Subjective and objective cue probability interact to shape perceptual decisions

Anonymous CogSci submission

Abstract

Expectations, or the prior probability of a particular choice outcome, are powerful sources of information for improving quick decisions made on the basis of noisy and incomplete evidence. Most existing research on expectations in perceptual decision making provide observers with perfect knowledge of prior probabilities, and thus fail to capture the experience-dependent uncertainty characteristic of expectations learned and used outside of the lab. Here, we report data from a novel expectation-guided perceptual decision making task specifically designed to overcome these limitations. We find that the effects of expectations learned through experience are modulated by the accuracy with which observers estimate the true probability of a cue, and that this estimation accuracy exhibits substantive variation across individuals. Additionally, we find that confidence ratings about the accuracy of a cue's probability are primarily driven by an observer's subjective estimate of the cue's probability rather than its objective probability as defined by experience.

Keywords: evidence accumulation; perceptual decision making; metacognition; statistical learning

Introduction

Expectations in perceptual decision making are defined as the prior probability that one of two choice outcomes is correct or will be rewarded. The standard approach for measuring effects of expectations in the lab either explicitly instructs humans about prior probability or trains non-human primates on tens of thousands of trials in order to ensure they have learned that probability (e.g., Hanks et al., 2011). While this learning of prior probability has many statistical advantages, it fails to capture the experience-dependent uncertainty inherent of perceptual expectations acquired outside of the lab. The formation of these naturalistic expectations requires aggregating across related experiences stored in episodic memory, and flexibly retrieving an estimate of the prior probability of a particular choice outcome in a context-specific manner (CO-AUTHOR and et al., 2023).

When the assumption of perfect prior knowledge is relaxed, a new question arises: is choice behavior guided by objective probabilities reflecting implicit knowledge based on error-based learning or subjective probabilities that declarative memory systems infer from experience? We have previously shown that dynamic effects of expectations can be explained by a dynamic reliability-weighted integration process, where reliability estimates are governed jointly by subjective and objective probability information (Authors, 2022). This paper presents a novel experimental task inspired by predictions of that model, and uses a series of linear regression analyses to begin revealing how different sources of probability information (objective and subjective) contribute to expectation-guided perceptual decisions.

Methods

We designed a novel cue-guided perceptual decision task that (i) required that observers learn expectations through experience and (ii) obtained direct measures of each cue's subjective predictive probability, along with a confidence rating in that estimate. Participants then used these learned cues as memory-based expectations while they performed an evidence accumulation based perceptual decision-making task.

Stimuli and participants

Stimuli Visual evidence consisted of two grayscale scene images. There were two sets of possible scene images (i.e., four images total), with pairs of images and their mappings to keyboard responses randomized across participants. The probability of a given image being the 'dominant' image in a "stream" of visual evidence displayed to the observer on each trial varied across three possible conditions, which were communicated to the observer via a colored border that circumscribed the visual evidence during the stream. There were two sets of possible borders, with each set comprised of of triadic colors (set1 = red, blue, yellow; set2 = orange, green, purple). The set of border colors, along with the borders' assignments to dominance probabilities for particular images in the visual evidence stream, were randomized across participants.

Participants A sample of 18 undergraduate students (age M=20.5 years; SD=1.47 years; 13 female, 3 male, 2 non-binary) were recruited from the authors' university. One subject was excluded from analysis because of a coding error in the experiment, resulting in a sample size of n=17 for all reported analyses. Participants were compensated for their time either with course credit or a pro-rated cash payment of $15/hour. This study was approved by the Institutional Review Board at the authors' University and all subjects provided written informed consent.

Task design

Data were collected in a single experimental session ranging between 60-90 minutes. Each session began by instructing participants on the mapping between the two scene images and the 1 and 2 number keys on a US keyboard.

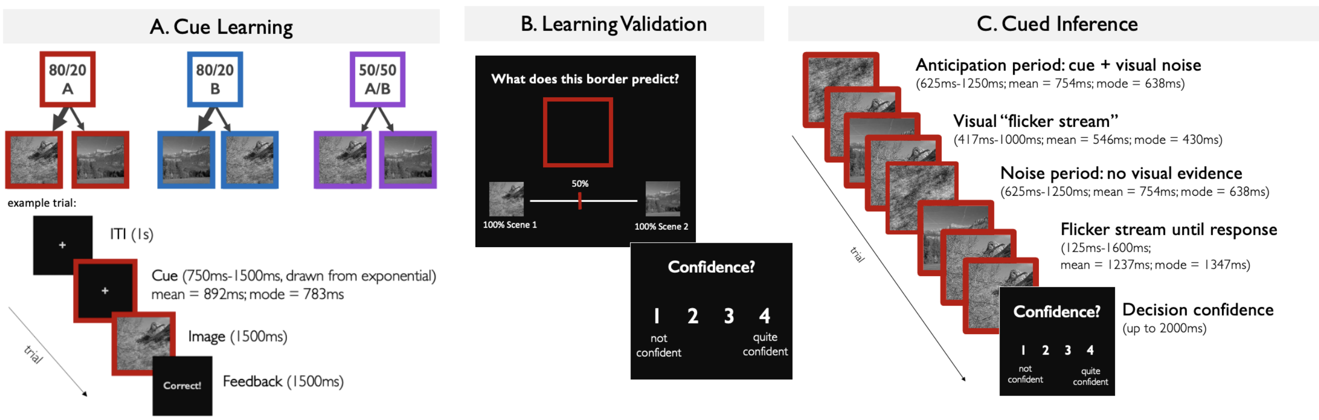

Figure 1: **Task design. (A)** Participants first learn that different colored borders make different predictions about the probability of observing one of two possible scene images. The objective probability of each cue is defined by the frequency with which it is followed by one of the two scene images. **(B)** After learning, participants use a slider to report their subjective estimate of each cue's probability followed by a confidence rating in that estimate. **(C)** In the last phase of the experiment, participants are presented with stochastic visual evidence inside of the colored borders. Their task to report which image is dominant (i.e., presented more frequently) in a 60Hz stream of visual evidence, followed by a confidence rating in that perceptual decision.

Calibration We used two interleaved QUEST staircases (Watson and Pelli, 1983) to identify values of visual evidence coherence that, for each image, resulted in 70% accuracy on the perceptual decision task (described in more detail below). Evidence coherence was defined as the proportion of signal frames in the visual evidence stream that contained the target stimulus. The calibration procedure ensured that decision difficulty in the Cued Inference phase would be identical for both target stimuli, effectively controlling for low-level visual differences between the images that might systematically bias choices toward one of the options. Participants completed 80 trials of the decision task (40 trials per staircase) during the calibration procedure and received feedback on their choices.

Cue Learning Next, participants learned the predictive probability of each cue (i.e., each colored border) by observing a series of cue-image pairings in which the cue was presented prior to the image it was probabilistically paired with (Figure 1A). To ensure active engagement—and to build up associative motor memories—participants were instructed to respond on each trial indicating which image appeared on screen after the cue using the previously-learned image-key mappings. Participants were told that there was a predictive relationship between the cues and scene images, and that their broader goal for this phase of the experiment was to learn that relationship. Finally, participants were also told that they were permitted to respond in the inter-stimulus interval (ISI) between the onset of the cue and scene image if they desired. Regardless of when participants responded, they received feedback on their response accuracy on each trial.

In order to maximally align learning and decision environments, we permitted the ISI between cue and image onset to vary across trials according to a truncated exponential distribution. This approach ensures a fixed hazard rate across learning trials, such that participants are maximally uncertain about the temporal onset of scene images across learning trials (CO-AUTHOR and collaborator, 2024). ISIs ranged from 750ms to 1500ms (mean = 892ms, mode = 783ms). After a scene image appeared on screen, participants had up to 1500ms to make their response. Post-trial feedback was displayed for 1500ms, and participants were told that they should respond faster on the next trial if the feedback screen appeared before they made a response.

Each cue was presented a total of 30 times and the order of cues was fully randomized, providing participants 90 total observations of cue-image pairings. The objective probability of each cue was defined as the frequency with which it preceded one of the two scene images. Two cues were randomly selected to have 80% predictive probabilities (i.e., one 80% cue for each image) and the third cue equally predicted both images (i.e., one 50% cue for both images).

Learning Validation Immediately after Cue Learning, participants were presented with each colored border and used a sliding scale to report their subjective estimates of each cue's predictive probability (Figure 1B). Their estimates were permitted to range from 50-100% and the slider was initialized to 50% on each trial. Both the subjective estimate and subsequent confidence rating (1-4; not confident-quite confident) were self-paced.

Cued Inference In the final phase of the experiment, participants observed a rapidly-alternating (60Hz) "stream" of the two scene images interleaved with pure noise frames (phase-scrambled superpositions of the images) (Figure 1C). Observers' task was to report which of the two scene images was presented more frequently (i.e., was the "target") on each trial. The proportion of target frames on each trial was defined by the calibrated coherence value for that target's trial, as estimated during the preceding Calibration phase. Crucially, this visual evidence was presented *inside* the colored borders, ensuring that information about the prior probability was always accessible to the observer. Participants were told that the predictive relationships they just learned between the colored borders and scene images also applied in this phase of the experiment (i.e., "the correct answer is usually the one predicted by the cue"). They were also instructed to respond as quickly and accurately as possible. However, because we also elicited decision confidence ratings on each trial, participants did not get feedback about their choice accuracy.

We incorporated two periods of stochastic visual noise into each Cued Inference trial. The durations of these noise periods, as well as the brief signal period in between them, were all drawn from separate truncated exponential distributions in order to guarantee a fixed hazard rate across trials (CO-AUTHOR and collaborator, 2024). The maximum duration of any trial was 3333ms, and any remaining time after the second noise period consisted of threshold-level visual evidence. Immediately after making a decision, participants had up to 3000ms to report their confidence in that decision's accuracy on a scale of 1-4 (not confidence-quite confident). Trials were separated by 1000ms intertrial interval.

Each subject completed 150 trials for each 80% cue and 75 trials for the 50% cue, thus completing 375 trials in total. The assignment of cue and target was fully randomized across trials, with the probability of a scene image being a target for a particular cue being defined by that cue's true probability. This means that, for 20% of trials with an 80% cue, the cue was *incongruent* with respect to the true trial target (i.e., its effective prediction was 0.2).

Analyses

Software All behavioral data were analyzed using R version 4.4.2. Regression models were fit using the 'lme4' package. Statistical tests on fitted coefficient values were performed using the 'lmerTest' package, and BIC values for fitted models were obtained using the 'compare_performance()' function from the 'performance' package in R. Specific models are described below.

Definitions Our Results section makes use of a number of different variable names. For clarity and convenience, we list each variable's definition here:

- *trueCue*: a cue's true/objective probability; takes on values 0.5 and 0.8
- *congruence*: the (mis)match between the target predicted by a cue and the true target of visual evidence on each trial; takes on values *congruent*, *incongruent*, and *neutral* (in the case of 50% cues)
- *subjectiveCue*: an observer's reported estimate of a cue's true probability; takes on values 0.5-1
- *cueConfidence*: an observer's confidence rating in their *subjectiveCue* report; takes on values 1-4
- *cueDiff*: a signed difference between *subjectiveCue* and *trueCue*
- *cueCorr*: a within-subject correlation coefficient between *subjectiveCue* and *trueCue*

Results

Our analyses aimed to answer the following questions:

1. **How well are subjects able to learn the predictive power of the cues during Cue Learning?** We answer this by examining the relationship between subjects' reported cue probabilities and their objective probabilities, as well as the influence of these factors on subjects' confidence in their learned probabilities.
2. **How does subjects' learning of the cues' predictiveness impact their choice behavior?** We answer this by examining both choices made by observers and their reaction times (RTs) during the Cued Inference phase, with reference to individual differences in the answer to Question 1.

Learning Validation: Accuracy and Confidence of Cue Estimates

Accuracy of cue estimates Figure 2 depicts cue- and observer-specific metrics of cue estimate accuracy. We computed the cue-specific metric (*cueDiff*) as the signed difference between a reported cue probability (*subjectiveCue*) and that cue's true objective probability (*trueCue*). As shown in Figure 2A, participants tended to overestimate the probability of the 50% cue while underestimating the probability of the 80% cues. A simple linear regression confirmed the statistical significance of this effect (β = -0.70882, t = -5.756, $p < .001$). We computed an observer-specific metric (*cueCorr*) as the linear correlation, across all cues, between the subjective cue probability (*subjectiveCue*) and its true probability (*trueCue*). Figure 2B illustrates the heterogeneity in cue estimation accuracy across participants. Whereas some participants reported subjective estimates that perfectly matched the true probability of the cue, others systematically mis-estimated probabilities across all of the cues.

Confidence in cue estimate Next, to investigate how subjective and objective probability information shape confidence in observers' cue estimates, we examined the success of 4 nested linear models in predicting *cueConfidence* (Table 1). The winning model used only a main effect of *subjectiveCue* to predict estimate confidence, and returned a significant main effect (β = 4.52, t = 5.968, $p < .001$). This suggests that confidence in a subjective cue estimate is primarily driven by the actual subjective estimate itself rather than the true objective cue probability.

To complement this analysis into the probability information shaping *cueConfidence* judgments, we investigated how our two measures of cue estimate accuracy–*cueDiff* and

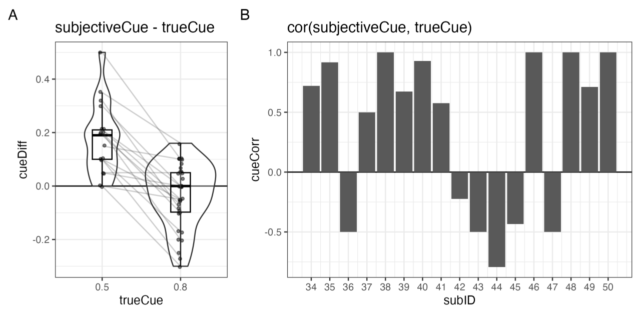

Figure 2: **Accuracy of subjective cue estimates. (A)** Cue-specific estimate accuracy (*cueDiff*) was better for 80% cues than 50% cues. **(B)** The cross-cue correlation of an observer's subjective probability estimate and the true probability of a cue (i.e., the *cueCorr* metric) across participants.

Effects Structure	BIC
trueCue	133.9
subjectiveCue	**110.7**
$trueCue + subjectiveCue$	114.1
$trueCue * subjectiveCue$	115.5

Table 1: **Comparing subjective and objective cue probability as predictors of cue estimate confidence**. The winning linear model (bolded) uses only subjective probability (*subjectiveCue*) to predict *cueConfidence*

cueCorr–related to confidence in subjective cue estimates. We fit a linear model that predicted *cueConfidence* using an interaction between *cueDiff* and *cueCorr*. Importantly, *cueDiff* is a predictor capturing the *difference* between a cue's true probability and the observer's estimate thereof, meaning that values closer to 0 indicate a higher degree of accuracy in subjective cue estimates. The model returned a a trending main effect of *cueDiff* only ($\beta_{cueDiff}$ = 1.50, t = 1.79, $p = .08$), indicating that confidence ratings in cue estimates increased linearly with the value of that cue estimate irrespective of how accurate the estimate actually was.

Cued Inference: Choice Behavior and Timing

We now turn to investigating how objective and subjective cue probability information contributed to choice behavior during Cued Inference. Here, we split our analyses to examine choice behavior (a) on the basis of expectations alone (i.e., during the anticipation period before visual evidence onset; see Methods & Figure 1), and (b) when memory-based expectations were integrated with incoming sensory information once visual evidence became available. In both cases, we examined both the timing and accuracy of choices.

Responses driven by expectations alone It has previously been shown that when a cue signals that upcoming visual evidence will be weakly informative *and* that the prior probability of a particular target is high (i.e., 80%), observers make

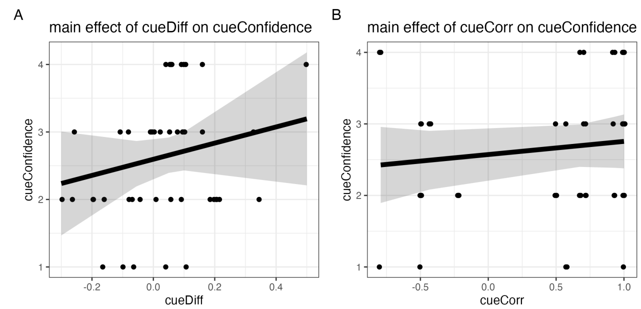

Figure 3: **Accuracy-based predictors of cue estimate confidence.** Shaded regions correspond to standard errors of model estimates. Individual points correspond to individual participant data. **(A)** Subjects' errors in estimating each individual cues' predictiveness (*cueDiff*) marginally predicted confidence in subjective cue estimates. **(B)** The overall correlation between estimated and true predictive probabilities (*cueCorr*) failed to predict *cueConfidence*.

a perceptual decision before observing any visual evidence at all (CO-AUTHOR and et al., 2023), which here we term an *early response*. Additionally, this effect was found to scale with the amount of elapsed time between the onset of a prior cue and visual evidence: as more time elapsed, observers became more likely to make an early response (CO-AUTHOR and et al., 2023).

We performed two model comparison analyses to investigate which of the five factors investigated above (*trueCue*, *subjectiveCue*, *cueConfidence*, *cueCorr*, and *cueDiff*) best predicted the probability of making an early response during the anticipation period (see Methods & Figure 1). Because this outcome variable is binary (i.e., a response was made during this period or not), we fit two sets of generalized linear models with family='binomial'.

The first round of model comparisons used each of these factors alone as a predictor. This comparison revealed that the observer-specific cue accuracy metric, *cueCorr*, was the best predictor of early responding in both sets of comparisons (Table 2, top comparison set).

The second round of model comparisons permitted each factor to interact with the duration of the anticipation period (*anticipationDuration*). This analysis revealed that the model permitting an interaction between *cueCorr* and *anticipationDuration* had the lowest BIC of all of the models—including all models in the previous comparison set—indicating that early responses are best captured by an interaction between observer-specific cue accuracy and the duration of the anticipation period (Table 2, bottom comparison set). This winning model returned a significant main effect of observer-specific estimate accuracy ($\beta_{cueCorr}$ = 13.574, z = 7.073, $p <.001$), a significant main effect of anticipation duration ($\beta_{anticipationDuration}$ = 0.159, z = 6.463, $p <.001$), and a significant interaction ($\beta_{cueCorr*anticipationDuration}$ = -0.133, z

Single-Factor Predictor	BIC
trueCue	2192.6
subjectiveCue	2182.6
cueConfidence	2146.0
cueCorr	**1820.6**
cueDiff	2231.0
Interaction Predictors	
$trueCue * anticipationDuration$	2176.7
$subjectiveCue * anticipationDuration$	2167.7
cueConfidence ∗ anticipationDuration	2138.1
$cueCorr * anticipationDuration$	**1786.0**
$cueDiff * anticipationDuration$	2217.0

Table 2: **Comparing predictors of early responses**. In the first comparison set (top), *cueCorr* was the best predictor. However, analyses permitting interactions with *anticipationDuration* (bottom) revealed that an interaction between this factor and *cueCorr* predicts probability of early response better than any other factor or combination. Winning models in each comparison set are bolded.

= -5.050, $p < .001$). Back-transformed model estimates for the interaction term are displayed graphically in Figure 4A, which shows that observers who better estimated the true cue probability (i.e., those with a *cueCorr* value closer to 1) were significantly more likely to make early responses. For trials at the extreme end of the anticipation duration distribution, however, this difference no longer remains significant.

Next, we investigated the probability that observers made a *correct* choice during the anticipation period. Accuracy in this case is defined as reporting the scene that *would have been* the target on that trial. We modeled the binary outcome variable using additive main effects of *trueCue*, *subjectiveCue*, and *cueCorr*. Of these predictors, *trueCue* exhibited the only trending effect on (visual-evidence blind) choice accuracy ($\beta_{trueCue}$ = 4.0338, z = 1.684, p = 0.092). Figure 4B shows how the effect of *trueCue* on early response accuracy corresponds to probability-matching behavior during the anticipation period.

Responses integrating memory and sensory evidence We next investigated the factors that best predict choices and reaction times (RTs) on trials when observers waited until after the onset of visual evidence to make a choice. Based on our findings displayed in Figure 4A, we investigated first whether *cueCorr* also fared best in predicting in response times on trials where observers did not respond early. To narrow down the space of possible comparisons, we focused on three conceptually-motivated model structures. We compared the predictive success of *cueCorr* alone against an interaction between *cueCorr* and *trueCue* and against an interaction between *cueCorr* and *cueConfidence*. Table 3 shows that RTs for choices made after the anticipation period were best predicted by this latter model, which returned a significant main effect of *cueCorr* ($\beta_{cueCorr}$ = 0.30874, t = 5.109, $p < .001$),

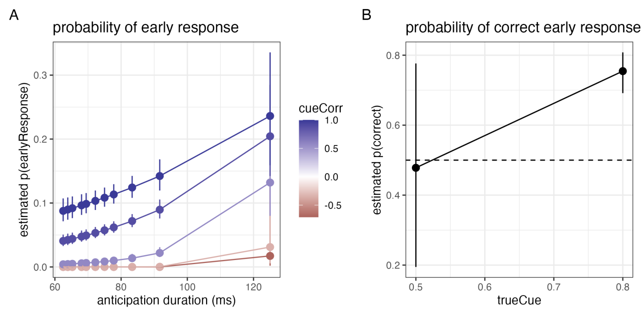

Figure 4: **Choice behavior driven by expectations alone.** rror bars correspond to standard errors of model estimates. **(A)** A linear model revealed a significant interaction between anticipation duration and *cueCorr* when predicting early responding. **(B)** The probability of making a correct response based on expectations alone reveals probability-matching behavior.

a significant main effect of *cueConfidence* ($\beta_{cueConfidence}$ = -0.06916, t = -4.719, $p < .001$), and a significant interaction ($\beta_{cueCorr*cueConf}$ = -0.06902, t = -3.363, p =.0001). Figure 5 graphically displays the interaction effect, which reveals that participants with greater *cueCorr* values generally took longer to make responses after the onset of visual evidence, especially when presented with cues whose estimates they were not confident in.

Single-Factor Predictor	BIC
cueCorr	17933.6
$cueCorr * trueCue$	17914.3
cueCorr ∗ cueConfidence	**16970.5**

Table 3: **Comparing predictors of RTs made after the anticipation period**. BIC comparisons reveal that RTs based on expectations and visual evidence are best predicted with an interaction between subjective estimate accuracy (*cueCorr*) and confidence in that estimate (*cueConfidence*)

Finally, we investigated how *cueCorr*, *trueCue*, and *cueConfidence* impact the accuracy of choices made after the onset of visual evidence. Previous analyses further motivated us to include *congruence*—the (mis)match between the true target on a trial and the target predicted by the objective cue—as an additional factor in this analysis. We began by comparing the performance of 9 different fixed effects structures in predicting binary choice accuracy on each trial where observers responded after the onset of visual evidence (Table 4). Despite its complexity, the model using a 3-way interaction among *cueCorr*, *cueConfidence*, and *congruence* provided the most parsimonious fit to accuracy for choices made after the onset of visual evidence.

The model returned a significant main effect of *cueCorr* (β = -0.86621, z = -2.257, $p < .001$), a significant effect of cue-evidence congruence (β = -0.72637, z = 4-2.235, p =.026), a

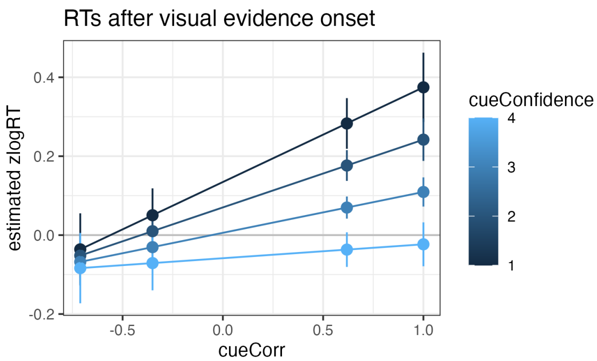

Figure 5: **RTs based on expectations and visual evidence**. Observers who made better overall estimates of cue probabilities (i.e., had high *cueCorr* values) generally took longer to respond after the onset of visual evidence, especially when on trials containing cues for which they were not confident in their estimates. Error bars correspond to standard errors of model estimates.

Effect Structure	BIC
trueCue	6619.2
subjectiveCue	6618.8
cueCorr	6585.4
cueConfidence	6218.4
trueCue ∗ *congruence*	6547.5
subjectiveCue ∗ *congruence*	6538.8
cueConfidence ∗ *congruence*	6084.1
trueCue ∗ *cueConfidence* ∗ *congruence*	6084.1
cueCorr*cueConfidence*congruence	**6038.1**

Table 4: **Comparing predictors of accuracy for choices made after visual evidence onset**. The winning model is bolded.

significant effect of cue-evidence incongruence (β = 0.86165, z = 2.105, p =.035). It also returned significant interactions between *cueCorr* and cue-evidence congruence (β = -0.97683, z = -2.164, p =.030), between *cueConfidence* and cue-evidence congruence (β = 0.30221, z = -2.164, p =.030), as well as a significant interaction between *cueConfidence* and cue-evidence *in*congruence (β = -0.52307, z = -3.562, p =.0004). Finally, the model returned a significant three-way interaction among *cueCorr*, *cueConfidence*, and cue-evidence congruence (β = 0.38579, z = 2.373, p =.018), which we display graphically in Figure 6.

The significant three-way interaction (central panel of Figure 6) can be interpreted as follows. When visual evidence indicates that the cue's prediction is correct on that trial (i.e., when cue & evidence are *congruent*), observers with higher *cueCorr* values become much more likely to make a correct response if presented with a cue whose probability they were confident in estimating. However, observers who did not display such sensitivity to the alignment of the subjective estimates with the cue's true probability (i.e., those with lower *cueCorr* values) were equally likely to make responses with high accuracy regardless of how confidently they estimated the probability of that trial's cue.

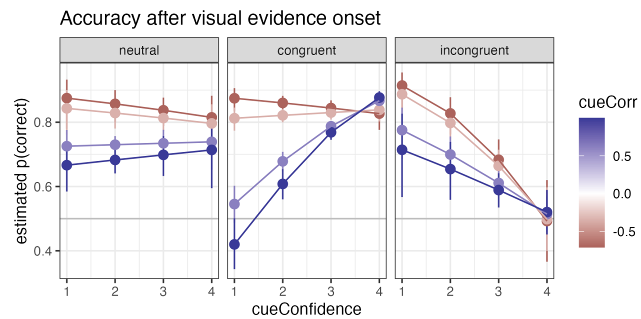

Figure 6: **Choice accuracy based on expectations and visual evidence**. Plot panels correspond to different configurations of cue-evidence congruence. Error bars correspond to standard errors of model estimates. An interpretation is provided in the main text.

Discussion

We measured human behavior on a novel expectation-guided perceptual decision task that required observers to learn expectations from experience, make explicit reports about their estimates of each cue's probability, and then use those cues to make perceptual decisions under uncertainty. Our findings indicate that, whereas choice behavior is driven by an interaction of objective and subjective cue probability, confidence ratings in subjective cue probabilities are primarily driven by the magnitude of the estimated probability itself. To the degree that subjective and objective estimates are generated by explicit and implicit memory systems, respectively, our results suggest that these two systems are integrated in the service of adaptively biasing choice behavior.

References

Authors. (2022). Precision-weighted evidence integration predicts time-varying influence of memory on perceptual decisions. *Cognitive Computational Neuroscience*.

CO-AUTHOR & collaborator. (2024). Achieving constant hazard rate in psychophysics experiments. *PsyArXiv*.

CO-AUTHOR & et al. (2023). Associative memory retrieval modulates upcoming perceptual decisions. *Cogn. Affect. Behav. Neurosci.*

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.*, *31*(17), 6339–6352.

Watson, A. B., & Pelli, D. G. (1983). QUEST: A bayesian adaptive psychometric method. *Percept. Psychophys.*, *33*(2), 113–120.