

Human confidence reports account for sensory uncertainty but in a non-Bayesian way

William T. Adler^{1*} & Wei Ji Ma^{1,2}

¹Center for Neural Science, ²Department of Psychology,
New York University, New York, NY

*Corresponding author: will.adler@nyu.edu

Humans can meaningfully report their confidence in a perceptual or cognitive decision. It is widely believed that these reports reflect the estimated probability that the decision is correct, but this belief is a hypothesis rather than an established fact. In a pair of perceptual categorization tasks, we tested whether explicit confidence reports reflect the Bayesian posterior probability of being correct, which would require subjects to take sensory uncertainty into account in a specific way. We find that subjects do take sensory uncertainty into account, but that they do so in a way that is inconsistent with the Bayesian hypothesis. Instead, heuristic models provide the best fit to confidence reports. This conclusion is robust to changes in the uncertainty manipulation, task, response modality, additional flexibility in the Bayesian model, and model comparison metric. Finally, we find that generic neural networks trained with error feedback produce confidence reports that are best fit by the same heuristic probabilistic models, suggesting that implementational constraints cause explicit confidence reports to deviate from being Bayesian.

Introduction

People often have a sense of a level of confidence about their decisions. Such a “feeling of knowing”^{1,2} may serve to improve performance in subsequent decisions³, learning¹, and group decision-making⁴. Much recent work has focused on identifying brain regions and neural mechanisms responsible for the computation of confidence in humans⁵⁻⁷, non-human primates⁸⁻¹⁰, and rodents¹¹. In the search for the neural correlates of confidence, the leading premise has been that confidence is Bayesian, i.e., the observer’s estimated probability that a choice is correct^{1,12-14}. In human studies, however, naïve subjects can give a meaningful answer when you ask them to rate their confidence about a decision¹⁵; thus, “confidence” intrinsically means something to people, and it is not a foregone conclusion that this intrinsic sense corresponds to the Bayesian definition. Therefore, we regard the above “definition” as a testable hypothesis about the way the brain computes explicit confidence reports; we use Bayesian decision theory to formalize this hypothesis.

Bayesian decision theory provides a general and often quantitatively accurate account of perceptual decisions in a wide variety of tasks¹⁶⁻¹⁸. According to this theory, the decision-maker combines knowledge about the statistical structure of the world with the present sensory input to compute a posterior probability distribution over possible states of the world. In principle, a confidence report might be derived from the same posterior distribution; this is the hypothesis described above, which we will call the Bayesian Confidence Hypothesis (BCH). The main goal of this paper is to test that hypothesis. Recent studies have attempted to test the BCH^{19,20} but, because of their experimental designs, are unable to meaningfully distinguish the Bayesian model from any other model of confidence.

We test the quantitative predictions of the BCH as we vary the quality of the sensory evidence and the task structure within individuals. We compare Bayesian models against a variety of alternative models, something that is rarely done but very important for the epistemological standing of Bayesian claims^{21,22}. We find that the BCH qualitatively describes human behavior, but that quantitatively, even the most flexible Bayesian model is outperformed by non-Bayesian probabilistic models. To better understand why, we trained neural networks to perform one of our tasks. We found that the same non-Bayesian models provided the best description of the output of the trained networks, suggesting that neural network architecture constrains the computation of confidence.

Results

Experiment 1

During each session, each subject completed two orientation categorization tasks, Tasks A and B. On each trial, a category C was selected randomly (both categories were equally probable), and a stimulus s was drawn from the corresponding stimulus distribution and displayed. The subject categorized the stimulus and simultaneously reported their confidence on a 4-point scale, with a single button press (**Fig. 1a**). The categories were defined by normal distributions on orientation, which differed by task (**Fig. 1b**). In Task A, the distributions had different means ($\pm\mu_C$) and the same standard deviation (σ_C); leftward-tilting stimuli were more likely to be from category 1. Variants of Task A are common in decision-making studies²³. In Task B, the distributions had the same mean (0°) and different standard deviations (σ_1, σ_2); stimuli around the horizontal were more likely to be from category 1. Variants of Task B are less common²⁴⁻²⁶ but have some properties of perceptual organization tasks; for example, a subject may have to detect when a stimulus belongs to a narrow category (e.g., in which two line segments are collinear) that is embedded in a broader category (e.g., in which two line segments are unrelated).

Subjects were highly trained on the categories; during training, we only used highest-reliability stimuli, and we provided trial-to-trial category correctness feedback. Subjects were then tested with 6 different reliability levels, which were chosen randomly on each trial. During testing, correctness feedback was withheld to avoid the possibility that confidence simply reflects a learned mapping between stimuli and the probability of being correct, something that no other confidence studies have done²⁶⁻²⁸.

Because we are interested in subjects' intrinsic computation of confidence, we did not instruct or incentivize them to assign probability ranges to each button (e.g., by using a scoring rule). If we had, we would have essentially been training subjects to use a specific model of confidence.

To ensure that our results were independent of stimulus type, we used two kinds of stimuli. Some subjects saw oriented drifting Gabors; for these subjects, stimulus reliability was manipulated through contrast. Other subjects saw oriented ellipses; for these subjects, stimulus reliability was manipulated through ellipse elongation (**Fig. 1c**). We found no major differences in model rankings between Gabor and ellipse subjects, therefore we will make no distinctions between the groups (Supplementary Information).

For modeling purposes, we assume that the observer's internal representation of the stimulus is a noisy measurement x , drawn from a Gaussian distribution with mean s and s.d. σ (**Fig. 1d,e**). In the model, σ is a function of stimulus reliability (Supplementary Information).

Bayesian model

A Bayes-optimal observer uses knowledge of the generative model to make a decision that maximizes the probability of being correct. Here, when the measurement on a given trial is x , this strategy amounts to

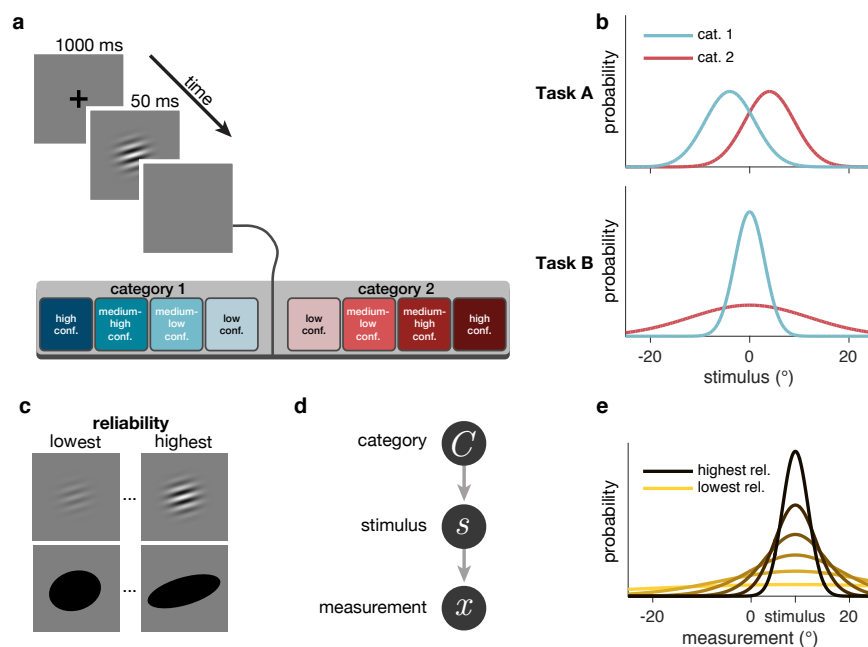


Figure 1: Task design and generative model. (a) Schematic of a test block trial. After stimulus offset, subjects reported category and confidence level with a single button press. (b) Stimulus distributions for Tasks A and B. (c) Examples of low and high reliability stimuli. Six (out of eleven) subjects saw drifting Gabors, and five subjects saw ellipses. (d) Generative model. (e) Example measurement distributions at different reliability levels. In all models (except Linear Neural), the measurement is assumed to be drawn from a Gaussian distribution centered on the true stimulus, with s.d. dependent on reliability.

choosing the category C for which the posterior probability $p(C | x)$ is highest. This is equivalent to reporting category 1 when the log posterior ratio, $d = \log \frac{p(C=1|x)}{p(C=2|x)}$, is positive.

In Task A, d is $d_A = \frac{2x\mu_C}{\sigma^2 + \sigma_C^2}$. Therefore, the ideal observer reports category 1 when x is positive; this is the structure of many psychophysical tasks²⁹. In Task B, however, d is $d_B = \frac{1}{2} \log \frac{\sigma^2 + \sigma_2^2}{\sigma^2 + \sigma_1^2} - \frac{\sigma_2^2 - \sigma_1^2}{2(\sigma^2 + \sigma_1^2)(\sigma^2 + \sigma_2^2)} x^2$; the observer needs both x and σ in order to make an optimal decision. See Supplementary Information for derivations of d_A and d_B .

From the point of view of the observer, σ is the trial-to-trial level of sensory uncertainty associated with the measurement³⁰. In a minor variation of the optimal observer, we allow for the possibility that the observer's prior belief over category, $p(C)$, is different from the true value of (0.5, 0.5); this adds a constant to d_A and d_B .

We introduce the Bayesian Confidence Hypothesis (BCH), stating that confidence reports depend on the internal representation of the stimulus (here x) only via d . In the BCH, the observer chooses a response by comparing d to a set of category and confidence boundaries. For example, whenever d falls within a certain range, the observer presses the “medium-low confidence, category 2” button. The BCH is thus an extension of the choice model described above, wherein the value of d is used to compute confidence as well as chosen category. Another way of thinking about this is: Bayesian models assume that subjects compute d in order to make an optimal choice. Assuming people compute d at all, are they able to use it to report confidence as well? We refer to the Bayesian model here as simply “Bayes.” (We tested several more constrained versions of this model, which are described in Supplementary Information.)

The observer's decision can be summarized as a mapping from a combination of a measurement and an uncertainty level (x, σ) to a response that indicates both category and confidence. We can visualize this

mapping as in **Figure 2**, first column. It is clear that the pattern of decision boundaries in the BCH is qualitatively very different between Task A and Task B. In Task A, the decision boundaries are quadratic functions of uncertainty; confidence decreases monotonically with uncertainty and increases with the distance of the measurement from 0. In Task B, the decision boundaries are neither linear nor quadratic and can even be non-monotonic functions of uncertainty.

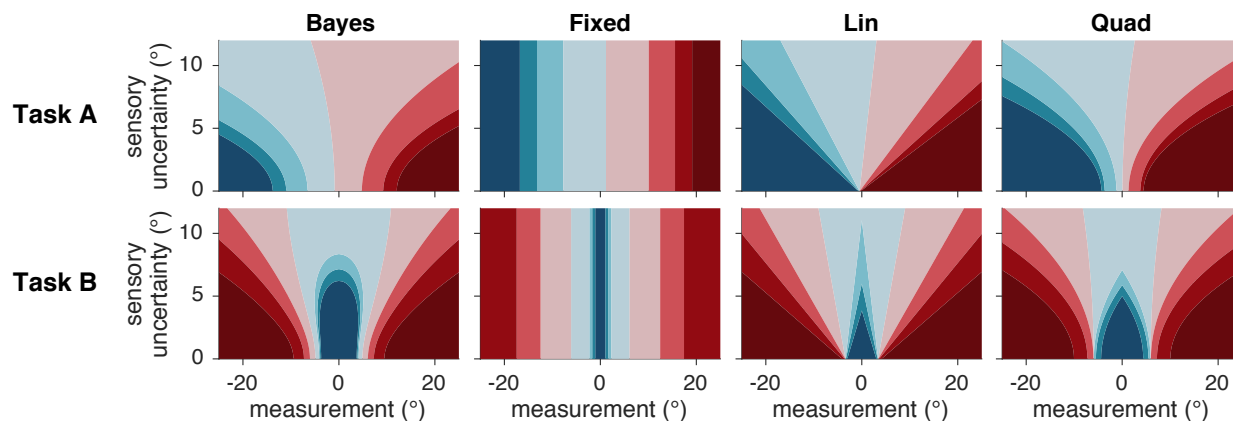


Figure 2: Decision rules/mappings in four models. Each model corresponds to a different mapping from a measurement and uncertainty level to a category and confidence response. Colors correspond to category and confidence response, as in **Figure 1a**. Plots were generated from the mean of subject 4's posterior distribution over parameters.

Alternative models

At first glance, it seems obvious that sensory uncertainty is relevant to the computation of confidence. However, this is by no means a given; in fact, a prominent proposal is that confidence is based on the distance between the measurement and the decision boundary, without any role for sensory uncertainty^{10,11,31}. Therefore, we tested a model (Fixed) in which the response is a function of the measurement alone (equivalent to a maximum likelihood estimate of the stimulus orientation), and not of the uncertainty of that measurement (**Fig. 2**, second column).

We also tested heuristic models in which the subject uses their knowledge of their sensory uncertainty but does not compute a posterior distribution over category. We have previously classified such models as probabilistic non-Bayesian³². In the Orientation Estimation model, subjects base their response on a maximum a posteriori estimate of orientation (rather than category), using the mixture of the two stimulus distributions as a prior distribution. In the Linear Neural model, subjects base their response on a linear function of the output of a hypothetical population of neurons.

We derived two additional probabilistic non-Bayesian models, Lin and Quad, from the observation that the Bayesian decision criteria are an approximately linear function of uncertainty in some measurement regimes and approximately quadratic in others. These models are able to produce approximately Bayesian behavior without actually performing any computation of the posterior. In Lin and Quad, subjects base their response on a linear or a quadratic function of x and σ , respectively (Supplementary Information). A comparison of the Lin and Quad columns to the Bayes column in **Figure 2** demonstrates that Lin and Quad can approximate the Bayesian mapping from (x, σ) to response despite not being based on the Bayesian decision variable. All of the models we tested were variants of the six models described so far (Bayes, Fixed, Orientation Estimation, Linear Neural, Lin, Quad).

Each trial consists of the experimentally determined orientation and reliability level and the subject's category and confidence response (an integer between 1 and 8). This is a very rich data set, which we summarize in **Figure 3**. We find the following effects: performance and confidence increase as a function of reliability

(Fig. 3a,b,h,i), and high-confidence reports are less frequent than low-confidence reports (Fig. 3e,f). Note Figure 3c,d especially; this is the projection of the data that we will use to demonstrate model fits for the rest of this paper. We use this projection because the vertical axis (mean button press) most closely approximates the form of the raw data. Additionally, because our models are differentiated by how they use uncertainty, it is informative to plot how response changes as a function of reliability, in addition to category and task.

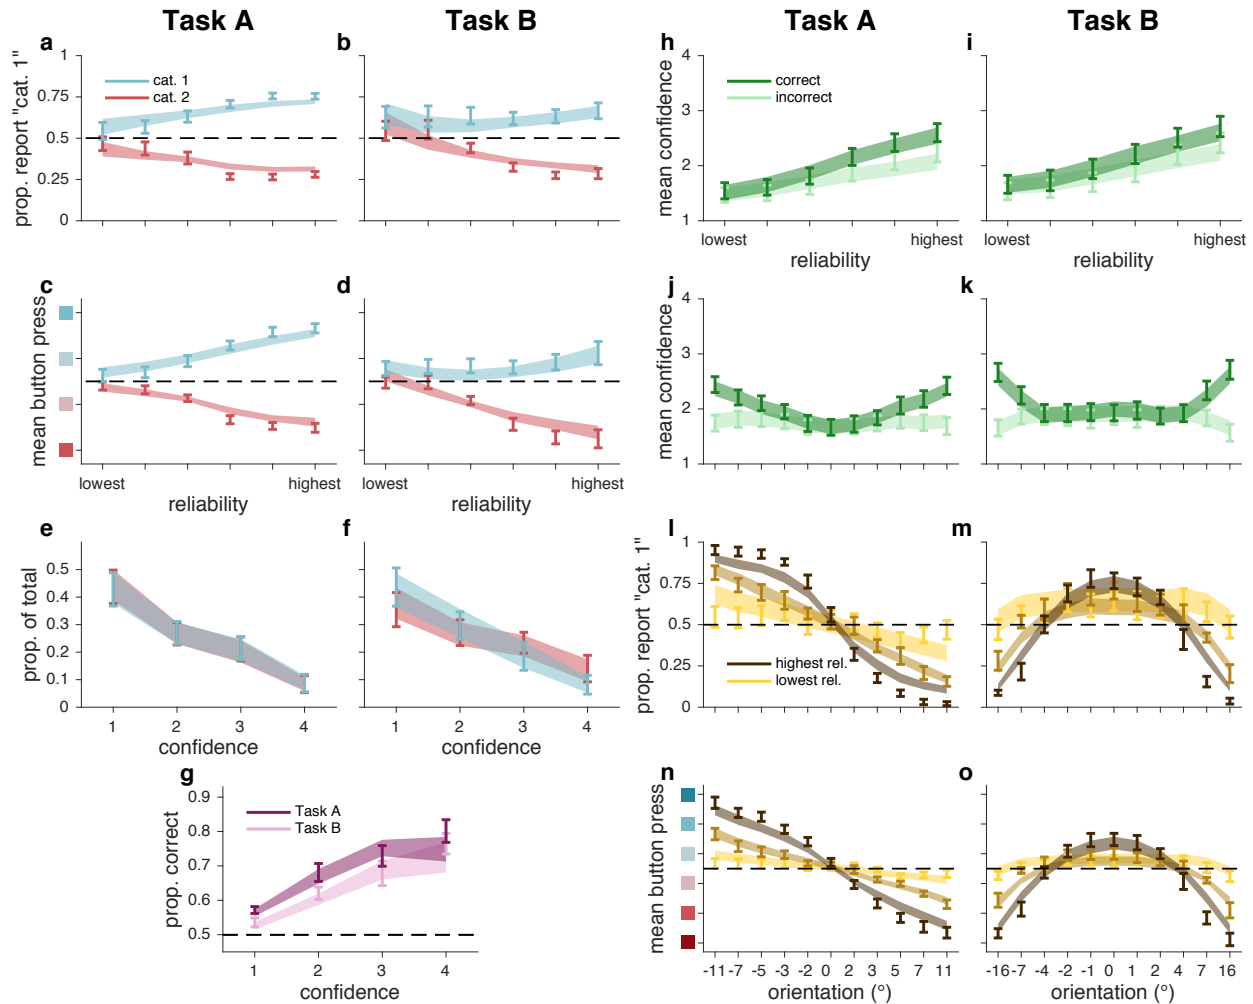


Figure 3: Behavioral data and fits from best model (Quad), experiment 1. Error bars represent ± 1 s.e.m. across 11 subjects. Shaded regions represent ± 1 s.e.m. on model fits (Supplementary Information). (a,b) Proportion "category 1" reports as a function of stimulus reliability and true category. (c,d) Mean button press as a function of stimulus reliability and true category. (e,f) Normalized histogram of confidence reports for both true categories. (g) Proportion correct category reports as a function of confidence report and task. (h,i) Mean confidence as a function of stimulus reliability and correctness. (j,k) Mean confidence as a function of stimulus orientation and reliability. (l,m) Proportion "category 1" reports as a function of stimulus orientation and reliability. (n,o) Mean button press as a function of stimulus orientation and reliability. (c,d,n,o) Vertical axis label colors correspond to button presses, as in Figure 1a. (l-o) For clarity, only 3 of 6 reliability levels are shown, although models were fit to all reliability levels.

Model comparison

We used Markov Chain Monte Carlo (MCMC) sampling to fit models to raw individual-subject data. To account for overfitting, we compared models using leave-one-out cross-validated log likelihood scores (LOO) computed with the full posteriors obtained through MCMC³³. A model recovery analysis ensured that our models are meaningfully distinguishable (Supplementary Information, **Fig. S3**). Unless otherwise noted, models were fit jointly to Task A and B category and confidence responses. Models are described in more detail in Supplementary Information.

Use of sensory uncertainty. We first compared Bayes to the Fixed model, in which the observer does not take trial-to-trial sensory uncertainty into account (**Fig. 4**). Fixed provides a poor fit to the data, indicating that observers use not only a point estimate of their measurement, but also their uncertainty about that measurement. Bayes outperforms Fixed by a summed LOO difference (median and 95% CI of bootstrapped sums across subjects) of 2265 [498, 4253]. For the rest of this paper, we will report model comparison results using this format.

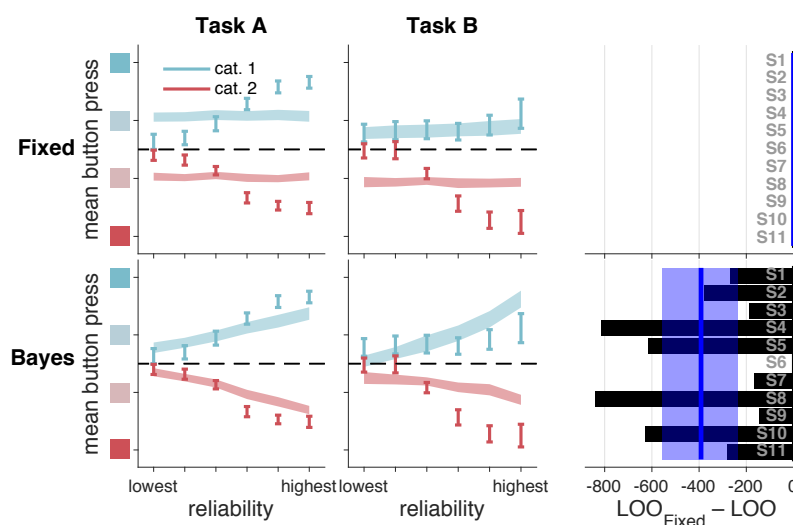


Figure 4: Model fits and model comparison for models Fixed and Bayes. Left and middle columns: model fits to mean button press as a function of reliability, true category, and task. Error bars represent ± 1 s.e.m. across 11 subjects. Shaded regions represent ± 1 s.e.m. on model fits, with each model on a separate row. Right column: LOO model comparison. Bars represent individual subject LOO scores for Bayes, relative to Fixed. Negative (leftward) values indicate that, for that subject, Bayes had a higher (better) LOO score than Fixed. Blue lines and shaded regions represent, respectively, medians and 95% CI of bootstrapped mean LOO differences across subjects. These values are equal to the summed LOO differences reported in the text divided by the number of subjects.

Although Bayes fits better than Fixed, it still shows systematic deviations from the data, especially at high reliabilities. (Because we fit our models to all of the raw data and because boundary parameters are shared across all reliability levels, the fit to high-reliability trials is constrained by the fit to low-reliability trials.)

Noisy log posterior ratio. To see if we could improve Bayes's fit, we tried a version that included decision noise, i.e. noise on the log posterior ratio d . We assumed that this noise takes the form of additive zero-mean Gaussian noise with s.d. σ_d . This is almost equivalent to the probability of a response being a logistic (softmax) function of d ³⁴. Adding d noise improves the Bayesian model fit by 804 [510, 1134] (**Table S1**).

For the rest of the reported fits to human behavior, we will only consider this version of Bayes with d noise, and will refer to this model as Bayes- dN . We will refer to Bayes- dN , Fixed, Orientation Estimation, Linear Neural, Lin, and Quad, when fitted jointly to category and confidence data from Tasks A and B, as our core models.

Heuristic models. Orientation Estimation performs worse than Bayes- dN by 2041 [385, 3623] (**Fig. 5**, second row). The intuition for one way that this model fails is as follows: at low levels of reliability, the MAP estimate is heavily influenced by the prior and tends to be very close to the prior mean (0°). This explains why, in Task B, there is a bias towards reporting “high confidence, category 1” at low reliability. Linear Neural performs about as well as Bayes- dN , with summed LOO differences of 1188 [-588, 2704], and the fits to the summary statistics are qualitatively poor (**Fig. 5**, third row).

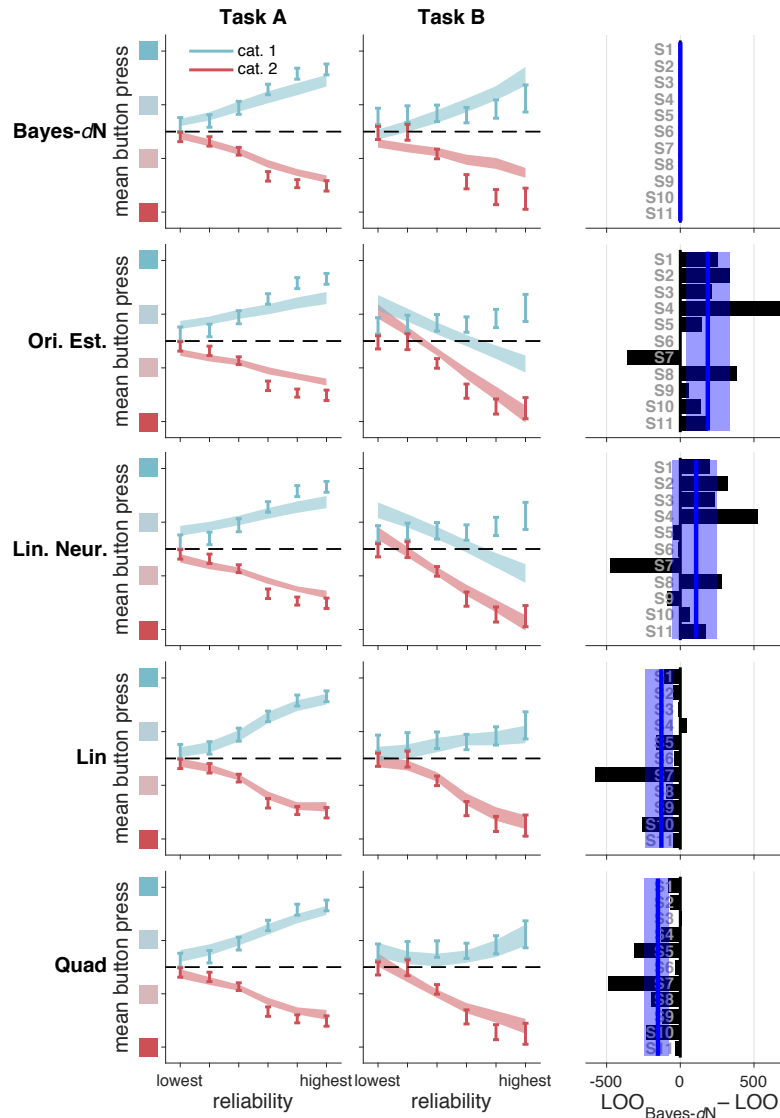


Figure 5: Model fits and model comparison for Bayes- dN and heuristic models. Left and middle columns: as in **Figure 4**. Right column: bars represent individual subject LOO scores for each model, relative to Bayes- dN . Negative (leftward) values indicate that, for that subject, the model in the corresponding row had a higher (better) LOO score than Bayes- dN . Blue lines and shaded regions: as in **Figure 4**.

Finally, Lin and Quad outperform Bayes- dN by 1398 [571, 2644] and 1667 [858, 2698], respectively. Both models provide qualitatively better fits, especially at high reliabilities (compare **Fig. 5**, first row, to **Fig. 5**, fourth and fifth rows), and strongly tilted orientations (compare **Fig. S7n,o** to **Fig. S11n,o** and **Fig. 3n,o**).

We summarize the performance of our core models in **Fig. 6**. Noting that a LOO difference of more than 5 is considered to be very strong evidence³⁵, the heuristic models Lin and Quad perform much better than Bayes- dN . Furthermore, we can decisively rule out Fixed. We will now test variants of our core models.

Non-parametric relationship between reliability and σ .

One potential criticism of our fitting procedure is that we assumed a parameterized relationship between reliability and σ (Supplementary Information). To see if our results were dependent on that assumption, we modified the models such that σ was non-parametric (i.e., there was a free parameter for σ at each level of reliability). With this feature added to our core models, Quad still fits better than Bayes- dN by 1676 [839, 2730], and better than Fixed by 6097 [4323, 7901] (**Table S1**). This feature improved Quad’s performance by 325 [141, 535]. For the rest of this paper, we will only report the fits of Bayes- dN , the best-fitting non-Bayesian model, and Fixed. See supplementary figures and tables for all other model fits.

Incorrect assumptions about the generative model. Suboptimal behavior can be produced by optimal inference using incorrect generative models, a phenomenon known as “model mismatch.”^{36–38} Up to now, Bayes- dN has assumed that observers have accurate knowledge of the parameters of the generative model. To test whether this assumption prevents Bayes- dN from fitting the data well, we tested a series of Bayesian models in which the observer has inaccurate knowledge of the generative model.

Bayes- dN assumed that, because subjects were well-trained, they knew the true values of σ_C , σ_1 , and σ_2 , the standard deviations of the stimulus distributions. We tested a model in which these values were free parameters, rather than fixed to the true value. We would expect these free parameters to improve the fit of Bayes- dN in the case where subjects were not trained enough to sufficiently learn the stimulus distributions. This feature improves Bayes- dN ’s fit by 908 [318, 1661], but it still underperforms Quad by 768 [399, 1144] (**Table S1**).

Previous models also assumed that subjects had full knowledge of their own measurement noise; the σ used in the computation of d was identical to the σ that determined their measurement noise. We tested models in which we fit $\sigma_{\text{measurement}}$ and $\sigma_{\text{inference}}$ as two independent functions of reliability³⁶. This feature improves Bayes- dN ’s fit by 1310 [580, 2175], but it still underperforms Quad by 362 [162, 602] (**Table S1**).

Weighted mixture of precision and the probability of being correct. A recent paper³⁹ assumed that confidence is a weighted mixture of the probability of being correct and sensory precision $\frac{1}{\sigma^2}$. We tested one such model, which fit better than Fixed by 3059 [758, 5528] but still underperforms Lin by 3478 [2211, 5020] (**Table S1**).

Separate fits to Tasks A and B. In order to determine whether model rankings were primarily due to differences in one of the two tasks, we fit our models to each task individually. In Task A, Quad fits better than Bayes- dN by 581 [278, 938], and better than Fixed by 3534 [2529, 4552] (**Fig. S13** and **Table S2**). In Task B, Quad fits better than Bayes- dN by 978 [406, 1756], and better than Fixed by 3234 [2099, 4390] (**Fig. S14** and **Table S3**).

Fits to category choice data only. In order to see whether our results were peculiar to combined category and confidence responses, we fit our models to the category choices only. Lin fits better than Bayes- dN by 595 [311, 927], and better than Fixed by 1690 [976, 2534] (**Fig. S15** and **Table S4**).

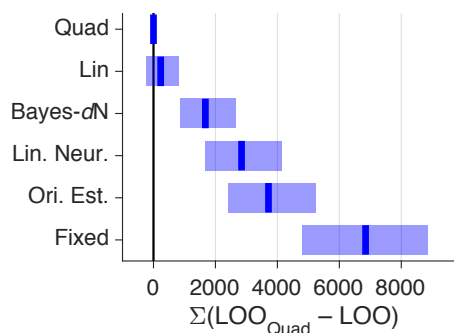


Figure 6: Comparison of core models, experiment 1. Models were fit jointly to Task A and B category and confidence responses. Blue lines and shaded regions represent, respectively, medians and 95% CI of bootstrapped summed LOO differences across subjects. LOO differences for these and other models are shown in **Figure S12a**.

Fits to Task B only, with noise parameters fitted from Task A. To confirm that the fitted values of sensory uncertainty in the probabilistic models are meaningful, we treated Task A as an independent experiment to measure subjects' sensory noise. The category choice data from Task A can be used to determine the four uncertainty parameters. We fitted Fixed with a decision boundary of 0° (equivalent to a Bayesian choice model with no prior), using maximum likelihood estimation. We fixed these parameters and used them to fit our models to Task B category and confidence responses. Lin fits better than Bayes- dN by 1773 [451, 2845], and better than Fixed by 5016 [3090, 6727] (**Fig. S16** and **Table S5**).

Experiment 2: Separate category and confidence responses. There has been some recent debate as to whether it is more appropriate to collect choice and confidence with a single motor response (as described above), or separate responses^{20,40–42}. Aitchison *et al.*¹⁹ found that confidence appears more Bayesian when subjects use separate responses. To confirm this, we ran a second experiment in which subjects chose a category by pressing one of two buttons, then reported confidence by pressing one of four buttons. Aitchison *et al.* also provided correctness feedback on every trial; in order to ensure that we could compare our results to theirs, we also provided correctness feedback in this experiment, even though this manipulation was not of primary interest. After fitting our core models, our results did not differ substantially from experiment 1: Lin fits better than Bayes- dN by 396 [186, 622], and better than Fixed by 2095 [1344, 2889] (**Fig. S17** and **Table S6**).

Experiment 3: Task B only. It is possible that subjects behave suboptimally when they have to do multiple tasks in a session; in other words, perhaps one task “corrupts” the other. To explore this possibility, we ran an experiment in which subjects completed Task B only. Quad fits better than Bayes- dN by 1361 [777, 2022], and better than Fixed by 7326 [4905, 9955] (**Fig. S18** and **Table S7**). In experiments 2 and 3, subjects only saw drifting Gabors; we did not use ellipses.

We also fit only the choice data, and found that Lin fits about as well as Bayes- dN , with summed LOO differences of 117 [-76, 436], and better than Fixed by 1084 [619, 1675] (**Fig. S19** and **Table S8**). This approximately replicates our previously published results²⁶.

Model comparison metric. None of our model comparison results depend on our choice of metric: in all three experiments, model rankings changed negligibly if we used AIC, BIC, AICc, or WAIC instead of LOO (Supplementary Information).

Neural network

Taken together, the model comparisons in experiments 1 to 3 convince us that there is no obvious way to explain human confidence ratings as Bayesian. Does this mean that the normative framework must be entirely rejected? We should instead consider that implementational constraints may restrict the brain's ability to perform fully Bayesian computation^{21,22}. To explore this possibility, we adopted an approach related to work by Yamins *et al.* (2014), who maximize a network's categorization performance without constraining it to neural data; the resulting model was highly predictive of cortical response⁴³. Here, we trained networks as if they were naïve human subjects, without constraining them to human behavioral data; the output from the networks have behavioral model rankings that are similar to those of the human subjects.

We trained biologically plausible feedforward neural networks (**Fig. 7a**) to perform Task B with online category correctness feedback⁴⁴. The trained networks perform the task in a way that is qualitatively similar to the subjects (**Fig. 7b,c**). We then fit their output with the same models that we used to fit subject data, and found that Lin is the best-fitting model, outperforming Bayes by summed AIC differences of 11662 [10011, 13522] (**Fig. 7d**, blue). The overall rankings of all models fit to the trained networks (**Fig. S5**, blue) was very similar to that of the human models (**Fig. S14a**), with a Spearman's rank correlation coefficient of 0.85. This convergence of neural network and human behavior suggests that neural architecture may impose constraints on the type of behavior that can be produced (Supplementary Information).

An alternative explanation of this result is that Bayes is too inflexible to fit any behavioral dataset based on neural activity. To rule out this possibility, we decoded optimal posterior probabilities from input unit activity on a per-trial basis, mapped these onto button presses using quantiles, and fit the behavioral models. Bayes provides the best fit, fitting these datasets better than Lin by 5739 [3935, 8045] and better than Quad by 800 [479, 1412] (**Fig. 7d**, black). Thus, the fact that Lin wins is not due to Bayes being generally inflexible.

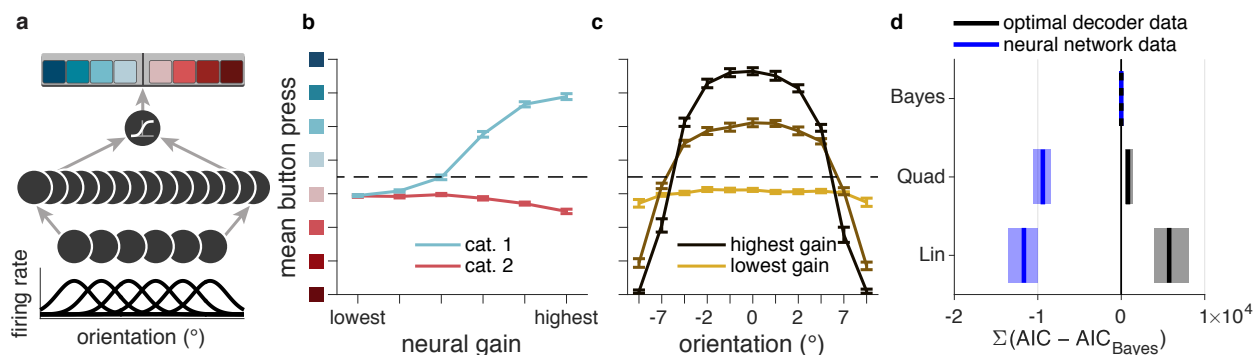


Figure 7: Neural network architecture, model comparison, and task behavior. (a) Feedforward neural network architecture. Input units were independent Poisson neurons with Gaussian tuning curves that were evenly spaced and of identical width. Input units were connected, all-to-all, to hidden rectified linear units. Hidden units were connected to a sigmoidal output unit. The output was mapped onto a category and confidence response (as in **Fig. 1a**) using eight quantiles (Supplementary Information). (b) Task behavior of trained neural networks. Mean button press as a function of neural gain and true category. Compare to **Figure 3d**. (c) Mean button press as a function of stimulus orientation and neural gain. Compare to **Figure 3o**. (d) Model comparison (as in **Fig. 6**) for data generated by the optimal decoder, and by trained neural networks. Fixed, Orientation Estimation, and Linear Neural were also fit to those data, but are not shown because they all fit more poorly than Bayes (**Fig. S5**).

Discussion

Although people can report subjective feelings of confidence, the computations that produce this feeling are not well-understood. Confidence has been defined as the observer’s computed posterior probability that a decision is correct. However, this hypothesis has not been fully tested. We used model comparison to investigate the computational underpinnings of human confidence reports, fitting a total of 75 models from 6 distinct model families. We also trained neural networks to perform a perceptual task, treating the network output as if it were subject-generated data for the purpose of model comparison⁴⁴. We carried out a strong and comprehensive test of a set of cognitive models, varying task components such as stimulus reliability and stimulus distributions²⁷.

Our first finding is that, like the optimal observer, subjects use knowledge of their sensory uncertainty when reporting confidence in a categorical decision; models in which the subject ignores their sensory uncertainty provided a poor fit to the data (**Fig. 4**). Our second finding is that, unlike the optimal observer, subjects do not appear to use knowledge of their sensory uncertainty in a Bayesian way. Instead, heuristic models that approximate Bayesian computation—but do not compute a posterior probability over category—outperform the Bayesian models in a variety of experimental contexts (**Fig. 5**, compare top row to bottom two rows). This result continued to hold after we relaxed assumptions about the relationship between reliability and noise, and about the subject’s knowledge of the generating model. We accounted for the fact that our models had different amounts of flexibility by using a wide array of model comparison metrics and by showing that our models were meaningfully distinguishable (Supplementary Information).

We trained neural networks to perform one of our tasks. Although the training procedure necessarily differed

from that of the humans, we found that the trained networks produced confidence responses that, like the human data, were best fit by heuristic models. This convergent result suggests that the structure of the neural network—and by extension, the structure of the brain—limits its ability to produce accurate posterior estimates in categorization tasks.

While we believe that we have tested a large proportion of the models that could reasonably be considered “Bayesian” for these tasks, there are many other heuristic models that we did not test that may describe the data better. So our main message is not that Lin and Quad are general descriptions of human confidence, but that human behavior is best described by models that are non-Bayesian. Some may argue that non-Bayesian models should be rejected because they are not generalizable. Bayesian models derive their generalizability from their normative nature: in any task, one can determine the performance-maximizing strategy. Although this property makes Bayesian models attractive and powerful, it is not clear that this property should override a bad fit in all cases.

Moreover, performance maximization is only one of several ecologically relevant organizing principles. The brain is also limited by the kinds of operations that neurons can perform and the ways by which organisms learn^{21,22}; our neural network analysis suggests that architectural or learning constraints may cause the brain to deviate from Bayesian computations. Future work could investigate the possibility that the brain is near-optimal under implementational constraints; this would connect Marrian levels within a single rational framework⁴⁵.

We will now describe how our results relate to recently published experimental findings. Rahnev *et al.*³¹ reported that subjective decision criteria are fixed across conditions of uncertainty. However, their study did not test models in which the criteria were a function of uncertainty, so they cannot make this conclusion very strongly. Additionally, their study used visibility ratings, which differ from confidence ratings⁴⁶. Finally, their results may be specific to the case where sensory uncertainty is a function of attention rather than stimulus reliability. We leave this question for future work.

Sanders *et al.*²⁰ reported that confidence has a “statistical” nature. However, their experiment was unable to determine whether confidence is probabilistic or Bayesian¹⁷, because the stimuli vary along only one dimension. As noted by Aitchison *et al.*, to distinguish models of confidence, the experimenter must use stimuli that are characterized by two dimensions (e.g., contrast and orientation)¹⁹. This is because, when fitting models that map from an internal variable to an integer confidence rating, it is impossible to distinguish between two internal variables that are monotonically related (in the case of Sanders *et al.*, the measurement and the posterior probability of being correct). Therefore, the only alternative model proposed by Sanders *et al.* is based on reaction time, rather than on the presented stimuli.

Like the present study, Aitchison *et al.*¹⁹ found evidence that confidence reports may emerge from heuristic computations. However, they sampled stimuli from only a small region of their two-dimensional space, where model predictions may not vary greatly. Therefore, their stimulus set did not allow for the models to be strongly distinguished. Furthermore, although they tested for *Bayesian* computation, they did not test for *probabilistic* computation (i.e., whether observers take sensory uncertainty into account on a trial-to-trial basis). Such a test requires that the experimenter vary the reliability of the stimulus feature of interest.

In this study, we only consider explicit confidence ratings, which differ from the implicit confidence that can be gathered from non-human animals¹³ (e.g., by measuring how frequently they decline to make a difficult choice⁸, or how long they will wait for a reward¹¹). It is possible that implicit confidence might be more Bayesian⁴⁷. However, to our knowledge, there are no studies that both measure implicit confidence and allow for comparison of the models presented here; this is an opportunity for future research.

What do our findings tell us about the neural basis of confidence? Previous studies have found that neural activity in some brain areas (e.g., human medial temporal lobe⁷ and prefrontal cortex⁴⁸, monkey lateral intraparietal cortex⁸ and pulvinar¹⁰, rodent orbitofrontal cortex¹¹) is associated with behavioral indicators of confidence, and/or with the distance of a stimulus to a decision boundary. However, such studies mostly

used stimuli that vary along a single dimension (e.g., net retinal dot motion energy, mixture of two odors). Because measurement is indistinguishable from the probability of being correct in these classes of tasks,¹⁹ neural activity associated with confidence may represent either the measurement or the probability of being correct. In addition to the recommendation of Aitchison *et al.* to distinguish between these possibilities by varying stimuli along two dimensions, we recommend fitting both Bayesian and non-Bayesian probabilistic models to behavior. Many physiological studies of decision-making focus on correlating neural activity to parameters of behavioral models. This approach only makes sense when the behavioral model is a good description of the behavior. Our results suggest that the Bayesian model is a relatively poor description of confidence behavior. Therefore, the proposal to do this kind of correlational analysis with parameters of the Bayesian confidence models¹² should be viewed with skepticism.

ACKNOWLEDGEMENTS

The authors would like to thank A. Emin Orhan for assistance on the neural network analysis, and Luigi Acerbi for helpful ideas and tools related to model fitting and model comparison. We would also like to thank Luigi Acerbi, Rachel N. Denison, Andra Mihali, Bas van Opheusden, and Aspen H. Yoo for helpful conversations and comments on the manuscript. This material was based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1342536.

AUTHOR CONTRIBUTIONS

W.T.A. and W.J.M. designed the experiments, analyzed the data, and wrote the manuscript. W.T.A. performed the experiments.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

References

- [1] Meyniel, F., Sigman, M. & Mainen, Z. F. Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron* **88**, 78–92 (2015).
- [2] Brown, A. S. A review of the tip-of-the-tongue experience. *Psychol. Bull.* **109**, 204–223 (1991).
- [3] Persaud, N., McLeod, P. & Cowey, A. Post-decision wagering objectively measures awareness. *Nat. Neurosci.* **10**, 257–261 (2007).
- [4] Bahrami, B., Olsen, K., Latham, P. E. & Roepstorff, A. Optimally interacting minds. *Science* (2010).
- [5] Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J. & Rees, G. Relating Introspective Accuracy to Individual Differences in Brain Structure. *Science* **329**, 1541–1543 (2010).
- [6] Fleming, S. M. & Dolan, R. J. The neural basis of metacognitive ability. *Phil. Trans. R. Soc. B* **367**, 1338–1349 (2012).
- [7] Rutishauser, U. *et al.* Representation of retrieval confidence by single neurons in the human medial temporal lobe. *Nat. Neurosci.* **18**, 1041–1050 (2015).
- [8] Kiani, R. & Shadlen, M. N. Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science* **324**, 759–764 (2009).
- [9] Fetsch, C. R., Kiani, R., Newsome, W. T. & Shadlen, M. N. Effects of cortical microstimulation on confidence in a perceptual decision. *Neuron* **83**, 797–804 (2014).
- [10] Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T. & Miyamoto, A. Responses of pulvinar neurons reflect a subject’s confidence in visual categorization. *Nat. Neurosci.* **16**, 749–755 (2013).

- [11] Kepecs, A., Uchida, N., Zariwala, H. A. & Mainen, Z. F. Neural correlates, computation and behavioural impact of decision confidence. *Nature* **455**, 227–231 (2008).
- [12] Pouget, A., Drugowitsch, J. & Kepecs, A. Confidence and certainty: distinct probabilistic quantities for different goals. *Nat. Neurosci.* **19**, 366–374 (2016).
- [13] Kepecs, A. & Mainen, Z. F. A computational framework for the study of confidence in humans and animals. *Phil. Trans. R. Soc. B* **367**, 1322–1337 (2012).
- [14] Drugowitsch, J., Moreno-Bote, R. & Pouget, A. Relation between Belief and Performance in Perceptual Decision Making. *PLoS ONE* **9**, e96511 (2014).
- [15] Peirce, C. S. & Jastrow, J. On Small Differences in Sensation. *Memoirs of the National Academy of Sciences* **3**, 73–83 (1884).
- [16] Knill, D. C. & Richards, W. *Perception as Bayesian Inference* (Cambridge University Press, 1996).
- [17] Ma, W. J. & Jazayeri, M. Neural coding of uncertainty and probability. *Annu. Rev. Neurosci.* (2014).
- [18] Kording, K. Decision Theory: What "Should" the Nervous System Do? *Science* **318**, 606–610 (2007).
- [19] Aitchison, L., Bang, D., Bahrami, B. & Latham, P. E. Doubly Bayesian Analysis of Confidence in Perceptual Decision-Making. *PLoS Comput. Biol.* **11**, e1004519–23 (2015).
- [20] Sanders, J. I., Hangya, B. & Kepecs, A. Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron* **90**, 499–506 (2016).
- [21] Bowers, J. S. & Davis, C. J. Bayesian just-so stories in psychology and neuroscience. *Psychol. Bull.* **138**, 389–414 (2012).
- [22] Jones, M. & Love, B. C. Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behav. Brain. Sci.* **34**, 169–188 (2011).
- [23] Britten, K. H., Shadlen, M. N., Newsome, W. T. & Movshon, J. A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* **12**, 4745–4765 (1992).
- [24] Liu, Z., Knill, D. C. & Kersten, D. Object classification for human and ideal observers. *Vis. Res.* **35**, 549–568 (1995).
- [25] Sanborn, A. N., Griffiths, T. L. & Shiffrin, R. M. Uncovering mental representations with Markov chain Monte Carlo. *Cogn. Psychol.* **60**, 63–106 (2010).
- [26] Qamar, A. T. *et al.* Trial-to-trial, uncertainty-based adjustment of decision boundaries in visual categorization. *PNAS* **110**, 20332–20337 (2013).
- [27] Maloney, L. T. & Mamassian, P. Bayesian decision theory as a model of human visual perception: testing Bayesian transfer. *Vis. Neurosci.* **26**, 147–155 (2009).
- [28] Körding, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244–247 (2004).
- [29] Green, D. M. & Swets, J. A. *Signal Detection Theory and Psychophysics* (1966).
- [30] Ma, W. J. Signal detection theory, uncertainty, and Poisson-like population codes. *Vis. Res.* **50**, 2308–2319 (2010).
- [31] Rahnev, D. *et al.* Attention induces conservative subjective biases in visual perception. *Nat. Neurosci.* **14**, 1513–1515 (2011).
- [32] Ma, W. J. Organizing probabilistic models of perception. *Trends. Cogn. Sci.* **16**, 511–518 (2012).

- [33] Vehtari, A., Gelman, A. & Gabry, J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *arXiv* (2015). [1507.04544v4](https://arxiv.org/abs/1507.04544v4).
- [34] Keshvari, S., van den Berg, R. & Ma, W. J. Probabilistic Computation in Human Perception under Variability in Encoding Precision. *PLoS ONE* **7**, e40216–9 (2012).
- [35] Kass, R. E. & Raftery, A. E. Bayes Factors. *J. Am. Stat. Assoc.* **90**, 773–795 (2012).
- [36] Acerbi, L., Vijayakumar, S. & Wolpert, D. M. On the Origins of Suboptimality in Human Probabilistic Inference. *PLoS Comput. Biol.* **10**, e1003661 (2014).
- [37] Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E. & Pouget, A. Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron* **74**, 30–39 (2012).
- [38] Orhan, A. E. & Jacobs, R. A. Are Performance Limitations in Visual Short-Term Memory Tasks Due to Capacity Limitations or Model Mismatch? *arXiv* (2014). [1407.0644v1](https://arxiv.org/abs/1407.0644v1).
- [39] Navajas, J. *et al.* The idiosyncratic nature of confidence. *bioRxiv* (2017).
- [40] Kiani, R., Corthell, L. & Shadlen, M. N. Choice Certainty Is Informed by Both Evidence and Decision Time. *Neuron* **84**, 1329–1342 (2014).
- [41] Navajas, J., Bahrami, B. & Latham, P. E. Post-decisional accounts of biases in confidence. *Curr. Opin. Behav. Sci.* **11**, 55–60 (2016).
- [42] Wilimzig, C., Tsuchiya, N., Fahle, M., Einhäuser, W. & Koch, C. Spatial attention increases performance but not subjective confidence in a discrimination task. *Journal of Vision* **8**, 7.1–10 (2008).
- [43] Yamins, D. L. K. *et al.* Performance-optimized hierarchical models predict neural responses in higher visual cortex. *P Natl Acad Sci Usa* **111**, 8619–8624 (2014).
- [44] Orhan, A. E. & Ma, W. J. The Inevitability of Probability: Probabilistic Inference in Generic Neural Networks Trained with Non-Probabilistic Feedback. *arXiv* (2016). [1601.03060v1](https://arxiv.org/abs/1601.03060v1).
- [45] Marr, D. *Vision. A Computational Investigation Into the Human Representation and Processing of Visual Information* (MIT Press, 1982).
- [46] Rausch, M. & Zehetleitner, M. Visibility Is Not Equivalent to Confidence in a Low Contrast Orientation Discrimination Task. *Front. Psychol.* **7**, 47–15 (2016).
- [47] Chen, S. Y., Ross, B. H. & Murphy, G. L. Implicit and explicit processes in category-based induction: Is induction best when we don't think? *Journal of Experimental ...* (2014).
- [48] Fleming, S. M., Huijgen, J. & Dolan, R. J. Prefrontal Contributions to Metacognition in Perceptual Decision Making. *J. Neurosci.* **32**, 6117–6125 (2012).
- [49] Pelli, D. G. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10**, 437–442 (1997).
- [50] Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–436 (1997).
- [51] Naka, K. I. & Rushton, W. A. S-potentials from luminosity units in the retina of fish (Cyprinidae). *The Journal of Physiology* **185**, 587–599 (1966).
- [52] DiMattina, C. Comparing models of contrast gain using psychophysical experiments. *Journal of Vision* **16**, 1 (2016).
- [53] Ma, W. J., Beck, J. M., Latham, P. E. & Pouget, A. Bayesian inference with probabilistic population codes. *Nat. Neurosci.* **9**, 1432–1438 (2006).

- [54] Girshick, A. R., Landy, M. S. & Simoncelli, E. P. Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nat. Neurosci.* **14**, 926–932 (2011).
- [55] Acerbi, L., Wolpert, D. M. & Vijayakumar, S. Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. *PLoS Comput. Biol.* **8**, e1002771 (2012).
- [56] Neal, R. M. Slice sampling. *Annals of statistics* (2003).
- [57] Gelman, A., Hwang, J. & Vehtari, A. Understanding predictive information criteria for Bayesian models. *arXiv* (2013). [1307.5928v1](https://arxiv.org/abs/1307.5928v1).
- [58] Acerbi, L. Personal communication. (2017).
- [59] Rigoux, L., Stephan, K. E., Friston, K. J. & Daunizeau, J. Bayesian model selection for group studies - revisited. *NeuroImage* **84**, 971–985 (2014).
- [60] Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J. & Friston, K. J. Bayesian model selection for group studies. *NeuroImage* **46**, 1004–1017 (2009).
- [61] van den Berg, R., Awh, E. & Ma, W. J. Factorial comparison of working memory models. *Psychol. Rev.* (2014).
- [62] Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
- [63] Beck, J. M. *et al.* Marginalization in neural circuits with divisive normalization. *J. Neurosci.* **31**, 15310–15319 (2011).

Supplementary Information

1 Experiment 1

1.1 Apparatus and stimuli

Apparatus. Subjects were seated in a dark room, at a viewing distance of 32 cm from the screen, with their chin in a chinrest. Stimuli were presented on a gamma-corrected 60 Hz 9.7-inch 2048-by-1536 display. The display (LG LP097QX1-SPA2) was the same as that used in the 2013 iPad Air (Apple); we chose it for its high pixel density (264 pixels/inch). The display was connected to a Windows desktop PC using the Psychophysics Toolbox extensions^{49,50} for MATLAB (Mathworks).

Stimuli. The background was mid-level gray (199 cd/m²). The stimulus was either a drifting Gabor (Subjects 3, 6, 8, 9, 10, and 11) or an ellipse (Subjects 1, 2, 4, 5, and 7). The Gabor had a peak luminance of 398 cd/m² at 100% contrast, a spatial frequency of 0.5 cycles per degrees of visual angle (dva), a speed of 6 cycles per second, a Gaussian envelope with a standard deviation of 1.2 dva, and a randomized starting phase. Each ellipse had a total area of 2.4 dva², and was black (0.01 cd/m²). We varied the contrast of the Gabor and the elongation (eccentricity) of the ellipse (**Section 1.2**).

Categories. In Task A, stimulus orientations were drawn from Gaussian distributions with means $\mu_1 = -4^\circ$ (category 1) and $\mu_2 = 4^\circ$ (category 2) and standard deviations $\sigma_1 = \sigma_2 = 5^\circ$. In Task B, stimulus orientations were drawn from Gaussian distributions with means $\mu_1 = \mu_2 = 0^\circ$, and standard deviations $\sigma_1 = 3^\circ$ (category 1) and $\sigma_2 = 12^\circ$ (category 2) (**Fig. 1b**). We chose these category means and standard deviations such that the accuracy of an optimal observer would be around 80%.

1.2 Procedure

Each subject completed 5 sessions. Each session consisted of two parts; the subject did Task A in the first part, followed by Task B in the second part, or vice versa (chosen randomly each session). Each part started with instruction and was followed by alternating blocks of 96 category training trials and 144 testing trials, for a total of three blocks of each type, with a block of 24 confidence training trials immediately after the first category training block. Combining all sessions and both tasks, each subject completed 2880 category training trials, 240 confidence training trials, and 4320 testing trials; we did not analyze category training or confidence training trials.

Instruction. At the start of each part of a session, subjects were shown 30 (72 in the first session) exemplar stimuli from each category. Additionally, we provided them with a printed graphic similar to **Figure 1b**, and explained how the stimuli were generated from distributions. We answered any questions.

Category training. To ensure that subjects knew the stimulus distributions well, we gave them extensive category training. Each trial proceeded as follows (**Fig. 1a**): Subjects fixated on a central cross for 1 s. category 1 or category 2 was selected with equal probability. The stimulus orientation was drawn from the corresponding stimulus distribution (**Fig. 1b**). Gabors had 100% contrast, and ellipses had 0.95 eccentricity (elongation). The stimulus appeared at fixation for 300 ms, replacing the fixation cross. Subjects were asked to report category 1 or category 2 by pressing a button with their left or right index finger, respectively. Subjects were able to respond immediately after the offset of the stimulus, at which point verbal correctness feedback was displayed for 1.1 s. The fixation cross then reappeared.

Confidence training. To familiarize subjects with the button mappings, they completed a short confidence training block at the start of every task. We told subjects that in this block, it would be harder to tell

what the stimulus orientation was, there would be no correctness feedback, and they would be reporting their confidence on each trial in addition to their category choice. We provided them with a printed graphic similar to the buttons pictured in **Figure 1a**, indicating that they had to press one of eight buttons to indicate both category choice and confidence level, the latter on a 4-point scale. The confidence levels were labeled as “very high,” “somewhat high,” “somewhat low,” and “very low.” Gabors had 0.4%, 0.8%, 1.7%, 3.3%, 6.7%, or 13.5% contrast, and ellipses had 0.15, 0.28, 0.41, 0.54, 0.67, or 0.8 eccentricity, chosen randomly with equal probability on each trial (**Fig. 1c**). Stimuli were only displayed for 50 ms. Trial-to-trial feedback consisted only of a message telling them which category and confidence level they had reported. Other than these changes, the trial procedure was the same as in category training.

Subjects were not instructed to use the full range of confidence reports, as that might have biased them away from reporting what felt most natural. Instead, they were simply asked to be “as accurate as possible in reporting their confidence” on each trial.

Testing. The trial procedure in testing blocks was the same as in confidence training blocks, except that trial-to-trial feedback was completely withheld. At the end of each block, subjects were required to take at least a 30 sec break. During the break, they were shown the percentage of trials that they had correctly categorized. Subjects were also shown a list of the top 10 block scores (across all subjects, indicated by initials) for the task they had just done. This was intended to motivate subjects to score highly, and to reassure them that their scores were normal, since it is rare to score above 80% on a block.

1.3 Subjects

11 subjects (2 male), aged 20-42, participated in the experiment. Subjects received \$10 per 40-60 minute session, plus a completion bonus of \$15. The experiments were approved by the University Committee on Activities Involving Human Subjects of New York University. Informed consent was given by each subject before the experiment. All subjects were naïve to the purpose of the experiment. No subjects were fellow scientists.

1.4 Descriptive statistics

The following statistical differences were assessed using repeated-measures ANOVA.

In Task A, there was a significant effect of true category on category choice ($F_{1,10} = 285, P < 10^{-7}$). There was no main effect of reliability, which took 6 levels of contrast or ellipse elongation, on category choice ($F_{5,50} = 0.27, P = 0.88$). In other words, subjects were not significantly biased to respond with a particular category at low reliabilities. There was a significant interaction between reliability and true category, which is to be expected ($F_{5,50} = 59.6, P < 10^{-15}$) (**Fig. 3a**).

In Task B, there was again a significant effect of true category on category choice ($F_{1,10} = 78.3, P < 10^{-5}$). There was no main effect of reliability ($F_{5,50} = 2.93, P = 0.051$). There was again a significant interaction between reliability and true category ($F_{5,50} = 28, P < 10^{-12}$) (**Fig. 3b**).

In Task A, there was a significant effect of true category on response ($F_{1,10} = 136, P < 10^{-6}$). There was no main effect of reliability ($F_{5,50} = 0.61, P = 0.642$). There was a significant interaction between reliability and true category ($F_{5,50} = 58.7, P < 10^{-13}$) (**Fig. 3c**).

In Task B, there was a significant effect of true category on response ($F_{1,10} = 54.2, P < 10^{-6}$). There was a significant effect of reliability ($F_{5,50} = 4.84, P = 0.0128$). There was a significant interaction between reliability and true category ($F_{5,50} = 29.2, P < 10^{-8}$) (**Fig. 3d**).

In Task A, there was a main effect of confidence on the proportion of reports ($F_{3,30} = 7.75, P < 10^{-3}$); low-

confidence reports were more frequent than high-confidence reports. There was no significant effect of true category ($F_{1,10} = 0.784, P = 0.397$) and no interaction between confidence and category on proportion of responses ($F_{3,30} = 1.45, P = 0.25$) (**Fig. 3e**).

In Task B, there was a main effect of confidence on the proportion of reports ($F_{3,30} = 4.36, P = 0.012$). There was no significant effect of category ($F_{1,10} = 0.22, P = 0.64$), although there was an interaction between confidence and category ($F_{3,30} = 8.37, P = 0.003$). This is likely because for task B, category 2 has a higher proportion of “easy” stimuli (**Fig. 3f**).

In both tasks, reported confidence had a significant effect on performance ($F_{3,30} = 36.9, P < 10^{-3}$). Task also had a significant effect on performance ($F_{1,10} = 20.1, P = 0.001$); although we chose the category parameters such that the performance of the optimal observer is matched, subjects were significantly better at Task A. There was no interaction between task and confidence ($F_{3,30} = 0.878, P = 0.436$) (**Fig. 3g**).

Fig. 3l-m shows psychometric choice curves for both tasks, at all 6 levels of reliability. Each point represents roughly the same number of trials.

Fig. 3n-o shows a similar set of psychometric curves. These curves differ from **Fig. 3l-m** in that they represent the mean button press rather than mean category choice.

In Task A (**Fig. 3l,n**), mean category choice and mean button press depend monotonically on orientation, with a slope that increases with reliability. In Task B (**Fig. 3m,o**), the mean category choice and mean button press tends towards category 1 when stimulus orientation is near horizontal, and tends towards category 2 when orientation is strongly tilted; this reflects the stimulus distributions.

Since our models do not include any learning effects, we wanted to ensure that task performance was stable. For all tasks and experiments, we found no evidence that performance changed significantly as a function of the number of trials. For each experiment and task (the 5 lines in **Fig. S1**), we fit a logistic regression to the binary correctness data for each subject, obtaining a set of slope coefficients. We then used a t-test to determine whether these sets of coefficients differed significantly from zero. In no group did the slopes differ significantly from zero; across all 5 groupings the minimum p -value was 0.077 (Task A, experiment 2), which would not be significant even before correcting for multiple comparisons.




Figure S1: Performance as a function of number of trials, for both tasks and for all experiments. Performance was computed as a moving average over test trials (200 trials wide). Shaded regions represent ± 1 s.e.m. over subjects. Performance did not change significantly over the course of each experiment.

1.5 Effect of stimulus type on results: Gabor vs. ellipse

Since some subjects only saw Gabors and some only saw ellipses, we used Spearman’s rank correlation coefficient to measure the similarity of the two groups’ model rankings. Spearman’s rank correlation coefficient between Gabor and ellipse subjects for the summed LOO scores of the model groupings in **Figure 6** and **Figure S12** was 0.952 and 0.944, respectively (a value of 1 would indicate identical rankings). In both model groupings, the identities of the lowest- and highest-ranked models was the same for both Gabor and

ellipse subjects. This indicates that the choice of stimulus type did not have a systematic effect on model rankings.

2 Experiment 2: Separate category and confidence responses and testing feedback

This control experiment was identical to experiment 1 except for the following modifications:

- Subjects first reported choice by pressing one of two buttons with their left hand, and then reported confidence by pressing one of four buttons with their right hand.
- Subjects reported confidence in category training blocks, and received correctness feedback after reporting confidence.
- There were no confidence training blocks.
- In testing blocks, subjects received correctness feedback after each trial.
- Subjects completed a total of 3240 testing trials.
- 8 subjects (0 male), aged 19–23, participated. None were participants in experiment 1, and again, none were fellow scientists.
- Drifting Gabors were used; no subjects saw ellipses.

3 Experiment 3: Task B only

This experiment was identical to experiment 1 except for the following modifications:

- Subjects completed blocks of Task B only.
- Subjects completed a total of 3240 testing trials.
- 15 subjects (7 female), aged 19–30, participated. None were participants in experiments 1 or 2.
- Drifting Gabors were used; no subjects saw ellipses.

4 Modeling

4.1 Measurement noise

For models (such as our core models) where the relationship between reliability (i.e., contrast or ellipse eccentricity) and noise was parametric, we assumed a power law relationship between reliability c and measurement noise variance σ^2 : $\sigma^2(c) = \gamma + \alpha c^{-\beta}$. We have previously²⁶ used this power law relationship because it encompasses a large family of monotonically decreasing relationships using only three parameters. The relationship is also consistent with a form of the Naka-Rushton function^{51,52} commonly used to describe the mapping from reliability to neural gain g : $g = \frac{\gamma c^\beta}{c^\beta + \alpha}$. The power law relationship then holds under the assumption that measurement noise variance is inversely proportional to gain⁵³.

For all models except the Bayesian model with additive precision, we assumed additive orientation-dependent noise in the form of a rectified 2-cycle sinusoid, accounting for the finding that measurement noise is higher at non-cardinal orientations⁵⁴. The measurement noise s.d. comes out to

$$\sigma(c) = \sqrt{\gamma + \alpha c^{-\beta}} + \psi \left| \sin \frac{\pi s}{90} \right|. \quad (1)$$

4.2 Response probability

We coded all responses as $r \in \{1, 2, \dots, 8\}$, with each value indicating category and confidence. For all models except the Linear Neural model, the probability of a single trial i is equal to the probability mass of the measurement distribution $p(x | s_i) = \mathcal{N}(x; s_i, \sigma_i^2)$ in a range corresponding to the subject's response r_i . Because we only use a small range of orientations, we can safely approximate measurement noise as a normal distribution, rather than a Von Mises. We find the boundaries $(b_{r_i-1}(\sigma_i), b_{r_i}(\sigma_i))$ in measurement space, as defined by the fitting model and parameters θ , and then compute the probability mass of the measurement distribution between the boundaries:

$$p_{m,\theta}(r_i | s_i, \sigma_i) = \int_{b_{r_i-1}}^{b_{r_i}} \mathcal{N}(x; s_i, \sigma_i^2) dx. \quad (2)$$

For Task A, $b_0 = -\infty^\circ$ and $b_8 = \infty^\circ$. For Task B, $b_0 = 0^\circ$ and $b_8 = \infty^\circ$; since the task is symmetric around 0° , we only use $|s|$ in our computation of the log likelihood.

To obtain the log likelihood of the dataset, given a model with parameters θ , we compute the sum of the log probability for every trial i , where t is the total number of trials:

$$\log p(\text{data} | \theta) = \sum_{i=1}^t \log p(r_i | \theta) = \sum_{i=1}^t \log p_\theta(r_i | s_i, \sigma_i). \quad (3)$$

4.3 Model specification

4.3.1 Bayesian

Derivation of d_A and d_B . The log posterior ratio d is equivalent to the log likelihood ratio plus the log prior ratio:

$$d = \log \frac{p(C = 1 | x)}{p(C = 2 | x)} = \log \frac{p(x | C = 1)}{p(x | C = 2)} + \log \frac{p(C = 1)}{p(C = 2)}. \quad (4)$$

To get d_A and d_B , we need to find the task-specific expressions for $p(x | C)$. The observer knows that the measurement x is caused by the stimulus s , but has no knowledge of s . Therefore, the optimal observer marginalizes over s :

$$p(x | C) = \int p(x | s)p(s | C) ds.$$

We substitute the expressions for the noise distribution and the stimulus distribution, and evaluate the integral:

$$p(x | C) = \int \mathcal{N}(s; x, \sigma^2) \mathcal{N}(s; \mu_C, \sigma_C^2) ds = \mathcal{N}(x; \mu_C, \sigma^2 + \sigma_C^2).$$

Plugging in the task- and category-specific μ_C and σ_C , and substituting these expressions back into equation (4), we get:

$$d_A = \frac{2x\mu_1}{\sigma^2 + \sigma_1^2} + \log \frac{p(C=1)}{p(C=2)} \quad (5)$$

$$d_B = \frac{1}{2} \log \frac{\sigma^2 + \sigma_2^2}{\sigma^2 + \sigma_1^2} - \frac{\sigma^2 - \sigma_1^2}{2(\sigma^2 + \sigma_1^2)(\sigma^2 + \sigma_2^2)} x^2 + \log \frac{p(C=1)}{p(C=2)}. \quad (6)$$

The 8 possible category and confidence responses are determined by comparing the log posterior ratio d to a set of decision boundaries $\mathbf{k} = (k_0, k_1, \dots, k_8)$. k_4 is equal to the log prior ratio $\log \frac{p(C=1)}{p(C=2)}$, which functions as the boundary on d between the 4 category 1 responses and the 4 category 2 responses; k_4 is the only boundary parameter in models of category choice (and not confidence). k_0 is fixed at $-\infty$ and k_8 is fixed at ∞ . In all models, the observer chooses category 1 when d is positive.

Because the decision boundaries are free parameters, our models effectively include a large family of possible cost functions. A different cost function would be equivalent to a rescaling of the confidence boundaries \mathbf{k} . To see this, it is probably easiest to consider category choice alone; there, asymmetric costs for getting either category wrong would translate into a different value of k_4 , the category decision boundary (i.e., the observer's prior over category). For us, this boundary (and all other boundaries) are free parameters.

The posterior probability of category 1 can be written as $p(C=1 | x) = \frac{1}{1 + \exp(-d)}$.

Levels of strength. The Bayesian model is unique in that it is possible to formulate a principled version with relatively few boundary parameters. In principle, it is possible that such a model could perform better than more flexible models, if those models are overfitting. We formulated several levels of strength of the BCH, with weaker versions having fewer assumptions and more sets of mappings between the posterior probability of being correct and the confidence report (**Fig. S2**). In the *ultrastrong BCH*, confidence is a function solely of the posterior probability of the chosen category. In the *strong BCH*, it is additionally a function of the current task. In the *weak BCH*, it is additionally a function of the identity of the chosen category.

Most studies cannot distinguish between the ultrastrong and strong BCH because they test subjects in only one task. Furthermore, the weak BCH is only justifiable in tasks where the categories have different distributions of the posterior probability of being correct; the subject may then rescale their mappings between the posterior and their confidence. Here, one can see that Task B has this feature by observing that, in the bottom row of **Figure S2**, the distributions of posterior probabilities are different for the two categories). Most experimental tasks are like Task A, where the distributions are identical. We compared Bayesian models (Bayes_{Ultrastrong}, Bayes_{Strong}) corresponding to each of these versions of the BCH.

In Bayes_{Ultrastrong}, \mathbf{k} is symmetric across k_4 : $k_{4+j} - k_4 = k_4 - k_{4-j}$ for $j \in \{1, 2, 3\}$. Furthermore, in Bayes_{Ultrastrong}, $\mathbf{k}_A = \mathbf{k}_B$. So Bayes_{Ultrastrong} has a total of 4 free boundary parameters: k_1, k_2, k_3, k_4 . Bayes_{Ultrastrong} consists of the observer determining the response by comparing d_A and d_B to a single symmetric set of boundaries (**Fig. S2**, left column).

Bayes_{Strong} is identical to Bayes_{Ultrastrong} except that \mathbf{k}_A is allowed to differ from \mathbf{k}_B . So Bayes_{Strong} has a total of 8 free boundary parameters: $k_{1A}, k_{2A}, k_{3A}, k_{4A}, k_{1B}, k_{2B}, k_{3B}, k_{4B}$. Bayes_{Strong} consists of the observer determining the response by comparing d_A to a symmetric set of boundaries, and d_B to a different symmetric set of boundaries (**Fig. S2**, middle column).

Bayes_{Weak} is identical to Bayes_{Strong} except that symmetry is not enforced for \mathbf{k}_B . So Bayes_{Weak} has a total of 11 free boundary parameters: $k_{1A}, k_{2A}, k_{3A}, k_{4A}, k_{1B}, k_{2B}, k_{3B}, k_{4B}, k_{5B}, k_{6B}, k_{7B}$. Bayes_{Weak} consists of the observer comparing d_A to a symmetric set of boundaries, and d_B to a different non-symmetric set of boundaries (**Fig. S2**, right column).

We did not include Bayes_{Strong} and Bayes_{Ultrastrong} in the core models reported in the main text, because Bayes_{Weak} provided a much better fit to the data. Because it was not necessary in the main text to distinguish

the three strengths of Bayesian models, we refer to Bayes_{Weak} there simply as Bayes. The model comparison results for all models are available in the supplemental tables. Furthermore, we included these models in our model recovery analysis (Section 4.9).




Figure S2: Distributions of posterior probabilities of being correct, with confidence criteria for Bayesian models with three different levels of strength. Solid lines represent the distributions of posterior probabilities for each category and task in the absence of measurement noise. Dashed lines represent confidence criteria, generated from the mean of subject 4's posterior distribution over parameters. Each model has a different number of sets of mappings between posterior probability and confidence report. In Bayes_{Ultrastrong}, there is one set of mappings. In Bayes_{Strong}, there is one set for Task A, and another for Task B. In Bayes_{Weak}, as in the non-Bayesian models, there is one set for Task A, and one set for each reported category in Task B. Plots were generated from the mean of subject 4's posterior distribution over parameters as in Figure 2.

Decision boundaries. In the Bayesian models without d noise, we translate boundary parameters \mathbf{k} from log posterior ratio boundaries \mathbf{b} to measurement boundaries corresponding to fitted noise levels σ . To do this, we use the parameters \mathbf{k} as the left-hand side of equations (5) and (6) and solve for x at the fitted levels of σ . These values were used as the measurement boundaries $\mathbf{b}(\sigma)$.

In the Bayesian models with d noise, we assume that, for each trial, there is an added Gaussian noise term on d , $\eta_d \sim p(\eta_d)$, where $p(\eta_d) = \mathcal{N}(0, \sigma_d^2)$, and σ_d is a free parameter. We pre-computed 101 evenly spaced draws of η_d and their corresponding probability densities $p(\eta_d)$. We used equations (5) and (6) to compute a lookup table containing the values of d as a function of x , σ , and η_d . We then used linear interpolation to find sets of measurement boundaries $\mathbf{b}(\sigma)$ corresponding to each draw of η_d ⁵⁵. We then computed 101 response probabilities for each trial (Section 4.2), one for each draw of η_d , and computed the weighted average according to $p(\eta_d)$.

4.3.2 Probability correct with additive precision

We tested a model where the decision variable was a weighted mixture of precision (equivalent in this case to the Fisher information of the measurement variable x) and the probability of being correct³⁹. In this model, the decision variable is $\frac{\omega}{\sigma^2} + \frac{1}{1+\exp(-|d|)}$, where ω is a free parameter. To find the measurement boundaries $\mathbf{b}(\sigma)$, we substituted equations (5) or (6) for d , and set the whole value equal to parameters \mathbf{k} , solving for x at the fitted levels of σ . This model can be considered a hybrid Bayesian-heuristic model. Like Bayes_{Ultrastrong}, it has 4 free boundary parameters. Although the model is a hybrid Bayesian-heuristic model, not a strictly Bayesian one, we refer to it as Bayes_{Ultrastrong} + precision in Figure S12 and Table S1.

4.3.3 Fixed

In Fixed, the observer compares the measurement to a set of boundaries that are not dependent on σ . We fit free parameters \mathbf{k} , and use measurement boundaries $b_r = k_r$.

4.3.4 Lin and Quad

In Lin and Quad, the observer compares the measurement to a set of boundaries that are linear or quadratic functions of σ . We fit free parameters \mathbf{k} and \mathbf{m} , and use measurement boundaries $b_r(\sigma) = k_r + m_r\sigma$ (Lin) or $b_r(\sigma) = k_r + m_r\sigma^2$ (Quad).

4.3.5 Orientation Estimation

In Orientation Estimation, the observer uses their knowledge of the stimulus distributions to compute a maximum a posteriori estimate of the stimulus:

$$\begin{aligned}\hat{s} &= \underset{s}{\operatorname{argmax}} p(s | x) \\ &= \underset{s}{\operatorname{argmax}} p(x | s)p(s) \\ &= \underset{s}{\operatorname{argmax}} [\mathcal{N}(s; x, \sigma^2)(p(s | C = 1) + p(s | C = 2))].\end{aligned}\quad (7)$$

The observer then compares \hat{s} to a set of boundaries \mathbf{k} to determine category and confidence response.

Decision boundaries. To find the decision boundaries in measurement space, we used *gmm1max_n2_fast* from Luigi Acerbi's [gmm1](#) Gaussian mixture model toolbox to solve equation (7), computing a lookup table containing the value of \hat{s} as a function of x and σ ³⁶. We then found, using linear interpolation, the values of x corresponding to σ and the free parameters \mathbf{k} . These values were used as the measurement boundaries $\mathbf{b}(\sigma)$.

4.3.6 Linear Neural

In this section, \mathbf{r} refers to neural activity, not button responses.

This model is different from all other models in that the generative model does not include measurement x . The model can be derived as follows.

All neurons have Gaussian tuning curves with variance σ_{TC}^2 and gain $g = \frac{1}{\sigma^2}$. Tuning curve means are contained in the vector of preferred stimuli $\tilde{\mathbf{s}}$. The number of spikes in the population is $\mathbf{r} \sim \text{Poisson}(g\mathcal{N}(s; \tilde{\mathbf{s}}, \sigma_{\text{TC}}^2))$. Neural weights are a linear function of the preferred stimuli: $\mathbf{w} = a\tilde{\mathbf{s}}$.

On each trial, we get some quantity that is a weighted sum of each neuron's activity, $z = \mathbf{w} \cdot \mathbf{r}$. $\mathbb{E}[z | s] = \mathbf{w} \cdot \mathbb{E}[\mathbf{r} | s] = ag \sum_j \tilde{s}_j \exp\left(-\frac{(s-\tilde{s}_j)^2}{2\sigma_{\text{TC}}^2}\right)$.

Rather than sum over all neurons, we assume an infinite number of neurons uniformly spanning all possible preferred stimuli \tilde{s} . This allows us to replace the sum with an integral. The expected value of z is $ag \int \tilde{s} \exp\left(-\frac{(s-\tilde{s})^2}{2\sigma_{\text{TC}}^2}\right) d\tilde{s} = ags\sqrt{2\pi\sigma_{\text{TC}}^2}$. The variance of z is

$$\sum_j w_j^2 f_j(s) = ag \int \tilde{s}^2 \exp\left(-\frac{(s-\tilde{s})^2}{2\sigma_{TC}^2}\right) d\tilde{s} = ag\sqrt{2\pi\sigma_{TC}^2}(\sigma_{TC}^2 + s^2).$$

Now that we have the mean and variance of z , we are going to assume that z is normally distributed. This is equivalent to assuming that there are a high number of spikes, because the Poisson distribution approximates the normal distribution as the rate parameter becomes high. To compute response probability, we fit neural activity boundaries \mathbf{k} , and replace equation (2) with

$$p_{\theta}(r_i | s_i, \sigma_i) = \int_{k_{r_i-1}}^{k_{r_i}} \mathcal{N}(z; ags_i\sqrt{2\pi\sigma_{TC}^2}, ag\sqrt{2\pi\sigma_{TC}^2}(\sigma_{TC}^2 + s_i^2)) dz.$$

4.4 Lapse rates

In confidence and category models, we fit three different types of lapse rate. On each trial, there is some fitted probability of:

- A “full lapse” in which the category report is random, and confidence report is chosen from a distribution over the four levels defined by λ_1 , the probability of a “very low confidence” response, and λ_4 , the probability of a “very high confidence” response, with linear interpolation for the two intermediate levels.
- A “confidence lapse” $\lambda_{\text{confidence}}$ in which the category report is chosen normally, but the confidence report is chosen from a uniform distribution over the four levels.
- A “repeat lapse” λ_{repeat} in which the category and confidence response is simply repeated from the previous trial.

In category choice models, we fit a standard category lapse rate λ , as well the above “repeat lapse” λ_{repeat} .

4.5 Parameterization

Because of tradeoffs when directly fitting parameters γ, α, β , we re-parameterized equation (1) as

$$\sigma(c) = \sqrt{\sigma_L^2 + \frac{(\sigma_L^2 - \sigma_H^2)(c^{-\beta} - c_L^{-\beta})}{c_L^{-\beta} - c_H^{-\beta}}} + \psi \left| \sin \frac{\pi s}{90} \right|,$$

where c_L and c_H were the values of the lowest and highest reliabilities used. This way, σ_L and σ_H were free parameters that determined the s.d. of the measurement distributions for the lowest and highest reliabilities, and β was a free parameter determining the curvature of the function between the two reliabilities. For models where the relationship between reliability and noise was non-parametric, the first term in equation (1) was replaced with free s.d. parameters ($\sigma_{\text{rel. } 1}, \dots, \sigma_{\text{rel. } 6}$) corresponding to each of the six reliability levels.

For models where subjects had incorrect knowledge about their measurement noise, we fitted two sets of uncertainty-related parameters. One set was for the generative noise (used in equation (2)), and the other set was for the subject’s believed noise (used in equations (5), (6), and (7)).

All parameters that defined the width of a distribution ($\sigma_L, \sigma_H, \sigma_d, \sigma_1, \dots$) were sampled in log-space and exponentiated during the computation of the log likelihood. See *model_parameters.xls* for a complete list of each model’s parameters.

4.6 Model fitting

Rather than find a maximum likelihood estimate of the parameters, we sampled from the posterior distribution over parameters, $p(\theta \mid \text{data})$; this has the advantage of maintaining a measure of uncertainty about the parameters, which can be used for both model comparison and for plotting model fits. To sample from the posterior, we use an expression for the unnormalized log posterior

$$\log p(\theta \mid \text{data}) = \log p(\text{data} \mid \theta) + \log p(\theta) + \text{constant}, \quad (8)$$

where $\log p(\text{data} \mid \theta)$ is given in equation (3). We assumed a factorized prior over each parameter j :

$$\log p(\theta) = \sum_{j=1}^n \log p(\theta_j),$$

where j is the parameter index and n is the number of parameters. We took uniform (or, for parameters that were standard deviations, log-uniform) priors over reasonable, sufficiently large ranges³⁶, which we chose before fitting any models.

We sampled from the probability distribution using a Markov Chain Monte Carlo (MCMC) method, slice sampling⁵⁶. For each model and dataset combination, we ran between 4 and 7 parallel chains with random starting points. For each chain, we took 40,000 to 600,000 total samples (depending on model computational time), discarded the first third of the samples, and kept 10,000 of the remaining samples, randomly selected. All samples with log posteriors less than 40 below the log posterior of the sample with the highest log posterior were discarded. Marginal probability distributions of the sample log likelihoods were visually checked for convergence across chains. In total we had 842 model and dataset combinations, with a median of 26,668 kept samples (IQR = 13,334).

4.7 Model comparison

4.7.1 Model groupings

We used 8 groupings of model-subject combinations where it made sense to consider the models as being on equal footing for the purpose of model comparison. The model-subject combinations were grouped by: experiment (which corresponded to subject population), data type (category response only vs. category and confidence response), task type (Task A, B, or both fit jointly). The 8 groupings correspond to **Figures S12–S19** and **Tables S1–S8**.

4.7.2 Metric choice

A more complex model is likely to fit a dataset better than a simpler model, even if only by chance. Since we are interested in our models' predictive accuracy for unobserved data, it is important to choose a metric for model comparison that takes the complexity of the model into account, avoiding the problem of overfitting. Roughly speaking, there are two ways to compare models: information criteria and cross-validation.

Most information criteria (such as AIC, BIC, and AICc) are based on a point estimate for θ , typically θ_{MLE} , the θ that maximizes the log likelihood of the dataset (equation (3)). For instance, AIC adds a correction for the number of parameters n to the log likelihood of the dataset: $\text{AIC} = -2 \sum_{i=1}^t \log p(r_i \mid \theta_{\text{MLE}}) + 2n$.

WAIC is a more Bayesian approach to information criteria that adds a correction for the effective number of parameters⁵⁷. Because WAIC is based on samples from the full posterior of θ (equation (8), typically sampled via MCMC), it takes into account the model’s uncertainty landscape.

Although information criteria are computationally convenient, they are based on asymptotic results and assumptions about the data that may not always hold⁵⁷. An alternative way to estimate predictive accuracy for unobserved data is to cross-validate, fitting the model to training data and evaluating the fit on held out data. Leave-one-out cross-validation is the most thorough way to cross-validate, but is very computationally intensive; it requires that you fit your model t times, where t is the number of trials. Here we use a method (PSIS-LOO, referred to here simply as LOO) proposed by Vehtari *et al.*³³ for approximating leave-one-out cross-validation that, like WAIC, uses samples from the full posterior of θ :

$$\text{LOO} = \sum_{i=1}^t \log \frac{\sum_u w_{i,u} p(r_i | \theta_u)}{\sum_u w_{i,u}}$$

where θ_u is the u -th sampled set of parameters, and $w_{i,u}$ is the importance weight of trial i for sample u . Pareto smoothed importance sampling provides an accurate and reliable estimate of the weights. LOO is currently the most accurate approximation of leave-one-out cross-validation⁵⁸. Conveniently, it has a natural diagnostic that allows the user to know when the metric may be inaccurate³³.

We tried to determine whether our results were dependent on our choice of metric. We computed AIC, BIC, AICc, WAIC, and LOO for all models in the 8 model groupings, multiplying the information criteria by $-\frac{1}{2}$ to match the scale of LOO. For AIC, BIC, and AICc, we used the parameter sample with the highest log likelihood as our estimate of θ_{MLE} . Then we computed Spearman’s rank correlation coefficient for every possible pairwise comparison of model comparison metrics for all model and dataset combinations, producing 80 total values (8 model groupings \times 10 possible pairwise comparisons of model comparison metrics). All values were greater than 0.998, indicating that, had we used an information criterion instead of LOO, we would not have changed our conclusions. Furthermore, there are no model groupings in which the identities of the lowest- and highest-ranked models are dependent on the choice of metric. The agreement of these metrics strengthens our confidence in our conclusions.

4.7.3 Metric aggregation

Summed LOO differences. In all figures where we present model comparison results (e.g., **Fig. 5**, right column), we aggregate LOO scores by the following procedure. Choose a reference model (usually the one with the lowest mean LOO score across subjects). Subtract all LOO scores from the corresponding subject’s score for the reference model; this converts all scores to a LOO “difference from best” score, with higher scores being worse. Repeat the following standard bootstrap procedure 10,000 times: Choose randomly, with replacement, a group of datasets equal to the total number of unique datasets, and take the sum over subjects of their “difference from best” scores for each model. Plots indicate the median and 95% CI of these bootstrapped summed “difference from best” scores. This approach implicitly assumes that all data was generated from the same model.

Group level Bayesian model selection. We also used LOO scores to compute two metrics that allow for model heterogeneity across the group. The first metric was “protected exceedance probability,” the posterior probability that one model occurs more frequently than any other model in the set⁵⁹, above and beyond chance (e.g., **Fig. S12b**). The second was the expected posterior probability that a model generated the data of a randomly chosen dataset⁶⁰ (e.g., **Fig. S4d**, **Fig. S12c**). The latter metric assumes a uniform prior over models, which is a function of the total number of datasets. We used the **SPM12** software package to compute these metrics.

In all but one of the 8 model groupings, all three methods of metric aggregation identify the same overall best model. For example, in **Fig. S12**, one model (Quad + non-param. σ) has the lowest summed LOO differences, the highest protected exceedance probability, and the highest expected posterior probability.

4.8 Visualization of model fits

Model fits were plotted by bootstrapping synthetic group datasets with the following procedure: For each task, model, and subject, we generated 20 synthetic datasets, each using a different set of parameters sampled, without replacement, from the posterior distribution of parameters. Each synthetic dataset was generated using the same stimuli as the ones presented to the real subject. We randomly selected a number of synthetic datasets equal to the number of subjects to create a synthetic group dataset. For each synthetic group dataset, we computed the mean output (e.g., button press, confidence, performance) per bin. We then repeated this 1,000 times and computed the mean and standard deviation of the mean output per bin across all 1,000 synthetic group datasets, which we then plotted as the shaded regions. Therefore, shaded regions represent the mean ± 1 s.e.m. of synthetic group datasets.

For plots with orientation on the horizontal axis (e.g., **Figure 3j–o**), stimulus orientation was binned according to quantiles of the task-dependent stimulus distributions so that each point consisted of roughly the same number of trials. For each task, we took the overall stimulus distribution $p(s) = \frac{1}{2}(p(s | C = 1) + p(s | C = 2))$ and found bin edges such that the probability mass of $p(s)$ was the same in each bin. We then plotted the binned data with linear spacing on the horizontal axis.

4.9 Model recovery

We performed a model recovery analysis⁶¹ to test our ability to distinguish our 6 core models, as well as the 2 stronger versions of the Bayesian model (**Section 4.3.1**). We generated synthetic datasets from each of the 8 models for both Tasks A and B, using the same sets of stimuli that were originally randomly generated for each of the 11 subjects. To ensure that the statistics of the generated responses were similar to those of the subjects, we generated responses to these stimuli from 4 of the randomly chosen parameter estimates obtained via MCMC sampling (as described in **Section 4.6**) for each subject and model. In total, we generated 352 datasets (8 generating models \times 11 subjects \times 4 datasets). We then fit all 8 models to every dataset, using maximum likelihood estimation (MLE) of parameters by an interior-point constrained optimization (MATLAB's *fmincon*), and computed AIC scores from the resulting fits.

We found that the true generating model was the best-fitting model, on average, in all cases (**Fig. S3**). Overall, AIC “selected” the correct model (i.e., AIC scores were lowest for the model that generated the data) for 86.6% of the datasets, indicating that our models are distinguishable.

Ideally, we would have evaluated our model recovery fits using LOO, as we evaluated the fits to human data. However, LOO can only be obtained when fitting with MCMC sampling, which takes orders of magnitudes longer than fitting with MLE. It would be impossible to fit all 352 synthetic datasets in a short amount of time using the same procedure and sampling quality standards described in **Section 4.6** (i.e., a large number of samples, across multiple converged chains). Furthermore, we do not believe that our results are dependent on how the models are fit and the fits are evaluated; we found that AIC and LOO scores gave us near-identical model rankings for data from real subjects (**Section 4.7.2**).

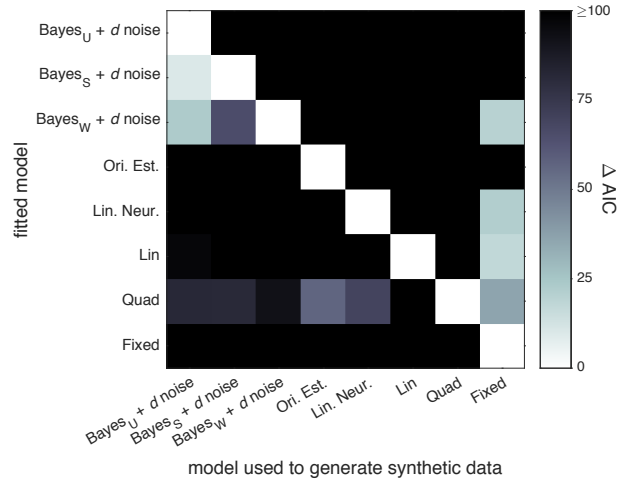


Figure S3: Model recovery analysis. Shade represents the difference between the mean AIC score (across datasets) for each fitted model and for the one with the lowest mean AIC score. White squares indicate the model that had the lowest mean AIC score when fitted to data generated from each model. The squares on the diagonal indicate that the true generating model was the best-fitting model, on average, in all cases.

5 Neural networks

5.1 Architecture

In this section, r and \mathbf{r} refer to neural activity, not button responses.

We trained 3-layer feedforward neural networks to perform Task B⁴⁴. The architecture, described below, is pictured in **Figure 7a**. The input units were 50 independent Poisson neurons. The mean number of spikes per trial was determined by Gaussian tuning curves with baselines, such that the neurons had spike count

$$\mathbf{r}_{\text{input}} \sim \text{Poisson}(g\mathcal{N}(s, \tilde{\mathbf{s}}, \sigma_{\text{TC}}^2) + \zeta).$$

$\tilde{\mathbf{s}}$ is the vector of preferred stimuli, which were linearly spaced from -40° to 40° . All neurons had tuning curve width $\sigma_{\text{TC}}^2 = 100$ and baseline $\zeta = 0.025$ (to limit the number of trials with zero spikes). Gain g varied from trial-to-trial, and was derived from fits to subject data. Input units were connected, all-to-all, to 200 hidden rectified linear units with responses

$$\mathbf{r}_{\text{hidden}} = \max(0, \mathbf{W}_{\text{input}}\mathbf{r}_{\text{input}}),$$

where $\mathbf{W}_{\text{input}}$ was the weight matrix applied to the input units. Both input and hidden layers included a bias unit with a constant response of 1, which, when multiplied by the fitted weights, effectively adds a fitted bias to the hidden units and output unit. Hidden units were connected to a sigmoidal output unit with response

$$r_{\text{output}} = \frac{1}{1 + \exp(-\mathbf{w}_{\text{hidden}} \cdot \mathbf{r}_{\text{hidden}})},$$

where $\mathbf{w}_{\text{hidden}}$ was the weight vector applied to the hidden units.

5.2 Training networks and generating datasets

Stimuli s were drawn from the same distributions used for the human experiments in Task B. Gains differed from trial-to-trial; to ensure that the reliability of the information available to the networks was similar to the reliability of the subjects' sensory information, the gains were derived from the fits to the Task A choice data in experiment 1⁵³. Indeed, the performance of the networks roughly matched the performance of the subjects (**Fig. S4c**). We used 15 different values for the number of training trials, ranging from 10 to 4.6×10^5 , logarithmically spaced (**Figure 7b–d** only depicts results from the most highly trained networks).

We used standard back-propagation⁶² to minimize cross-entropy between network output and category labels; as with the human subjects, the networks did not receive probabilistic feedback during training. Weights were initialized to small random values drawn from a zero-mean Gaussian distribution with s.d. 0.05. We used mini-batch gradient descent with a batch size of 10, over a single epoch. We used L2 regularization with regularization term $\alpha = 10^{-4}$.

We decoded the optimal posterior $p(C = 1 | \mathbf{r}_{\text{input}})$, which allowed us to compute fractional information loss. Fractional information loss was defined as the KL-divergence between $p(C = 1 | \mathbf{r}_{\text{input}})$ and r_{output} , normalized by the mutual information between the category labels and $\mathbf{r}_{\text{input}}$ ^{44,63}.

Learning rate η decreased as a function of the batch number j :

$$\eta_j = \frac{\eta_0}{1 + \tau j}.$$

We used a constrained pattern search optimization (MATLAB's *patternsearch*) to find, for the gains associated with each subject, the η_0 and τ that minimized fractional information loss on a validation set. We used *patternsearch* because, unlike *fmincon*, it is well-suited for optimizing stochastic objective functions.

From each trained network, we generated a test set consisting of 2160 trials, the same number of Task B trials completed by subjects in experiment 1. r_{output} was mapped onto the 8 category and confidence responses using quantiles. We produced datasets from 4 separately trained networks for gains associated with each subject, generating 660 datasets in total (15 numbers of training trials \times 11 subject-derived sets of gains \times 4 datasets).

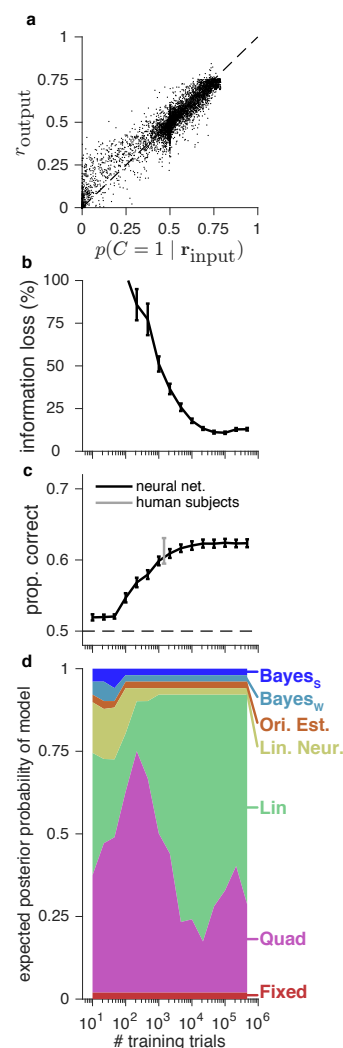


Figure S4: Neural network optimality, task performance, and model comparison. (a) Posterior probabilities decoded optimally from input unit activity, scattered against network output. Each point is a test trial from a neural network that was trained on the maximum number of training trials. For clarity, a randomly selected subset of test trials is plotted. (b) Fractional information loss as a function of the number of training trials. Error bars represent ± 1 s.e.m. across the means of datasets generated with the gains derived from each subject. (c) Black line indicates network test performance as a function of the number of training trials. Gray error bar indicates ± 1 s.e.m. for the Task B performance of subjects in experiment 1; all subjects completed 1440 training trials. (d) Expected posterior probability that a model generated a randomly chosen dataset⁶⁰, as a function of the number of training trials.

We found that r_{output} was a fairly good approximation of the optimal posterior $p(C = 1 | \mathbf{r}_{\text{input}})$, with some positive bias when the posterior was low (**Fig. S4a**). We also found that information loss and performance went down as the number of training trials increased (**Fig. S4b–c**). Both information loss and performance appear to reach asymptote around 2×10^4 trials; therefore, it is unlikely that our results are due to insufficient training.

5.3 Behavioral models

We fit the 660 network-generated datasets, obtaining AIC scores for each dataset and model. The fitted models were identical to the Task B models described in the text (and shown in **Figure S14** and **Table S3**), except that we removed

the following mechanisms that we knew not to be present in the neural network generative process:

- All lapse rates except for a uniform lapse rate over all 8 responses
- Orientation-dependent noise
- d noise (applicable to Bayesian models only)

As with model recovery (**Section 4.9**), we used MLE and AIC (rather than MCMC and LOO) for computational efficiency, due to the large number of datasets being fitted.

After fitting, we computed the expected posterior probability distribution over models at each number of training trials (**Section 4.7.3**). We found that Lin and Quad fit the data best, with Quad fitting best for data generated from less well-trained networks, and Lin fitting best for data generated from highly trained networks (**Fig. S4d**). This transition is consistent with previous results⁴⁴. We plotted the summed AIC differences for data generated from the most highly trained networks in **Figure S5**, blue bars.

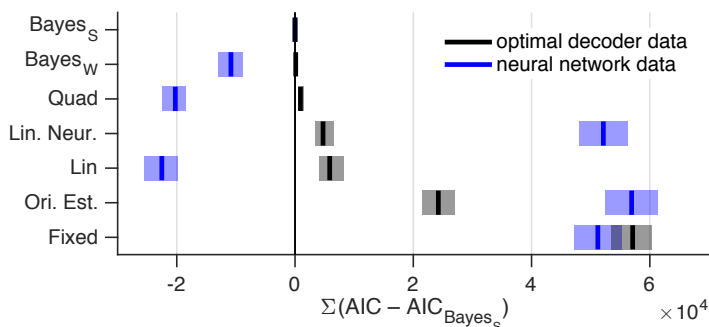


Figure S5: Model comparison for data generated from an optimal decoder of spikes, and from neural networks trained with 4.6×10^5 training trials. Models are ordered by quality of fit to optimal decoder data. As in **Figure 7d**, but for all models.

5.4 Control

To ensure that our results were not due to overfitting, we performed a model recovery analysis. This was similar to the previously described model recovery, except that we used input layer activity $\mathbf{r}_{\text{input}}$ rather than x and σ . For the $\mathbf{r}_{\text{input}}$ used in the test set of the 44 most highly trained networks (11 subject-derived sets of gains \times 4 datasets), we decoded the per-trial optimal posteriors $p(C = 1 | \mathbf{r}_{\text{input}})$. We then mapped the posterior onto the 8 category and confidence responses using quantiles (this is also how we mapped the trained networks' r_{output} onto responses, see above). We then fit these datasets with the same models used to fit the datasets produced by the trained networks. We found that Bayes_{Strong} was the best-fitting

model (**Fig. S5**, black). This suggests that the architecture or the training procedure of the neural networks constrains the type of behavior that can be produced.

6 Data and code availability

All data and code used for running experiments, model fitting, and plotting is available on [GitHub](#).

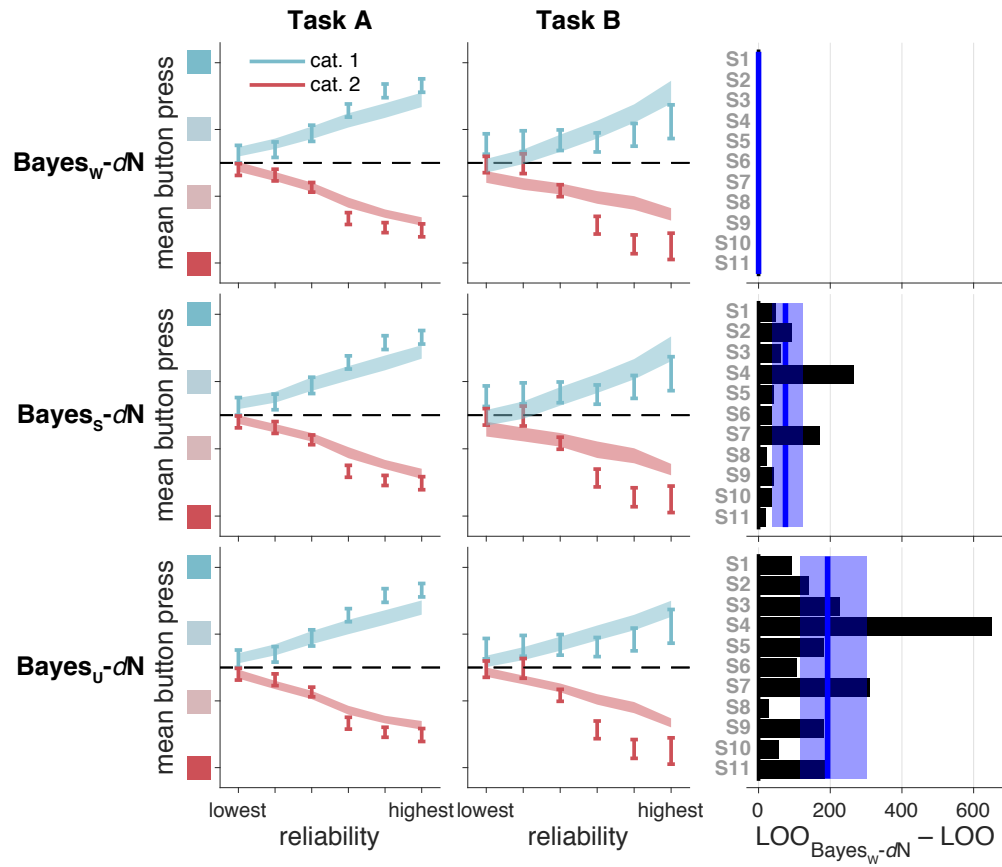


Figure S6: Model fits and model comparison for the three strengths of the Bayesian model, as in Figure 5. In the main text, $\text{Bayes}_{\text{Weak}}-dN$ is referred to simply as Bayes.

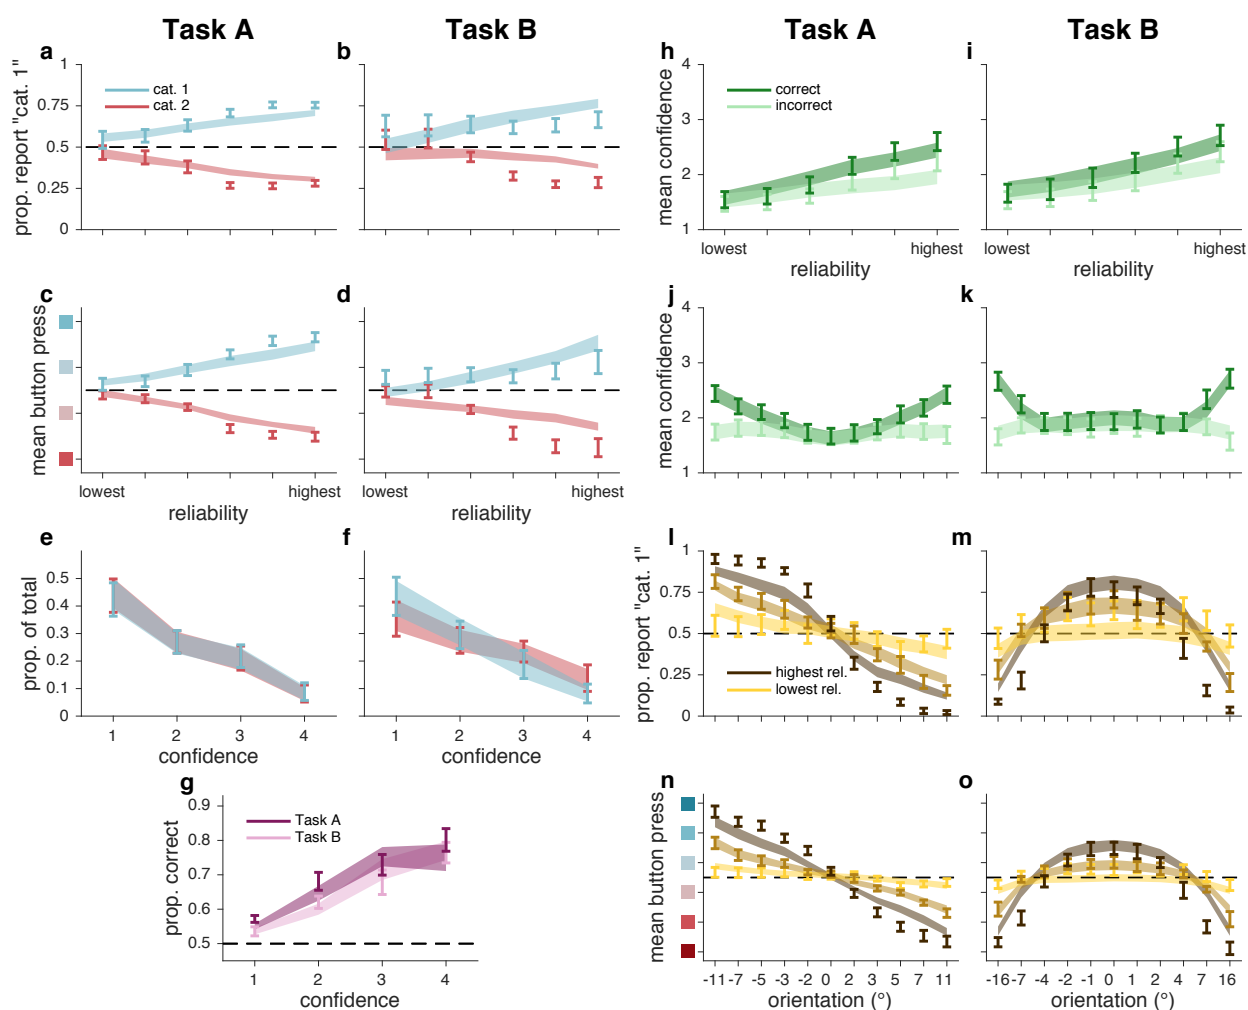


Figure S7: Bayes_{Weak}-dN fits, as in Figure 3. In the main text, Bayes_{Weak}-dN is referred to simply as Bayes.

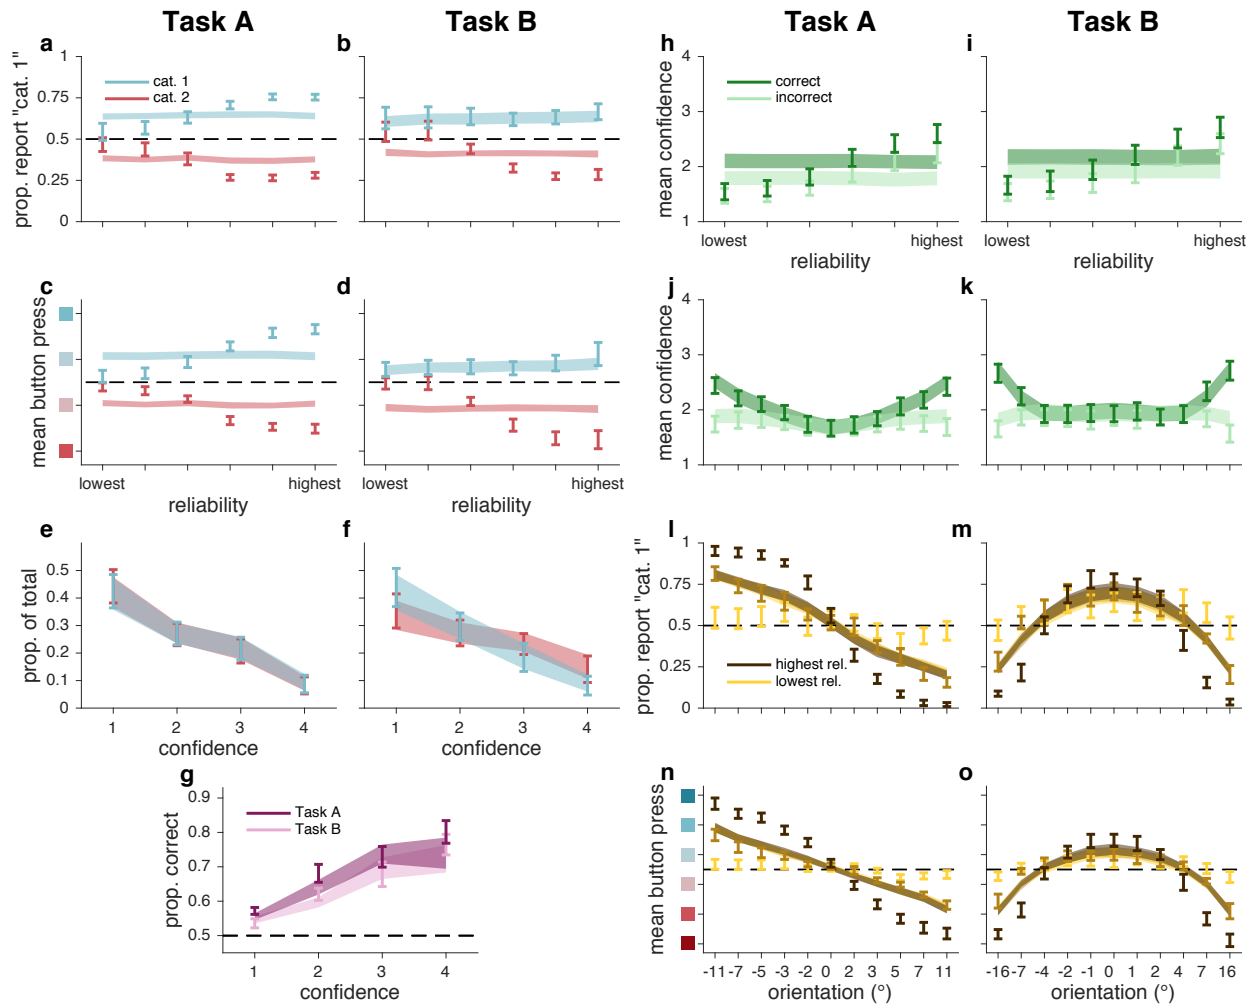


Figure S8: Fixed fits, as in **Figure 3**.

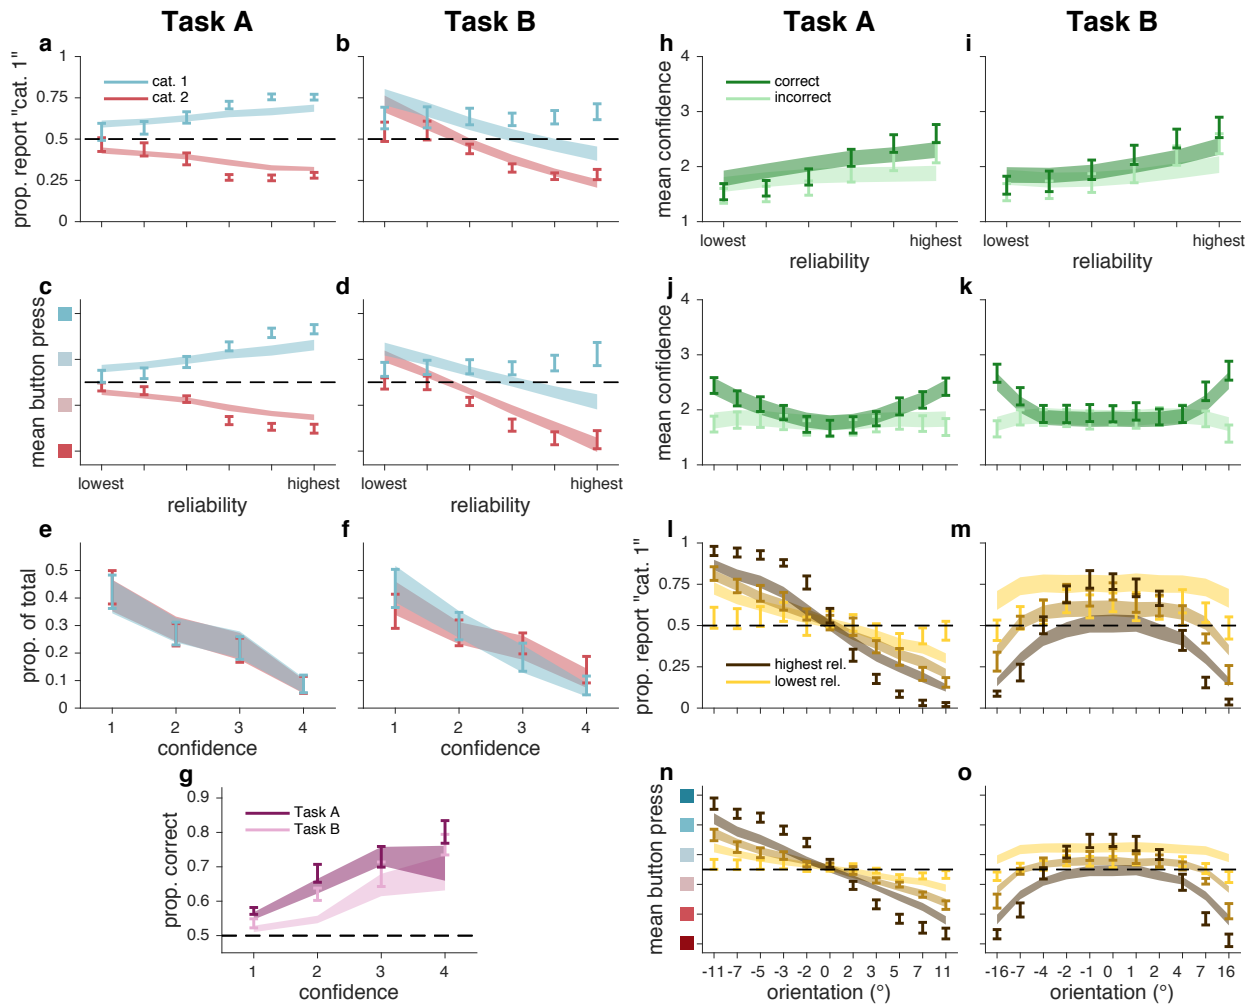


Figure S9: Orientation Estimation fits, as in **Figure 3**.

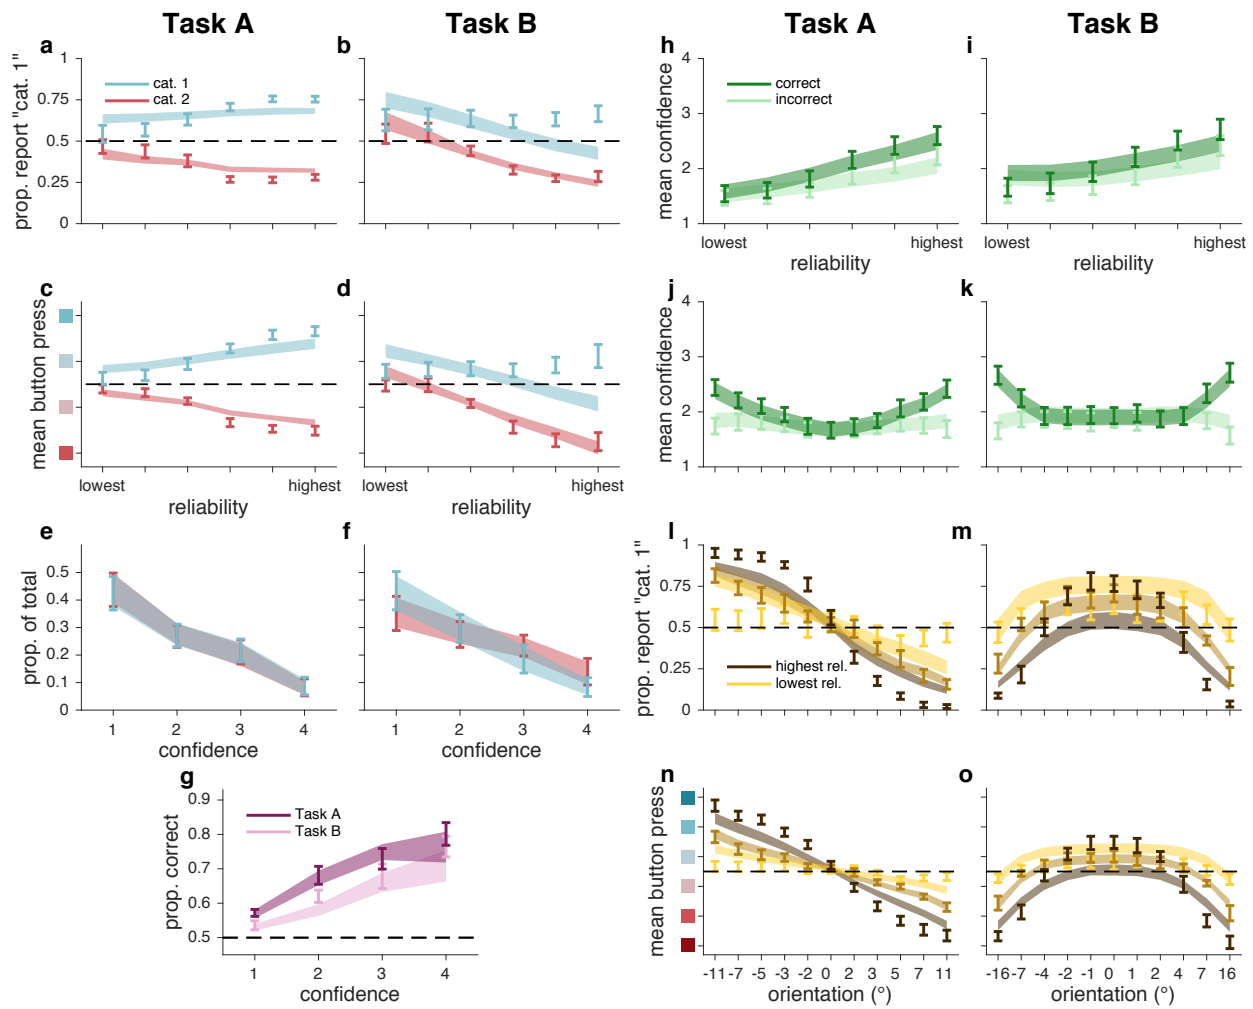


Figure S10: Linear Neural fits, as in Figure 3.

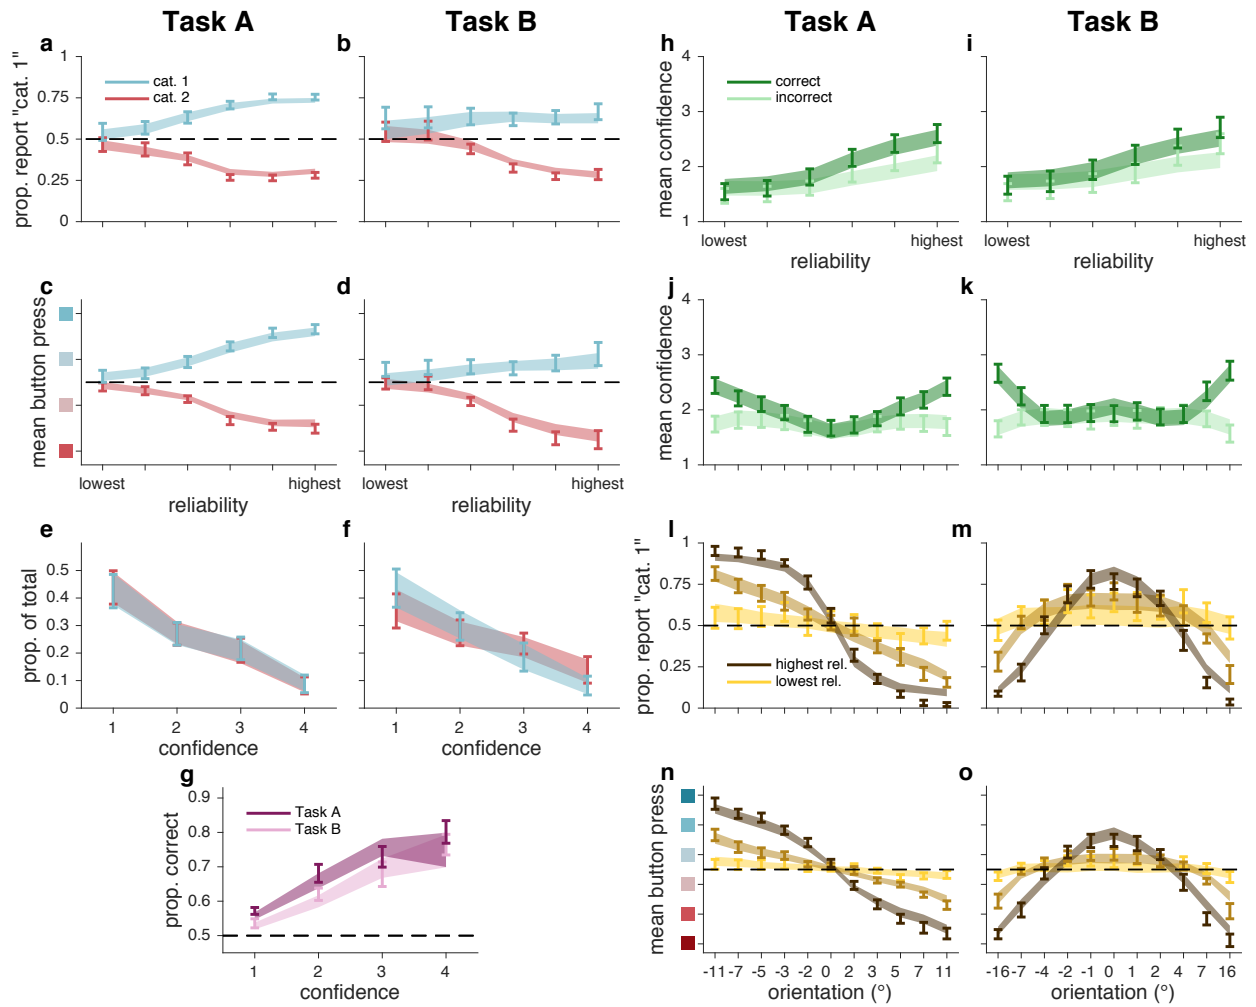


Figure S11: Lin fits, as in **Figure 3**.

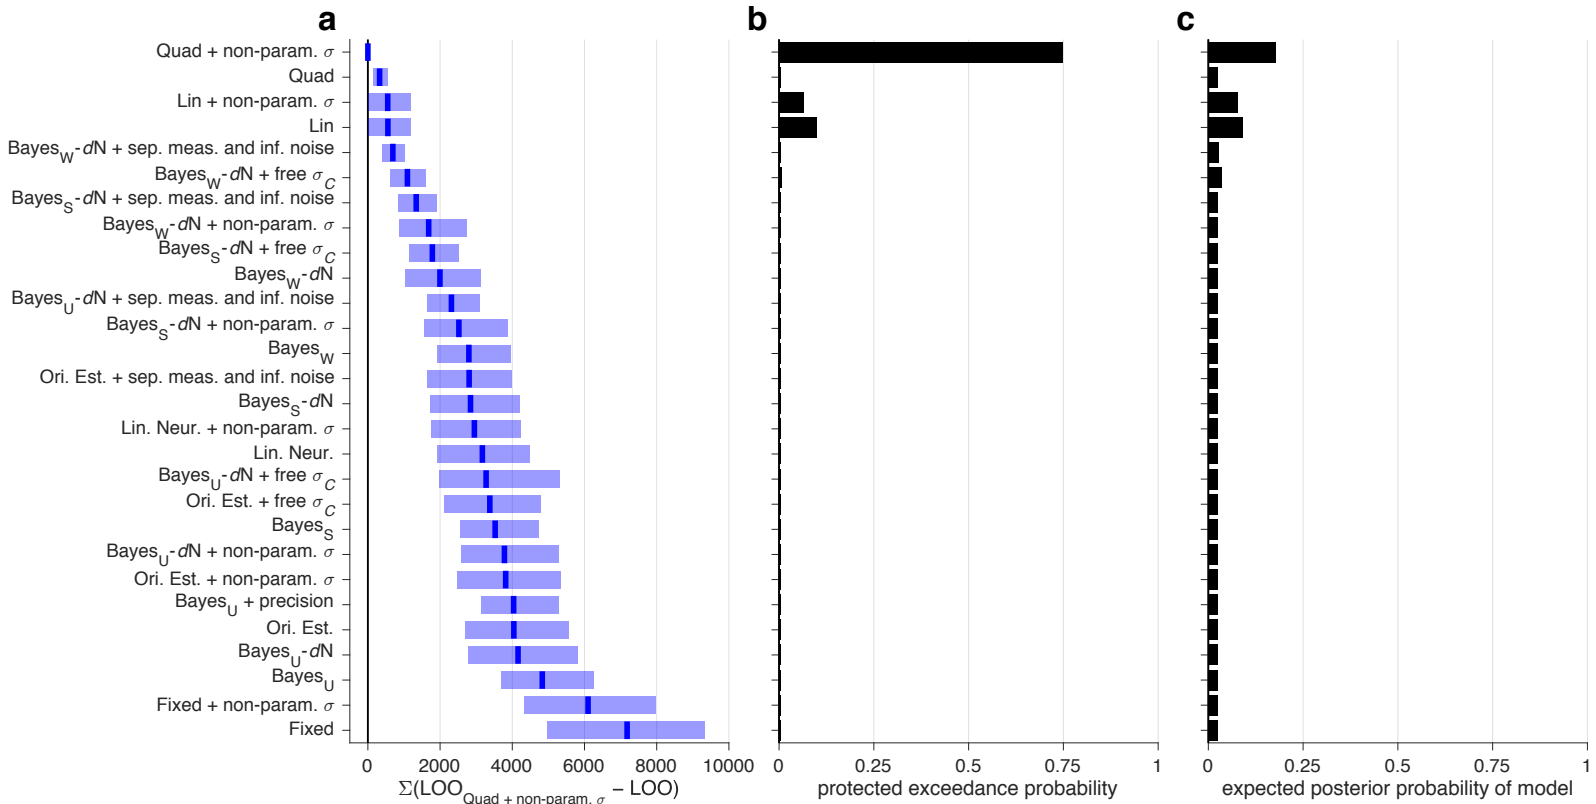


Figure S12: Model comparison, experiment 1. Models were fit jointly to Task A and B category and confidence responses. **(a)** Medians and 95% CI of bootstrapped sums of LOO differences, relative to the best model. Higher values indicate worse fits. **(b)** The protected exceedance probability, i.e., the posterior probability that a model occurs more frequently than the others⁵⁹. **(c)** The expected posterior probability that a model generated the data of a randomly chosen subject⁶⁰.

	10 pms	12 pms	16 pms	20 pms	24 pms	27 pms	30 pms	34 pms	36 pms	39 pms	42 pms	45 pms	48 pms	50 pms	53 pms	56 pms	59 pms	62 pms	65 pms	68 pms	71 pms	74 pms	77 pms	80 pms	83 pms	86 pms	89 pms	92 pms	95 pms	98 pms		
Quad + non-param. σ	0.000 (0.000, 0.000)																															
Quad	1000 (1000, 1000)	0.000 (0.000, 0.000)																														
Lin + non-param. σ	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																													
Lin	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																												
Bayes _W -dN + sep. meas. and inf. noise	4000 (4000, 4000)	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																											
Bayes _W -dN + free σ_C	5000 (5000, 5000)	4000 (4000, 4000)	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																										
Bayes _S -dN + sep. meas. and inf. noise	6000 (6000, 6000)	5000 (5000, 5000)	4000 (4000, 4000)	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																									
Bayes _W -dN + non-param. σ	7000 (7000, 7000)	6000 (6000, 6000)	5000 (5000, 5000)	4000 (4000, 4000)	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																								
Bayes _S -dN + free σ_C	8000 (8000, 8000)	7000 (7000, 7000)	6000 (6000, 6000)	5000 (5000, 5000)	4000 (4000, 4000)	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																							
Bayes _W -dN	9000 (9000, 9000)	8000 (8000, 8000)	7000 (7000, 7000)	6000 (6000, 6000)	5000 (5000, 5000)	4000 (4000, 4000)	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																						
Bayes _U -dN + sep. meas. and inf. noise	10000 (10000, 10000)	9000 (9000, 9000)	8000 (8000, 8000)	7000 (7000, 7000)	6000 (6000, 6000)	5000 (5000, 5000)	4000 (4000, 4000)	3000 (3000, 3000)	2000 (2000, 2000)	1000 (1000, 1000)	0.000 (0.000, 0.000)																					
Bayes _S -dN + non-param. σ																																
Bayes _W																																
Ori. Est. + sep. meas. and inf. noise																																
Bayes _S -dN																																
Lin. Neur. + non-param. σ																																
Lin. Neur.																																
Bayes _U -dN + free σ_C																																
Ori. Est. + free σ_C																																
Bayes _S																																
Bayes _U -dN + non-param. σ																																
Ori. Est. + non-param. σ																																
Bayes _U + precision																																
Ori. Est.																																
Bayes _U -dN																																
Bayes _U																																
Fixed + non-param. σ																																
Fixed																																

Table S1: Cross comparison of all models in **Figure S12**. Cells indicate medians and 95% CI of bootstrapped summed LOO score differences. A negative median indicates that the model in the corresponding row had a higher score (better fit) than the model in the corresponding column. For readability, see instead *table_S1.xls*. For the parameters of each model, see *model_parameters.xls*.

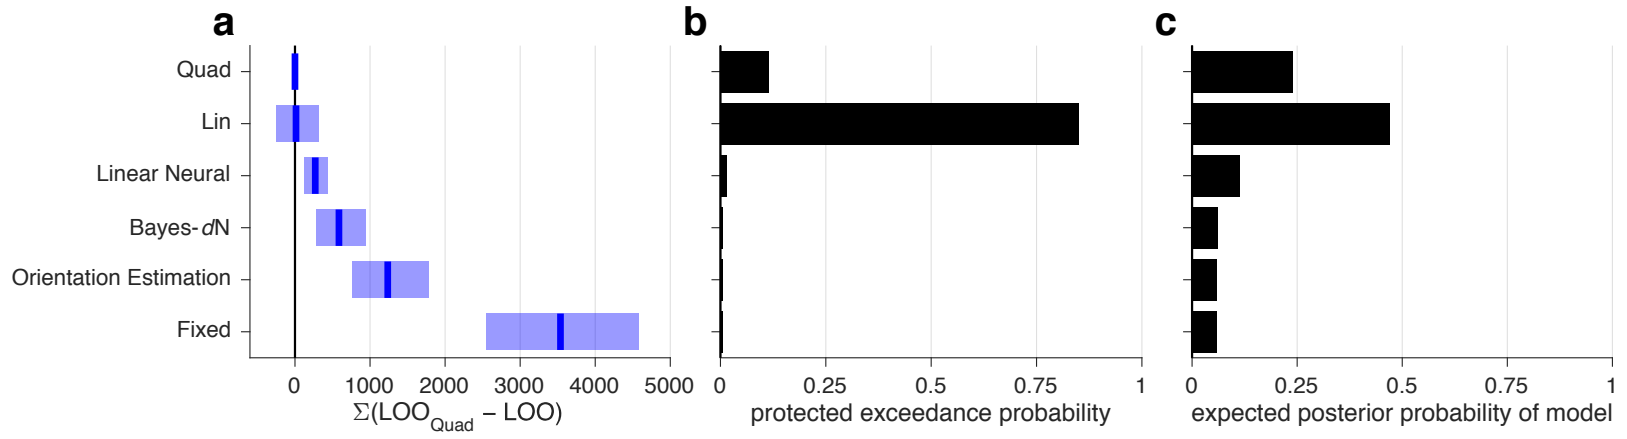


Figure S13: Model comparison, experiment 1. Models were fit to Task A category and confidence responses. See **Figure S12** caption.

		12 pars. Fixed	13 pars. Bayes- <i>dN</i>	12 pars. Ori. Est.	13 pars. Lin. Neur.	16 pars. Lin
16 pars.	Quad	-3534 [-4552, -2529]	-581 [-938, -278]	-1241 [-1798, -767]	-270 [-436, -117]	-14 [-325, 246]
16 pars.	Lin	-3532 [-4353, -2651]	-572 [-799, -339]	-1232 [-1566, -863]	-259 [-609, 124]	
13 pars.	Lin. Neur.	-3255 [-4343, -2231]	-313 [-724, 75]	-972 [-1599, -412]		
12 pars.	Ori. Est.	-2302 [-2881, -1705]	651 [425, 885]			
13 pars.	Bayes- <i>dN</i>	-2956 [-3723, -2163]				

Table S2: Cross comparison of all models in **Figure S13**. See **Table S12** caption.

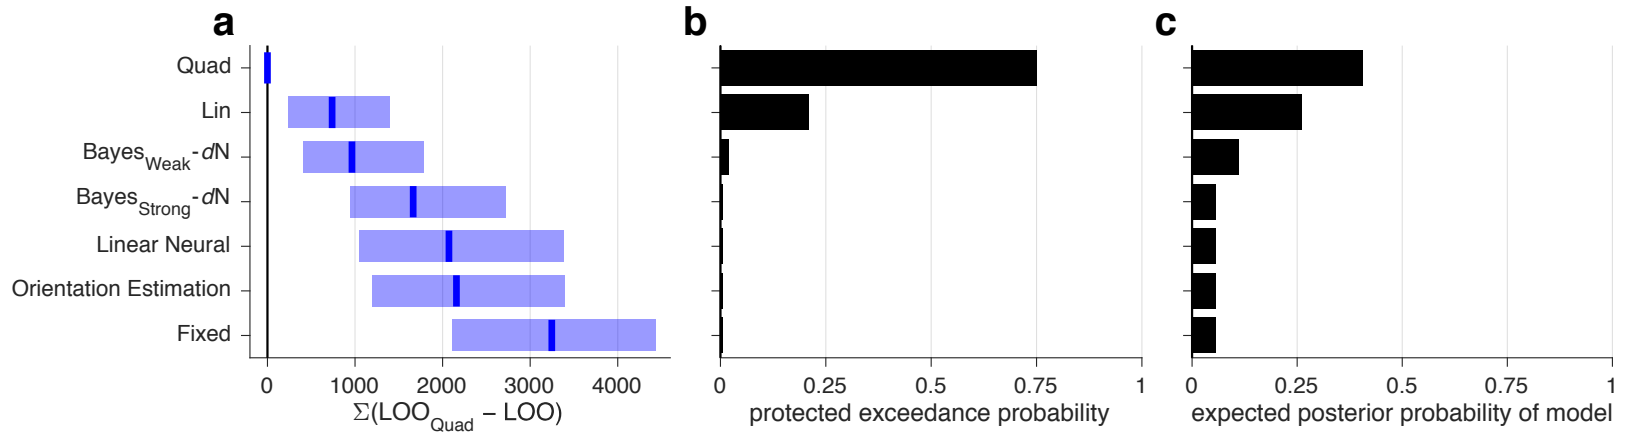


Figure S14: Model comparison, experiment 1. Models were fit to Task B category and confidence responses. See **Figure S12** caption.

		15 pars. Fixed	13 pars. Bayes _S -dN	16 pars. Bayes _W -dN	15 pars. Ori. Est.	16 pars. Lin. Neur.	22 pars. Lin
22 pars.	Quad	-3234 [-4390, -2099]	-1664 [-2698, -958]	-978 [-1756, -406]	-2156 [-3352, -1192]	-2060 [-3368, -1037]	-744 [-1387, -224]
22 pars.	Lin	-2480 [-3323, -1645]	-919 [-1788, -279]	-232 [-900, 346]	-1415 [-2439, -439]	-1326 [-2442, -337]	
16 pars.	Lin. Neur.	-1117 [-2093, -349]	421 [-1095, 1689]	1106 [-374, 2583]	-80 [-222, 62]		
15 pars.	Ori. Est.	-1043 [-1962, -273]	502 [-934, 1693]	1184 [-202, 2588]			
16 pars.	Bayes _W -dN	-2230 [-3239, -1307]	-691 [-1082, -390]				
13 pars.	Bayes _S -dN	-1534 [-2425, -634]					

Table S3: Cross comparison of all models in **Figure S14**. See **Table S12** caption.




Figure S15: Model comparison, experiment 1. Models were fit jointly to Task A and B category choices. See **Figure S12** caption.

		8 pars. Fixed	9 pars. Bayes- <i>dN</i>	8 pars. Ori. Est.	9 pars. Lin. Neur.	10 pars. Lin
10 pars.	Quad	-1319 [-2541, -611]	-236 [-1072, 358]	-154 [-772, 613]	-729 [-1365, -38]	323 [-423, 1127]
10 pars.	Lin	-1690 [-2534, -976]	-595 [-927, -311]	-492 [-1023, 238]	-1087 [-1690, -245]	
9 pars.	Lin. Neur.	-591 [-2068, 460]	486 [-504, 1211]	591 [406, 789]		
8 pars.	Ori. Est.	-1190 [-2614, -144]	-114 [-1026, 579]			
9 pars.	Bayes- <i>dN</i>	-1095 [-1657, -629]				

Table S4: Cross comparison of all models in **Figure S15**. See **Table S12** caption.

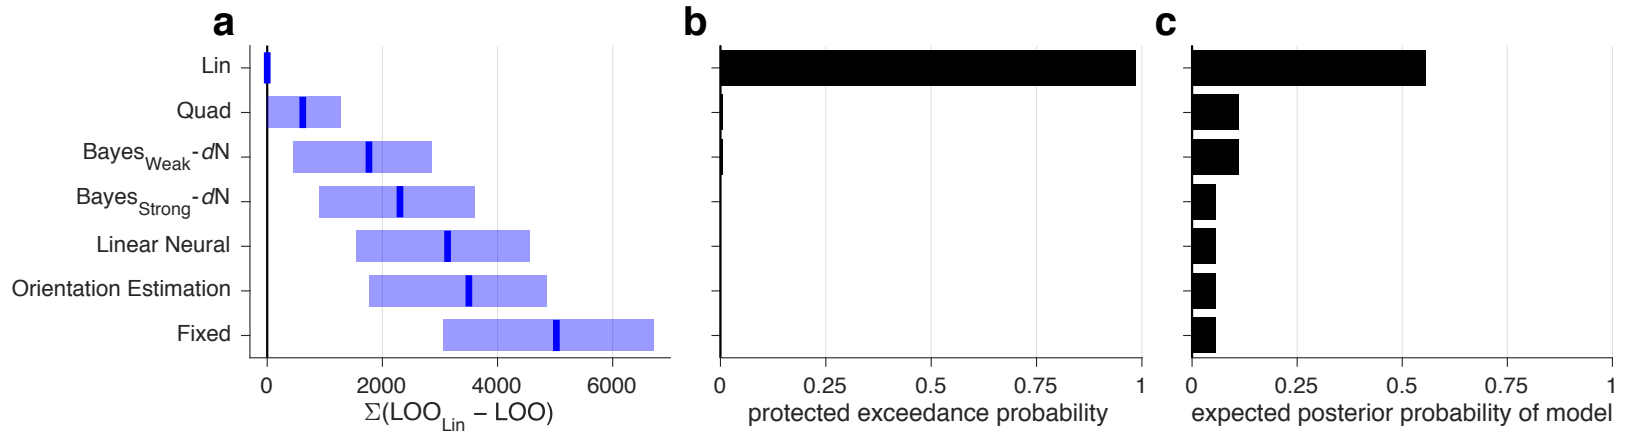


Figure S16: Model comparison, experiment 1. Noise parameters were fit to Task A category choices and then fixed during the fitting of Task B category and confidence responses. See **Figure S12** caption.

		15 pars. Fixed	13 pars. Bayes _S -dN	16 pars. Bayes _W -dN	15 pars. Ori. Est.	16 pars. Lin. Neur.	22 pars. Lin
22 pars.	Quad	-4367 [-6304, -2391]	-1670 [-3268, -39]	-1135 [-2501, 333]	-2836 [-4544, -1122]	-2449 [-4200, -969]	606 [6, 1269]
22 pars.	Lin	-5016 [-6727, -3090]	-2303 [-3578, -921]	-1773 [-2845, -451]	-3497 [-4860, -1817]	-3127 [-4549, -1575]	
16 pars.	Lin. Neur.	-1837 [-3566, -378]	846 [-609, 2092]	1386 [-13, 2732]	-345 [-933, 262]		
15 pars.	Ori. Est.	-1498 [-2877, -420]	1184 [23, 2129]	1724 [575, 2808]			
16 pars.	Bayes _W -dN	-3257 [-3965, -2494]	-533 [-920, -283]				
13 pars.	Bayes _S -dN	-2704 [-3351, -2027]					

Table S5: Cross comparison of all models in **Figure S16**. See **Table S12** caption.

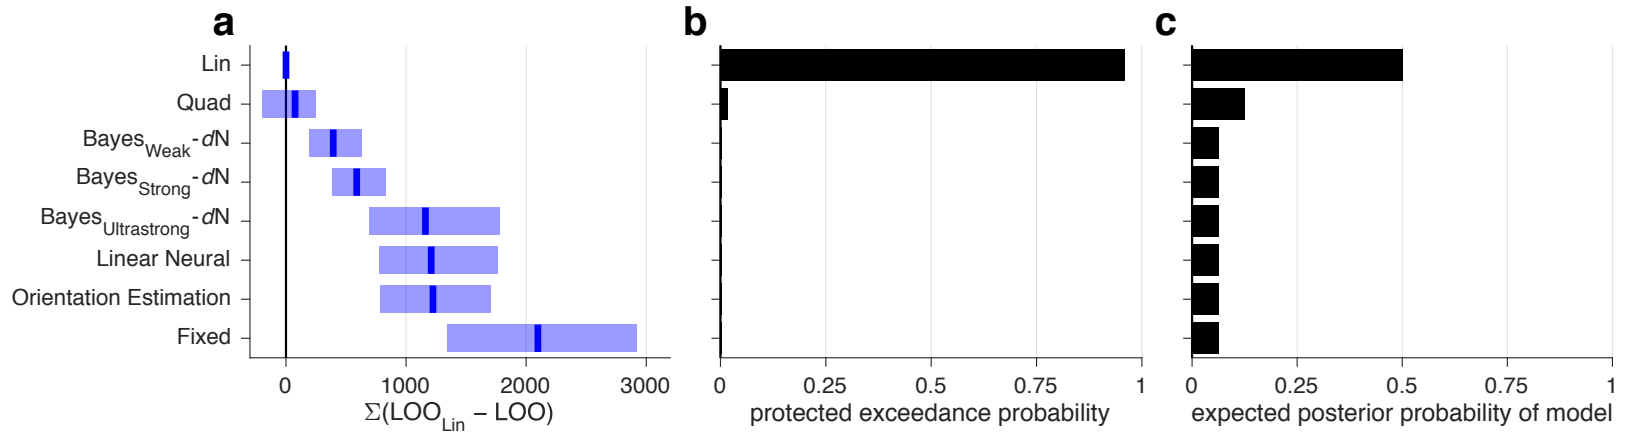


Figure S17: Model comparison, experiment 2. Models were fit jointly to Task A and B category and confidence responses. See **Figure S12** caption.

		19 pars. Fixed	13 pars. Bayes _U -dN	17 pars. Bayes _S -dN	20 pars. Bayes _W -dN	19 pars. Ori. Est.	20 pars. Lin. Neur.	30 pars. Lin
30 pars.	Quad	-2014 [-3036, -1186]	-1096 [-1807, -530]	-523 [-893, -220]	-331 [-562, -109]	-1136 [-1815, -638]	-1124 [-1922, -613]	74 [-195, 252]
30 pars.	Lin	-2095 [-2889, -1344]	-1160 [-1780, -694]	-589 [-841, -375]	-396 [-622, -186]	-1218 [-1680, -791]	-1205 [-1757, -785]	
20 pars.	Lin. Neur.	-876 [-1401, -395]	55 [-711, 693]	623 [99, 1184]	801 [253, 1491]	-7 [-219, 216]		
19 pars.	Ori. Est.	-872 [-1297, -487]	56 [-561, 574]	620 [218, 1089]	813 [356, 1394]			
20 pars.	Bayes _W -dN	-1678 [-2542, -1007]	-767 [-1262, -386]	-190 [-363, -82]				
17 pars.	Bayes _S -dN	-1490 [-2210, -886]	-565 [-1032, -266]					
13 pars.	Bayes _U -dN	-907 [-1572, -365]						

Table S6: Cross comparison of all models in **Figure S17**. See **Table S12** caption.

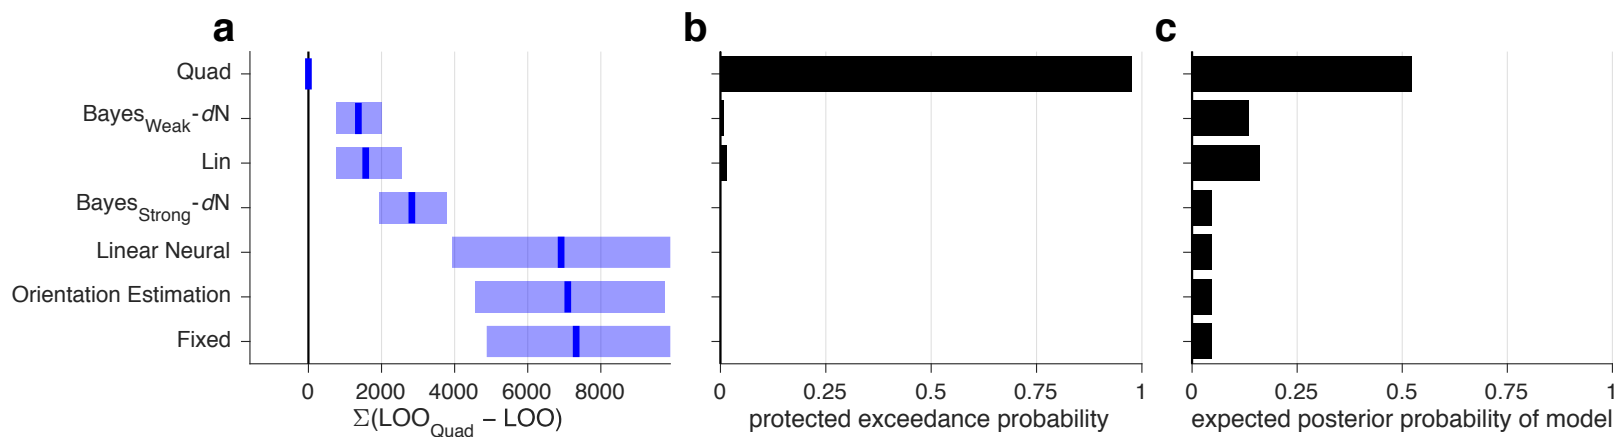


Figure S18: Model comparison, experiment 3. Models were fit to Task B category and confidence responses. See **Figure S12** caption.

		15 pars. Fixed	13 pars. Bayes _S -dN	16 pars. Bayes _W -dN	15 pars. Ori. Est.	16 pars. Lin. Neur.	22 pars. Lin
22 pars.	Quad	-7326 [-9955, -4905]	-2833 [-3807, -1926]	-1361 [-2022, -777]	-7120 [-9838, -4636]	-6902 [-10376, -3981]	-1577 [-2562, -750]
22 pars.	Lin	-5759 [-7866, -3694]	-1240 [-2567, 65]	226 [-812, 1246]	-5530 [-7707, -3539]	-5337 [-8191, -2846]	
16 pars.	Lin. Neur.	-450 [-1535, 1290]	4114 [733, 7796]	5552 [2338, 9135]	-214 [-1176, 1256]		
15 pars.	Ori. Est.	-256 [-841, 423]	4311 [1432, 7134]	5727 [3067, 8527]			
16 pars.	Bayes _W -dN	-5967 [-8702, -3369]	-1454 [-2179, -835]				
13 pars.	Bayes _S -dN	-4505 [-7282, -1816]					

Table S7: Cross comparison of all models in **Figure S18**. See **Table S12** caption.




Figure S19: Model comparison, experiment 3. Models were fit to Task B category choices. See **Figure S12** caption.

		7 pars. Fixed	8 pars. Bayes- <i>dN</i>	7 pars. Ori. Est.	8 pars. Lin. Neur.	8 pars. Lin
8 pars.	Quad	-777 [-1361, -359]	162 [-290, 670]	-531 [-1059, -23]	-988 [-1526, -549]	290 [-138, 793]
8 pars.	Lin	-1084 [-1675, -619]	-117 [-436, 76]	-830 [-1317, -334]	-1294 [-1825, -778]	
8 pars.	Lin. Neur.	215 [-566, 827]	1174 [535, 1772]	457 [254, 685]		
7 pars.	Ori. Est.	-255 [-987, 369]	707 [119, 1259]			
8 pars.	Bayes- <i>dN</i>	-964 [-1290, -663]				

Table S8: Cross comparison of all models in **Figure S19**. See **Table S12** caption.