Investigating implicit and explicit expectations in perceptual decision making

Ari Khoudary (ari.khoudary@uci.edu)

Aaron Bornstein* (aaron.bornstein@uci.edu)

Megan Peters* (megan.peters@uci.edu)
Department of Cognitive Sciences
University of California, Irvine

Abstract

Expectations, or the prior probability of a choice outcome, are powerful sources of evidence for improving decision making under uncertainty. Most expectations in the real world are learned *implicitly* on the basis of statistical properties of observers' environments. However, most studies investigating effects of expectations on perceptual decisions *explicitly* instruct observers on prior probabilities within the experiment, and thus fail to capture the experience-dependent uncertainty of real-world expectation learning. Here, we report data from a novel expectation-guided perceptual decision making task specifically designed to address this gap. Human observers (n=21) learned, through experience, probabilistic relationships between cues and images. Then, they explicitly reported both an estimate of each cue's prediction and a confidence rating in that estimate before performing a cued perceptual decision task. We find that, although these measurements are highly correlated, confidence in an explicit report is the primary factor that interacts with implicit expectations to shape perceptual decisions.

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

Introduction

Expectations in many perceptual decision making studies are operationalized as the prior probability that one of two possible 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 animals on thousands of trials in order to ensure they have learned that probability (e.g., Hanks et al., 2011). While this approach to ensuring a stable estimate of prior probability has many practical advantages, it may obscure the role of experience-dependent uncertainty of the sort inherent to perceptual expectations acquired outside of the lab. These sorts of more naturalistic expectations may influence decisions by a dynamic estimation process, such as aggregating across related experiences stored in memory to infer the prior probability of a particular choice outcome in a context-specific manner (Bornstein et al., 2023).

Indeed, previous studies using both the standard and more naturalistic approaches have identified neural and behavioral evidence of expectations that affect perceptual decisions dynamically over the course of a single choice (Hanks et al., 2011, Bornstein et al., 2023). Although each study explained their results using different evidence accumulation models, we recently showed that both effects can be explained by a single model that generates decisions by performing dynamic precision-weighted integration of parallel streams of memory and sensory evidence (Khoudary et al., 2022). This *dynamic integration theory* posits that when observers have uncertainty about (i) the difficulty of an upcoming perceptual decision, (ii) which prior probability to use for that decision, and (iii) what the true value of that prior probability is, they perform an automatic reliability estimation process to determine how much to weight information from each evidence source (expectation and sensation).

To complement our existing simulation-based support for the theory, we designed a novel expectation-guided perceptual decision task to test its key prediction: that the effect of expectations derived from memory increases at points in time when visual evidence is highly uncertain. This paper presents a series of regression analyses investigating how our novel measurements of uncertainty about learned expectations (1) relate to the true value of those implicit expectations and (2) impact subsequent expectation-guided perceptual decisions.

Methods

In our task, observers first learn implicit expectations by observing an interleaved series of probabilistic cue-image pairings. Immediately after learning, they explicitly report an estimate of each cue's predictive probability followed by a confidence rating in that estimate. After making subjective reports for all learned cues, observers perform a cued perceptual decision task that stochastically manipulates the reliability of sensory evidence in a trial-by-trial manner.

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.

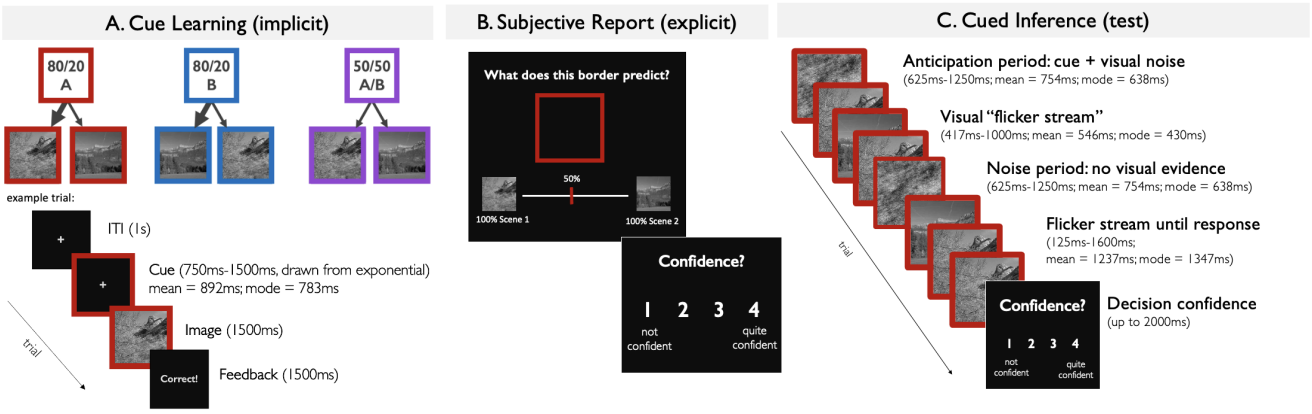

Figure 1: **Task design. (A)** Participants first learn that different colored borders make different predictions about the probability of observing one of two possible scene images. The objective probability of each cue is defined by the frequency with which it is followed by one of the two scene images. **(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.

Participants A sample of 22 undergraduate students (age M=20.5 years; SD=1.86 years; 16 female, 4 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=21 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.

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 (Peters and Maniscalco, 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: each 80% cue thus preceded its dominant image 24/30 times it was presented and the 50% cue preceded each image 15 times. Colors were randomly

assigned to cue probabilities.

Subjective Report of Learned Probabilities To obtain explicit reports of learned cue probabilities, participants were presented with each colored border and used a sliding scale to report their best estimate of a cue's predictive probability based on the associations observed in the immediately preceding Cue Learning phase (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 (Peters and Maniscalco, 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 coefficient values were performed using the `lmerTest` package, marginal means and pairwise contrasts of fitted models were obtained using the `emmeans` package, correlation coefficients and p-values were obtained using the `Hmisc` package, and performance metrics for fitted models (BIC & R^2) were obtained using the `compare_performance()` function from the `performance` package. Correlation coefficients were z-transformed prior to being used as predictors in regression analyses, and all regression models included a random intercept for each subject.

Results

Our analyses aimed to answer the following questions:

1. **How well do explicit reports correspond to implicitly learned probabilities?** We answer this by computing two metrics of the accuracy of explicit reports relative to the true value of implicitly learned probabilities, and then examining which factors drive confidence in explicit probability reports.
2. **How do implicit and explicit expectations impact perceptual decisions under uncertainty?** We answer this by examining effects of each expectation type, along with metrics of explicit expectation accuracy, on behavior in the Cued Inference phase.

1. Learned Probability Reporting: Accuracy and Confidence of Reported Cue Probabilities

Accuracy of cue estimates Figure 2 depicts two complementary metrics of the accuracy of explicit reports of implicitly learned cue probabilities. The $cueDiff$ metric quantifies, for each cue, the difference between the true probability as defined by the structure of the Cue Learning phase ($trueCue$) and an observer's explicit report of what they learned that probability to be ($subjectiveCue$). As shown in Figure 2A, participants tended to overestimate the probability of the 50% cue relative to its true value and underestimate the probability of the 80% cues relative to their true values. Two independent t-tests against 0 confirmed the significance of these deviations (50% cue: $t_{41} = 6.33$, $p < .001$; 80% cue: $t_{83} = 8.92$, $p < .001$). The $cueCorr$ metric quantifies, for each observer, the linear correlation between $trueCue$ and $subjectiveCue$. 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 We turned next to examining the factors driving observers' confidence ratings in their $subjectiveCue$ reports (i.e., $cueConfidence$). To do this, we conducted formal model comparison on 4 nested linear models. As shown by the BIC values in Table 1, a model estimating $cueConfidence$ using only $subjectiveCue$ (the observer's report of a cue's probability) best trades off complexity with goodness-of-fit. The marginal increase in R^2 values as a

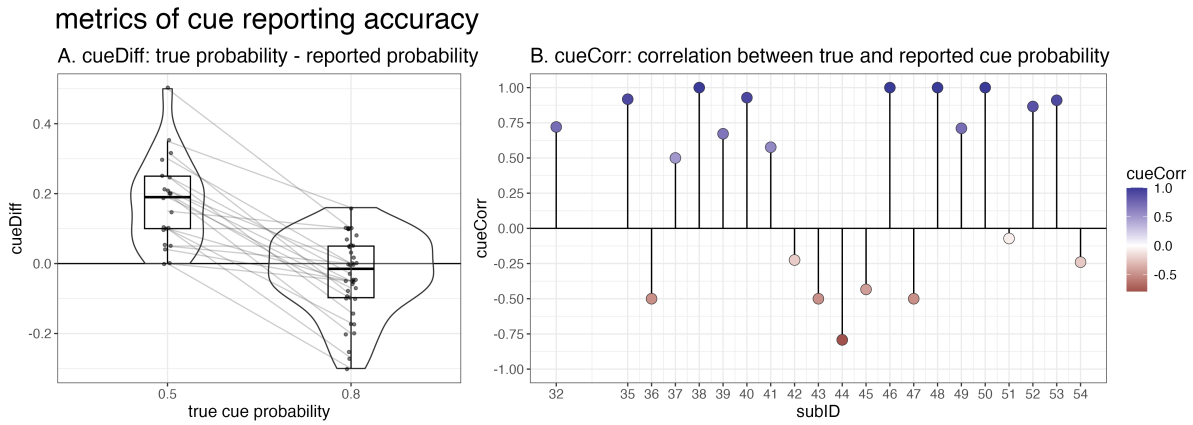

Figure 2: **Metrics quantifying the accuracy of reported cue probabilities. (A)** Participants overestimated the true probability of the 50% cue and underestimated the true probability of the 80% cue. **(B)** The linear correlation between a cue's true predictive probability as defined in the Cue Learning phase (*trueCue*) and subjects' report of that probability (*subjectiveCue*) varies across individual subjects. Cool colors represent positive *cueCorr* values, warm colors represent negative *cueCorr* values.

function of model complexity further indicate the strength of *subjectiveCue* as a predictor for variability in *cueConfidence* judgments. The winning model reveals a strong positive relationship between *subjectiveCue* and *cueConfidence* (β = 0.602, $t_{58} = 6.654$, $p < .001$), indicating that confidence in a probability report scaled positively with the magnitude of the reported probability itself.

Effects Structure	BIC	R^2
trueCue	139.3	0.162
subjectiveCue	**115.9**	0.433
trueCue + *subjectiveCue*	117.2	0.458
trueCue ∗ *subjectiveCue*	120.5	**0.465**

Table 1: **Comparing predictors of confidence in reported cue probabilities**. Best values for each metric are bolded. The winning linear model uses only the reported probability (*subjectiveCue*) to predict *cueConfidence* and returns a strong positive relationship between *cueConfidence* and *subjectiveCue* magnitude.

2. Cued Inference: Choice Behavior and Timing

We now turn to investigating how implicitly learned cue probabilities (*trueCue*), explicit subjective reports of those probabilities (*subjectiveCue*), and the overall accuracy of those reports (*cueDiff* and *cueCorr*) interact to shape behavior during perceptual decision making. 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.

A. Responses driven by expectations alone We began our analysis by computing linear correlations among five variables of interest: (1) *trueCue*: each cue's true predictive

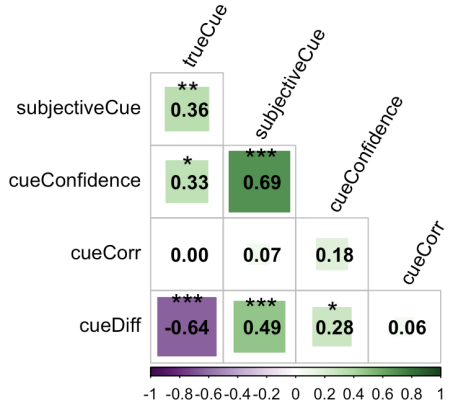

Figure 3: **Linear correlations among possible predictors of early responding.** Stars indicate significance of a t-test against 0; * $p < .05$, ** $p < .01$, *** $p < .001$

probability as learned implicitly during Cue Learning; (2) *subjectiveCue*: observers' explicit report of each cue's probability; (3) *cueConfidence*: observers' confidence in their *subjectiveCue* report; (4) *cueDiff*: the difference between *subjectiveCue* and *trueCue*; and (5) *cueCorr*: the overall correlation between an observer's *subjectiveCue* reports and their corresponding *trueCue* values. Figure 3 shows that *cueConfidence* and *cueDiff* both exhibit significant correlations with *trueCue* and *subjectiveCue*, whereas *cueCorr* does not appear strongly correlated with any other variables. We used this correlation analysis to narrow down the space of possible fixed effects structures and avoid issues of multicollinearity when fitting regression models.

The first round of regression analyses examined which fixed effects structure best explains the probability that observers made an early response, which occurred on 373/7682 (4.85%) Cued Inference trials. Table shows that a linear combination of *cueCorr* (z-transformed) and *subjectiveCue* explained the most variance in the data (marginal R^2 =

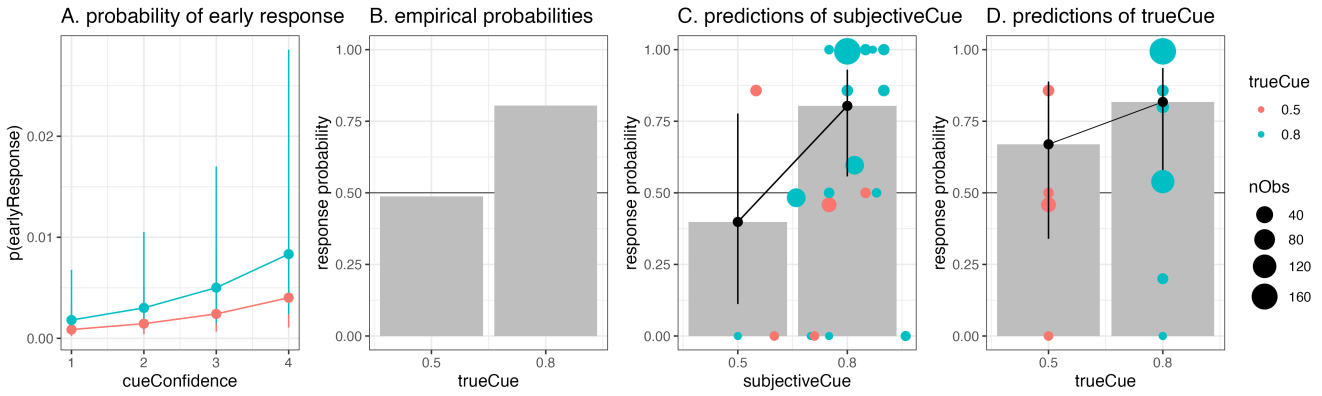

Figure 4: **Choice behavior on the basis of expectation alone.** Errorbars correspond to standard error of model estimates. In C & D, the size of individual points corresponds to how many observations contributed to that point. **(A)** The probability of making a choice without perceiving any visual evidence increases independently as a function of *cueConfidence* and *trueCue*. **(B)** Empirical biases for choices based only on expectations accord with predictions of optimal decision theory. **(C)** Choice bias predictions based on explicit expectations (preferred by BIC). **(D)** Choice bias predictions based on implicit expectations.

0.191), whereas a linear combination of *cueConfidence* and *trueCue* best balanced complexity with goodness-of-fit (BIC=1600.9). Further, the second-best model by BIC used only *cueConfidence* to predict early responding (BIC=1605.1), underscoring the utility of this factor for predicting early responding. The winning model by BIC returned a significant main effect both for *cueConfidence* ($\beta = 0.462$, $z_{7354} = 4.25, p < .001$) and *trueCue* ($\beta = 2.4436$, $z_{7354} = 3.504, p < .001$), such that observers were more likely to make early responses when *trueCue* was 80% *or* when they had higher confidence in their *subjectiveCue* estimate for that 80% cue (Figure 4A).

The next round of regression analyses investigated how the same set of fixed effects structures fared in predicting RT trends for early responses. Reaction times were log-transformed and z-scored prior to model fitting. As shown in Table 2, several fixed effects structures exhibited highly similar performance both in terms of BIC and R^2. The best model by BIC consisted of a fixed effect only of *cueConfidence* (BIC=220.9), whereas the model with the greatest R^2 value used an additive combination of *cueCorr*, *trueCue*, and *subjectiveCue* ($R^2 = 0.006$). The winning model by BIC returned a null effect of *cueConfidence* on RTs for early responses ($\beta = 0.024, t_{204} = 0.782, p = 0.43$), and the winning model by R^2 returned null effects for all three predictors ($\beta_{trueCue} = -0.145, t_{358} = -0.786, p = 0.432$; $\beta_{subjectiveCue} = 0.219, t_{366} = 0.809, p = 0.419$; $\beta_{cueCorr} = -0.0105, t_{9.55} = -0.262, p = 0.799$), suggesting that factors other than those investigated in these analyses (e.g., decision threshold) were the primary drivers of RT variability for responses made before the onset of visual evidence.

The final set of regressions investigated whether observers exhibited systematic biases in the choices made before the onset of visual evidence. Optimal decision theory states that, as the prior probability of one option approaches 1 and the inter-stimulus interval is sufficiently short, observers should forego any evidence accumulation and always respond according to the more likely option (Simen et al., 2009). In the case of our experiment, early responses should be biased toward the dominant prediction of 80% cues and not exhibit any bias toward a particular option for 50% cues. Figure 4B shows that aggregated response probabilities qualitatively accord with this prediction. Interestingly, the BIC metric favored *subjectiveCue* as a better predictor of early choice biases than models that used *trueCue* on its own or in addition to *subjectiveCue* (Table 2; $\Delta BIC = 5.1$). The performance advantage of *subjectiveCue* as a predictor is displayed in Figures 4C and D. Both models confirmed the statistical significance of the choice biases displayed in early responding behavior ($\beta_{subjectiveCue} = 6.071, z_{370} = 3.02, p = 0.003$; $\beta_{trueCue} = 2.644, z_{370} = 0.051$).

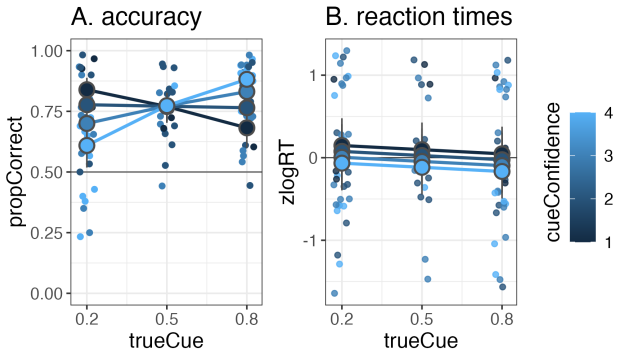

Figure 5: **Choice behavior on the basis of integrated expectations and sensory evidence.** Errorbars correspond to standard error of model estimates **(A)** Accuracy was best explained by an interaction between *trueCue* and *cueConfidence*. **(B)** Reaction times were best explained by an additive combination of *trueCue* and *cueConfidence*.

B. Responses integrating memory and sensory evidence We next used formal model comparisons to investigate (i) factors driving choices made on the basis of integrated expectation and sensory evidence as well as (ii) the effect of

Expectations Only		
p(earlyResponse)		
Effect Structure	BIC	Marginal R^2
trueCue	1649.8	0.021
subjectiveCue	1643.0	0.043
cueCorr	1682.5	0.149
cueConfidence	1605.1	0.040
cueDiff	1686.0	0.002
cueCorr + *trueCue*	1653.6	0.167
cueCorr + *subjectiveCue*	1647.4	**0.191**
trueCue + *subjectiveCue*	1643.1	0.038
cueConfidence + *trueCue*	**1600.9**	0.040
cueCorr + *trueCue* + *subjectiveCue*	1647.5	0.182
reaction time		
Effect Structure	BIC	Marginal R^2
trueCue	222.4	1.54e-04
subjectiveCue	222.4	2.24e-04
cueCorr	222.4	2.24e-04
cueConfidence	**220.9**	4.16e-04
cueDiff	221.8	4.97e-04
cueCorr + *trueCue*	228.4	0.005
cueCorr + *subjectiveCue*	228.4	0.005
trueCue + *subjectiveCue*	231.7	5.82e-04
cueConfidence + *trueCue*	226.0	7.92e-04
cueCorr + *trueCue* + *subjectiveCue*	233.7	**0.006**
choice bias		
Effect Structure	BIC	Marginal R^2
trueCue	305.9	0.009
subjectiveCue	**300.8**	0.031
trueCue + *subjectiveCue*	304.08	**0.041**
Expectations + Sensory Evidence		
accuracy		
Effect Structure	BIC	Marginal R^2
trueCue	7398.4	0.015
subjectiveCue	7458.7	0.004
cueCorr	7470.8	1.35e-04
cueConfidence	7067.5	0.007
cueDiff	7469.9	2.73e-04
cueCorr + *trueCue*	7407.3	0.015
cueCorr + *subjectiveCue*	7467.5	0.016
trueCue + *subjectiveCue*	7401.5	0.017
cueConfidence + *trueCue*	7010.4	0.021
cueConfidence * *trueCue*	**6930.5**	**0.042**
cueCorr + *trueCue* + *subjectiveCue*	7410.3	0.017
reaction time		
Effect Structure	BIC	Marginal R^2
trueCue	15255.4	0.002
subjectiveCue	15238.9	0.005
cueCorr	15282.5	1.58e-04
cueConfidence	14704.4	0.005
cueDiff	15282.4	2.24e-06
cueCorr + *trueCue*	15264.4	0.002
cueCorr + *subjectiveCue*	15247.9	0.005
trueCue + *subjectiveCue*	15228.6	0.005
cueConfidence + *trueCue*	**14692.0**	**0.006**
cueConfidence * *trueCue*	14694.0	**0.006**
cueCorr + *trueCue* + *subjectiveCue*	15237.6	0.005

Table 2: **Regression model comparisons**. Bolded values indicate the winning value for each metric. All models included a random intercept for subjects.

the stochastic noise period on this behavior. This introduces *cue validity*–whether the cue's prediction was correct relative to the visual evidence stream–as another dimension that can modulate behavior. We model cue validity by adding a level of 0.2 to the *trueCue* variable, which captures the trials for which the 80% cue made an invalid prediction with respect to the true answer as given by visual evidence.

Table 2 (bottom) shows that accuracy for choices made after the anticipation period were best captured by an interaction between *cueConfidence* and *trueCue*, whereas reaction times were best captured by an additive effect of *cueConfidence* and *trueCue*. Figure 5A shows how confidence in an explicit cue estimate significantly interacted with *trueCue* to modulate choice accuracy ($\beta_{trueCue*cueConfidence} = 1.3635, z_{6984} = 9.245, p < .001$): high *cueConfidence* enhanced accuracy on valid cue trials ($trueCue = 0.8$) but impaired accuracy on invalid cue trials ($trueCue = 0.2$), whereas the opposite pattern obtained on trials for which observers had low confidence in their explicit cue estimate. Figure 5B shows how *cueConfidence* and *trueCue* are combined in guiding reaction times. RTs were faster when observers had high confidence in their explicit probability estimate ($\beta_{cueConfidence} = -0.071, t_{6979} = -6.50, p < .001$), and also became faster as a function of that cue's predictiveness about the correct answer ($\beta_{trueCue} = -0.168, t_{6966} = -4.60, p < .001$).

These findings, together with those from early responding behavior, demonstrate a key role for both implicit and explicit cue probabilities in perceptual decisions. Specifically, our results suggest that confidence in an explicit probability estimate–a quantity that is highly correlated with, but conceptually distinct from, the magnitude of the explicit estimate–is even more useful than the explicit expectation itself for capturing quantitative trends in the data.

Discussion

We measured human behavior on a task that required observers to learn expectations from experience, make explicit reports and confidence judgments about their estimates of each cue's probability, and then use those cues to make perceptual decisions under uncertainty. We used a series of regression models to investigate which factors best accounted for quantitative trends in the behavioral data, and found that most behavior is best described by a combination of implicitly learned and explicitly reported probability information. These findings indicate a key role for the experience-dependent properties of implicit expectations that previous tasks were not sensitive enough to measure.

We also found a substantive amount of variability in the accuracy of explicit expectation reports, and did not have the statistical power to detect effects of expectations on reaction times. Future work will use develop sequential sampling models to investigate the factors driving variability in each of these outcome variables in order to deepen understanding of implicit and explicit probability information are dynamically combined in the service of adaptive behavior.

Acknowledgements

The authors thank Nadeja Jackson, Nurazmera Mossa, and Allison Krueger for their assistance in collecting the human subjects data.

References

Bornstein, A. M., Aly, M., Feng, S. F., Turk-Browne, N. B., Norman, K. A., & Cohen, J. D. (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.

Khoudary, A., Peters, M. A. K., & Bornstein, A. M. (2022). Precision-weighted evidence integration predicts time-varying influence of memory on perceptual decisions. *Cognitive Computational Neuroscience*.

Peters, M. A. K., & Maniscalco, B. (2024). Achieving constant hazard rate in psychophysics experiments. *PsyArXiv*.

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *J. Exp. Psychol. Hum. Percept. Perform.*, *35*(6), 1865–1897.

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