

A Switching Observer for Human Perceptual Estimation

Highlights

- Human motion direction estimate summary statistics match Basic Bayesian predictions
- However, bimodality of estimate distributions violated Basic Bayesian predictions
- An observer that switches between prior and likelihood better explains estimates
- This suggests a switching heuristic to achieve nearly optimal estimation behavior

Authors

Steeve Laquittaine, Justin L. Gardner

Correspondence

jlg@stanford.edu

In Brief

Statistically optimal perceptual inference behavior is obtained by multiplicatively combining priors and likelihoods. Here Laquittaine and Gardner find evidence for a human inference strategy that approximates Bayesian optimality by a switching heuristic that forgoes combining priors and evidence.

A Switching Observer for Human Perceptual Estimation

Steeve Laquitaine^{1,2} and Justin L. Gardner^{1,2,3,*}

¹Department of Psychology, Stanford University, Stanford, CA 94305, USA

²Laboratory for Human Systems Neuroscience, RIKEN Brain Science Institute, Wako-shi, Saitama 351-0198, Japan

³Lead Contact

*Correspondence: jlg@stanford.edu

<https://doi.org/10.1016/j.neuron.2017.12.011>

SUMMARY

Human perceptual inference has been fruitfully characterized as a normative Bayesian process in which sensory evidence and priors are multiplicatively combined to form posteriors from which sensory estimates can be optimally read out. We tested whether this basic Bayesian framework could explain human subjects' behavior in two estimation tasks in which we varied the strength of sensory evidence (motion coherence or contrast) and priors (set of directions or orientations). We found that despite excellent agreement of estimates mean and variability with a Basic Bayesian observer model, the estimate distributions were bimodal with unpredicted modes near the prior and the likelihood. We developed a model that switched between prior and sensory evidence rather than integrating the two, which better explained the data than the Basic and several other Bayesian observers. Our data suggest that humans can approximate Bayesian optimality with a switching heuristic that forgoes multiplicative combination of priors and likelihoods.

INTRODUCTION

There is apparent tension between theories that propose that human decision making can be modeled as optimal Bayesian inference and theories that suggest that human decision making is fundamentally heuristic and subject to bias. Many aspects of human cognition, such as language acquisition and processing (Chater and Manning, 2006; Tenenbaum et al., 2006), action selection (Chater and Manning, 2006; Körding and Wolpert, 2004; Tenenbaum et al., 2006; Wolpert et al., 1995), prediction (Griffiths and Tenenbaum, 2006), reasoning (Tenenbaum et al., 2011), and sensory inference (Fetsch et al., 2013; Girshick et al., 2011; Knill and Richards, 1996; Stocker and Simoncelli, 2006; Weiss et al., 2002), have been modeled as optimal read-outs of statistical inference processes, for example, starting from the notion that visual input offers inherently ambiguous information about the state of the world (von Helmholtz, 1924). Bayesian models propose optimal solutions to incorporating

prior knowledge to resolve this ambiguity (Knill and Richards, 1996). Indeed, human visual perceptual estimates have been shown to follow Bayesian predictions for estimates of motion speed (Jogan and Stocker, 2015; Stocker and Simoncelli, 2006; Vintch and Gardner, 2014; Weiss et al., 2002) and direction (Chalk et al., 2010), spatial location (Berniker et al., 2010; Tassinari et al., 2006; Vilares et al., 2012), ambiguous 3D structure (Adams et al., 2004; Freeman, 1994), depth (Ban et al., 2012; Kim et al., 2015; Ramachandran, 1988), and edge orientation (Girshick et al., 2011).

While this Bayesian framework is elegant and appealing, it is well known that human decisions can deviate from normative standards (Cheadle et al., 2014; Nassar et al., 2016; Neiman and Loewenstein, 2011; Rahnev and Denison, 2016; Sharot et al., 2011; Tversky and Kahneman, 1974) by showing inappropriate adherence to strategies that are non-optimal in the current context (Abrahamyan et al., 2016; Acerbi et al., 2014; Beck et al., 2012; Fischer and Whitney, 2014; Raviv et al., 2014) or use simpler, computationally frugal strategies to incorporate priors and improve cognition (Nassar et al., 2010; Rahnev and Denison, 2016; Raviv et al., 2012; Vul et al., 2014; Wilson et al., 2013).

However, optimal and heuristic frameworks need not be at odds when considering the difference between computational goals of a system and constraints on implementation to achieve these goals (Marr, 1982; Nassar et al., 2010; Rahnev and Denison, 2016; Vul et al., 2014; Wilson et al., 2013). In this view, optimality theory precisely describes the target for what a system should achieve, but attaining this goal might be subject to constraints of implementation that require suboptimal or heuristic solution. Take, for instance, the central computation in Bayesian inference in which priors and likelihoods are multiplied to form a posterior distribution. This is the optimal solution to incorporating these two pieces of information, and while proposals exist for how these might be computed by neural populations (Beck et al., 2008; Girshick et al., 2011; Jazayeri and Movshon, 2006; Ma et al., 2006; Pouget et al., 2013; Wei and Stocker, 2015), less complex computations may be easier to instantiate in neural circuits. An attractive alternative is that simple shortcuts or heuristics can achieve behavior that meets or approximates the normative goals of Bayesian calculations without complex computations (Moreno-Bote et al., 2011; Raviv et al., 2012; Wilson et al., 2013). Simpler heuristic approaches that approximate Bayesian inference without integrating prior and likelihood would reduce the computational complexity required to achieve nearly optimal behavior.

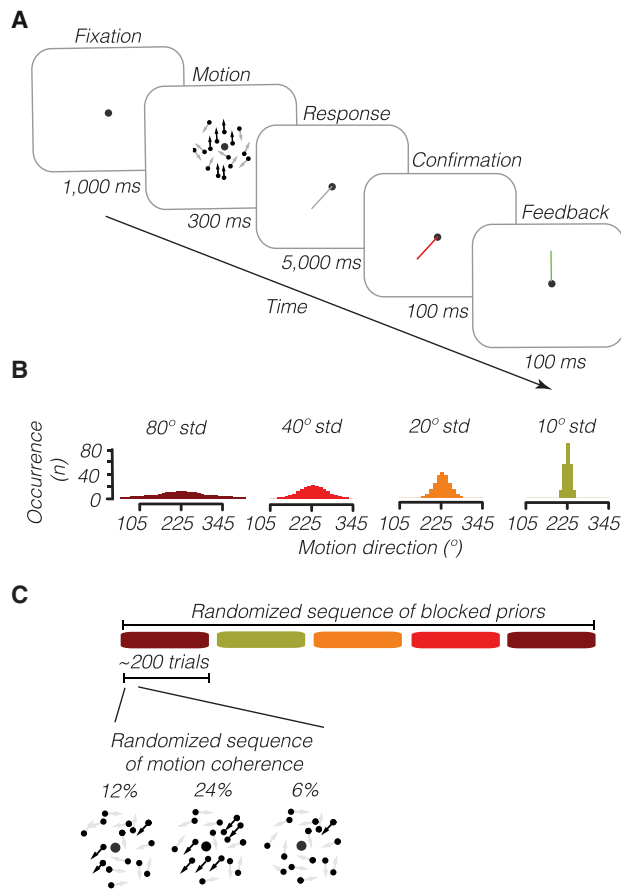


Figure 1. Motion Direction Estimation Task in which Sensory Evidence and Priors Were Independently Manipulated

In each trial (A), subjects reported the direction of motion of a random dot stimulus. The strength of sensory evidence was manipulated by changing the percentage of dots (6%, 12%, or 24% in different randomly interleaved trials) moving coherently (dark arrows, not actually shown in stimulus) with the remaining dots moving in random directions (gray arrows). The strength of the priors was manipulated by controlling the width of the distribution of directions in a block of trials (B; brown, red, orange, and green distributions are priors with standard deviations of 80°, 40°, 20°, and 10°, respectively) while keeping the mean, 225°, the same. Priors were randomized across 5 blocks of about 200 trials (C), while sensory evidence was randomized within each block.

We examined human behavior in a motion direction estimation task and found that while subjects approximated an optimal Bayesian observer in summary statistics, their trial-by-trial estimates suggested a heuristic implementation that did not require integration of prior and likelihood. Our estimation task was designed to test conformity with a commonly used formulation of a Bayesian observer (Girshick et al., 2011; Stocker and Simoncelli, 2006; Wei and Stocker, 2015), which we call the Basic Bayesian observer. We found evidence that behavior whose summary statistics matched predictions of the Basic Bayesian observer was actually better modeled as a process that switches between prior and sensory likelihood rather than integrates the two as would be expected by the optimal Basic Bayesian computation. That is, the Basic Bayesian observer could fit the mean and variance of direction estimates, but not the shape of

estimate distributions, which were bimodal with peaks near the prior mean and the presented direction. We captured this bimodality in a Switching observer model, which qualitatively and quantitatively provided better fits to the behavioral data than the Basic Bayesian observer, including variants that relaxed various assumptions about the shape of priors and likelihoods as well as the method of reading out the posterior. Thus, by using the Basic Bayesian framework as a normative benchmark, we found that humans in our task show behavior that approaches the goals of optimality in summary statistics but was best characterized as being implemented by a simpler computational heuristic of switching, rather than multiplicatively combining, priors and likelihoods.

RESULTS

Basic Bayesian Observer Accounts for Subject Estimate Mean and Variability

We independently manipulated the strength of sensory evidence and priors in a motion direction estimation task (Figure 1) to test whether subjects' behavior conformed to a Basic Bayesian observer model of sensory inference (Girshick et al., 2011). On each trial, a brief (300 ms) presentation of random dot motion was presented while subjects fixated on a central fixation point. Subjects were instructed to report the direction of motion of the dots by turning an electronic paddle wheel to move a white line from a random initial orientation to point in the direction of motion that they had perceived and then confirmed that report with a button press. Subjects were not given explicit feedback as to whether they were correct or not; instead, a green feedback line that pointed in the true direction of motion would be shown afterward. We dissociated the effects of sensory evidence and priors by independently manipulating the strength of motion by changing the percentage of dots moving in the same direction (motion coherence) and the strength (sharpness) of the distribution of motion directions over trials within blocks of 202–226 trials. We reasoned that we should be able to test whether our subjects' behavior conformed to the predictions of a Basic Bayesian observer by examining the mean and variability of their estimates as a function of the strength of sensory evidence and priors.

We compared subjects' behavior to that of a basic model of sensory Bayesian inference, similar to other models that have been used to explain orientation and speed biases (Figure 2A; Girshick et al., 2011; Stocker and Simoncelli, 2006; see STAR Methods for details). Briefly, the Basic Bayesian observer formalizes prior and sensory likelihood as von Mises distributions, which are multiplied together to give a posterior distribution from which the estimated direction is chosen as the one with the maximum a posteriori probability and selected with some motor noise and probability of lapsing. On each trial, a particular motion direction, θ_{true} , is displayed and the Basic Bayesian observer simulates a sensory likelihood distribution (black distribution, Figure 2A) as a von Mises. This distribution is centered on θ_e , a draw from a sensory evidence distribution centered on θ_{true} . Thus, on average, the direction with the highest sensory likelihood is θ_{true} but, on any given trial, can differ from the true direction. Prior distributions (brown distribution, Figure 2A) were von Mises, same as the true prior distribution,

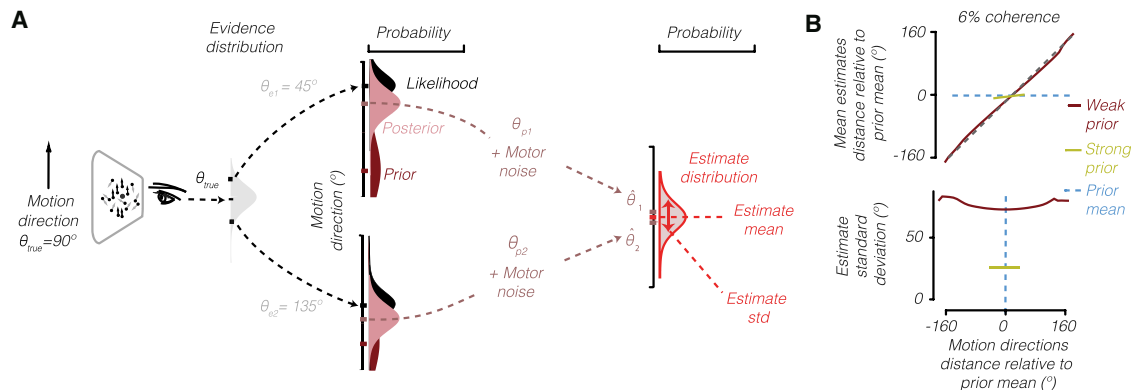


Figure 2. Basic Bayesian Observer Model

Example Bayesian integration for two trial instances of the same motion direction (A). To estimate the motion direction, θ_{true} , of a stimulus (e.g., 90°), the Basic observer obtains a visual measurement, θ_e , from an evidence distribution that simulates trial-to-trial variability of sensory evidence (e.g., $\theta_e = 45^\circ$ in trial 1, top row, or $\theta_e = 135^\circ$ in trial 2, bottom row). Sensory likelihood is modeled as a distribution around θ_e (black distributions) and then multiplied against the prior distribution (brown distributions) into a posterior (pink distributions). Percepts are obtained by reading out the posterior mode and are then transformed by a noisy motor process into an estimate $\hat{\theta}$. Over many trials, the Basic Bayesian observer generates an estimate distribution (red outlined distribution). The Basic Bayesian observer predicts that when the prior is strong (B; green lines) the mean of the estimate distributions (top) will be more biased toward the prior mean (blue dashed line) and will be less variable (i.e., the estimate distribution will have lower standard deviation, bottom) than when the prior is weak (brown lines).

whose maximum was the actual prior direction but width was a fit parameter. After multiplying prior and sensory likelihood, the resulting posterior (pink distribution, Figure 2A) mode was taken to be the predicted percept θ_p . To simulate lapses, on a small proportion of trials p_r , we set the observer's percept to a random direction. The distribution of percepts formed over trial instances of the motion was convolved with motor noise formalized as a von Mises distribution of concentration k_m . The width of the sensory likelihood and estimate distribution for each motion coherence (three fit parameters k_e), width for each prior distribution (four fit parameters k_p), and the motor noise (k_m), as well as the lapse rate (p_r), were fit to the behavioral data on a trial-by-trial basis using maximum likelihood estimation.

The Basic Bayesian observer predicts that subject mean direction estimates will be biased toward the mean of the prior and that this effect is influenced by both the strength of sensory evidence and the strength of the prior (Figure 2B). These predictions were borne out in the behavioral data. When sensory evidence was strong (motion coherence of 24%), subject mean estimates followed the actual direction presented (diagonal one-to-one line, Figure 3A, left column). When sensory evidence was weak (motion coherence 12% and 6%, center and right column), subject mean estimates systematically deviated from the veridical (diagonal one-to-one line) toward the prior mean (blue horizontal line at 0°). This bias was stronger for stronger priors in which the distributions of motion directions presented had smaller standard deviations (from 80° to 10° , brown to green). The bias was also the largest for directions displayed nearer to the prior, producing a nonlinear relationship between displayed and estimated directions (curved fit lines, Figure 3A, right column). This effect is due to the circularity of the direction space: when the displayed direction is opposite from the mean of the prior, the posterior estimate is pulled in both directions away from the likelihood, resulting in no mean effect of the prior. These effects on the mean estimates conformed to the expecta-

tions of the Basic Bayesian observer whose fits to the data are shown as solid lines (Figure 3A).

The Basic Bayesian observer also predicts that subject estimate variability should be reduced as both the sensory evidence and prior become stronger. These predictions were also borne out in the behavioral data. When sensory evidence was strong, subject estimate variability was low (Figure 3B, left column) and relatively unaffected by prior strength. When sensory evidence was weak (center and right column), variability was overall larger but reduced as priors became stronger (brown to green), which the observer model predicts because stronger priors will increase confidence in estimates and result in larger biases but lower variability, effectively minimizing expected deviation from true motion directions over trials. The variability was also lower for directions displayed nearer to the prior mean. This effect, similar to effects on mean estimates, is a result of the circular space of direction. Thus, the Basic Bayesian observer can account for the effect of the prior and its strength on both mean and variance of subjects' estimates. Moreover, model fits verify that subjects' estimates incorporated prior knowledge about the motion direction distributions as the full Basic Bayesian observer model provided a better fit to the behavioral data than one that exclusively relied on sensory likelihood (uniform priors, the average Akaike information criterion [AIC] difference [$AIC_{Maximum\ likelihood} - AIC_{Basic\ Bayesian}$], or ΔI , between the maximum likelihood observer and the Basic Bayesian observer was 13,950, 95% CI [10,247 17,653] over subjects [$n = 12$ subjects], in favor of the Basic Bayesian observer for all 12 subjects, where a difference of 2 is equivalent to adding another parameter to the model).

Basic Bayesian Observer Fails to Account for Bimodality of Estimate Distributions

While the Basic Bayesian model could account for the summary statistics of subjects' estimate distributions, it did a poor job of

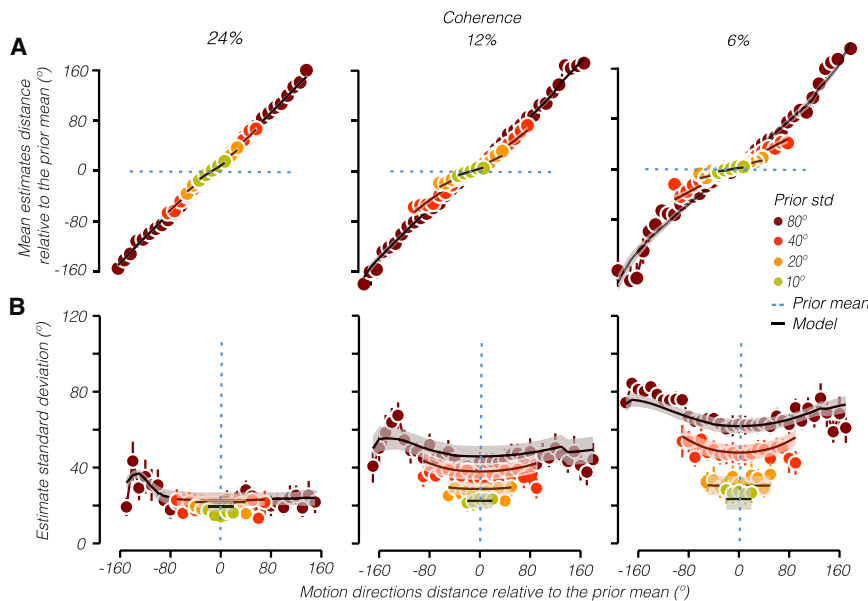


Figure 3. Basic Bayesian Observer Model Fit to Mean and Variability of Subjects' Estimates

Subject estimates mean (A) and standard deviation (B) averaged over 12 subjects with the Basic Bayesian observer's best fit (lines, fit on a subject-by-subject basis and averaged over subjects) for four prior conditions (80° to 10° SD, brown to green colors) and three motion coherences (left, middle and right panels). Errors for data and model predictions are SEM over subjects. Note that SEM are sometimes too small to be visible. Blue dashed lines are prior means.

governed by the widths of sensory likelihood and prior, no more parameters are needed for the Switching observer compared to the Basic Bayesian observer. The switching estimation process is similar to a process where two noisy representations of sensory likelihood and prior compete in a winner-take-all manner.

predicting the full distribution shapes, which had a bimodal rather than a unimodal shape (pink histogram with red contour, Figure 2A). Contrary to the Basic Bayesian observer prediction that estimate distributions should have one peak (black curves, subject-averaged data shown for 6% coherence, Figure 4; see Figures S1 and S2 for subject-by-subject distributions) lying somewhere between the mean of the prior and sensory likelihood, subject estimate distributions were characterized by two peaks: one at the motion direction (gray tick mark) and one at the prior mean (blue line). The two peaks were clearest when the stimulus direction was the furthest from the prior mean (e.g., 140° clockwise of the prior mean tagged with the number 1). The peak at the prior mean was also higher than the peak at the motion direction when the prior was stronger (top brown versus red histograms tagged with numbers 1 and 2). Thus, while the mean and standard deviation of estimates conformed to the expectations of the Basic Bayesian observer, the full shapes of the distributions being bimodal rather than unimodal were not predicted.

To account for the bimodal estimate distribution, we developed a Switching observer. The model represents sensory likelihood and priors (black and brown histograms, Figure 5A) just like the Basic Bayesian observer, with the same fit parameters (three κ_e that quantify sensory evidence strengths, four κ_p for prior strengths, one ρ_r for the probability of random estimation, and one κ_m for motor precision), but rather than forming a posterior by multiplying the two distributions, switches between estimates chosen from the prior distribution (brown histogram) and the sensory likelihood (black histogram), producing bimodal instead of unimodal estimate distributions (Figure 5A versus Figure 2A, pink histogram with red contours). Guided by the observation that the subject estimate distribution peaks at the prior mean were higher than the peak at the motion direction when the prior was stronger, we defined switching as a competitive mechanism controlled by the strength of prior and sensory evidence (Figure 5A; see STAR Methods, Equation 6). As this switching probability is

Fits of the Switching observer to the behavioral data qualitatively reproduced the mean, variability, and bimodality of the estimate distribution data and quantitatively fit better than the Basic Bayesian observer (see Figure S3A for switching probability estimates and S3B for all model fit parameters). When coherence was high, the Switching observer chooses estimates from the sensory likelihood more often, and therefore, mean estimates were veridical (one-to-one line, Figure 5B, left column) and variability was low (Figure 5C, left column, solid lines are model fits). When coherence was low, the Switching observer chooses more estimates from the prior distribution, resulting in more bias toward the mean of the prior (Figure 5B, right column) but less variability (from brown to green, Figure 5C, right column) as more estimates are chosen from only the prior rather than both prior and sensory likelihood. As designed, the Switching observer created two peaks in the estimate distributions as observed in the data (Figure 5D, solid curves). A direct model comparison of the Basic Bayesian and Switching observer fit to the subjects' data showed that switching outperformed the Bayesian model for 8 out of 12 subjects (Figure 5E, the average AIC difference, $AIC_{\text{Basic Bayesian}} - AIC_{\text{Switching}}$, was 395, 95% CI [96 695] over subjects in favor of switching).

While there is clearly variation in how well the Switching observer fits individual subject's data, examination of estimation distributions for two well and two poorly fit subjects (Figure 5F, top and bottom, respectively; see Figures S1 and S2 for subject-by-subject distributions) indicates that bimodality (distributions near the prior and sensory evidence, blue line and gray tick mark) is evident in subject-by-subject data. Moreover, poorly fit subjects tend to have more random estimates, not a stronger tendency to form a single distribution near the expected posterior of the Basic Bayesian observer. This increased randomness could be due to subjects not having a completely fixed and noiseless estimate of the prior mean, which is assumed by the Switching observer. Indeed, if we relax this assumption by allowing there to be trial-to-trial variation of the prior mean by

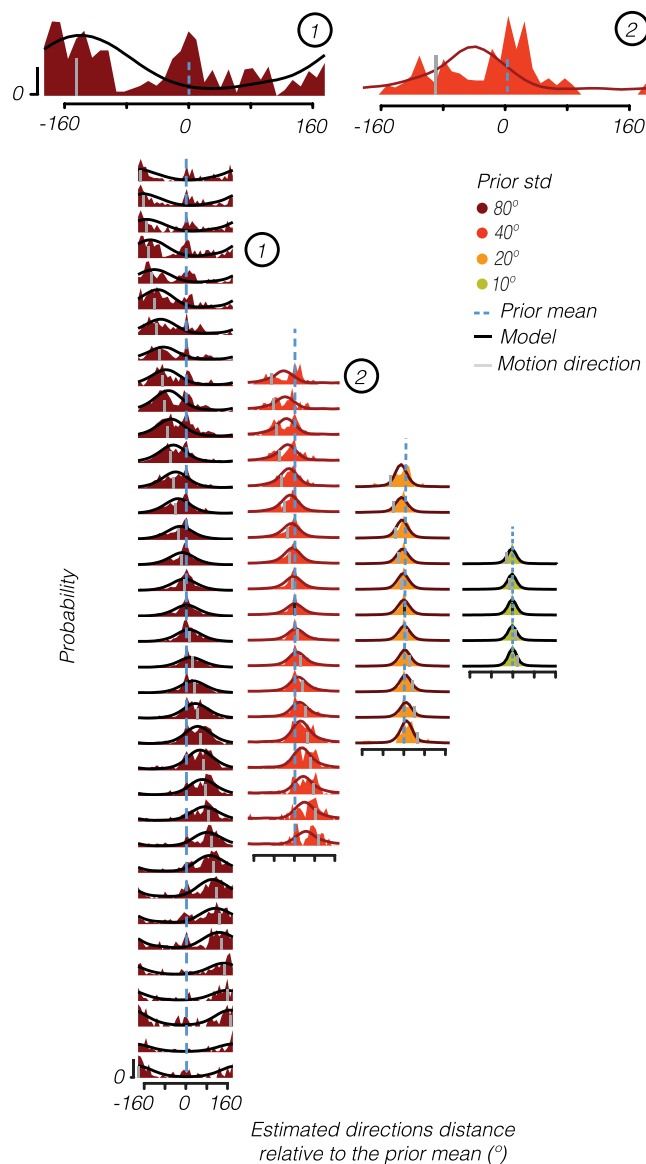


Figure 4. Basic Bayesian Observer Model Fails to Predict Bimodality of Subject Estimates

Subjects' estimate distributions (10° bins) show bimodality with one peak near the prior and another at the presented motion direction (different colors represent different prior strengths, and data are shown for 6% motion coherence). Estimate distributions were calculated for each subject, then averaged over subjects (see Figures S1 and S2 for subject-by-subject estimate distributions). Bimodality is most evident for conditions in which the prior and presented motions are very different (two example histograms are shown enlarged at the top; circled numbers indicate which distributions these are). The Basic Bayesian observer predicts a single peak in between the prior and the presented motion direction (curves, average fits over subject). Blue dashed line is the mean of the prior. The ordinates of all axes were scaled between 0 and the maximum of each histogram to maximize visibility.

sampling from the prior distribution, this Switching with prior sampling observer outperformed the Basic Bayesian for all 12 subjects with an average AIC difference of 478, 95% CI [262 693].

Not only did the Switching observer better fit the behavioral data, but it did so with fit parameters that showed that subjects had more accurate estimates of prior strength than the Basic Bayesian observer. We compared the median over subjects of the four subjective prior strength parameters fitted by the Switching observer to each subject's data to veridical strength and found that although the Switching observer's median fitted prior over subjects were significantly weaker than the 80° and 40° veridical priors (91° median SD, $W = 7$, $p = 0.009$ and 72° median SD, $W = 0$, $p = 0.005$, respectively, $n = 12$ subjects, two-tailed one-sample Wilcoxon test), the Switching observer's fitted prior standard deviations for the 20° and 10° priors did not significantly differ from the veridical prior standard deviations (31° median SD, $W = 15$, $p = 0.06$ and 18° median SD, $W = 28$, $p = 0.42$ for the 20° and 10° priors, $n = 12$ subjects, two-tailed one-sample Wilcoxon test). In contrast, the Basic Bayesian observer's median fitted priors over subjects were significantly weaker than the veridical priors for all priors except the 20° prior (96° median SD, $W = 0$, $p = 0.0005$; 73° median SD, $W = 1$, $p = 0.0009$; 38° median SD, $W = 16$, $p = 0.08$; 23° median SD, $W = 8$, $p = 0.012$; for the 80°, 40°, 20°, and 10° priors, respectively, $n = 12$ subjects, two-tailed one-sample Wilcoxon test), which would suggest that the observers were unable to make accurate estimates of the prior distribution.

Alternative Models

The bimodality in estimate distributions that the Switching observer captures could have been a result of learning through the trial block, but an analysis of early and late trials suggests otherwise. Rather than switching between prior and sensory likelihood as prescribed by the Switching observer, subjects could have relied primarily on sensory evidence in the early phase of learning before the prior condition is learned and then primarily on the learnt priors in later trials, resulting in bimodal estimate distributions across a block. We used the slopes of estimates versus displayed directions (Figure 3A) as a measure of how biased subjects were by the prior (veridical is a one-to-one slope and complete bias to the prior is a slope of 0). Comparison of these slopes in early versus late trials of prior blocks suggested that priors were learned quickly, consistent with observations from other studies (Berniker et al., 2010). The slope in the first hundred trials were well correlated with the last 100 trials' slopes ($r = 0.84$, 95% CI [0.73 0.94] over subjects, $n = 12$ subjects, all $p < 0.01$: $p = 1.7e-6$, $p = 4.2e-5$, $p = 4.1e-6$, $p = 1.05e-6$, $p = 7e-3$, $p = 3e-4$, $p = 2.9e-4$, $p = 0.33$, $p = 2.1e-5$, $p = 7e-3$, $p = 9.1e-8$, $p = 9.8e-5$, $n = 12$ combinations of priors and coherences, respectively, for subjects 1 to 12, two-tailed Pearson correlation). Moreover, a modified version of the Basic Bayesian observer confirmed that priors were learned quickly. Specifically, prior strength was allowed to grow exponentially across trials with a learning rate parameter, τ (see STAR Methods). This analysis showed that, on average, 95% of prior strength was learnt within 45 trials of the typical 8,000 trials that subjects undergo during the experiment ($\tau = 14.9$, 95% CI [6.1 23.7] over subjects, $n = 12$ subjects). In addition, we fitted the Switching observer to the first and last halves of each subject's data and compared the model fit parameters between halves. The model goodness of fit, the AICs, did not significantly

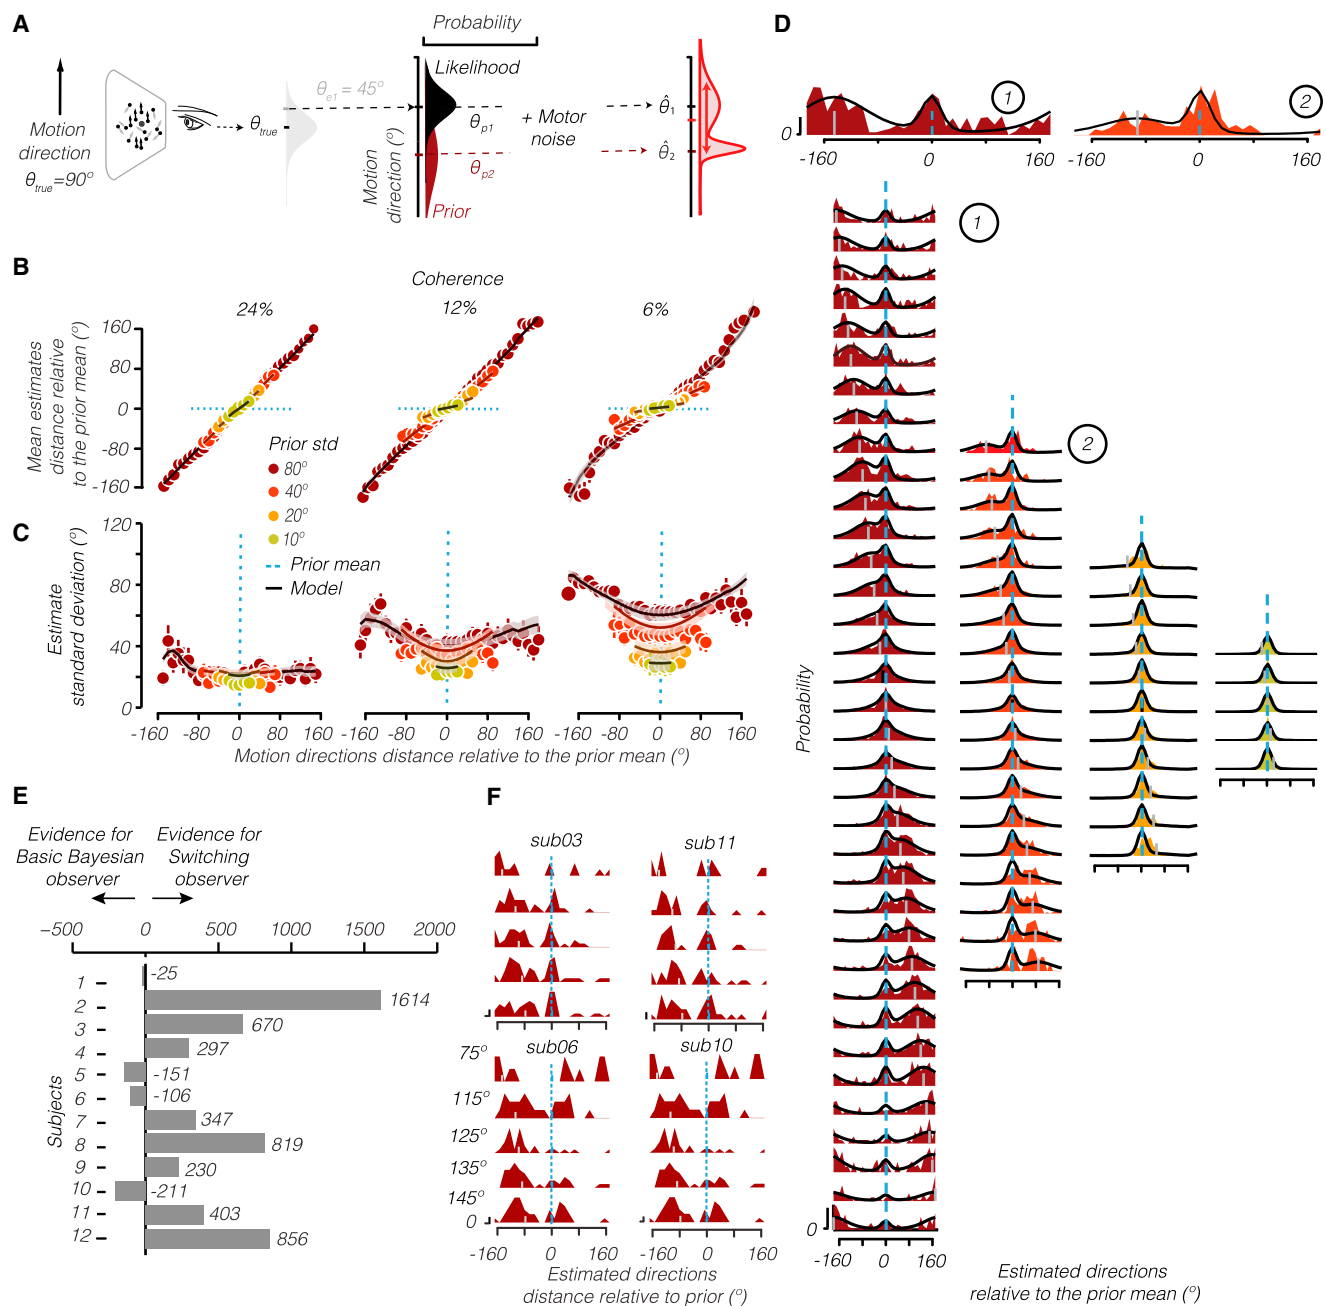


Figure 5. Switching Observer Fits of Summary Statistics and Bimodality of Subject Estimate Distributions

(A–D) The Switching observer (A) is based on the same likelihood and prior distributions as the Basic Bayesian observer but, rather than multiplying the two together, switches between estimates from the prior mean and likelihood according to their relative strengths. The Switching observer fits subject estimates mean (B) and standard deviation (C) as well as the bimodal estimate distributions (D). All conventions for (A)–(D) are the same as Figures 2, 3, and 4.

(E) AIC differences between the Basic Bayesian and the Switching observers for each subject.

(F) Four example subjects representing good to poorer fits of the Switching model. Bimodality is evidenced by distributions at the prior (blue line) and sensory evidence (gray tick mark), for five motion directions distant from the prior, at 6% coherence and for the 80° prior. See Figure S3 for switching probabilities and Switching observer's best-fit parameters and Figures S4 and S5 for results in the motion direction experiment with 80° prior only. Conventions for (F) are as in Figure 4 except that a histogram bin of 15° was used instead of 10° to reduce noise.

differ between the first and second halves of the experiment for each subject ($Z = 32$, $p = 0.62$, Wilcoxon signed rank test, $n = 12$ subjects). The prior strength best-fit parameters also did not

significantly change for each subject between the first and the second halves of the experiment ($Z = 37$, $p = 0.90$ for 80° prior; $Z = 25$, $p = 0.30$ for 40° prior; $Z = 23$, $p = 0.23$ for 20° prior;

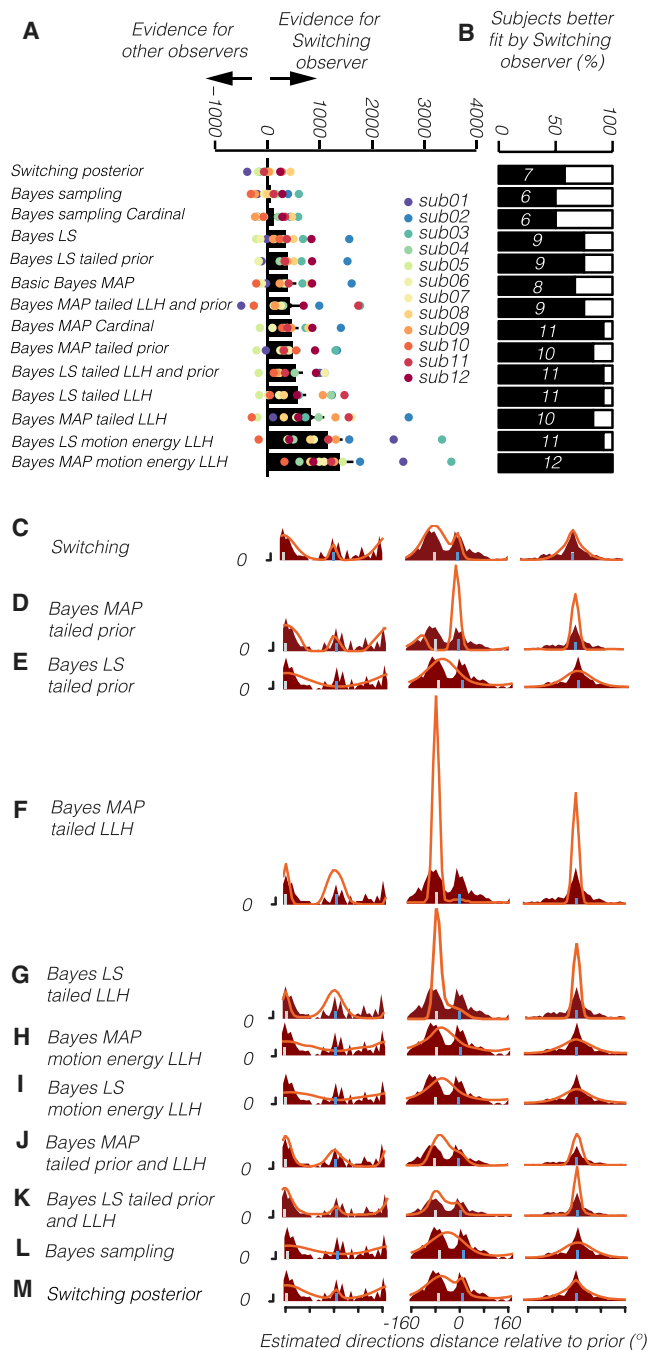


Figure 6. Alternative Model Comparisons

(A and B) The AIC difference between alternative Bayesian observers and the Switching observer averaged over subjects (A) and number of participants better fit by the Switching observer (B) evidence stronger support for the Switching observer.

(C–M) Example subject estimate distributions (brown histograms) with observer model simulated predictions (red curves) demonstrate failure modes of alternative models: Switching (C), Bayes MAP long-tailed prior (D), Bayes LS long-tailed prior (E), Bayes MAP long-tailed LLH (F), Bayes LS long-tailed LLH (G), Bayes MAP motion energy LLH (H), Bayes LS motion energy LLH (I), Bayes MAP tailed prior and LLH (J), Bayes LS tailed prior and LLH (K), Bayes sampling (L), and Switching posterior (M). Estimate distributions are for

$Z = 18$, $p = 0.10$ for 10° prior, $n = 12$ subjects, two-tailed Wilcoxon signed rank test). These results indicate that subjects quickly learnt the prior strengths and switched across the experiment.

The Basic Bayesian observer may have failed to capture the bimodality of estimate distributions because of the assumptions we made about the shape of the prior distribution. In particular, we assumed a peaked, bell-shaped distribution (von Mises) matching the actual distribution of directions, but subjects may have priors that allow for higher probability of random directions. For example, subjects may have learned a prior with long tails in which there is some probability that motion directions are chosen not from the actual prior, but from a random direction. The multiplication of such a longer-tailed prior distribution could result in two peaks and thus better fit the bimodal behavioral data. We therefore tested a variation of the Basic Bayesian observer in which the prior representation could have different shapes as parameterized by a mixture distribution of uniform and von Mises (see [STAR Methods](#)). This Bayesian observer with long-tailed priors did not fit the behavioral data better than the Switching observer (Figure 6A, Bayes maximum a posteriori [MAP] tailed prior). The average AIC difference was 483, 95% CI [219 756] over subjects, in favor of switching for 10 out of 12 subjects (Figure 6B). The estimate distributions produced by the model were bimodal but qualitatively (Figure 6D) and quantitatively (Figure 6A) differed from the subject estimate distributions. In particular, priors that have a strong enough effect to create a bimodal distribution when the presented direction and the mean of the prior are far away (Figure 6D, left column) are too strong for when likelihood and prior mean are close together (Figure 6D, center column). This causes there to be a much larger peak than present in the data at the prior mean in these conditions, which were well accounted for by the Switching observer (Figure 6C).

Our models did not take into account the stochasticity of the stimulus, which could give rise to bimodal estimate distributions without explicitly modeling switching. For each stimulus presentation, some dots moved coherently while others moved in random directions. This amounts to sampling dot directions from a long-tailed delta generative distribution (mixture of delta and uniform), which peaks at the coherent direction and whose peak amplitude is the coherence. This sensory likelihood distribution shape could result in a bimodal estimate distribution because of changes in the trial-by-trial stimulus statistics. When random sampling produces a low proportion of coherent dots, the likelihood is closer to uniform, resulting in a large bias toward the prior. When more dots are coherent, the likelihood will be more peaked and estimates will be biased toward the presented direction. To test this hypothesis, we replaced the von Mises sensory likelihoods in the Basic Bayesian observer with sensory likelihoods based directly on the stimulus characteristics (see [STAR Methods](#)). This indeed produced bimodal estimate distributions when sensory likelihood and mean of the prior were distant (Figure 6F, left column). However, analogous

6% motion coherence and the weakest prior (80° standard deviation) for motion directions (gray tick mark) far from the prior mean (left), near the prior mean (middle), and at the prior mean (blue vertical bar, right). LS, least-squares readout; MAP, Maximum a posteriori readout; LLH, likelihood.

to the long-tailed prior model, the likelihood was too strong when prior mean and presented direction were close together, resulting in poorer fits for these conditions (Figure 6F, center column). The Switching observer provided a better quantitative fit of the data (Figure 6A, Bayes MAP tailed likelihood [LLH]). The average AIC difference was 829, 95% CI [342 1316] over subjects, in favor of switching for 10 out of 12 subjects (Figure 6B). Model variations in which sensory likelihood was derived from biologically plausible motion energy filters (Adelson and Bergen, 1985) also failed to reproduce the bimodality observed in the data (Figure 6H). The average AIC difference was 1,377, 95% CI [872 1881] over subjects, in favor of switching for all 12 subjects (Figures 6A and 6B, Bayes MAP motion energy LLH).

We also combined the two above observer models to create a model that included both a long-tail prior and likelihood function. This observer again performed well when the presented direction was far from the prior (Figure 6J, left column) but was outperformed by the Switching observer when presented direction and prior mean were close together (Figure 6J, center and right columns). In particular, the peak produced became too large for directions nearly matched to the prior mean, and it quantitatively failed to produce better fits (Figure 6A, Bayes MAP tailed LLH and prior). The average AIC difference was 429, 95% CI [12 847] over subjects, in favor of switching for 9 out of 12 subjects (Figure 6B).

While both the Basic Bayesian and Switching observers are predicated on learning the prior distribution, subjects may have used information just from the previous trial (Abrahamyan et al., 2016; Fründ et al., 2014), for example, biasing estimations on the previous motion direction, which would be a sample from the prior (Fischer and Whitney, 2014; Raviv et al., 2012; Verstynen and Sabes, 2011) and thus approximate Bayesian inference. If so, then trial estimates should be biased toward the previous motion direction even if that direction is opposite to the mean of the prior. However, when trials were sorted such that current displayed directions were positioned between prior mean and previous directions, subject estimates were significantly biased toward the prior mean and not the previous motion direction (the average estimate over all subject trials was significantly biased toward the prior mean, distant from the current direction by 4.35° , 95% CI [3.14 5.57], $n = 7,504$ trials, which is significantly larger than 0), suggesting that they learnt priors over stimulus history and did not solely rely on the previous stimulus. The bias toward the prior was significant for 4 out of 12 subjects (the average distance between the estimate and the current direction was 2.93° , 95% CI [0.03 5.82] for subject 1, $n = 758$ trials; 21.40° , 95% CI [17.30 25.51] for subject 2, $n = 691$ trials; 7.29° , 95% CI [4.55 10.02] for subject 3, $n = 801$ trials; 12.24° , 95% CI [6.70 17.78] for subject 10, $n = 548$ trials). The remaining 8 out of 12 subjects were not significantly biased in either direction (the average distance between the estimate and the current direction was -0.56° , 95% CI [-5.66 4.55] for subject 4, $n = 387$ trials; -2.25° , 95% CI [-7.39 2.88] for subject 5, $n = 553$ trials; -0.43° , 95% CI [-4.18 3.32] for subject 6, $n = 699$ trials; 1.30° , 95% CI [-2.75 5.35] for subject 7, $n = 530$ trials; -1.99° , 95% CI [-7.59 3.62] for subject 8, $n = 508$ trials; 3.97° , 95% CI [0.57 7.38] for subject 9, $n = 782$ trials; 1.57° , 95% CI [-2.34 5.49] for

subject 11, $n = 656$ trials; 1.38° , 95% CI [-3.97 6.72] for subject 12, $n = 589$ trials).

The Basic Bayesian observer makes the assumption that subjects optimally read out the posterior (i.e., use its mode). We relaxed this assumption with a Bayesian sampling observer in which estimates were produced by sampling the posterior (Moreno-Bote et al., 2011). This Bayesian sampling observer failed to qualitatively describe subject estimate distributions as it produced a unimodal distribution (Figure 6L) similar to the Basic Bayesian observer. The average AIC difference was 66, 95% CI [-84 217] over subjects (Figure 6A, Bayes sampling), in favor of switching for 6 out of 12 subjects (Figure 6B). We note that when the subject estimate distributions were particularly noisy, the Sampling observer could sometimes provide fits that quantitatively appeared competitive with other models by placing a single posterior over both prior and likelihood (Figure 6I), even though it failed to reproduce the bimodality of the estimates.

Least-squares readout, a common rule that uses the posterior mean rather than the mode, gave similar results. In particular, a Basic Bayesian model with least-squares readout (Figure 6A, Bayes LS) failed to outperform the Switching observer. The average AIC difference was 363, 95% CI [85 640] over subjects, in favor of switching for 9 out of 12 subjects (Figure 6B). Least-squares readout variations of all other models (Figure 6A, models with LS) either made similar predictions to the MAP readouts, presumably because the posterior distributions were not significantly skewed (Figures 6G, 6I, and 6K), or would tend to lose bimodality because of averaging across the full posterior rather than picking the maximum value (Figure 6E).

Bimodal estimate distributions might also be produced by motion direction oblique effects (Ball and Sekuler, 1980; Gros et al., 1998) similar to those that have been shown to explain biases in orientation discrimination (Girshick et al., 2011). We defined a cardinal direction prior as a mixture of four von Mises distributions that peaked at the cardinal motion directions (0° , 90° , 180° , 270°), and this additional prior was integrated in each trial with the priors learnt during the task and the sensory likelihood. We fit a Basic Bayesian observer with cardinal priors and either a maximum a posteriori readout or sampling readout to subject trial estimates and found that the Switching observer quantitatively (Figure 6A, Bayes MAP cardinal and Bayes sampling cardinal) and qualitatively outperformed these two observers, which both predicted more than two peaks for all motion directions. The Switching observer fitted the data better than the Bayesian observer with cardinal prior. The average AIC difference was 471, 95% CI [234 708] over subjects, in favor of switching for 11 out of 12 subjects (Figure 6B). The switching observer also fitted better than the Bayesian sampling observer with cardinal prior. The average AIC difference was 123, 95% CI [-44 290] over subjects, in favor of switching for 6 out of 12 subjects (Figure 6B).

Bimodal distributions of estimates might be due to a switching process not between prior and likelihood, but prior and posterior. That is, an observer might engage in optimal Bayesian inference on some trials but might strategically report the prior mean on other trials when sensory uncertainty is too high. We tested an observer model (Switching posterior) that chooses estimates at the prior mode with probability, p_{prior} , and otherwise combines sensory likelihood and prior into a posterior and switches to

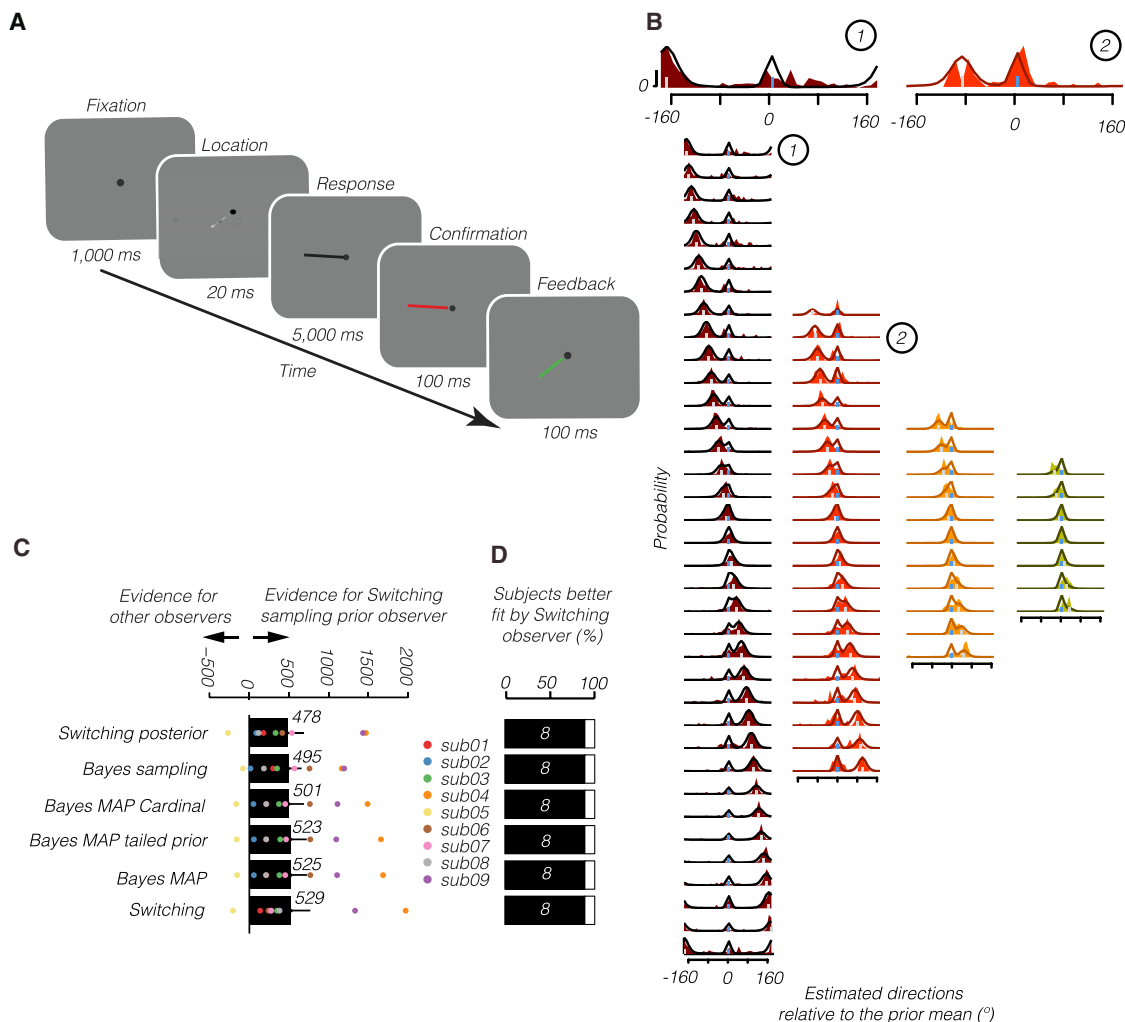


Figure 7. Estimate Distributions for Spatial Orientation Task Also Show Bimodality

Subjects' estimate the spatial orientation of a bar (A) whose sensory evidence strength is controlled by contrast and prior by the distribution of orientations shown in a block. Estimate distributions (B) were bimodal and well fit (C and D) by a Switching observer with unstable prior mean (Switching sampling prior observer). All conventions are the same as Figure 5. See also Figures S6 and S7 for subject-by-subject estimate distributions.

the posterior mode (see STAR Methods, Equation 6). While this Switching posterior model produces bimodal estimates and does a reasonable job of accounting for the data, the original Switching observer better fit the data (Figure 6A, Switching posterior), outperforming the Switching posterior observer with an average AIC difference of 21, 95% CI [−103 145] over subjects, in favor of our Switching observer for 7 out of 12 subjects (Figure 6B). The Switching posterior observer predicted estimate distribution peaks at the posterior, i.e., between the motion direction and prior, which deviates from the peak at the motion direction observed in our data (Figure 6M).

Bimodality of estimate distributions could be a result of a particular strategy specific to having been exposed to blocks of trials with very narrow priors. We therefore had five naïve subjects perform the motion direction estimation task in which they were only exposed to the wide (80° standard deviation) prior. Subject-averaged estimate distributions (Figure S4A) and sub-

ject-by-subject estimate distributions (Figure S5) showed clear bimodality. Quantitative fits of all the models showed that the Switching observer outperformed all the other models (Figures S4B and S4C). Thus, the Switching behavior was not a consequence of exposure to very sharp prior distributions.

Switching between priors and sensory evidence based on their strengths could be a generalized mechanism by which humans can realize perceptual inference. We therefore asked whether subjects would show similar switching behavior in other tasks. We had nine new subjects perform a spatial orientation estimation task that was similar to the motion direction estimation task except rather than using a patch of random dots and have subjects estimate motion direction, a thin bar filled with filtered noise appeared in 36 different possible spatial orientations ranging from 5° to 355° in steps of 10°, and subjects had to estimate its spatial orientation following the same task procedures as in the motion direction estimation task (Figure 7A). We manipulated

the sensory likelihood by changing the contrast (15.6% and 100%) of the stimulus; the lower the contrast, the more uncertain the spatial orientation should be. At low contrast (15.6%), subjects' data evidenced bimodal estimate distributions around the prior mean and likelihood in subject-averaged data (Figure 7B) and for individual subjects (Figures S6 and S7). However, compared to the direction estimation task, we found that subjects in the spatial orientation task were less certain of the prior mean. Thus, the Switching observer restricted to making estimates at the prior mean did not fit as well as a Switching observer in which instability to the prior mean was modeled by sampling the prior (see STAR Methods). This Switching prior sampling observer outperformed the Basic Bayesian observer with an average AIC difference of 525, 95% CI [150 900], in favor of Switching sampling prior for 8 out of 9 subjects (Figures 7C and 7D). In the direction estimation task, it outperformed all other models with average AIC difference of 82, 95% CI [−41 206]; 103, 95% CI [−58 264]; 149, 95% CI [67 231]; 205, 95% CI [131 279]; 445, 95% CI [241 649]; 477, 95% CI [269 684]; 478, 95% CI [262 693]; 511, 95% CI [185 839]; 554, 95% CI [392 716]; 570, 95% CI [379 760]; 627, 95% CI [403 850]; 657, 95% CI [377 939]; 911, 95% CI [464 1,359]; 1,233, 95% CI [746 1,720]; 1,459, 95% CI [995 1,923], for Switching, Switching posterior, Bayes sampling, Bayes sampling cardinal, Bayes LS, Bayes LS tailed prior, Bayes MAP, Bayes MAP tailed LLH and Prior, Bayes MAP cardinal, Bayes MAP tailed prior, Bayes LS tailed LLH and Prior, Bayes LS tailed LLH, Bayes MAP tailed LLH, Bayes LS motion energy LLH, and Bayes MAP motion energy LLH, respectively.

DISCUSSION

We used a motion direction estimation task in which we manipulated motion strength and motion direction distributions to test whether humans combine prior knowledge and sensory evidence into perceptual judgments optimally as prescribed by a Basic Bayesian observer. We found that a Switching observer, which approximates optimal Bayesian integration but switches between prior mean and sensory evidence, best accounted for subjects' data compared to other models.

In agreement with statistical optimality, when evidence was weak (low motion coherence) and prior strong (narrow motion direction distribution), subject's average estimates were biased toward the prior mean and the bias increased when priors became stronger and sensory evidence weaker, supporting findings from motion perception (Chalk et al., 2010; Hanks et al., 2011) and sensorimotor tasks (Berniker et al., 2010; Körding and Wolpert, 2004; Tassinari et al., 2006; Vilares et al., 2012). Furthermore, we confirmed the Bayesian predictions that estimate variability should decrease when prior and evidence become stronger, in line with previous findings that the variability of human motion direction estimates increases when stimulus contrast becomes lower (Chalk et al., 2010) and that human time interval judgments are less variable when time interval distributions become narrower (Jazayeri and Shadlen, 2010).

However, in apparent contradiction to the prediction of a Basic Bayesian observer, subject estimate distributions showed two peaks: at the prior mean and at the actual motion direction. We note that this observation was possible because the estima-

tion task that we used allowed us to see the full distributions of estimates that would have been obscured by a two-alternative forced choice task. We developed a Switching observer that could account for this finding by choosing between estimates from the prior mean or evidence based on their relative strengths. The Switching observer not only provided better qualitative and quantitative fits to estimate distributions than the Basic Bayesian observer or recency heuristics (Fischer and Whitney, 2014; Körding and Wolpert, 2004; Raviv et al., 2012; Wolpert et al., 1995), but could also approximate average estimate mean and variability. We showed that the estimation bimodality apparent in the motion direction estimation task generalized to a spatial orientation estimation task, for which the Switching observer with instability in the prior estimate provided a better account of subject estimates than the Basic Bayesian observer.

Our study substantiates other reports that subjects can learn quickly and accurately about both the prior mean and its strength. Our analysis showed that subjects could learn prior strengths quickly and stably within 100 trials, in line with previous works (Berniker et al., 2010; Chalk et al., 2010; Krishnamurthy et al., 2017). The Switching observer best fit prior strengths, suggesting that priors were learnt accurately, which supports previous assumptions of optimality in prior representation (Berniker et al., 2010; Tenenbaum et al., 2011) and the idea that people can learn and operate on an accurate representation of priors (Chalk et al., 2010; Fetsch et al., 2013; Girshick et al., 2011; Grifiths and Tenenbaum, 2006; von Helmholtz, 1924; Knill and Richards, 1996; Körding and Wolpert, 2004; Krishnamurthy et al., 2017; Stocker and Simoncelli, 2006; Weiss et al., 2002). That subjects' behavior was quickly and stably affected by the distribution of trials in a session demonstrates that they have access to information about the strength of the prior needed to make optimal decisions according to Bayesian theory. Indeed, the Switching observer model fits suggest that the width of the learned priors were reasonably accurate. Therefore, the deviance from the expectation of the Basic Bayesian observer that we found in the bimodality of estimate distributions was not due to an inability to correctly learn and use information about prior strength. The Switching model suggests that subjects can use prior strength in a sophisticated and sensible way—biasing choices across trials more toward the prior mean when they are sure of the prior and more toward the presented direction when they are uncertain of the prior.

That the Switching observer fit the behavioral data better than all the Bayesian observers we tested should not be taken as evidence against our subjects using a Bayesian strategy per se. While model comparison statistics and qualitative assessment all point to the Switching observer as being the best description of the estimation behavior, a few Bayesian observers come close. In particular, the Switching posterior observer would suggest that observers mandatorily form a posterior on some trials but, on other trials, give up and choose the prior. This observer could be construed as a lazy Bayesian, heuristically choosing the prior to avoid the cost of computing the full posterior when appropriate. Alternatively, a hierarchical Bayesian observer could be formulated in which perceptual judgments are determined by hierarchical beliefs, or hyper-priors

(Lee and Mumford, 2003; Sato et al., 2007; Tenenbaum et al., 2011) that motion directions were drawn from either a uniform or a peaked distribution at each trial, with probabilities determined by the ratio of likelihood and prior strengths. This would effectively amount to a reformulation of our Switching observer in Bayesian terminology. Regardless of these alternative accounts, or future efforts that might uncover better models for the bimodality in our data, our results clearly document bimodality of estimates that are at odds with the standard, basic Bayesian observer.

Some estimation behaviors appear to deviate from optimal but still might be the consequence of optimal considerations. For example, biased estimates of direction of motion could result from optimal read out of neural populations (Jazayeri and Movshon, 2007; Wei and Stocker, 2015). Estimates repelled rather than attracted to the prior can be accounted for by asymmetric likelihood functions (Wei and Stocker, 2015) that could be the result of neural tuning functions whose statistics are matched to those of the environment (Ganguli and Simoncelli, 2010). We have similarly explored different forms of likelihood and prior distributions, namely mixtures of uniform and von Mises distributions, and found that while these could produce bimodal estimate distributions, the Switching model provided better fits of the data. We cannot preclude the possibility that there may be some distributional form that would allow a Bayesian model with multiplicative integration to better explain the data. Nonetheless, the Switching model demonstrates that it is possible to approximate the summary statistics of a Bayesian observer without multiplicative integration and match the actual estimate distributions produced by human observers.

The Switching observer could be implemented in the brain by probabilistic population codes (Beck et al., 2008; Ma et al., 2006; Pouget et al., 2013) or through representations of priors or expectancy (Girshick et al., 2011; Gorlin et al., 2012; Kok et al., 2013, 2012; Summerfield and Koehlin, 2008; Vintch and Gardner, 2014; Yang et al., 2012) and be less costly than full Bayesian integration. Bayesian integration requires computing the product of likelihood and prior probabilities over a large number of hypothetical stimulus states, which has been proposed to occur through the summation of two neural population responses (Ma et al., 2006). Switching, however, provides a mechanism of inference without the costs of integration: the brain only needs to read out the neurons that show the greatest responses from the two populations instead of integrating all neurons' responses.

Optimal and heuristic solutions are not mutually exclusive options. Thus, our results do not imply that statistically optimal inference behaviors can never be achieved by humans. Indeed, in different tasks and contexts, it has been proposed that humans use a wide variety of strategies with some being characterized as heuristic (Gigerenzer and Gaissmaier, 2011; Nassar et al., 2010; Rahnev and Denison, 2016; Raviv et al., 2012; Tversky and Kahneman, 1974; Vul et al., 2014; Wilson et al., 2013) and others as statistically optimal (Ernst and Banks, 2002; Girshick et al., 2011; Jazayeri and Shadlen, 2010; Körding and Wolpert, 2004; Marr, 1982; Najemnik and Geisler, 2005; Nassar et al., 2010; Rahnev and Denison, 2016; Stocker and Simoncelli, 2006; Vul et al., 2014; Wei and Stocker, 2015; Weiss

et al., 2002; Wilson et al., 2013). Humans may use different heuristics in different tasks based on the adaptive advantage that they provide in different environments (Gigerenzer and Gaissmaier, 2011). For example, one would gain from relying on the previous stimulus in a volatile environment (Fischer and Whitney, 2014; Raviv et al., 2012; Summerfield et al., 2011; Wilson et al., 2013) while acquiring a prior over stimulus history and using it will reduce average estimate error when the environment is stable (Girshick et al., 2011; Stocker and Simoncelli, 2006; Summerfield et al., 2011; Wei and Stocker, 2015). When the space of hypothetical stimuli relevant to a task becomes too large or the shape of the prior probability distribution over that space becomes too complex, one might have to sample the posterior, thus reducing the cost of the computation (Vul et al., 2014). What drives the brain to select a particular strategy may depend on processing constraints among which the complexity of prior features might play an important role (Acerbi et al., 2014).

The switching behavior that we have documented here suggests that behaviors can have qualities consistent with both optimal theory and heuristics; across trials, mean and standard deviation of estimates matched well with optimal predictions despite the fact that closer examination showed behavior that switched between prior and likelihood. More generally, human behavior may best be understood in a framework that views the goals of behavior as being derived by optimality theory but constraints of implementation playing an important role in determining actual behavior (Simon, 1955, 1979). Rather than viewing optimal and heuristic theories as being in apparent tension, viewed together they give a more complete understanding of the forces that shape human behavior.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
 - Human Subjects
- METHOD DETAILS
 - Stimuli
 - Motion Direction Estimation Task
 - Spatial Orientation Task
 - Analyses
 - Basic Bayesian Observer Model
 - Prior Learning Model
 - Switching Observer Model
 - Bayesian Models with Variations in Prior
 - Bayesian Models with Variations in Likelihood
 - Sampling Bayesian Observer Model
 - Bayesian Models with Least-Square Readout
 - Bayesian Models with Learnt and Cardinal Priors
 - Model Fit Convergence
 - Model Comparison
- QUANTIFICATION AND STATISTICAL ANALYSIS
- DATA AND SOFTWARE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information includes seven figures and can be found with this article online at <https://doi.org/10.1016/j.neuron.2017.12.011>.

ACKNOWLEDGMENTS

This research was funded by a Grants-in-Aid for Scientific Research 24300146 (to J.L.G.) from the Japanese Ministry of Education, Culture, Sports, Science and Technology (MEXT), the Hellman Faculty Scholar Fund, and a Low Vision Research Award from Research to Prevent Blindness and Lions Clubs International Foundation. We particularly thank Ilias Rentzeperis and Cameron McKenzie for continuous discussions and feedback. We thank Kenji Haruhana and Toshiko Ikari for administrative assistance and Arman Abrahamyan, Sivar-amakrishnan R. Kaveri, Ansgar Koene, Dylan Cable, and Daniel Birman for their feedback. We also thank Wei Ji Ma for advice on modeling stimulus stochasticity and Matteo Carandini, David Heeger, Tirin Moore, Ewart Thomas, and Noah Goodman for helpful discussions and ideas on analyses and modeling.

AUTHOR CONTRIBUTIONS

S.L. and J.L.G. designed research; S.L. performed research; S.L. and J.L.G. contributed unpublished analytic tools; S.L. analyzed data; S.L. and J.L.G. wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: May 3, 2017

Revised: October 23, 2017

Accepted: December 5, 2017

Published: December 28, 2017

REFERENCES

- Abrahamyan, A., Silva, L.L., Dakin, S.C., Carandini, M., and Gardner, J.L. (2016). Adaptable history biases in human perceptual decisions. *Proc. Natl. Acad. Sci. USA* **113**, E3548–E3557.
- Acerbi, L., Vijayakumar, S., and Wolpert, D.M. (2014). On the origins of suboptimality in human probabilistic inference. *PLoS Comput. Biol.* **10**, e1003661.
- Adams, W.J., Graf, E.W., and Ernst, M.O. (2004). Experience can change the 'light-from-above' prior. *Nat. Neurosci.* **7**, 1057–1058.
- Adelson, E.H., and Bergen, J.R. (1985). Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A* **2**, 284–299.
- Akaike, H. (1974). A new look at statistical model identification. *IEEE Transact Automatic Control* **19**, 716–723.
- Ball, K., and Sekuler, R. (1980). Models of stimulus uncertainty in motion perception. *Psychol. Rev.* **87**, 435–469.
- Ban, H., Preston, T.J., Meeson, A., and Welchman, A.E. (2012). The integration of motion and disparity cues to depth in dorsal visual cortex. *Nat. Neurosci.* **15**, 636–643.
- Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., and Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142–1152.
- Beck, J.M., Ma, W.J., Pitkow, X., Latham, P.E., and Pouget, A. (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* **74**, 30–39.
- Berniker, M., Voss, M., and Kording, K. (2010). Learning priors for Bayesian computations in the nervous system. *PLoS ONE* **5**, e12686.
- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning* (Springer).
- Chalk, M., Seitz, A.R., and Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *J. Vis.* **10**, 2.
- Chater, N., and Manning, C.D. (2006). Probabilistic models of language processing and acquisition. *Trends Cogn. Sci.* **10**, 335–344.
- Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Hecce Castañón, S., and Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron* **81**, 1429–1441.
- Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429–433.
- Fetsch, C.R., DeAngelis, G.C., and Angelaki, D.E. (2013). Bridging the gap between theories of sensory cue integration and the physiology of multisensory neurons. *Nat. Rev. Neurosci.* **14**, 429–442.
- Fischer, J., and Whitney, D. (2014). Serial dependence in visual perception. *Nat. Neurosci.* **17**, 738–743.
- Forbes, C., Evans, M., Hastings, N., and Peacock, B. (2011). *von Mises Distribution*. In *Statistical Distributions*, Fourth Edition (John Wiley & Sons), pp. 191–192.
- Freeman, W.T. (1994). The generic viewpoint assumption in a framework for visual perception. *Nature* **368**, 542–545.
- Fründ, I., Wichmann, F.A., and Macke, J.H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *J. Vis.* **14**, 9.
- Ganguli, D., and Simoncelli, E.P. (2010). Implicit encoding of prior probabilities in optimal neural populations. *Adv. Neural Inf. Process. Syst.* **2010**, 658–666.
- Georgopoulos, A.P., Taira, M., and Lukashin, A. (1993). Cognitive neurophysiology of the motor cortex. *Science* **260**, 47–52.
- Ghahramani, Z., Wolpert, D.M., and Jordan, M.I. (1997). Computational models of sensorimotor integration. *Adv. Psychol.* **119**, 117–147.
- Gigerenzer, G., and Gaissmaier, W. (2011). Heuristic decision making. *Annu. Rev. Psychol.* **62**, 451–482.
- Girshick, A.R., Landy, M.S., and Simoncelli, E.P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nat. Neurosci.* **14**, 926–932.
- Gorlin, S., Meng, M., Sharma, J., Sugihara, H., Sur, M., and Sinha, P. (2012). Imaging prior information in the brain. *Proc. Natl. Acad. Sci. USA* **109**, 7935–7940.
- Griffiths, T.L., and Tenenbaum, J.B. (2006). Optimal predictions in everyday cognition. *Psychol. Sci.* **17**, 767–773.
- Gros, B.L., Blake, R., and Hiris, E. (1998). Anisotropies in visual motion perception: a fresh look. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **15**, 2003–2011.
- Hanks, T.D., Mazurek, M.E., Kiani, R., Hopp, E., and Shadlen, M.N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* **31**, 6339–6352.
- Hansen, N., and Kern, S. (2004). Evaluating the CMA evolution strategy on multimodal test functions. In *Parallel Problem Solving from Nature - PPSN VIII*, X. Yao, E. Burke, J.A. Lozano, J. Smith, J.J. Merelo-Guervós, J.A. Bullinaria, J. Rowe, P. Tino, A. Kabán, and H.-P. Schwefel, eds. (Springer Berlin Heidelberg), pp. 282–291.
- Jazayeri, M., and Movshon, J.A. (2006). Optimal representation of sensory information by neural populations. *Nat. Neurosci.* **9**, 690–696.
- Jazayeri, M., and Movshon, J.A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* **446**, 912–915.
- Jazayeri, M., and Shadlen, M.N. (2010). Temporal context calibrates interval timing. *Nat. Neurosci.* **13**, 1020–1026.
- Jogan, M., and Stocker, A.A. (2015). Signal integration in human visual speed perception. *J. Neurosci.* **35**, 9381–9390.
- Kim, H.R., Angelaki, D.E., and DeAngelis, G.C. (2015). A novel role for visual perspective cues in the neural computation of depth. *Nat. Neurosci.* **18**, 129–137.
- Knill, D.C., and Richards, W. (1996). *Perception as Bayesian Inference* (Cambridge University Press).
- Kok, P., Jehee, J.F.M., and de Lange, F.P. (2012). Less is more: expectation sharpens representations in the primary visual cortex. *Neuron* **75**, 265–270.

- Kok, P., Brouwer, G.J., van Gerven, M.A.J., and de Lange, F.P. (2013). Prior expectations bias sensory representations in visual cortex. *J. Neurosci.* 33, 16275–16284.
- Körding, K.P., and Wolpert, D.M. (2004). Bayesian integration in sensorimotor learning. *Nature* 427, 244–247.
- Krishnamurthy, K., Nassar, M.R., Sarode, S., and Gold, J.I. (2017). Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nat. Hum. Behav.* Published online May 8, 2017. <https://doi.org/10.1038/s41562-017-0107>.
- Lee, T.S., and Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 20, 1434–1448.
- Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438.
- Marr, D. (1982). *Vision* (MIT Press).
- Moreno-Bote, R., Knill, D.C., and Pouget, A. (2011). Bayesian sampling in visual perception. *Proc. Natl. Acad. Sci. USA* 108, 12491–12496.
- Najemnik, J., and Geisler, W.S. (2005). Optimal eye movement strategies in visual search. *Nature* 434, 387–391.
- Nassar, M.R., Wilson, R.C., Heasley, B., and Gold, J.I. (2010). An approximately Bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *J. Neurosci.* 30, 12366–12378.
- Nassar, M.R., Bruckner, R., Gold, J.I., Li, S.-C., Heekeren, H.R., and Eppinger, B. (2016). Age differences in learning emerge from an insufficient representation of uncertainty in older adults. *Nat. Commun.* 7, 11609.
- Neiman, T., and Loewenstein, Y. (2011). Reinforcement learning in professional basketball players. *Nat. Commun.* 2, 569.
- Pouget, A., Beck, J.M., Ma, W.J., and Latham, P.E. (2013). Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16, 1170–1178.
- Rahnev, D., and Denison, R. (2016). Suboptimality in perception. *bioRxiv*. <https://doi.org/10.1101/060194>.
- Ramachandran, V.S. (1988). Perception of shape from shading. *Nature* 337, 163–166.
- Rao, V., DeAngelis, G.C., and Snyder, L.H. (2012). Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. *J. Neurosci.* 32, 10063–10074.
- Rauber, H.J., and Treue, S. (1998). Reference repulsion when judging the direction of visual motion. *Perception* 27, 393–402.
- Raviv, O., Ahissar, M., and Loewenstein, Y. (2012). How recent history affects perception: the normative approach and its heuristic approximation. *PLoS Comput. Biol.* 8, e1002731.
- Raviv, O., Lieder, I., Loewenstein, Y., and Ahissar, M. (2014). Contradictory behavioral biases result from the influence of past stimuli on perception. *PLoS Comput. Biol.* 10, e1003948.
- Sato, Y., Toyoizumi, T., and Aihara, K. (2007). Bayesian inference explains perception of unity and ventriloquism aftereffect: identification of common sources of audiovisual stimuli. *Neural Comput.* 19, 3335–3355.
- Sharot, T., Korn, C.W., and Dolan, R.J. (2011). How unrealistic optimism is maintained in the face of reality. *Nat. Neurosci.* 14, 1475–1479.
- Simon, H.A. (1955). A behavioral model of rational choice. *Q. J. Econ.* 69, 99–118.
- Simon, H.A. (1979). Rational decision-making in business organizations. *Am. Econ. Rev.* 69, 493–513.
- Stocker, A.A., and Simoncelli, E.P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nat. Neurosci.* 9, 578–585.
- Summerfield, C., and Koechlin, E. (2008). A neural representation of prior information during perceptual inference. *Neuron* 59, 336–347.
- Summerfield, C., Behrens, T.E., and Koechlin, E. (2011). Perceptual classification in a rapidly changing environment. *Neuron* 71, 725–736.
- Tassinari, H., Hudson, T.E., and Landy, M.S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *J. Neurosci.* 26, 10154–10163.
- Tenenbaum, J.B., Griffiths, T.L., and Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* 10, 309–318.
- Tenenbaum, J.B., Kemp, C., Griffiths, T.L., and Goodman, N.D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science* 331, 1279–1285.
- Tversky, A., and Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science* 185, 1124–1131.
- Verstynen, T., and Sabes, P.N. (2011). How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J. Neurosci.* 31, 10050–10059.
- Vilares, I., Howard, J.D., Fernandes, H.L., Gottfried, J.A., and Kording, K.P. (2012). Differential representations of prior and likelihood uncertainty in the human brain. *Curr. Biol.* 22, 1641–1648.
- Vintch, B., and Gardner, J.L. (2014). Cortical correlates of human motion perception biases. *J. Neurosci.* 34, 2592–2604.
- von Helmholtz, H. (1924). *Helmholtz's Treatise on Physiological Optics* (The Optical Society of America).
- Vul, E., Goodman, N., Griffiths, T.L., and Tenenbaum, J.B. (2014). One and done? Optimal decisions from very few samples. *Cogn. Sci.* 38, 599–637.
- Wei, X.-X., and Stocker, A.A. (2015). A Bayesian observer model constrained by efficient coding can explain 'anti-Bayesian' percepts. *Nat. Neurosci.* 18, 1509–1517.
- Weiss, Y., Simoncelli, E.P., and Adelson, E.H. (2002). Motion illusions as optimal percepts. *Nat. Neurosci.* 5, 598–604.
- Wilson, R.C., Nassar, M.R., and Gold, J.I. (2013). A mixture of delta-rules approximation to bayesian inference in change-point problems. *PLoS Comput. Biol.* 9, e1003150.
- Wolpert, D.M., Ghahramani, Z., and Jordan, M.I. (1995). An internal model for sensorimotor integration. *Science* 269, 1880–1882.
- Yang, J., Lee, J., and Lisberger, S.G. (2012). The interaction of bayesian priors and sensory data and its neural circuit implementation in visually guided movement. *J. Neurosci.* 32, 17632–17645.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited Data		
Behavioral data	This paper	https://doi.org/10.17632/nxkvtrj9ps.1
Software and Algorithms		
MATLAB	MathWorks	Matlab_R2013b
R	R Development Core Team (2008)	http://www.R-project.org
Custom code (experiment, model, analyses)	This paper	https://github.com/steevelaquittaine/projInference

CONTACT FOR REAGENT AND RESOURCE SHARING

All resources, including data and codes used for the analyses on this paper, are publicly available (see [Data Software Availability](#) and [Key Resources Table](#)). Further information and requests for resource sharing should be directed to and will be fulfilled by the Lead Contact, Justin L. Gardner (jlg@stanford.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Human Subjects

Twelve healthy adults (3 female, age 29 - 40) with normal or corrected-to-normal vision participated in the motion direction estimation study, five other healthy adults (3 females, age 18 to 40) participated in the second motion direction study with only the 80° prior and nine other healthy adults (6 females, age 19 - 25) participated in the spatial orientation study. Experimental procedures were approved in advance by the RIKEN Functional MRI Safety and Ethics Committee and by the Stanford University Institutional Review Board. All subjects provided informed written consent before participating in experiments. All subjects except one, the first author (subject 1), were not aware of the purpose of the study.

METHOD DETAILS

Stimuli

Motion stimuli consisted of fields of white dots moving on a dark background in one of 36 directions (5 to 355° in steps of 10°) and with one of three motion coherences (6, 12 and 24%). Individual dots had a diameter of 3 pixels (0.1°) and were presented with a density of 16.7 dots/(°)² within a 5° diameter circular aperture. On the first video frame of motion presentation, dots were randomly positioned within the aperture. On each subsequent video frame, 6, 12 or 24% of the dots were randomly picked and displaced in the same direction by 0.028° (speed of 2.8°/sec), while the remaining dots moved in random directions. Dots that moved outside the aperture were placed at the opposite side of aperture.

Throughout each trial a gray circular fixation point with a radius of 0.2° was shown at the location of the center of the aperture. A circular fixation point was used instead of a fixation cross to prevent subjects from using the vertical and horizontal axes of the cross as references to estimate motion direction ([Berniker et al., 2010](#); [Rauber and Treue, 1998](#)).

Visual stimuli were generated using MGL (<http://justingardner.net/mgl>) running on MATLAB (The MathWorks, Inc., Natick, MA, USA) and presented on a CRT monitor (Trinitron 21-inch flat screen; Dell Computer Company). The monitor had a resolution of 1280 × 960 pixels, a refresh rate of 100 Hz, was placed 50.5 cm from the subject and had a linearized gamma.

Motion Direction Estimation Task

Subjects viewed the motion stimuli and had to report motion direction in repeated trials ([Figure 1A](#)). Subjects were instructed to estimate the motion direction as fast and accurately as possible. At trial start subjects fixated the central gray dot and had to maintain fixation throughout the trial. After 1 s, a motion stimulus appeared for 0.3 s (motion phase). Subjects then had 5 s to report the direction of motion by adjusting the orientation of a line with an electronic paddle wheel (response phase) until it matched their perceived direction and then press “1” on a keyboard. They were then provided with a brief visual confirmation of their estimated direction in which the line became red for 0.1 s (confirmation phase). At trial end, the true motion direction was displayed (green line displayed for 0.1 s at feedback phase). Subjects were not explicitly told whether their response was correct or not.

Priors and sensory likelihood were independently manipulated in the task. Priors were manipulated in each experimental block of ~200 trials by randomly drawing motion directions from one of four discrete von Mises distributions centered on 225 degrees with different widths (concentrations were 33.3, 8.7, 2.8, 0.7 equivalent to 10, 20, 40, and 80° of standard deviation). Each block of trials therefore had the same prior mean, but different strength of priors. The strength of sensory evidence was manipulated on a trial-by-trial basis by displaying the motion stimulus on each trial with a pseudo-randomly selected motion coherences of 6, 12 or 24%. Each experimental session contained four to five pseudo-randomly selected blocks and all subjects performed at least 5 sessions (> 1000 trials per prior) of the task (8, 8, 9, 5, 6, 7, 6, 6, 8, 6, 6, 6 sessions for subjects 1 to 12 respectively) to ensure that priors were well learnt (Abrahamyan et al., 2016; Berniker et al., 2010; Fründ et al., 2014). Sessions were performed on separate days.

Motion coherences and durations were chosen to encourage rapid learning of priors while avoiding potential confounds that might bias subject estimates. Motion stimuli were displayed for a short duration of 0.3 s so that even with signal integration over time the sensory evidence would be sufficiently weak enough to reveal prior-induced biases (Fischer and Whitney, 2014; Rao et al., 2012; Raviv et al., 2012; Verstynen and Sabes, 2011). The true motion directions were displayed at the end of each trial to help subjects quickly learn the experimental priors. Importantly, we avoided reward-induced biases by not delivering reward feedback which might interfere with prior learning. We note that the initial position of the response line was randomizing on each trial to prevent bias. To minimize potentially confounding biases induced by involuntary eye movement we used slow motions (2.8°/sec).

Subjects were instructed as follows: “The basic task is

- Fixate and then view a motion of noisy dots.
- Report the direction you saw as best as possible, as accurately as possible, and as fast as possible by adjusting the wheel. Press keyboard button 1 when you have made a choice. You have a maximum of 5 s to report your choice.
- Then your response will be confirmed (green line).
- You will then get feedback about the true direction of the motion (red line).

Important !!

- Please fixate the fixation point during the entire experiment.
- Your objective is to estimate and report the direction of the motion as best as possible.
- In the ideal case, your choice (green line) will match perfectly the feedback (true direction of the motion).”

Spatial Orientation Task

The spatial orientation task was similar to the motion direction estimation task except a thin bar filled with filtered noise appeared for 20 ms in one of 36 different possible spatial orientations (5° to 355° in 10° steps) and with one of two possible contrasts (15.6% and 100%) on a gray background rather than a patch of white random dots. The stimulus noise was generated by filling a 2.5° radius, 5° angle span wedge with random intensities low pass-filtered with a 2-D Gaussian filter (30° std, centered at frequency zero).

Analyses

All computations were done in the circular space of directions. Prior and likelihoods were modeled as Von Mises distributions of the following form:

$$\mathcal{V}(\theta; \mu, \kappa) = \frac{e^{\kappa \cos(\theta - \mu) - \kappa}}{2\pi I_0(\kappa)} \quad (1)$$

where θ is the circular space of directions, μ and κ are the mean and strength (or concentration parameter) of the von Mises and I_0 is the modified Bessel function of order 0. We note that when the concentration parameter becomes large, the von Mises equation becomes numerically unstable, so we used the form above (with $-\kappa$ in the numerator) and normalized to insure that the distributions summed to one. The concentration parameters of the von Mises distributions were converted into standard deviations, s , as follows:

$$s = \sqrt{\sum_{\theta=1}^{360} p(\phi(\theta, \mu)) \phi(\theta, \mu)^2} \quad (2)$$

where $\phi(\theta, \mu)$ are the signed angular distances between the displayed directions θ and the mean μ of the von Mises distributions. All probability distributions were computed over directions ranging from 1 to 360° discretized to 1°. We chose the von Mises distribution because it is regarded as the circular analog of the Gaussian distribution (Forbes et al., 2011; Moreno-Bote et al., 2011) which is a popular choice for modeling the shape of the sensory likelihood because of mathematical convenience (Ball and Sekuler, 1980; Ernst and Banks, 2002; Ghahramani et al., 1997; Gros et al., 1998; Weiss et al., 2002). Displayed directions and direction estimates were expressed as signed angular distances from the prior mean (225°, e.g., Figure 3A).

Basic Bayesian Observer Model

Following Girshick et al. (2011), the Basic Bayesian observer made maximum *a posteriori* estimates of motion direction from a distribution formed by multiplying prior and likelihood distributions and these estimates were assumed to be corrupted by motor noise and lapses. Each trial's displayed direction θ_{true} can produce sensory evidence θ_e with probabilities $p(\theta_e | \theta_{true})$ described by a von Mises distribution $\mathcal{V}(\theta_e; \theta_{true}, \kappa_e)$ centered on θ_{true} , with concentration κ_e . Each trial is modeled as a draw θ_{e_i} from the evidence distribution where i is the trial number, e.g., $\theta_{e_1} = 45^\circ$ (gray distribution, Figure 2A). The sensory likelihood for that trial, $p(\theta_{e_i} | \theta_{true})$ was a von Mises distribution with the same concentration parameter as the evidence distribution centered on the draw θ_{e_i} for that trial from the evidence distribution. We assumed that during each prior block, subjects formed a prior belief about the stimulus coherent motion directions and formalized the “learnt” prior $p(\theta_{true})_L$ as a von Mises distribution with mean $\theta_{\mu_{prior}}$ set at the experimental prior's true mean 225° , and concentration κ_{prior} fitted to the data (brown distributions, Figure 2A). The observer multiplicatively integrates the sensory likelihood and the learnt prior into a posterior $p(\theta_{true} | \theta_{e_i})$ (pink distributions) which is biased toward the strongest distribution (top: likelihood, bottom: prior). Thus, more generally, for a trial in which any sensory evidence θ_{e_i} was sampled from the evidence distribution, the posterior is given by:

$$p(\theta_{true} | \theta_{e_i}) = \frac{p(\theta_{e_i} | \theta_{true})p(\theta_{true})_L}{p(\theta_{e_i})} \quad (3)$$

and the posterior mode was taken to be the predicted percept:

$$\theta_p = \underset{\theta_{true}}{\operatorname{argmax}} (p(\theta_{true} | \theta_{e_i})) \quad (4)$$

From the equations above we can compute $p(\theta_p | \theta_{true})$ which is the distribution of percepts θ_p the model would generate given θ_{true} . Our model then assumes that the percepts θ_p are corrupted by motor noise and lapses to generate the distribution of actual estimates $\hat{\theta}$. We modeled lapses as random estimation with probability p_r and motor noise was described as a von Mises probability distribution with mean 0 and strength κ_m . The distribution of estimates for a given motion direction was given by:

$$p(\hat{\theta} | \theta_{true}) = \mathcal{V}(\hat{\theta}; 0, \kappa_m) * \left[(1 - p_r)p(\theta_p | \theta_{true}) + \frac{p_r}{360} \right] \quad (5)$$

where $*$ denotes the discrete circular convolution.

We separately fitted the model to subject trial estimates with a maximum likelihood procedure. We used the Nelder-mead algorithm to search for the model parameters that minimized the negative sum of the log probability of the model generating all subject estimates. The results hold when fitting the data with the covariance matrix adaptation evolution strategy algorithm (CMA-ES; Chalk et al., 2010; Hanks et al., 2011; Hansen and Kern, 2004) which can fit the data better than Nelder-Mead when the objective function space is noisy and ill-conditioned, that is when the negative log likelihood has many local maxima. This technique was useful for models described below where there are random fluctuations in the distribution of the stimulus dot directions from trial-to-trial. The model had nine fit parameters: three likelihood strengths, κ_e , that were allowed to vary with motion coherence, four prior strengths, κ_p , that varied with experimental prior strength, p_r and κ_m . The model predictions of estimate mean and variability are the circular mean and circular standard deviation (Figures 3A and 3B) of $p(\hat{\theta} | \theta_{true})$ (black curves, Figure 4).

Prior Learning Model

We modified the Basic Bayesian observer by adding a learning rate parameter, τ , that allowed us to fit how fast prior strengths were learnt within each prior block (κ_{prior} were replaced by $\kappa_{prior}(1 - e^{-t/\tau})$).

Switching Observer Model

The Switching observer was identical to the Basic Bayesian observer, except that rather than forming a posterior by multiplying prior and likelihood, the observer switched between prior mean and evidence based on their relative strength. Specifically, the probabilities p_{prior} and p_e that the percepts were respectively drawn from the prior mean and the sensory evidence distribution were:

$$p_{prior} = \frac{\kappa_{prior}}{\kappa_{prior} + \kappa_e} \text{ and } p_e = 1 - p_{prior}$$

On a small proportion of trials p_r the observer's percept was set to a random direction and the resulting distribution of percepts was convolved with motor noise formalized as a von Mises distribution of concentration κ_m . The probability of the estimates given the displayed motion direction was thus:

$$p(\hat{\theta} | \theta_{true}) = \mathcal{V}(\hat{\theta}; 0, \kappa_m) * \left[(1 - p_r) \left(p_e p(\theta_e | \theta_{true}) + p_{prior} \delta(\theta_p - \theta_{\mu_{prior}}) \right) + \frac{p_r}{360} \right] \quad (7)$$

where $\delta()$ is a delta function over percepts θ_p that peaks at the prior mean $\theta_{\mu_{prior}}$.

We also tested two variations of the Switching observer. One which we call the Switching prior sampling observer, in which the prior direction was not taken as a delta function as above, but instead was sampled from the prior distribution, thus allowing there

to be uncertainty about the exact prior. Another variation of the Switching observer, which switches but uses Bayesian computations, the Switching posterior observer, was also tested. The Switching posterior combines sensory likelihood and prior into a posterior and switches between the posterior mode and the prior mean.

Bayesian Models with Variations in Prior

We tested models based on the Basic Bayesian observer in which we relaxed the assumption that the prior was a von Mises distribution. We called this model the Bayesian observer with long-tailed prior. For this observer, any direction has some small probability of occurring. Priors were thus formalized as mixtures of a von Mises and a circular uniform distribution. The model had 11 fit parameters (three likelihood strengths κ_e , four prior strengths κ_{prior} , a unique mixture coefficient for all priors that controlled the weight of the tail relative to the von Mises, ω_{tail} , motor precision κ_m , and lapse rate p_r).

Bayesian Models with Variations in Likelihood

We tested several plausible variations in the likelihood representation that could allow the Basic Bayesian observer to produce bimodal distributions. We modeled the stochasticity in the stimulus to obtain a more accurate description of the stimulus, which in effect relaxes the simplifying assumption that evidence distributions were von Mises-shaped and peaked at stimulus displayed motion direction. Instead we used the distribution of stimulus individual dot directions to calculate an evidence distribution. We called the model the Bayesian observer with long-tailed likelihood as the derived evidence distributions and sensory likelihood displayed long tails. The motion coherence of the stimulus was controlled by moving a percentage of the dots (6, 12 or 24%) in the same directions while other dots were moved in random directions. At each trial, the stimulus displayed on the screen was composed of N dots ($N = 328$), that moved in the directions θ_i where i is an integer from 1 to N. The directions θ_i were randomly sampled from the mixture of a uniform (the distribution of the directions of the incoherent dots) and a delta function (the proportion of dots that moved in the stimulus coherent direction) such that the probability of the dot direction θ_i given the coherent motion direction θ_{true} with the coherence c is:

$$p(\theta_i | \theta_{true}, c) = c\delta(\theta_i - \theta_{true}) + \frac{(1 - c)}{360} \quad (8)$$

where $\delta()$ is a delta function that peaks at the stimulus coherent motion direction, θ_{true} . Because of sensory noise, each single dot direction θ_i is considered to provide sensory evidence distribution $p(\theta_{e_i} | \theta_i)$ which can be described by a von Mises distribution $\mathcal{V}(\theta_{e_i}; \theta_i, \kappa)$ that peaks at the dot direction θ_i and has strength κ . Thus the evidence distribution for each dot is computed as:

$$\begin{aligned} p(\theta_{e_i} | \theta_{true}, c) &= \sum_{\theta_i} p(\theta_{e_i} | \theta_i) p(\theta_i | \theta_{true}, c) \\ &= c\mathcal{V}(\theta_{e_i}; \theta_{true}, \kappa) + \frac{1 - c}{360} \end{aligned} \quad (9)$$

Given the N sensory evidence that were drawn, the sensory likelihood under the hypothesized coherent motion directions θ_{true} of the stimulus is inferred by integrating the probability of the sensory evidence over the N dots:

$$\begin{aligned} p(\theta_e | \theta_{true}, c) &= \prod_{i=1}^{N=328} p(\theta_{e_i} | \theta_i, c) \\ &= \prod_{i=1}^{N=328} \left(c\mathcal{V}(\theta_{e_i}; \theta_{true}, \kappa) + \frac{1 - c}{360} \right) \end{aligned} \quad (10)$$

and the sensory likelihood inherits the long tails of the distribution from which the dot directions were sampled (Equation 9). The sensory likelihood and the prior were multiplicatively combined into a posterior, which is bimodal due to the sensory likelihood's tail, and the percepts were produced by reading out the posterior mean. The percepts were then converted into estimates that can deviate from the percepts due to motor noise and lapses. The model had seven fit parameters (the strength of evidence for each dot κ , four prior strengths κ_{prior} , motor precision κ_m and lapse rate p_r).

The Bayesian observer with long-tailed prior and likelihood, multiplicatively combined long-tailed integrated likelihood with long-tailed prior into a bimodal posterior, and readout posterior mean as percept. The model has eight fit parameters (the strength of evidence for each dot κ , four prior strengths κ_{prior} , a unique mixture coefficient that controlled prior tail weight relative to von Mises, ω_{tail} , motor precision κ_m , and lapse rate p_r).

We also constructed a model variation in which sensory likelihood was derived from motion-energy filters (Adelson and Bergen, 1985). We constructed 24 motion-energy filters with direction peak selectivity ranging from 0° to 345° in steps of 15° . Filters were tuned for speed near the stimulus speed of $2.8^\circ/\text{sec}$ with peak temporal frequency selectivity of 1 and 1.4 cycles/sec and peak spatial frequency selectivity of 0.46 and 0.5 cycle/ $^\circ$ producing peak speed selectivity of 2; 2.17; 2.8; $3.04^\circ/\text{sec}$ and filter mean responses were averaged over time. Four thousands random dot stimuli (all in the same direction, different directions were inferred based on symmetry) for each of the three coherences used in the experiment (6, 12 and 24%) were simulated. Motion-energy filter

responses to these stimuli were converted into direction evidence with population vector decoding which weights each filter's preferred motion direction by its response magnitude (Georgopoulos et al., 1993). Sensory likelihoods were then derived from the distribution of different motion direction evidence across different random dot stimuli.

Sampling Bayesian Observer Model

The sampling Bayesian observer represents and integrates the sensory likelihood and the prior in the same way as the Basic Bayesian observer, but the percepts are not created by reading out the posterior mode but by randomly sampling the posterior (Berniker et al., 2010; Körding and Wolpert, 2004; Moreno-Bote et al., 2011; Tassinari et al., 2006; Vilares et al., 2012). This model has the same nine fit parameters as the Basic Bayesian model.

Bayesian Models with Least-Square Readout

Model variations of Basic Bayesian and other observers were implemented with a least-squares readout, which consists in reading out the circular mean of the posterior instead of its mode.

Bayesian Models with Learnt and Cardinal Priors

This model was identical to the Basic Bayesian observer, except that it also used a prior over cardinal directions $p(\theta_{true})_c$. The additional prior was a mixture of four equally weighted von Mises which peaked at 0, 90, 180 and 270°, with the same concentration parameter κ_c which was a fit parameter (Chalk et al., 2010; Girshick et al., 2011). Thus, for a trial i in which any sensory evidence θ_{ei} was sampled from the evidence distribution, the two priors and the likelihood were integrated as follows:

$$p(\theta_{true} | \theta_{ei}) = \frac{p(\theta_{ei} | \theta_{true})p(\theta_{true})_L p(\theta_{true})_c}{p(\theta_{ei})} \quad (11)$$

Model Fit Convergence

We used the following nine sets of initial parameters to maximize the likelihood of convergence of parameter estimates for all observers. We manually adjusted parameters to match each model estimate distributions with the actual estimates distributions for each subject. We then also chose all eight combinations of prior and likelihood widths that were eight times stronger or weaker than the manually adjusted parameters. We also chose another set in which prior and sensory likelihood widths were all set to a single value: the average of the manually fit prior and likelihood widths. We started fitting with each of these ten initial starting values and used the Nelder-Mead algorithm to maximize the log-likelihood of each observer model producing the measured data. The fits obtained for all initial parameters converged when changing the fit parameters during the Nelder-Mead search did not result in change in the log likelihood and fit parameters above the very strict function and parameter tolerances of 1e-4. Using a stricter tolerance did not substantially change the fit, while more relaxed tolerance worsened the fit. We also fitted the data with a second non-linear optimization algorithm: the covariance matrix adaptation search algorithm (CMAES) which produced similar results.

Model Comparison

Model comparisons were made using Akaike information criterion (AIC) which accounts for overfitting by penalizing models with greater number of parameters (Akaike, 1974; Bishop, 2006). AIC was calculated as follows:

$$AIC = 2 \left[n - \log \left(p(\hat{\theta} | \text{model}) \right) \right] \quad (12)$$

where n is the number of fit parameters of the model and $p(\hat{\theta} | \text{model})$ is the likelihood of the trial data given the model best fit parameters. We calculated Δ_i which is the AIC difference between any two models for all subjects. The reported AIC difference was always calculated by subtracting the Switching or Switching sampling prior observer's AICs (such that positive differences indicate more evidence for the Switching observer), except when the Basic Bayesian observer's fit was compared to the maximum likelihood observer's fit. In that case, the Basic Bayesian observer's AIC was subtracted from the maximum likelihood AIC. A difference of 2 is equivalent to adding another parameter to the model.

QUANTIFICATION AND STATISTICAL ANALYSIS

Two-tailed one sample Wilcoxon tests were used to compare the median over subjects, $n = 12$, of the prior strengths parameters fitted by the Switching observer to subject's data to veridical strength. A non-parametric test was chosen because fitted prior strength values were not all normally distributed (two-tailed Lilliefors' test with a statistical threshold $\alpha = 0.05$).

Two-tailed Pearson correlation was used to compare the slopes of each subject's estimated versus displayed directions for the twelve prior x coherence conditions in the first hundred versus late hundred trials of prior blocks. The slopes were calculated for

each task condition by pooling the first or last hundred trials of each prior block over all blocks and by fitting estimated versus displayed directions with a line. Two-tailed Pearson correlation coefficient r and p -value quantified for each subject how well the twelve slopes calculated during early trials matched late trials slopes.

Confidence intervals were calculated as the $\Delta_l \pm 1.96s\sqrt{n}$ where s is the standard deviation of Δ_l calculated over subjects and n is the number of subjects.

DATA AND SOFTWARE AVAILABILITY

The data and the analysis scripts are available on <https://github.com/steevelaquittaine/projInference> and <https://doi.org/10.17632/nxkvtrj9ps.1>.