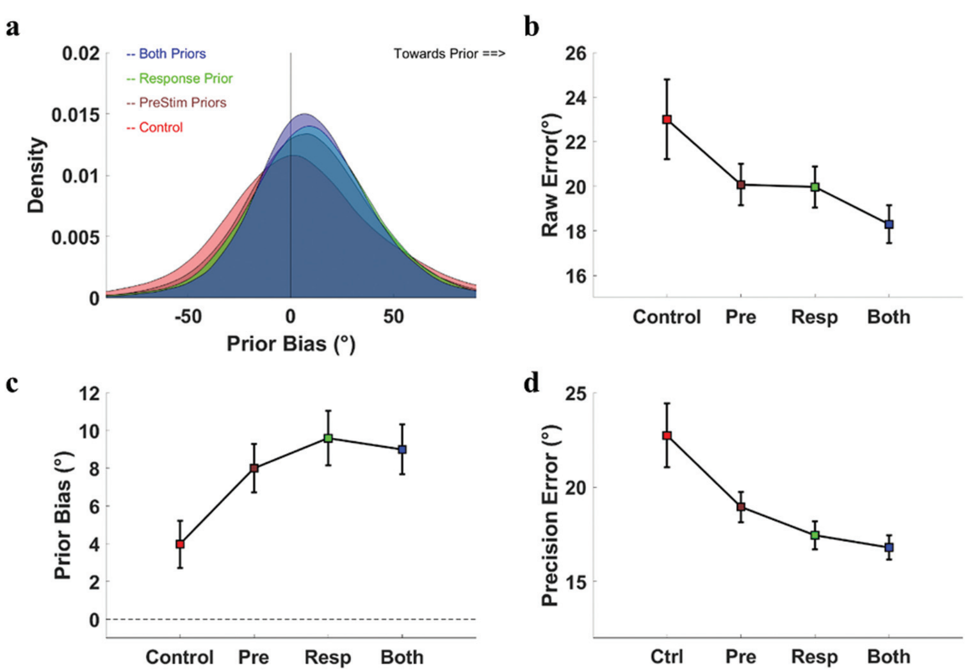

Note. a: Error distribution across participants. Positive error (x-axis) indicates a bias towards the prior. b: Raw errors (°) for the four conditions. These errors were split into (c) bias and (d) precision errors. Red denotes the control condition, brown denotes the prestimulus prior condition, green denotes the response prior condition and blue denotes the both priors condition. Error bars indicates one standard error, and the markers indicate the means across subjects.

condition did not differ from the response prior condition, $t(19) = .12$, $p = .90$, $\eta^2 < .01$, $BF = .23$. We also examined whether adding a prestimulus prior improved performance above and beyond the response prior: The response prior condition had significantly larger errors than the both priors condition ($M = 18.3°$, $SD = 3.8°$), $t(19) = 2.63$, $p = .017$, $\eta^2 = .28$, $BF = 11.5$.

Bias Toward the Prior

The prestimulus prior ($M = 8.0°$, $SD = 5.7°$) was significantly more biased toward the prior than the control condition ($M = 4.0°$, $SD = 5.6°$), $t(19) = 4.06$, $p < .001$, $\eta^2 = .48$, $BF = 33.0$. The prestimulus prior condition was not significantly more or less biased than the response priors condition ($M = 9.6°$, $SD = 6.4°$), $t(19) = 1.70$, $p = .106$, $\eta^2 = .14$, $BF = .40$. Bias for the response prior significantly differed from the control condition, $t(19) = 5.99$, $p < .001$, $\eta^2 = .65$, $BF = 66.8$, but did not significantly differ from than the both priors condition ($M = 9.0°$, $SD = 5.9°$), $t(19) = .79$, $p = .438$, $\eta^2 = .03$, $BF = .31$. Although prestimulus priors alone are sufficient to cause a prior bias, this might just be reflecting the influence of the prior at response rather than at perception, because there is no additional bias from having a prior at prestimulus in addition to at response.

Precision (Bias-Adjusted)

The control condition ($M = 22.8°$, $SD = 7.5°$) was estimated significantly less precisely than the prestimulus prior ($M = 19.0°$, $SD = 3.6°$), $t(19) = 2.56$, $p = .019$, $\eta^2 = .27$, $BF = 3.12$. The prestimulus prior condition was associated with marginally less precision than the response prior condition ($M = 17.5°$, $SD = 3.3°$), $t(19) = 2.05$, $p = .054$, $\eta^2 = .19$, $BF = 4.21$. The response prior condition had significantly better precision than the control condition, $t(19) = 3.68$, $p = .002$, $\eta^2 = .42$, $BF = 17.0$), but did not significantly differ from than the both priors condition ($M = 16.8°$, $SD = 2.9°$), $t(19) = 1.11$, $p = .281$, $\eta^2 = .06$, $BF = .40$. As with Experiment 1, there was inconclusive evidence for whether or not having an additional prestimulus prior affects precision. This emphasizes the need for a larger dataset.

Discussion

Replicating the previous experiment, having a prior poststimulus, at response, was sufficient to cause Bayesian-like effects: Estimates were both more biased to the prior and also more precise. Comparing the response prior to the both priors condition showed at best mixed evidence of additional performance improvements for the both priors condition. Although there was a small effect in raw error, this effect disappeared when we attempted to separate the effect into bias and precision. There is little apparent effect of adding the prestimulus prior (i.e., of enabling effects at encoding), even though unlike Experiment 1, participants could no longer just depend on the response prior in the Both Priors condition. This finding suggests that the priors have their bulk of their effect on postperceptual decision-making.

Comparing the prestimulus prior condition to the response prior condition also supports this hypothesis. Participants were clearly affected by the prestimulus prior yet the effects of the prestimulus prior largely resembled the response-prior, mirroring the result of Bang and Rahnev (2017) on orientations, where the postcue accounted for the effects of the precue. In this case it suggests that there is little effect of the prestimulus prior on encoding. However, the prestimulus prior would still be expected to have an effect during the decision stage. The slighter weaker bias and precision effects compared to the both priors condition could be explained as participants having to rely on the memory of the prior in the prestimulus prior condition, whereas it was on-screen in the both priors condition.

This issue of uneven memory demands for the prestimulus cue versus the response cue (which did not have to be remembered since it persisted on-screen) was why the duration of the prestimulus cue was left as self-paced. Because the presentation duration of the prestimulus priors were self-paced additional analyses were done to determine how long the participants spent on that screen. Although the response prior was left on-screen for the whole duration of the response (ranged from about 2 to 4 s), participants only spent approximately 400 ms looking at the prestimulus priors. This raises the concern that the time spent looking at the prior could be playing a role in our findings, as the duration was unequal across conditions. This was examined in the next study.

Experiment 3: Timed Versus Self-Paced Priors

Experiment 3 aimed to explore whether manipulating the time spent looking at the prior played a role and also to equate this across conditions. To achieve this, we ran the experiment on two more separate groups of participants. In one group, the prior cue was, as in Experiments 1 and 2, self-paced. With the other group, we fixed the duration of the cue, both at prestimulus and at response, to 1,000 ms.

Method

Two additional groups of twenty participants were recruited from the same pool. Experiment 3 was identical to Experiment 2 except for two details. First, the prestimulus prior only trials were dropped (the both priors condition remained), with the three remaining conditions being equiprobable. Unlike Experiment 1, the trials were mixed, not blocked. For half the participants (seven males, 13 females, median age = 20), the duration of the prestimulus prior was fixed at 1 second, as was the duration of the response prior. The other participants (eight males, 12 females, median age = 21) had the same stimulus timings as in Experiments 1 and 2.

Results

A graphical depiction of the raw error, bias and precision results is given in Figure 5b–5d. Figure 5a shows the distribution of errors, with errors toward the prior coded as positive.

Between Groups: Self-Paced Versus Timed Duration (1 s)

There were no significant differences for any of the matched conditions across the self-paced and timed groups. For raw error, there is no significant difference between the control condition of the self-paced group and the control condition of the timed group, $t(38) = 1.11$, $p = .275$, $\eta^2 = .03$, $BF = .54$. There is also no significant difference in raw error across groups for the response prior condition, $t(38) = .29$, $p = .776$, $\eta^2 < .01$, $BF = .31$, or the both priors condition, $t(38) = .89$, $p = .380$, $\eta^2 = .02$, $BF = .43$.

Figure 5
Experiment 3 Results

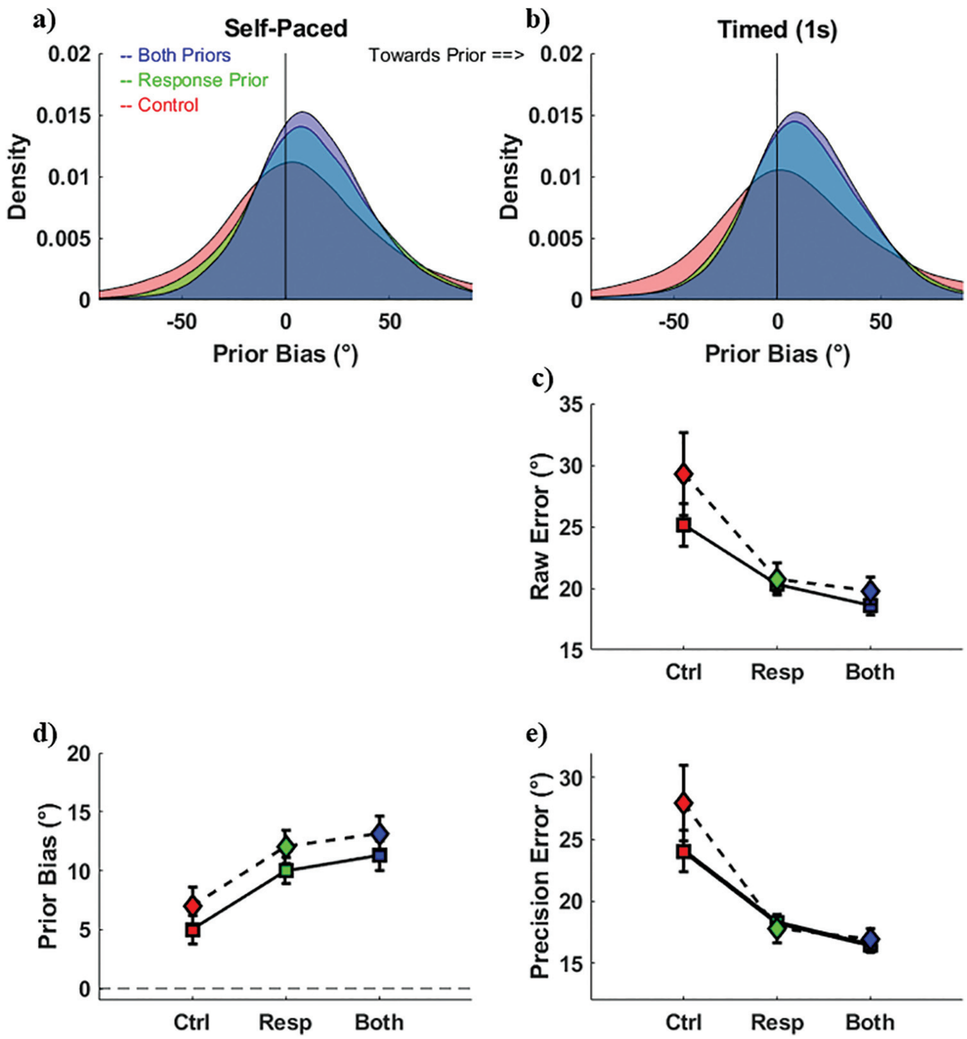

Note. a: Error distribution across participants for the self-paced group. Positive error (x-axis) indicates a bias towards the prior. b: Error distribution across participants for the timed group. c: Raw errors (°) for the three conditions, for the two groups. These errors were split into (d) bias and (e) precision errors. Red denotes the control condition, green denotes the response prior condition, and blue denotes the both priors condition. Error bars indicates one standard error, and the markers indicate the means across subjects. The solid line indicates the self-paced priors group and the dashed line indicate the timed (1 s) priors group.

The lack of between group effects also held true when decomposing the raw errors into the bias (all *p*s > .05) and precision metrics (all *p*s > .05). *BF* analyses (comparing the effect of response prior [response prior - control, *BF* = .31] and of both priors [both prior - control, *BF* = .31] on bias across the two group conditions) support the null hypothesis. This subtraction was done because there might be baseline performance differences across the two groups of participants. Similar *BF* analysis on the effect on precision due to the response and both prior was less conclusive, although there was more support for the null over the alternate hypothesis (*BF*s = .60 and .49, respectively). These results suggest that the duration for which the cue is presented did not significantly modulate the effect of the prior.

Raw Error

To further examine if time spent on prior was playing a role, we separately analyzed the raw error, bias, and precision for each timing condition. Critically, we replicate the main findings of the previous studies and find no divergence between the two prior durations.

For the self-paced group, the control condition (M = 25.2°, SD = 7.6°) had significantly larger errors than the response prior condition (M = 20.3°, SD = 3.1°), $t(19) = 3.35$, $p = .003$, $\eta^2 = .38$, BF = 15.6). The response prior condition did not significantly differ in raw error magnitude from the both priors condition (M = 18.6°, SD = 3.3°), $t(19) = 1.85$, $p = .080$, $\eta^2 = .16$, $BF = 1.10$.

For the timed group, the control condition (M = 29.3°, SD = 15.0°) had significantly larger errors than the response prior condition (M = 20.8°, SD = 5.8°), $t(19) = 3.12$, $p = .006$, $\eta^2 = .35$, BF = 12.9. The response prior condition did not significantly differ in raw error magnitude from the both priors condition (M = 19.8°, SD = 4.9°), $t(19) = 1.22$, $p = .239$, $\eta^2 = .08$, BF = .36.

Bias Toward the Prior

For the self-paced group, the response prior condition (M = 10.0°, SD = 5.0°) had significantly larger biases than the control condition (M = 5.0°, SD = 5.4°), $t(19) = 4.65$, $p < .001$, $\eta^2 = .55$, BF = 207. The response prior condition was not significantly more or less biased than the both priors condition (M = 11.3°, SD = 5.6°), $t(19) = 1.11$, $p = .279$, $\eta^2 = .06$, BF = .40.

For the timed group, the response prior condition (M = 12.1°, SD = 6.3°) was also significantly more biased toward the prior than the control condition (M = 7.0°, SD = 7.2°), $t(19) = 4.15$, $p <$ .001, $\eta^2 = .49$, BF = 31.5. The response prior condition was again not significantly more or less biased than the both priors condition (M = 13.2°, SD = 6.4°), $t(19) = 1.55$, $p = .138$, $\eta^2 = .12$, BF = .65). Therefore, although the response prior increases bias with respect to the control condition, there is ambiguous evidence whether the prestimulus prior has any additional effect on either group.

Precision (Bias-Adjusted)

For the self-paced group, the control condition (M = 24.1°, SD = 7.6°) was significantly less precise than the response prior condition (M = 18.3°, SD = 3.0°), $t(19) = 3.55$, $p = .002$, $\eta^2 = .41$, BF = 34.0. The response prior condition was in turn significantly less precise than the both priors condition (M = 16.5°, SD = 2.7°), $t(19) = 2.36$, $p = .029$, $\eta^2 = .24$, BF = 2.12, although the BF test gave an ambiguous result.

For the timed group, the control condition (M = 27.9°, SD = 13.7°) was significantly less precise than the response prior condition (M = 17.8°, SD = 5.1°), $t(19) = 3.51$, $p = .002$, $\eta^2 = .41$, BF = 17.1. Precision for the response prior condition was not significantly different from the both priors condition (M = 17.0°, SD = 3.9°), $t(19) = 1.61$, $p = .125$, $\eta^2 = .13$, BF = .69, although the BF test again gave an ambiguous result.

Discussion

Taken as a whole, Experiment 3 suggests that time on prior was likely not a confounding issue. Furthermore, regardless of the prior duration, we replicate the previous finding that the bulk of the effect of priors occurs after encoding. The additional effects of the prestimulus prior on either bias or precision for either group was again inconclusive, which highlights the need for a large-N study.

Experiment 4: Dual Mechanisms of Priors?

Although the previous studies convincingly demonstrate that having priors at response is sufficient to cause a Bayesian-like effect, the additional effects of having priors available at perception are still unclear. Chiefly, the previous results could suggest either there is no additional effect either on bias or precision, or there is a minor effect only on precision. To distinguish

Table 1
Statistics for Experiment 4 (n = 80)

Error	Control	PreStim	Resp
Raw error			
Control M = 25.1°, SD = 7.0°	—	—	—
PreStim M = 23.4°, SD = 5.9°	$t(79) = 3.49$, $p = .001$, $\eta^2 = .13$, BF = 29.69	—	—
Resp M = 22.6°, SD = 5.3°	$t(79) = 4.06$, $p < .001$, $\eta^2 = .17$, BF = 175.34	$t(79) = 1.47$, $p = .146$, $\eta^2 = .03$, BF = 0.34	—
Both M = 21.5°, SD = 4.9°	$t(79) = 5.50$, $p < .001$, $\eta^2 = .28$, BF = 31,457.99	$t(79) = 3.46$, $p = .001$, $\eta^2 = .13$, BF = 26.98	$t(79) = 2.23$, $p = .029$, $\eta^2 = .06$, BF = 1.26
Bias error			
Control M = 5.2°, SD = 6.0°	—	—	—
PreStim M = 8.2°, SD = 7.5°	$t(79) = 3.70$, $p < .001$, $\eta^2 = .15$, BF = 55.56	—	—
Resp M = 8.2°, SD = 6.9°	$t(79) = 4.27$, $p < .001$, $\eta^2 = .19$, BF = 357.06	$t(79) = 0.02$, $p = .981$, $\eta^2 < .01$, BF = 0.12	—
Both M = 9.0°, SD = 5.9°	$t(79) = 5.42$, $p < .001$, $\eta^2 = .27$, BF = 23,044.93	$t(79) = 1.16$, $p = .250$, $\eta^2 = .02$, BF = 0.23	$t(79) = 1.15$, $p = .255$, $\eta^2 = .02$, BF = 0.23
Precision error			
Control M = 24.8°, SD = 6.8°	—	—	—
PreStim M = 21.7°, SD = 5.7°	$t(79) = 5.25$, $p < .001$, $\eta^2 = .26$, BF = 11,799.97	—	—
Resp M = 21.3°, SD = 5.7°	$t(79) = 5.29$, $p < .001$, $\eta^2 = .26$, BF = 14,016.20	$t(79) = 0.70$, $p = .487$, $\eta^2 = .01$, BF = 0.16	—
Both M = 19.9°, SD = 4.8°	$t(79) = 7.25$, $p < .001$, $\eta^2 = .40$, BF = 40594940	$t(79) = 3.55$, $p = .001$, $\eta^2 = .14$, BF = 35.94	$t(79) = 2.84$, $p = .006$, $\eta^2 = .09$, BF = 5.05

Note. PreStim = prestimulus; Resp = response; BF = Bayes factors. $BF > 3$ suggests evidence for the alternate hypothesis, whereas $BF < .33$ suggests evidence for the null.

between these possibilities, we initially combined the data from Experiments 1, 2, and 3 (80 data sets). Although there was no difference in terms of bias between the both priors condition (M = 10.8°, SD = 6.2°) and the response prior condition (M = 10.2°, SD = 6.0°), $t(79) = 1.31$, $p = .195$, $\eta^2 = .02$, $BF = .28$, there was one for precision: Both priors condition (M = 16.6°, SD = 3.2°) was more precise than the response prior condition (M = 17.8°, SD = 3.8°), $t(79) = 3.61$, $p < .001$, $\eta^2 = .14$, $BF = 42.23$. This would suggest that although priors at decision produce large benefits on precision and bias, that the benefit of a response prior is limited only to an effect on response precision.

This aggregate analysis has the disadvantage of combining data across experiments that differ in methodological details. Therefore, we conducted a preregistered (https://osf.io/d4xjk) study of 80 participants using the design of Experiment 2. We chose to replicate Experiment 2 as this was the study that best emphasized the importance of the prestimulus prior. In Experiments 1 and 3, participants could potentially ignore the prestimulus prior since there was no condition in which the prior was not shown during the decision stage.

Method

This replication experiment was conducted online via Prolific (https://prolific.co), with the experiment being coded in HTML canvas/Javascript instead of Matlab/Psychophysics Toolbox. The stimuli colors, sizes and timings were kept the same as the original experiment, assuming participants were using a set-up similar to what we used in the lab (e.g., apparent stimulus sizes might vary depending on how far participants situate themselves from the screen, participant screen calibration might affect the relative luminance of the colors used, etc.). Participants were credited 5.20 GBP per hour, and the maximum bonus was 5GBP. We also removed the need for the confidence arc judgements. We ran this online replication on 80 participants (29 female, 69 right-handed, median age = 23 years old). This N was prechosen because a power analysis on the previous data suggested that 80 participants would give us an 80% power to detect the small additional effect of prestimulus priors on precision. Further, 80 participants were sufficient to show this effect strongly in the aggregate analysis on Experiments 1–3.

Results

For the sake of brevity, below we report the comparisons that matters the most. For the interested, Table 1 has the t-statistics and BFs for all possible comparisons.

Prestimulus Priors Additionally Increase Precision

The both priors condition (M = 19.9°, SD = 4.8°) was estimated significantly more precisely than the response prior condition (M = 21.3°, SD = 5.7°), $t(79) = 2.84$, $p = .006$, $\eta^2 = .09$. In addition, BF analysis on this data returned a factor of 5.05, which is evidence in favor of the alternate hypothesis.

Prestimulus Priors Do Not Additionally Increase Bias

We ran the identical analyses on bias, where we observe no significant effect between the both priors condition (M = 9.0°, SD = 5.9°) and the response prior condition (M = 8.2°, SD = 6.9°), $t(79) = 1.15$, $p = .255$, $\eta^2 = .02$. This was not because of a lack of experimental power: BF analysis on this data returned a factor of .23, suggesting that the available data is evidence in favor of the null hypothesis. Therefore, Experiment 4 supports the claim that prestimulus priors increases precision (BF = 5.05) while not affecting bias (BF = .23). Note that this pattern is also consistent with the aggregate data combining across Experiments 1–3 ($BF_{precision} = 42.23$, $BF_{bias} = .28$).

Rather than solely relying on just incoming data (the stimuli) *or* just the available prior, an optimal Bayesian observer is one that integrates both pieces of information. Given the assumption that each participant has a fixed uncertainty width across trials, that they fully internalized the prior width/*mu* per trial, and that they integrate this information in a Bayesian manner, we can estimate what the average bias/precision across trials should be (see online supplemental materials for more details). What is clear from Figure 6e is that an optimal observer tends to be increasingly biased toward the prior as the distance from the prior to the target increases. This is logical, since the further away from the prior the actual target is, the greater is the distance it can be pulled toward the prior without overshooting it. Likewise, a target that is already at the prior, by definition cannot be biased toward the prior (any estimation error in this case is going to cause it to be biased away from the prior). Note that this correlation between bias and distance from prior will also occur if participants only respond with the (center of the) prior. However, this hypothetical scenario overpredicts the extent of the bias, and also predicts that all the effects will be on bias. Clearly, participants were not adopting this strategy.

For both bias and precision there are clear differences between the prior conditions and control conditions, showing that participants are using the prior. Of interest is that for bias, the both priors and response prior condition showed similar trends, where the further the (color) distance from the prior, the greater is the bias toward the prior (which is exactly what is predicted by the optimal Bayesian fit). Also, of note is that the participant performance is quite close to that of this Bayesian optimal, suggesting that participants were behaving in a Bayesian-like way. The benefit of having the prior available prestimulus (and at response) did not cause participants to be over-Bayesian or overutilize the prior. That there is a consistent prior bias observed in the control condition, might also suggest that participants could sometimes decipher and integrate the prior even when the prior was not explicitly communicated (note that even in the Control condition, the targets were drawn from a 30° SD Gaussian prior). Alternatively, it could be that people are biased toward the ensemble mean (Brady & Alvarez, 2011; Chetverikov et al., 2016; Dakin & Watt, 1997) rather than the having integrated the prior per se. Importantly, the nonrandom nature of the stimuli were equivalent across conditions.

Of greater interest is that, for precision, there is a clear divergence between just having response priors versus adding a prestimulus prior (the difference between the green and blue lines). Furthermore, the improvement in precision appears to happen regardless of the stimulus-to-prior distance. Unlike bias effects which increase with stimulus-to-prior distance, the first bin already shows this difference in precision significantly, $t(79) = 2.49$, $p = .015$, $BF = 3.92$, and this difference remains at Bin 4, $t(79) = 3.27$, $p = .002$, $BF = 15.95$. The implication here is that the priors, if they are available before or at perception, improves the encoding precision not only of the prior, but also of colors further away

Figure 6

Experiment 4 is a Replication of Experiment 2 (n = 80)

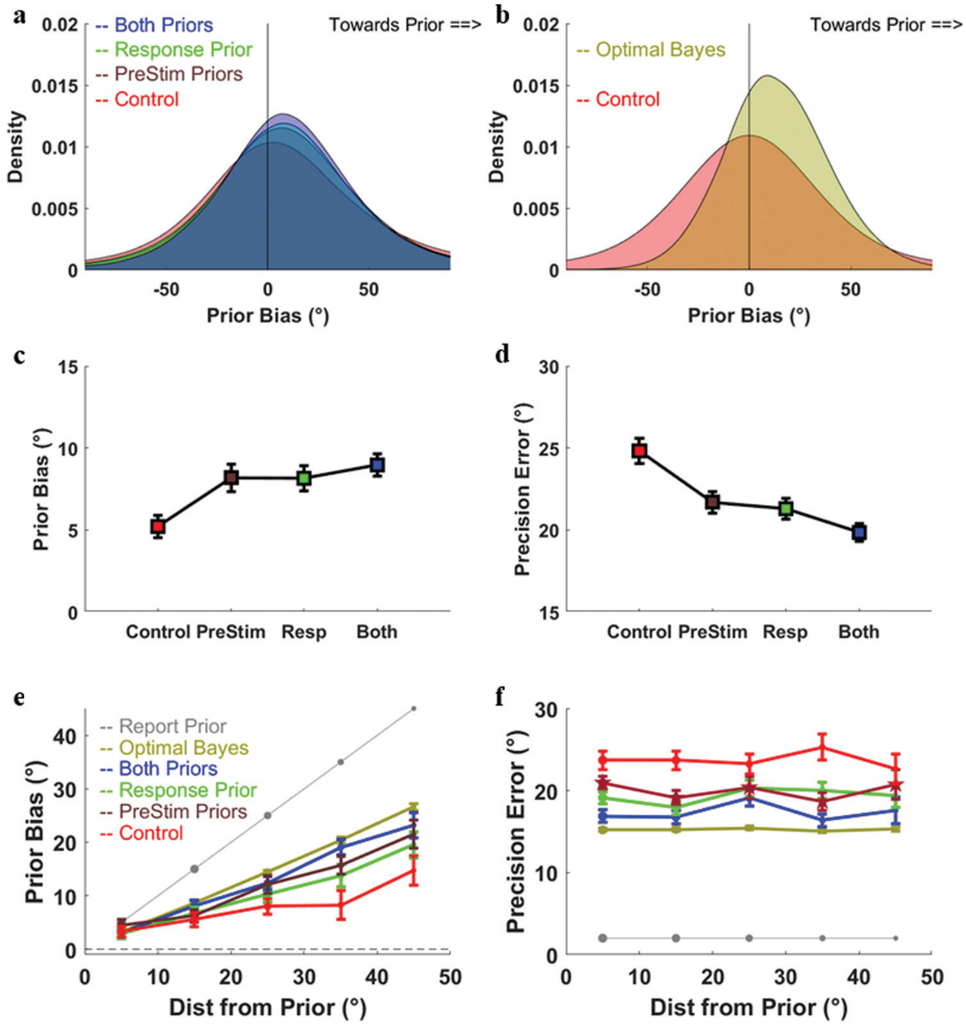

Note. a: Error distribution across participants. Positive error (x-axis) indicates a bias towards the prior. c: mean bias and (d) mean precision errors. e: bias across distance from prior and (f) precision errors across distance from prior. Red denotes the control condition, brown denotes the pre-stimulus prior condition, green denotes the response prior condition, and blue denotes the both priors condition. Mustard denotes the Optimal Bayesian given the participants' errors in the control condition (see the online supplemental materials). The gray line represents the situation where participants only responds with the (center of the prior). The bins are in intervals of 10° of distance from prior. The bins, from the first to the fifth, contains 28.5%, 25.3%, 20.1%, 14.3%, and 9.1% of the data, respectively. Error bars indicate one standard error, and the markers indicate the means across subjects (marker sizes indicate relative proportion of data going into those bins).

from the prior. This also suggests that the additional effect here is not due to something like color/repetition priming (Kristjánsson & Campana, 2010; Shurygina et al., 2019), because if that were the case, the colors further away from the prior should be less expected and primed to a lesser extent (see Figure 2 for an example of the cue). Perceptual history also tends to result in biases away from what is shown (i.e., biases away from the prior; Fritsche et al., 2017), which was not observed: Addition of the prestimulus prior has no apparent effect on bias.

Discussion

The data highlights that there are two possible mechanisms for the effect of priors. First, even without the opportunity to affect encoding, priors will affect perceptual decision-making by biasing responses toward the prior, while also improving the precision or limiting the spread of the responses. When further allowed to interact with perceptual encoding, priors then have a *small* additional effect on precision, without a corresponding effect on bias.

That the optimal Bayesian observer closely matched the participant trends is not necessarily indicating that participants are being Bayesian optimal. It could be the case that the process leading to integration of the prior is less than optimal, but they happened to compensate for this by underestimating the width of the prior (narrower prior widths are predicted to cause bigger effects of priors). That people are less than optimal integrators could be why the 'both priors' case seems nearer to the optimal than the response prior case is. However, as these effects are not necessarily Bayesian, but merely compatible with mechanisms capable of mimicking Bayesian effects (Bowers & Davis, 2012), we can conclude only that the effect of priors, particularly at the response stage, are Bayesian-like in the effects that are produced. In fact, the possibility that people are using the priors twice, at perception and again at decision-making, makes it difficult to reconcile with a pure Bayesian account of human behavior.

One could argue that the prestimulus priors are having an effect only at decision and not encoding. We agree that it is dangerous to assume that a prior shown at perception can only have an effect during perception. Participants could be remembering the prior and accessing this information during response. However, even acknowledging the possibility, the results still demonstrate dissociable effects of priors presented at distinct stages of a task. Further, the scenario that the perceptual prior is only affecting decision-making is unlikely for several reasons. First, we manipulated prior duration in Experiment 3 and observed no effect, suggesting that more time with the prior is not sufficient to produce noticeable effects on decision-making. Further, if there were to be an effect of longer/redundant cue presentation, we would expect it to be just an exaggerated effect of the effect of response priors: *Both* precision and bias should be affected. Alternatively, such presentation could lead people to be overly reliant on the prior, causing increased bias without necessarily an associated increase in precision. Neither of these predict a case of increased precision without an associated increase in bias. Finally, there is evidence from other lines of work for perceptual effects from nonrandom stimulus properties (e.g,, Jabar et al., 2017). Therefore, the evidence is at least suggestive that there is something specific about the predecisional timing of the prior cue that is leading to its unique effects.

Thus far we have only considered the case where the prior is changing from trial to trial. It could perhaps be the case that the effect that priors have on perception requires time to develop, that that is why the effects thus far has mainly been driven by the prior's influence on response. Thus, we ran Experiment 5, comparing fixed versus constantly changing priors to determine whether there are differences in how the two types of priors affect behavior.

Experiment 5: What If Priors Were Fixed?

We ran another version of the experiment on a new set of (in-lab) participants. Each participant went through three blocks (order was counterbalanced). In one block the prior was, as with the previous experiments, changing from trial to trial. In the other block, the prior was fixed. The final block served as a control condition. If an effect on encoding does require time to develop, then there should be differences in error when comparing between fixed priors and changing priors. Particularly, the fixed prior condition should result in greater prior biases and/or greater precision than the changing priors condition.

Method

Sixteen additional participants (seven male, nine female, median age = 21) were recruited from the same pool. This experiment consisted of three blocks, similar to Experiment 1. There was the control block where stimuli on each trial were drawn from an unadvertised prior that was randomly chosen. There was a change block where the prior was advertised and changed on each trial. Finally, there was a fixed block where the prior was advertised, but remained consistent across the 100 trials in the block (this fixed prior was chosen at random per participant). None of the trials had a prestimulus prior since, in the fixed prior block, the cue given in the response/feedback phase of the preceding trial already provided the same information a prestimulus cue would have. Hence, when participants reset their mouse to the central position to begin the next trial, the cue disappeared, followed by a 500-ms central fixation, followed by the 100-ms stimulus duration. The details of this experiment were otherwise identical to experiment 1. Block order was counterbalanced across participants.

Results

A graphical depiction of the raw error, bias and precision results is given in Figure 7b–7d. Figure 7e shows how the three measures change over time on task.

Raw Error

The control condition (no advertised prior, M = 22.1°, SD = 5.9°) had significantly larger errors than the changing priors condition (M = 16.8°, SD = 3.0°), $t(15) = 3.46$, $p = .004$, $\eta^2 = .46$, BF = 18.1, and the fixed prior condition (M = 16.4°, SD = 4.9°), $t(15) = 5.11$, $p < .001$, $\eta^2 = .65$, BF = 458). The changing priors condition did not differ in raw error magnitude from the fixed prior condition, $t(15) = .31$, $p = .760$, $\eta^2 < .01$, BF = .26).

Bias Toward the Prior

The changing priors condition (M = 6.3°, SD = 4.5°) had significantly more bias toward the prior than the control condition (M = −.6°, SD = 4.2°), $t(15) = 4.09$, $p < .001$, $\eta^2 = .54$, BF = 64.0), and the changing priors condition did not differ in bias from the fixed prior condition (M = 7.8°, SD = 6.0°), $t(15) = 1.02$, $p = .321$, $\eta^2 = .07$, BF = .31), which met the $BF < .33$ criteria we adopted for this study.

Precision (Bias-Adjusted)

The control condition (M = 21.8°, SD = 6.1°) was estimated significantly less precisely than the changing priors condition (M = 15.6°, SD = 3.4°), $t(15) = 3.99$, $p = .001$, $\eta^2 = .53$, BF = 33.8, and the precision for the changing priors condition did not significantly differ from the fixed prior condition (M = 15.4°, SD = 4.9°), $t(15) = .15$, $p = .883$, $\eta^2 < .01$, BF = .28.

Effects Across Trials

Here we tested a fixed versus changing prior to see if having a constant prior produces expectation effects that are greater than changing priors. Another way of probing this question is to see if having a fixed prior produced expectation effects that develop over time. To investigate this, we separated the fixed prior blocks

Figure 7
Experiment 5 Results

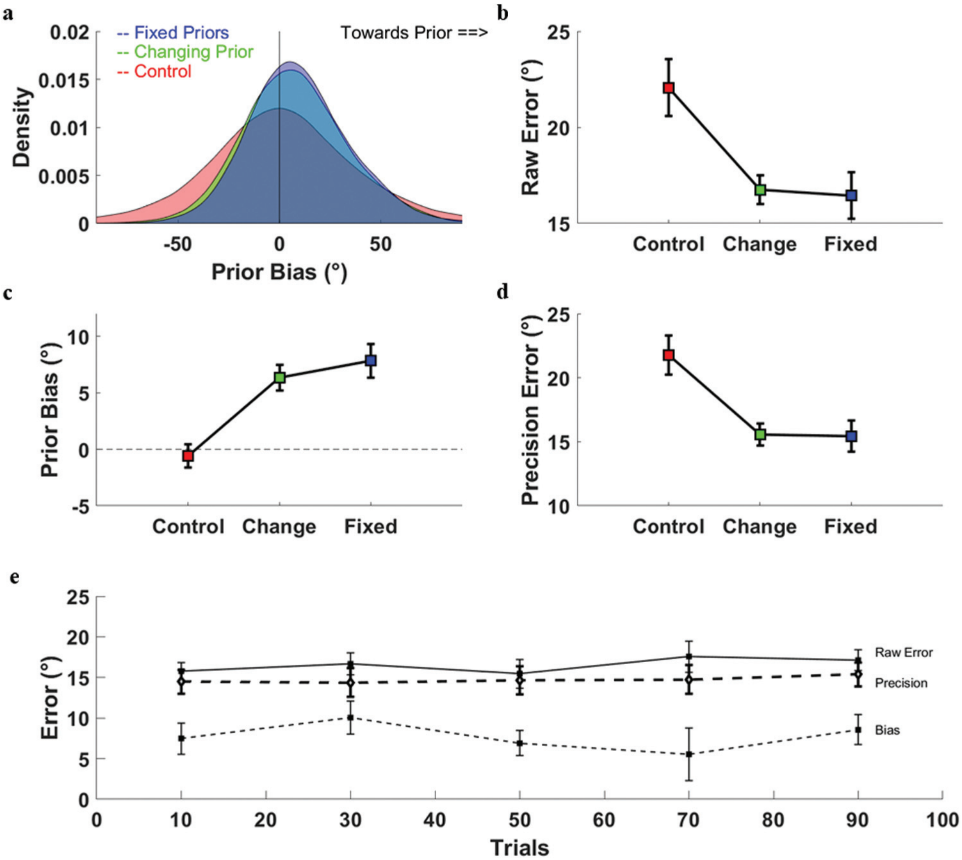

Note. a: Error distribution across participants. Positive error (x-axis) indicates a bias towards the prior. b: Raw errors (°) for the three conditions. These errors were split into (c) bias and (d) precision errors. Red denotes the control condition, green denotes the changing priors condition, and blue denotes the fixed prior condition. Error bars indicates one standard error, and the markers indicate the means across subjects. e: Error measures for the fixed prior condition broken down into 20-trial segments.

into 20-trial bins and examined performance changes over time (Figure 7e). For the raw errors, there was no significant difference between any of the bins (all *ps* $>$ .05), for example trials 1–20 had on average the same error magnitude as trials 81–100. This was also true for the precision and bias measures (all *ps* $>$ .05). In other words, there did not appear to be any measurable effect of time when priors were kept constant. This was true regardless of block order (e.g., whether the fixed prior block was done before the changing priors block or vice-versa).

Discussion

Our previous studies used priors that changed on a per trial basis. In contrast, some researchers use priors/expectations that are consistent across many trials (e.g., Jabar et al., 2017; Zhou et al., 2020). Are there major differences in the effects of fixed priors, particularly in regard to effects during perception? If a prior was fixed, alterations to the perceptual system could gradually accrue occur over time, causing greater perceptual effects. Alternatively, it could be that past experience needs to be accumulated to result in an attentional set (Leber & Egeth, 2006). However, Experiment 5 demonstrated that having a fixed prior did little to change the effect. We should note that while the *BF* analyses between the changing and fixed prior trials supported the null (for both bias and precision effects), we still view that Experiment 5 is likely underpowered for concluding whether fixing versus changing the prior over trials has zero effect. What we think this does demonstrate though, is that the effects observed in Experiments 1–4 are unlikely to have been greatly altered had we used a fixed prior.

The lack of effects here might appear contradictory to previous evidence of perceptual learning across blocks (Jabar & Anderson, 2015; Jabar et al., 2017) and that probability learning in general occurs relatively quickly (Hon & Jabar, 2018; Hon et al., 2013). Of note though is that for these studies, participants were not explicitly informed about the distributions. What this current lack of learning-related effects is suggesting might be that being explicitly given priors precludes the need to acquire the information piecemeal. A single exposure to the prior cue might be sufficient to internalize the probability information (or at least, mostly internalize it).

General Discussion

Expectations about the environment play a large role in shaping behavior. There has been considerable research demonstrating how establishing priors changes the way participants respond in a wide range of tasks (Summerfield & de Lange, 2014). What is less well-understood is the underlying mechanism of how these priors influence behavior (and during which stage of processing). One critical question is whether these priors change the perceptual/memory representations that people form or just change how decisions are made based on the encoded representations. In most studies, priors are available at both perception and decision-making, making this critical question impossible to untangle.

The present study manipulated when priors were presented. Priors given after encoding has occurred would only have the opportunity to affect decision-making, yet when a prior is presented during the response and decision stage (after perception) we see traditional Bayesian-like effects where responses become more certain, biased toward an expectation, and the amount of bias depends on the deviation from the expected value, suggesting that people are incorporating the two pieces of information in the response. Specifically, the magnitude of the bias shift increases with increased stimulus-prior distance and is in line with the predictions of an optimal Bayesian model (Geisler, 2011). This is also consistent with decision processes acting in a Bayesian fashion when making inferences about the environment from perceptual information (e.g., Fougnie et al., 2015; Maloney & Mamassian, 2009; Mamassian & Landy, 2001; Norton et al., 2019).

In contrast, if the prior is additionally presented during perceptual processing, there is an increase in response precision that suggests improved encoding of information. We use the term 'encoding' here to be agnostic on whether the effect is happening at perception or during encoding of the information into working memory, which are difficult to disentangle. Both perception and working memory likely rely on similar representations, as there is considerable evidence that working memory representations are perceptual representations that are maintained by internal attention (Chun, 2011; Courtney et al., 1997; Kiyonaga & Egner, 2013). However, there could be additional selection between perception and working memory. Representations might be biased against perceptual distractors (Rutman et al., 2010; Sreenivasan & Jha, 2007), and prestimulus expectations/priors conceivably could aid in this process by defining the target set (e.g., what targets to maintain, and what distractors to suppress). However, this would predict that target colors further away from the cued color should suffer some performance cost. Instead, the prestimulus cue increases precision in a way that showed no evidence of being dependent on the stimulus' value relative to the prior, suggesting that a bias from distractors account cannot explain the findings.

One possible mechanism for the additional effect of prestimulus priors is feature-based attention, which involves the selective gain or tuning of population responses (Ling et al., 2009) coding for that target feature. Single-unit studies corroborate this hypothesis, as enhanced neural response and synchrony are seen whenever a stimulus matched a target feature (Bichot et al., 2005). Increased neural firing due to feature-based attention can be seen even in regions with no stimuli (Serences & Boynton, 2007). Any increase in neural sensitivity or prefiring will likely improve how well the stimulus end up being encoded, particularly when the stimulus duration is limited (in this case 100ms). Alternatively, changes in perceptual precision also could be explained in terms of more efficient predictive coding (de Lange et al., 2018). This account posits less recurrent processing/feedback (O'Brien & Raymond, 2012) due to the fact that knowledge of the prior minimizes the mismatch between expectations and reality. This can boost performance if the time to process the stimulus is limited, by increasing the efficiency of stimulus processing. Future work is necessary to adjudicate between possible accounts, and to explore how these effects differ from effects during decision-making.

Do the present findings help us interpret the discrepant findings across studies? The Bayesian-like effect on decision-making accounts for the bulk of the influence of expectations, as compared to the effect on encoding. That the encoding effect is small is likely why some studies fail to find evidence for cuing/expectation effects on sensory processing (Bang & Rahnev, 2017; Rungratsameetaweemana et al., 2018), whereas others find effects (Kok et al., 2012; Zhou et al., 2020). Bootstrap analyses suggested that about 80 participants were required to show this effect reliably. Another possible explanation for discrepant findings is due to methodological differences. For example, the Bang and Rahnev study used cues that indicated the likely response and may have increased effects on decision-making and reduced encoding effects (the cue was only indirectly informative of stimulus identity). Thus, to observe encoding effects it may be necessary to directly influence participants' expectations of stimulus identity. Finally, while the lack of an effect of expectations on early sensory processing-related ERPs (Rungratsameetaweemana et al., 2018) could also be due to the effect at encoding being very small, another argument could be made about the choice of ERP components; the visual negativity (peaks 150-300ms poststimulus onset) and prepeak centro-parietal positive potential (200–750 ms poststimulus onset) used are rather late components. In contrast, the Jabar et al. (2017) study on visual evoked potentials found that probability reliably affected the C1 component (peaks 50–100 ms poststimulus), without observable differences in the later P1/N1 components. Later components are more likely affected by attentional or recurrent feedback (Di Russo et al., 2003), which may mask earlier differences in sensory processing (especially if they were small to begin with).

Another important question relevant to the present work is *when* explicit awareness of priors is necessary (Vadillo et al., 2020). In the present case the priors were explicitly advertised to participants, but there is evidence that stimulus probabilities can affect the visual cortex without participants being able to explicitly state what the distribution is (Jabar et al., 2017). This suggests that implicitly-learned priors can at least affect perceptual encoding. If attention capture can influence color reports without participant awareness (Chen et al., 2019), explicitly shown priors could possibly have similar effects. One suggestion is that explicit awareness is necessary for top–down expectations to have an effect, but not for the generation of bottom-up stimulus-evoked prediction errors (Meijs et al., 2018). Future work is necessary to determine whether explicit awareness of priors is necessary for the Bayesian-like effects that arise during decision stages. Concretely, could estimates still be biased toward the prior in cases where participants are unable to explicitly state what the prior is?

The current study provides some answers toward the role of priors, but at the same time, raises other questions. Nevertheless, the insights from these studies should help integrate models of perception and decision-making, as well as constrain theory on how humans use expectations to interact with the environment. With fields like behavioral/neuro economics increasingly incorporating Bayesian frameworks to predict how human beings make decisions (Bach, 2016; Chater, 2015; Gilboa, 2015), it becomes important to understand how we leverage our acquired information to compensate for noisy sensory inputs.

References

Anderson, B. (2014). Probability and the changing shape of response distributions for orientation. *Journal of Vision*, *14*(13), 15. https://doi.org/10.1167/14.13.15

Bach, D. R. (2016). Decision-making under uncertainty. In M. Reuter & C. Montage (Eds.), *Neuroeconomics* (pp. 99–111). Springer. https://doi.org/10.1007/978-3-642-35923-1_6

Bang, J. W., & Rahnev, D. (2017). Stimulus expectation alters decision criterion but not sensory signal in perceptual decision making. *Scientific Reports*, *7*(1), 17072. https://doi.org/10.1038/s41598-017-16885-2

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., Shadlen, M. N., Latham, P. E., & Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*, *60*(6), 1142–1152. https://doi.org/10.1016/j.neuron.2008.09.021

Bichot, N. P., Rossi, A. F., & Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. *Science*, *308*, 529–534. https://doi.org/10.1126/science.1109676

Bitzer, S., Bruineberg, J., & Kiebel, S. J. (2015). A Bayesian attractor model for perceptual decision making. *PLoS Computational Biology*, *11*(8), e1004442. https://doi.org/10.1371/journal.pcbi.1004442

Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, *138*(3), 389–414. https://doi.org/10.1037/a0026450

Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychological Science*, *22*(3), 384–392. https://doi.org/10.1177/0956797610397956

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. https://doi.org/10.1163/156856897X00357

Chater, N. (2015). Can cognitive science create a cognitive economics? *Cognition*, *135*, 52–55. https://doi.org/10.1016/j.cognition.2014.10.015

Chen, J., Leber, A. B., & Golomb, J. D. (2019). Attentional capture alters feature perception. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(11), 1443–1454. https://doi.org/10.1037/xhp0000681

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2016). Building ensemble representations: How the shape of preceding distractor distributions affects visual search. *Cognition*, *153*, 196–210. https://doi.org/10.1016/j.cognition.2016.04.018

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407–1409. https://doi.org/10.1016/j.neuropsychologia.2011.01.029

Courtney, S. M., Ungerleider, L. G., Keil, K., & Haxby, J. V. (1997). Transient and sustained activity in a distributed neural system for human working memory. *Nature*, *386*(6625), 608–611. https://doi.org/10.1038/386608a0

Dakin, S. C., & Watt, R. J. (1997). The computation of orientation statistics from visual texture. *Vision Research*, *37*(22), 3181–3192. https://doi.org/10.1016/S0042-6989(97)00133-8

de Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, *22*(9), 764–779. https://doi.org/10.1016/j.tics.2018.06.002

Di Russo, F., Martínez, A., & Hillyard, S. A. (2003). Source analysis of event-related cortical activity during visuo-spatial attention. *Cerebral Cortex*, *13*(5), 486–499. https://doi.org/10.1093/cercor/13.5.486

Fougnie, D., Cormiea, S. M., Zhang, J., Alvarez, G. A., & Wolfe, J. M. (2015). Winter is coming: How humans forage in a temporally structured environment. *Journal of Vision*, *15*(11), 1. https://doi.org/10.1167/15.11.1

Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, *3*, 1229. https://doi.org/10.1038/ncomms2237

Fritsche, M., Mostert, P., & de Lange, F. P. (2017). Opposite effects of recent history on perception and decision. *Current Biology*, *27*(4), 590–595. https://doi.org/10.1016/j.cub.2017.01.006

Geisler, W. S. (2011). Contributions of ideal observer theory to vision research. *Vision Research*, *51*(7), 771–781. https://doi.org/10.1016/j.visres.2010.09.027

Gilboa, I. (2015). Rationality and the Bayesian paradigm. *Journal of Economic Methodology*, *22*(3), 312–334. https://doi.org/10.1080/1350178X.2015.1071505

Hon, N., & Jabar, S. B. (2018). Learning in the target prevalence tffect. *Perception*, *47*(7), 789–798. https://doi.org/10.1177/0301006618773081

Hon, N., Yap, M. J., & Jabar, S. B. (2013). The trajectory of the target probability effect. *Attention, Perception, & Psychophysics*, *75*(4), 661–666. https://doi.org/10.3758/s13414-013-0429-x

Honig, M., Ma, W. J., & Fougnie, D. (2020). Humans incorporate trial-to-trial working memory uncertainty into rewarded decisions. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *117*, 8391–8397. https://doi.org/10.1073/pnas.1918143117

Jabar, S. B., & Anderson, B. (2015). Probability shapes perceptual precision: A study in orientation estimation. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(6), 1666–1679. https://doi.org/10.1037/xhp0000121

Jabar, S. B., Filipowicz, A., & Anderson, B. (2017). Tuned by experience: How orientation probability modulates early perceptual processing. *Vision Research*, *138*, 86–96. https://doi.org/10.1016/j.visres.2017.07.008

Jabar, S., & Fougnie, D. (2020). *How do expectations change behavior? Investigating the contributions at encoding versus decision-making*. https://osf.io/cnzmj/

Kiyonaga, A., & Egner, T. (2013). Working memory as internal attention: Toward an integrative account of internal and external selection processes. *Psychonomic Bulletin & Review*, *20*(2), 228–242. https://doi.org/10.3758/s13423-012-0359-y

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, *36*, 1–16. https://pdfs.semanticscholar.org/04d4/7572cec08b7a582a9366e5ac61dcfd633f2a.pdf

Kok, P., Brouwer, G. J., van Gerven, M. A., & de Lange, F. P. (2013). Prior expectations bias sensory representations in visual cortex. *The Journal of Neuroscience*, *33*(41), 16275–16284. https://doi.org/10.1523/JNEUROSCI.0742-13.2013

Kok, P., Failing, M. F., & de Lange, F. P. (2014). Prior expectations evoke stimulus templates in the primary visual cortex. *Journal of Cognitive Neuroscience*, *26*(7), 1546–1554. https://doi.org/10.1162/jocn_a_00562

Kok, P., Jehee, J. F., & de Lange, F. P. (2012). Less is more: Expectation sharpens representations in the primary visual cortex. *Neuron*, *75*(2), 265–270. https://doi.org/10.1016/j.neuron.2012.04.034

Kristjánsson, A., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention,* *Perception, & Psychophysics*, *72*(1), 5–18. https://doi.org/10.3758/APP.72.1.5

Leber, A. B., & Egeth, H. E. (2006). Attention on autopilot: Past experience and attentional set. *Visual Cognition*, *14*(4-8), 565–583. https://doi.org/10.1080/13506280500193438

Ling, S., Liu, T., & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Research*, *49*(10), 1194–1204. https://doi.org/10.1016/j.visres.2008.05.025

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*, 1432–1438. https://doi.org/10.1038/nn1790

Maloney, L. T., & Mamassian, P. (2009). Bayesian decision theory as a model of human visual perception: Testing Bayesian transfer. *Visual Neuroscience*, *26*(1), 147–155. https://doi.org/10.1017/S0952523808080905

Mamassian, P., & Landy, M. S. (2001). Interaction of visual prior constraints. *Vision Research*, *41*(20), 2653–2668. https://doi.org/10.1016/S0042-6989(01)00147-X

Meijs, E. L., Slagter, H. A., de Lange, F. P., & van Gaal, S. (2018). Dynamic Interactions between top–down expectations and conscious awareness. *The Journal of Neuroscience*, *38*, 2318–2327. https://doi.org/10.1523/JNEUROSCI.1952-17.2017

Morey, R. D., Rouder, J. N., & Jamil, T. (2015). Package 'BayesFactor'. http://cran/r-projectorg/web/packages/BayesFactor/BayesFactor

Norton, E. H., Acerbi, L., Ma, W. J., & Landy, M. S. (2019). Human online adaptation to changes in prior probability. *PLoS Computational Biology*, *15*, e1006681. https://doi.org/10.1371/journal.pcbi.1006681

O'Brien, J. L., & Raymond, J. E. (2012). Learned predictiveness speeds visual processing. *Psychological Science*, *23*(4), 359–363. https://doi.org/10.1177/0956797611429800

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*(2), 225–237. https://doi.org/10.3758/PBR.16.2.225

Rungratsameetaweemana, N., & Serences, J. T. (2019). Dissociating the impact of attention and expectation on early sensory processing. *Current Opinion in Psychology*, *29*, 181–186. https://doi.org/10.1016/j.copsyc.2019.03.014

Rungratsameetaweemana, N., Itthipuripat, S., Salazar, A., & Serences, J. T. (2018). Expectations do not alter early sensory processing during perceptual decision-making. *The Journal of Neuroscience*, *38*(24), 5632–5648. https://doi.org/10.1523/JNEUROSCI.3638-17.2018

Rutman, A. M., Clapp, W. C., Chadick, J. Z., & Gazzaley, A. (2010). Early top-down control of visual processing predicts working memory performance. *Journal of Cognitive Neuroscience*, *22*(6), 1224–1234. https://doi.org/10.1162/jocn.2009.21257

Schoups, A., Vogels, R., Qian, N., & Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. *Nature*, *412*(6846), 549–553. https://doi.org/10.1038/35087601

Serences, J. T., & Boynton, G. M. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, *55*(2), 301–312. https://doi.org/10.1016/j.neuron.2007.06.015

Shams, L., Iwaki, S., Chawla, A., & Bhattacharya, J. (2005). Early modulation of visual cortex by sound: An MEG study. *Neuroscience Letters*, *378*(2), 76–81. https://doi.org/10.1016/j.neulet.2004.12.035

Shen, S., & Ma, W. J. (2019). Variable precision in visual perception. *Psychological Review*, *126*(1), 89–132. https://doi.org/10.1037/rev0000128

Shurygina, O., Kristjánsson, Á., Tudge, L., & Chetverikov, A. (2019). Expectations and perceptual priming in a visual search task: Evidence from eye movements and behavior. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(4), 489–499. https://doi.org/10.1037/xhp0000618

Sreenivasan, K. K., & Jha, A. P. (2007). Selective attention supports working memory maintenance by modulating perceptual processing of distractors. *Journal of Cognitive Neuroscience*, *19*(1), 32–41. https://doi.org/10.1162/jocn.2007.19.1.32

Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, *13*(10), 9. https://doi.org/10.1167/13.10.9

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. https://doi.org/10.1038/nrn3838

Summerfield, C., & Egner, T. (2016). Feature-based attention and feature-based expectation. *Trends in Cognitive Sciences*, *20*(6), 401–404. https://doi.org/10.1016/j.tics.2016.03.008

Summerfield, C., & Koechlin, E. (2010). Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. *Frontiers in Human Neuroscience*, *4*, 208. https://doi.org/10.3389/fnhum.2010.00208

Vadillo, M. A., Linssen, D., Orgaz, C., Parsons, S., & Shanks, D. R. (2020). Unconscious or underpowered? Probabilistic cuing of visual attention. *Journal of Experimental Psychology: General*. Advance online publication. https://doi.org/10.1037/xge0000632

van den Berg, R., Shin, H., Chou, W. C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(22), 8780–8785. https://doi.org/10.1073/pnas.1117465109

Wei, X. X., & Stocker, A. A. (2015). A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nature Neuroscience*, *18*(10), 1509–1517. https://doi.org/10.1038/nn.4105

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235. https://doi.org/10.1038/nature06860

Zhou, Y. J., Pérez-Bellido, A., Haegens, S., & de Lange, F. P. (2020). Perceptual expectations modulate low-frequency activity: A statistical learning magnetoencephalography study. *Journal of Cognitive Neuroscience*, *32*(4), 691–702. https://doi.org/10.1162/jocn_a_01511

Received March 19, 2020
Revision received April 26, 2021
Accepted April 27, 2021 ■