Article

Neuron

Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality

Highlights

- Human decisions based on multiple ambiguous cues are typically suboptimal
- Sensory noise and response selection cannot explain choice suboptimality alone
- Imperfections in mental inference cause a dominant fraction of choice suboptimality
- Most of choice suboptimality arises from imprecise rather than biased computations

Authors

Jan Drugowitsch, Valentin Wyart, Anne-Dominique Devauchelle, Etienne Koechlin

Correspondence

jan_drugowitsch@hms.harvard.edu (J.D.), valentin.wyart@ens.fr (V.W.)

In Brief

Decisions made under uncertainty show a suboptimal variability whose origin is usually ascribed to the peripheries of the decision process. Drugowitsch et al. show that computational imprecisions in the decision process itself account for a dominant fraction of choice suboptimality.

Drugowitsch et al., 2016, Neuron *92*, 1398–1411
December 21, 2016
http://dx.doi.org/10.1016/j.neuron.2016.11.005

CellPress

CellPress

Neuron
Article

Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality

Jan Drugowitsch,[1,2,3,4,5,*] Valentin Wyart,[1,4,*] Anne-Dominique Devauchelle,[1] and Etienne Koechlin[1]
[1]Laboratoire de Neurosciences Cognitives, Inserm unit 960, Département d'Études Cognitives, École Normale Supérieure, PSL Research University, 75005 Paris, France
[2]Département des Neurosciences Fondamentales, Université de Genève, CH-1211 Geneva, Switzerland
[3]Department of Neurobiology, Harvard Medical School, Boston, MA 24615, USA
[4]Co-first author
[5]Lead Contact
*Correspondence: jan_drugowitsch@hms.harvard.edu (J.D.), valentin.wyart@ens.fr (V.W.)
http://dx.doi.org/10.1016/j.neuron.2016.11.005

SUMMARY

Making decisions in uncertain environments often requires combining multiple pieces of ambiguous information from external cues. In such conditions, human choices resemble optimal Bayesian inference, but typically show a large suboptimal variability whose origin remains poorly understood. In particular, this choice suboptimality might arise from imperfections in mental inference rather than in peripheral stages, such as sensory processing and response selection. Here, we dissociate these three sources of suboptimality in human choices based on combining multiple ambiguous cues. Using a novel quantitative approach for identifying the origin and structure of choice variability, we show that imperfections in inference alone cause a dominant fraction of suboptimal choices. Furthermore, two-thirds of this suboptimality appear to derive from the limited precision of neural computations implementing inference rather than from systematic deviations from Bayes-optimal inference. These findings set an upper bound on the accuracy and ultimate predictability of human choices in uncertain environments.

INTRODUCTION

Humans show a remarkable ability to make efficient decisions in a wide range of situations, from perceptual decisions based on sensory cues to social and economic decisions based on preferences and reward. Identifying the computations underlying decisions is essential for understanding human cognition and its neural substrates. In many situations, making decisions requires combining multiple pieces of ambiguous or even conflicting information from external cues (Gold and Shadlen, 2007; Shadlen and Kiani, 2013), a task that is optimally performed by Bayesian probabilistic inference (Knill and Pouget, 2004; Ma et al., 2006; Yang and Shadlen, 2007). This inference process is at the heart of many forms of decisions, ranging from the categorization of an ambiguous sensory stimulus (e.g., Roitman and Shadlen, 2002; Weiss et al., 2002; Ernst and Banks, 2002) to the reinforcement of a reward-yielding action (e.g., Behrens et al., 2007; Daw et al., 2005, 2011), i.e., any task that requires inferring a "latent state" from noisy or ambiguous cues. Although human perceptual and reward-guided choices have been shown to resemble Bayes optimality in such conditions (Navalpakkam et al., 2010; Ma et al., 2011), they typically exhibit a large variability beyond what can be explained by the provided evidence and are thus inherently suboptimal (Beck et al., 2012).

Most empirical studies have attributed such choice suboptimality to imperfections in peripheral stages of decision making at the input and output of the inference process. In perceptual categorization tasks, errors are typically attributed to task-independent noise in sensory processing preceding the task-dependent inference process (Osborne et al., 2005; Brunton et al., 2013; Kaufman and Churchland, 2013). In reward-guided learning and abstract reasoning tasks, by contrast, suboptimal choices are often ascribed to stochasticity in the response selection following the decision process (Sutton and Barto, 1998; Daw et al., 2006; Griffiths and Tenenbaum, 2006; Vul and Pashler, 2008; Vul et al., 2009). In these tasks, however, it is theoretically impossible to distinguish between these peripheral imperfections and errors arising from the inference process itself (Tsotsos, 2001; Whiteley and Sahani, 2012; Beck et al., 2012; Renart and Machens, 2014; Dayan, 2014), such as systematic deviations or *biases* from Bayes-optimal inference (Beck et al., 2012) or unstructured variability or *noise* in underlying neural computations (Renart and Machens, 2014). Indeed, perceptual categorization tasks make no distinction between the relevant feature(s) of sensory cues and the information they provide for the decision. For example, in the click-count discrimination task introduced by Brunton et al. (2013), the noisy perception of clicks (i.e., imperfect sensory processing) is behaviorally indistinguishable from the noisy accumulation of noiseless click percepts (i.e., imperfect inference). Similarly, in reward-guided learning and abstract reasoning tasks, the number of presented cues remains constant across decisions, making it impossible to

distinguish between variability in inference and response selection. Consequently, the origin and structure of human choice suboptimality remain largely unclear.

To address this issue, we devised a computational modeling approach and an experimental protocol derived from a canonical evidence accumulation task, such that imperfections in sensory processing, inference, and response selection have separable statistical signatures in observed choices. We then decomposed the statistical signature of inferential imperfections in terms of a "bias-variance" tradeoff between (1) systematic deviations from Bayes-optimal inference and (2) unstructured variability stemming from the limited computational precision of mental inference. This experimental framework revealed that in contrast to current views, the suboptimality of human choices made under ambiguity arises dominantly from imperfections in inference, of which two-thirds could not be accounted for by *any* systematic deviation from Bayes optimality and thus resulted from the limited precision of neural computations implementing inference. In absolute terms, this limited precision of mental inference caused a substantial loss of 30% of the theoretical information provided by each evidence sample, which is not accounted for by existing theoretical models of decision making.

RESULTS

Probabilistic Inference Task and Human Choice Behavior

To identify the origin of human choice suboptimality, we created a probabilistic cue combination task derived from the well-known "weather prediction" task (Knowlton et al., 1996; Poldrack et al., 2001; Gluck et al., 2002; Yang and Shadlen, 2007), in which key properties of the underlying inference process (i.e., extracting and accumulating the evidence provided by successive cues) could be manipulated independently of sensory processing and response selection: (1) the *length* of inference (i.e., the number of cues that need to be combined) and (2) the *dimensionality* of inference (i.e., the number of alternatives to choose from). This experimental protocol afforded the comparison of human choice behavior to Bayes-optimal choices and the quantification of the extent to which inferential imperfections alone are responsible for the observed choice suboptimality.

In every trial, participants observed a sequence of 2–16 oriented patterns or *cards* and were then asked to indicate which category or *deck* among two or three possible ones they judged the sequence to have been drawn from (Figure 1A; see Experimental Procedures). As expected, participants' choices reflected the combination of the pieces of information conveyed by successive cards presented within each trial; indeed, choice accuracy (Figure 2A) increased gradually with the number of presented cards in both two- and three-category conditions (repeated-measures ANOVA, both $F_{7,147} > 45.2$, $p < 0.001$). Although categories were randomly selected across successive trials, human choices in both conditions were slightly biased toward the category drawn on the previous trial and shown as feedback (Figure 2B; t test against zero, both $t_{21} > 3.5$, $p < 0.002$), but not toward the previous choice (both $t_{21} < 0.5$, $p > 0.2$) or categories drawn on earlier trials (all $t_{21} < 1.8$, $p > 0.05$). Furthermore, choice accuracy remained stable over the course of the experiment (Figures S1A and S1B; first versus second half: both $t_{21} < 0.5$, $p > 0.5$).

Estimating the Separate Contributions to Choice Suboptimality

In this task, the optimal decision-making model integrates ambiguous information across successive cards through Bayesian inference by accumulating evidence in favor of each deck and then chooses the deck associated with the largest total evidence (Figure 1B). As for participants, the model predicts choice accuracy to increase with the number of drawn cards. However, the model choice accuracy systematically and substantially exceeds human performance in both conditions (Figure 2A). This suboptimality of human choices may stem from imperfections in sensory processing, inference, or response selection, each of which affects the decision-making process in different ways (Figure 1B). Furthermore, in our experimental protocol, these three sources of choice suboptimality have distinct statistical signatures on the variability of human choices compared to Bayes-optimal ones (Figure 1C; see Experimental Procedures).

We first observed that the variability of human choices increased quasi-linearly with the number of presented cards in both the two- and three-category conditions (Figure 3A, dots; both $F_{7,147} > 15.2$, $p < 0.001$). Assuming that this variability derived from a unique source of imperfections (sensory, inferential, or selection based), we then parameterized this source in an otherwise optimal decision-making model and fitted the predicted variance structure to the observed choice variability (Figure 3A, lines; see Experimental Procedures). In both conditions, fitting the sensory source model led to postulated orientation discrimination thresholds of an order of magnitude larger than those tabulated in the literature for comparable stimuli (Burbeck and Regan, 1983; Webster et al., 1990; Burr and Wijesundra, 1991) (two-category condition: 19.0 ± 0.8 degrees [deg.]; three-category condition: 23.4 ± 0.9 deg., mean ± SEM). Even more problematic, the fitted thresholds differed significantly between the two- and three-category conditions, despite involving the exact same stimuli (Figure 3A, left; paired t test, $t_{21} = 4.8$, $p < 0.001$). Besides, the linear increase of human choice variability with the number of presented cards rules out imperfections in response selection as a possible source (Figure 3A, right).

Bayesian model comparison across the different models consistently revealed that the inferential source model (Figure 3A, middle) explained human choices decisively better than alternative accounts in both conditions (Figure 3B; both Bayes factors $>10^{12}$, both exceedance $p > 0.99$). Here, the inferential source model postulates imperfections in the interpretation of the evidence provided by each card in favor of each category and/or in the accumulation of evidence across cards (see Figure 1B and Experimental Procedures), such that we refer to both under the umbrella term of "inferential" imperfections. We confirmed the validity of our model comparison among the three hypothesized sources of choice variability through a validation procedure using synthetic choice data from each model (Figures 3C and S2; see Supplemental Experimental Procedures). Additional analyses revealed that the magnitude of the inferential imperfections measured separately

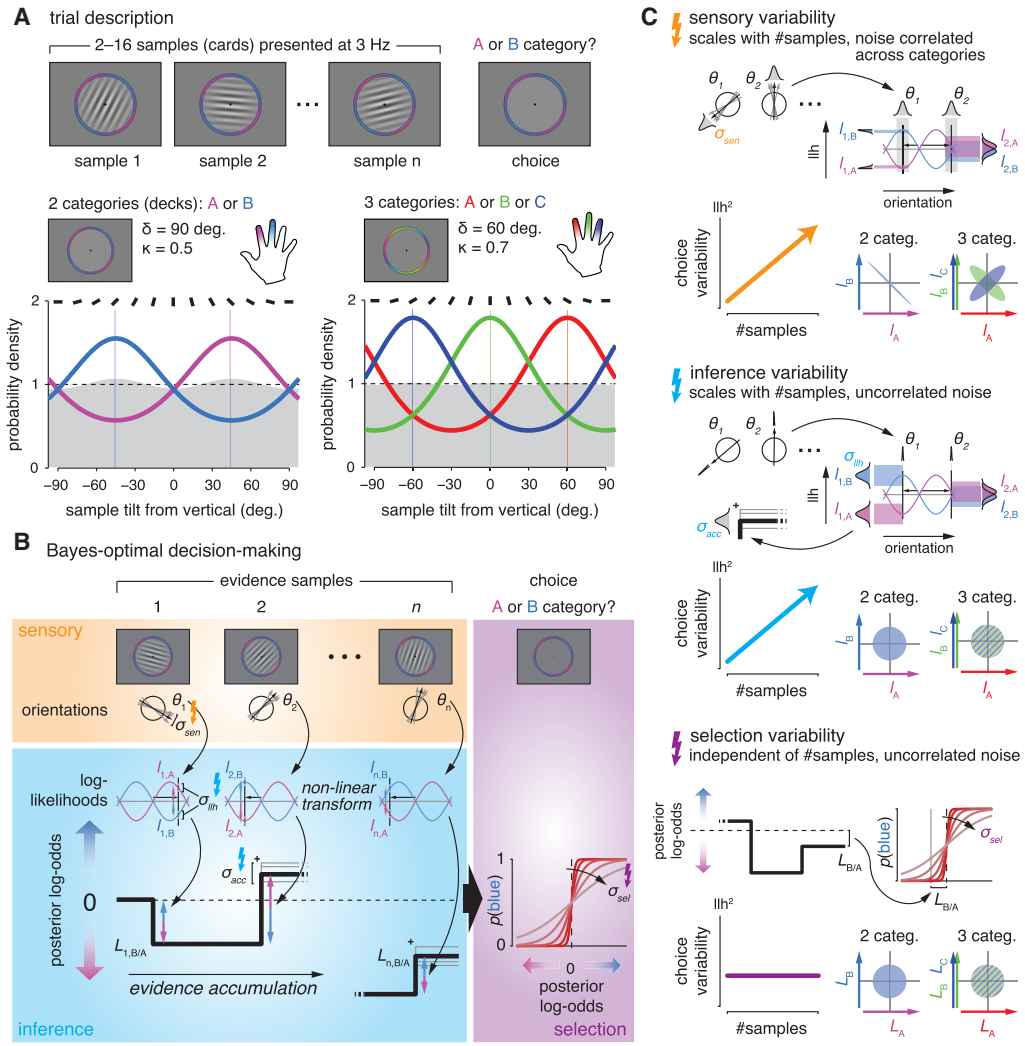

Figure 1. Experimental Protocol and Theoretical Model of Choice

(A) Trial description, generative distributions, and experimental conditions. Each trial consisted of a sequence of 2–16 cues (cards) presented at approximately 3 Hz, drawn from a generative probability distribution (deck) centered on one among two or three cardinal orientations. At sequence offset, participants were prompted to indicate the deck from which they believed the cards were drawn from by pressing a button. The drawn category was indicated following each choice by a transient color change of the fixation point. In the two-category condition (bottom left), sequences of cards were drawn from one of two color-coded circular Gaussian distributions of low concentration $\kappa = 0.5$ centered on $+45^\circ$ (pink) and -45° (cyan), whereas the three-category condition (bottom right) consisted of three distributions of low concentration $\kappa = 0.7$ centered on $+60^\circ$ (red), 0° (green), and -60° (blue).

(B) Optimal decision making in this task is achieved by accumulating evidence for each category conveyed by each orientation sample θ_j (log likelihoods, $l_{j,A/B}$), then accumulating evidence $L_{j,A/B}$ for each category by summing log likelihoods and finally contrasting the accumulated evidence across categories (log posterior) to choose the most likely category given the observed orientation samples. In the three-category condition, such evidence accumulation is independent across categories. In the two-category condition depicted here, it is sufficient to accumulate log-likelihood ratios and decide based on the sign of the resulting log posterior. Sensory variability is task independent and modeled by adding Gaussian noise (σ_{sen}) to the orientation percepts. Inferential variability is task dependent and modeled by adding independent Gaussian noise to the log likelihoods, with respect to each category (σ_{llh}), and/or to each step of the evidence accumulation (σ_{acc}). Variability in response selection corresponds to stochastic action selection by adding Gaussian noise to the final log posterior (σ_{sel}).

(C) Distinct signatures of sensory, inference, and selection variability on the suboptimality of resulting choices. Sensory variability (top panel) alters the perceived orientation of each sample, which causes orientation-dependent (e.g., θ_1 versus θ_2) and correlated variability in the resulting log likelihoods l_A, l_B and l_C, which, due to the non-linear mapping between orientation and log likelihood, becomes non-Gaussian. Because each sample is affected independently, the resulting choice variability scales with the number of presented samples. Inference variability (middle panel) alters the interpretation and accumulation of the evidence in favor of each category. Similar to sensory variability, the resulting choice variability scales with the number of presented samples. However, it affects individual log likelihoods and their accumulation independently (l_A, l_B, and l_C uncorrelated and independent of θ_j), leading to different trial-by-trial choice predictions that can be distinguished through model comparison. Selection variability (bottom panel) alters the "readout" of the accumulated evidence during response selection, leading to uncorrelated variability in the accumulated evidence across categories (L_A, L_B, and L_C uncorrelated). In contrast to the two other sources of variability, the resulting "probability-matching"-like choice variability is independent of the number of presented samples.

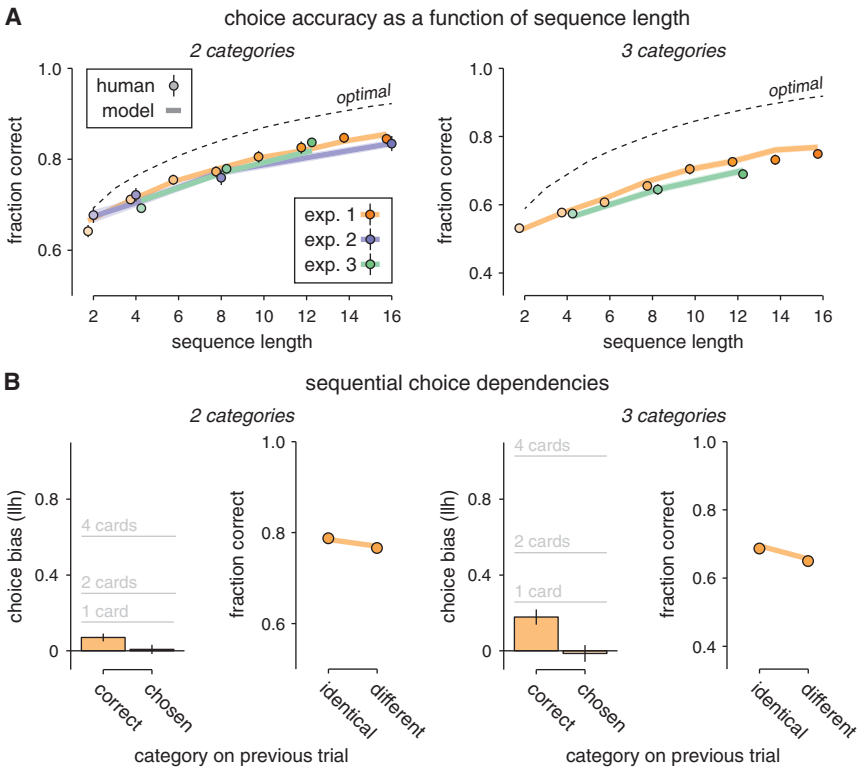

Figure 2. Human Choice Behavior

(A) Fraction correct with respect to sequence length in the two-category (left) and three-category (right) conditions in the three experiments (dots, mean ± SEM, measured across participants unless noted otherwise). Lines show the predicted fraction correct for the best-fitting model, assuming variability in probabilistic inference. The fraction correct for the normative Bayesian decision maker, corresponding to the maximally attainable (optimal) performance, is shown as dashed lines.

(B) Sequential choice dependencies in the two-category (left) and three-category (right) conditions of the first experiment, (error bars, mean ± SEM) measured in terms of choice bias estimates from a logistic regression of choice in trial n against the correct and chosen category in trial $n-1$ (left panel). Participants are slightly biased toward the previously correct category but indifferent to their previous choice. For comparison, gray lines show the average ideal (noise-free) choice information provided by one, two, and four cards. The fraction correct measured (dots) and predicted (colored lines) for trials in which the drawn category is identical or different from the one drawn on the previous trial shows a subtle effect of the choice bias in the direction of the previously correct category (right panel).

for the first and second halves of each experiment was similar (t_{21} = 1.0, p > 0.2) and correlated significantly across participants (r = 0.67, df = 20, p < 0.001), thereby reflecting a stable and robust feature of the participants' decision performance (Figure S1).

Quantifying Sensory Noise Independently of Probabilistic Inference

We have so far distinguished sensory and inferential imperfections on the basis of their statistical signatures on choice variability. In a second experiment, performed by a new set of participants (Figure 4A), we segregated between these two sources of choice suboptimality more directly by explicitly measuring an upper bound on participants' orientation discrimination thresholds in our protocol. The experiment interleaved two-category *probabilistic inference* task trials used in the first experiment ("accumulation" trials) with *orientation categorization* task trials performed only on the last card of the sequence ("last-card" trials; see Experimental Procedures). As expected, participants' choices in last-card trials were based solely on the orientation of the last pattern (logistic regression of choice against last card tilt from between-category boundary, t_{16} = 6.4, right-tailed p < 0.001) and not on the evidence provided by preceding cards (logistic regression of choice against Bayes-optimal accumulated evidence: t_{16} < 0, right-tailed p > 0.5).

To reliably estimate an upper bound of the orientation discrimination threshold in last-card trials, we deliberatively chose last card orientations close to the category boundaries in these trials (1–8 deg. from the horizontal and vertical axes; Figure 4B). Fitting the optimal decision-making model with only sensory imperfections to participants' choices in last-card trials, we obtained a mean discrimination threshold of 2.4 ± 0.2 deg. (± SEM), matching values measured for similar stimuli (Burbeck and Regan, 1983; Webster et al., 1990; Burr and Wijesundra, 1991). However, these estimated thresholds predicted choice accuracies in accumulation trials that dramatically exceeded those of participants (Figure 4D). If we reversed the procedure and estimated orientation discrimination thresholds using sensory imperfections now fitted to participants' choices in accumulation trials, we recovered (as in the first experiment) implausible thresholds of a higher order of magnitude (21.0 ± 0.9 deg., paired t test, t_{16} = 21.1, p < 0.001), which were uncorrelated to threshold estimates from last-card trials (Figure 4C; linear correlation across participants, r^2 = 0.02, df = 15, p > 0.5).

Bayesian model comparison further showed that a model featuring only sensory imperfections of identical magnitudes in accumulation and last-card trials explained participants' behavior decisively worse than a model including additional inferential imperfections in accumulation trials (Bayes factor > 10^{48}, exceedance p > 0.99). Thus, sensory imperfections alone could not successfully explain the suboptimality of human choices observed in accumulation trials. Finally, comparing models with imperfections in sensory processing, inference, and response selection in those trials led to the same qualitative and quantitative conclusion as in the first experiment; inferential imperfections emerged as the prominent source of human choice suboptimality in our task (Figures 4E and 4F).

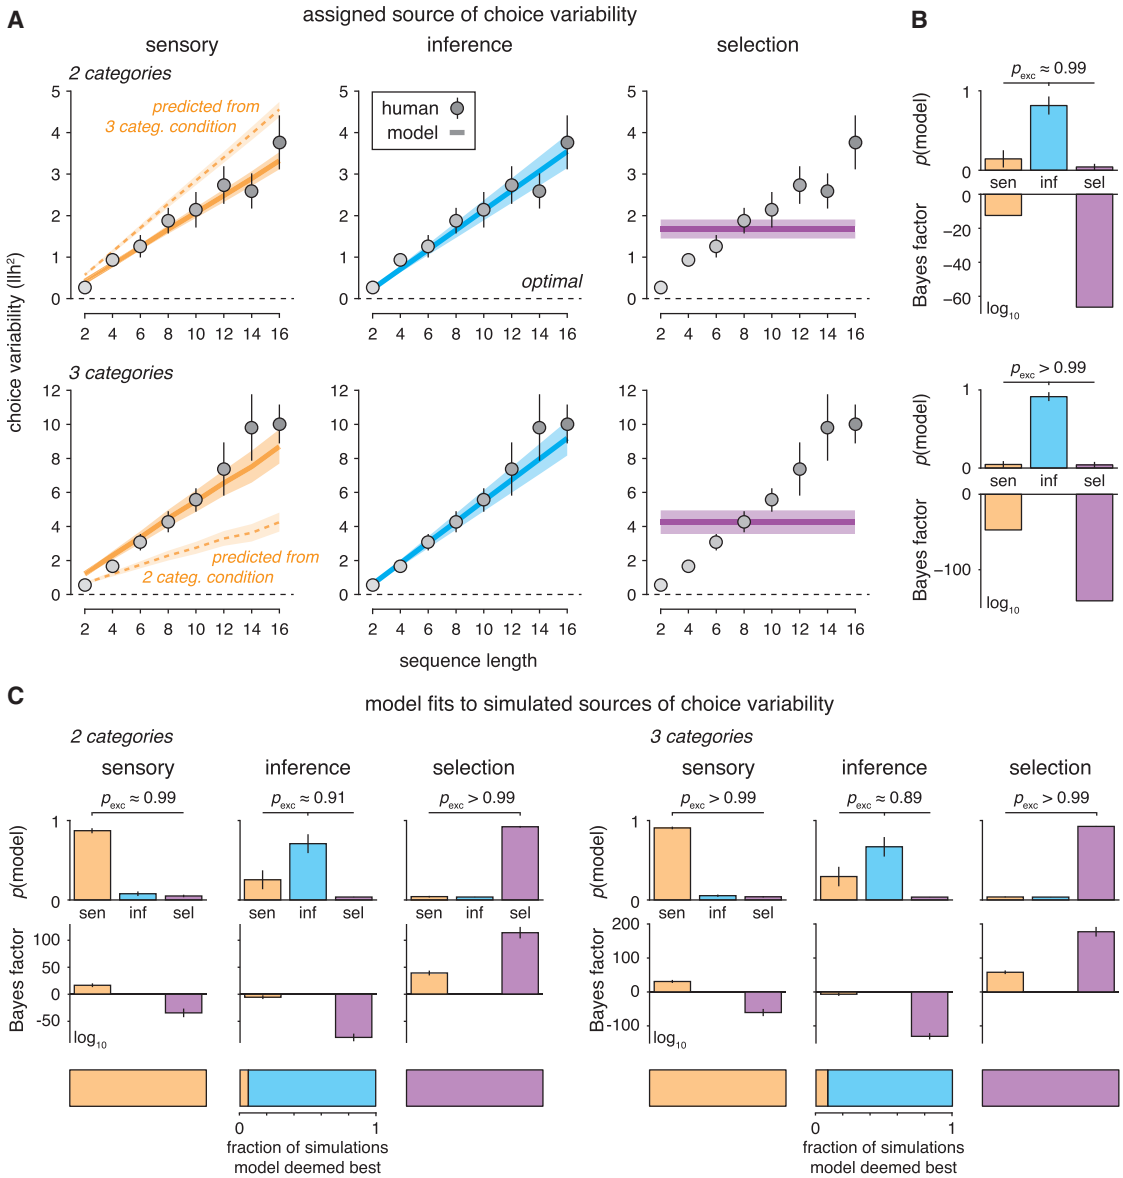

Figure 3. Inference-Driven Source of Human Choice Suboptimality

(A) Estimated choice suboptimality (dots, mean ± SEM) grows with sequence length in the two-category (top row) and three-category (bottom row) conditions. This suboptimality is measured as the squared variability of human choices around choices predicted by the Bayes-optimal decision maker (dashed lines). It has units of squared log likelihoods, where the log likelihood measures the amount of information that each sample provides about the generative deck. Model predictions are shown as lines (shaded error bars: mean ± SEM). For the model assuming sensory variability, predictions from the best-fitting model in each condition are shown in the other condition as dashed lines.

(B) Bayesian model comparison between candidate sources of choice suboptimality in the two-category (top row) and three-category (bottom row) conditions: sensory variability (left bars), inference-driven variability (middle bars), and selection-driven variability (right bars). For each condition, the top panel depicts the results of random-effect comparisons (in terms of the probability of sampling each model, mean ± SD; p_{exc}, exceedance probability), the bottom panel depicts the results of fixed-effects comparisons (in terms of the Bayes factor).

(C) Model fits to simulated sources of choice suboptimality in the two-category (left) and three-category (right) conditions: simulating sensory variability (left column), inference-driven variability (middle column), and selection-driven variability (right column). Same conventions as in (B). The bottom row shows the fraction of simulations, in which each model is deemed best for each simulated source of choice suboptimality.

Considering Multiple Simultaneous Sources of Choice Suboptimality

We next relaxed the assumption that choice suboptimality stems from a unique source of imperfections at the sensory processing, inference, or response selection stage by considering a model that features imperfections at all three stages simultaneously. We fitted this combined model to participants' choices in the two-category condition of both the first and second experiment

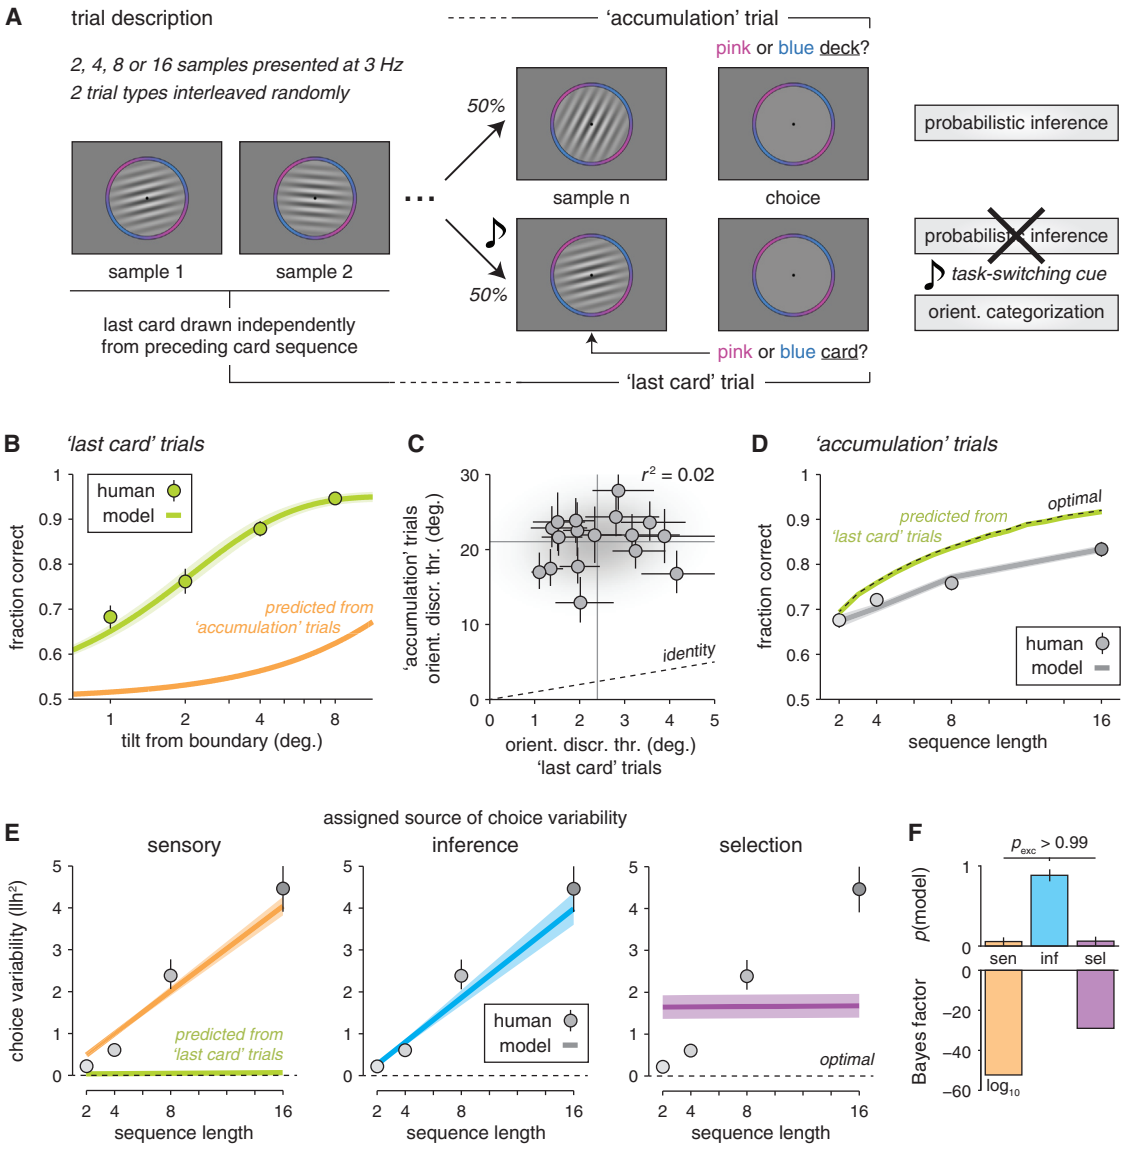

Figure 4. Ruling Out Sensory Noise as a Significant Source of Human Choice Suboptimality

(A) Trial description for experiment 2. As in experiment 1, each trial consists of a sequence of 2–16 cues presented at approximately 3 Hz and drawn from a wide circular Gaussian distribution centered on +45° (pink category) or −45° (blue category). Unpredictably, in half of the trials, the last cue was presented simultaneously with a tone, which prompted participants to respond not to the judged category of the sequence ("accumulation" trials, top row) but to the category of the last card ("last card" trials, bottom row), which, on these trials, was drawn independently from the category of the preceding cue sequence.

(B) Fraction correct in the "last card" trials with respect to the tilt of the last card from the closest category boundary (dots, mean ± SEM) fitted by a model assuming sensory variability, with an orientation discrimination threshold of 2.4 ± 0.2 deg. (lime line), and by a model assuming sensory variability estimated from "accumulation" trials (orange line).

(C) Comparison between orientation discrimination thresholds estimated using a model assuming sensory variability in the "last card" trials (horizontal axis) and "accumulation" trials (vertical axis). Dots show estimates from individual participants (error bars: full widths at half maximum likelihood); the dashed line is the identity line.

(D) Fraction correct in the "accumulation" trials, with respect to sequence length (dots, mean ± SEM). The gray line shows the predicted fraction correct choices for the best-fitting model, assuming variability in probabilistic inference, whereas the lime line shows the predicted fraction correct choices for the model, assuming sensory variability and fitted to "last card" trials.

(E) Measured and predicted choice suboptimality estimates (dots, mean ± SEM) in the "accumulation" trials. Same conventions and results as in Figure 3A. The lime line shows the predicted choice suboptimality from a model assuming sensory variability and fitted to "last card" trials.

(F) Bayesian model comparison between candidate sources of choice suboptimality. Same conventions as in Figure 3B.

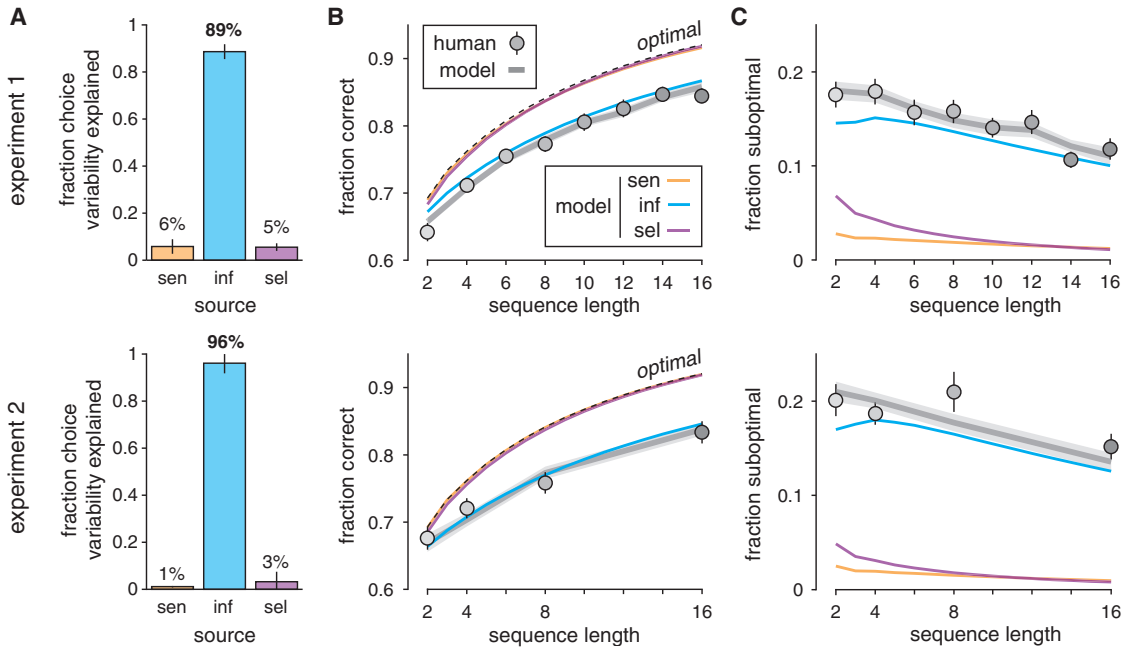

Figure 5. Decomposition of Human Choice Suboptimality

(A) Decomposition of human choice suboptimality into sensory (left bars), inference-driven (middle bars), and selection-driven (right bars) sources of variability in the two-category condition of experiment 1 (top row) and experiment 2 (bottom row). 89% (experiment 1) and 96% (experiment 2) of the measured choice suboptimality were assigned uniquely to variability in probabilistic inference. Error bars show SEM.

(B) Predicted impacts of sensory (orange lines), inference-driven (blue lines), and selection-driven (purple lines) variability on fraction correct (dots, mean ± SEM) in the two-category condition of experiment 1 (top row) and experiment 2 (bottom row).

(C) Predicted impacts of sensory, inference-driven, and selection-driven variability on the fraction of suboptimal choices (i.e., inconsistent with the normative Bayesian decision maker) in the two-category condition of experiment 1 (top row) and experiment 2 (bottom row). Same conventions as in (B). Note that the fraction of suboptimal choices decreases with the number of shown cards for all three sources of choice variability because the signal-to-noise ratio of the decision variable grows with the number n of shown cards (the signal grows with n, whereas the noise SD grows with $\sqrt{n}$ for sensory and inference-driven variability or remains constant for selection-driven variability).

(excluding last-card trials) to compute the relative contributions of these three sources of imperfections to human choice suboptimality (Figure 5A; see Experimental Procedures). In the first and second experiment, the fitted model assigned 89% and 96%, respectively, of the observed choice variability uniquely to inferential imperfections (experiment 1: sensory: 5.8% ± 3.1%, inference: 88.7% ± 3.2%, selection: 5.5% ± 1.6%; experiment 2: sensory: 1.1% ± 0.3%, inference: 96.1% ± 4.3%, selection: 3.2% ± 4.2%, mean ± SEM; see Figure S2B). Validating our model fits, we found that the fitted magnitude of sensory imperfections in this combined model corresponded to orientation discrimination thresholds (experiment 1: 2.6 ± 1.0 deg.; experiment 2: 2.2 ± 0.3 deg., mean ± SEM), which were virtually identical to the ones directly computed from last-card trials in the second experiment (2.4 ± 0.2 deg.).

To quantify the behavioral impact of each of these imperfections, we simulated the choice suboptimality due uniquely to each of these imperfections by setting, in turn, each of them at their best-fitting value while holding the two other imperfections at zero (Figure 5B). In the first experiment, we found that imperfections in sensory processing, inference, and response selection caused 2%, 13%, and 3%, respectively, of suboptimal choices (2%, 16%, and 3%, respectively, in the second experiment) across tested sequence lengths (Figures 5B and 5C). These results thus confirm that inferential imperfections are the main source of choice suboptimality in our task.

Partitioning Choice Suboptimality in Terms of a Bias-Variance Tradeoff

Next, we investigated the structure of observed imperfections in mental inference. Inferential imperfections could correspond to systematic deviations from Bayes-optimal inference, often coined as biases in psychology. Such *deterministic biases* may arise for various reasons, e.g., because the agent performs Bayes-optimal inference based on wrong assumptions about the generative structure of the task (Beck et al., 2012; Dayan, 2014) or because the agent uses heuristics that deviate from Bayes-optimal inference (Tsotsos, 2001; Whiteley and Sahani, 2012). Alternatively, these inferential imperfections might correspond to the limited precision of neural mechanisms implementing inference (Renart and Machens, 2014), thereby generating *unstructured fluctuations* in resulting choices. These two causes have opposing choice effects in response to identical sequences of cues: deterministic biases would make the agent consistently repeat the same choice, whereas unstructured fluctuations would cause inconsistent and unrelated choices in response to identical sequences.

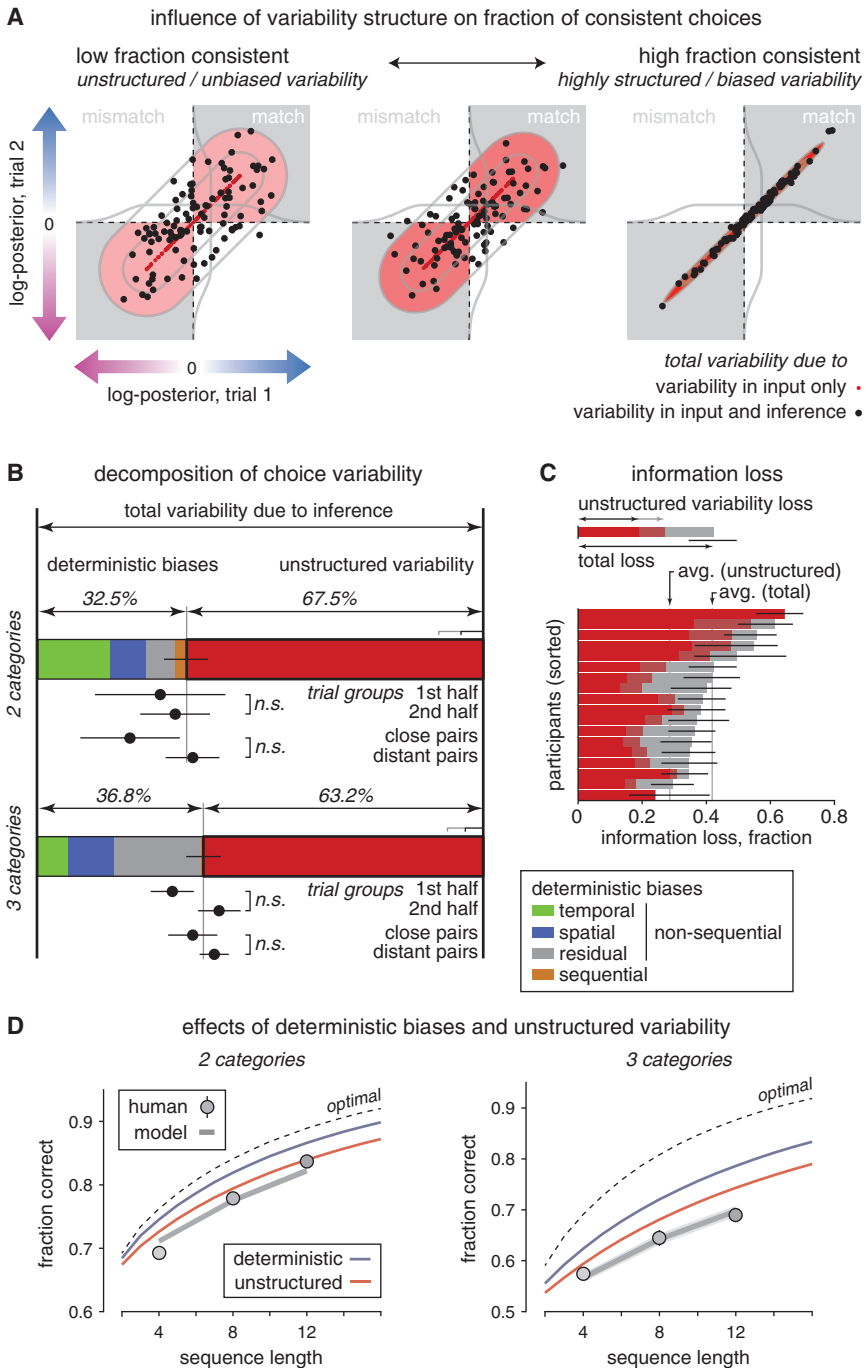

Figure 6. Bias-Variance Partitioning of Inference-Driven Suboptimality

(A) Theoretical relationship between the measured fraction of consistent choices for paired card sequences and the underlying variability structure. Two-dimensional plane representing the noisy log posteriors accumulated in two repetitions of the exact same sequence (each dot corresponds to a pair of trials). Across-repetition noise correlation (or structure) grows with the measured fraction consistent (in red, from left to right). For illustration purposes, we have assumed the noise-free log posterior to be uniformly distributed across trials. The left column corresponds to a fully stochastic observer (with no deterministic bias), whereas the right column corresponds to an almost perfectly deterministic observer (fraction consistent choices ≈1). The middle column corresponds to the bias-variance decomposition estimated in the human data.

(B) Decomposition of human choice suboptimality into deterministic biases (left side) and unstructured "noise" (right side, red) for the two-category (top row) and three-category (bottom row) conditions. The contribution of deterministic biases is split into modeled spatial (blue), modeled temporal (green), sequential (orange), and residual (gray) components. Dots indicate estimates of this bias-variance median split for trial subgroups: first versus second half and close versus distant pairs. Error bars show SEM.

(C) Impact of unstructured variability in mental inference on information loss (loss of mutual information between category and presented orientation per cue; information loss 1 = no mutual information; see Figure S3D and Supplemental Experimental Procedures for details on how loss is computed) in the two-category condition. Bars show estimates of the fraction of information loss sorted across participants (mode ± 95% credible intervals), with the contribution of unstructured variability highlighted in red (different shadings: min/max contribution). Small arrows indicate the average total (~40%) and unstructured (~25%) fraction of information loss.

(D) Predicted impacts of deterministic biases (blue lines) and unstructured variability (red lines) on fraction correct in the two-category (left panel) and three-category (right panel) conditions (dots, mean ± SEM).

To evaluate the relative contributions of these two causes of inferential imperfections, we conducted a third experiment with a task seemingly identical to the first experiment, including both two- and three-category conditions. However, and unbeknownst to participants, every card sequence was presented twice in distinct trials occurring at different points in time throughout the experiment, which allowed us to measure the consistency of choices across repeated sequences. We then used the bias-variance decomposition approach from the Estimator Theory to quantify the relative contributions of deterministic biases and unstructured fluctuations in inference, which partitions choice suboptimality into (1) a "bias" term arising from all possible systematic deviations from Bayes-optimal inference at play in our task and (2) a "variance" term that captures the *intrinsic* variability of choices unrelated to *any* systematic bias (Figure 6A). Importantly, this bias-variance decomposition captures *all* deterministic biases without requiring their explicit description and thus their knowledge as long as they

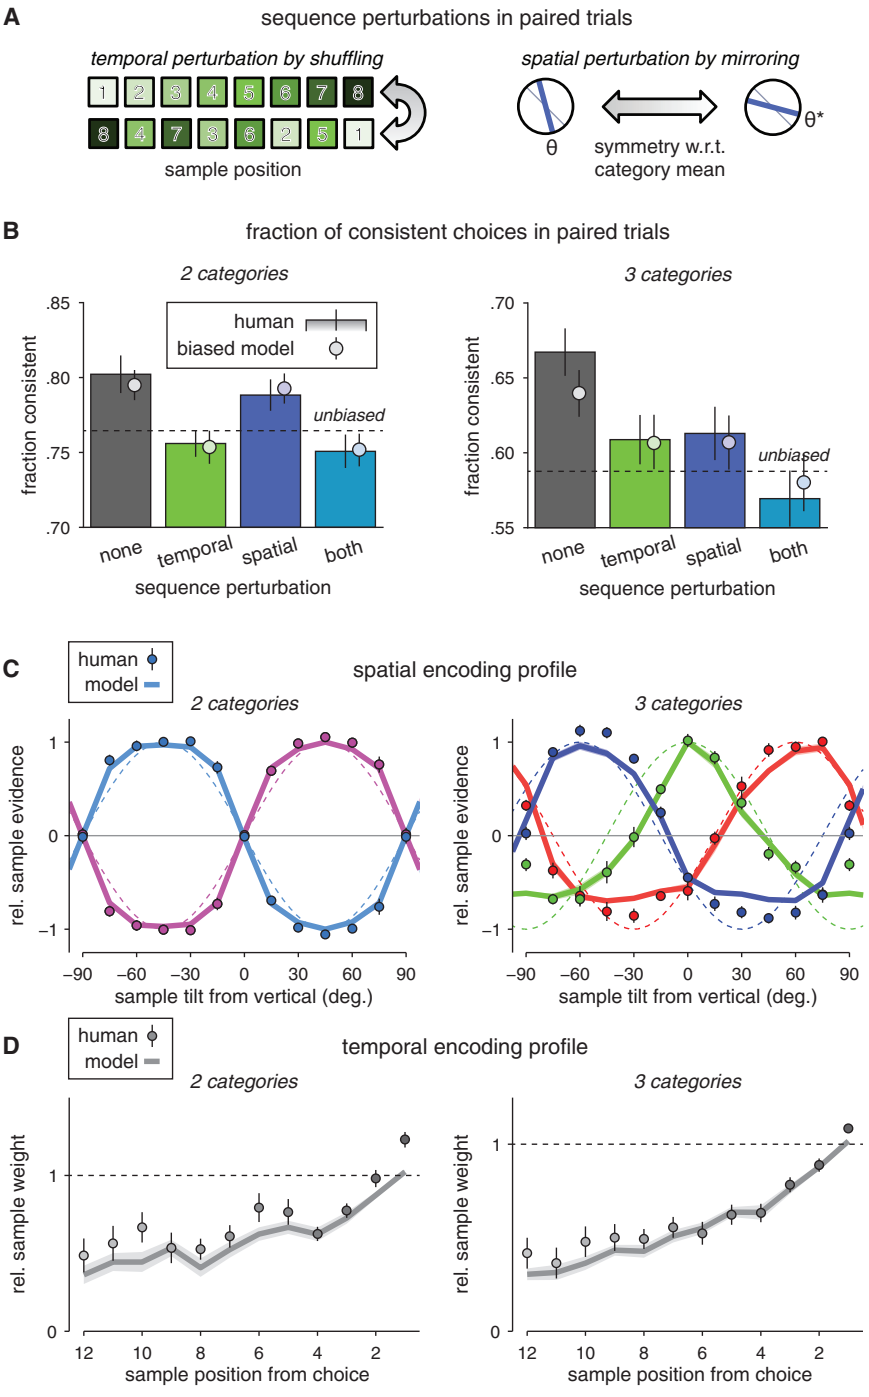

Figure 7. Contribution of Deterministic Biases to Human Choice Suboptimality

(A) Illustration of selective sequence perturbations in the temporal and spatial dimensions of card sequences repeated at different times throughout experiment 3. Left: temporal perturbations were triggered by "shuffling" card positions. Right: spatial perturbations were triggered by "mirroring" sample tilts, with respect to their generative deck orientations. Neither perturbation influences the predictions of the optimal Bayesian decision maker in terms of accuracy or fraction of consistent choices.

(B) Measured fraction of consistent choices for paired sequences with identical (gray bar) or perturbed (colored bars) characteristics in the spatial or temporal dimensions of stimulation in the two-category (left) and three-category (right) conditions. Dots show the predicted fraction of consistent choices for the biased model, whose spatial and temporal encoding profiles are shown in (C) and (D). Error bars show SEM. Dashed lines show the predicted fraction of consistent choices for an unbiased model comprising only unstructured variability.

(C) Empirical spatial encoding profiles in the two-category (left) and three-category (right) conditions. Relative cue evidence (proportional to the maximum cue log likelihood) for human participants (dots), biased (thick lines), and unbiased (dashed lines) models fitted to human choices.

(D) Empirical temporal encoding profiles in the two-category (left) and three-category (right) conditions. Relative cue weights for human participants (dots), biased (thick lines), and unbiased (dashed lines) models fitted to human choices. Error bars and shades indicate SEM.

remain stable across trials (see Experimental Procedures and Figure S3A).

In both conditions, participants made a substantial fraction of inconsistent choice to repeated card sequences (Figure 7B, gray bars), confirming that choice suboptimality did not arise exclusively from deterministic biases. Furthermore, the fraction of consistent choices in response to identical card sequences exceeded that of a model featuring only unstructured fluctuations in inference (Figure 7B, dashed lines; both $t_{17} > 7.3$, $p < 0.001$).

Quantitatively, the bias-variance decomposition attributed only 33% and 37% of the total choice suboptimality to deterministic biases in the two- and three-category conditions, respectively (Figure 6B), even when taking into account sequential dependencies across consecutive choices that contributed 7% and 1% of deterministic biases in the two conditions (Figure 2B). This means that about two-thirds of human choice suboptimality was unrelated to *any* deterministic bias and was consequently uniquely attributable to unstructured fluctuations in inference (with minor contributions of sensory noise and stochasticity in the response selection that amount together to less than 10% of the total choice suboptimality). In the two-category condition, these unstructured fluctuations alone amounted to a 29% loss in every piece of incoming information (Figure 6C). Interestingly, although the total choice suboptimality increased in the three-category condition relative to the two-category condition, the proportion with which unstructured fluctuations contributed to the total choice suboptimality was statistically

indistinguishable between conditions ($t_{17} = 0.3$, $p > 0.5$). In other words, the distinct contributions of deterministic biases and unstructured fluctuations to choice suboptimality grew in the same proportion when choosing among three categories instead of two categories. Using the same validation procedure as before (see Supplemental Experimental Procedures), we confirmed that our decomposition technique correctly estimated the bias-variance tradeoff on synthetic choice data from models featuring known proportions of deterministic biases and unstructured fluctuations (Figure S2C).

Characterizing the Nature of Deterministic Biases

Our bias-variance decomposition confirmed the presence of deterministic biases in mental inference. Some of them certainly affect the "temporal" accumulation of evidence, whereas others concern the "spatial" mapping of card orientations into the category space. To quantify the respective influences of these two types of biases, each card sequence was not only repeated identically, but also after shuffling the card presentation order and/or after mirroring pattern orientations relative to the mean orientation of the deck they were drawn from (Figure 7A). Importantly, these spatiotemporal perturbations preserved the total accumulated evidence in terms of Bayes-optimal inference. Consequently, the fraction of consistent choices should decrease when comparing temporally perturbed pairs of sequences to identical pairs only in the presence of temporal biases. The exact same prediction holds for spatial biases across spatially perturbed pairs of sequences. As expected, we found the fraction of consistent choices to decrease significantly for both temporally and spatially perturbed compared to identical sequences (Figure 7B, colored bars; spatial perturbation: $F_{1,17} > 6.7$, $p < 0.05$; temporal perturbation: $F_{1,17} > 52.2$, $p < 0.001$; interaction: $F_{1,17} < 0.9$, $p > 0.2$), indicating the presence of spatiotemporal biases in mental inference.

To characterize these spatiotemporal biases, we first fitted the contributions of cards presented at different positions in the sequence and of different orientations to participants' choices (Figures 7C and 7D, dots). We found that compared to Bayes-optimal inference, a card positioned closer to the choice and card orientations closer to vertical and horizontal orientations over-contributed to the inference process. These distortions corresponded to well-known cognitive biases overweighting more recent pieces of information (Usher and McClelland, 2001; Ossmy et al., 2013) and cardinal orientations (Girshick et al., 2011; Wei and Stocker, 2015). We then considered a model featuring Bayes-optimal inference, along with explicit accounts of these cognitive biases (Figure S5), as well as additional unstructured fluctuations in the inference process. Fitting this model to participants' choices (Figures 7C and 7D, lines) showed that these biases account for the decrease in choice consistency observed across temporally and spatially perturbed sequences (Figure 7B, dots). Furthermore, a bias-variance decomposition accounting for these particular biases indicated that they could explain 72% and 46% of all deterministic biases in the two- and three-category conditions, respectively (Figure 6B). Importantly, we verified that modeling these biases explicitly during bias-variance decomposition did not modify our main conclusion: unstructured fluctuations in inference remained the dominant source of choice suboptimality (Figure S6).

DISCUSSION

Making decisions often requires combining multiple pieces of ambiguous information from external cues (Gold and Shadlen, 2007; Shadlen and Kiani, 2013). In such conditions, human choices derive from covert mental processes that resemble probabilistic inference (Roitman and Shadlen, 2002; Weiss et al., 2002; Ernst and Banks, 2002; Behrens et al., 2007; Yang and Shadlen, 2007), but are typically highly variable and consequently often suboptimal (Beck et al., 2012). Previous decision-making studies, ranging from perceptual categorization to reward-guided learning, have classically attributed this choice suboptimality to either noisy sensory processing in perceptual tasks featuring weak, ambiguous, or noisy sensory evidence (Osborne et al., 2005; Brunton et al., 2013; Kaufman and Churchland, 2013) or stochastic response selection in tasks featuring open questions (Griffiths and Tenenbaum, 2006; Vul and Pashler, 2008; Vul et al., 2009), sequential learning (Acerbi et al., 2014), or volatile contingencies (Daw et al., 2006; Behrens et al., 2007). However, the tasks used in these studies make it theoretically impossible to distinguish these peripheral sources of choice suboptimality from imperfections in probabilistic inference. By contrast, in the present study, imperfections in sensory processing, inference, and response selection were distinguishable by altering human choices with distinct statistical signatures. This allowed us to show that imperfections in inference alone accounted for about 90% of human suboptimal choices in our task, whereas imperfections in sensory processing and response selection together form a negligible fraction. Thus, inferential imperfections constitute an important source of human choice suboptimality, which may have been underestimated by previous studies and confounded with imperfections in sensory processing or response selection.

The analysis of choice consistency across repeated pseudo-random sequences of cues further reveals that only one-third of inferential imperfections stemmed from deterministic distortions of Bayes-optimal inference (Beck et al., 2012). Importantly, our bias-variance decomposition approach further allowed us to establish the fraction of suboptimal choices that derives from *all* possible deterministic deviations from Bayes-optimal inference, only part of which could be identified explicitly. Unidentified biases possibly include adaptive gain coding, range normalization, and anchoring effects across successive cues or choice alternatives (Usher and McClelland, 2001; Albantakis and Deco, 2009; Louie et al., 2013; Cheadle et al., 2014).

Consequently, our bias-variance decomposition indicates that two-thirds of inferential imperfections are neither related to the structure of our task (and *any* deterministic distortion associated with our task) nor to *any* misspecification of our computational model. We thus propose that this dominant fraction of human suboptimal choices reflects random fluctuations (i.e., noise) in inference, arising primarily from intrinsic variability in the computational and coding precision of elicited variables represented in populations of neurons (Renart and Machens, 2014). In information terms, this intrinsic variability yielded a surprisingly large loss of about 30% in every piece of information in the canonical two-category condition, which raises the issue of whether it arises from a lack of attention or motivation, insufficient training, or

increasing fatigue over the course of the experiment. None of these factors, however, appears to be a plausible explanation. First, a "lapse" probability parameter in our models captured random choices unrelated to stimuli (Figures S2A and S7), thus factoring out lapses in attention from our estimation of inferential variability. Second, the orientation discrimination thresholds measured in our task are consistent with values tabulated in the literature (Burbeck and Regan, 1983; Webster et al., 1990; Burr and Wijesundra, 1991), indicating that participants were focused on the task. Third, accuracy did not rise over the course of the experiment, contrary to what insufficient training would predict. Fourth, the stable accuracy observed within each experimental session did not argue in favor of increasing fatigue. Consequently, this intrinsic variability appears to reflect the near-maximal precision of neural codes and computations implementing inference.

At the neural level, inference has been hypothesized (Knill and Pouget, 2004; Ma et al., 2006) and shown to engage parietal and prefrontal regions that encode and accumulate evidence along category-defining stimulus features (Freedman and Assad, 2006; Fitzgerald et al., 2011; Mante et al., 2013). The inferential variability we measured here may arise at three stages of neural processing implementing inference: (1) in the neural mapping from sensory regions encoding relevant stimulus features (orientation in our case) to parietal and prefrontal regions encoding abstract categories (Soltani and Wang, 2010), (2) in the neural coding of abstract categories in these associative regions (e.g., Mante et al., 2013), and (3) in the neural updating of category-selective representations following a new piece of information (Yang and Shadlen, 2007; Cain et al., 2013). Importantly, the inferential variability we estimated from behavior measures the *effective* precision of neural processing at the level of large ensembles of neurons implementing inference rather than the neural variability of individual neurons or synapses (Renart and Machens, 2014), which presumably average out over large populations of neurons (Beck et al., 2012). One possibility is that this limited precision is the result of computational constraints on neural circuits implementing inference. In particular, the maintenance of information over time in neural circuits (up to several seconds in our paradigm) might suffer from "temporally diffusive" noise, which grows with elapsed time in the decision process (Burak and Fiete, 2012). Alternatively, the processing of successive cues might depend on a slow "cognitive bottleneck," which limits the processing resources that can be allocated to cues presented in rapid succession (Wyart et al., 2012, 2015). These two accounts differ in terms of their dependency on presentation rate: a faster presentation rate should decrease choice variability for "temporally diffusive" noise, but increase choice variability in case of a "cognitive bottleneck." Another possibility is that this limited precision reflects the neural implementation of sampling processes for realizing probabilistic inference (Haefner et al., 2016). Such processes are effective algorithms for approximating Bayes-optimal inference, which rapidly become intractable in real-world situations (Fiser et al., 2010; Brooks et al., 2011). They produce sequences of samples from posterior distributions over task-relevant variables like categories in the present study. Drawing few samples results in noisy posterior distributions (Lengyel et al., 2015), corresponding to the intrinsic variability we identified here. Disentangling these different neural accounts of intrinsic variability in inference is beyond the scope of the present study and remains an open question for future work.

An important question concerns the potential function of inferential imprecisions: why do humans feature such a large degree of random fluctuations in inference if these fluctuations cause such a sizable fraction of additional, avoidable behavioral errors? The hypothesis of computational constraints on neural circuits implementing inference raises the issue of whether extensive practice of our task featuring few and ambiguous cues over days or weeks would induce the recruitment of additional neural resources and improve the effective precision of mental inference. Another possibility is that in volatile (changing) environments, a low computational precision comes at little cost because in such circumstances, a "diffusion" of belief distributions over task-relevant variables becomes the normative strategy (Behrens et al., 2007), something that can be achieved by random fluctuations in inference. In this context, the limited precision of inference might even explain the presence of deterministic biases in decision making. For instance, if the precision of inference is low, it becomes advantageous to overweight the pieces of information that have been the least neurally processed, typically the most recent ones. Although "recency" biases are typically suboptimal in the context of noise-free computations, they may have evolved as optimizing inference in neural circuits implementing approximate computations.

In summary, identifying the origin of human choice suboptimality under uncertainty and elucidating the structure of the underlying variability is critical for understanding decision making and its neural substrates. Beyond noise in sensory processing and stochasticity in response selection, we found that imperfections in probabilistic inference are a significant contributor to human choice suboptimality. Critically, two-thirds of this choice suboptimality derive from random fluctuations in inference rather than from biased computations. This intrinsic variability reflects the effective precision of neural computations underlying inference. This computational precision sets a previously unsuspected low upper bound on the accuracy and ultimate predictability of human choices in uncertain environments, which needs to be accounted for in theoretical models of decision making.

EXPERIMENTAL PROCEDURES

Participants

63 healthy participants took part in the three experiments, with no overlap among the participants tested in each experiment. All had normal or corrected-to-normal vision and no history of neurological or psychiatric disorders (35 females, mean age: 24). All human participants provided informed written consent prior to the experiment, and both experiments were approved by the local ethics committee (Comité de Protection des Personnes, Ile-de-France VI, Inserm approval #C07-28, DGS approval #2007-0569, IDRCB approval #2007-A01125-48). Our sample sizes are similar to those generally employed for comparable studies.

Stimuli and Task Design

The task was a variant of the "weather prediction" task, in which participants were asked to infer the generative category (deck) of a sequence of stimuli (cards) among two (or three, see below) alternatives that differed in terms of their generative distributions (Figure 1A). Stimuli depicted high-contrast, noise-free Gabor patterns of varying orientation, presented at fixation at an

average rate of 3 Hz (see Supplemental Experimental Procedures). In each trial, the orientation of successive cards varied according to a circular Gaussian distribution (von Mises distribution; concentration κ = 0.5 and 0.7 for the two- and three-category conditions, respectively) centered on a mean orientation characterizing each deck. The mapping rule between orientation and categories (represented as colors) was provided throughout the whole experiment by a static colored annulus surrounding the stimuli. Sequence lengths (i.e., the number of stimuli per sequence) varied unpredictably from 2 to 16 stimuli across trials. At sequence offset, participants provided their response by pressing one out of two (or three) keys with their right hand. Following each response, feedback about the true generative category of the sequence was provided via a transient color change in the fixation point.

Experiment 1 (25 participants) consisted of two conditions (choice among two or three categories, Figure 1A), corresponding to two 50-min sessions taking place on different days. Both conditions were divided into short blocks of approximately 50 trials (each lasting about 5 min), such that participants could take short rest periods between them. Three participants were excluded from analyses because they misunderstood the task instructions; their decisions were based only on the last card of each sequence (as revealed by a standard logistic regression of choices).

Experiment 2 (20 participants) was identical to the two-category condition of experiment 1, except that the "accumulation" trials used in experiment 1 were randomly and unpredictably mixed with "last card" trials in equal proportions. Last-card trials were identical to "accumulation" trials, except that the last card was paired with a single auditory tone that instructed participants to report the category of this last card, irrespective of preceding ones. The last sample was drawn independently of preceding samples at 1, 2, 4, or 8 deg. from either of the two category boundaries corresponding to the horizontal and vertical axes. The experiment was divided into short blocks of 72 trials (each lasting about 7 min), corresponding to two 60-min sessions taking place on different days. Trials included sequences of 2, 4, 8, or 16 samples. Three participants were excluded from analyses because of chance-level performance in either accumulation or last-card trials.

Experiment 3 (20 participants) was similar to experiment 1 and again included the two- and three-category conditions. Experiment 3 consisted of two 50-min sessions taking place on different days. As in experiment 1, both conditions were divided into short blocks of approximately 50 trials. In each condition and unbeknownst to participants, each pseudo-random sequence containing 4, 8, or 12 samples was repeatedly presented at five occasions consisting of (1) an exact repetition of the original sequence, (2) the original sequence, in which the presentation order was shuffled, (3) the original sequence, in which sample orientations were mirrored with respect to the mean orientation of the generative category, and (4) the original sequence with both the abovementioned perturbations (Figure 7A). The temporal distance (in number of intervening trials) among these different presentations of the same sequence was controlled for within and across participants. No participant reported any repetition effects at the end of the experiment. Two participants were excluded from the analyses because they did not perform one of the two sessions.

Modeling and Fitting Sources of Choice Suboptimality

Here, we provide a summary of the used models. More details and discussion can be found in the Supplemental Experimental Procedures.

In trial n, after T_n observed sample orientations $\theta_{n1}, \ldots, \theta_{nT_n}$, the ideal decision maker accumulates log likelihoods ℓ_{ntk} to form the log posterior $z_{nT_nk} = \sum_{t=1}^{T_n} \ell_{ntk}$, with respect to each category k, and chooses the category x_n associated with the largest of these log posteriors, $x_n = \mathrm{argmax}_k(z_{nT_nk})$ (Figure 1B). Each sample is generated by draws from a von-Mises distribution on a half-circle, with mean μ_k for category k, such that the log likelihood of the sample at position t, with respect to category k, is given by $\ell_{ntk} = \kappa \cos(2(\theta_{nt} - \mu_k))$. In each trial, the Bayes-optimal choice is deterministically related to the sequence of samples, such that all choice variability across trials stems from variability in presented samples.

To introduce additional choice variability, we considered various hypotheses (Figures 1B and 1C). First, we introduced noisy orientation percepts (sensory noise) in log likelihoods $\hat{\ell}_{ntk} = \kappa \cos(2(\theta_{nt} + \varepsilon_{nt} - \mu_k))$, where $\varepsilon_{nt} \sim \mathrm{N}(0, \sigma^2_{sen})$ are zero-mean Gaussian variables, independent across samples, with sensory noise variance σ^2_{sen}. Second, we introduced variability at the inference stage by noisy log-likelihood estimates $\hat{\ell}_{ntk} = \ell_{ntk} + \varepsilon_{ntk}$, where $\varepsilon_{ntk} \sim \mathrm{N}(0, \sigma^2_{inf})$ are independent zero-mean Gaussian variables, with inference noise variance σ^2_{inf}. In contrast to sensory noise, this noise causes uncorrelated Gaussian noise in the log posteriors, leading to different predictions of individual choices (see Supplemental Experimental Procedures). We modeled two variants of variability at the inference stage: the first assumes additive noise to each log likelihood ("likelihood" variability), resulting in T_n noise terms, and the second assumes additive noise to the combination of log likelihoods ("accumulation" variability), resulting in $T_n - 1$ noise terms. The accumulation variability model provided better fits to the observed behavior (two-category: Bayes factor $>10^2$, exceedance p $\approx$ 0.81; three-category: Bayes factor $>10^{12}$, exceedance p > 0.99; Figures S6B and S6C), and is thus the "inference" model discussed in the main text. Third, we introduced stochasticity in response selection by drawing samples from the posterior belief taken to some power, $p(x_n = k \mid \theta_{n1:T_n}) \propto \exp(\beta z_{nT_nk})$, where β is a free parameter that is fit to the observed behavior ($\beta = 0$, random choices; $\beta = 1$, posterior sampling; $\beta \to \infty$, deterministic choices). This response strategy is indistinguishable from adding single constant-variance Gaussians added to the log posterior of each category (see Supplemental Experimental Procedures). Thus, in contrast to sensory noise and inference variability, stochasticity in response selection postulates that the magnitude of the additional variability is independent of the number of samples, T_n, in the sequence. Finally, we considered a model in which all variability is introduced by a variable initial state of the log-posterior accumulator (Figure S6).

To distinguish between these different sources of choice variability, we fitted for each subject separately for each model the choices of all trials combined by a maximum-likelihood procedure that, on one hand, avoids local maxima and, on the other hand, estimates parameter uncertainty by the width of the posterior. We avoided confounders due to occasional random responses and response biases by adding a "lapse" probability parameter and $K - 1$ response bias parameters (Figure S7). Model comparison (both fixed effects and random effects) was based on approximating the model evidence by the Bayesian information criterion. More details on model fitting can be found in the Supplemental Experimental Procedures. In all model fits, the concentration parameter κ was used as a scaling parameter by setting it to its true value. Participants could have misestimated this concentration, but such a misestimation would not qualitatively change our conclusions (see Supplemental Experimental Procedures).

Figure 3A shows the scaling of choice variability in two ways. The solid lines indicate the maximum-likelihood model predictions averaged across fitted participants. The dots show noise parameters re-fitted separately for each sequence length and each participant. In both cases, we did not include lapses and response biases in these fits to avoid model-dependent parameter biases across different sequence lengths. The model comparison in Figure 3B was based on model fits that included a lapse probability and response biases.

The fraction correct measured in "last card" trials presented in experiment 2 was fitted using a model featuring only sensory noise, precisely because no probabilistic inference was required to solve these trials (in the sense of combining the information provided by successive samples). We inferred the orientation discrimination threshold Δ_θ (Figure 4C) per participant and condition from $0.75 = \Phi(\Delta_\theta / \sqrt{2\sigma^2_{sen}})$, where Φ is the standard cumulative Gaussian, corresponding to the difference in orientations that results in 75% correct choices in a two-alternative forced-choice task, given a Gaussian orientation percept with variance σ^2_{sen}.

The fraction choice variability in Figure 5A was inferred per participant for the two-category condition by fitting a model that combined variability in sensory processing, inference, and response selection, resulting in a separate log-posterior difference variance estimate for each type of variability that contributes additively to the behavioral variability and could thus be decomposed. We could not apply the same rationale to the three-category condition because choices are based on two log-posterior differences rather than one.

The performance predictions in Figures 5B and 5C resulted from averaging over generated choices for 10^6 virtual trials for each participant and sequence length, with the same sequence statistics as those presented to the participants. This allowed us to provide predictions for all sequence lengths,

including ones not used in the experiment. Choices were generated either according to the Bayes-optimal model or by models including single sources of variability whose magnitudes matched those shown in Figure 5A. The full model combined all three sources of variability and provides predictions for the actual sequences presented to the participants to make them directly comparable to participants' choice behavior.

Unless noted otherwise, statistical analyses of differences of scalar model-free measures (fraction correct, fraction suboptimal, and fraction consistent) or model-based quantities (e.g., estimates of choice variability) between conditions of interest (e.g., sequence length, two versus three categories) relied on standard, two-tailed parametric tests (e.g., paired t test, repeated-measures ANOVA) across tested participants (≈20 in each experiment), i.e., outside the small-sample regime, thereby matching the core assumptions of the applied statistical parametric tests.

Modeling and Fitting the Structure of Choice Suboptimality

We modeled the bias-variance structure by assuming that inference variability results from a noisy likelihood $\widehat{\ell}_{ntk} = \ell_{ntk} + f_k(\theta_{nt}) + \varepsilon_{ntk}$ that can be decomposed into (1) the correct likelihood ℓ_{ntk}, (2) a sequence-dependent deterministic bias $f_k(\theta_{nt})$, and (3) additive random fluctuations ε_{ntk}. This model provided joint likelihoods for trial pairs, in which identical card sequences where shown with separate contributions of deterministic biases and random fluctuations. Thus, we could estimate these contributions by fitting participants' choice in pairs using the same maximum likelihood procedure as before. We quantified the contribution of explicit spatial and temporal perturbations and sequential choice dependencies to deterministic biases in Figure 6B by measuring how much of these biases could be "explained away" by modeling these perturbations and dependencies explicitly. Furthermore, we performed multiple control analyses to determine whether participants' choices featured slowly drifting biases, which the bias-variance decomposition could have mistaken for random fluctuations (Figure S4). None of the control analyses suggested this to be the case. Details of the model, fitting procedures, and control analyses can be found in the Supplemental Experimental Procedures.

The performance predictions in Figure 6D were generated similarly to those in Figure 5B. The predictions for Bayes-optimal inference, those for a model with deterministic biases only, and those for one with unstructured fluctuations only were simulated for 10^6 virtual trials with the same sequences statistics as those seen by the participants. Those for the full model were based on the actual sequences presented to the participants to make them directly comparable to participants' choice behavior. The variability magnitudes used for simulation were those estimated from fits of the full model to the choice behavior of individual participants.

Assumed Deviations of Bayes-Optimal Inference

"Spatial" distortions introduce deterministic perturbations into the mapping between the orientation of the presented sample and the associated log likelihoods, with respect to each category. We assumed three such perturbations that could occur in isolation or in combination. The "orientation" bias assumes that the perceived orientation was tilted by a constant angle (Figure S5A). The "confirmation" bias assumes a likelihood that is increased if positive and decreased if negative by the same value, and as such introduced an imbalance in how each sample contributed to the log posterior, with respect to different categories (Figure S5B). The "oblique" effect assumes a perturbation of the perceived orientation either toward or away from oblique orientations, depending on its parameter (Figure S5C).

When modeling "temporal" distortions from Bayes-optimal inference, we assumed that samples within an observed sequence contributed with different weights to the final choice. We considered two variants. The first assume that with each new sample, the log posterior up to, but excluding, this sample is multiplied (i.e., either discounted or amplified) by a constant multiplicative factor. Overall, depending on the factor (lower or higher than one), this leads to "recency" or "primacy" effects (Figure S5D). The second variant assumes this factor to change linearly with the position of the current sample within the sequence and is thus more general than the first variant (Figure S5E). The mathematical description of each of these distortions can be found in the Supplemental Experimental Procedures.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and seven figures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2016.11.005.

AUTHOR CONTRIBUTIONS

J.D., V.W., and E.K. designed the experiments. V.W. and A.-D.D. conducted the experiments. J.D. and V.W. designed the models and performed the analyses. J.D., V.W., A.-D.D., and E.K. discussed the results and wrote the paper.

ACKNOWLEDGMENTS

This work was supported by an advanced research grant from the European Research Council awarded to E.K. (ERC-2009-AdG-250106), a young investigator award from the Fyssen Foundation to V.W., a junior researcher grant from the French National Research Agency awarded to V.W. (ANR-14-CE13-0028), and two department-wide grants from the French National Research Agency (ANR-10-LABX-0087 and ANR-10-IDEX-0001-02).

Received: May 31, 2016
Revised: August 4, 2016
Accepted: October 28, 2016
Published: December 1, 2016

REFERENCES

Acerbi, L., Vijayakumar, S., and Wolpert, D.M. (2014). On the origins of suboptimality in human probabilistic inference. PLoS Comput. Biol. *10*, e1003661.

Albantakis, L., and Deco, G. (2009). The encoding of alternatives in multiple-choice decision making. Proc. Natl. Acad. Sci. USA *106*, 10308–10313.

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.

Behrens, T.E.J., Woolrich, M.W., Walton, M.E., and Rushworth, M.F.S. (2007). Learning the value of information in an uncertain world. Nat. Neurosci. *10*, 1214–1221.

Brooks, S., Gelman, A., Jones, G.L., and Meng, X.-L. (2011). Handbook of Markov Chain Monte Carlo (CRC Press).

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98.

Burak, Y., and Fiete, I.R. (2012). Fundamental limits on persistent activity in networks of noisy neurons. Proc Natl Acad Sci. *109*, 17645–17650.

Burbeck, C.A., and Regan, D. (1983). Independence of orientation and size in spatial discriminations. J. Opt. Soc. Am. *73*, 1691–1694.

Burr, D.C., and Wijesundra, S.A. (1991). Orientation discrimination depends on spatial frequency. Vision Res. *31*, 1449–1452.

Cain, N., Barreiro, A.K., Shadlen, M., and Shea-Brown, E. (2013). Neural integrators for decision making: a favorable tradeoff between robustness and sensitivity. J. Neurophysiol. *109*, 2542–2559.

Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Herce Castañón, S., and Summerfield, C. (2014). Adaptive gain control during human perceptual choice. Neuron *81*, 1429–1441.

Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat. Neurosci. *8*, 1704–1711.

Daw, N.D., O'Doherty, J.P., Dayan, P., Seymour, B., and Dolan, R.J. (2006). Cortical substrates for exploratory decisions in humans. Nature *441*, 876–879.

Daw, N.D., Gershman, S.J., Seymour, B., Dayan, P., and Dolan, R.J. (2011). Model-based influences on humans' choices and striatal prediction errors. Neuron *69*, 1204–1215.

Dayan, P. (2014). Rationalizable irrationalities of choice. Top. Cogn. Sci. *6*, 204–228.

Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. Nature *415*, 429–433.

Fiser, J., Berkes, P., Orbán, G., and Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. Trends Cogn. Sci. *14*, 119–130.

Fitzgerald, J.K., Freedman, D.J., and Assad, J.A. (2011). Generalized associative representations in parietal cortex. Nat. Neurosci. *14*, 1075–1079.

Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85–88.

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.

Gluck, M.A., Shohamy, D., and Myers, C. (2002). How do people solve the "weather prediction" task?: individual variability in strategies for probabilistic category learning. Learn. Mem. *9*, 408–418.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Griffiths, T.L., and Tenenbaum, J.B. (2006). Optimal predictions in everyday cognition. Psychol. Sci. *17*, 767–773.

Haefner, R.M., Berkes, P., and Fiser, J. (2016). Perceptual decision-making as probabilistic inference by neural sampling. Neuron *90*, 649–660.

Kaufman, M.T., and Churchland, A.K. (2013). Cognitive neuroscience: sensory noise drives bad decisions. Nature *496*, 172–173.

Knill, D.C., and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. Trends Neurosci. *27*, 712–719.

Knowlton, B.J., Mangels, J.A., and Squire, L.R. (1996). A neostriatal habit learning system in humans. Science *273*, 1399–1402.

Lengyel, M., Koblinger, Á., Popović, M., and Fiser, J. (2015). On the role of time in perceptual decision making. arXiv, arXiv:1502.03135, https://arxiv.org/abs/1502.03135.

Louie, K., Khaw, M.W., and Glimcher, P.W. (2013). Normalization is a general neural mechanism for context-dependent decision making. Proc. Natl. Acad. Sci. USA *110*, 6139–6144.

Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. Nat. Neurosci. *9*, 1432–1438.

Ma, W.J., Navalpakkam, V., Beck, J.M., Berg, Rv., and Pouget, A. (2011). Behavior and neural basis of near-optimal visual search. Nat. Neurosci. *14*, 783–790.

Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78–84.

Navalpakkam, V., Koch, C., Rangel, A., and Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. Proc. Natl. Acad. Sci. USA *107*, 5232–5237.

Osborne, L.C., Lisberger, S.G., and Bialek, W. (2005). A sensory source for motor variation. Nature *437*, 412–416.

Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., and Donner, T.H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. Curr. Biol. *23*, 981–986.

Poldrack, R.A., Clark, J., Paré-Blagoev, E.J., Shohamy, D., Creso Moyano, J., Myers, C., and Gluck, M.A. (2001). Interactive memory systems in the human brain. Nature *414*, 546–550.

Renart, A., and Machens, C.K. (2014). Variability in neural activity and behavior. Curr. Opin. Neurobiol. *25*, 211–220.

Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. *22*, 9475–9489.

Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791–806.

Soltani, A., and Wang, X.-J. (2010). Synaptic computation underlying probabilistic inference. Nat. Neurosci. *13*, 112–119.

Sutton, R.S., and Barto, A.G. (1998). Reinforcement Learning (MIT Press).

Tsotsos, J.K. (2001). Complexity, vision and attention. In Vision and Attention, L. Harris and M. Jenkin, eds. (Cambridge University Press), pp. 105–128.

Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. Psychol. Rev. *108*, 550–592.

Vul, E., and Pashler, H. (2008). Measuring the crowd within: probabilistic representations within individuals. Psychol. Sci. *19*, 645–647.

Vul, E., Hanus, D., and Kanwisher, N. (2009). Attention as inference: selection is probabilistic; responses are all-or-none samples. J. Exp. Psychol. Gen. *138*, 546–560.

Webster, M.A., De Valois, K.K., and Switkes, E. (1990). Orientation and spatial-frequency discrimination for luminance and chromatic gratings. J. Opt. Soc. Am. A *7*, 1034–1049.

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.

Whiteley, L., and Sahani, M. (2012). Attention in a bayesian framework. Front. Hum. Neurosci. *6*, 100.

Wyart, V., de Gardelle, V., Scholl, J., and Summerfield, C. (2012). Rhythmic fluctuations in evidence accumulation during decision making in the human brain. Neuron *76*, 847–858.

Wyart, V., Myers, N.E., and Summerfield, C. (2015). Neural mechanisms of human perceptual choice under focused and divided attention. J. Neurosci. *35*, 3485–3498.

Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. Nature *447*, 1075–1080.

Neuron, Volume *92*

Supplemental Information

Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality

Jan Drugowitsch, Valentin Wyart, Anne-Dominique Devauchelle, and Etienne Koechlin

Contents

1 Supplemental Figures

List of Figures

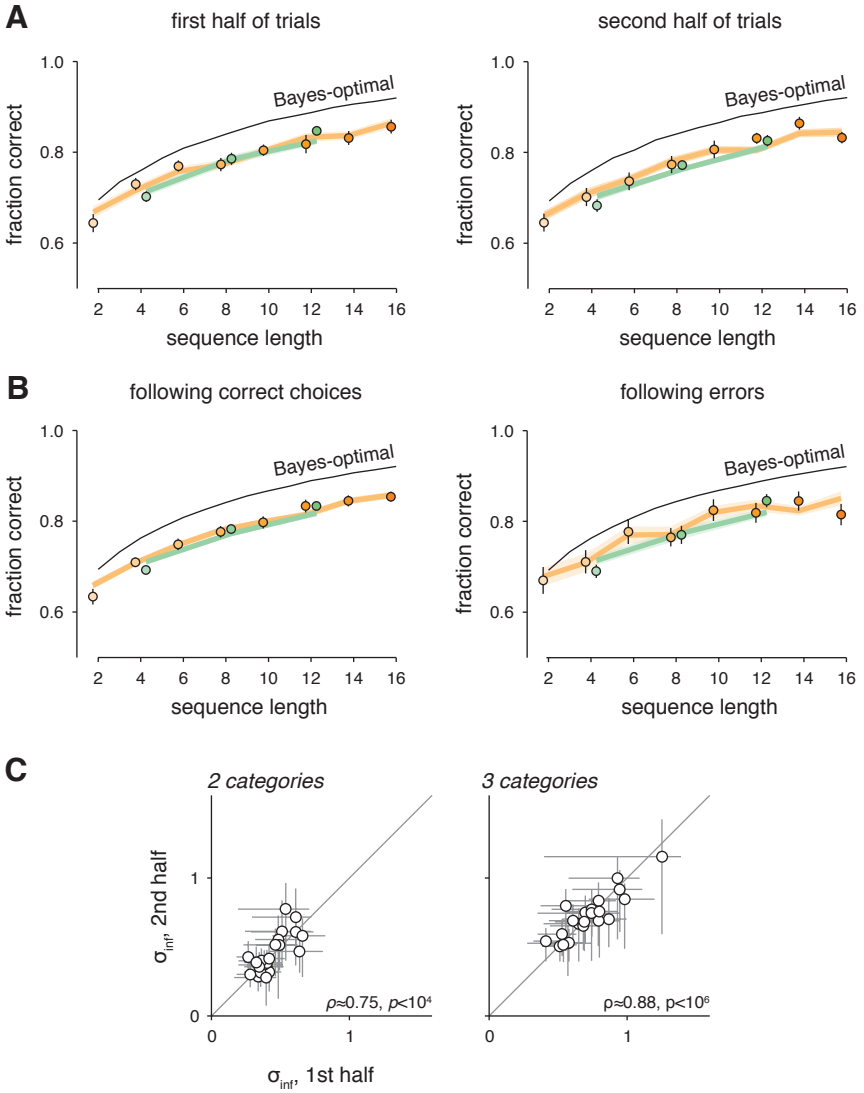

Figure S1. Relates to Figure 2. Psychometric functions and variability for different trial subgroups. (A) Left panel: psychometric function for the two-category condition of exp. 1 (orange) and 3 (green) in the first half of trials. Right panel: psychometric function for the two-category condition of exp. 1 and 3 in the second half of trials. (B) Left panel: psychometric function for the two-category condition of exp. 1 and 3 in trials following correct choices. Right panel: psychometric function for the two-category condition of exp. 1 and 3 in trials following errors. In both (A) and (B) error bars indicate s.e.m., and the black curve indicates the theoretical psychometric function of the normative, Bayes-optimal observer. Due to the lower number of available trials, we did not include the data and fits from experiment 2. (C) The two panels show for each subject the inferred variability magnitude (mode $\pm$ 95% credible intervals) when fitting the inference variability model to either only the first half or the second half of all trials of the first experiment. In no case was there a significant difference across subjects between the variability magnitude for the two trial subgroups (2 categories, $t_{21} = 1.1, p = 0.3$; 3 categories, $t_{21} = 0.5, p = 0.6$). In both cases, the inferred magnitude was significantly correlated across trial subgroups (ρ and p in plots denote Pearson product-moment correlation coefficient statistics).

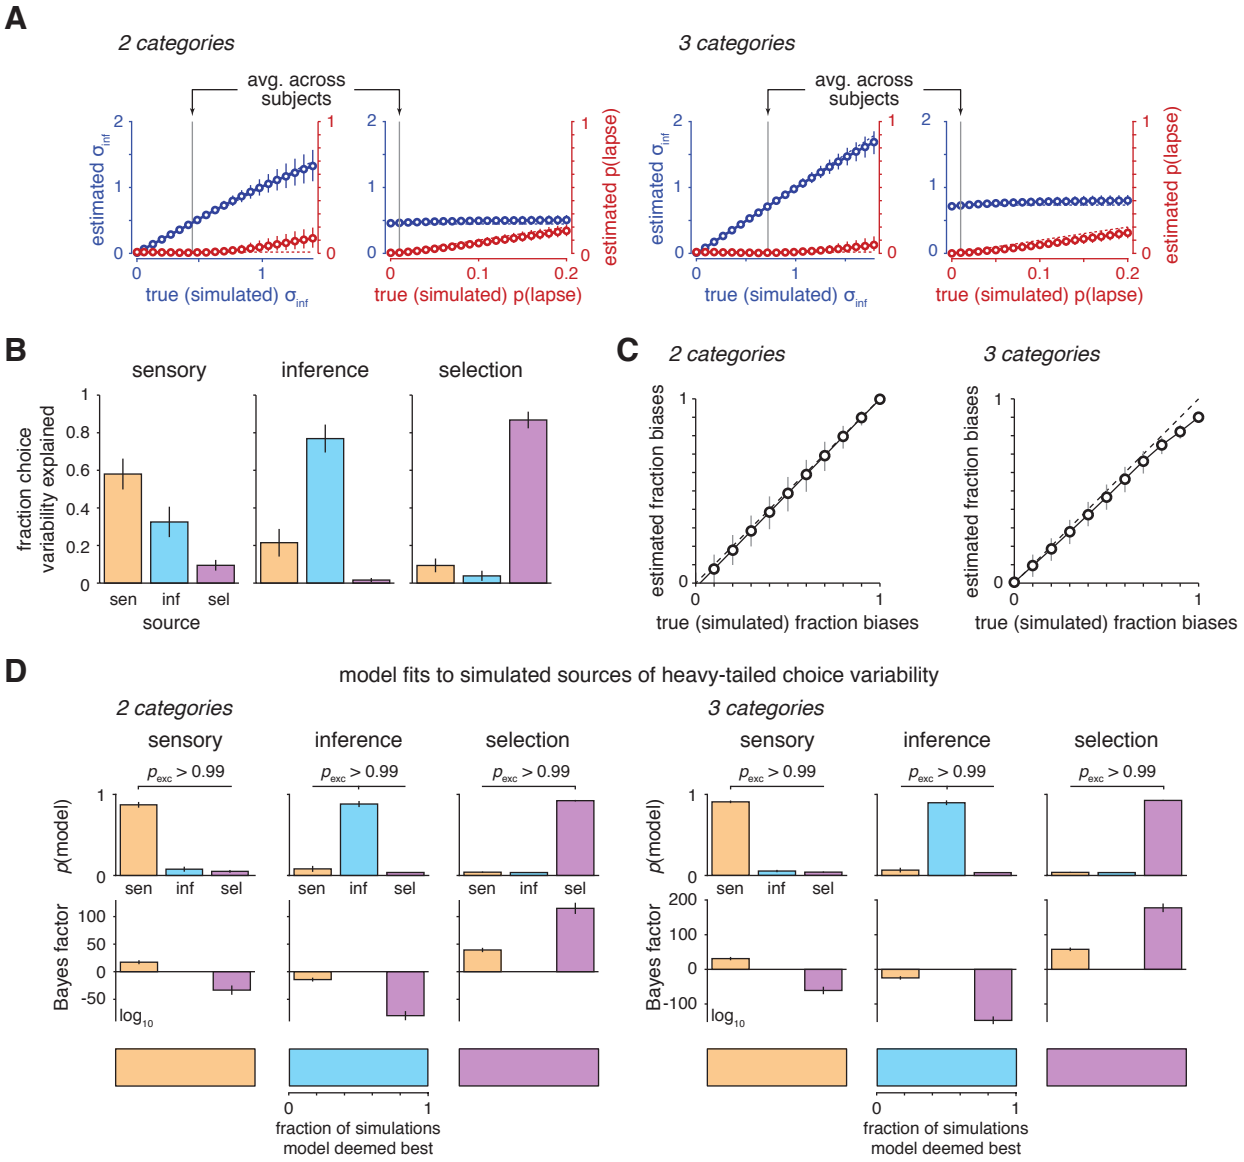

Figure S2. Relates to Figures 3, 5, and 6. Recovering model parameter fits and estimates from simulated behavior. (A) To test how well we could distinguish between inference variability and random lapses, we simulated subject behavior assuming inference variability on the exact same trials that the humans subjects saw in the first experiment. To see how well we recover either of the two parameter, we varied one parameter while keeping the other constant. Per parameter combination we simulated behavior separately for each of the 22 subjects, and computed mean parameter estimates across subjects for 1000 repetitions of this procedure (showing mean, error bars = 2.5th to 97th percentile across repetitions). The dashed lines (mostly below the estimated values) show the true parameter values used for the simulations. For both conditions, there is a "spill-over" between the estimated variability and lapse rate, but this bias is negligible for parameter ranges that we recover from fitting the human subject data (grey lines). (B) To validate the variability decomposition, we applied the same procedure as in (A), but this time performed for each of the 22 subjects three simulations, each introducing variability at a different point in the decision-making process (sensory/inference/selection; indicated by top labels). The magnitude of this variability was set to match the human subjects' observed performance. The results show means and SDs across 1000 repetitions of this procedure, using the exact same procedure as for Fig. 5A in the main text to compute the variability decomposition, but this time on the simulated behavior. (C) To test how well we could distinguish deterministic biases from unstructured variability, we simulated behavior with variability due to inference and deterministic biases according to the model described in Sec. 2.8. The behavior was simulated for the exact same trial sequences observed by the 18 human subjects performing the third experiment. The plots show the estimated average fraction of deterministic biases across subjects (mean, error bars from 2.5th to 97.5th percentile across 1000 repetitions) when fitting our model to simulated behavior. The plots illustrate that, in particular for the range of fractions observed for human subjects (around 0.3), the model correctly recovers the correct fraction. (D) To test how sensitive our model fits are to assuming Gaussian inference variability, we here repeated the analysis of Fig. 3A (main text) described in Sec. 2.9.1, while using heavy-tailed Student's T distributed inference variability with 2.5 degrees of freedom (df) to simulate the behavior, while still assuming Gaussian variability when fitting the different models. $df > 2$ was chosen to have a well-defined mean and variance.

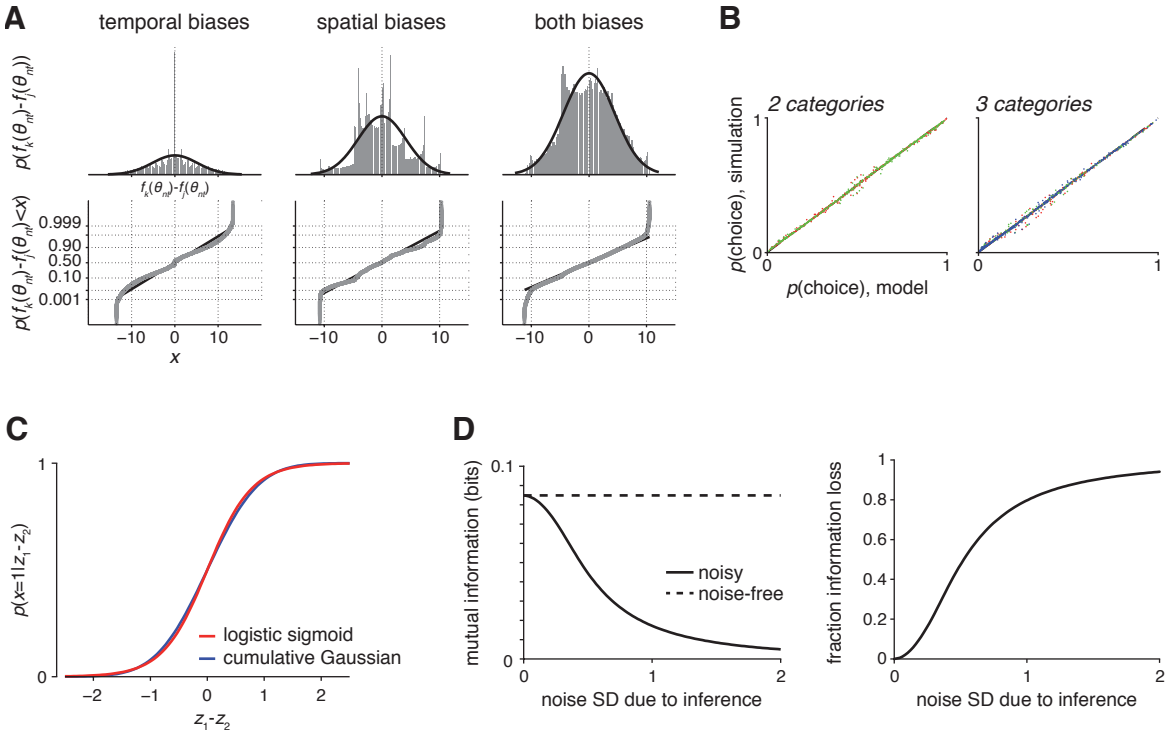

Figure S3. Relates to Figure 6. Illustrating and validating model assumptions and approximations. (A) Here we show by simulation that temporal and spatial biases from Sec. 2.5 cause the bias term differences $f_j(\theta_{nt}) - f_k(\theta_{nt})$ to roughly follow a zero-mean Gaussian distribution, as assumed in in Sec. 2.4. The above shows their empirical distribution, found by simulating 10^5 trials with different sample sequences, as observed by subjects in the three-category condition of the first experiment. For each trial, we computed the contribution of each sample to the log-posteriors while adding the temporal and/or spatial biases discussed in Sec. 2.5, with parameters $\gamma_\theta = 5°$, $\gamma_\kappa = 1$, $\gamma_c = 0.1$ (spatial biases), and $\alpha_a = e^{-0.1}$, $\alpha_b = 0.01$ (temporal biases). Denoting these biased contributions by $\hat{\ell}_{ntk}$ (including their temporal weighting and spatial perturbation) and their unbiased counterparts by ℓ_{ntk}, we found for each sample θ_{nt} in each trial n the bias difference by $f_k(\theta_{nt}) - f_j(\theta_{nt}) = \left(\hat{\ell}_{ntk} - \hat{\ell}_{ntj}\right) - (\ell_{ntk} - \ell_{ntj})$ for each $k \neq j$. The distribution of differences is shown from left to right for temporal biases and spatial biases only, and for both biases in combination. The top panels of (A) show the histogram of differences and the best-fit zero-mean Gaussian (arbitrarily scaled). The bottom panel show the cumulative data distribution in grey, scaled vertically such that the cumulative of a Gaussian becomes a line. The black line connects the 1st and 3rd quantile. The figure shows that, except for the tails, $f_k(\theta_{nt}) - f_j(\theta_{nt})$ is well approximated by a Gaussian. (B) To test if the Gaussian approximation for the sensory variability model did not strongly perturb the derived choice probabilities, we compared the choice probabilities computed with the approximate expressions to those found by simulation. Assuming $\sigma_{sen} = 30°$ ($= \pi/6$ radians), we tested the match on 500 typical trials, by, for each trial, simulating 50.000 instantiations of the sensory noise to get an empirical estimate of these choice probabilities. In each trial we computed one choice probability per category, each corresponding to a different color in the above plots. We deliberately chose a very high noise magnitude, as in this noise regime the Gaussian approximation is more likely to break down. (C) To illustrate the similarity between the logistic sigmoid and the cumulative Gaussian, we here plot the cumulative function of a zero-mean Gaussian with unit variance, and the matching logistic sigmoid with β determined by Eq. (S44). The variance of the Gaussian results from the difference of two Gaussian variables with individual variances $\sigma^2 = \frac{1}{2}$. The latter is the variance used to compute β. (D) The left panel shows the mutual information between the sample log-likelihood ratio and the generative category for the noisy and noise-free log-likelihood ratio for a single stimulus orientation (see Sec. 2.10 for the derivation). In the noise-free case, this mutual information is approximately 0.085 bits per oriented stimulus (1 bit = revealing correct category). In the noisy case, the mutual information drops monotonically with the standard deviation of the noise, matching the noise-free case only for $\sigma_{inf} = 0$. The right panel shows the fraction information loss of the noisy case when compared to the noise-free case. This information loss is computed as one minus the ratio between noisy and noise-free mutual information, as shown in (A). At $\sigma_{inf} = 2$, this loss reaches above 0.9, indicating that at this level of variability, less than 10% of the original information remains.

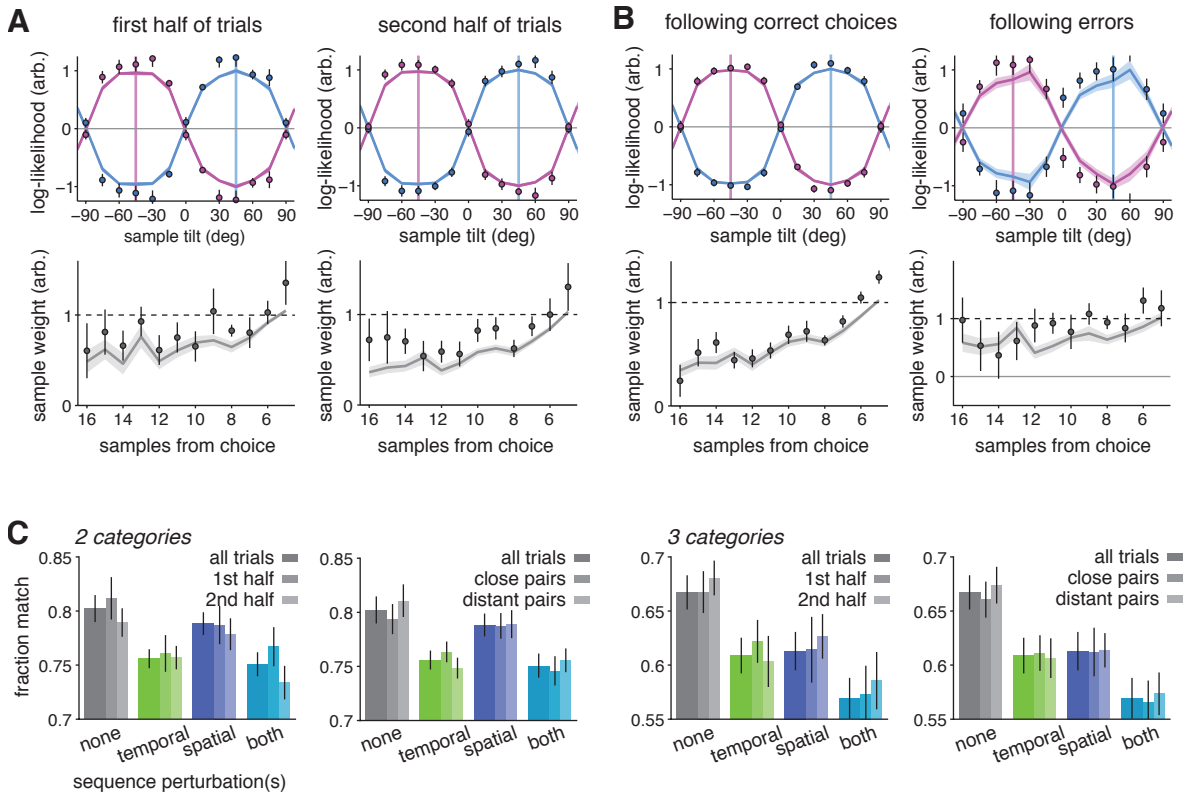

Figure S4. Relates to Figure 7. Deterministic biases and fraction match for different trial subgroups. (A) Left panel: feature (top row) and temporal (bottom row) encoding curves for the two-category condition of exp. 1 in the first half of trials. Right panel: feature (top) and temporal (bottom) encoding curves for the two-category condition of exp. 1 in the second half of trials. (B) Left panel: feature (top row) and temporal (bottom row) encoding curves for the two-category condition of exp. 1 in trials following correct choices. Right panel: feature (top) and temporal (bottom) encoding curves for the two-category condition of exp. 1 in trials following errors. In (A) and (B), dots and error bars indicate human data and s.e.m., and curves and shaded error bars indicate best-fitting model predictions including feature (top) and temporal (bottom) deterministic biases. (C) Fraction matched choices for trial subgroups, experiment 3. The fraction of matched choices ($\pm$ SEM across subjects) are compared for paired trials when computed across all trials (thick bars) to the fraction computed for trial subgroups (thin bars) for different trial pairings (none = exact same sample sequence, temporal = one trial is shuffled version of other, spatial = one trial is mirrored version of other, both = temporal + spatial). On one hand (left panels), we computed this fraction match for the first and second half of trials separately, while excluding trial pairs that spun both trial subgroups. On the other hand (right panels), we split trial pairs into those that appeared closer together within the trial sequence, and those that were more distant. In no case did we find a significant effect of trial sub-grouping on the measured fraction of matched choices (2-way repeated-measures ANOVAs, 1st/2nd half, 2 categories: trial subgroup $F_{1,17} = 1.2, p = 0.3$, perturbation $F_{3,51} = 14.7, p < 0.001$, trial subgroup $\times$ perturbation $F_{3,51} = 1.0, p = 0.4$; 3 categories: trial subgroup $F_{1,17} = 0.1, p = 0.74$, perturbation $F_{3,51} = 18.6, p < 0.001$, trial subgroup $\times$ perturbation $F_{3,51} = 0.8, p = 0.5$; close/distant pairs, 2 categories: trial subgroups $F_{1,17} = 0.3, p = 0.6$, perturbation $F_{3,51} = 27.8, p < 0.001$, trial subgroups $\times$ perturbation $F_{3,51} = 1.4, p = 0.3$; 3 categories: trial subgroups $F_{1,17} = 0.4, p = 0.5$, perturbation $F_{3,51} = 30.6, p < 0.001$, trial subgroups $\times$ perturbation $F_{3,51} = 0.4, p = 0.8$).

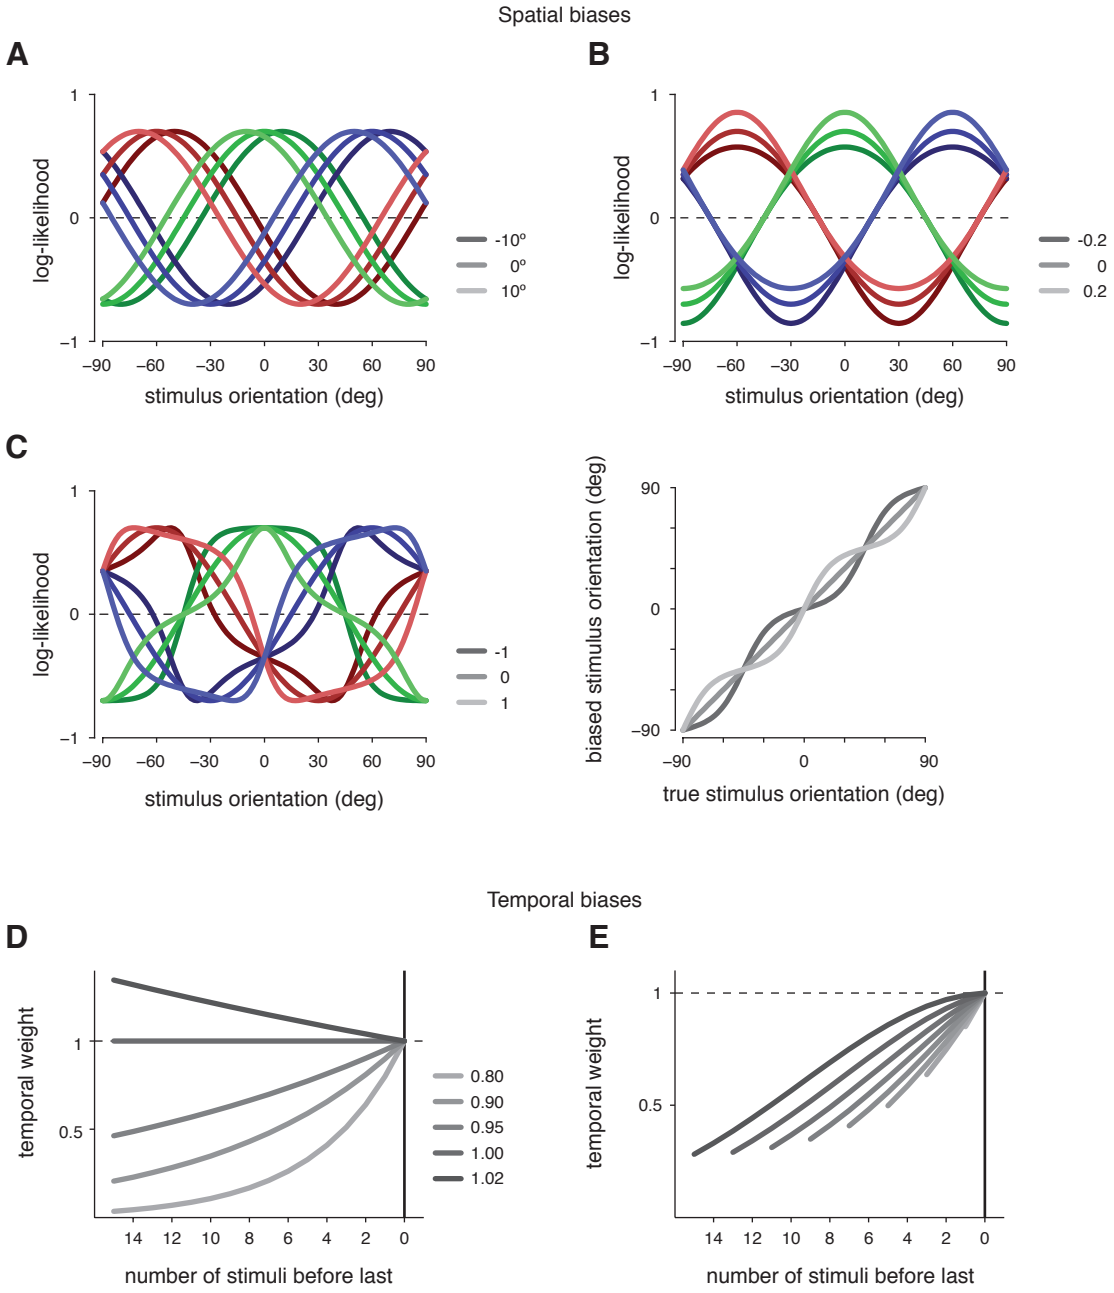

Figure S5. Relates to Figures 6 and 7. Spatial and temporal biases. (A) illustrates the orientation bias for different values of γ_θ, causing a shift in the log-likelihoods. (B) shows the change in log-likelihoods due to the confirmation bias for different values of γ_c. (C) illustrates both the change in log-likelihoods (right panel) and the biased stimulus orientation (left panel) due to the oblique effect for different values of γ_κ. In these panels, the different colors correspond to different categories in the 3-categories condition, and the saturation of these colors to different parameter values. Panels (D) and (E) show the weight λ_{nt} applied to the log-likelihood of samples at different points within the samples sequence, aligned to the last sample in this sequence. (D) shows the temporal bias introduced by time-invariant exponential weighting for different values of the parameter α. As can be seen, $\alpha < 1$ introduces a recency effect, and $\alpha > 1$ a primacy effect. (E) shows the bias for the linearly changing exponential weighting scheme with parameters $\alpha_a = 0.85$ and $\alpha_b = 0.01$. The different shadings correspond to weightings of sequences of different lengths, and illustrate that introducing a linear dependency in the trial-by-trial weighting causes the overall weighting to depend on the sequence length.

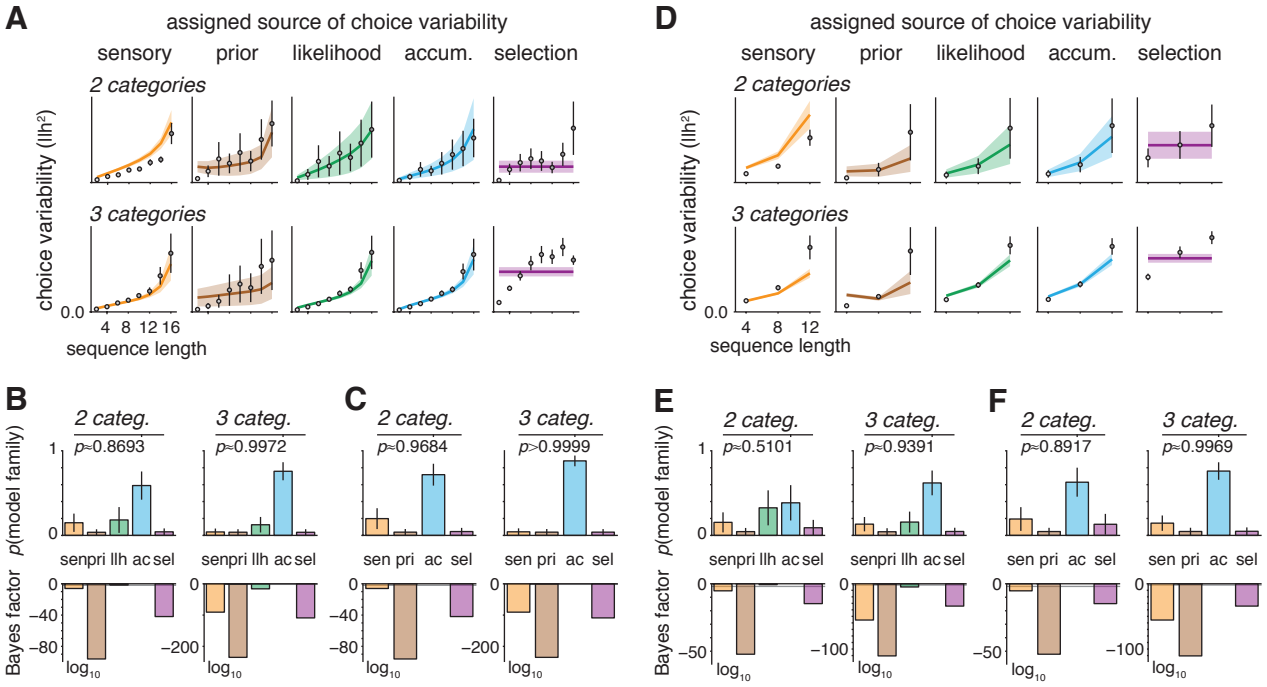

Figure S6. Relates to Figures 3 and 4. Model comparison for different variability models, taking into account temporal and spatial biases. This figure differs from Fig. 3A/B in the main text in that it is based on models with temporal and spatial biases. Panels (A)-(C) shows the comparison based on the data of experiment one, and (D)-(F) based on that of experiment three. Taking into account these biases does not change our main conclusions. We considered all possible bias combinations, resulting in 24 models per variability type. The use of temporal biases allows us to distinguish between prior and selection variability (see Sec. 2.3.4) (A)/(D) Model prediction (grey line and shaded area, mean $\pm$ SEM across subjects) vs. per-sequence length fit (dots with error bars, mean $\pm$ SEM across subjects) of how the noise variance changes with sequence length for five variability models and both task conditions. These predictions are shown for each variability type for the combination of biases that best fitted the subjects' behavior. The larger number of parameters of models with biases resulted in a better per-subject fit, but in a larger across-subject variability in the fitted parameters, as reflected by the larger error bars than in Fig. 3A in the main text. (B)/(E) FFX and RFX model comparison for different models and conditions. The RFX comparison (top) compares model families, where each family features all models of a specific variability type with all possible combinations of spatial and temporal bias (sen = sensory, pri = prior, llh = likelihood, ac = accumulation, sel = selection; mean probability $\pm$ SD). The exceedance probability p is the probability with which the likelihood variability model family is more likely than any other model family. The FFX comparison (bottom) shows for each variability model family the Bayes factor for each family compared to the accumulation variability family. This factor is based on the model within each family with the highest model evidence. The grey line at $2^{10} = 100$ is the threshold at which the evidence in favor of the accumulation model is considered decisive. (C)/(F) Same as (B)/(E), but without the likelihood variability model family. This comparison was included to avoid the sharing of probability mass between too-similar models (Stephan et al., 2009).

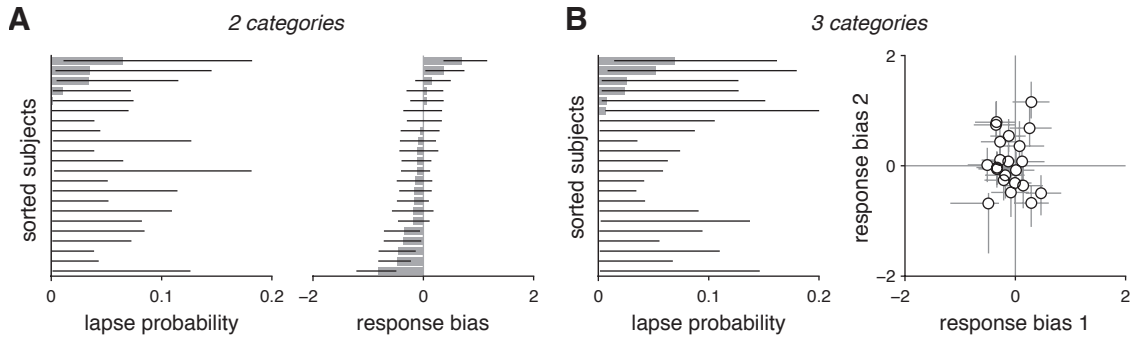

Figure S7. Relates to Figure 3. Lapse and response bias parameters for model assuming inference variability. All panels show parameter modes $\pm$ 95% credible intervals per subject. The subjects are not aligned across panels. (A) Lapses and biases for the 2-categories condition of the first experiment. The left panel shows the lapse probability per subject (mode across subjects not significantly different from zero, $t_{21} = 1.9$, $p = 0.075$). The right panel shows the response bias added to category 1 per subject (mode across subjects not significantly different from zero, $t_{21} = -1.7$, $p = 0.099$). (B) The same as in (A), but for the 3-categories condition. The right panel shows the response biases added to categories 1 and 2 per subject. Across subjects, the mode of only the lapse probability is marginally different from zero (lapse probability, $t_{21} = 2.1$, $p = 0.022$; response bias 1, $t_{21} = -1.5$, $p = 0.14$; response bias 2, $t_{21} = 0.6$, $p = 0.6$).

2 Supplemental Experimental Procedures

Section 2.1 provides additional details about the stimulus and task. Sections 2.2 to 2.5 provide further details about used models. All sections from Section 2.6 onwards provide more in-depth mathematical details, but are not required reading to understand the essence of the models.

2.1 Stimulus and task design details

Stimuli were high-contrast, noise-free Gabor patterns (diameter: 4 degrees of visual angle, spatial frequency: 2 cycles per degree of visual angle, Michelson contrast: 75%) of varying orientation, presented at fixation for 100 ms at an average stimulation rate of 3 Hz with small amount of uniform jitter (± 33 ms). Sequence lengths (i.e., the number of stimuli per sequence) varied uniformly and unpredictably from 2 to 16 stimuli.

Each Gabor pattern was presented on top of a luminance pedestal, and each sequence began with two luminance pedestals presented in rhythm with the following Gabor patterns so that the onset of the first stimulus was predictable in time. At sequence offset, participants were prompted by a luminance pedestal in the fixation point (go signal) for a choice regarding the most likely generative category of the sequence. Participants provided their response by pressing one out of two (or three, depending on the condition) keys with their right hand. The mapping between categories (represented as colors) and response keys was fixed and explained to the participant before the start of the experiment. If no response was provided 1 s following the go signal, the trial was aborted and the participant was informed by a beep that his response was too slow which represented less than 1% of trials for all tested participants.

Following each response, feedback about the true generative category of the sequence was provided via a transient change in the color of the fixation point (e.g., from grey to pink if the sequence was generated from the pink category). We did not provide feedback as to whether participants picked the most likely category given the limited evidence provided by the sequence. Consequently, as in the weather prediction task, participants accuracy was bounded by the randomness of presented sequences around their generative means. Furthermore, this means that participants could learn not only the means, but also the spreads of generative distributions.

2.2 Optimal decision making without added variability

2.2.1 Generative model

The generative model of the task is as follows. In each trial n of N independent trials, the experimenter picks one category $k \in \{1, \dots, K\}$ (called *decks* in the main text; $K \in \{2, 3\}$ for 2 or 3 categories) with associated category mean μ_k. Based on this mean, T_n orientations, $\theta_{n1}, \dots, \theta_{nT_n}$, called *samples* (or *cards* in the main text), are generated by drawing them independently and identically distributed (i.i.d.) from a von Mises distribution over the half-ciricle $[0, \pi)$, centered on μ_k and with concentration κ. That is, each sample is independently drawn from

$$p\left(\theta_{nt}|\mu_k\right) = \frac{e^{\kappa \cos(2(\theta_{nt} - \mu_k))}}{\pi I_0(\kappa)}, \tag{S1}$$

where $I_0(\cdot)$ is the modified Bessel function of order 0.

2.2.2 Optimal decision-making

To derive the optimal decision-making strategy, we assume a uniform prior over categories, and a 0-1 loss function, charaterized by a gain (loss) of 0 (1) for correct (incorrect) decisions. For such a loss structure, it is optimal to choose the option associated with the most likely generative category (Berger, 1993). To find this most likely category, we assume subjects to know the category means, $\mu_1, \dots, \mu_K$, and the concentration parameter κ of the generative density, Eq. (S1). Based on this and the uniform category prior, $p(\mu_k) \propto 1$, the posterior probability of category k having generated $\theta_{n1:T_n} = \{\theta_{n1}, \dots, \theta_{nT_n}\}$ in trial n is by Bayes rule,

$$p(\mu_k|\theta_{n1:T_n}) \propto p(\theta_{n1}, \dots, \theta_{nT_n}|\mu_k) = \prod_{t=1}^{T_n} p(\theta_{nt}|\mu_k). \tag{S2}$$

If x_n denotes the chosen category in trial n, then optimal decision-making is performed by choosing the $x_n = k$ for which the above posterior is a maximum, that is $x_n = \text{argmax}_k\, p(\mu_k|\theta_{n1:T_n})$.

The optimal strategy is implemented incrementally with each additional sample by tracking the log-posteriors z_{ntk} for all k's in trial n after the tth sample by

$$z_{ntk} = z_{n,t-1,k} + \ell_{ntk}, \tag{S3}$$

where, initially, $z_{n0k} = 0$, and ℓ_{ntk} is the decision-relevant component of the unnormalized log-likelihood for the tth sample,

$$\ell_{ntk} = \log p\left(\theta_{nt}|\mu_k\right) + \text{const.} = \kappa \cos\left(2(\theta_{nt} - \mu_k)\right). \tag{S4}$$

With these log-postieriors, optimal decision-making in trial n is achieved by choosing

$$x_n = \underset{k}{\text{argmax}}\, z_{nT_nk}. \tag{S5}$$

2.3 Introducing variability

Before discussing individual models of variability, we relate two different general models of variable choices. Assume K log-posteriors $z_1, \ldots, z_K$, to each of which we add i.i.d. zero-mean noise ε_k, with $\langle \varepsilon_k \rangle = 0$ for all k. Choices are again based on picking the largest, but this time noise-perturbed, log-posterior, that is

$$x = \underset{k}{\text{argmax}} \left(z_k + \varepsilon_k \right), \tag{S6}$$

With this choice rule, the probability of picking category k becomes (Fig. 1D in main text)

$$p\left(x = k | z_{1:K}\right) = \int p\left(\forall j \neq k : z_j + \varepsilon_j < z_k + \varepsilon_k | z_{1:K}, \varepsilon_{1:K}\right) \prod_k p(\varepsilon_k) \mathrm{d}\varepsilon_{1:K}. \tag{S7}$$

For Gaussian noise the above choice distribution become the cumulative function of a unimodal ($K = 2$) or bimodal ($K = 3$) Gaussian. If the noise is Gumbel-distributed instead, the choice distribution is given by a logistic sigmoid ($K = 2$), or its multidimensional generalization, the softmax function ($K = 3$) (see Sec. 2.6).

From the empirical point-of-view, the choice distributions resulting from either Gaussian or Gumbel-distributed noise are barely distinguishable. Specifically, a choice distribution resulting from Gumbel-distributed noise with scale β^{-1} will appear like one resulting from Gaussian noise with variance $\sigma^2 \approx \pi^2/(6\beta^2)$ (see Sec. 2.6). Thus, we can use these two distributions interchangeably when modeling the subjects' choices. We use this property for two purposes. First, we can relate choice predictions from different models of variability, even if they predict different forms of the choice distribution. Second, we will fit models using the computationally simpler softmax choice function even for models that assume Gaussian noise.

2.3.1 Variability at the selection stage

A possible source of variability in the subjects' choices is that, at the selection stage, they perform these choices by drawing samples from their belief about the correctness of either choice. Such sampling corresponds to choosing option k in trial n with probability

$$p\left(x_n = k | \theta_{n1:T_n}\right) \propto p\left(\mu_k | \theta_{n1:T_n}\right)^\beta \propto \frac{e^{\beta z_{nT_n k}}}{\sum_j e^{\beta z_{nT_n j}}} \tag{S8}$$

where we have used Eqs. (S2) and (S3). With $\beta = 1$, the above constitutes strict posterior sampling. We consider a slightly more general form by allowing β to take any non-negative value, allowing from completely random choices ($\beta = 0$), over strict posterior sampling ($\beta = 1$), to optimal decision making without any added variability ($\beta \to \infty$, leading to Eq. (S5)). In either case, the choice distribution corresponds to the softmax function with fixed inverse temperature β.

This model predicts that the magnitude of the variability added to each of the log-posteriors does not depend on the sequence length T_n. This is because the softmax parameter β in Eq. (S8) is independent of the sequence length T_n that resulted in each of the log-posteriors, $z_{nT_n 1}, \ldots, z_{nT_n K}$. By the relation between the softmax and Gaussian cumulative density (see Sec. 2.6), β can be translated into the Gaussian noise variance $\sigma^2 \approx \pi^2/(6\beta^2)$, which will also be independent of T_n. Thus, mechanistically, the above choice rule can be implemented by adding a single zero-mean Gaussain noise term with variance σ^2 to each of the K log-posteriors.

2.3.2 Variability at the inference stage

We consider two possibilities of how the inference process introduces variability in the log-posteriors. The *likelihood variability* model perturbs each log-likelihood by adding zero-mean Gaussian noise. The *accumulation variability* model assumes noise to affect the accumulation of evidence itself, and is implemented by

adding zero-mean Gaussian noise as soon as a new log-likelihood is added to the current log-posterior. In the main text, both types of variability are discussed, but most inference variability models only use accumulation variability, which is more strongly supported by Bayesian model comparison.

The difference between these two models of variability is subtle, and best understood by considering a trial in which a choice is made after observing a sequence of only two samples, in which case each log-posterior equals the sum of two log-likelihoods. In the likelihood variability model variability is added to each of these log-likelihoods, such that each log-posterior is perturbed by two noise terms. In the accumulation variability model, in contrast, only a single noise term perturbs each log-posterior, as only a single sum has been performed. Here, we consider forming the log-posterior after the first sample as initialization of the posterior rather than a sum. In general, the likelihood variability model will always add one more noise term to the log-posteriors than the accumulation noise model. Thus, the difference in prediction between these two models will be most pronounced for short sequence lengths, and they are hard to tell apart in general. Therefore, we refer to both types of variability by the umbrella term 'inference' variability.

Formally, the variability for both models follows zero-mean Gaussian noise, $\varepsilon_{ntk} \sim \mathcal{N}(0, \sigma^2_{inf})$, that is independent across trials n, samples t in the sequence, and categories k. For the likelihood variability model, this noise is added to each likelihood,

$$\hat{\ell}_{ntk} = \ell_{ntk} + \epsilon_{ntk}, \tag{S9}$$

where $\hat{\ell}_{ntk}$ denotes the noise-perturbed log-likelihood. As the log-posterior with respect to each category sums up these noisy log-likelihoods, it is distributed as

$$p\left(z_{nT_nk}|\theta_{n1:T_n}\right) = \mathcal{N}\left(z_{nT_nk}|\sum_{t=1}^{T_n}\ell_{ntk}, T_n\sigma^2_{inf}\right). \tag{S10}$$

As can be seen, the noise only influences the log-posterior variance, which scales linearly with sequence length T_n.

The accumulation variability model adds one less noise term to the log-posterior. Thus, for this model, each log-posterior is distributed as

$$p\left(z_{nT_nk}|\theta_{n1:T_n}\right) = \mathcal{N}\left(z_{nT_nk}|\sum_{t=1}^{T_n}\ell_{ntk}, (T_n-1)\sigma^2_{inf}\right). \tag{S11}$$

Both models predict that the magnitude of the variability added to each of the log-posteriors increases linearly with sequence length.

2.3.3 Variability at the sensory stage

We assume that variability at the sensory stage perturbs each sensory percept θ_{nt} by additive zero-mean Gaussian noise $\varepsilon_{nt} \sim \mathcal{N}(0, \sigma^2_{sen})$, after which the noisy θ_{nt} is re-mapped onto its original half-circular domain $[0, \pi)$ by modular arithmetic. This results in the noisy log-likelihoods to be given by

$$\hat{\ell}_{ntk} = \kappa \cos\left(2(\theta_{nt} + \varepsilon_{nt} - \mu_k)\right). \tag{S12}$$

As for likelihood or accumulation variability models, the magnitude of sensory variability increases with sequence length T_n.

Adding Gaussian noise to the orientation percept rather than the log-likelihoods has the following effects. First, as the Gaussian noise is passed through a non-linearity (the cosine in the log-likelihood function), its distribution won't be Gaussian in the log-posterior. Second, as the log-likelihoods for different categories k are all affected by the same noise term ε_{nt}, their variability due to this noise will be highly correlated. These correlations will only be apparent for the 3-category condition in which choices are determined by two log-posterior differences rather than only one. In the 2-category condition, the only factor

that distinguishes sensory variability from likelihood or accumulation variability is the structure of the variability in the log-likelihoods (see Sec. 2.7.4).

To fit models that assume variability at the sensory stage, we deviate from Eq. (S7), and instead approximate the densities of the relevant log-likelihood differences by matching the moments of a multivariate Gaussian. See Section 2.7 for the resulting expressions.

2.3.4 Variability at the prior stage

Given that in each trial a series of log-likelihoods need to be accumulated to form a log-posterior, another source of variability in this log-posterior might be in the initial state of the accumulator. Assuming this variability to be i.i.d. zero-mean Gaussian noise, $\varepsilon_k \sim \mathcal{N}\left(0, \sigma_{pri}^2\right)$ for the accumulator of each category k, this results in the log-posteriors to be distributed as

$$p\left(z_{nT_nk}|\theta_{n1:T_n}\right) = \mathcal{N}\left(z_{nT_nk}|\sum_{t=1}^{T_n}\ell_{ntk}, \sigma_{pri}^2\right), \tag{S13}$$

with variability independent of sequence length. Perfect accumulation of evidence makes this model indistinguishable from one that assumes variability at the selection stage. These two models can only be distinguished if subjects feature temporal biases (to be introduced in Sec. 2.5.2), which cause them to weight noise early in the sequence differently from noise late in the sequence (see Fig. S6).

2.3.5 Variability at multiple stages

In order to model variability at multiple stages of the decision process, we note that the variability at each stage is well captured by log-posteriors, z_{nT_nk}, or their difference between categories, that are perturbed by (potentially correlated) Gaussian noise. Thus, variability introduced at multiple stages corresponds to Gaussian noise with a covariance matrix that is the sum of the covariance matrices corresponding to the individual stages. The log-posterior means are usually unperturbed, expect for when some of the variability is introduced at the sensory stage (see Sec. 2.7). In this case, these means are the ones from the sensory variability model, with covariances that again sum up across all stages.

2.3.6 Misspecification of subjective generative concentration κ

All models described so far were fitted using an implicit scaling parameter – corresponding to the subjective concentration κ of the generative distributions of orientation, set at its true value ($\kappa = 0.5$ in the two-category condition, $\kappa = 0.7$, in the three-category condition). In other words, we have expressed all noise parameters in the models with respect to these true theoretical values. Although this use of a scaling parameter does not affect any of our main conclusions about the type of variability (sensory, inference or response selection) which causes choice suboptimality, we discuss below how parameter estimates in the different noise models change if subjects did not estimate correctly κ – that is, the spread of the generative distributions. For simplicity, we consider the bias-free case, even though the same principles apply if such biases are included.

For the 2-categories condition, the choice in trial n is fully determined by the log-posterior difference $z_{nT_n2} - z_{nT_n1}$, which by Eqs. (S4), (S10), and Eqs. (S11) is distributed as

$$z_{nT_n2} - z_{nT_n1}|\theta_{n1:T_n} \sim \mathcal{N}\left(\kappa\sum_{t=1}^{T_n}\left(\cos\left(2(\theta_{nt}-\mu_2)\right) - \cos\left(2(\theta_{nt}-\mu_1)\right), 2S_n\sigma_{inf}^2\right)\right), \tag{S14}$$

with $S_n = T_n$ and $S_n = T_n - 1$ for the likelihood and accumulation variability models, respectively. The probability of choosing option $x_n = 1$, which occurs if $z_{nT_n2} - z_{nT_n1} < 0$, is thus given by

$$p(x_n = 1|\theta_{n1:T_n}) = \Phi\left(\frac{\kappa}{\sqrt{\sigma_{inf}^2}}\frac{\sum_{t=1}^{T_n}\left(\cos\left(2(\theta_{nt}-\mu_2)\right) - \cos\left(2(\theta_{nt}-\mu_1)\right)\right)}{\sqrt{2S_n}}\right), \tag{S15}$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard Gaussian. For the 3-categories condition, the choice probability has a different form due to the larger number of possible choices. What remains the same, and what is essential here, is that κ in this choice probability again appears as the fraction $\frac{\kappa}{\sqrt{\sigma_{inf}^2}}$. Thus, the magnitude of the estimated σ_{inf}^2 depends on the chosen κ.

What happens if the subject's assumed $\tilde{\kappa}$ differs from the κ we have used to estimate σ_{inf}^2? If subjects have over-estimated $\tilde{\kappa} > \kappa$ then we have underestimated σ_{inf}^2, as in this case $\sigma_{inf}^2 = \frac{\kappa^2}{\tilde{\kappa}^2}\tilde{\sigma}_{inf}^2 < \tilde{\sigma}_{inf}^2$, where $\tilde{\sigma}_{inf}^2$ is the true noise variance featured by the subject. On the other hand, if subjects have under-estimated $\tilde{\kappa}$, then our estimated σ_{inf}^2's are too large. None of this affects our main conclusions about the type of variability, which depends on variability structure and scaling with sequence length rather than absolute magnitudes.

The assumed size of κ only has an effect when comparing variability magnitudes between conditions. In the main text we discuss that the estimated inference variability variance per category in the 3-categories condition is significantly larger than that in the 2-categories condition. This comparison relies on assuming that subjects use $\kappa_2 = 0.5$ and $\kappa_3 = 0.7$ in the 2-categories and 3-categories condition, respectively. However, even if subjects have assumed $\tilde{\kappa}_2 = \tilde{\kappa}_3$, and we adjust our estimates for this assumption by $\tilde{\sigma}_{inf,3} = \frac{\kappa_2}{\kappa_3}\sigma_{\mathrm{inf},3}$, the variability per category in the 3-categories condition is still significantly larger than in the 2-categories condition (two-sided, paired $t_{21} = 2.4$, $p = 0.027$). In fact, the difference remains significant at the 0.05 level up to $\tilde{\kappa}_3 \approx 0.70\tilde{\kappa}_2$ (variability per category) or $\tilde{\kappa}_3 \approx 0.58\tilde{\kappa}_2$ (total variability). Thus, as long as subjects did not assume the likelihood in the 3-categories condition to be less concentrated than in the 2-categories condition, they featured significantly larger per-category variability in the 3-categories condition.

The question of a mis-specified κ disappears if we focus on sensory instead of inference variability. In this case, κ acts as a multiplicative factor in the noisy log-likelihood, Eq. (S12), and thus scales mean and standard deviation of the noisy log-posterior difference equivalently. This causes κ to cancel out when taking the ratio of these quantities, as is done to predict the sensory variability choice probability (see Sec. 2.7 for equations). As a result, the choice probabilities predicted by the sensory variability model are insensitive to any mis-specification of κ. This fact, together with the observation that we find the sensory variability magnitude to grow when moving from 2 to 3 categories (see main text), further supports our claim that the variability magnitude grows with the difficulty of the task.

2.4 Decomposing variability into deterministic biases and residual terms

So far, variability at the inference stage was assumed to be to be additive in the log-posteriors. Here, we provide a more fine-grained description of its structure by splitting the noise terms into a deterministic bias and a residual variability component. Specifically, $\hat{\ell}_{ntk}$ denotes the internal estimate of ℓ_{ntk}, which is composed of

$$\hat{\ell}_{ntk} = \ell_{ntk} + f_k(\theta_{nt}) + \varepsilon_{ntk}, \tag{S16}$$

where $f_k(\cdot)$ is a deterministic, but unknown, function of the perceived orientation θ_{nt}, or a *deterministic bias*, and ε_{ntk} is zero-mean Gaussian noise, or *unstructured variability*. All of what is described below applies also if the deterministic component is a function of the whole sample sequence rather than only single samples. To keep the presentation simple, the derivation is only shown for the single-sample case.

For the 2-category case, bias $f_k(\theta_{nt})$ and residual variability ϵ_{ntk} are across trials n and samples t assumed to be distributed as

$$f_2(\theta_{nt}) - f_1(\theta_{nt}) \sim \mathcal{N}\left(0, 2\sigma_b^2\right), \tag{S17}$$

$$\epsilon_{ntk} \sim \mathcal{N}\left(0, \sigma_v^2\right), \qquad k = 1, 2. \tag{S18}$$

For $f_2(\theta_{nt}) - f_1(\theta_{nt})$, this distribution is induced by the randomization of θ_{nt} across trials and samples (see Fig. S3A for a justification of the Gaussian assumption). We do not need to specify the distribution of individual $f_k(\theta_{nt})$'s, as only their differences matter for the choice distributions. For the 3-category case we make conceptually similar assumptions, but the derivations are more burdensome and are provided in detail in Sec. 2.8.2.

With the above, the log-likelihood difference estimates are Gaussian with moments

$$\left\langle \hat{\ell}_{nt2} - \hat{\ell}_{nt1} \right\rangle = \ell_{nt2} - \ell_{nt1} \tag{S19}$$

$$\mathrm{var}\left(\hat{\ell}_{nt2} - \hat{\ell}_{nt1}\right) = \mathrm{var}\left(f_2(\theta_{nt}) - f_1(\theta_{nt})\right) + \mathrm{var}\left(\varepsilon_{nt2} - \varepsilon_{nt1}\right) = 2(\sigma_b^2 + \sigma_v^2). \tag{S20}$$

Setting $\sigma_b^2 + \sigma_v^2 = \sigma_{inf}^2$, and using $z_{nT_nk} = \sum_{t=1}^{T_n} \hat{\ell}_{ntk}$ recovers our original likelihood variability formulation, Eqs. (S9) and (S10).

In the above we have assume additive noise to each of the T_n log-likelihoods, just like for the likelihood variability model. In the corresponding accumulation variability model, we only add $T_n - 1$ bias and residual variability terms. For what follows, we for simplicity assume variability in the likelihoods. The same concepts apply for accumulation variability.

2.4.1 The bias-variance decomposition

To see the influence of the bias terms, $f_k(\cdot)$, consider the case in which multiple trials feature the exact same sample sequence. In this case, the tth sample θ_{nt} and the associated $f_k(\theta_{nt})$ and $f_j(\theta_{nt})$ terms are the same across all these trials, while the ε_{ntk} terms vary. Then, the variance of the log-likelihood difference estimate

around its true value decomposes into

$$
\begin{aligned}
&\mathrm{var}\left(\hat{\ell}_{ntk} - \hat{\ell}_{ntj} | f_k(\theta_{nt}), f_j(\theta_{nt})\right) \\
&\quad = \left\langle \left((\ell_{ntk} - \ell_{ntj}) - \left(\hat{\ell}_{ntk} - \hat{\ell}_{ntj}\right)\right)^2 \right\rangle_{p\left(\hat{\ell}_{ntk} - \hat{\ell}_{ntj} | f_k(\theta_{nt}), f_j(\theta_{nt})\right)} \\
&\quad = (\ell_{ntk} - \ell_{ntj})^2 - 2\left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle (\ell_{ntk} - \ell_{ntj}) + \left\langle \left(\hat{\ell}_{ntk} - \hat{\ell}_{ntj}\right)^2 \right\rangle \\
&\quad = \left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle^2 - 2\left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle (\ell_{ntk} - \ell_{ntj}) + (\ell_{ntk} - \ell_{ntj})^2 \\
&\qquad + \left\langle \left(\hat{\ell}_{ntk} - \hat{\ell}_{ntj}\right)^2 \right\rangle - 2\left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle \left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle + \left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle^2 \\
&\quad = \left(\left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle - (\ell_{ntk} - \ell_{ntj}) \right)^2 + \left\langle \left(\left(\hat{\ell}_{ntk} - \hat{\ell}_{ntj}\right) - \left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle \right)^2 \right\rangle \qquad \text{(S21)} \\
&\quad = \left(f_k(\theta_{nt}) - f_j(\theta_{nt})\right)^2 + \left(\mathrm{var}(\varepsilon_{ntk}) + \mathrm{var}(\varepsilon_{ntj})\right), \qquad \text{(S22)}
\end{aligned}
$$

where all expectations are implicitly conditional on $f_k(\theta_{nt})$ and $f_j(\theta_{nt})$. The third equality is based on adding and subtracting $2\left\langle \hat{\ell}_{ntk} - \hat{\ell}_{ntj} \right\rangle^2$, and the last equality uses the above definition of $\hat{\ell}_{ntk}$, Eq. (S16). In Eq. (S21), the first term in the sum is the square distance between the mean log-likelihood difference estimate and its true value. Thus, it is a measure of the estimate's deterministic bias. The second term is the unstructured variability of the log-likelihood difference estimate around its mean, and thus measures the estimate's variance, irrespective of its bias. Therefore, this decomposition is commonly known as the *bias-variance decomposition* (Bishop, 2006). Re-expressing these two measures in terms of the decomposition of $\hat{\ell}_{ntk}$ in Eq. (S22) confirms that $f_k(\theta_{nt})$ introduces bias, whereas ε_{ntk} introduces variance.

2.4.2 Estimating the contribution of deterministic biases to the overall variability

To estimate the contribution of deterministic biases to the overall variability, we fit a model that contains this contribution as an explicit parameter. We do so by grouping trials m and n in which the same sample sequence has been shown to the subjects, that is, in which $\theta_{n1:T_n} = \theta_{m1:T_m}$, and by modeling the choice probabilities in both trial in combination. Here, we only discuss the 2-categories condition. Handling the 3-categories condition is conceptually similar, but mathematically more burdensome due to having to specify the posterior over four log-posterior differences simultaneously. The mathematical details for both the 2 and the 3-categories condition are provided in Sec. 2.8.

In the 2-categories condition, the (noise-perturbed) log-posterior differences $z_{nT_n2} - z_{nT_n1}$ and $z_{mT_m2} - z_{mT_m1}$ fully determine the choices in trials n and m. Their joint probability is given by

$$
\begin{aligned}
&p\left(z_{nT_n2} - z_{nT_n1}, z_{mT_m2} - z_{mT_m1}\right) \\
&\quad = \mathcal{N}\left(\begin{pmatrix} z_{nT_n2} - z_{nT_n1} \\ z_{mT_m2} - z_{mT_m1} \end{pmatrix} \middle| \begin{pmatrix} \sum_{t=1}^{T_n} (\ell_{nt2} - \ell_{nt1}) \\ \sum_{t=1}^{T_n} (\ell_{nt2} - \ell_{nt1}) \end{pmatrix}, 2T_n\sigma_{inf}^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right), \qquad \text{(S23)}
\end{aligned}
$$

where we have used $\sigma_{inf}^2 = \sigma_b^2 + \sigma_v^2$, and $\rho = \sigma_b^2/\sigma_{inf}^2$ is the fraction of variance that deterministic biases contribute to the overall noise variance. The correlation ρ between the log-posterior differences is introduced by the bias terms shared by trials n and m. Without these bias terms, the differences would be completely uncorrelated. These shared terms boost the probability of performing the same choice in both grouped trials (see Fig. 6A in main text), irrespective of its correctness.

The above defines our recipe for estimating to which degree deterministic biases contribute to the overall noise variance. We do so by fitting the full model to the behavior of each subject for each condition separately to find the parameters that best explain this behavior (see Methods for the fitting procedure). However, instead of modeling each trial in isolation, we group trials in which the same sample sequence

has been shown, and model their choices jointly by use of Eq. (S23). This allows us to estimate ρ, which directly quantifies the amount of bias contribution. A similar procedure, described in Sec. 2.8, leads to the ρ-estimates for the 3-categories condition.

2.5 Introducing explicit temporal and spatial biases

We introduce biases both on the way the sensory percepts are converted into log-likelihoods, as well as on how these log-likelihoods are accumulated over time. We refer to the first kind of biases as *spatial biases*, and to the second kind as *temporal biases*. These two types of biases do not interact directly, and so can be implemented completely separately by first applying the spatial biases to compute the log-likelihoods, and then the temporal biases to find the log-posterior predictions and associated choice distribution.

2.5.1 Spatial biases

All spatial biases influence the mapping from sensory percept to log-likelihoods, but different biases deal with different aspects of this mapping. If several biases are combined, they are applied in the order they are presented below.

Orientation bias. The orientation bias causes the stimulus orientation to be perceived with a bias γ_θ (Fig. S5A). With this biases, the log-likelihoods is given by

$$\ell_{ntk} = \kappa \cos\left(2(\theta_{nt} + \gamma_\theta - \mu_k)\right), \tag{S24}$$

instead of Eq. (S4).

Oblique effect. The oblique effect causes a bias of shifting the observed orientations towards the cardinal directions, or away from them (Fig. S5C). It is realized conceptually by applying the cumulative of a von Mises distribution on $[0, \pi/2)$ and $[\pi/2, \pi)$, such that the biased orientation is perceived as

$$\tilde{\theta} = \frac{\pi}{2} \int_0^{4\theta} \text{VonMises}\left(a|0, \gamma_\kappa\right) \mathrm{d}a, \tag{S25}$$

where γ_κ determines the strength of this bias.

Specified as above, γ_κ needs to be non-negative by definition, such that only biases towards the diagonal directions are allowed. To also support biases towards cardinal directions, we use the series expansion of the cumulative,

$$\tilde{\theta} = \theta + \frac{1}{2I_0(\gamma_\kappa)} \sum_{j=1}^{\infty} \frac{\sin\left(4\theta j\right) I_j(\gamma_\kappa)}{j}, \tag{S26}$$

where $I_j(\cdot)$ is the modified Bessel function of order j. This series expansion allows for negative γ_κ's and thus a bias towards cardinal directions. In practice, we approximate the infinite series by its first ten terms.

Confirmation bias. The last bias over-emphasizes high likelihoods and thus introduces a bias towards over-weighting the category supported by the current sample and under-weighting the other categories (Fig. S5B). Its implementation is based on scaling the generative log-likelihood, Eq. (S4) by its exponential, that is

$$\ell_{ntk} = \kappa \cos\left(2(\theta_{nt} - \mu_k)\right) e^{\gamma_c \cos(2(\theta_{nt} - \mu_k))}. \tag{S27}$$

Here, $\gamma_c > 0$ causes a bias towards confirming the supported category, $\gamma_c < 0$ causes a bias away from it, and $\gamma_c = 0$ leaves the log-likelihood unchanged.

2.5.2 Temporal biases

Temporal biases are introduced by modulating the degree to which each likelihood contributes to the posterior upon which the decision is based. This allows us to implement both recency and primacy effects.

Formally, we denote the weight on the tth noisy log-likelihood $\hat{\ell}_{ntk}$ in trial n by λ_{nt}, such that the log-posterior with respect to category k is given by

$$z_{nT_nk} = \sum_{t=1}^{T_n} \lambda_{nt} \hat{\ell}_{ntk}. \tag{S28}$$

This weighting has the following effect on the log-posteriors. If we only consider single trials, and assume likelihood variability as in Eq. (S9), then the log-posteriors that include the temporal bias are distributed as

$$p\left(z_{nT_nk} | \theta_{n1:T_n}\right) = \mathcal{N}\left(z_{nT_nk} | \sum_{t=1}^{T_n} \lambda_{nt} \ell_{ntk}, \sigma_{inf}^2 \sum_{t=1}^{T_n} \lambda_{nt}^2\right). \tag{S29}$$

As can be seen, the weights also appear in the variance term. This is due to these weights scaling both the log-likelihood and the additive Gaussian noise.

Other models feature a similar mean of the log-posterior, but a different variance. For the accumulation variability model we assume that the temporal weight is applied before new evidence is added, such that the log-posterior is distributed as for the likelihood variability model, only with its variance replaced by $\sigma_{inf}^2 \sum_{t=2}^{T_n} \lambda_{nt}^2$. Thus, the only change is that λ_{n1} does not modulate the variability. In the prior variability model, noise only influences the first log-likelihood, such that its associated log-posteriors are distributed again as for the likelihood variability model, but with variance $\sigma_{pri}^2 \lambda_{n1}^2$. The model that assumes variability at the selection stage remains unchanged, as by assumption, $\lambda_{nT_n} = 1$, always.

If we consider trials groups, n and m, in which the same sample sequence was presented, then the log-posterior joint density for the 2-categories condition changes from Eq. (S23) to

$$\begin{aligned} &p\left(z_{nT_n2} - z_{nT_n1}, z_{mT_m2} - z_{mT_m1}\right) \\ &\quad = \mathcal{N}\left(\begin{pmatrix} z_{nT_n2} - z_{nT_n1} \\ z_{mT_m2} - z_{mT_m1} \end{pmatrix} \middle| \begin{pmatrix} \sum_{t=1}^{T_n} \lambda_{nt} \left(\ell_{nt2} - \ell_{nt1}\right) \\ \sum_{t=1}^{T_n} \lambda_{nt} \left(\ell_{nt2} - \ell_{nt1}\right) \end{pmatrix}, 2\sigma_{inf}^2 \sum_{t=1}^{T_n} \lambda_{nt}^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right) \end{aligned} \tag{S30}$$

Thus, the mean is weighted as one would intuitively expect, and the covariance matrix is simply scaled by the weight. The correlation structure itself remains unchanged, as the temporal weighting affects all log-posterior differences equally. Similar changes apply to the joint distribution for the 3-categories condition. Specifically, the means are replaced by their weighted equivalents, and the covariance matrix is scaled by $\sum_{t=1}^{T_n} \lambda_{nt}^2$ rather than by T_n.

We consider two different ways to parameterize the λ_{nt}'s, which are discussed in turn.

Time-invariant exponential weighting. Time-invariant exponential weighting is based on multiplying the previous log-posterior with a constant α upon the addition of each new likelihood (Fig. S5D). For $\alpha < 1$, this can be interpreted as down-weighting the past upon perceiving new information. Formally, this leads to

$$z_{nT_nk} = \alpha(\alpha(\dots) + \hat{\ell}_{n,T_n-1,k}) + \hat{\ell}_{nT_nk} = \sum_{t=1}^{T_n} \alpha^{T_n - t} \hat{\ell}_{ntk}, \tag{S31}$$

resulting in the weights

$$\lambda_{nt} = \alpha^{T_n - t}. \tag{S32}$$

With this scheme, $\alpha < 1$ corresponds to a recency effect, putting more weight on log-likelihoods towards the end of the sample sequences, and $\alpha > 1$ leads to a primacy effect that puts more focus on samples early in the sequence. When fitting behavioral data, we used $-\infty < \log(\alpha) < \infty$ rather than α directly.

Linearly changing exponential weighting. In the previous section we have assumed α to remain constant across all samples in a sequence. Here, we introduce a slight generalization of this scheme that allows α to change linearly over samples according to

$$\alpha_t = \alpha_a + \alpha_b(t-1), \tag{S33}$$

where $t = 1, \ldots, T_n$ is the index in the sample sequence (Fig. S5E). This leads to the log-posteriors

$$z_{nT_nk} = (\alpha_a + \alpha_b(T_n - 2)) \left((\alpha_a + \alpha_b(T_n - 1)) (\ldots) + \hat{\ell}_{n,T_n-1,k} \right) + \hat{\ell}_{nT_nk} = \sum_{t=1}^{T_n} \hat{\ell}_{ntk} \prod_{s=t-1}^{T_n-2} (\alpha_a + \alpha_b s), \tag{S34}$$

with corresponding weights

$$\lambda_{nt} = \prod_{s=t-1}^{T_n-2} (\alpha_a + \alpha_b s). \tag{S35}$$

As before, for fitting behavioral data, we parameterized the model by $-\infty < \log(\alpha_a) < \infty$ rather than fitting α_a directly.

2.6 The relation between the softmax and the cumulative Gaussian

In this section, we derive choice distributions for Gaussian and Gumbel-distributed additive noise, and discuss their relation. Specifically, assume a set of K log-posteriors $z_{1:K}$, that are perturbed by i.i.d. zero-mean additive noise $\varepsilon_{1:K}$. A choice x is made by choosing the category associated with the largest of these noisy log-posteriors, that is

$$x = \underset{k}{\text{argmax}} \left(z_k + \varepsilon_k\right) \tag{S36}$$

In the following we derive the choice distribution $p(x|z_{1:K})$ for both Gaussian and Gumbel-distributed noise, for both $K = 2$ and $K = 3$, and show how they relate to each other.

2.6.1 Gaussian noise

First, assume that each ε_k is a zero-mean Gaussian with variance σ^2, that is $\varepsilon_k \sim \mathcal{N}(0, \sigma^2)$. Then, for $K = 2$, option 1 is chosen if $z_1+\varepsilon_1 > z_2+\varepsilon_2$, or equivalently, if $\varepsilon_2-\varepsilon_1 < z_1-z_2$. Noting that $\Delta_\varepsilon = \varepsilon_2-\varepsilon_1 \sim \mathcal{N}(0, 2\sigma^2)$, the choice distribution is thus given by

$$\begin{aligned} p(x = 1|z_1, z_2) &= \int p\left(\Delta_\varepsilon < z_1 - z_2 | z_1, z_2, \Delta_\varepsilon\right) \mathcal{N}\left(\Delta_\varepsilon|0, 2\sigma^2\right) \mathrm{d}\Delta_\varepsilon \\ &= \int_{-\infty}^{z_1-z_2} \mathcal{N}\left(\Delta_\varepsilon|0, 2\sigma^2\right) \mathrm{d}\Delta_\varepsilon \\ &= \Phi\left(\frac{z_1 - z_2}{\sigma\sqrt{2}}\right), \end{aligned} \tag{S37}$$

where $\Phi(a) = \int_{-\infty}^{a} \mathcal{N}(b|0, 1)\mathrm{d}b$ is the cumulative distribution function of a standard Gaussian. As this function is increasing in its argument, the likelihood of choosing option 1 increases in z_1 and decreases in z_2, as one would expect.

For $K = 3$ options, option k is chosen if $\forall j \neq k : z_k + \varepsilon_k > z_j + \varepsilon_j$, or equivalently, if $\forall j \neq k : \varepsilon_j - \varepsilon_k < z_k - z_j$. Without loss of generality we assume $k = 1$ and define $\Delta_2 = \varepsilon_2 - \varepsilon_1$ and $\Delta_3 = \varepsilon_3 - \varepsilon_1$, which are jointly Gaussian and distributed as

$$p(\Delta_2, \Delta_3) = \mathcal{N}\left(\begin{pmatrix} \Delta_2 \\ \Delta_3 \end{pmatrix} \middle| \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2\sigma^2 & \sigma^2 \\ \sigma^2 & 2\sigma^2 \end{pmatrix}\right). \tag{S38}$$

Option 1 is chosen if $\Delta_2 < z_1 - z_2$ and $\Delta_3 < z_1 - z3$, resulting in

$$p(x = 1|z_{1:3}) = \int_{-\infty}^{z_1-z_2} \int_{-\infty}^{z_1-z_3} p(\Delta_2, \Delta_3)\mathrm{d}\Delta_3\mathrm{d}\Delta_2, \tag{S39}$$

which is the cumulative distribution function of a bivariate Gaussian. There is no closed form for this function, and thus it needs to be computed numerically.

2.6.2 Gumbel-distributed noise

For Gumbel-distributed noise the resulting choice distribution is very similar to that resulting from Gaussian noise, but is mathematically more appealing. In particular, noise variables are distribution as

$$p(\varepsilon_k) = \beta e^{-\beta\varepsilon_k - e^{-\beta\varepsilon_k}}, \qquad p(\varepsilon_k < c) = e^{-e^{-\beta c}}, \tag{S40}$$

where β is the inverted scale parameter of the distribution.

Assuming $K = 2$ and proceeding as for the Gaussian case, it can be shown that $\Delta_\varepsilon = \varepsilon_2 - \varepsilon_1$ follows a zero-mean Logistic distribution with scale β^{-1}, that is $\Delta_\varepsilon \sim \text{Logistic}\left(0, \beta^{-1}\right)$. The choice probability for option 1 is the cumulative of this distribution at $z_1 - z_2$ and is thus given the cumulative of this distribution,

$$p(x = 1|z_1, z_2) = \frac{1}{1 + e^{-\beta(z_1-z_2)}}, \tag{S41}$$

which is the logistic sigmoid.

For $K > 2$, option k is chosen if $\forall j \neq k : \varepsilon_j < z_k - z_j + \varepsilon_k$, resulting in the choice distribution

$$\begin{aligned} p(y = k|z_{1:K}) &= \int_{-\infty}^{\infty} \mathrm{d}\varepsilon_k p(\varepsilon_k) \prod_{j \neq k} \int_{-\infty}^{z_k - z_j + \varepsilon_k} \mathrm{d}\varepsilon_j p(\varepsilon_j) \\ &= \int_{-\infty}^{\infty} \mathrm{d}\varepsilon_k p(\varepsilon_k) \prod_{j \neq j} p\left(\varepsilon_j < z_k - z_j + \varepsilon_k | z_k, z_j, \varepsilon_k\right), \end{aligned} \tag{S42}$$

which is a product of cumulative probabilities marginalized over different values of ε_k. Substituting the Gumbel distribution for the $p(\varepsilon)$'s results in

$$\begin{aligned} p(y = k|z_{1:K}) &= \int_{-\infty}^{\infty} \beta e^{-\beta\varepsilon_k - e^{-\beta\varepsilon_k}} \prod_{j \neq k} e^{-e^{-\beta(z_k - z_j + \varepsilon_k)}} \mathrm{d}\varepsilon_k \\ &= \int_{-\infty}^{\infty} \beta e^{-\beta\varepsilon_k} e^{-e^{-\beta\varepsilon_k} \sum_j e^{-\beta(z_k - z_j)}} \mathrm{d}\varepsilon_k \\ &= -\int_{\infty}^{0} e^{-y \sum_j e^{-\beta(z_k - z_j)}} \mathrm{d}y \\ &= \frac{1}{\sum_j e^{-\beta(z_k - z_j)}} \\ &= \frac{e^{\beta z_k}}{\sum_j e^{\beta z_j}}, \end{aligned} \tag{S43}$$

where the third equality is based on $y = e^{-\beta\varepsilon_k}$, and the fourth equality uses $\int e^{cy}\mathrm{d}y = c^{-1}e^{cy}$ with $c = -\sum_j e^{-\beta(z_k - z_j)}$. This shows that, unlike for Gaussian noise, assuming Gumbel-distributed noise results in the choice distribution for $K > 2$ to be given by the easy-to-evaluate softmax function. For $K = 2$ this function reduces to the logistic sigmoid.

2.6.3 Relating Gaussian and Gumbel-distributed noise

For $K = 2$, the choice distribution resulting from assuming Gaussian noise is a cumulative Gaussian, whereas that resulting from noise following the Gumbel distribution is a logistic sigmoid. From an experimental point of view, these two distributions are extremely similar and can only be distinguished with a large amount of data (see Fig. S3C). With only around 500 trials per subject, we can treat these two distributions as interchangeable.

Specifically, for most variability models we assume Gaussian noise, but fit the individual subject's behavior with a logistic sigmoid. To recover the standard deviation of the Gaussian noise, we relate the parameters of these two choice distributions by moment matching: the logistic sigmoid emerges from the cumulative of a logistic random variable with scale β^{-1}, which has variance $\pi^2/(3\beta^2)$. The cumulative Gaussian, in contrast, is based on a Gaussian with variance $2\sigma^2$. Equating these two variances results in

$$\sigma^2 \approx \frac{\pi^2}{6\beta^2}, \tag{S44}$$

which provides a good match between the two distributions (Fig. S3C).

For $K = 3$, the choice distribution is again determined by pairwise differences between noise variables, which are now correlated. In case of Gaussian noise, this results in the hard-to-evaluate cumulative distribution function of a bivariate Gaussian, whereas Gumbel-distributed noise results in the analytic softmax function. Thus, to simplify computation, we use the softmax function to fit the subject's behavior, even if the underlying model assumes Gaussian noise, using again Eq. (S44) to recover the noise standard deviation.

To test the quality of the approximate relationship in Eq. (S44) we simulated a large number of trials with statistics similar to these shown to the subjects. Furthermore, we added Gaussian noise to the log-posteriors in each trial, with a variance matching that observed in human data, and found the predicted choice distribution by evaluating Eq. (S37) or (S39) for 2 or 3 categories, respectively. Finally, we found the β that minimized the Kullback-Leibler divergence between these choice probabilities and those predicted by the sigmoid Eq. (S41) or softmax Eq. (S43). As shown below, the β's that provided the best match (when minimizing the Kullback-Leibler divergence) deviated slightly from those predicted by Eq. (S44). Thus, for any conversion between σ^2 and β (e.g., model fits, parameter estimates) we used the below values rather than Eq. (S44), using $\beta = \frac{C}{\sigma}$.

	C	$\frac{C}{C_{MM}}$
Moment matching, C_{MM}	$\frac{\pi}{\sqrt{6}} \approx 1.283$	–
2 categories	1.274	0.993
3 categories	1.303	1.016

2.7 Sensory variability choice distributions

In this section, we derive the choice distributions predicted by the sensory variability model. This model assumes that the perception of sample θ_{nt} is perturbed by zero-mean additive Gaussian noise ε_{nt} with variance σ_{sen}^2, that is $\varepsilon_{nt} \sim \mathcal{N}(0, \sigma_{sen}^2)$. This results in the noise-perturbed log-likelihoods to be given by

$$\hat{\ell}_{ntk} = \kappa \cos\left(2(\theta_{nt} + \varepsilon_{nt} - \mu_k)\right). \tag{S45}$$

Decision-making is performed by accumulating these log-likelihoods to form a log-posterior, followed by choosing the option associated with the largest of these log-posteriors.

2.7.1 Choice distribution for 2-categories condition

In the 2-categories condition, the decision-related quantity is the difference between the log-posterior for category 1 and that for category 2, given by

$$z_{nT_n1} - z_{nT_n2} = \sum_{t=1}^{T_n} (\hat{\ell}_{nt1} - \hat{\ell}_{nt2}) = 2\kappa \sum_{t=1}^{T_n} \cos\left(2(\theta_{nt} + \varepsilon_{nt} - \mu_1)\right), \tag{S46}$$

where we have used the fact that $\mu_2 = \mu_1 + \pi/2$, such that $\cos(2(\theta - \mu_2)) = -\cos(2(\theta - \mu_1))$. A positive difference, $z_{nT_n1} - z_{nT_n2} > 0$, leads to the choice of option 1.

To find the probability of choosing option 1, we assume the log-likelihood differences to be roughly Gaussian, with moments

$$\begin{aligned}
\left\langle \hat{\ell}_{nt1} - \hat{\ell}_{nt2} \right\rangle &= 2\kappa \int \cos\left(2(\theta_{nt} + \varepsilon_{nt} - \mu_1)\right) \mathcal{N}\left(\varepsilon_{nt} | 0, \sigma_{sen}^2\right) \mathrm{d}\varepsilon_{nt} \\
&= 2\kappa e^{-2\sigma_{sen}^2} \cos\left(2(\theta_{nt} - \mu_1)\right), && \text{(S47)} \\
\left\langle \left(\hat{\ell}_{nt1} - \hat{\ell}_{nt2}\right)^2 \right\rangle &= 4\kappa^2 \int \cos\left(2(\theta_{nt} + \varepsilon_{nt} - \mu_1)\right)^2 \mathcal{N}\left(\varepsilon_{nt} | 0, \sigma_{sen}^2\right) \mathrm{d}\varepsilon_{nt} \\
&= 2\kappa^2 \left(1 + e^{-8\sigma_{sen}^2} \cos\left(4(\theta_{nt} - \mu_1)\right)\right), && \text{(S48)} \\
\mathrm{var}\left(\hat{\ell}_{nt1} - \hat{\ell}_{nt2}\right) &= \left\langle \left(\hat{\ell}_{nt1} - \hat{\ell}_{nt2}\right)^2 \right\rangle - \left\langle \hat{\ell}_{nt1} - \hat{\ell}_{nt2} \right\rangle^2, && \text{(S49)}
\end{aligned}$$

where the solutions to the above integrals have been found by Mathematica (Wolfram Research, Inc., 2010). As the log-posterior difference is the sum of the log-likelihood differences, it will also be approximately Gaussian, given by

$$p\left(z_{nT_n1} - z_{nT_n2} | \theta_{n1:T_n}\right) \approx \mathcal{N}\left(z_{nT_n1} - z_{nT_n2} | \sum_{t=1}^{T_n} \left\langle \hat{\ell}_{nt1} - \hat{\ell}_{nt2} \right\rangle, \sum_{t=1}^{T_n} \mathrm{var}\left(\hat{\ell}_{nt1} - \hat{\ell}_{nt2}\right)\right). \tag{S50}$$

As a result, the probability for choosing option 1 is given by

$$p\left(x_n = 1 | \theta_{n1:T_n}\right) = p\left(z_{nT_n1} - z_{nT_n2} \geq 0 | \theta_{n1:T_n}\right) \approx \Phi\left(\frac{\sum_{t=1}^{T_n} \left\langle \hat{\ell}_{nt1} - \hat{\ell}_{nt2} \right\rangle}{\sqrt{\sum_{t=1}^{T_n} \mathrm{var}\left(\hat{\ell}_{nt1} - \hat{\ell}_{nt2}\right)}}\right), \tag{S51}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard Gaussian.

2.7.2 Choice distribution for 3-categories task

In the 3-categories condition, the choice is fully determined by the two log-posterior differences, $z_{nT_n1} - z_{nT_n2}$ and $z_{nT_n1} - z_{nT_n3}$. To derive their moments, we again take into account that the category means are related by $\mu_2 = \mu_1 - \pi/3$ and $\mu_3 = \mu_1 - 2\pi/3$. Let us define

$$a_{nt} = \frac{\hat{\ell}_{nt1} - \hat{\ell}_{nt2}}{\kappa} = \cos\left(2(\theta_{nt} + \varepsilon_{nt} - \mu_1)\right) - \cos\left(2\left(\theta_{nt} + \varepsilon_{nt} - \mu_1 + \frac{\pi}{3}\right)\right), \tag{S52}$$

$$b_{nt} = \frac{\hat{\ell}_{nt1} - \hat{\ell}_{nt3}}{\kappa} = \cos\left(2(\theta_{nt} + \varepsilon_{nt} - \mu_1)\right) - \cos\left(2\left(\theta_{nt} + \varepsilon_{nt} - \mu_1 + \frac{2\pi}{3}\right)\right), \tag{S53}$$

such that

$$z_{nT_n1} - z_{nT_n2} = \kappa \sum_{t=1}^{T_n} a_{nt}, \qquad \text{and} \qquad z_{nT_n1} - z_{nT_n3} = \kappa \sum_{t=1}^{T_n} b_{nt}. \tag{S54}$$

As for the 2-categories condition, we assume that the a_{nt}'s and b_{nt}'s are roughly Gaussian, with moments

$$\langle a_{nt} \rangle = e^{-2\sigma_{sen}^2}\left(\cos\left(2(\theta_{nt} - \mu_1)\right) + \sin\left(2\left(\theta_{nt} - \mu_1\right) + \frac{\pi}{6}\right)\right), \tag{S55}$$

$$\left\langle a_{nt}^2 \right\rangle = \frac{1}{4}\left(6 + 3e^{-8\sigma_{sen}^2}\left(\cos\left(4(\theta_{nt} - \mu_1)\right) + \sqrt{3}\sin\left(4(\theta_{nt} - \mu_1)\right)\right)\right), \tag{S56}$$

$$\mathrm{var}(a_{nt}) = \left\langle a_{nt}^2 \right\rangle - \langle a_{nt} \rangle^2, \tag{S57}$$

$$\langle b_{nt} \rangle = e^{-2\sigma_{sen}^2}\left(\cos\left(2(\theta_{nt} - \mu_1)\right) + \cos\left(2(\theta_{nt} - \mu_1) + \frac{\pi}{3}\right)\right), \tag{S58}$$

$$\left\langle b_{nt}^2 \right\rangle = \frac{1}{4}\left(6 + 3e^{-8\sigma_{sen}^2}\left(\cos\left(4(\theta_{nt} - \mu_1)\right) - \sqrt{3}\sin\left(4(\theta_{nt} - \mu_1)\right)\right)\right), \tag{S59}$$

$$\mathrm{var}(b_{nt}) = \left\langle b_{nt}^2 \right\rangle - \langle b_{nt} \rangle^2, \tag{S60}$$

$$\langle a_{nt} b_{nt} \rangle = \frac{1}{4}\left(3 + 6e^{-8\sigma_{sen}^2}\cos\left(4(\theta_{nt} - \mu_1)\right)\right), \tag{S61}$$

$$\mathrm{cov}(a_{nt}, b_{nt}) = \langle a_{nt} b_{nt} \rangle - \langle a_{nt} \rangle \langle b_{nt} \rangle. \tag{S62}$$

where the involved integrals have again been solved by Mathematica (Wolfram Research, Inc., 2010). Using these moments, the joint distribution of the log-posterior differences is given by

$$p\left(z_{nT_n1} - z_{nT_n2}, z_{nT_n1} - z_{nT_n3} | \theta_{n1:T_n}\right) \tag{S63}$$
$$\approx \mathcal{N}\left(\begin{pmatrix} z_{nT_n1} - z_{nT_n2} \\ z_{nT_n1} - z_{nT_n3} \end{pmatrix} \middle| \kappa \begin{pmatrix} \sum_{t=1}^{T_n} \langle a_{nt} \rangle \\ \sum_{t=1}^{T_n} \langle b_{nt} \rangle \end{pmatrix}, \kappa^2 \begin{pmatrix} \sum_{t=1}^{T_n} \mathrm{var}(a_{nt}) & \sum_{t=1}^{T_n} \mathrm{cov}(a_{nt}, b_{nt}) \\ \sum_{t=1}^{T_n} \mathrm{cov}(a_{nt}, b_{nt}) & \sum_{t=1}^{T_n} \mathrm{var}(b_{nt}) \end{pmatrix}\right).$$

Based on this joint distribution we find the choice distribution by

$$p(x_n = 1 | \theta_{n1:T_n}) = p\left(z_{nT_n1} - z_{nT_n2} \geq 0, z_{nT_n1} - z_{nT_n3} \geq 0 | \theta_{n1:T_n}\right) \tag{S64}$$

$$p(x_n = 2 | \theta_{n1:T_n}) = p\left(z_{nT_n1} - z_{nT_n2} < 0, z_{nT_n1} - z_{nT_n2} < z_{nT_n1} - z_{nT_n3} | \theta_{n1:T_n}\right), \tag{S65}$$

$$p(x_n = 3 | \theta_{n1:T_n}) = 1 - p(x_n = 1 | \theta_{n1:T_n}) - p(x_n = 2 | \theta_{n1:T_n}). \tag{S66}$$

Finding the above involves the numerical computation of the cumulative distribution function of a multivariate Gaussian, which we perform by using the function `mvncdf` in SciPy (Jones et al., 2001).

2.7.3 Accuracy of the Gaussian approximation

For both the 2-categories and the 3-categories condition we approximate the distribution of the log-likelihood differences by Gaussians. If this approximation is too coarse, the model fits are guided by the approximation rather than by the features of the model. To make sure that this is not the case, we tested its accuracy

by comparing for 500 typical trials the choice probability estimated by 50.000 instantiations of this noise to that predicted by the above derivation. As shown in Fig. S3B, this results in a very good match between theory and simulation, confirming that the applied approximations are adequate and do not confound our conclusions.

2.7.4 The relation between sensory and inference variability

In the 2-category condition, the log-posterior difference is Gaussian with variance $2(T_n - 1)\sigma_{inf}^2$, such that the probability of choosing option 1 is for the accumulation variability model given by

$$\Phi\left(\frac{-2\kappa \sum_{t=1}^{T_n} \cos\left(2(\theta_{nt} - \mu_1)\right)}{\sqrt{(T_n - 1)2\sigma_{inf}^2}}\right). \tag{S67}$$

For sensory variability, we find with the above expressions that

$$\sum_{t=1}^{T_n} \text{var}\left(\hat{\ell}_{nt1} - \hat{\ell}_{nt2}\right) = 4\kappa^2 e^{-4\sigma_{sen}^2}\left(\frac{T_n}{2}\left(e^{4\sigma_{sen}^2} - e^{-4\sigma_{sen}^2}\right) - \left(1 - e^{-4\sigma_{sen}^2}\right)\sum_{t=1}^{T_n} \cos^2\left(2(\theta_{nt} - \mu_1)\right)\right), \tag{S68}$$

where we have used $\cos(4\alpha) = 2\cos^2(2\alpha) - 1$. This leads to the probability of choosing option 1 to be given by

$$\Phi\left(\frac{-2\kappa \sum_{t=1}^{T_n} \cos\left(2(\theta_{nt} - \mu_1)\right)}{2\kappa\sqrt{\frac{T_n}{2}\left(e^{4\sigma_{sen}^2} - e^{-4\sigma_{sen}^2}\right) - \left(1 - e^{-4\sigma_{sen}^2}\right)\sum_{t=1}^{T_n} \cos^2\left(2(\theta_{nt} - \mu_1)\right)}}\right). \tag{S69}$$

Thus, the main difference to inference variability is that the denominator also depends on the sample sequence, through $\sum_t \cos^2\left(2(\theta_{nt} - \mu_1)\right)$. This term measures sums over the angular distances between the orientations and the generative mean. In particular, each term in the sum is largest if the orientation matches one of the generative means, μ_k, such that the sum is smallest if all orientations in the sequence are furthest away from these means. As the term sums over all orientations in the sequence, it will vary most for short sequences, and will converge to a common value for long sequences. Therefore, short sequences are best to distinguish between models postulating sensory and those postulating inference variability.

2.8 Choice distributions with deterministic biases

This section describes how to find the choice distribution under the assumption of a biased log-posterior estimate, for both the 2-categories and 3-categories condition. In both cases, we assume that the exact same sample sequence $\theta_{n1:T_n}$ has been presented in both trials n and m. Our aim is to derive the joint probability of choosing categories x_n and x_m in these trials.

We assume the log-posterior to be based on noisy and biased log-likelihoods, $\hat{\ell}_{ntk} = \ell_{ntk} + f_k(\theta_{nt}) + \varepsilon_{ntk}$, where ℓ_{ntk} is the true log-likelihood of sample t in trial n with respect to category k. Our aim is to distinguish the contribution of bias $f_k(\theta_{nt})$ and variance ε_{ntk} to the overall variability. In the below, we assume the biases to act on single orientation percepts, and to be independent across samples and trials. The same results follow if we were to assume biases to act on the whole sequence of orientations.

2.8.1 Choice distribution for the 2-categories condition

In the 2-categories condition, we assume assume bias and variance terms, $f_k(\theta_{nt})$ and ε_{ntk}, the be across trials and samples distributed as $f_2(\theta_{nt}) - f_1(\theta_{nt}) \sim \mathcal{N}(0, 2\sigma_b^2)$ and $\varepsilon_{ntk} \sim \mathcal{N}(0, \sigma_v^2)$ for $k \in \{1, 2\}$ (see Fig. S3A for a justification of this assumption). We only consider bias differences, as the log-posterior differences $z_{nT_n2} - z_{nT_n1}$ and $z_{mT_m2} - z_{mT_m1}$ fully determine the choice in trials n and m. Conditional on $f_{n1:T_n1}$ and $f_{n1:T_n2}$, where $f_{n1:T_nk} = \{f_k(\theta_{n1}), \ldots, f_k(\theta_{T_n1})\}$, their joint distribution is

$$\begin{aligned} &p\left(z_{nT_n2} - z_{nT_n1}, z_{mT_m2} - z_{mT_m1} | f_{n1:T_n1}, f_{n1:T_n2}\right) \\ &\quad = \mathcal{N}\left(\begin{pmatrix} z_{nT_n2} - z_{nT_n1} \\ z_{mT_m2} - z_{mT_m1} \end{pmatrix} \middle| \begin{pmatrix} \sum_{t=1}^{T_n} \left(\ell_{nt2} - \ell_{nt1} + f_2(\theta_{nt}) - f_1(\theta_{nt})\right) \\ \sum_{t=1}^{T_n} \left(\ell_{nt2} - \ell_{nt1} + f_2(\theta_{nt}) - f_1(\theta_{nt})\right) \end{pmatrix}, \begin{pmatrix} 2T_n\sigma_v^2 & 0 \\ 0 & 2T_n\sigma_v^2 \end{pmatrix}\right). \end{aligned} \tag{S70}$$

Thus, they share the same mean, including the bias terms. However, due to assuming knowledge of all bias terms, they are uncorrelated and thus independent.

The log-posterior difference become dependent as soon as we average across many trial grouping. Specifically, averaged over many such grouping with different sample sequences, we have $f_2(\theta_{nt}) - f_1(\theta_{nt}) \sim \mathcal{N}\left(0, 2\sigma_b^2\right)$, such their sum over T_n samples is distributed as $\mathcal{N}\left(0, 2T_n\sigma_b^2\right)$. As a result, if we replace knowledge of the bias terms by their averages (that is, by marginalizing them out), the resulting joint distribution becomes

$$\begin{aligned} &p\left(z_{nT_n2} - z_{nT_n1}, z_{mT_m2} - z_{mT_m1}\right) \\ &\quad = \iint p\left(z_{nT_n2} - z_{nT_n1}, z_{mT_m2} - z_{mT_m1} | f_{nT_n}\right) p\left(f_{n1:T_n1}, f_{n1:T_n2}\right) \mathrm{d}f_{n1:T_n1} \mathrm{d}f_{n1:T_n2} \\ &\quad = \mathcal{N}\left(\begin{pmatrix} z_{nT_n2} - z_{nT_n1} \\ z_{mT_m2} - z_{mT_m1} \end{pmatrix} \middle| \begin{pmatrix} \sum_{t=1}^{T_n} \left(\ell_{nt2} - \ell_{nt1}\right) \\ \sum_{t=1}^{T_n} \left(\ell_{nt2} - \ell_{nt1}\right) \end{pmatrix}, 2T_n\sigma_{inf}^2 \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right), \end{aligned} \tag{S71}$$

where the log-posterior difference are now correlated with coefficient ρ. To find the above, we have used

$$\mathrm{var}(z_{nT_n2} - z_{nT_n1}) = 2T_n\sigma_b^2 + 2T_n\sigma_v^2 = 2T_n\sigma_{inf}^2, \tag{S72}$$

$$\mathrm{cov}\left(z_{nT_n2} - z_{nT_n1}, z_{mT_m2} - z_{mT_m1}\right) = \sum_{t=1}^{T_n} \left\langle \left(f_2(\theta_{nt}) - f_1(\theta_{nt})\right)^2 \right\rangle = 2T_n\sigma_b^2, \tag{S73}$$

$$\rho = \frac{2T_n\sigma_b^2}{2T_n\sigma_v^2 + 2T_n\sigma_b^2} = \frac{\sigma_b^2}{\sigma_{inf}^2}. \tag{S74}$$

This shows that bias terms introduce correlations in log-posterior differences of grouped trials, and as a result also in their choices. Furthermore, these correlations are always non-negative, which results in an increase in the probability of making the same choice (correct or incorrect) in both of these trials. At last, by the definition of the correlation coefficient $\rho = \sigma_b^2/\sigma_{inf}^2$, this coefficient determines the fraction of noise variance that the biases contribute to the overall inference noise variance. We estimated this ρ for each subject by fitting a model in which the above joint distribution determined the subjects' choices in each trial pair in which the same sample sequence was presented.

2.8.2 Choice distribution for the 3-categories condition

The same principles apply in the 3-categories condition, but their formulation is more complex due to choices being based on two log-posterior differences per trial rather than one. This causes the joint choice probability of trials n and m to be given by the cumulative of a 4-dimensional multivariate Gaussian. For now, we will focus on the covariance matrix of this Gaussian. We will come back to its mean in a later section.

To fully specify the covariance between all four random variables, we need to additionally introduce the parameters ρ_{21}, ρ_{31}, and ρ_{32}, based on which we define

$$\begin{aligned}
\text{var}\left(f_2(\theta_{nt}) - f_1(\theta_{nt})\right) &= 2\sigma_{inf}^2\rho_{21}, && \text{(S75)}\\
\text{var}\left(f_3(\theta_{nt}) - f_1(\theta_{nt})\right) &= 2\sigma_{inf}^2\rho_{31}, && \text{(S76)}\\
\text{var}\left(f_3(\theta_{nt}) - f_2(\theta_{nt})\right) &= 2\sigma_{inf}^2\rho_{32}. && \text{(S77)}
\end{aligned}$$

Furthermore, we define the total average variance, $3^{-1}\sum_{k\neq l}\text{var}\left(f_k(\theta_{nt}) - f_l(\theta_{nt}) + \varepsilon_k - \varepsilon_l\right)$ to equal $2\sigma_{inf}^2$, where each ε_k is an independent zero-mean Gaussian with variance $\left\langle\varepsilon_k^2\right\rangle = \sigma_v^2$. We allow the three correlation parameters to become negative, in which case, the above terms cease to be variances (which cannot be negative). Even then, however, we have $\sigma_v^2 = (1-\bar{\rho})\sigma_{inf}^2$, where $\bar{\rho} = 3^{-1}\sum_{k\neq l}\rho_{kl}$ is the average correlation coefficient, such that $-1 \leq \bar{\rho} \leq 1$, and consequently, $\sigma_v^2 \geq 0$, leaving the problem well-defined.

Based on the above definitions it is easy to show that the relevant covariances are given by

$$\begin{aligned}
\text{cov}\left(f_2(\theta_{nt}) - f_1(\theta_{nt}), f_3(\theta_{nt}) - f_1(\theta_{nt})\right) &= \sigma_{inf}^2\left(\rho_{21} + \rho_{31} - \rho_{32}\right), && \text{(S78)}\\
\text{cov}\left(f_2(\theta_{nt}) - f_1(\theta_{nt}), f_3(\theta_{nt}) - f_2(\theta_{nt})\right) &= \sigma_{inf}^2\left(-\rho_{21} + \rho_{31} - \rho_{32}\right), && \text{(S79)}\\
\text{cov}\left(f_3(\theta_{nt}) - f_1(\theta_{nt}), f_3(\theta_{nt}) - f_2(\theta_{nt})\right) &= \sigma_{inf}^2\left(-\rho_{21} + \rho_{31} + \rho_{32}\right), && \text{(S80)}
\end{aligned}$$

Overall, the whole model is specified by the four parameters σ_{inf}^2, ρ_{21}, ρ_{31}, and ρ_{32}.

The base covariance matrix. Here we derive the base covariance matrix, which is the one corresponding to the choice distribution in which option 1 is chosen in both trials n and m. As will be shown in a later section, all other covariance matrices can be derived from this one. Choosing option 1 in both trials corresponds to $z_{nT_n2} - z_{nT_n1} \leq 0$ and $z_{nT_n3} - z_{nT_n1} \leq 0$ for trial n, and the analogous inequalities for trial m. To specify the full covariance matrix over the four log-posterior differences, we require the following individual variances and covariances:

$$\begin{aligned}
\text{var}\left(\hat{\ell}_{nt2} - \hat{\ell}_{nt1}\right) &= 2\sigma_{inf}^2\rho_{21} + 2\sigma_v^2 = 2\sigma_{inf}^2\left(1 + \frac{2}{3}\rho_{21} - \frac{1}{3}\rho_{31} - \frac{1}{3}\rho_{32}\right), && \text{(S81)}\\
\text{var}\left(\hat{\ell}_{nt3} - \hat{\ell}_{nt1}\right) &= 2\sigma_{inf}^2\left(1 - \frac{1}{3}\rho_{21} + \frac{2}{3}\rho_{31} - \frac{1}{3}\rho_{32}\right), && \text{(S82)}\\
\text{cov}\left(\hat{\ell}_{nt2} - \hat{\ell}_{nt1}, \hat{\ell}_{nt3} - \hat{\ell}_{nt1}\right) &= \sigma_{inf}^2\left(\rho_{21} + \rho_{31} - \rho_{32}\right) + \sigma_v^2 = 2\sigma_{inf}^2\left(\frac{1}{2} + \frac{1}{3}\rho_{21} + \frac{1}{3}\rho_{31} - \frac{2}{3}\rho_{32}\right) && \text{(S83)}\\
\text{cov}\left(\hat{\ell}_{nt2} - \hat{\ell}_{nt1}, \hat{\ell}_{mt2} - \hat{\ell}_{mt1}\right) &= 2\sigma_{inf}^2\rho_{21}, && \text{(S84)}\\
\text{cov}\left(\hat{\ell}_{nt2} - \hat{\ell}_{nt1}, \hat{\ell}_{mt3} - \hat{\ell}_{mt1}\right) &= 2\sigma_{inf}^2\left(\frac{1}{2}\rho_{21} + \frac{1}{2}\rho_{31} - \frac{1}{2}\rho_{32}\right), && \text{(S85)}\\
\text{cov}\left(\hat{\ell}_{nt3} - \hat{\ell}_{nt1}, \hat{\ell}_{mt3} - \hat{\ell}_{mt1}\right) &= 2\sigma_{inf}^2\rho_{31}. && \text{(S86)}
\end{aligned}$$

Multiplied by T_n for the likelihood variability model or $T_n - 1$ for the accumulation variability model, these terms form the covariance matrix of the four log-posterior differences.

The covariance matrix is by definition positive semi-definite, which introduces constraints on the possible parameter combinations $(\rho_{21}, \rho_{31}, \rho_{32})$. Positive semi-definiteness is guaranteed as long as all eigen-

values of the matrix are non-negative. Its four eigenvalues can be shown to be given by

$$\frac{1}{6}\left\{\begin{array}{c} 3-\rho_{21}-\rho_{31}-\rho_{32}, \\ 9-3\rho_{21}-3\rho_{31}-3\rho_{32}, \\ 6+4(\rho_{21}+\rho_{31})-2\rho_{32} \\ -\sqrt{61(\rho_{21}^2+\rho_{31}^2)+\rho_{31}(30-70\rho_{32})+\rho_{21}(30-22\rho_{31}-70\rho_{32})+(3-7\rho_{32})^2}, \\ 6+4(\rho_{21}+\rho_{31})-2\rho_{32} \\ +\sqrt{61(\rho_{21}^2+\rho_{31}^2)+\rho_{31}(30-70\rho_{32})+\rho_{21}(30-22\rho_{31}-70\rho_{32})+(3-7\rho_{32})^2} \end{array}\right\} \tag{S87}$$

The first, second, and fourth eigenvalue are guaranteed to be non-negative for all possible correlation coefficient combinations. The third eigenvalue, however, can become negative for some points within the cube $(\rho_{21}, \rho_{31}, \rho_{32})$ of correlation coefficients. In particular, it becomes negative for parameter combinations outside of the cone defined by

$$5(\rho_{21}^2+\rho_{31}^2+\rho_{32}^2)-2(\rho_{21}+\rho_{31}+\rho_{32})-6(\rho_{21}\rho_{31}+\rho_{21}\rho_{32}+\rho_{31}\rho_{32})-3=0, \tag{S88}$$

In what follows, we re-parameterize the correlation coefficients to ensure that the used coefficients always remain inside this cone.

Re-parameterizing the correlation coefficients. The cone that describes the boundary within which the covariance remains positive semi-definite has its peak at $(-1,-1,-1)$ and its axis along the line $\rho_{21}=\rho_{31}=\rho_{32}$. First, we introduce a re-parameterization $(\bar{\rho}, a, b)$ based on rotation and scaling,

$$\rho_{21}=\bar{\rho}-\frac{a}{6}-\frac{b}{2\sqrt{3}}, \qquad \rho_{31}=\bar{\rho}-\frac{a}{6}+\frac{b}{2\sqrt{3}}, \qquad \rho_{32}=\bar{\rho}+\frac{a}{3}. \tag{S89}$$

Here $\bar{\rho}$ determines the coordinate along the axis of the cone, and a and b the coordinates in the plane orthogonal to the cone's axis. With these parameters, the cone reduces to $9(1+\bar{\rho})^2=4(a^2+b^2)$, showing that a and b need to lie within a circle with radius $3(1+\bar{\rho})/2$. A further re-parameterization

$$a=\frac{3}{2}(1+\bar{\rho})\alpha\sqrt{1-\frac{\beta^2}{2}}, \qquad b=\frac{3}{2}(1+\bar{\rho})\beta\sqrt{1-\frac{\alpha^2}{2}}, \tag{S90}$$

ensures that a and b remain within this circle for all $\alpha \in [-1,1]$ and $\beta \in [-11]$. This leads to the final mapping

$$\rho_{21}=\bar{\rho}-\alpha'-\sqrt{3}\beta', \qquad \rho_{31}=\bar{\rho}-\alpha'+\sqrt{3}\beta', \qquad \rho_{32}=\bar{\rho}+2\alpha', \tag{S91}$$

with

$$\alpha'=\frac{\bar{\rho}+1}{4}\alpha\sqrt{1-\frac{\beta^2}{2}}, \qquad \beta'=\frac{\bar{\rho}+1}{4}\beta\sqrt{1-\frac{\alpha^2}{2}}. \tag{S92}$$

Thus, as long as $\bar{\rho} \in [-1,1]$, $\alpha \in [-1,1]$ and $\beta \in [-1,1]$, we guarantee to find parameters $(\rho_{21}, \rho_{31}, \rho_{32})$ within or on the surface of the cone. What this parameterization does not guarantee is that all correlation coefficients are bounded by -1 and 1, which needs to be checked additionally. Overall, this leads to the base covariance matrix

$$\mathrm{cov}\begin{pmatrix} \hat{\ell}_{nt2}-\hat{\ell}_{nt1} \\ \hat{\ell}_{nt3}-\hat{\ell}_{nt1} \\ \hat{\ell}_{mt2}-\hat{\ell}_{mt1} \\ \hat{\ell}_{mt3}-\hat{\ell}_{mt1} \end{pmatrix} = 2\sigma_{inf}^2 \begin{pmatrix} 1-\alpha'-\sqrt{3}\beta' & \frac{1}{2}-2\alpha' & \bar{\rho}-\alpha'-\sqrt{3}\beta' & \frac{1}{2}\bar{\rho}-2\alpha' \\ \frac{1}{2}-2\alpha' & 1-\alpha'+\sqrt{3}\beta' & \frac{1}{2}\bar{\rho}-2\alpha' & \bar{\rho}-\alpha'+\sqrt{3}\beta' \\ \bar{\rho}-\alpha'-\sqrt{3}\beta' & \frac{1}{2}\bar{\rho}-2\alpha' & 1-\alpha'-\sqrt{3}\beta' & \frac{1}{2}-2\alpha' \\ \frac{1}{2}\bar{\rho}-2\alpha' & \bar{\rho}-\alpha'+\sqrt{3}\beta' & \frac{1}{2}-2\alpha' & 1-\alpha'+\sqrt{3}\beta' \end{pmatrix}. \tag{S93}$$

In the main text, only the average correlation coefficient, $\bar{\rho}$, is reported.

The choice distribution for arbitrary trial groups. Based on the above, we can compute the choice probability of choosing x_n and x_m in grouped trials n and m. This is based on using the statistics

$$\begin{pmatrix} z_{nT_n2} - z_{nT_n1} \\ z_{nT_n3} - z_{nT_n1} \\ z_{mT_m2} - z_{mT_m1} \\ z_{mT_m3} - z_{mT_m1} \end{pmatrix} \sim \mathcal{N}\left(\sum_{t=1}^{T_n} \begin{pmatrix} \ell_{nt2} - \ell_{nt1} \\ \ell_{nt3} - \ell_{nt1} \\ \ell_{mt2} - \ell_{mt1} \\ \ell_{mt3} - \ell_{mt1} \end{pmatrix}, \sum_{t=1}^{T_n} \mathrm{cov} \begin{pmatrix} \hat{\ell}_{nt2} - \hat{\ell}_{nt1} \\ \hat{\ell}_{nt3} - \hat{\ell}_{nt1} \\ \hat{\ell}_{mt2} - \hat{\ell}_{mt1} \\ \hat{\ell}_{mt3} - \hat{\ell}_{mt1} \end{pmatrix} \right), \tag{S94}$$

which we will denote $\boldsymbol{z}_{n,m} \sim \mathcal{N}(\boldsymbol{\mu}_{n,m}, \boldsymbol{\Sigma}_{n,m})$ (for the accumulation variability model, the sum's lower limit for the variance is $t = 2$ rather than $t = 1$). For choice $x_n = x_m = 1$, the choice probability is simply the probability of $\boldsymbol{z}_{n,m} \leq 0$ (element-wise), which can be easily computed from the above normal density.

For choices other than $x_n = 1$ or $x_m = 1$, we can derive the choice probability through a linear transformation of the above density. In particular, define

$$\boldsymbol{A}_1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad \boldsymbol{A}_2 = \begin{pmatrix} -1 & 0 \\ -1 & 1 \end{pmatrix}, \qquad \boldsymbol{A}_3 = \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}, \tag{S95}$$

such that these matrices map the base case into the ones of interest for the corresponding choice, that is

$$\boldsymbol{A}_2 \begin{pmatrix} z_2 - z_1 \\ z_3 - z_1 \end{pmatrix} = \begin{pmatrix} z_2 - z_1 \\ z_3 - z_2 \end{pmatrix}, \qquad \boldsymbol{A}_3 \begin{pmatrix} z_2 - z_1 \\ z_3 - z_1 \end{pmatrix} = \begin{pmatrix} z_1 - z_3 \\ z_2 - z_3 \end{pmatrix}. \tag{S96}$$

With these mapping we can define the 4×4 matrix

$$\boldsymbol{B}_{n,m} = \begin{pmatrix} \boldsymbol{A}_{x_n} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{A}_{x_m} \end{pmatrix}. \tag{S97}$$

This allows us to compute the corresponding choice probability by the mass of $\boldsymbol{B}_{n,m}\boldsymbol{z}_{n,m} \leq 0$, where

$$\boldsymbol{B}_{n,m}\boldsymbol{z}_{n,m} \sim \mathcal{N}\left(\boldsymbol{B}_{n,m}\boldsymbol{\mu}_{n,m}, \boldsymbol{B}_{n,m}\boldsymbol{\Sigma}_{n,m}\boldsymbol{B}_{n,m}^T\right). \tag{S98}$$

Thus, computing the cumulative of the above multivariate Gaussian reveals the choice probability for arbitrary trial groups.

2.8.3 Fitting models with deterministic biases

We estimated the balance between bias and variance by fitting subjects choices in pairs of trials of the second experiment in which the same sample sequence was shown. Fitting was performed by maximum-likelihood, using the same hybrid posterior sampling/gradient ascent approach as to fit all other models (see Sec. 2.9). This resulted in the overall bias/variance decomposition reported in the main text and in Fig. 6B. We performed the decomposition into temporal and spatial (i.e., orientation perception) deterministic distortions by fitting three additional models, either only assuming temporal or spatial distortions, or a combination of both. The contribution of each type of distortion was measured by the amount of bias we were able to explain away by adding these distortions, and is shown in Fig. 6B (green and blue portion). A two-way repeated-measures ANOVA with spatial and temporal biases as separate factors revealed a significant effect of both bias types on the residual bias (two categories: spatial $F_{1,17} = 55.9$, $p < 0.001$, temporal $F_{1,17} = 28.0$, $p < 0.001$; three categories: spatial $F_{1,17} = 14.2$, $p = 0.002$, temporal $F_{1,17} = 4.6$, $p = 0.046$) but not significant interaction between them (spatial $\times$ temporal: two categories $F_{1,17} = 4.4$, $p \approx 0.051$; three categories $F_{1,17} < 0.1$, $p > 0.5$). The lack of interaction allowed us to interpret their contributions as additive.

2.8.4 The contribution of sequential biases

We analyzed the effect of sequential choice dependencies, as shown in Fig. 2B, on the choice variability decomposition as follows. First, we introduced an additional bias towards choosing the previous feedback

(i.e., the previously drawn deck) to the variability structure model, whose magnitude was a free parameter in units of log-posterior odds. Second, we fitted this sequential bias-augmented model as before to the participants behavior. Any kind of variability due to such sequential biases was reflected in a reduction of unstructured variability when compared to model fits that did not include this bias. Thus, third, we attributed any such reduction to sequential deterministic bias, as illustrated by the orange portion in Fig. 6B. We repeated this analysis by adding a second sequential bias towards choosing the same deck as in the previous. This second, choice repetition regressor: 1) did not differ significantly from zero across human participants (Fig. 2B), and 2) did not result in a measurable reduction of unstructured variability when compared to model fits that did not include this bias. We therefore safely omitted it when computing the overall impact of sequential deterministic bias depicted in Fig. 6B.

2.8.5 Stability of deterministic biases

A drift in deterministic biases throughout an experimental session would be captured by the unbiased variability term ε_{ntk}, and would inflate our measure of σ_v^2. We performed several tests to detect such a drift. First, a change in bias might be reflected by a change in performance throughout an experimental session. Comparing behavior between the first and second half of trial did not reveal any significant change in neither performance (Fig. S1A) not inferred behavioral variability (Fig. S1C). Second, neither did we find any significant change in spatial or temporal biases between the first and second half of trials (Fig. S4A). Third, a strategy employed by the subjects might be influenced by the outcome of the previous choice. However, as discussed in the main text, we did not find any influence of a trials outcome on successive choices. Estimating spatial and temporal biases conditional on this outcome revealed a qualitative, but insignificant difference in the temporal biases, but not the spatial biases (Fig. S4B). Fourth, a drift in deterministic biases might change the fraction of matched choices in trial pairs in which the same sample sequence was shown. Comparing this fraction between first and second half of trials in the third experiment did not reveal any significant difference across trial subgroups (Fig. S4C). Fifth, a drift in deterministic biases predicts a larger measure of these biases for close than for distant trials. Estimating the contribution of bias to total variability separately for different trial subgroups (same trial groupings as in Fig. S4C) did not reveal any significant change in bias magnitudes between either trial subgroups (Fig. 6B). Finally, we found that in both the two- and three-decks condition, participants choice consistency across repeated sequences was unrelated to the strongly varying temporal distance between the paired sequences (logistic regression, both $t_{17} < 0.9$, $p > 0.2$; trial distance mean = 72 trials, $\approx$ 7min.; interquartile range = 42 trials, $\approx$ 3 to 12min.). This validated our assumption that deterministic biases are systematic across trials. In combination, this makes it unlikely that our measure of the contribution of deterministic biases to overall variability was strongly influence by a drift in these biases across trials.

2.9 Model fitting, fit validation, and number of parameters

To distinguish between the different hypotheses, we fitted each model to the choice behavior of each participant separately by maximum-likelihood, following a two-step procedure. First, we drew 1,000,000 samples (+ 50,000 burn-in; 100,000 samples + 5,000 burn-in for the above noise structure models) from the parameter posterior (assuming a uniform prior over a limited parameter range) using a population Markov Chain Monte Carlo (MCMC) method (Goodman and Weare, 2010) implemented in Python (Foreman-Mackey et al., 2012), and then chose the sample associated with the largest posterior probability as a starting point for pseudo-gradient ascent on this posterior. The found mode resulted in the maximum-likelihood parameters. We performed this two-step procedure: 1) to avoid getting stuck in localized maxima, 2) to estimate the parameter uncertainty by the width of the posterior, and 3) to ensure finding the posterior mode rather than a close-by set of parameters.

Wherever shown, Bayesian credible intervals on parameter estimates were found in two steps. First, we ensured good sample coverage of the posterior by re-sampling all posterior samples by drawing them with replacement from the set of MCMC samples with a probability proportional to their posterior probability. Second, we computed the x% credible intervals by the $x/2$th and $(1 - x/2)$th percentile of the re-sampled set.

The log-likelihood per subject, model, and condition resulted from summing the log-likelihoods for the subjects choices for individual trials (or trial pairs for the noise structure models). Per trial, the choice likelihoods were given by the softmax function for selection, likelihood and accumulation noise models, and the cumulative Gaussian for all other models. To avoid confounds due to occasional random responses and response biases, we added to each model a lapse probability (choosing uniformly random rather than according to choice model) and $K - 1$ response biases (fixed offset added to final log-posterior for all but one category, see Fig. S7). The resulting total number of parameters per model is provided further below.

Based on these model fits, we performed two types of Bayesian model comparisons. In both cases, we used the Bayesian model evidence (approximated by the Bayesian Information Criterion) as a measure of model fit quality. This model evidence adequately takes into account the model complexity by penalizing models with an overly large number of parameters. The fixed-effects comparison assumes all participants to have used the same underlying model to generate their behavior, such that the overall model evidence for model $\mathcal{M}$ is proportional to the product of model evidences for model $\mathcal{M}$ for all participants. Based on this model evidence, we compared different models by computing their Bayes factor as the ratio of model evidences of the compared model. A Bayes factor exceeding 100 was considered decisive evidence supporting the model in the nominator of this factor (Kass and Raftery, 1995). The random-effects comparison is more conservative in allowing different participants to use different models to generate their behavior, and aims at inferring the distribution over models that participants draw from (Stephan et al., 2009). For this comparison, we computed support for the most likely model by the exceedance probability (by MCMC sampling, 20,000 samples + 10,000 burn-in), which is the (second-order) probability that participants were more likely to choose this model to generate behavior than any alternative model.

2.9.1 Validating the model fitting procedure

To assess the validity of our model comparison approach in discriminating the three hypothesized sources of choice variability, we performed the following validation procedure: we generated synthetic choice data from the sensory, inferential or selective source models with variability magnitudes corresponding to the best-fitting values predicting human choice variability. Specifically, For each of the 22 subjects and both two/three-category conditions we performed three simulations, each introducing variability at a different point in the decision-making process (sensory/inference/selection), using the models described further above. We then applied our model fitting procedure to the synthetic choice data for estimating the discriminative power of our model comparison procedure (Fig. S2D and Fig. 3C in main text). We found that in both conditions, synthetic choice variability reflecting sensory and selective imperfections were classified as stemming from sensory and selective imperfections, respectively, with more than 99% of correct classification in terms of both fixed- and random-effects statistics. Synthetic choice variability reflecting inferential

imperfections were classified as stemming from inferential imperfections with 90% of correct classification (Fig. 3C in main text). Moreover, Fig. S2A shows that our procedure was able to accurately differentiate inferential from attentional sources of choice variability. Lastly, repeating the same analysis with synthetic choice data using heavy-tailed instead of Gaussian inference variability revealed that model-fitting was not sensitive to the specifics of the assumed variability structure (Fig. S2D). Overall, this indicates that our approach was able to accurately recover the sources of choice suboptimality, and was even conservative in detecting inferential imperfections by tending to attribute choice variability stemming from inferential imperfections to peripheral imperfections rather than the converse.

2.9.2 Number of parameters

The following table shows the total number of parameters, depending on the model type and the added biases. The number of base parameters depends on the number of categories and if trials are modeled individually (experiments 1, 2 and 3) or in pairs (experiment 3). It does not depend on the assumed type of variability (e.g. inference variability, sensory variability, ...). Each bias adds additional parameters to the model. As the linearly changing exponential temporal bias is strictly more general than the time-invariant type, only one of the two temporal biases can be added at one time.

	individual trials		trial pairs	
	2 categories	3 categories	2 categories	3 categories
Variability model	1	1	2	4
Response biases	1	2	1	2
Lapse probability	1	1	1	1
Spatial biases				
Orientation bias		+1		
Oblique effect		+1		
Confirmation bias		+1		
Temporal biases				
Time-invariant exponential		+1		
Linearly changing exponential		+2		
Total	3 – 8	4 – 9	4 – 9	7 – 12

2.10 Computing the information loss for the 2-categories condition

We quantified the information loss by computing the mutual information between generative category and log-likelihood ratio for a single stimulus orientation for both the noise-free and the noisy case. The information loss itself is the mutual information for the noisy case divided by that of the noise-free case. By the data processing inequality, this fraction is guaranteed to be at most one. We also decomposed this loss into contributions due to unstructured variability and deterministic biases. Due to the non-linear mapping between variability and information loss, this decomposition was not unique, such that we computed the bounds on the loss decomposition within which it is compatible with the estimated variability decomposition. These bounds are shown in Fig. 6C (main text).

Denoting the generative category by x and the noise-free log-likelihood ratio by Δ_ℓ, the mutual information can be written as

$$\mathrm{MI}(x;\Delta_\ell) = \sum_x p(x) \left(\int p(\Delta_\ell|x) \log p(\Delta_\ell|x) \mathrm{d}\Delta_\ell - \int p(\Delta_\ell|x) \log p(\Delta_\ell) \mathrm{d}\Delta_\ell \right). \tag{S99}$$

The expression for the noisy case is analogous, with Δ_ℓ's replaced by the noisy log-likelihood ratio, $\tilde{\Delta}_\ell$.

To compute the mutual information, we first need to find an expression for $p(\Delta_\ell|x)$. For a given stimulus orientation θ, the log-likelihood ratio of category one vs. category two is given by

$$\Delta_\ell = \ell_{\cdot 1} - \ell_{\cdot 2} = 2\kappa \cos\left(2(\theta - \mu_1)\right), \tag{S100}$$

where we have used the fact that, in the 2-categories condition, we have $\mu_2 = \mu_1 + \frac{\pi}{2}$. The above shows that Δ_ℓ is a deterministic function of θ that is bounded by $-2\kappa \leq \Delta_\ell \leq 2\kappa$. Thus we can use the generative model to find $p(\Delta_\ell|x)$ by

$$p(\Delta_\ell|x) = p(\theta|x) \left| \frac{\mathrm{d}\theta}{\mathrm{d}\Delta_\ell} \right| = \frac{e^{\kappa \cos(2(\theta-\mu_x))}}{\pi I_0(\kappa)} \left| \frac{1}{-4\kappa \sin\left(2(\theta-\mu_1)\right)} \right|, \tag{S101}$$

where the second equality is based on replacing $\frac{\mathrm{d}\theta}{\mathrm{d}\Delta_\ell}$ by one over the evaluated derivative $\frac{\mathrm{d}\Delta_\ell}{\mathrm{d}\theta}$. What remains is to re-express the θ's in the above in terms of Δ_ℓ's. Considering first $x = 1$, we find that

$$\theta = \begin{cases} \mu_1 + \frac{1}{2}\cos^{-1}\left(\frac{\Delta_\ell}{2\kappa}\right) & \text{if } \theta \leq \mu_1, \\ \mu_1 - \frac{1}{2}\cos^{-1}\left(\frac{\Delta_\ell}{2\kappa}\right) & \text{otherwise} \end{cases} \tag{S102}$$

Using this expression, the evenness of $\cos(\cdot)$, and $\sin\left(\pm\cos^{-1}(a)\right) = \pm\sqrt{1-a^2}$ reveals after a few lines of algebra that

$$p(\Delta_\ell|x=1) = \frac{e^{\frac{\Delta_\ell}{2}}}{4\kappa\pi I_0(\kappa)\sqrt{\kappa^2 - \frac{\Delta_\ell^2}{4}}}. \tag{S103}$$

An analogous derivation for $x = 2$ leads to

$$p(\Delta_\ell|x=2) = \frac{e^{-\frac{\Delta_\ell}{2}}}{4\kappa\pi I_0(\kappa)\sqrt{\kappa^2 - \frac{\Delta_\ell^2}{4}}}. \tag{S104}$$

With the above, the marginal $p(\Delta_\ell)$ is given by

$$p(\Delta_\ell) = \frac{1}{2}\left(p(\Delta_\ell|x=1) + p(\Delta_\ell|x=2)\right), \tag{S105}$$

where we have used $p(x=1) = p(x=2) = \frac{1}{2}$. Based on these expressions, we have computed $\mathrm{MI}(x;\Delta_\ell)$ by numerically evaluating the required integrals in Eq. (S99) over $\Delta_\ell \in [-2\kappa, 2\kappa]$, using the `integral` function of MATLAB (The MathWorks Inc., 2013).

To find the mutual information for the inference noise model, we have used

$$p\left(\tilde{\Delta}_\ell|\Delta_\ell\right) = \mathcal{N}\left(\tilde{\Delta}_\ell|\Delta_\ell, 2\sigma_{inf}^2\right), \tag{S106}$$

where σ_{inf}^2 is the estimated inference noise variance that appears twice, as the log-likelihood ratio is the difference between two noisy log-likelihoods. This results in the conditional densities

$$p(\tilde{\Delta}_\ell|x) = \int_{-2\kappa}^{2\kappa} p\left(\tilde{\Delta}_\ell|\Delta_\ell\right) p(\Delta_\ell|x)\mathrm{d}\Delta_\ell, \tag{S107}$$

whose values we again found by numerical integration using the `integral` function. The marginal $p\left(\tilde{\Delta}_\ell\right)$ was found as before. Based on these expressions we computed $\mathrm{MI}\left(x; \tilde{\Delta}_\ell\right)$ again numerically, using the `integral` function, but this time over $\tilde{\Delta}_\ell \in [-2\kappa - 4\sqrt{2}\sigma_{inf}, 2\kappa + 4\sqrt{2}\sigma_{inf}]$ to capture four standard deviations of the Gaussian inference noise. How information loss depends on σ_{inf}^2 is shown in Fig. S3D.

Supplemental References

Berger, J. O. (1993). *Statistical Decision Theory and Bayesian Analysis*. Springer Series in Statistics. Springer, 2nd edition.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J. (2012). emcee: The MCMC hammer. arXiv:1202.3665 [astro-ph].

Goodman, J. and Weare, J. (2010). Ensemble samplers with affine invariance. *Communications in Applied Mathematics and Computational Science*, 5(1):65–80.

Jones, E., Oliphant, T., Peterson, P., et al. (2001). SciPy: Open source scientific tools for Python.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., and Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, 46:1004–1017.

The MathWorks Inc. (2013). MATLAB. Version 8.1.0 (R2013a).

Wolfram Research, Inc. (2010). Mathematica. Version 8.0.