

# Discrete confidence levels revealed by sequential decisions

Matteo Lisi

Centre for Applied Vision Research, City, University of London, UK

Gianluigi Mongillo

&

Andrei Gorea

Université Paris Descartes, Sorbonne Paris Cité, Paris, France

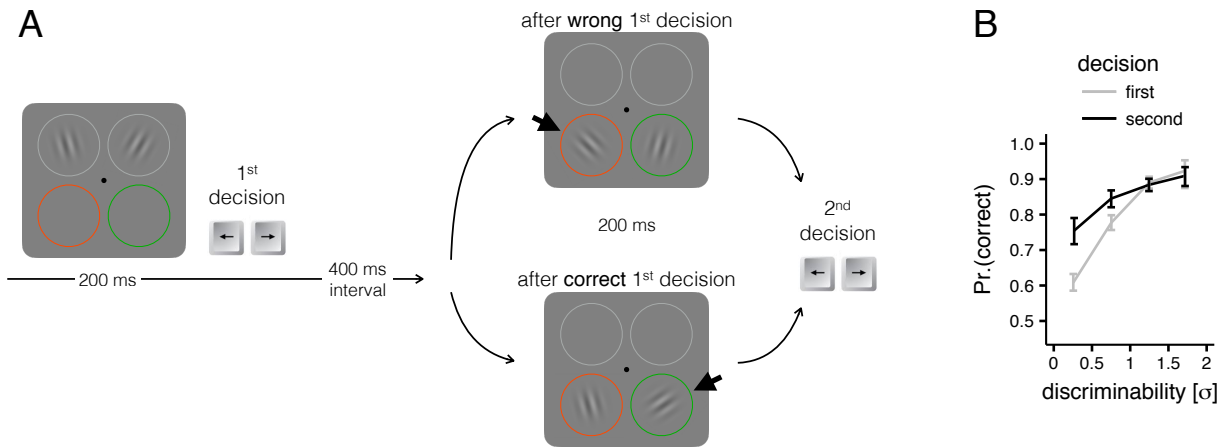
## Abstract

Whether humans are optimal decision makers is still a debated issue in the realm of perceptual decisions. Taking advantage of the direct link between an optimal decision-making and the confidence in that decision, we offer a new dual-decisions method of inferring such confidence without asking for its explicit valuation. Our method circumvents the well-known miscalibration issue with explicit confidence reports as well as the specification of the cost-function required by 'opt-out' or post-decision wagering methods. We show that observers' inferred confidence in their first decision and its use in a subsequent decision (conditioned upon the correctness of the first) fall short of both the ideal Bayesian strategy, as well as of an under-sampling approximation or a modified Bayesian strategy augmented with an additional bias term to accommodate global miscalibration of confidence. The observed data are instead significantly better fitted by a model positing that observers use only few confidence levels or states, at odds with the continuous confidence function of stimulus level prescribed by a normative behavior. These findings question the validity of normative-Bayesian accounts of subjective confidence and metaperceptual judgments.

## Introduction

How do we select behavioral responses when interacting with a given environment? According to a popular theory – the rational choice theory – human agents evaluate, for each decision to be taken, the expected utility of each possible course of action, and select thereafter the one yielding the highest rank. Despite the fact that the *descriptive* adequacy of rational choice theory has long been challenged on empirical as well as on theoretical grounds, mainly questioning its biological/psychological plausibility [16, 17, 20, 29, 42], the debate is far from settled [5, 6, 19]. In order to maximize expected utility, an agent has to be able to associate a *subjective* probability to each possible consequence of its actions. The characterization of this process requires (i) the quantitative assessment of the agent's ability to attach probabilities to events, and (ii) the appraisal of the measured subjective probabilities against those of a Bayesian observer, with the same prior knowledge as the agent, would assign to the same events. The experimental instantiation of such comparisons has been hindered by serious methodological problems with assessing subjective probabilities.

Subjective probability is tantamount to the agent's confidence in the occurrence of an event  $E$  [10]. It is formally defined as a marginal rate of substitution [10, 40]: the agent's subjective probability about the event  $E$  occurring or having occurred would be  $p(E)$  if the agent is indifferent to gaining one unit of utility contingent on  $E$  against gaining  $p(E)$  units of utility for sure. Current methods for measuring subjective probabilities in, for example, perceptual decisions (that is, the confidence about choices being correct) using opt-out or post-decision wagering techniques are straightforward operationalizations of the above definition. A major, well and long known problem with these methods is that they rely on unverifiable assumptions about the utility function of the participant (see, e.g., [40]). Such methods cannot disentangle subjective probability from factors such as opportunity cost in waiting-time paradigms. More frequently used, methods



**Figure 1: Experimental paradigm and results.** **A.** Orientation-orientation (O-O) paradigm. Participants had to make two consecutive decisions on which of two Gabor-patches was more tilted away from the vertical (hence the signal on which the participant had to base the decision was the signed difference in orientation between the two Gabors). While the absolute difference in tilt between the two Gabors (the discriminability) is drawn independently from a uniform distribution in each of the two Gabor-pairs, the location of the more tilted Gabor in the second Gabor-pair (left/right; indicated by the short thick black arrows in the right-hand panel) was made dependent on the accuracy of participant's first response. **B.** Performance (proportion of correct responses) plotted as a function of discriminability (measured in units of internal noise) separately for 1<sup>st</sup> and 2<sup>nd</sup> decisions (averaged over all conditions and experiments). Note how performance in the 2<sup>nd</sup> decision is higher, especially for trials where the sensory evidence (or discriminability, expressed in standard deviations of internal noise) is small and, in the absence of prior information, performance would be at or near chance. Error bars represent bootstrap 95% CI.

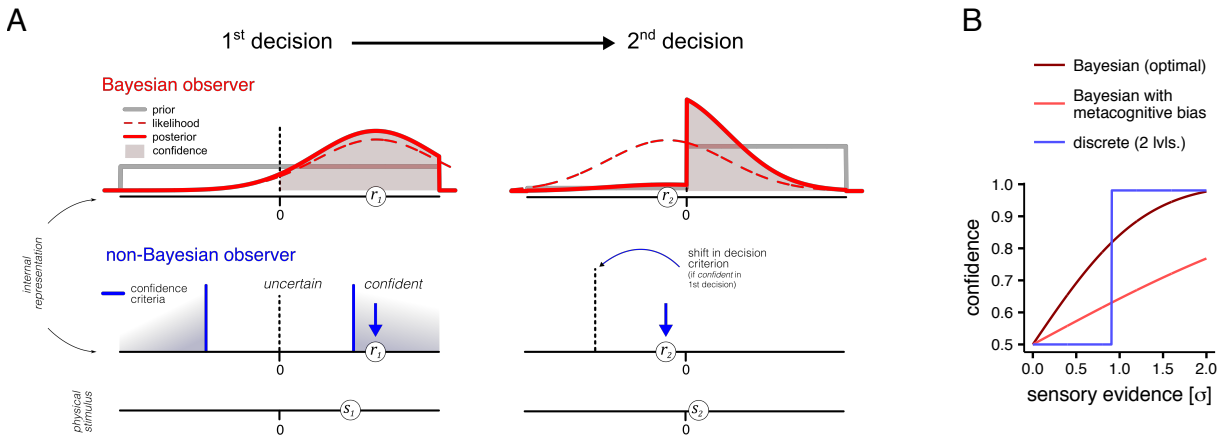
1 requiring *explicit* confidence valuation by the decision-maker suffer from well-documented miscalibrations  
 2 and response biases (see, e.g., [14, 15, 28]).

3 Here, we present a novel approach of estimating subjective probabilities which overcomes the problems  
 4 above. Human participants were presented with two consecutive signals and asked to decide whether they  
 5 were above or below some reference value. The key innovation was that the statistics of the second signal was  
 6 made contingent upon the decision-maker having made a correct decision on the first signal (explicit feedback  
 7 is not provided on a trial-by-trial basis): correct/incorrect first decisions resulted into signals above/below  
 8 the reference value for the second decision (a diagram of the experimental protocol is shown in Fig. 1 in the  
 9 context of an orientation discrimination task; the signal here is the difference in orientation between the two  
 10 Gabors). Differences in performance between the second and the first decisions, at signal parity, allow the  
 11 estimation of the subjective probability of being correct on the first decision (i.e., the confidence). Using  
 12 this approach we show that humans are quite accurate in assessing confidence, yet they exhibit systematic  
 13 deviations from optimality, such as global under-confidence and high-confidence errors. These systematic de-  
 14 viations cannot be accounted for either by a biased Bayesian observer (which perform Bayesian computations  
 15 using a biased estimate of the variability of the internal signals) or by a sample-based approximation of the  
 16 optimal Bayesian strategy (see Supplemental information). An alternative non-Bayesian model, character-  
 17 ized by a finite number of discrete confidence levels, provides the best and most parsimonious description of  
 18 the empirical patterns of humans' choices. These results suggest that the evaluation of subjective confidence  
 19 in human observers is not based on a full posterior distribution, and provide new cues about which kind of  
 20 shortcuts and approximation may be used by the brain to handle meta-perceptual uncertainty.

## 21 Results

### 22 The dual-decision paradigm

23 The paradigm is illustrated in Figure 1A with an example of the orientation discrimination task (orientation-  
 24 orientation or O-O condition, see Material and methods and Supplemental information for details). Partici-  
 25 pants are presented with two consecutive Gabor-pairs, and for each pair must decide which of the two Gabors  
 26 was more tilted with respect to the vertical; correct/incorrect first decisions result into displaying the second



**Figure 2: Bayesian and a class of non-Bayesian models of confidence and of sequential decision-making.** **A.** Internal beliefs of the Bayesian (top panels) and non-Bayesian observer (bottom panels), for 1<sup>st</sup> and 2<sup>nd</sup> decision (from left to right). The lowest abscissae represent the state of the world (e.g. the physical difference in orientation between the two Gabors in the orientation task of Fig. 1) while the internal representations of the decision variable (the sensory evidence) are referred to the upper two abscissae. The first stimulus in this ad-hoc trial ( $s_1$ ) evokes a noise contaminated internal response ( $r_1 = s_1 + \eta$ , where  $\eta \sim \mathcal{N}(0, \sigma^2)$ ). The Bayesian observer (top panels) knows about the statistics of its perceptual system and is able to compute the likelihood function  $p(r_1|s_1)$ . Because for the 1<sup>st</sup> decision the prior is flat (grey lines in the upper panels), the likelihood (dashed curves) largely overlaps with the posterior (red continuous curves). The Bayesian probability that represents the subjective confidence  $c_1$  that the true value of the signal lies in the positive semi-axis corresponds to the area of the posterior distribution shaded in red ( $c_1 = p(\hat{s}_1 > 0|s_1)$ ). The observer knows that the second signal will be drawn from the positive interval (right-hand side) after a correct 1<sup>st</sup> decision, and from the negative semi-interval after an error. Hence the ideal observer in the 2<sup>nd</sup> decision assigns a prior probability (grey line) equal to the confidence in the first choice to positive signal values. Inasmuch as the response to the first task was correct, the true value of  $s_2$  for the 2<sup>nd</sup> decision is positive (lower mid-panel).  $s_2$  being however small (as illustrated), it evokes by chance (due to internal noise) a negative internal response  $r_2$  that would (in the absence of prior information) lead to the erroneous conclusion that  $s_2$  was negative. Nevertheless, given the asymmetrical prior, the posterior distribution is still favouring the correct choice that  $s_2$  is positive. The lower left panel illustrates how a non-Bayesian observer (who does not have knowledge about the statistics of the internal noise, and only perceives point estimates) could make the same choices as the Bayesian observer by comparing the internal response  $r_1$  with a confidence criterion (the two vertical blue lines symmetrically placed about the decision criterion, dashed vertical line). By classifying decisions in two discrete levels, 'confident' vs. 'uncertain' or 'non-confident', and by shifting the 2<sup>nd</sup> decision criterion only after 'confident' 1<sup>st</sup> decision, the non-Bayesian 'discrete' observer can also increase the frequency of correct second choices. **B.** Relationship between confidence and internal sensory evidence for the Bayesian and non-Bayesian models, and for another model of a Bayesian observer that has a biased estimate of the internal variability, and can therefore display over or under confidence (relative to optimal). For the discrete and the biased Bayesian models the curve is computed by averaging MLE estimates of the parameters across subjects. Note that the discrete model (blue curve) does not actually compute probabilities, however it adjusts the criterion for the second decision. This criterion shift can be transformed into the equivalent confidence level of the ideal Bayesian observer.

1 pair with the more tilted Gabor in the right/left place-holder, respectively, and the participants are told  
 2 so. The very same experimental format was used with a duration discrimination task (duration-duration or  
 3 D-D condition) where participants had to decide which of two Gaussian blob flashes (presented sequentially)  
 4 was displayed for a longer duration. These two conditions were tested both in a first experiment where the  
 5 difficulties of first and second decision were independently drawn from a uniform distribution (random-pairs  
 6 experiment) and, on a different group of subjects, in another experiment where these difficulties were not  
 7 independent (correlated-pairs experiment): more specifically difficult/easy first decisions were more likely  
 8 to be followed by easy/difficult second decisions, respectively (see Material and methods). This correlation  
 9 was introduced to encourage participants to rely more on information provided by the first pair of stimuli  
 10 and exploit the statistical dependence between the signals in the two decision. Additionally, we also tested  
 11 a condition where the two tasks were combined, duration-orientation or D-O condition, to test within our  
 12 paradigm the proposal that confidence may work as a 'common currency' between different perceptual judg-  
 13 ments [12, 13].

14

**Table 1:** Results of the logistic analysis that measured the influence of the correctness of the first decision on the probability of choosing 'right' in the second decision (after accounting for the effect of stimulus values). The reported odds-ratio is between the odds of having the observer choosing 'right' (in the second decision) after a correct first decision and after a wrong one (given the same set of stimuli). The column  $p(\text{correct})$  displays the mean and standard deviations of the proportion of correct responses in the second decision, computed across observers. The last two columns are 95% CI of the odds-ratio for each condition.

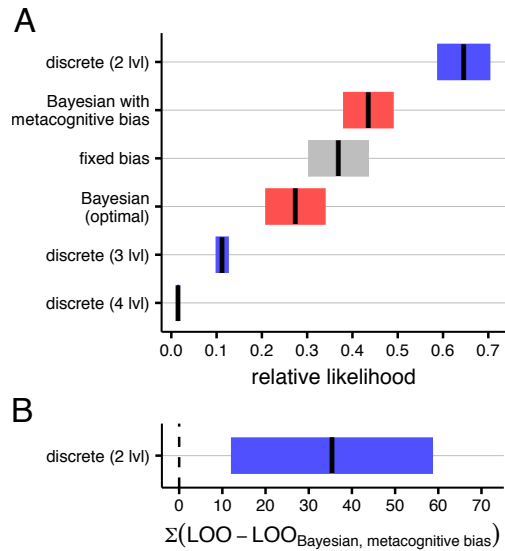
Condition	Experiment	$p(\text{correct})$	$\chi^2$	$p$	odds-ratio	95%LCL	95%UCL
O-O	correlated-pairs	0.86(0.05)	12.50	0.0004	6.21	3.09	12.47
O-O	random-pairs	0.88(0.05)	1.23	0.27	3.47	0.43	27.97
D-D	correlated-pairs	0.88(0.06)	14.16	0.0001	10.12	4.88	21.00
D-D	random-pairs	0.87(0.04)	16.02	<.0001	32.75	15.82	67.76
D-O	correlated-pairs	0.85(0.07)	14.37	0.0001	9.68	4.59	20.42

## 1 Computational models

2 Given the specific stimuli presented in the first decision and participant's trial-by-trial responses, participant's  
3 performance in the dual-decision paradigm can be compared with that of an ideal observer that accurately  
4 estimates the probability of being correct in the first decision and uses it as prior information for the second  
5 decision, according to the rules of Bayesian decision theory (see Figure 2A top). Assuming that the sensory  
6 noise is adequately described by a Gaussian distribution, updating prior expectations amounts to shifting the  
7 decision criterion for the second decision by the same distance observed between the first signal and criterion,  
8 expressed in standard deviations of sensory noise,  $\sigma$  (see Supplemental information for the full derivation).  
9 We have also compared human's behavior with that of three other classes of models: (i) a biased Bayesian  
10 observer model, which also performs probabilistic computations to estimate its confidence, albeit using a  
11 biased estimate of the variability of its internal response (this allows the model to reproduce patterns of  
12 over or under confidence, which would correspond respectively to under or over estimation of the variability  
13 of the internal noise); (ii) a heuristic, or non-Bayesian observer characterized by a finite number of discrete  
14 confidence levels (this model implements the idea that observers may make perceptual decisions using only  
15 point estimates of the decision variable, rather than the full probability distribution, see Figure 2A bottom);  
16 and (iii) a model representing a sample-based approximation of the Bayesian model (where the posterior  
17 distribution is approximated by finite number of samples; see Supplemental information). These models  
18 made different predictions of the relationship between sensory evidence and confidence (see Figure 2B) and  
19 can be reliably discriminated in synthetic datasets (see Supplemental information, Figure S3). Compared to  
20 the optimal Bayesian model, the biased Bayesian model is characterized by one additional free parameter,  
21 a scaling factor that indicates by how much the participants over or under estimate the variability of their  
22 sensory noise. The non-Bayesian models instead are characterized by two additional free parameters for  
23 each additional discrete confidence level: the first parameter indicates the location of the confidence criterion  
24 (vertical blue lines, Figure 2A bottom), and the second parameter indicates by how much the criterion for the  
25 second decision is shifted following a 'confident' first decision (black dashed vertical line, in Figure 2A bottom-  
26 right). A summary of the estimated values of model parameters is reported in Supplemental information,  
27 Table S2.

## 28 Model-free analyses

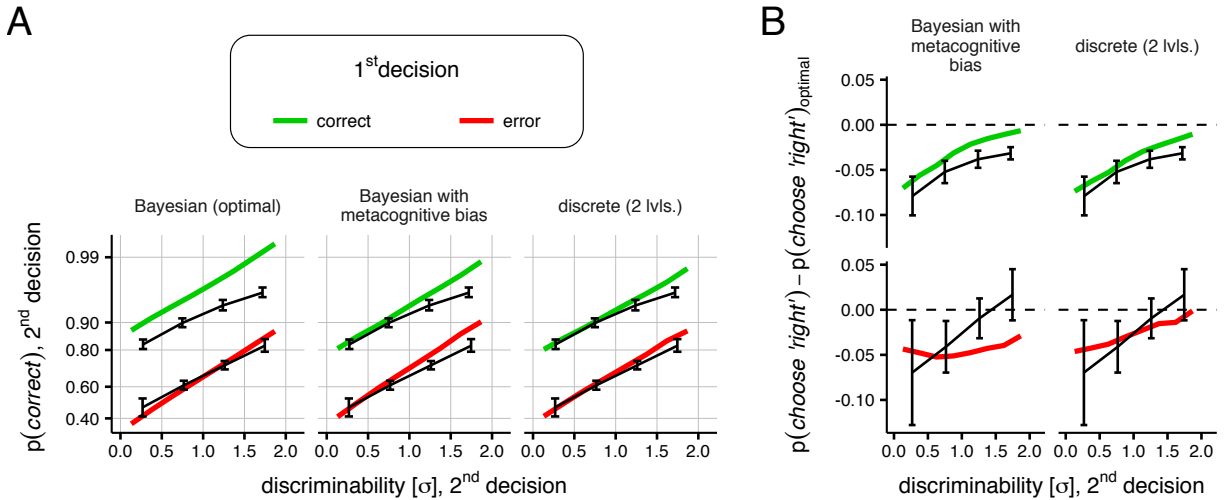
29 Figure 1B shows the proportion of correct responses as a function of discriminability (measured in units of  
30 internal noise) in the first and second decision (see Figure S1 for separate plots for each experiment and  
31 condition). As it can be seen, second decision performance (darker traces) is higher than first decision  
32 performance, especially for the most difficult trials (where the difference between the two stimuli in a pair  
33 was very small and first decision performance was close to chance). Since the more tilted or longer stimuli  
34 appeared more often in the right placeholder for the second decision (given that task difficulty was adjusted  
35 so as to yield average first-decision performance above chance), an increase in second-decision performance  
36 could be the result of a fixed bias (i.e., participants having chosen a 'right' response more frequently), without  
37 necessarily involving a trial-by-trial monitoring of uncertainty. To control for this possibility we performed a



**Figure 3: Model comparison.** **A.** Relative likelihoods  $\mathcal{L}(m|\text{data})$ , averaged across participants (see main text for details). Models depicted in red are probabilistic models which use full probabilistic distributions over the decision variable; in blue are models that only use point-estimates, but still evaluate the uncertainty heuristically on a trial-by-trial basis; the model in grey represents an observer that does not monitor uncertainty but simply adopts a fixed bias, corresponding to the belief of being more frequently correct than wrong in the first decision. Black lines represent mean likelihoods and error bars represent bootstrapped standard errors of the mean (SEM). **B.** Leave-one-out (LOO) cross-validated log-likelihood differences (discrete 2lvl minus Bayesian with a metacognitive bias) summed over subjects and task (the error bar is the bootstrapped standard error of the sum).

1 logistic analysis to measure the influence of the correctness of the first decision on the probability of choosing  
2 'right' on the second decision (i.e. reporting that the signal was drawn from the positive semiaxis, that is  
3  $s_2 > 0$ ). For each experiment (correlated- and random-pairs) and condition (O-O, D-D, and D-O) we fitted  
4 a multilevel (mixed-effects) logistic regression, using **R** [36] and the `lme4` package [2], with the absolute  
5 difference between the two stimuli (in units of  $\sigma$ ) and the accuracy of the first response as fixed effect  
6 predictors, and the participant as a grouping factor. We evaluated statistically the effect of the correctness  
7 of the first response with a likelihood ratio test between the fitted model and a reduced model where the  
8 effect of the first response was set to 0. This test was significant for all the experiments and conditions (see  
9 Table 1), with the exception of the O-O condition in the random-pairs experiment. In order to check whether  
10 a simpler fixed-bias model would really suffice to describe performance in this latter case, we performed a  
11 Monte Carlo simulation. For each trial we estimated the expected probability of a 'right' second response on  
12 the basis of the stimuli presented and the psychometric function fitted to the first decision responses. This  
13 results in a set of Bernoulli trials with different probabilities of success, which we simulated  $10^5$  times in  
14 order to estimate, using the percentile method, a 99.5% confidence interval (corresponding to the Bonferroni  
15 corrected alpha level 0.005) on the expected proportion of second responses 'right' given the stimuli and the  
16 first response. The results revealed that after a wrong first response none of the participants responded 'right'  
17 more often than what would be expected given the stimuli: all the observed proportions were within the  
18 confidence intervals. Instead, after a correct response, the observed proportion of responses 'right' exceeded  
19 the confidence interval for 3 out of the 5 participants in this experiment. This result suggests that also in  
20 this condition, where the fixed bias hypothesis could not be rejected at the group level, the evidence favors  
21 the hypothesis that most participants have monitored the confidence in their first response on a trial-by-trial  
22 basis.

23 Overall, choice accuracies in the second decision revealed only small differences across task (see the  
24 also table 1). We investigated also whether our manipulation of the statistical relationship between the  
25 difficulties of first and second decisions did affected performance. We run a multilevel logistic regression,  
26 with the accuracy of the second decision as dependent variable, and the stimulus differences (absolute values)  
27 together with the experiment (random vs. correlated pairs) as predictors. This analysis did not reveal any  
28 significant difference in performance,  $\chi^2(1) = 0.55$ ,  $p = 0.46$ , therefore data from the two experiments  
29 were collapsed in subsequent analyses. We used the same method to compare conditions where the two  
30 decisions were made in the same modality (O-O and D-D) with the condition where they were made on  
31 different modality (D-O; for this comparison only data from the correlated pairs experiment were used).  
32 Interestingly, we found that the difference in accuracy between these two conditions, although small ( $\approx$   
33 2% see table 1), resulted statistically significant  $\chi^2(1) = 8.24$ ,  $p = 0.004$ . We will discuss the possible  
34 implications of this finding in the Discussion.



**Figure 4: Comparison of model predictions and empirical data.** **A.** The predicted (colored lines) and observed (black lines, with error bars representing  $\pm 1$  bootstrapped SEM) proportion of correct choices (i.e. reporting  $s_2 > 0$ ), split according to the correctness of the first decision and averaged across all observers and conditions, are plotted on a probit scale as a function of the discriminability of the second pair of stimuli (expressed in units of internal noise,  $\sigma$ ). **B.** Differences in the proportion of choices 'right' of the biased Bayesian model (left panels), and of the discrete model with 2 levels of confidence (right panels). The upper panels demonstrate the under-confidence (relative to optimal), present in the data and reproduced by both models. The lower panels display choices after a wrong first decision, demonstrating the pattern of high confidence errors: after a wrong first response, subjects do not show the same pattern of under-confidence as after a correct first response. Instead, they tend to choose 'right' (and therefore commit an error) almost as often as the ideal, particularly for easy second decisions.

## 1 Model comparison

2 We compared the models using the Akaike Information Criterion, AIC [1], to account for the different number  
 3 of free parameter. For ease of interpretation, AIC scores were converted to a probability scale by taking the  
 4 difference with respect to the AIC of the best fitting model (separately for each participant) and transforming  
 5 the result into relative likelihoods of the model given the data,  $\mathcal{L}(m|\text{data}) \propto \exp[-\frac{1}{2}(\text{AIC}_m - \min \text{AIC})]$  [8].  
 6 We also included into the model comparison a fixed-bias model, which corresponds to the hypothesis that  
 7 observers did not monitor their own uncertainty on a trial-by-trial basis, but simply increased the frequency  
 8 of responses 'right' in the second decision by shifting their decision criterion by a fixed amount. Although  
 9 this hypothesis was rejected by model-free analyses, it is useful to see how it compares to the other models.  
 10 We find that the model with the higher relative likelihood was the non-Bayesian observer with 2 discrete  
 11 confidence levels (Figure 3A). In particular, the relative likelihood (averaged across tasks and conditions)  
 12 of the non-Bayesian observer with 2 confidence levels was higher than the ideal Bayesian observer for 13  
 13 out of 14 observers. Most interestingly, it was also higher than the second best model (the Bayesian model  
 14 with metacognitive bias) for 11 out of 14 observers. As an additional test we run a leave-one-out (LOO)  
 15 cross-validation of the two best fitting models, and found again that the non-Bayesian model was on average  
 16 better at predicting the actual behavior of the participants (Figure 3B).

17 Given the results of the model comparison, it is important to understand the differences between the  
 18 predictions of the Bayesian and non-Bayesian models. Figure 4A displays on a probit scale the observed pro-  
 19 portion of correct second decision choices as a function of the difficulty of the first decision, averaged across  
 20 observers, tasks and conditions. The examination of this plot reveals several interesting features. First, the  
 21 left panel in Figure 4A indicates that the main difference between the data and the prediction of the ideal  
 22 Bayesian observer is that human observers make more frequent errors than the ideal after a correct first deci-  
 23 sion, which implies that they choose 'right' less frequently than ideal. Because in our task the bias favouring  
 24 'right' responses in the second decision is a function of the confidence in the first decision (see Supplemental  
 25 information and Figure 2), this pattern denotes *under confidence*, i.e. human observers consistently under-  
 26 estimate their probability of being correct in the first decision. The finding of a global under-confidence bias  
 27 allows us to reject the alternative models based on a sampling scheme, in which probability distributions

1 are approximated by a finite number of samples. We find that sampling-based approximations of the ideal  
2 Bayesian model cannot adequately describe the data, at least not without a number of additional assump-  
3 tions, because for a small number of samples they produce a marked overconfidence bias (see Supplemental  
4 information), opposite to the one found in the data. Note that as the number of samples increases the  
5 sampling models converge to the ideal Bayesian model. The biased Bayesian observer instead can account  
6 for this pattern, by assuming that participants use a biased estimate of their sensory noise (figure 4A, middle  
7 panel); maximum likelihood estimation of this parameter indicates that the under confidence bias in the data  
8 could be explained by assuming that they over-estimate their sensory noise approximately by a factor of 3.

9 The biased Bayesian observer is less confident than the ideal also after wrong first decisions, leading to  
10 less frequent errors in the second decision<sup>1</sup>. This implies that after an error, human observers respond more  
11 frequently 'right' than predicted by the biased Bayesian model, indicating that they make more frequent  
12 high-confidence errors (or alternatively that they are on average more confident when they make a wrong  
13 choice). This can be appreciated better in Figure 4B, where the predicted and observed rates of 'right' second  
14 decision choices are represented as differences from the ideal. The upper panels show the under-confidence  
15 pattern (after correct first decisions). This is predicted by both the biased-Bayesian and the alternative, non-  
16 Bayesian discrete model. The lower panels of Figure 4B show that the under-confidence is not systematically  
17 present after a wrong first decision; instead, observers deviate from the predictions of the biased Bayesian  
18 observer and respond more frequently 'right'; in particular they chose 'right' as frequently as the ideal for  
19 trials where the intensity of the second signal was high. This implies that they were less sensitive than the  
20 biased Bayesian observer to the evidence provided by the second pair of stimuli when their first decision was  
21 wrong.

22 Only the discrete model predicts rates of 'right' responses that are closer to the observed ones, simulta-  
23 neously showing underconfidence after correct first decisions, and more high-confidence errors. According to  
24 the discrete model, high-confidence errors are to be expected whenever the internal signal not only has oppo-  
25 site sign with respect to the physical stimulus (due to random noise fluctuations), but is also large enough to  
26 exceed the confidence threshold (see Figure 2) and consequently elicit a 'confident' wrong decision. The rela-  
27 tionship between the internal signal, or sensory evidence, and the confidence predicted by the ideal Bayesian  
28 observer and the two best fitting models is plotted in Figure 2B: high confidence errors occur whenever the  
29 signal upon which the wrong decision is based falls in the high-confidence region of the discrete model (blue,  
30 stepwise curve). With our behavioral task we cannot measure the trial-by-trial internal signals, however we  
31 can compute their expected value as the mean of their probability distribution conditioned on the stimulus  
32 and the choice (this amounts to taking the mean of a truncated Gaussian distribution centered on the true  
33 value of the stimulus with upper or lower boundary at 0, depending on whether the observer decided 'left' or  
34 'right', respectively) and compare it to the individual confidence thresholds estimates. This analysis revealed  
35 that even for wrong first decisions, the internal signal was expected to exceed the confidence criterion in  $\approx$   
36 42% of the trials (and in  $\approx$  77% of correct first decisions), supporting the idea that a substantial proportion  
37 of high confidence errors is to be expected given the stimuli and the observed pattern of choices.

## 38 Discussion

39 We developed a novel approach where in a sequence of two dependent perceptual decisions humans could  
40 improve their second decision performance by taking advantage of the fact that the statistics of the stimuli  
41 presented for this second decision depended on the correctness of their first decision. This experimental  
42 protocol can be regarded as a laboratory proxy of more complex environments in which confidence is used  
43 to guide behavior in cases where a current decision depends on the unknown outcome of a previous one. We  
44 developed a normative Bayesian observer model for this task, i.e., an ideal observer who performs Bayesian  
45 inference to estimate the posterior probability of being correct in the first decision (the confidence), and uses  
46 it optimally (i.e., maximizes the probability of a correct second decision) as a prior probability distribution  
47 for making the second decision. By comparing participants' performance with the predictions of the Bayesian  
48 model, we were able to test the hypothesis that human observers are able to evaluate probabilistically their

---

<sup>1</sup>Since observers are, overall, more frequently correct than wrong in the first decision (observers made  $\approx$  80% correct first decisions), the ideal Bayesian model attains an overall better performance (based on the totality of trials).

1 own uncertainty and make meta-perceptual judgments according to a Bayesian-probabilistic strategy.

2 The present results revealed clearly that the participants engaged in this task exhibited systematic devia-  
3 tions from the predictions of the normative Bayesian model. While they were clearly able to take trial-by-trial  
4 uncertainty into account (a simple fixed bias model was not sufficient to account for the observed behavior:  
5 see Results, Model-free analyses), their pattern of second decisions revealed the presence of a marked, global  
6 miscalibration of confidence, more frequently in the form of under-confidence, as indexed by lower rates of  
7 responses 'right' in the second decision with respect to ideal. We developed an alternative 'biased' Bayesian  
8 model, designed to reproduce global miscalibration of confidence such as the observed under-confidence by  
9 assuming that observers use a Bayesian-probabilistic strategy but may have a biased estimate of the vari-  
10 ability of the internal signals. This model is in line with recent work showing that transcranial magnetic  
11 stimulation (TMS) on the visual cortex can degrade performance in a visual discrimination task without  
12 affecting comparative judgments of uncertainty made over multiple decisions [34]. Peters and colleagues [34]  
13 explained the results by assuming that human observers rely on a learned statistical model of their own per-  
14 ceptual system that may become invalid under specific conditions, such as the external disturbances elicited  
15 by means of the TMS. In that case subjects are metacognitively blind to the additional noise introduced by  
16 the TMS and therefore under-estimate their internal variability. In our case, the biased Bayesian observer  
17 model can reproduce the under-confidence bias by assuming that, on average, participants over-estimate  
18 their internal noise approximately by a factor of 3.

19 Although the biased Bayesian outperformed the ideal, unbiased Bayesian model in predicting the actual  
20 behavior, we found additional discrepancies between predictions and empirical data that question the biased-  
21 Bayesian's descriptive accuracy. Specifically, we found that on average the biased Bayesian model predicts  
22 less frequent (with respect to behavior) 'right' choices after a wrong first decision, see Figure 4, indicating  
23 that human observers made more frequent high confidence errors. Overall, our results show that (i) human  
24 observers show less confidence than the ideal Bayesian observer after a correct first decision, but at the same  
25 time (ii) after a first wrong decision are more confident than the alternative Bayesian model augmented  
26 with a bias parameter to account for the underconfidence (relative to ideal). Taken together, these findings  
27 indicate that the relationship between the internal signals (the sensory evidence) and the confidence does not  
28 have the shape that should be expected if subjective confidence were computed as by a Bayesian observer,  
29 namely as the posterior probability of being correct.

30 The pattern found in our data could be accounted for by our non-Bayesian model, in particular the simple  
31 model with only two discrete confidence levels<sup>2</sup>. From a psychological point of view, the non-Bayesian class  
32 of models posits that confidence is discretized in a number of distinct levels, as also suggested by previous  
33 work [44]. For example, in the one-criterion variant of our model there would be two discrete confidence  
34 states: confident - i.e. more likely that the response was correct - vs. non-confident - i.e., equally likely  
35 that the response was correct or wrong (see Figure 2). This one confidence criterion variant of our model  
36 provided the best and most parsimonious description of the empirical data (see Results, Model comparison).  
37 It is important to note that because this model does not assume any knowledge about the variability of the  
38 internal decision variable, these confidence levels are defined only on an ordinal scale and do not convey a  
39 precise numerical information about the probabilities involved.

40 Although our heuristic class of models may seem too simplistic to represents an accurate algorithmic  
41 description of the mechanisms underlying subjective confidence valuation, we argue that the specific way  
42 in which the relationship between sensory evidence and confidence is realized in this class of models - as a  
43 discrete step function - must capture some aspects of the actual mechanism implemented by the brain. More  
44 specifically, we propose that confidence grows with sensory evidence according to a piecewise function: more  
45 slowly than ideal for a certain range of values, or sub-domains, of sensory evidence (e.g., for values close to  
46 the decision boundary 0, contributing to produce the global under confidence bias), but at the same time it  
47 also grows faster than ideal in other sub-domains, resulting in more frequent high confidence wrong decisions.  
48 This proposal implies that subjective confidence does not match the Bayesian posterior probability of the  
49 ideal Bayesian observer, but instead is based on different computational processes that may require only  
50 point estimates, as opposed to full probability distribution. This pattern where subjective probabilities both  
51 overestimate and underestimate the objective probability in distinct sub-domains of the same task resemble

---

<sup>2</sup>Our analysis indicates that the increase in log-likelihood obtained by adding more than two confidence levels is small and does not compensate the parallel increase in model complexity (number of free parameters), suggesting overfitting.



1 the distortions of subjective probability that have been found also in non-perceptual task, such as those  
2 involving decisions from experience [45].

3 Given that participants' performance was slightly lower in the D-O condition (see Results, Model-free  
4 analyses), where first and second decisions were made in different modalities, it is reasonable to ask whether  
5 differences of the observed data from the predictions of the Bayesian models were mostly driven by this  
6 specific condition. Interestingly, looking at the model likelihoods separately for each condition, see Figure S4  
7 it becomes evident that this was not the case as the D-O condition was instead the only one where the  
8 likelihood of the biased-Bayesian resulted slightly higher than that of the non-Bayesian model with two  
9 confidence level. Taken together, both the pattern of model likelihoods, and the slightly lower performance,  
10 suggest that participants may have used a slightly different strategy in this condition. While our current  
11 data do not allow to pinpoint exactly what caused these different outcomes, they do suggest that claims of  
12 confidence as a 'common currency' between different perceptual judgments [12,13] should be re-evaluated  
13 using a more diverse set of experimental protocols.

14 Taken together, our behavioral and modelling results suggest that humans can use sensory evidence to  
15 perform probabilistic judgments, but ultimately cannot assign numerically precise subjective probabilities  
16 to perceptual interpretations. These comparative probability judgments, not linked to precise numerical  
17 values, are the essence of qualitative probability reasoning [39], a weakened but more pragmatic and intuitive  
18 counterpoint of classical probability theory. The notion of qualitative probability lies at the foundations of  
19 the notion of subjective (or Bayesian) probability since its early formulation [11,39]. Much of the early  
20 work on the theory of Bayesian probability has been dedicated to identifying the conditions that allow the  
21 departure from the qualitative probability toward the quantitative (numerically precise) probability, defined  
22 according to the classical Kolmogorov axioms [24]. Several propositions have been put forward, but all  
23 of them ultimately assume that the decision-maker's knowledge allows the partition the probability space  
24 associated with the event space into a uniform and arbitrarily large collection of disjoint states or events [39].  
25 While this assumption can be used to provide a formal framework for exact reasoning under uncertainty,  
26 it may be too fine-grained for a realistic, biological decision-maker. Indeed, in the real world assigning  
27 exact numerical probabilities is often difficult or impossible, and the ability to compare the likelihood of two  
28 events without having to provide exact probabilities may be sufficient. Our results support this idea, as the  
29 performance of the heuristic observer with two discrete confidence levels, despite its relative simplicity, shows  
30 a marked increase in the probability of correct second decisions relatively to a simpler model which does  
31 not consider any prior information (the expected proportion of second correct choices in the absence of any  
32 criterion shift was  $\approx 80\%$ ), while the ensuing performance benefit obtained by the Bayesian model may not  
33 be enough to justify its increased computational complexity [25] and costs. Indeed, it is known that efficient  
34 adaptive decisions can be taken also by means of simple heuristics, which may outperform more rigorous  
35 Bayesian strategies once the costs of information acquisition and processing are taken into account [18].

36 Our findings suggest that subjective confidence does not match the Bayesian posterior probability of being  
37 correct, even after taking into account the possibility that observers have a biased estimate of the internal  
38 noise or approximate the subjective probabilities with a finite set of samples. These results have important  
39 implications for the current investigation of the neural substrate of metacognition, given that many of the  
40 relevant studies assume, more or less explicitly, that subjective confidence corresponds to the Bayesian  
41 posterior probability (e.g., [23,30]). Although Bayesian decision theory provides a general and logic way  
42 of processing information, the evidence supporting the hypothesis that it is algorithmically implemented  
43 in the brain is debated [26,37]. We argue that to achieve real progress in the understanding of human  
44 decision-making and metacognition, it is fundamental that researchers consider also alternative hypotheses  
45 and models. Moreover, the view that Bayesian decision theory is always a normative description of rational  
46 decision making (sometimes referred to as "Bayesianism") has been criticized also on purely theoretical  
47 grounds [3]: while Bayesian decision theory is demonstrably optimal in the context of "small" worlds [4,39]  
48 where all relevant alternatives, consequences and probabilities can be meaningfully estimated and assigned,  
49 it is unclear to what extent it can be considered rational in the context of "large" worlds, where not all  
50 the alternatives, consequences and probabilities are known. While simple perceptual inference problems  
51 can, under certain conditions, be effectively considered "small world" problems, there is no real reason for  
52 why this should be the case of perception in general. This line of reasoning thus undermines the proposal  
53 that subjective confidence should necessarily be operationalized as the Bayesian posterior probability, and

1 motivates the need of empirical studies and novel approaches to test the descriptive accuracy, as well as the  
2 ecological rationality [9] of applying Bayesian decision theory as a *process* model [27] to explain the behavior  
3 of imperfect, biological decision makers. In the present study we have not only found further evidence  
4 that brain processes subtending meta-perception do not conform to the ideal benchmark represented by the  
5 Bayesian observer, but also provided a novel, general experimental protocol that we believe will be a valuable  
6 tool for future investigations of confidence and metacognition. The availability of a rich set of new protocols  
7 for assessing cognitive functions is also important for translating advances in neuroscience and cognitive  
8 science into concrete, theory-driven, clinical applications [21].

## 9 Acknowledgments

10 This work was supported by Grant ANR-12-BSH2-0005 from the French National Research Agency to A.  
11 Gorea.

## 12 Material and methods

13 Participants (5 for the random-pairs experiment, and 9 for the correlated-pairs experiment) performed a sequence of two decisions  
14 on each trial, with the statistics of the second decision stimuli being conditioned on the correctness of the first decision, with  
15 this dependence made explicit to the observers. All experiments were performed in accordance with French regulations and  
16 the requirements of the Helsinki Convention. The protocol of the experiments was approved by the Paris Descartes University  
17 Ethics Committee for Non-invasive Research (CERES).

18 We describe here the general structure of the protocol, with the details of the implementation and of the analysis being provided  
19 in the Supplemental information. At the beginning of each trial, two stimuli were presented in two placeholders on the left and  
20 right of the fixation point. The two stimuli differed along one physical dimension [orientation of the Gabor-patches – spatial  
21 frequency 1.5 cycles/dva (degrees of visual angle), standard deviation of the envelope 0.7 dva, contrast 25%, – or duration –  
22 Gaussian blobs with a peak intensity of  $\approx 25.3$  cd/m<sup>2</sup> and a standard deviation of 0.65 dva]. The participant was required to  
23 indicate which of the two stimuli was characterized by a higher value along the given dimension by pressing the left/right arrow  
24 keys. The difference between the two stimuli was uniformly distributed within 2 JNDs (just noticeable differences), measured  
25 in preliminary sessions (see below). 400 ms after providing the first response, a second pair of stimuli was presented and the  
26 participant was again asked to indicate which of the two has a higher value. The difference in value between the stimuli in  
27 the second pair was also randomly sampled from a uniform distribution. However the location of the higher-value stimulus  
28 depended this time on the correctness of the first response: if the first response was correct, the higher-value stimulus was  
29 presented on the right, and on the left otherwise. Participants were informed about this rule, and were asked to use it in order  
30 to achieve the best possible second decision accuracy. Before starting the experimental trials, participants were explained the  
31 rule and were familiarized with a version of the task where the difference between the two stimuli to be compared could go  
32 up to very high values (up to 45° in the orientation task, and up to 1 second in the duration task, uniformly distributed).  
33 The large differences in the practice session were intended to make the rule clear and unambiguous for all participants. In  
34 one first experiment discrimination difficulties in the first and second decisions were drawn independently (random-pairs). In a  
35 second experiment we biased the probability of association between discrimination difficulties in the first and second decision  
36 (correlated-pairs, see Supplemental information, Figure S5). Specifically, when the difference in intensity in the first pair was  
37 less than 1 JND, there was a 0.7 probability that the difference in the second pair would be larger than 1 JND, and vice versa.  
38 This was intended to encourage participants to make use of the rule. The random-pairs experiment was tested in two conditions,  
39 run in different sessions on different days (order balanced), where the two decisions involved both an orientation discrimination  
40 (O-O) or a discrimination of duration (D-D). Each of these sessions comprised 500 trials. The correlated-pairs experiment was  
41 declined in three different conditions, O-O, D-D and D-O (where the two decisions involved a duration discrimination followed  
42 by an orientation discrimination). Each of the three sessions comprised 300 trials, each consisting of two consecutive perceptual  
43 decisions. The different number of total trials in the correlated- and random- pairs experiments was designed so that they  
44 resulted in similar number of easy (difficult) decisions followed by difficult (easy) decisions. Each testing session was divided  
45 in 10 blocks of trials. At the end of each block participants were given a feedback about the overall accuracy of their second  
46 decisions in that block. Additionally, to help participants keep track of their performance, starting with the end of the second  
47 block they were also informed on whether their accuracy had increased or decreased with respect to the previous block.

## 48 Analysis

49 For each participant we estimated the standard deviation of the internal noise,  $\sigma$ , by fitting on the proportion of 'right' choices in  
50 the first decision four different psychometric models that made different assumptions about whether participants were biased or  
51 made frequent attention lapses (i.e., stimulus independent errors), and by combining the 4 estimates of the standard deviation  
52 by weighting them according to the Akaike weight of each model [8] (see Supplemental information for details). Next, we used  
53 the estimated  $\sigma$  to transform the values of the stimulus from raw units (e.g. degrees and seconds) to units of internal noise.  
54 Finally, we fitted the non-Bayesian models using maximum likelihood estimation, and compared the predictions of Bayesian and  
55 non-Bayesian models with the observed pattern of second decisions (all models make similar predictions for the first decision).  
56 Because these models differ in the number of free parameters, in order to prevent overfitting we compared the models on the

1 basis of the Akaike Information Criterion [1]. Additionally we performed a leave-one-out cross-validation of the two best-fitting  
2 models. All analyses were performed in the open-source software **R** [36]; the data and the code of the analysis are available  
3 upon request. The mathematical details of the computational models are provided in Supplemental information.

## 4 References

- 5 [1] H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723,  
6 dec 1974.
- 7 [2] D. Bates, M. Maechler, B. Bolker, and S. Walker. *lme4: Linear mixed-effects models using Eigen and S4*, 2014.
- 8 [3] K. Binmore. *Rational Decisions*. Princeton University Press, 2009.
- 9 [4] K. Binmore. On the Foundations of Decision Theory. *Homo Oeconomicus*, 2017.
- 10 [5] J. S. Bowers and C. J. Davis. Bayesian just-so stories in psychology and neuroscience. *Psychological bulletin*, 138(3):389–  
11 414, 2012.
- 12 [6] J. S. Bowers and C. J. Davis. Is that what Bayesians believe? reply to Griffiths, Chater, Norris, and Pouget (2012).  
13 *Psychological Bulletin*, 138(3):423–426, 2012.
- 14 [7] D. H. Brainard. The Psychophysics Toolbox. *Spatial Vision*, 10(4):433–436, jan 1997.
- 15 [8] K. P. Burnham and D. R. Anderson. *Model Selection and Multimodel Inference: A Practical Information-Theoretic*  
16 *Approach*. Springer New York, New York, US, 2nd editio edition, 2002.
- 17 [9] N. Chater, T. Felin, D. C. Funder, G. Gigerenzer, J. J. Koenderink, J. I. Krueger, D. Noble, S. A. Nordli, M. Oaksford,  
18 B. Schwartz, K. E. Stanovich, and P. M. Todd. Mind, rationality, and cognition: An interdisciplinary debate. *Psychonomic*  
19 *Bulletin and Review*, pages 1–34, 2017.
- 20 [10] B. De Finetti. La prévision: ses lois logiques, ses sources subjectives. In *Annales de l’institut Henri Poincaré*, volume 7,  
21 pages 1–68, 1937.
- 22 [11] B. de Finetti. *Theory of Probability*. Wiley, New York, 1975.
- 23 [12] V. De Gardelle, F. Le Corre, and P. Mamassian. Confidence as a common currency between vision and audition. *PLoS*  
24 *ONE*, 11(1), 2016.
- 25 [13] V. de Gardelle and P. Mamassian. Does Confidence Use a Common Currency Across Two Visual Tasks? *Psychological*  
26 *Science*, 25(6):1286–1288, 2014.
- 27 [14] J. Drugowitsch, R. R. Moreno-Bote, and A. Pouget. Relation between Belief and Performance in Perceptual Decision  
28 Making. *PLoS ONE*, 9(5):e96511, 2014.
- 29 [15] S. M. Fleming and H. C. Lau. How to measure metacognition. *Frontiers in human neuroscience*, 8(443):1–9, 2014.
- 30 [16] S. J. Gershman, E. J. Horvitz, and J. B. Tenenbaum. Computational rationality: A converging paradigm for intelligence  
31 in brains, minds, and machines. *Science*, 349(6245):273–278, 2015.
- 32 [17] G. Gigerenzer and W. Gaissmaier. Heuristic decision making. *Annual Review of Psychology*, 62:451–482, 2011.
- 33 [18] G. Gigerenzer, P. M. Todd, and ABC Research Group. *Simple Heuristics That Make Us Smart*. Oxford University Press,  
34 2000.
- 35 [19] T. L. Griffiths, N. Chater, D. Norris, and A. Pouget. How the Bayesians got their beliefs (and what those beliefs actually  
36 are): comment on Bowers and Davis (2012). *Psychological Bulletin*, 138(3):415–422, 2012.
- 37 [20] R. J. Herrnstein and D. Prelec. Melioration: A theory of distributed choice. *The Journal of Economic Perspectives*,  
38 5(3):137–156, 1991.
- 39 [21] Q. J. Huys, T. V. Maia, and M. J. Frank. Computational psychiatry as a bridge from neuroscience to clinical applications.  
40 *Nature Neuroscience*, 19(3):404–413, 2016.
- 41 [22] C. Kaernbach. Simple adaptive testing with the weighted up-down method. *Perception & Psychophysics*, 49(3):227–229,  
42 may 1991.
- 43 [23] A. Kepecs and Z. F. Mainen. A computational framework for the study of confidence in humans and animals. *Philosophical*  
44 *Transactions of the Royal Society B*, 367(1594):1322–1337, 2012.
- 45 [24] A. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, Berlin, 1933.
- 46 [25] J. Kwisthout, T. Wareham, and I. Van Rooij. Bayesian Intractability Is Not an Ailment That Approximation Can Cure.  
47 *Cognitive Science*, 35(5):779–784, 2011.
- 48 [26] S. Laquitaine and J. L. Gardner. A Switching Observer for Human Perceptual Estimation. *Neuron*, 97(2):462–474.e6, jan  
49 2018.
- 50 [27] L. T. Maloney and P. Mamassian. Bayesian decision theory as a model of human visual perception: Testing Bayesian  
51 transfer. *Visual Neuroscience*, 26(01):147, 2009.
- 52 [28] P. Mamassian. Visual Confidence. *Annual Review of Vision Science*, 2(1):annurev-vision-111815-114630, 2016.

- 1 [29] G. F. Marcus and E. Davis. How Robust Are Probabilistic Models of Higher-Level Cognition? *Psychological Science*,  
2 24(12):2351–2360, 2013.
- 3 [30] F. Meyniel and S. Dehaene. Brain networks for confidence weighting and hierarchical inference during probabilistic learning.  
4 *Proceedings of the National Academy of Sciences*, 114(19):E3859–E3868, 2017.
- 5 [31] J. C. Nash. On Best Practice Optimization Methods in {R}. *Journal of Statistical Software*, 60(2):1–14, 2014.
- 6 [32] J. C. Nash and R. Varadhan. Unifying Optimization Algorithms to Aid Software System Users: {optimx} for {R}. *Journal*  
7 *of Statistical Software*, 43(9):1–14, 2011.
- 8 [33] D. G. Pelli. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial vision*,  
9 10(4):437–442, jan 1997.
- 10 [34] M. A. Peters, J. Fesi, N. Amendi, J. D. Knotts, H. Lau, and T. Ro. Transcranial magnetic stimulation to visual cortex  
11 induces suboptimal introspection. *Cortex*, 93:119–132, aug 2017.
- 12 [35] M. J. D. Powell. The BOBYQA algorithm for bound constrained optimization without derivatives. Technical report,  
13 Centre for Mathematical Sciences, University of Cambridge, Cambridge, UK, 2009.
- 14 [36] R Core Team. R: A language and environment for statistical computing, 2015.
- 15 [37] D. Rahnev and R. N. Denison. Suboptimality in Perceptual Decision Making. *Behavioral and Brain Sciences*, pages 1–107,  
16 feb 2018.
- 17 [38] A. N. Sanborn and N. Chater. Bayesian Brains without Probabilities. *Trends in Cognitive Sciences*, 20(12):883–893, dec  
18 2016.
- 19 [39] L. J. Savage. *The foundations of statistics*. John Wiley & Sons, Inc., 1954.
- 20 [40] L. J. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*,  
21 66(336):783–801, 1971.
- 22 [41] H. Shteingart and Y. Loewenstein. The Effect of Sample Size and Cognitive Strategy on Probability Estimation. *Decision*,  
23 2(2):107–117, 2015.
- 24 [42] H. A. Simon. A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1):99, feb 1955.
- 25 [43] E. Vul, N. Goodman, T. L. Griffiths, and J. B. Tenenbaum. One and done? Optimal decisions from very few samples.  
26 *Cognitive Science*, 38(4):599–637, 2014.
- 27 [44] H. Zhang, N. D. Daw, and L. T. Maloney. Human representation of visuo-motor uncertainty as mixtures of orthogonal  
28 basis distributions. *Nature Neuroscience*, 18(8):1152–1158, 2015.
- 29 [45] H. Zhang and L. T. Maloney. Ubiquitous log odds: A common representation of probability and frequency distortion in  
30 perception, action, and cognition. *Frontiers in Neuroscience*, 6(JAN):1–14, 2012.

# 1 Supplemental information

## 2 Material and methods

### 3 Apparatus

4 Participants sat in a quiet, dimly lit room, with the head positioned on a chin rest at a distance of 60 cm from the display  
5 screen, a gamma-linearized Mitsubishi Diamond Plus 230SB CRT monitor (screen resolution 1600x1200, vertical refresh rate  
6 85 Hz). Stimuli were generated by a computer running Matlab (Mathworks) with the Psychophysics Toolbox [7,33].

### 7 Stimuli and procedure

8 All experiments were run with the same visual display, consisting of a central fixation point and four placeholders, continu-  
9 ously visible on a uniform gray background (luminance  $\approx 13.6$  cd/m<sup>2</sup>). The four placeholders were circles measuring 2.8 dva  
10 (degrees of visual angle) in diameter, whose centers were placed at 1.8 dva from the fixation point. Two placeholders were  
11 grey ( $\approx 15.5$  cd/m<sup>2</sup>), and were placed above the horizontal mid-line; the other two were placed below the mid-line and were  
12 colored in red the one on the left, and in green the one on the right (their luminance was matched with the grey placeholders,  
13  $\approx 15.5$  cd/m<sup>2</sup>).

14 In the orientation (O-O) task the stimuli were two Gabor gratings (sinusoidal luminance modulation presented within a Gaussian  
15 contrast envelope) of different orientations, presented for 200 ms. The spatial frequency of the Gabors was set at 1.5 cycles/dva,  
16 the phase was drawn randomly, and the standard deviation of the Gaussian envelope was 0.7 dva. The Gabor displayed in the  
17 left placeholder was always tilted to the left, and the one appearing in the right placeholder was always tilted to the right. The  
18 task of the participants was to indicate which Gabor was more tilted from the vertical; the less tilted of the two Gabors was  
19 always tilted by 15°; the minimum difference was 0.1°.

20 In the duration (D-D) task the two stimuli consisted of white Gaussian blobs (standard deviation 0.65 dva), presented sequen-  
21 tially in the two placeholders (left/right). The order of presentation (left/right stimulus first) was balanced with respect to  
22 the longer/shorter duration of presentation. Participants were asked to indicate the location (left/right) of the longer duration  
23 blob. The shorter duration was always set to 600 ms, and the difference between shorter and longer durations was discretized  
24 in bins determined by the vertical refresh of the monitor ( $\approx 12$  ms). The minimum duration difference was one single monitor  
25 refresh interval.

### 26 Pre-test JND measurement

27 Before the orientation and duration task, we measured individual JNDs using a weighted up-down staircase procedure [22].  
28 The purpose of this pre-test was to quickly obtain a measure of the JND in order to adapt the range of stimuli in the main  
29 experiment to individual sensitivities. The staircase procedure continued until 30 reversals were counted. The initial step size  
30 (the size of the decrease/increase of the difference between the two stimuli) was 2° in the orientation task and 4 refresh intervals  
31 in the duration task ( $\approx 50$ ms), and was diminished to 0.5° and 1 refresh after the second reversal. Stimuli in these pre-test  
32 measurements were presented only in the top placeholders.

### 33 Participants

34 5 subjects (2 female; mean age 30.8, standard deviation 3.1) participated in both conditions of the random-pairs experiment  
35 (D-D, and O-O; see Main text, Material and methods). 9 participants (4 female, 2 authors; mean age 33.9, standard deviation  
36 9.6) participated in the 3 conditions of the correlated-pairs experiment (D-D, O-O, D-O). All participants (except the author)  
37 were naïve to the specific purpose of the experiment. All conditions were performed in separate session on different days. The  
38 order of the D-D and O-O conditions was counterbalanced across subjects, while the D-O condition was always performed in  
39 the last session. All participants had normal or corrected-to-normal vision and gave their informed consent to perform the  
40 experiments.

## 41 Computational models

### 42 Bayesian observer

43 To model the performance in our task, we considered that the observer to make a decision estimates the difference in the  
44 intensity between the left and right stimuli,  $s_i = s_i^{(\text{right})} - s_i^{(\text{left})}$ , where the subscript  $i = 1, 2$  indicates the decision at hand  
45 (first and second) and  $s_i^{(\text{right})}$  and  $s_i^{(\text{left})}$  denotes the stimuli (unsigned deviations from vertical of the two Gabor gratings in the  
46 orientation task, or durations of the two blobs in the duration task). Hence the task amounts to deciding whether  $s_i$  is greater  
47 or less than 0. We assumed that the observer has only access to a corrupted version of  $s_i$ ,  $r_i = s_i + \eta$ , where  $\eta$  is Gaussian  
48 noise with variance  $\sigma^2$ . The ideal Bayesian observer has full knowledge of the statistics of the internal noise, and to make a  
49 decision computes a posterior probability over the variable  $s$ . Since the internal noise is assumed to be Gaussian, the likelihood  
50 function giving the probability of observing  $r_i$  given  $s_i$  is  $p(r_i | s_i) = \mathcal{N}(s_i, \sigma^2)$ , where  $\mathcal{N}(s_i, \sigma^2)$  is a normal distribution with  
51 mean  $s_i$  and variance  $\sigma^2$ .

In first decision, the prior probability of  $s_1$  is uniform with the range  $(-R, R)$ , that is  $p(s_1) = 1/(2R)$  if  $-R \leq s \leq R$  and 0 otherwise. Combining the prior and the likelihood function, the unconditioned probability of observing  $r_1$  can be expressed as:

$$p(r) = \int_{-R}^R p(r_1|s_1) p(s_1) ds_1 \quad (1)$$

$$= \frac{1}{4R} \left[ \operatorname{erf} \left( \frac{R+r_1}{\sigma\sqrt{2}} \right) + \operatorname{erf} \left( \frac{R-r_1}{\sigma\sqrt{2}} \right) \right]$$

1 The posterior probability of  $s_1$  after having observed  $r_1$  is obtained applying Bayes rule:

$$p(s_1 | r_1) = \frac{p(r_1 | s_1) p(s_1)}{p(r_1)} \quad (2)$$

2 Finally, the decision variable  $c_1^+(r_1)$ , corresponding to the probability that the  $s_1$  was greater than 0 is given by

$$c_1^+(r_1) = \int_0^R p(s_1|r_1) ds_1 = \frac{\operatorname{erf} \left( \frac{r_1}{\sigma\sqrt{2}} \right) + \operatorname{erf} \left( \frac{R-r_1}{\sigma\sqrt{2}} \right)}{\operatorname{erf} \left( \frac{R+r_1}{\sigma\sqrt{2}} \right) + \operatorname{erf} \left( \frac{R-r_1}{\sigma\sqrt{2}} \right)} \quad (3)$$

3 When  $c_1^+(r_1) \geq 1/2$  the observer chooses 'right' (i.e. he decides to report that  $s_1$  was positive, an outcome hereafter indicated  
4 with the notation  $d_1^+$ ) otherwise he chooses 'left' (i.e. he reports that  $s_1$  was negative, notated as  $d_1^-$ ). The posterior probability  
5 of being correct in the first decision, that is the confidence of the ideal observer, is given by  $c_1 = \max [c_1^+(r_1), 1 - c_1^+(r_1)]$ .  
6 The ideal observer would use the probability  $c_1$  to adjust prior expectations for the second decision, specifically by assigning a  
7 prior probability equal to  $c_1$  to the possibility that  $s_2$  will be drawn from the positive interval:

$$p(s_2) = \begin{cases} \frac{c_1}{R}, & \text{if } 0 < s_2 \leq R \\ \frac{1-c_1}{R}, & \text{if } -R \leq s_2 < 0 \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

8 By applying the same calculation as above with the updated prior  $p(s_2)$ , one obtains the decision variable  $c_2^+(r_2)$  for the second  
9 decision:

$$c_2^+(r_2) = \frac{c_1 \left[ \operatorname{erf} \left( \frac{r_2}{\sigma\sqrt{2}} \right) + \operatorname{erf} \left( \frac{R-r_2}{\sigma\sqrt{2}} \right) \right]}{(1-c_1) \operatorname{erf} \left( \frac{R+r_2}{\sigma\sqrt{2}} \right) + c_1 \operatorname{erf} \left( \frac{R-r_2}{\sigma\sqrt{2}} \right) + (2c_1-1) \operatorname{erf} \left( \frac{r_2}{\sigma\sqrt{2}} \right)} \quad (5)$$

10 This equation reduces to the one for the first decision variable when  $c_1 = 1/2$ , as it should. The range  $R$  on which  $s_1$  and  $s_2$   
11 takes values is immaterial. It is possible to simplify Eqs (3) and (5) by taking the limit  $R \rightarrow \infty$ . In this limit one obtains:

$$c_1^+(r_1) = \frac{1}{2} \left[ 1 + \operatorname{erf} \left( \frac{r_1}{\sigma\sqrt{2}} \right) \right] \quad (6)$$

12

$$c_2^+(r_2) = \frac{c_1 \left[ 1 + \operatorname{erf} \left( \frac{r_2}{\sigma\sqrt{2}} \right) \right]}{1 + (2c_1-1) \operatorname{erf} \left( \frac{r_2}{\sigma\sqrt{2}} \right)} \quad (7)$$

13 The decision rules described by these equations amount to comparing the internal signal  $r_i$  to a criterion  $\theta$ , and decide accordingly  
14 (i.e., if  $r_i \geq \theta$ , chose  $d_2^+$ , otherwise chose  $d_2^-$ ). The criterion for the first decision,  $\theta_1$ , can be expressed as:

$$c_1^+(\theta_1) = \frac{1}{2} \Rightarrow \operatorname{erf} \left( \frac{\theta_1}{\sigma\sqrt{2}} \right) = 0 \quad (8)$$

15 which is satisfied for  $\theta_1 = 0$ . Similarly, the criterion for the second decision,  $\theta_2$ , can be expressed as:

$$c_2^+(\theta_2) = \frac{1}{2} \Rightarrow \operatorname{erf} \left( \frac{\theta_2}{\sigma\sqrt{2}} \right) = 1 - 2c_1 \quad (9)$$

16 which indicates that  $\theta_2$  is a function of  $c_1$  (see Main text, Figure 2B).  
17

The likelihood of the Bayesian observer choosing  $d_2^+$  (after having chosen  $d_1^+$  in the first decision) can be expressed as:

$$p(d_2^+ | s_2, d_1^+, s_1) = \int_0^{+\infty} p(d_2^+ | \theta_2 = -r_1, s_2) p(r_1 | d_1^+, s_1) dr_1$$

$$= \frac{1}{2} + \frac{1}{2p(d_1^+ | s_1)} \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \int_0^{+\infty} \operatorname{erf} \left( \frac{r_1 + s_2}{\sigma\sqrt{2}} \right) \exp \left[ -\left( \frac{r_1 - s_1}{\sigma\sqrt{2}} \right)^2 \right] dr_1 \quad (10)$$

When the observer has chosen  $d_1^-$  in the first decision instead the likelihood takes the form:

$$p(d_1^+ | s_2, d_1^-, s_1) = \int_{-\infty}^0 p(d_2^+ | \theta_2 = r_1, s_2) p(r_1 | d_1^-, s_1) dr_1$$

$$= \frac{1}{2} - \frac{1}{2p(d_1^- | s_1)} \cdot \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^0 \operatorname{erf} \left( \frac{r_1 - s_2}{\sigma\sqrt{2}} \right) \exp \left[ -\left( \frac{r_1 - s_1}{\sigma\sqrt{2}} \right)^2 \right] dr_1 \quad (11)$$

18 The integrals in the two equations 10 and 11 do not have closed form solutions, so they were evaluated numerically using the  
19 adaptive quadrature algorithm as implemented in the function `integrate()` in **R** [36].

## 1 *Biased-Bayesian observer*

2 Similarly to the ideal Bayesian observer presented in the previous section, the biased-Bayesian observer makes decisions based  
3 on an internal signal corrupted by Gaussian noise with variance  $\sigma^2$ . However he does not know exactly the value of  $\sigma$  so when  
4 asked to assess his confidence he uses  $\hat{\sigma} = m\sigma$  (with  $m > 0$ ). Therefore, after having chosen  $d_1^+$  in the first decision, his  
5 confidence can be expressed as:

$$c_1 = \frac{1}{2} \left[ 1 + \operatorname{erf} \left( \frac{r_1}{m\sigma\sqrt{2}} \right) \right] \quad (12)$$

6 The adjusted criterion for the second decision is given by the same relationship that was found for the Bayesian observer

$$\operatorname{erf} \left( \frac{\theta_2}{\sigma\sqrt{2}} \right) = 1 - 2c_1 \quad (13)$$

7 However in this case solving for  $\theta_2$  gives

$$\theta_2 = -\frac{r_1}{m} \quad (14)$$

8 This means that when  $m > 1$  the internal noise  $\sigma$  is overestimated ( $\hat{\sigma} > \sigma$ ) and the criterion is shifted away from zero less  
9 than in the optimal Bayesian model, leading to an under-confidence bias (i.e. in the second decisions the observer choses in  
10 agreement with the belief that the first decision was correct more frequently than what expected if performance was optimal).  
11 Conversely the criterion is shifted more than ideal when  $\sigma$  is underestimated ( $m < 1$ ), leading to an opposite over-confidence  
12 bias.

## 13 *Non-Bayesian observer*

14 As an alternative to the optimal Bayesian model we considered a class of models that do not assume any knowledge about the  
15 nature of the internal stochastic process linking the stimulus  $s_i$  with the internal observation  $r_i$ . These non-Bayesian models  
16 perform similarly to the Bayesian model for the first decision, that is when  $r_1 \geq 0$  they chose  $d_1^+$ , and  $d_1^-$  otherwise. However,  
17 they cannot estimate a full probability distribution over the values of  $s_i$ , and therefore can assess confidence only by comparing  
18 the internal response  $r_i$  (which can be described as a point-estimate) to a set of one or more fixed criteria (Main text, Figure 2).  
19 In the case of a single confidence criterion, the non-Bayesian observer is confident in the response when the internal signal exceed  
20 the confidence criterion, and uncertain otherwise. When confident about the first response, he shifts the decision criterion for  
21 the second decision by a fixed amount (thereby increasing the probability of choosing  $d_2^+$ ). If only one criterion is used, then  
22 the model has 2 discrete confidence levels (e.g., uncertain vs confident). In such a model, if the confidence criterion is  $w_1$ , the  
23 probability of the observer being confident about his first decision, after having responded  $d_1^+$ , can be expressed as:

$$p(\text{confident} \mid s_1, d_1^+) = \frac{1 - \Phi \left( \frac{s_1 - w_1}{\sigma} \right)}{1 - \Phi \left( \frac{-w_1}{\sigma} \right)} \quad (15)$$

24 Note that this probability do not denote the confidence of the observer about his choice, which instead is assumed here to be a  
25 discrete binary state. Applying the law of total probability, the probability of the observer reporting  $d_2^+$  in the second decision  
26 can be expressed as:

$$p(d_2^+ \mid s_2) = p(\text{confident} \mid s_1, d_1^+) \Phi \left( \frac{s_2 - \theta_2}{\sigma} \right) + [1 - p(\text{confident} \mid s_1, d_1^+)] \Phi \left( \frac{s_2}{\sigma} \right) \quad (16)$$

27 where  $\theta_2$  is the shift in criterion for the second decision applied by the observer when he is confident in his first decision.  $w_1$   
28 and  $\theta_1$  are free parameters that we fit to the data by maximum likelihood estimation. It is straightforward to extend the  
29 model in order to have more than two confidence levels. In our analysis we considered models with 2, 3, 4 discrete levels of  
30 confidence, which had 2,4 and 6 free parameters, respectively. The parameters  $w_1, w_2, w_3$  and  $\theta_2, \theta_3, \theta_4$  were constrained so  
31 that  $0 \leq w_1 \leq w_2 \leq w_3$  and  $0 \geq \theta_2 \geq \theta_3 \geq \theta_4$ .

## 32 *Model fitting*

33 In the first decision  $s_1$  is uniformly chosen from an interval centered around 0 (i.e.,  $s_1$  will be above 0 with probability 1/2),  
34 and there is no prior information about the sign of  $s_1$ . Therefore, for all the models presented in the previous sections, the  
35 probability  $p(d_1^+ \mid s_1)$  of the observer choosing  $d_1^+$  can be expressed as:

$$p(d_1^+ \mid s_1) = \frac{1}{2} \left[ 1 + \operatorname{erf} \left( \frac{s_1}{\sigma\sqrt{2}} \right) \right] = \Phi \left( \frac{s_1}{\sigma} \right) \quad (17)$$

36 where  $\Phi$  is the cumulative distribution function of the standard normal distribution. We fitted this function by maximum  
37 likelihood to estimate  $\hat{\sigma}$ , for each participant and task (duration and orientation), taking into account only the first decisions. We  
38 estimated also other psychometric models, that extended this simple model to account for the possibility that the observer were  
39 biased,  $p(d_1^+ \mid s_1) = \Phi \left( \frac{s_1 - b}{\sigma} \right)$  (where  $b$  indicate the bias), or made stimulus-independent errors (e.g. attention lapses) with non-  
40 zero probability,  $p(d_1^+ \mid s_1) = \lambda + (1 - 2\lambda)\Phi \left( \frac{s_1}{\sigma} \right)$  (where  $\lambda$  is the lapse probability); or both  $p(d_1^+ \mid s_1) = \lambda + (1 - 2\lambda)\Phi \left( \frac{s_1 - b}{\sigma} \right)$ .  
41 To avoid committing to any of these models, for each participant we averaged the estimates of  $\hat{\sigma}$  according to the Akaike  
42 weights [8] of each psychometric model, and used the model-averaged estimate for transforming the stimuli from physical units

1 to units of internal noise. The same was done for the estimates of the bias term,  $\hat{b}$ , which was then subtracted from the stimuli  
 2 value to take into account subjective decision biases in the computation of models' likelihoods. The remaining free parameters  
 3 for the biased-Bayesian and the non-Bayesian observer were estimated by numerically maximizing the (log) likelihood of the  
 4 data (the observed patterns of second decisions). Maximum likelihood fits were obtained via the BOBYQA algorithm [35] as  
 5 implemented in the `optimx` package [31, 32] in **R** [36].

## 6 Sampling-based approximation of Bayesian observer

7 An interesting alternative to the models presented in the previous sections is represented by models where the observer does not  
 8 have access to the full probability distribution of his internal signals, but bases his decision on a limited number of samples. In  
 9 these models the posterior probability that the choice is correct (the confidence) is approximated on the basis of a fixed number  
 10  $n$  of samples  $x_1, x_2, \dots, x_n$  drawn from the posterior distribution  $p(s_i | r_i)$ . The performance of these sampling-based models  
 11 will approach the optimal Bayesian model as  $n \rightarrow \infty$ , however they are expected to display systematic biases and deviations  
 12 from the optimal model for small number of samples [38]. Here we show, and confirm by simulation, that the sampling-based  
 13 approximation of the Bayesian observer will display a systematic over-confidence bias, that is in the opposite direction with  
 14 respect to what we found for most of our subjects, and is thus inconsistent with our behavioral results.

### 15 Sampling bias and over-confidence

16 It has been demonstrated that when a probability  $p$  is estimated from a small sample as the empirical frequency of successes  $k$   
 17 out of  $n$  random trials,  $\hat{p} = k/n$ , the probability of overestimation, that is when  $\hat{p} > p$ , or underestimation,  $\hat{p} < p$ , depends in  
 18 a complex way on both the probability  $p$  and the sample size  $n$  [41]. This is however for estimating a single fixed probability  
 19  $p$ . What would be instead the expected bias, over many repeated estimations, when the probability  $p$  varies randomly within a  
 20 given range? In our experiments confidence is the posterior probability that a binary choice is correct, and as such it varies from  
 21 complete uncertainty,  $p = 0.5$ , to complete certainty,  $p = 1$ , hence  $p \in [0.5, 1]$ . We show here that when a set of probabilities  
 22  $p_1, \dots, p_m$  uniformly distributed in the interval  $[0.5, 1]$  is estimated using a limited number of samples  $n$ , the predominant bias  
 23 is one of over-estimation.

24 For a given  $n$  and  $p$  the probabilities of over- and under- estimation can be expressed as:

$$p(\hat{p} > p | n, p) = \sum_{k=\lfloor np \rfloor + 1}^n \binom{n}{k} p^k (1-p)^{n-k} \quad (18)$$

25

$$p(\hat{p} < p | n, p) = \sum_{k=1}^{\lceil np \rceil} \binom{n}{k} p^k (1-p)^{n-k} \quad (19)$$

26 Where  $\lceil \cdot \rceil$  and  $\lfloor \cdot \rfloor$  are the ceiling and the floor operators, i.e. functions that map a number to the smallest following integer or  
 27 the largest previous integer, respectively. Following Shteingart and Loewenstein [41] we consider the difference between these  
 28 two, denoted as *probability estimation bias*, which takes the form:

$$\Delta p = p(\hat{p} > p | n, p) - p(\hat{p} < p | n, p) \quad (20)$$

29 When  $\Delta p$  is positive, it indicates that the probability  $p$  is more likely to be overestimated than underestimated, and viceversa  
 30 for negative values. Assuming that all values of  $p$  in the interval are equally likely, the expected bias can be computed by  
 31 integrating  $\Delta p$  over the range of  $p$  (that is  $[0.5, 1]$ ):

$$\mathbf{E}[\Delta p] = \int_{0.5}^1 \Delta p \, dp \quad (21)$$

32 A positive value of the expected probability estimation bias (that is  $\mathbf{E}[\Delta p] > 0$ ) indicates that, on average, the probabilities in  
 33 this interval are more frequently overestimated rather than underestimated. This integral can be evaluated numerically, and in  
 34 Figure S2A we plotted the expected probability estimation bias as a function of the number of samples, for two different ranges.  
 35 When  $p$  varies within the range of confidence,  $[0.5, 1]$ , the value of the probability estimation bias is always positive, although  
 36 modulated by the number of samples, indicating that in the range  $[0.5, 1]$  overestimation is more likely than underestimation.

### 37 Fixed- $n$ Bayesian sampler

38 Here we confirm by simulation the intuition presented in the previous paragraph. We consider a fixed- $n$  policy, where the  
 39 observer draws a fixed number of samples for each decision. Although alternative decision policies are possible (such as an  
 40 accumulator policy, where the decision is taken after a minimum number of samples is accumulated in favor of one of the  
 41 options), these have been shown elsewhere to result in very similar performances as the fixed- $n$  policy [43].

42 We start by providing the mathematical details of the model. Similarly to the previous cases, we assume that the observer  
 43 has only access to  $r_1$ , a corrupted version of the stimulus  $s$ ,  $r_1 = s_1 + \eta$ , where  $\eta$  is Gaussian noise with variance  $\sigma^2$ . In the  
 44 first decision the prior is flat and, taking the limit of the stimuli range  $R \rightarrow \infty$ , the posterior distribution  $p(s|r)$  results in a  
 45 Gaussian distribution centered on the internal observation  $r_1$ . The probability that a sample from this distribution is above 0  
 46 (the criterion for the first decision) can be computed as:

$$p(x_n > 0 | r_1) = \frac{1}{2} \left[ 1 + \operatorname{erf} \left( \frac{r_1}{\sigma\sqrt{2}} \right) \right] = \Phi \left( \frac{r_1}{\sigma} \right) \quad (22)$$



1 The conditional probability that the observer chooses  $d_1^+$  given  $r_1$  is obtained by summing the probability of all the set of  
2 samples with at least  $\lceil \frac{n}{2} \rceil$  samples above 0 and can be expressed as:

$$p(d_1^+ | r_1) = \sum_{k=\lceil \frac{n}{2} \rceil}^n \binom{n}{k} \Phi\left(\frac{r_1}{\sigma}\right)^k \left[1 - \Phi\left(\frac{r_1}{\sigma}\right)\right]^{n-k} \quad (23)$$

3 where  $p(x_i > 0|r)$  is the probability that a single sample  $x_i$  is above 0. In other words the observer chooses  $d_1^+$  when the majority  
4 of samples is above 0. The observer's confidence in his decision,  $c_1 = \max[c_1^+, 1 - c_1^+]$ , is given by the proportion of samples  
5 in favor of the choice made, which can be expressed as:

$$c^+ = \frac{1}{n} \sum_{i=1}^n 1_{x_i > 0} \quad (24)$$

6 where  $1_{x_i > 0}$  denotes the indicator function ( $1_{x > 0} = 1$  if  $x > 0$ , and 0 otherwise). The probability that the observer chooses  $d_1^+$   
7 given  $s_1$  can be calculated by integrating the probability  $p(d_1^+ | r_1)$  over all possible values of  $r_1$ :

$$p(d_1^+ | s_1) = \sum_{k=\lceil \frac{n}{2} \rceil}^n \binom{n}{k} \int_{-\infty}^{\infty} \Phi\left(\frac{r_1}{\sigma}\right)^k \left[1 - \Phi\left(\frac{r_1}{\sigma}\right)\right]^{n-k} p(r_1 | s_1) dr_1 \quad (25)$$

8 In the second decision the prior probability of the stimulus  $s_2$  is different for stimuli above or below 0. From the point of  
9 view of the observer, the prior probability that the stimulus for the second decision  $s_2$  is above 0 corresponds to the confidence  
10  $c_1$  that the first decision was correct and, conversely,  $p(s_2 < 0) = 1 - c_1$ . Updating the prior probability for the second decision  
11 amounts to a shift in the decision criterion, as demonstrated for the full Bayesian model. The shift in criterion for the second  
12 decision  $\theta_2$  is a function of the confidence in the sampling model as in the full Bayesian model

$$\theta_2 = \sigma\sqrt{2} \operatorname{erf}^{-1}(1 - 2c_1) \quad (26)$$

13 However, and differently from the full Bayesian model, the confidence is limited to a finite set of values determined by the  
14 number of samples  $n$ . The number of possible confidence levels is  $\lceil \frac{n+1}{2} \rceil$ , and each of these corresponds to a value of  $\theta$ . If the  
15 observer's first decision was  $d_1^+$ , then the probability of the confidence level  $c_j$

$$p(c_j | s_1, d_1^+) = \binom{n}{k} \int_{-\infty}^{\infty} \Phi\left(\frac{r_1}{\sigma}\right)^k \left[1 - \Phi\left(\frac{r_1}{\sigma}\right)\right]^{n-k} p(r_1 | s_1) dr_1 \quad (27)$$

16 where  $j \in (1, \dots, \lceil \frac{n+1}{2} \rceil)$  is an index linked to the number of samples above the criterion,  $k$ , according to  $k = \lfloor \frac{n-1}{2} \rfloor + j$ . A level  
17 of confidence  $c_j$  would result in a shift in decision criterion  $\theta_j$ , calculated according to Eq (26). The probability of choosing  
18 (+) in the second decision is

$$p(d_2^+ | s_2, s_1, c_j) = \sum_{k=\lceil \frac{n}{2} \rceil}^n \binom{n}{k} \int_{-\infty}^{\infty} \Phi\left(\frac{r_2 - \theta_j}{\sigma}\right)^k \left[1 - \Phi\left(\frac{r_2 - \theta_j}{\sigma}\right)\right]^{n-k} p(r_2 | s_2) dr_2 \quad (28)$$

19 Taking everything together, the probability of the observer choosing  $d_2^+$  given  $s_2$ ,  $s_1$  and the first decision can be expressed as:

$$p(d_2^+ | s_2, d_1^+, s_1) = \sum_{j=1}^{\lceil \frac{n+1}{2} \rceil} p(c_j | s_1) p(d_2^+ | s_2, d_2^+, s_1, c_j) \quad (29)$$

## 20 Simulation

21 We simulated the model for values of  $n$  ranging from 2 to 9. In order to compare the model with the full Bayesian observer,  
22 the value of  $\sigma$  for each  $n$  are adjusted so as to obtain the same proportion of correct responses in the first decision. For each  
23 value of  $n$  and for 5000 iterations we: (1) generated a random set of stimuli for 500 trials; (2) simulated the Bayesian model on  
24 those trials; (3) estimated  $\sigma$  for the sampling models based on the set of first responses produced by the Bayesian model (this  
25 was done using maximum likelihood estimation and equation 25); (4) simulated the sampling model. The average values of  $\sigma$   
26 obtained are shown in table S1.

27 This approach ensured that all models resulted in similar proportion of first correct decisions, see figure S2B. The proportion  
28 of responses 'right' ( $d_2^+$ ) in the second decision are plotted in Figure S2C. As expected, they show a pattern of marked over-  
29 confidence: all sampling models tended to respond  $d_2^+$  more often than the optimal model, despite similar accuracy in the first  
30 decision. The bias is larger for models with smaller number of samples, and decreases approaching the optimal Bayesian model  
31 as  $n$  increases. Importantly, this bias is incompatible with the observed behavioral data, which showed on average a marked  
32 under-confidence bias (see Main text, Results).

**Table S1:** Estimated values of  $\sigma$  (and their standard errors) that result in similar performance as the Bayesian model with  $\sigma = 1$ .

n. of samples	$\sigma$	se
2	0.74	0.06
3	0.83	0.06
4	0.84	0.06
5	0.88	0.07
6	0.89	0.07
7	0.91	0.07
8	0.91	0.07
9	0.93	0.07

**Table S2:** Maximum likelihood parameter estimates. For each experiment, condition and model the table shows the mean and the standard deviation (within parentheses) of the parameter estimates across participants.

Condition	Experiment	Model	$m$	$w_1$	$\theta_2$	$w_2$	$\theta_3$	$w_3$	$\theta_4$
D-D	correlated-pairs	biased-Bayesian	2.06 (1.58)						
D-O	correlated-pairs	biased-Bayesian	2.13 (3.01)						
O-O	correlated-pairs	biased-Bayesian	3.01 (2.47)						
D-D	random-pairs	biased-Bayesian	1.27 (0.68)						
O-O	random-pairs	biased-Bayesian	5.97 (3.75)						
D-D	correlated-pairs	fixed-bias			-0.67 (0.24)				
D-O	correlated-pairs	fixed-bias			-0.9 (0.94)				
O-O	correlated-pairs	fixed-bias			-0.63 (0.3)				
D-D	random-pairs	fixed-bias			-0.84 (0.46)				
O-O	random-pairs	fixed-bias			-0.23 (0.22)				
D-D	correlated-pairs	discrete (2 lvl)		0.81 (0.66)	-2.11 (1.52)				
D-O	correlated-pairs	discrete (2 lvl)		0.64 (0.45)	-1.55 (1.12)				
O-O	correlated-pairs	discrete (2 lvl)		1.06 (0.97)	-2.34 (1.77)				
D-D	random-pairs	discrete (2 lvl)		0.55 (0.32)	-1.66 (0.86)				
O-O	random-pairs	discrete (2 lvl)		1.63 (0.47)	-2.85 (2.28)				
D-D	correlated-pairs	discrete (3 lvl)		0.58 (0.56)	-0.73 (0.71)	0.56 (0.54)	-2.11 (2.1)		
D-O	correlated-pairs	discrete (3 lvl)		0.43 (0.5)	-0.83 (0.71)	1.04 (0.83)	-1.56 (1.39)		
O-O	correlated-pairs	discrete (3 lvl)		0.61 (0.53)	-1.04 (1.31)	0.92 (1.35)	-1.59 (1.68)		
D-D	random-pairs	discrete (3 lvl)		0.44 (0.4)	-0.76 (0.27)	0.47 (0.73)	-1.04 (0.74)		
O-O	random-pairs	discrete (3 lvl)		1.08 (0.27)	-0.07 (0.16)	0.69 (0.43)	-2.78 (2.27)		
D-D	correlated-pairs	discrete (4 lvl)		0.53 (0.55)	-0.61 (0.51)	0.54 (0.5)	-1.06 (0.9)	0.43 (0.47)	-1.42 (1.53)
D-O	correlated-pairs	discrete (4 lvl)		0.42 (0.51)	-0.84 (0.59)	0.76 (0.8)	-0.76 (1.07)	0.7 (0.67)	-1.08 (1.12)
O-O	correlated-pairs	discrete (4 lvl)		0.94 (0.91)	-0.67 (0.58)	0.8 (1.62)	-1.46 (1.72)	0.44 (0.66)	-1.05 (1.63)
D-D	random-pairs	discrete (4 lvl)		0.43 (0.41)	-0.67 (0.24)	0.4 (0.69)	-0.93 (0.78)	1.3 (1.52)	-0.6 (0.65)
O-O	random-pairs	discrete (4 lvl)		1.31 (0.44)	-0.64 (1.44)	1.11 (0.9)	-1.62 (1.86)	0.39 (0.55)	-1.65 (2.18)

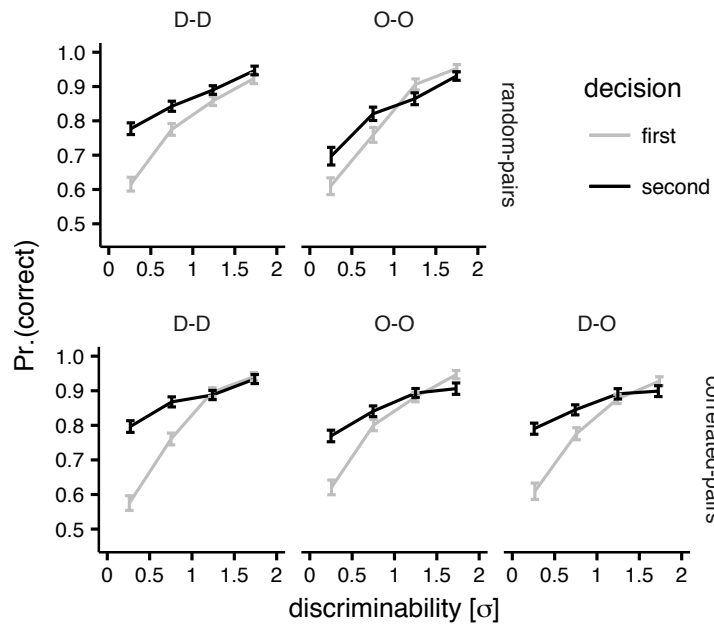


Figure S1: Performance. Same conventions as Figure 1B.

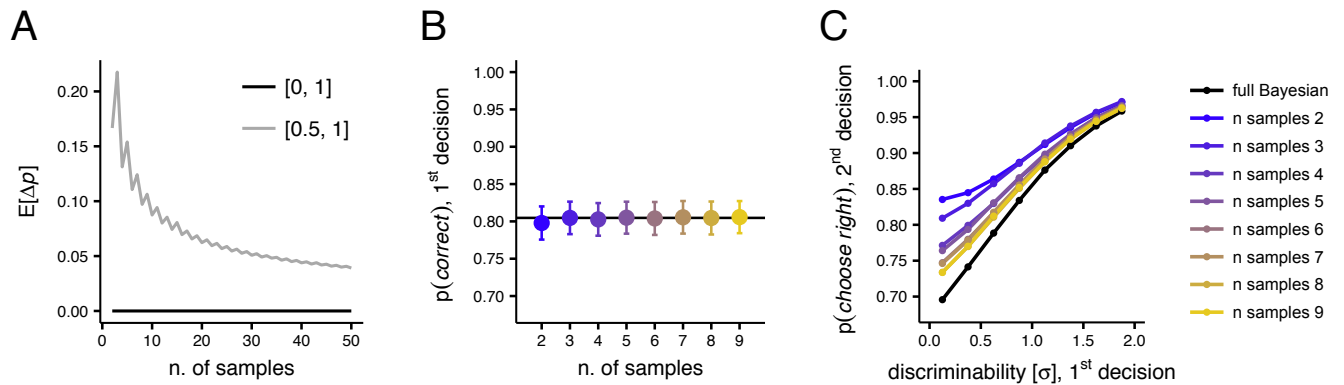
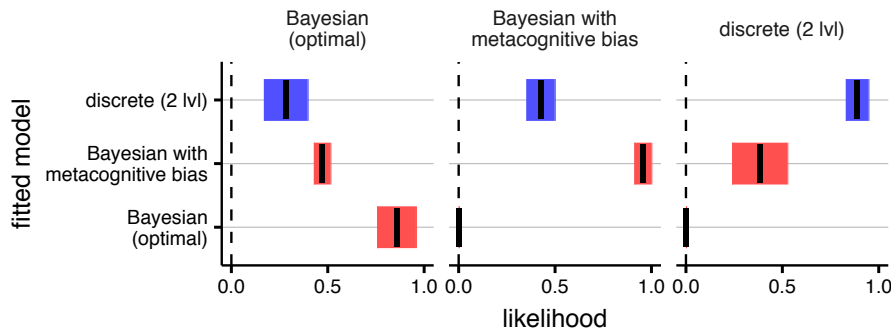



Figure S2: Sampling-based approximation of the Bayesian model. Expected probability estimation bias  $\Delta p$  plotted as a function of the number of samples when the probability  $p$  varies randomly (uniformly) either in the range [0.5, 1], plotted in grey, or in the range [0, 1] (A). It can be seen that when  $p$  varies within the range of confidence, from chance to certainty [0.5, 1] the predominant bias is one of over-estimation (because  $\Delta p$  is always positive, grey line). Only when  $p$  varies over the whole domain of probability, [0, 1], the expected bias is on average zero and over-estimation and under-estimation are equally likely (black line). Proportion of correct first decision in the sampling model as a function of the number of samples  $n$ ; the horizontal black line indicates the average performance of the full Bayesian model; error bars represents SEM (B). Proportions of responses 'right' ( $d_2^+$ ) in the second decision as a function of the difficulty of the first decision, as predicted by the full Bayesian model (black line) and the fixed- $n$  Bayesian sampler model (C).



**Figure S3: Model recovery analysis.** In order to ensure that the models were distinguishable, and to verify that their implementation did not contain errors, we performed a model recovery analysis for our three main computational models: Bayesian (optimal), biased-Bayesian and discrete (2 lvl.). We generated synthetic data (10 simulated observes, for 1000 trials each), with parameters randomly generated from a Gaussian distribution with mean and standard deviations set to the mean and standard deviation of the parameters fitted to our empirical data. Each panel indicate a different generative model, while different lines represent the mean and bootstrapped standard error of the models fit to the synthetic dataset. In each case the model with the highest relative likelihood is the one that generated the data, indicating that the model are correctly recovered and confirm that they are distinguishable.



**Figure S4: Model likelihoods split according to condition and experiment.** Same conventions as Figure 3 in the main text..