Supplementary Information

Supplemental Results, Experiment 2 Binning Analysis

To determine if faster responses (*i.e.*, responses made earlier in the accumulation of visual information) produced larger cueing effects in Experiment 2, we split the data into two bins, based on the midpoint of the 500ms response deadline. This splitting produced a bin with RTs <= 250ms (faster bin) and a bin with RTs > 250ms (slower bin). The results of a 2x2x2 ANOVA with cue-validity, cue-onset-probability, and RT bin as repeated-measures factors are reported in Table S1. We found a significant three-way interaction ($F(1,95) = 4.25$, $p = 0.042$), with a significantly larger cueing effect difference in the faster bin compared to that of the slower bin.

	Proportion Correct	
	***F* (1, 95)**	***p***
Cue validity	**76.914**	**< 0.001**
Cue probability	0.618	0.434
Response time bin	**58.296**	**< 0.001**
validity X probability	**8.699**	**0.004**
validity X bin	0.338	0.562
probability X bin	0.121	0.729
validity X probability X bin	**4.245**	**0.042**

Table S1. ***Statistics, Experiment 2 binning analysis***. 2x2x2 ANOVA with cue-validity, cue-onset-probability, and response time bin as repeated-measures factors.

To follow up on this significant three-way interaction, we assessed the cue-validity X cue-onset-probability effect separately for each RT bin. We found a significant cue-validity X cue-onset-probability interaction in the faster bin (mean Δ Δ proportion correct = 0.0424, $t(95) = 2.82$, two-tailed $p = 0.0058$), but not in the slower bin (mean Δ Δ proportion correct = 0.0086, $t(95) = 1.12$, two-tailed $p = 0.2640$). To better understand the observed interaction in the faster bin, we first assessed the cue-validity effect separately for each cue-onset-probability. We found significant cueing effects for both low-probability cues (mean Δ proportion correct = 0.0817, $t(95) = 5.31$, two-tailed $p < 0.0001$) and high-probability cues (mean Δ proportion correct = 0.0392, $t(95) = 4.80$, two-tailed $p < 0.0001$). We then compared the impact of cue-onset-probability separately for each cue-validity condition within the faster bin. Aligning with Experiment 1 and the original analysis of Experiment 2, we found a significant effect of onset probability for valid cues (mean Δ proportion correct = 0.0230, $t(95) = 2.49$, two-tailed $p = 0.0147$), but no evidence for an impact of onset probability for invalid cues (mean Δ proportion correct = -0.0195, $t(95) = -1.61$, two-tailed $p = 0.1113$).

Supplemental Results, Experiment 2 Cue-Absent Comparisons

Here, we compared RT (mean of correct trials) and proportion correct between the cue-absent condition and each validity condition of Experiment 2. Relative to cue-absent trials, participants were significantly faster on both valid (mean Δ RT = -14.72ms, $t(95) = -16.29$, two-tailed $p < 0.0001$) and invalid trials (mean Δ RT = -13.25ms, $t(95) = -16.39$, two-tailed $p < 0.0001$). The shared directionality of this effect, with faster responses in both validity conditions relative to absent, suggests that the cues provided temporal information about upcoming target presentation. However, participants had significantly worse proportion correct scores relative to cue-absent trials in both the valid (mean Δ proportion correct =

-0.0104, $t(95) = -2.94$, two-tailed $p = 0.0041$) and invalid (mean Δ proportion correct = -0.0582, $t(95) = -15.39$, two-tailed $p < 0.0001$) conditions. This pattern of results, with faster yet less accurate responses in cue-presented conditions, aligns with a classic speed-accuracy tradeoff.

Supplemental Results, Additional Online Experiments

In Tables S2-S5, we present the results from four additional online experiments. Experiment 3 (N=132) required responses within 700ms. Experiment 4 (N=45) allowed participants to respond freely (but trials timed out after 20s). The final two experiments had a response *delay* of 1067ms. In Experiment 5 (N=47), this delay was unenforced, with early attempts to respond having no effect on the trial sequence. In Experiment 6 (N=45), this delay was enforced, with early response attempts causing the trial to end early, followed by a 1500ms reminder that participants must wait to make a response.

In sum, the results are consistent with what we observed in Experiment 1: no consistent cue-validity X cue-onset-probability interactions when RTs were allowed or required to be made after 500ms. In the accuracy domain, we observed main effects of cue validity in all four experiments, with no evidence for a main effect of cue probability or validity X probability interaction. In the RT domain, we did observe significant validity X probability interactions in Experiment 5 and Experiment 6, but they ran in opposite directions. In Experiment 5, the mean change in RT (valid minus invalid) for low-probability cues was 1.1481 ms, and the mean change in RT (valid minus invalid) for high-probability cues was -3.8334 ms. In Experiment 6, the mean change in RT (valid minus invalid) for low-probability cues was -13.0673 ms, and the mean change in RT (valid minus invalid) for high-

probability cues was -7.9938 ms. In both Experiment 5 and Experiment 6, participants were required to delay their responses, complicating the interpretation of any differences in RT.

	Proportion Correct		Response Time	
	***F* (1, 131)**	***p***	***F* (1, 131)**	***p***
Cue validity	**100.7**	**< 0.0001**	**90.13**	**< 0.0001**
Cue probability	0.282	0.597	0.273	0.602
validity X probability	0.028	0.866	0.234	0.629

Table S2. ***Statistics, Experiment 3, 700ms response deadline***. 2x2 ANOVA with cue-validity and cue-probability as repeated-measures factors.

	Proportion Correct		Response Time	
	***F* (1,44)**	***p***	***F* (1,44)**	***p***
Cue validity	**36.78**	**< 0.0001**	**75.74**	**< 0.0001**
Cue probability	0.77	0.385	0.966	0.331
validity X probability	0.644	0.426	0.482	0.491

Table S3. ***Statistics, Experiment 4, free-response.*** 2x2 ANOVA with cue-validity and cue-probability as repeated-measures factors.

	Proportion Correct		Response Time	
	***F* (1, 46)**	***p***	***F* (1, 46)**	***p***
Cue validity	**65.86**	**< 0.0001**	0.719	0.401
Cue probability	1.909	0.174	0.001	0.97
validity X probability	0.249	0.62	**4.429**	**0.0408**

Table S4. ***Statistics, Experiment 5, 1067ms response delay (unenforced)***. 2x2 ANOVA with cue-validity and cue-probability as repeated-measures factors.

	Proportion Correct		Response Time	
	***F* (1, 44)**	***p***	***F* (1, 44)**	***p***
Cue validity	**23.69**	**< 0.0001**	**39.27**	**< 0.0001**
Cue probability	1.658	0.205	0.221	0.641
validity X probability	0	0.988	**4.207**	**0.0463**

Table S5. ***Statistics, Experiment 6, 1067ms response delay (enforced)***. 2x2 ANOVA with cue-validity and cue-probability as repeated-measures factors.

Supplemental Results, Prior Predictive Check.

Before fitting the model to the Experiment 2 data, we want to ensure that the priors we have chosen do not privilege any particular outcome. We do this via a prior predictive check in which we take 1000 random draws from each prior distribution in line 4 (*i.e.*, from δ_0, δ_1, δ_2, and δ_3 for each individual, which yields 4*96=384 distributions of 1000 parameters each). We use the resulting 384 distributions to compute *d* for each experimental condition, for each individual. We then compute the interaction term for each individual and plot the 95% credible interval calculated as the highest posterior density interval (HPDI) of the 1000 samples (**Figure S1**). Thus before seeing the data, the model allows for a wide range of possible interaction magnitudes and does not favor a particular direction (*i.e.*, the means of each distribution are centered at zero).

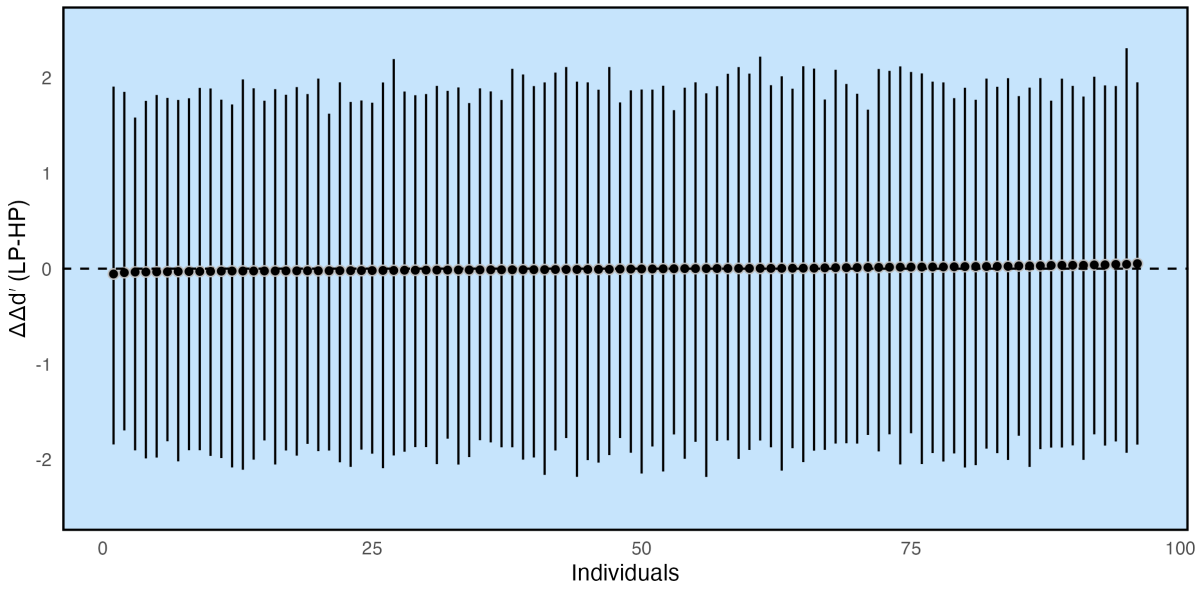

Figure S1. ***Prior Predictive Check, Model 1.*** Black lines indicate the bounds of the 95% highest posterior density interval (HPDI). Circles mark the mean of each distribution, and individuals are plotted in ascending order. See text for additional details.

Supplemental Results, Posterior Predictive Check.

The dashed line in Figure 6 marks the mean interaction term computed directly from the Experiment 2 data, which falls neatly in the center of the model-based interaction distribution. This suggests that our data are well described by the model. Another way to assess model fit is to perform a *posterior predictive check*. We know *numTrials* for each condition, and we have a distribution of posterior parameter estimates for each condition. We can then make a random draw from a binomial distribution to generate a simulated value for *numCCW*. Using the simulated values for each condition, we compute the interaction term for each individual and plot the 95% HPDI (**Figure S2**). For all the participants, the actual interaction term, computed directly from their respective hit and false alarm rates, are indicated with circles and all fall within the bounds of the 95% HPDIs. This means that data simulated from the updated model aligns well with what was actually observed.

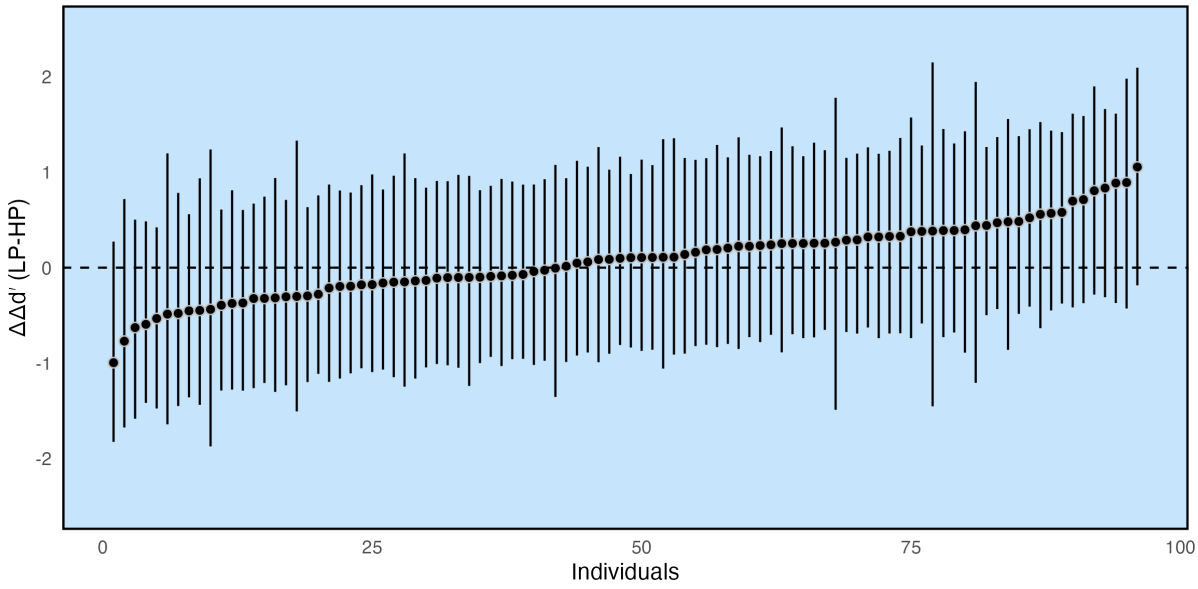

Figure S2. ***Posterior Predictive Check, Model 1.*** Black lines indicate the bounds of the 95% highest posterior density interval (HPDI). Circles mark the empirically observed interaction value, with individuals plotted in ascending order. See text for additional details.

Supplemental Methods, Experiment 1 Minutiae

- *Fixation squares.* The height and width of the fixation squares were 0.25 DVA; each was presented at 0.25 DVA from the screen center on the horizontal meridian.
- *Exogenous cue*. The exogenous cue was a white circle, with a radius of 0.25 DVA, presented 2.5 DVA above the center of the subsequent target grating (valid trial) or distractor grating (invalid trial).
- *Target and Distractor Stimuli*. The target and distractor stimuli were sinusoidally-modulated, contrast-defined gratings in a Gaussian window (spatial frequency, 2 cycles per degree; contrast, 50%, diameter, 4 degrees of visual angle, DVA) presented at 6 degrees of visual angle from the screen center along the horizontal meridian.
    - Each experimental session had 840 trials, split into 15 blocks of 56 trials. The 56 unique trial types (target side, L/R; cue validity, V/I; onset probability, H/L; response delay condition, 33/67/133/267/533/1067/2133 ms) were presented in a random order during each block. At the end of each block, the tilt of the gratings (i.e., the rotation from vertical in angular degrees) was updated based on the orientation discrimination accuracy in the block. The tilt got smaller if the overall accuracy in the block was greater than 80%; it got larger if the overall accuracy in the block was less than 80%. The magnitude of the change depended on the deviation from 80% correct, in line the the following formula:
        ✳ $Tilt_{new} = Tilt_{current} - (Tilt_{current} * (p - 0.8))$, where p refers to proportion correct in the previous 56 trials
    - The purpose of updating the tilt is to keep the task appropriately difficult. For the two co-authors, the initial tilt in session 1, block 1 was 10°. For the remaining five participants, the initial tilt in session 1, block 1 was 15°. For the two co-authors, who

were the first two participants, there were two additional modifications. After session 11, the stimulus presentation time was reduced from 100ms to 50ms. After session 12, the tilt updating formula used 70% correct as the reference point (i.e., the tilt got smaller if the overall accuracy in the block was greater than 70%; it got larger if the overall accuracy in the block was less than 70%). This was done to make the the task more difficult for these participants. These mid-experiment updates were not made for the subsequent participants.

Supplemental Methods, Experiment 2 Minutiae

Experimental Setting

- Participants were recruited via Prolific (using the settings below); the experiment was hosted on Pavlovia and completed in a web browser.
- Study Details:
  - Which devices can participants use to take your study? Desktop only
  - Location: "all countries available"
  - Study distribution: Standard sample
  - Prescreen participants:
    - Age: Minimum Age: 18, Maximum Age: 50
    - Fluent languages: English
    - Approval rate: Minimum Approval Rate: 80, Maximum Approval Rate: 100

Experimental Flow - Session 1

- Consent
- Experimental instructions
- Practice trials (16 trials)
- Experiment (960 trials)
  - 10 blocks of 96 trials
- Compensation

Experimental Flow - Sessions 2 and 3 (to be completed on separate days)

- Consent
- Instructions reminder
- Experiment (960 trials)
  - 10 blocks of 96 trials
- Compensation

Experimental Details

- 960 total experimental trials in each session
- The task: Report the orientation of the target line
  - 2AFC: CW of vertical, CCW of vertical
  - CCW: left arrow key
  - CW: right arrow key
- Two lines are presented on each trial:
  - Color, white. Length, 100 pixels. Location, 400 pixels to the left or right of a central fixation circle, presented for 100ms.
  - The target is post-cued by a small horizontal line (20 pixels) displayed either 25 pixels to the left or 25 pixels to the right of fixation point.
  - At the beginning of each session, each line is independently and randomly rotated 10° CW or CCW from vertical. After trial 57 (due to a coding error), the tilt magnitude was automatically increased to 20°.
  - The target appears on the left and on the right an equal number of times.
- Each trial has 6 segments:
  1. Fixation: 1000 ms
  - The fixation circle (radii = 10 pixels) appears at the center of a gray screen.
  - The fixation circle is either red or green. Each color appears an equal number of times.
  - At the beginning of the experiment, one color is selected to signal a high probability (0.8) that a peripheral cue (see below) will appear, with the other color signaling a low probability (0.2) that a peripheral cue will appear.
  2. Cue: 67 ms
  - The fixation circle appears at the center of a gray screen.
  - Cues are presented probabilistically in line with the cue-probability conditions above. This results in two kinds of trials: *cue-presented* trials and *cue-absent* trials.
    - On *cue-presented* trials, an abrupt-onset cue (white circle, radii = 10 pixels) presented at one of two locations: (400,70) or (-400,70) relative to screen center
    - On *cue-absent* trials, no cue appears.
    - At the beginning of the session, half the trials are assigned to have a “valid” cue and half are assigned to have an “invalid” cue.
      - When the cue appears on the same side as the target line, it is valid; when it appears on the same side as the distractor line, it is invalid.
      - When the cue is absent, it is neither valid nor invalid.
  3. Inter-stimulus interval (ISI): 33 ms
  - The fixation circle appears at the center of a gray screen.
  4. Stimuli: 100 ms
  - The fixation circle appears at the center of a gray screen.
  - Line stimuli are presented on both sides of fixation
  5. Response: 500 ms or until response, whichever is shorter
  - The fixation circle appears at the center of a gray screen, now rendered in white.
  - Response cue presented indicating which line is target, keyboard response allowed in time interval
  6. Feedback: 500 ms
  - If the response is correct, a 👍 emoji appears.

- If the response is incorrect, a 👎 emoji appears.
- If the response is not made within 500ms, a ⏱️ emoji appears.

7. Inter-trial interval (ITI): 500 ms

- The fixation circle appears at the center of a gray screen, now rendered in white.