Suprathreshold perceptual decisions constrain models of confidence

Shannon M. Locke^{a*}, Michael S. Landy^{b,c}, Pascal Mamassian^a

a) Laboratoire des Systèmes Perceptifs, Département d'Études Cognitives, École Normale Supérieure, PSL University, CNRS, 75005 Paris, France

b) Department of Psychology, New York University, New York, NY, United States

c) Center for Neural Science, New York University, New York, NY, United States

* corresponding author: `shannon.locke@ens.fr`

Abstract

Perceptual confidence is an important internal signal about the certainty of our decisions and there is a substantial debate on how it is computed. We highlight three confidence metric types from the literature: observers either use 1) the full probability distribution to compute probability correct (Probability metrics), 2) point estimates from the perceptual decision process to estimate uncertainty (Evidence-Strength metrics), or 3) heuristic confidence from stimulus-based cues to uncertainty (Heuristic metrics). These metrics are rarely tested against one another, so we examined models of all three types on a suprathreshold spatial discrimination task. Observers were shown a cloud of dots sampled from a dot generating distribution and judged if the mean of the distribution was left or right of centre. In addition to varying the horizontal position of the mean, there were two sensory uncertainty manipulations: the number of dots sampled and the spread of the generating distribution. After every two perceptual decisions, observers made a confidence forced-choice judgement whether they were more confident in the first or second decision. Model results showed that observers were on average best-fit by a Heuristic model that used dot cloud position, spread, and number of dots as cues. However, almost half of the observers were best-fit by an Evidence-Strength model that uses the distance between the discrimination criterion and a point estimate, scaled according to sensory uncertainty, to compute confidence. This signal-to-noise ratio model outperformed the standard unscaled distance from criterion model favoured by many researchers and suggests that this latter simple model may not be suitable for mixed-difficulty designs. An accidental repetition of some sessions also allowed for the measurement of confidence agreement for identical pairs of stimuli. This N-pass analysis revealed that human observers were more consistent than their best-fitting model would predict, indicating there are still aspects of confidence that are not captured by our model. As such, we propose confidence agreement as a useful technique for computational studies of confidence. Taken together, these findings highlight the idiosyncratic nature of confidence computations for complex decision contexts and

the need to consider different potential metrics and transformations in the confidence computation.

Keywords: *confidence, decision-making, metacognition, perception, modelling*

1 Author summary

The feeling of confidence in what we see can influence our future behaviour and learning. Understanding how the brain computes confidence is important goal of researchers. As such, researchers have identified a host of potential models. Yet, they rarely test a wide range of models against each other to find the best at predicting choice behaviour. The study had human participants compare their confidence for pairs of easy perceptual decisions, reporting if they had higher confidence in the first or second decision. We tested twelve models, covering all three types of models proposed in previous studies, finding strong support for two models. The winning Heuristic model combines all three factors affecting choice uncertainty with an idiosyncratic weighting to compute confidence. The other winning model uses a transformation where the strength of the sensory signal is scaled according to sensory uncertainty. This latter model outperformed the frequently-used unscaled version of this model, suggesting that the scaling transformation should be adopted in future studies. We also assessed the agreement of confidence reports in identical decision scenarios. Humans had higher agreement than almost all model predictions. We propose using confidence agreement intentionally as a second performance benchmark of model fit.

2 Introduction

Perceptual confidence is a metacognitive judgement accompanying a perceptual judgement that is thought to reflect the observer’s belief about the quality or correctness of their perceptual decision. For example, a person is likely to have more confidence in their judgement of whether the road ahead bends left or right on a bright sunny day than on a foggy evening because the chances of making a mistake are higher in the latter scenario. Confidence is a decision about a decision, and so it is often referred to as a *Type 2* judgement to contrast it with the *Type 1* perceptual judgement (Clarke et al., 1959). Confidence judgements are ubiquitous in everyday life and are one of the many metacognitive evaluations that guide behaviour and learning (Van den Berg et al., 2016; Meyniel et al., 2015; Frömer et al., 2021).

Numerous models of perceptual confidence have been proposed by researchers, often to capture distinct aspects of decision-making behaviour. For example, some models focus on the relationship

between the speed of the perceptual decision and confidence (Kiani et al., 2014) and others on the degree to which confidence reports can distinguish correct from incorrect decisions (Maniscalco and Lau, 2012). The majority of confidence models are process models that extract the confidence decision variable from either the original Type 1 decision process, or a comparable decision process, such as a partially or completely independent reconstruction (Fleming and Daw, 2017) or further evolved Type 1 decision process (Pleskac and Busemeyer, 2010). The output of these confidence models is a confidence decision variable that can then be mapped to a behaviour (e.g., pressing a button on a keyboard). Confidence decision variables used by researchers in their modelling can be categorised into three main types (Table 1).

Probability confidence metrics are consistent with a common definition that confidence is the probability that the perceptual decision is correct (Pouget et al., 2016). This is also referred to as statistical confidence, or Bayesian confidence, if it is computed from a Bayesian posterior probability (Sanders et al., 2016; Adler and Ma, 2018; Li and Ma, 2020). To compute a probability metric, the observer must consider the probability of all the possible states of the stimulus consistent with their perceptual choice. For example, if they reported a motion direction as clockwise of vertical, this would be the probability of the stimulus being anywhere between $0 - 180^\circ$ clockwise of the decision boundary. Thus, computation of a full probability distribution is necessary for this confidence variable. Bayesian Probability-metric models are typically used as an upper benchmark of confidence performance, as they consider all available information to assess the probability of being correct for reporting their confidence.

In contrast, *Evidence-Strength* confidence metrics are derived from point estimates in the Type 1 decision process and so do not require the representation of full probability distributions. One common Evidence-Strength metric is the *Distance-From-Criterion* (DFC) metric typically used in extended Signal Detection Theory (SDT) models of confidence (Galvin et al., 2003; Fleming and Lau, 2014). In the extended SDT framework, confidence is monotonically related to the unsigned distance between the sensory measurement and the decision criterion. In the case of categorical confidence judgements (e.g., binary report, scale, etc.), additional confidence criteria delineate the mapping between distance and confidence (Maniscalco and Lau, 2012; Fleming and Daw, 2017; Locke et al., 2020a). Most DFC-metric models also allow this sensory measurement to differ from that used for the perceptual decision (Mamassian, 2016), either by additional confidence noise corrupting the measurement (Barrett et al., 2013; Maniscalco and Lau, 2016; Bang et al., 2019) and/or from altering the measurement with parallel decision processes (Fleming and Daw, 2017; Mamassian and de Gardelle, 2021). Variants of the DFC

Table 1: Some implementations of the three most common confidence metric types used for modelling perceptual confidence and example study that used this model.

Probability	Evidence-Strength	Heuristic
<ul style="list-style-type: none"> • Bayesian confidence: posterior probability of being correct (Adler and Ma, 2018) • Log-Probability-Ratio: of posterior probability or likelihood (Adler and Ma, 2018) • Two-best: posterior probability difference of the top two choice alternatives (Li and Ma, 2020) • Entropy: the uncertainty across the posterior distribution of all choice alternatives (Li and Ma, 2020) 	<ul style="list-style-type: none"> • Extended SDT: distance between measurement & criterion (Galvin et al., 2003) • Drift Diffusion: diffusion process state given elapsed time (Kiani et al., 2014) • Balance-of-Evidence: relative final state of separate accumulators (Vickers, 1979) • Supporting-evidence: strength of decision-consistent evidence only (Zylberberg et al., 2012; Maniscalco et al., 2016) 	<ul style="list-style-type: none"> • Stimulus Variability: over- or under-weighting of external noise as a cue (De Gardelle and Mamassian, 2015) • Reaction-time: between stimulus onset and response (Patel et al., 2012) • Other stimulus cues to task difficulty (Mole et al., 2018)

metric have been proposed to allow a point estimate of sensory uncertainty to remap the confidence criteria (Zylberberg et al., 2014; Adler and Ma, 2018; Denison et al., 2018), or allow biases in the distance measures to reflect consideration of only choice-congruent evidence (Maniscalco et al., 2016).

Another common Evidence-Strength metric is the *Accumulator* metric. In accumulation-to-bound models, the Type 1 decision process is represented with one or more decision variables that are updated according to incoming sensory evidence (Gold and Shadlen, 2007). In this manner, accumulation-to-bound models capture both the final choice and the temporal dynamics of that choice (e.g., reaction time). Thus, both the distance of the Type 1 decision variable from the initial decision state and total decision time are factored into the computation of the confidence decision variable. Often a mapping that corresponds to the probability of being correct is selected by experimenters (e.g., Kiani et al., 2014; Zylberberg et al., 2016). That is, the observer uses a point estimate of the decision process and a well-calibrated mapping function to extract the same confidence decision variable they would have got if they computed it from a full probability distribution. Other confidence decision variables of this type consider the relative states of two or more separate accumulators at the time of the decision (Vickers, 1979) or partially interacting accumulators (Zylberberg et al., 2012). Extended versions of the accumulation-to-bound models allow the final state of the decision variable for the perceptual judgement and the confidence judgement to differ based on additional evidence accumulation between the Type 1 and Type 2 reports (Pleskac and Busemeyer, 2010; Balsdon et al., 2020).

The third metric type is the *Heuristic* metric, an ever-expanding category of confidence decision

variables that have been created to capture behaviour beyond that predicted by the favoured standard models mentioned above. While in some cases, the heuristic label has been applied to variants of the Evidence-Strength metric (Aitchison et al., 2015; Maniscalco et al., 2016; Li and Ma, 2020), the majority of “heuristic” confidence models have focused on the use of stimulus cues to infer decision uncertainty. Support for this method of computing confidence comes from a series of studies demonstrating that observers over- or under-weight external noise in the stimulus when reporting confidence (De Gardelle and Mamassian, 2015; Spence et al., 2015; Boldt et al., 2017; Bertana et al., 2021). That is, they display a dissociation between Type 1 and Type 2 performance such that, if perceptual performance is matched for two stimuli with different levels of external noise, confidence is not equated. Other identified heuristic cues are reaction time (Patel et al., 2012) and task-difficulty variables (Mole et al., 2018). Importantly, Heuristic metrics do not require access to the Type 1 decision process, and reflect learned associations between stimulus features or difficulty cues and confidence, which may or may not be well calibrated to reality. However, we note that for all the evidence of heuristic cue use, participants do not always rely on stimulus-uncertainty cues when they are available (e.g., see Barthelmé and Mamassian, 2010; Locke et al., 2020b).

These variations in the confidence computation alter how stimulus strength and sensory uncertainty influence confidence. Taking the earlier example of judging the bend in the road, stimulus strength would be the magnitude of the bend and sensory uncertainty the amount of external noise in the scene. To understand the effect of sensory uncertainty on the confidence decision variable, we will consider two roads, one bending left observed through obscuring fog (i.e., high sensory uncertainty) and the other bending right observed on a sunny day (i.e., low sensory uncertainty). In both cases, the bends in the road are of the same magnitude and the observer correctly identifies the direction of the bends. But how does their sense of confidence compare in these two scenarios?

Figure 1A depicts how an observer using a probability metric would assign higher confidence in the sunny scenario, as the probability of being correct (i.e., the shaded region) is greater due to the difference in spread (i.e., sensory uncertainty). Note that the two probability distributions are centred on the two sensory measurements, which, in this case, are the most likely sensory measurements and so are equidistant from the unbiased decision criterion (i.e., are matched in terms of stimulus strength). The behaviour of the Probability metric for different stimulus strengths also depends on whether any transformations are applied. As shown in Figure 1D, the probability metric approaches a ceiling of 100% probability of being correct if stimulus strength is increased. Even with higher sensory uncertainty, this ceiling is approached relatively quickly with slightly greater stimulus strengths.

However, the probability of being correct for a binary Type 1 decision could also be expressed as a 137
Log-Probability Ratio (LPR), which is the probability of being correct for the selected Type 1 choice 138
divided by the probability of being correct if you had selected the other Type 1 option (see Peirce 139
and Jastrow, 1884; Galvin et al., 2003; Adler and Ma, 2018), In this scenario, the confidence metric 140

Figure 1: Sensory uncertainty and stimulus strength in the confidence computation. A-C) Extracting the confidence decision variable from the Type 1 decision process. Scenarios contrasted: left-bending road on a foggy evening (teal; high uncertainty) and a right-bending road, with the same bend-magnitude, on a sunny day (purple; low uncertainty). The observer correctly identifies the road-bend direction, as shown by their response in quotation marks. Vertical line: decision boundary. A) Probability metric. Curves: normalised likelihood functions of the bend angle, given the sensory measurement (marker; the most likely measurement selected for illustrative purposes). Shaded region: probability of the judgement being correct. The shaded region is greater for the sunny scenario, so the observer has higher confidence in this judgement. B) Scaled-distance metric. A signal-to-noise ratio transformation is applied to the distributions in A (i.e., rescaled to units of standard deviation while the areas under the curve on either side of the Type 1 criterion are preserved). The rescaled sensory measurement in the sunny scenario has a greater Distance-from-Criterion (DFC) and is judged as more confident. C) Influence of a centred prior (see inset). The posterior distributions (continuous curves), computed according to Bayes' Rule, are differentially shifted towards the centre from the likelihood function locations (dashed). The foggy scenario is shifted more because of its higher uncertainty. Consequently, Probability metrics and DFC metrics yield higher confidence for the sunny scenario. D-F) How the confidence metric changes with stimulus strength. The greater the stimulus strength, the larger the confidence metric. However, raw probability values asymptote at 100% confident (D), whereas the confidence decision variable is unbounded if a Log-Probability-Ratio transformation is applied to the probability of being correct (E) or a Unscaled- or Scaled-Distance metric is used (F). Note y -axes have been rescaled for illustrative purposes. G) Heuristic confidence metric computed from estimates of stimulus strength (bend angle) and sensory uncertainty (viewing conditions) without consideration of the Type 1 process. The observer sets the weights on these factors, w_1 and w_2 .

is unbounded, and the difference in the metric between low and high uncertainty increases for larger stimulus strengths (see Figure 1E). So while the relationship of greater confidence for lower uncertainty is typically preserved, the underlying confidence decision variables can differ dramatically depending on transformations in the computation.

An observer who uses the standard DFC Evidence-Strength metric would not behave the same as an observer using the Probability metric. As both measurements have the same DFC, the observer would have the same degree of confidence for the sunny and foggy scenarios. It is only if the observer scales this measurement according to the degree of uncertainty (i.e., computes the signal-to-noise (SNR) ratio) using a point-estimate of sensory uncertainty, would confidence be greater for the sunny scenario (Figure 1B). The effect of this scaling means the sensory measurement is now further from the decision boundary for the the sunny scenario compared to the foggy scenario. A similar reasoning would apply for an accumulation-to-bound Evidence-Strength metric, where the sensory measurement (i.e., final decision state) and decision time (proportional to uncertainty) jointly determine the degree of confidence. The value of a DFC Evidence-Strength metric scales linearly with stimulus strength. In the case of a SNR scaling, the sensory uncertainty affects the slope (Figure 1F).

Figure 1C depicts yet another way sensory uncertainty can influence confidence. If the prior expectation about the stimulus is concentrated at the decision boundary (e.g., the majority of roads are straight or with only a slight bend), the current observation can be biased towards this expected stimulus. This enhances the difference between the sunny and foggy scenarios for both a Probability metric and a DFC Evidence-Strength metric. In fact, even an unscaled DFC Evidence-Strength metric would now reflect greater confidence for the sunny scenario due to the prior biasing effect being stronger for measurement from the foggy scenario. Note that we are defining different states of the Type 1 decision process in Figures 1A-C, with the metric type reflecting different degrees of access or use of the decision-process information by the metacognitive system (i.e., full probability distributions or only point estimates). Thus, there is no conflict when an observer, who uses the prior for their Type 1 decision, only uses a point estimate of a biased sensory measurement to compute an unscaled DFC metric for confidence.

Finally, an observer computing a heuristic confidence metric considers the sunniness or fogginess of the day as well as the angle of the bend as independent inputs without considering an underlying Type 1 process. They then combine these inputs with some weighting scheme of their choice (Figure 1G). If they correctly apply the rule of lower confidence for foggy scenarios, even if it imperfectly captures the relationship between fog and external noise, there will be an effect of sensory uncertainty on

confidence.

Despite the diversity of potential confidence models, only a few studies have directly compared models of different metric types. One approach has been to compare the behavioural signatures of competing models. Studies using this approach have tended to compare a Probability metric with a Heuristic metric, finding support for the probability metric (Barthelmé and Mamassian, 2010; Sanders et al., 2016) or differing support for each metric depending on the perceptual task (Bertana et al., 2021). Formal model comparisons, on the other hand, have been focused on Probability versus Evidence-Strength metrics with mixed results (note that this is according to our metric definitions not those of the authors who used the term “heuristic” more liberally). Aitchison et al. (2015) investigated simultaneous versus sequential Type 1 and Type 2 reports, finding the Bayesian-confidence probability metric best fits sequential reports, but this metric and an Evidence-Strength metric could explain simultaneous reports equally well. Adler and Ma (2018) investigated categorisation behaviour when category means differed but variances were matched versus when category variances differed but means were matched. They found that an Evidence-Strength metric with quadratic sensory-uncertainty-dependent bounds best fits the behaviour in both tasks. Lisi et al. (2021) tested if confidence could be used in a subsequent decision where the perceptual evidence depended on the correctness of previous perceptual choice, and found that the model with a discrete Evidence-Strength metric outperformed the Bayesian Probability-metric model. Finally, Li and Ma (2020) investigated various Probability and Evidence-Strength metrics in the context of a three-alternative forced-choice task. They found a Probability metric that compares the posterior probabilities of the two best alternatives outperformed other metrics. Thus, even with direct comparison of the models, the evidence is mixed for the different metric types and further work is needed to understand the nature of the confidence decision variable. This need was highlighted recently by visual metacognition researchers who stated that determining how confidence is computed with detailed and falsifiable models is important for the field (see Rahnev et al., 2021).

The aim of the present study was to compare the fit of confidence models of all three metric types to the behaviour of humans performing a visual decision-making task. Specifically, observers had to infer if the mean of a dot-generating distribution was left or right of centre (Figure 2A). We employed a mixed-difficulty design, varying the position of the mean (i.e., stimulus strength), as well as the spread of the dot-generating distribution (the *quality* manipulation) and the number of dots drawn (the *quantity* manipulation). Both of these manipulations affected sensory uncertainty. This design allowed us to assess how the observer took sensory uncertainty into account in the computation of confidence.

We used the confidence forced-choice method (Mamassian, 2020; Mamassian and de Gardelle, 2021): after two consecutive perceptual decisions, the observer reports whether they had greater confidence in the first or second decision. This allowed us to investigate suprathreshold perceptual decisions where confidence is typically high without being concerned with ceiling effects (e.g., always reporting high confidence or the highest scale rating). We targeted this difficulty regime because the metrics are particularly divergent for more extreme stimulus strengths (Figure 1D-F). We selected seven base models (twelve models total considering variants in the prior distribution) that captured the diversity of potential confidence computations illustrated in Figure 1 (models described in Table 2). In addition to formal model comparison, we also investigated a new, qualitative, behavioural signature of confidence: *confidence agreement*. Due to a coding error, many sessions were identical repeats, in some cases as many as 5 repeats (Figure 2B). We compared the confidence agreement of the observers to the confidence agreement of model simulations using the best-fitting parameter values as a benchmark of model fit.

To preview the results, we found that the confidence judgements were affected by both the quantity and quality manipulations and all three metric types were supported by at least one observer. The best-fitting model was the Heuristic model, followed closely by a DFC Evidence-Strength model where

Figure 2: Experimental methods. A) Task design. Observers were shown dots drawn from a Gaussian generating distribution with seven possible horizontal spatial offsets between $\pm 4^\circ$ from the screen centre. They judged if the distribution mean was left or right of centre. After each pair of perceptual decisions (Interval 1 and Interval 2), they reported the interval in which they had higher confidence in their decision. There were six levels of stimulus uncertainty, defined by information quantity (number of dots: 2 or 5) and quality (sampling distribution SD: 1.5, 2, or 2.5°), presented in an interleaved design. B) Occurrence of unique stimulus sets (colour coded). Squares: stimulus set in a single session. Sets are ordered by type, not session order, and participants by frequency of stimulus set repeats (5-pass to 1-pass).

Table 2: Summary of the seven base models. Models spanned all three metric types and considered different implementations of confidence noise, prior distributions, and confidence-variable transformations. Standard unbiased Gaussian confidence noise is used unless noted otherwise. There are twelve distinct models for comparison when including the prior variants.

Model	Type	Description
Ideal Conf. Observer	Probability	Compares the $p(\text{correct})$ for each decision, computed from the posterior distribution. Only the centred-prior variant considered for this model.
Basic Probability	Probability	Compares the $p(\text{correct})$, but with early confidence noise (beta, constrained $[0, 1]$) applied before the comparison, interval bias, and two prior variants considered (flat & centred).
Probability Difference	Probability	Compares the $p(\text{correct})$, but with late confidence noise applied after the comparison, interval bias, and two prior variants considered (flat & centred).
Log Probability Ratio	Probability	Compares the $p(\text{correct})$, but with a LPR transformation applied, confidence noise, interval bias, and two prior variants considered (flat & centred).
Unscaled Distance	Evidence Strength	Compares the DFC of point-estimates of the Type 1 decision process, with confidence noise and interval bias. Two prior variants considered (flat & centred).
Scaled Distance	Evidence Strength	Compares the DFC of SNR-scaled point-estimates of the Type 1 decision process, with confidence noise and interval bias. Two prior variants considered (flat & centred).
Heuristic	Heuristic	Compares a weighted sum of the estimated difference in stimulus strength and estimate difference in sensory-uncertainty factors, with late confidence noise and interval bias.

evidence strength was scaled by the sensory uncertainty. Together, our results indicate observers were unlikely to compute full probability distributions for confidence and that they might rely on stimulus-based heuristic cues if available. For confidence agreement, the best-fitting model, on a per-participant basis, almost always underestimated the observers' confidence agreement. This finding suggests there is still room for improvement in the modelling of perceptual confidence.

3 Results

3.1 Confirming that the stimulus manipulations affected confidence

First, we examined whether the confidence reports meaningfully distinguished accuracy in the Type 1 spatial task (Figure 3A). Observers had an overall high level of accuracy of 0.91 ± 0.002 (mean $\pm$ SEM, $\mu = 0^\circ$ trials excluded from calculation). Despite this near-ceiling performance, the interval chosen as more confident was on average more likely to be correct than the declined interval in the pair: 0.95 ± 0.003 versus 0.85 ± 0.004 ($t_{15} = 14.96$, $p < 0.01$). Performance was even better if responses are scored according to the mean of the dots displayed on the screen (i.e., the centroid), removing the

factor of external noise on performance: 0.94 ± 0.003 for all, 0.98 ± 0.003 for chosen, and 0.90 ± 0.004 for declined ($t_{15} = 19.55$, $p < 0.01$). These results indicate that observers made meaningful confidence forced-choice judgements.

The effect of the quantity and quality manipulations on confidence can be seen in the raw response data (Figure 3B). For all trials, the further the dot cloud was from the screen centre, compared to the stimulus of the other interval, the more likely it was to be chosen as more confident. When the two intervals had a different number of dots (right panel), the interval that contained the greater quantity of dots was more likely to be chosen as more confident. Finally, when the two intervals differed in dot spread, the interval that contained the smaller spread (i.e., larger inverse spread) was more likely to be chosen as more confident. In the plots, positive difference values favour Interval 2 as more confident choices (blue) and negative differences Interval 1 choices (red). Note in the right panel, the interval

Figure 3: Confidence manipulation checks. A) Supra-threshold Type 1 task performance. The mean proportion of correct spatial judgements across observers is shown for different trial categories. The left-most bar shows performance for all trials, and the next two bars for trials sorted by confidence, either being in the interval chosen as more confident or declined. B) Raw confidence choices, sorted by stimulus properties across the two intervals of a confidence pair. The colour code represents the proportion of “Interval 2” as more confident choices averaged across observers. Confidence choices are plotted as a function of the difference in distance from centre across intervals ($I_2 - I_1$), and difference in inverse dot spread. Left panel: comparisons where the number of dots was the same in each interval. Right panel: comparisons where the number of dots differed, with stimulus information and confidence selectively flipped so that Interval 2 has more dots for plotting purposes. The distance from centre and dot spread were calculated using the empirical mean and SD of the dots displayed, and binned in the range $\pm 3^\circ$ for plotting. Gold line: the confidence-indifference contour, where the observer is equally likely to report Interval 1 or 2, calculated from the Full model in the nested logistic regression analysis. C) Model comparison for the nested logistic regression analysis. AICc scores are reported relative to the Full model (winner) that contained both quantity and quality predictors. The +Quantity and +Quality models only contained one of these predictors and the Basic model contained neither. The results show that both the quantity and quality manipulations affected confidence. Larger positive scores indicate a worse fit. Error bars: $\pm$ SEM.

with more dots was coded as Interval 2 for plotting purposes.

To quantitatively confirm both the quantity and quality manipulations affected confidence, we performed a nested logistic regression analysis. The full model, which had the difference in distance from the screen centre, inverse dot spread, and number of dots as predictors, outperformed simpler models without the quantity and/or quality predictors (Figure 3C). See the Supplementary Materials for more details on the models and model comparison. The gold confidence-indifference lines in Figure 3B provide a visualisation of the fit of the full model.

Together, these results suggest observers' confidence computations were affected by the stimulus strength as well as both sources of sensory uncertainty in this easy perceptual task. To better understand the computation of the confidence decision variable, we next examined several process models that capture the full decision process (i.e., both perceptual and confidence judgements).

3.2 Type 1 model comparison results

Four Type 1 models were considered. All models had three free parameters, a perceptual decision criterion (k_1), per-dot sensory noise (σ_{dot}), and a lapse rate (λ). The models differed in two ways: 1) whether the posterior distribution was computed with a flat or centred prior, and 2) if the Type 1 decision variable was the probability that the dot-cloud was on the left/right or the point estimate of the posterior's mode relative to the decision criterion (both metrics are always in 100% agreement). The two prior-variants are also often in agreement about the perceptual choice, except for cases close to a biased decision boundary. Thus, it was unlikely to see large differences in the Type 1 model-comparison results. As expected, we found identical fits for the probability and signed-distance metrics (Figure 4A), and near-identical fits for the prior-variants. The purpose of fitting all four models was to ensure that the Type 1 parameters could be fixed accordingly based on the Type-1 responses for the Type 2 model fits.

3.3 Type 2 model comparison results

Twelve Type 2 models were compared, with confidence computed with a Probability metric in seven models, a DFC Evidence-Strength metric in four models, and a Heuristic metric in one model (Figure 4B; models summarised in Table 2 and their equations detailed in Table 3). The Heuristic model had the best AICc score averaged across observers, with 7 of the 16 observers best fit by this model. The Scaled-Distance models were a strong competitor, with 4 observers best fit by the flat-prior variant and 3 by the centred-prior variant. The average AICc scores for these variants were respectively

10.8 $\pm$ 26.9 and 23.1 $\pm$ 20.7 higher than the Heuristic model. The flat-prior variant of the Probability-Difference model best fit 2 observers and had an average relative AICc score of 41.0 $\pm$ 23.8. Overall, the Ideal, Basic-Probability, and Unscaled-Distance models fit poorly. Model results were relatively unchanged with alternative prior distributions and likelihood functions (see Supplementary Materials).

Figure 4: Model fit results ($n = 16$). A) Relative AICc scores for the Type 1 models. Scores were compared to the winning flat-prior variant models. Bars: average relative AICc score. Markers: Individual participant results. Colour: flat-prior (orange) or centred-prior (blue) variant. B) Relative AICc scores for the Type 2 models, with fraction best-fit annotated. Models are grouped by metric type. Note the different y-axis scales between panels A and B. C) Average best-fitting Heuristic-model coefficients across all observers. Note that the position coefficient was always fixed at 1 in the model, but is graphed to illustrate the relative coefficient weights. D) The best-fitting quantity and quality coefficients per observer. Marker colour and type indicates the best-fitting model for that observer. Dashed lines: the position coefficient for comparison. E) Model recovery results. For each model and observer, 10 data sets were simulated using the participant's best-fitting parameters (1920 data sets total) and then fit by each of the 12 models. The best-fitting model was tallied per dataset per simulated model. Dark squares along the downward diagonal indicate high model recovery success. F) The relative AICc scores compared to the simulated model (downward diagonal of 0). Error bars: $\pm$ SEM.

3.4 Examining the Heuristic model fits

We used the best-fitting parameters of the Heuristic model to investigate choice behaviour further. Figure 4C shows the estimated coefficients, which can be compared directly because the predictors were z-scored. Confidence was most strongly determined by the stimulus strength (i.e., dot-cloud position). Of the two sensory-uncertainty manipulations, the quantity of dots was given more weight than the quality of the dots, as can be seen by the coefficients, $\beta_{quantity} = 0.68 \pm 0.10$ versus $\beta_{quality} = 0.38 \pm 0.09$ ($t_{15} = 3.75$, $p < 0.01$) respectively, which are also both significantly different from 0 ($t_{15} = 6.61$, $p < 0.01$ and $t_{15} = 4.34$, $p < 0.01$ respectively). These are contrasted per participant in Figure 4D. There appears to be clustering of coefficient values according to best-fitting model (examined in more detail in the Supplementary Materials), with only Heuristic-best-fit participants giving more weight to dot quantity than stimulus strength. There was also a slight but not significant confidence interval bias to choose Interval 2 as more confident ($\beta_{bias} = k_2 = 0.13 \pm 0.08$; $t_{15} = 1.63$, $p > 0.05$; Figure 4C) and confidence noise was $\sigma_{conf} = 1.20 \pm 0.17$, according to the best-fitting parameters for the Heuristic model across all participants.

3.5 Model recovery analysis

A model recovery analysis further supported our model-comparison results. Data were simulated according to a particular model, with identical experiment structure and parameters consistent with our participants' behaviour. Our simulated data were almost always best-fit by the model that generated them (Figure 4E). This indicates that the models are distinguishable from one another in model comparison. The centred-prior variant of the Basic-Probability model had the lowest recovery success with 78.75% datasets recovered. The Heuristic model was the highest with a 100% recovery rate. On average the recovery rate was $91.72 \pm 6.54\%$ (mean $\pm$ SD).

A comparison of the relative AICc scores in Figure 4F shows that some model-simulation fit pairs are more similar in model-fit quality than others, even though the simulated model was almost always correctly recovered. First, flat- and centred-prior variants tended to have similar AICc scores. A similar pattern emerges for the Scaled-Distance and LPR models. Then there are the unidirectional similarities. Datasets generated by simulating the Ideal-Confidence-Observer model could often be well fit by the centred-prior variants of other models. This was to be expected if the model converges to the Ideal-Confidence-Observer model when confidence noise and interval bias approach 0. But, by penalising model complexity, the Ideal-Confidence-Observer model is often recovered. The other

unidirectional similarity is between the Unscaled-Distance models and the Heuristic model. This is because the Heuristic model with no weight given to the quantity and quality predictors is the likelihood-variant Unscaled-Distance model. However, penalising the extra complexity of the Heuristic model helps to ensure the recovery of the Unscaled-Distance models.

3.6 Confidence agreement

Confidence agreement was quantified by counting the number of the most consistent confidence response on a per-trial basis (e.g, 4 “Interval 1” responses out of 5 passes is 80% agreement). The pattern of confidence agreement averaged across observers shows that 1) observers had high overall levels of confidence agreement, and 2) comparisons close to the confidence-indifference line are less consistent than those far from this line (Figure 5A). We then investigated the predicted confidence agreement according to each of the twelve models, simulated using the best-fitting parameters on a per-participant basis (results of an example participant shown Figure 5B and all participants in the Supplementary Materials). As expected, the Ideal-Confidence-Observer model always had the highest confidence agreement, because there is no confidence noise in this model. The Basic-Probability models then had the second and third highest levels of confidence agreement, followed by the remaining models, which tended to have more similar levels of confidence agreement. The higher confidence

Figure 5: **Confidence agreement results** ($n = 15$). Confidence agreement was calculated per trial as the proportion of the most-selected confidence choice for the participants who did a 3-, 4-, or 5-pass version of the experiment (Figure 2B). A) Heatmaps of average confidence agreement according to the properties of the two stimuli displayed, pooled across observers. Gold: the indifference lines where each interval is equally likely to be selected as more confident according to the preliminary analyses (Figure 3B). B) Comparing the 5-pass confidence agreement of a representative example participant (red; #11) with the predicted confidence agreement of the models. Green: the best-fitting model for this observer (Heuristic model). Black: other models (flat- and centred-prior variants had similar confidence agreement counts so only the flat-prior variant is shown). Grey: the Basic-Probability model with additional late noise (1% SD). Model predictions calculated from 100 simulated datasets using the participant-specific best-fitting parameters. Error bars: ± 2 SD.

agreement of the Basic-Probability models, however, was not robust. As expected, many of the confidence comparisons had values close to 1 (Figure 1D), with the computer simulations detecting very small differences in probability (e.g., 0.985 versus 0.998). A human observer is unlikely capable of such comparisons, and so to get a more realistic prediction of confidence agreement for the Basic-Probability model, we also simulated a version of the model that included 1% SD late noise (grey lines in Figure 5B). The confidence agreement of this model was much closer to the other non-ideal models.

To investigate if the best fitting model captured the confidence agreement of observers, we compared the proportion of trials for the highest possible agreement count. Note that the minimum and maximum agreement depended on whether the participant did a 3-, 4-, or 5-pass version of the task. As the minimum number of passes was 3 (e.g., 2/3 and 3/3 agreement), selecting only a single count value to compare was the only appropriate choice for a statistical comparison. For 14 of the 15 observers, their confidence behaviour was in closer agreement across the repeats than predicted by their best-fitting model. A Wilcoxon signed-rank test confirmed that this difference in confidence agreement is significant ($n = 15$, $z = 3.35$, $p < 0.01$).

4 Discussion

We fit confidence models of all three metric types to confidence forced-choice responses in a suprathreshold spatial-discrimination task. Each metric type was supported by at least one observer, with 44% observers best fit by the Heuristic model (i.e., Heuristic metric), 44% by the Scaled-Distance model (i.e., DFC Evidence-Strength metric with a SNR transformation), and 12% by the Probability-Difference model (i.e., Probability metric with late confidence noise). The Heuristic model was considered the winner overall at the group level. The modelling results suggest there is no universal computational strategy for taking sensory uncertainty into account when judging confidence. All four possibilities (i.e., a probability metric, SNR-scaling, use of a centred-prior, heuristic cue use; see Figure 1) were supported by at least one observer. Taken together, these results suggest that the computation of confidence is highly idiosyncratic in environments where decision uncertainty is influenced by multiple factors, further supporting the hypothesis of a highly individual nature to perceptual confidence (Graziano and Sigman, 2009; Ais et al., 2016; Navajas et al., 2017). However, within the metric types, there was a preference for some computations over others, which has implications for our understanding of the computation of perceptual confidence more generally.

4.1 Evidence-Strength metrics

354

Our results clearly show that for this mixed-difficulty experiment, the Unscaled-Distance model performed significantly worse than the Scaled-Distance model, which took into account sensory uncertainty by applying a SNR transformation to the point estimate used to compute the DFC. In other words, a single point estimate of the decision process was insufficient to capture confidence behaviour and a secondary point estimate of sensory uncertainty is used by the metacognitive system. This finding is particularly relevant for the Extended-SDT framework of perceptual confidence. It is often assumed that the sensory measurement is simply compared to static confidence criteria (Maniscalco and Lau, 2012; Mamassian, 2016), which is no issue due to many SDT tasks being of a fixed difficulty level (note that this is not true of tasks that staircase difficulty, e.g., Fleming et al., 2010; Bang et al., 2019). In mixed-difficulty designs, however, it has been proposed that confidence criteria are to some degree updated, sometimes on a trial-by-trial basis, according to the level of sensory uncertainty (Zylberberg et al., 2014; Adler and Ma, 2018; Rahnev, 2021). The reason an observer would do this is to avoid a confidence paradox of more readily assigning high confidence to stimuli with large sensory noise simply because the noise is likely to carry the measurement far from the perceptual decision criterion. Yet, it appears that human observers do not shift their criteria appropriately to avoid this paradox (Zylberberg et al., 2014; Rahnev, 2021). However, the failure to find appropriate scaling by the participants in these studies may be due to the technique of measuring confidence. Without an incentive structure, confidence ratings are essentially meaningless and there is little motivating accurate shifts, but this is not true for the confidence forced-choice judgements as relative comparisons will always have a sensible interpretation even in the absence of an incentive structure.

374

From a modelling perspective, shifting the criteria to account for sensory uncertainty and the scaled DFC measure have the same effects on the confidence computation. Previous studies likely favoured the shifting-criterion description simply because criterion plasticity is commonly accepted (Rahnev and Denison, 2018), whereas in the current study there were no confidence criteria to shift so only the scaling method was appropriate (see also Mamassian and de Gardelle, 2021). This highlights the need to better understand the metacognitive mechanism for accounting for sensory uncertainty. Our results suggest shifting of confidence criteria is not necessary as observers can scale DFC measures, however the “stickiness” of criteria has been used to explain suboptimal behaviour in decision contexts requiring shifting criteria (Gorea and Sagi, 2000; Rahnev et al., 2011; Zylberberg et al., 2014). It is unclear how these suboptimalities can be explained in a SNR version of Extended-SDT. If researchers

384

wish to apply this form of scaling to an Extended SDT model, we propose using the log-likelihood-ratio representation ($\ln \beta$), which rescales the decision axis by the amount of sensory noise for a given stimulus level (Macmillan and Creelman, 2005), as placing perceptual and confidence criteria in this space removes the need to model multiple shifted criteria.

At a minimum, computing the SNR requires point estimates of the signal strength and uncertainty, but could also be achieved with a full probability distribution representation. After all, Probability metrics could be fairly well approximated with these point estimates. Thus, it appears that the linear scaling of the DFC confidence metrics is a key aspect of the Scaled-Distance model’s success in contrast to the Probability-metric models (Figure 1D-F). Further research is needed to understand if this result is due to the way probabilistic computations are implemented in the metacognitive system (Pouget et al., 2013) versus task-specific influences. In particular, as this task often asked observers to compare two easy perceptual decisions, a scaled DFC allows participants to be certain but not indifferent for the confidence forced-choice judgements, whereas comparisons of the raw probability values leads to quibbling over two values close to 1 (Figure 1D).

4.2 Heuristic metrics

The winning Heuristic model is another model where observers are rarely indifferent for comparisons of easy trials. The model in this task had to consider three stimulus-based cues when assigning confidence: the position of the dot cloud, the number of dots, and the dot spread. The distance of the centroid from the centre was the predictor given the most weight, followed by the number of dots, then dot spread. Heuristic observers tended to give more weight to the number of dots displayed than the other observers, with 5/7 observers giving almost equal or more weight to this predictor than the position of the dot cloud (i.e., stimulus strength). If is unclear why the number of dots was particularly salient to these observers in this task, given the amount of previous evidence suggesting that variability, here dot-cloud spread, is a strong heuristic cue (De Gardelle and Mamassian, 2015; Spence et al., 2015; Boldt et al., 2017; Bertana et al., 2021). However, to our knowledge, no other perceptual confidence task has jointly manipulated both the quantity and quality of sensory information. This suggests that future work interested in investigating heuristic cues in confidence should consider targeting the quantity of information provided.

From a computational perspective, approximately half of all confidence forced-choice comparisons were between intervals that differed in the number of dots and all differed in empirical spread (assuming that was used instead of the unknowable generative spread), so stimulus-based heuristics could have

provided quicker and less computationally complex answers to confidence comparisons in most trials 416
of this task (Gigerenzer and Gaissmaier, 2011). In contrast, a Bayesian confidence observer would 417
often be faced with comparing similar subjective probabilities of being correct that could prolong 418
deliberation on the confidence report. Furthermore, as many perceptual scenarios in daily life are 419
similarly suprathreshold, it may be advantageous more generally to rely on stimulus-based heuristics 420
(or scaled-distance measurements) simply to avoid indifference between confidence-driven choices. For 421
example, when attempting to select the freshest salad mix at the supermarket, you can be assessing 422
two bags that look equally fresh but you might select the bag without spinach because it expires 423
quickly. This heuristic leads to a satisfactory result quickly. In this regard, the confidence forced- 424
choice technique with easy perceptual decisions proved a suitable method for investigating heuristic 425
confidence that could be employed with neural measures to increase our understanding of real-world 426
computations of confidence (Gardner, 2019). 427

4.3 Probability metrics 428

In addition to the Probability-Difference model, the Probability metric category contained the Ideal 429
Confidence Observer, the Basic-Probability model, and the LPR model. The centred-prior variants of 430
these models conform to the definition of “Bayesian Confidence” according to the Bayesian Confidence 431
Hypothesis (Sanders et al., 2016; Adler and Ma, 2018; Li and Ma, 2020), and were not a best-fitting 432
model of any participant. The flat-prior variants of these models also conform to the definition of 433
Bayesian Confidence, as the probability of being correct cannot be computed from an un-normalised 434
likelihood function, and were best fit for two participants. As such, our results do not strongly support 435
the Bayesian Confidence Hypothesis, in agreement with some previous studies (Adler and Ma, 2018; 436
Denison et al., 2018; Lisi et al., 2021) but not others (Aitchison et al., 2015; Li and Ma, 2020). 437
Despite this finding, the results of the current study should not be interpreted as against Bayesian 438
computations in the brain more generally. The averaging of the perceived dot locations that is the basis 439
of all models we tested is optimal Bayesian cue combination (Landy et al., 2011). Additionally, the 440
computation of the posterior mode in the centred-prior variants of the Unscaled- and Scaled-Distance 441
models must also involve a Bayesian computation. This is ambiguous in the case of the flat-prior 442
variants of Unscaled- and Scaled-Distance models, because these models are identical to a model that 443
uses the point estimate from the un-normalised likelihood function. The important distinction made 444
by the Bayesian Confidence Hypothesis is whether the metacognitive system has access to and makes 445
use of the full posterior distribution for reporting confidence and our results are inconclusive in this 446

regard.

A clear finding from the present study is that confidence noise is an important component of confidence models, adding support to the many previous studies that suggest that the perceptual decision variable and the confidence decision variable are not the same due to additional confidence noise (Maniscalco and Lau, 2016; Bang et al., 2019; Shekhar and Rahnev, 2021). This can be seen best in the Probability-metric models where confidence noise is a particularly important distinguishing factor among the models. For example, the Ideal-Confidence-Observer, Basic-Probability, and Probability-Difference models have no noise, early additive Beta noise, or late additive Gaussian noise respectively, which is applied to the estimated probability of being correct. This choice has dramatic consequences for the fit of the models (Figure 4B). According to the best-fitting parameters for the winning Heuristic model, there was also a small confidence interval bias towards reporting Interval 2 that was not significant at the group level, suggesting confidence interval bias plays less of a role than confidence noise in the participants' confidence forced-choice judgements.

Another consideration in the computation of confidence is the resolution of probability judgements for confidence evaluations. This can be best seen in the contrast of the Basic-Probability model and the LPR model. If the distribution means in both intervals are far from the centre, the raw probability values will be similarly close to 1 (Figure 1D), but if a LPR transformation is applied, small differences in the relative positions of the distribution means can translate into large differences in confidence (Figure 1E). Effectively, the LPR model has a very high resolution for extreme probability comparisons. Fortunately, this can be assessed in the confidence forced-choice method despite potential distortions in perceived probability (Kahneman and Tversky, 1979; Fox and Poldrack, 2009) because any distortion would apply to both Interval 1 and 2 and thus not change the relative judgement in confidence, as long as the distortion function is monotonic. In contrast, traditional methods of measuring confidence are likely to suffer from probability distortions in probability judgements or ceiling effects of ratings at high levels of confidence (Fleming and Lau, 2014). Our results are not conclusive on this point. While the Probability-Difference model was third best in terms of the number of observers best fit, the LPR model was overall a superior model at the group level and in the model recovery analysis showed some fit similarity to the second best-fitting Scaled-Distance model (Figure 4F). Reviewing previous studies, a high resolution for extreme probabilities does not match with findings of compression of continuous confidence probabilities to a few discrete levels for perceptual confidence (Lisi et al., 2021) or knowledge of motor uncertainty distributions (i.e., motor confidence Zhang et al., 2015). Further work is needed to understand the representation of confidence for extreme probabilities. It is possible

that our Probability-metric models were overly simplistic, and models with intermediate resolution, such as by allowing the confidence noise to vary with signal strength or including probability distortions in the mapping function (Zhang and Maloney, 2012), would improve the fit. However, the increased flexibility of such models is also likely to pose a challenge in distinguishing between candidate models.

4.4 Use of a centred prior

We were unable to reach a strong conclusion about the use of centred priors in the computation of confidence. Only three observers were best-fit by a centred-prior variant model (Figure 4B), but this increased to six if the stimulus prior assumed perfect knowledge of the three levels of dot spread tested (see Supplementary Materials). It is well known that observers do not always adapt perfectly to the experimental environment, causing them to use incorrect priors or priors matching environmental statistics (Ma, 2012; Gardner, 2019), and this may have occurred in the current study.

4.5 Confidence agreement

The confidence agreement analysis showed that the best-fitting model per observer almost always underestimated the degree of confidence agreement in the N-pass designs. From a modelling perspective, there are three factors that could have limited confidence agreement in the present experiment: 1) Type 1 sensory noise, 2) Type 2 confidence noise, and 3) the resolution of the confidence decision variable. The effect of sensory noise can be observed in the predicted confidence agreement of the Ideal Confidence Observer (Figure 5B), as this was the only factor relevant in the simulation of this model. The remaining models have lower predicted confidence agreement, which is due to the influence of the confidence noise. The Basic-Probability model had higher confidence agreement than the others, likely due to the model using a different noise distribution (Table 3). Examining the model simulations in more detail, we found that many of the confidence comparisons for the Basic-Probability model were between two extremely high probabilities and that the model is not robust to even small amounts of late decision noise. It is unlikely that a human observer has such a high resolution for the confidence decision variable, which is why resolution may be a third factor, effectively discretising confidence into various levels as suggested by Lisi et al. (2021). However, lowering the confidence resolution will decrease confidence agreement, so this alone cannot explain why the confidence models with non-discretised confidence variables under-predicted confidence agreement. It is also not due to an interval response bias, as this was included in the model simulations.

What the Basic-Probability model does illustrate is that the choice of the confidence noise model

affects confidence agreement. This suggests that future studies of confidence agreement could investigate different confidence-noise distributions (e.g., Shekhar and Rahnev, 2021), assumptions about noise being independent between the two intervals, partially or fully parallel confidence decision processes (Fleming and Daw, 2017), and serial-dependence effects on confidence (Rahnev et al., 2015). In general, the N-pass technique offers an interesting additional benchmark for assessing confidence models, just as it has proved useful in better understanding the computations of perceptual decision-making (Burgess and Colborne, 1988; Li et al., 2006; Hasan et al., 2012). Though it is important to keep in mind that the inputs to the confidence computation will never be identical due to sensory noise, so not all analyses developed for perceptual decision-making will be applicable to their confidence counterpart. Overall, confidence agreement appears to be a promising evaluation technique, with some researchers already leveraging the power of multiple presentations mostly in the form of trial “replays” to understand confidence (Balsdon et al., 2020; Charles et al., 2020), so we expect that the use of multiple passes in confidence experiments will only increase.

4.6 Strengths and limitations

A key strength of this study was the variety of metric types tested. We succeeded in fitting models of the Probability, Evidence-Strength, and Heuristic types while maintaining reasonable model identifiability (see the model recovery analysis in Figure 4E). In part, this was due to our novel approach of pairing the confidence forced-choice technique (Mamassian, 2020; Mamassian and de Gardelle, 2021) with an easy perceptual task. Specifically, the models are more divergent for easy trials and the confidence forced-choice method allowed us to probe the confidence decision variable in this range.

A downside of the confidence forced-choice technique is that there is only one confidence report for every two perceptual decisions, doubling the number of perceptual trials needed per participant. In our study, participants completed five hours of testing each on trials with very brief stimulus presentations. Using stimuli that require long presentation times (e.g., random-dot motion) would have a serious impact on the rate of data collection and introduce concerns over memory of the decision in the first interval. A consequence of using very brief stimulus presentations is that these stimuli are less suitable for studying the accumulation of evidence. Typically, researchers who investigate the temporal dynamics of decision-making use long presentation times, or let the observer decide when to terminate viewing the stimulus, using accumulation-to-bound models (e.g., Kiani et al., 2014; Zylberberg et al., 2016; Balsdon et al., 2020). For these reasons, we did not investigate accumulation-to-bound models in the present study. In general, we would not recommend the confidence forced-choice method for

evaluating these models or the temporal dynamics of confidence.

Regarding the use of mostly easy perceptual decisions, we have already highlighted some ways this may have affected the strategy of observers. To avoid indifference in the majority of confidence forced-choice judgements, observers may have used computational strategies that would give a preference to one interval over the other. Other studies found that the best-fitting model could change according to the task structure (Aitchison et al., 2015; Bertana et al., 2021), so it is possible that observers here adapted their computations of confidence to the decision context. If true, this could prevent researchers from reaching a consensus over the computational steps for confidence, and thus knowing the extent of computations possible (e.g., access to a full posterior distribution, point estimate of uncertainty, etc.). The key factors precipitating a switch in strategy (e.g., mixed-difficulty contexts, easy choices, etc.) should be the focus of computational efforts to study confidence.

The unintentional repetition of sessions resulting in an N-pass design to the experiment had both strengths and limitations. Obvious concerns are a reduction in the number of unique stimuli and stimulus pairings in the task, which affects the quality of the dataset for model fitting. There is also the possibility of the observer recognising the repetition and repeating responses, though none reported noticing the N-pass nature of the task. The clear advantage was the ability to conduct an exploratory analysis of confidence agreement, which revealed an explanatory deficit in the best-fitting models. We recommend researchers in confidence agreement design the N-pass carefully. For example, shuffling the order of repeated pairs to avoid order effects (Hasan et al., 2012). Another consideration is the difficulty of the task. Difficult trials are more likely to lead to Type 1 judgements that differ between passes, whereas in easier designs, such as in the present study, the majority of Type 1 judgements are the same. Finally, we see no reason why confidence-agreement experiments cannot be used with other types of confidence reports besides confidence forced-choice. As such, including at least a 2-pass design into any confidence experiment should be feasible.

4.7 Conclusion

By using the confidence forced-choice method, we have shown that observers are not indifferent for easy perceptual choices. Almost half of the observers took sensory uncertainty into account by computing the signal-to-noise ratio (SNR) in the Type 1 decision process, while a similar number used stimulus-based heuristic cues to compute confidence. Overall, this suggests that observers use the Distance-From-Criterion (DFC) Evidence-Strength and Heuristic metrics over Probability metrics, the most notable of which is the Bayesian Confidence Hypothesis. Furthermore, while heuristic cue use is likely

to vary according to the specifics of the individual experiment, the Scaled-Distance model is applicable 571
in any mixed-difficulty design and should be more widely considered. Our results suggest relying on the 572
simple unscaled metric typically used in extended Signal Detection Theory (SDT) could lead to worse 573
model fits or an unnecessary proliferation of confidence criteria to account for this transformation. 574
An accidental repetition of the presented stimuli also allowed us to capture the confidence agreement 575
of observers; a novel measure of model fit. This analysis revealed that observers are much more 576
consistent than would be predicted by any of the models except the Ideal Confidence Observer. We 577
propose confidence agreement as a readily accessible model-validation benchmark for future efforts in 578
modelling perceptual confidence. 579

5 Methods 580

5.1 Participants 581

Sixteen participants (21 – 43 years old, ten female) with normal or corrected-to-normal vision took 582
part in the study. All participants but two were naive to the design of the experiment. Testing was 583
conducted in accordance with the ethics requirements of the Institutional Review Board at New York 584
University. Participants received details of the experimental procedures and gave informed consent 585
prior to the experiment. 586

5.2 Apparatus 587

Stimuli were displayed on a Sony G400 CRT monitor (36 x 27 cm, 1024 x 768 pixel, 85 Hz). Participants 588
sat 55 cm from the monitor with their head stabilised by a chin rest. All responses were entered on a 589
standard computer keyboard. The experiment was conducted using custom-written code in MATLAB 590
version R2014a (The MathWorks, Natick, MA), using Psychtoolbox version 3.0.11 (Brainard, 1997; 591
Pelli, 1997; Kleiner et al., 2007). 592

5.3 Task 593

In this task, participants judged if the mean of an invisible dot-generating distribution, a 2D circular 594
symmetric Gaussian, was left or right of centre (i.e., a *Type 1*, perceptual judgement). The distribution 595
had seven possible spatial offsets of the mean (-4, -2, -1, 0, 1, 2, and 4 deg) and three possible standard 596
deviations (1.5, 2, or 2.5 deg). The distribution mean did not deviate vertically from the half-height 597
of the screen and the screen centre was indicated by a fixation cross prior to stimulus presentation. 598

Participants were presented with either two or five independently sampled dots (Gaussian-blobs with 0.1 deg SD). The white dots were simultaneously presented on a mid-grey background for 23.5 ms. The number of dots (i.e., *quantity* manipulation) and the spread of the generating distribution (i.e., *quality* manipulation), produced in six levels of sensory uncertainty. Fewer dots or larger spread made the distribution mean more difficult to localise. After every two stimulus presentations (denoted Interval 1 and Interval 2), the participant reported if they had greater confidence that their first decision or second decision was correct (i.e., a *Type 2*, metacognitive judgement). This confidence forced-choice technique avoids over- or under-confidence biases by having the observer report relative confidence (Mamassian, 2020; Mamassian and de Gardelle, 2021). An example trial pair is shown in Figure 2A. The stimulus location, number of dots, and distribution spread were randomised at the level of individual trials in an interleaved design, with trial pairings left to chance. Consequently, confidence comparisons could be between any combination of stimulus strength and sensory uncertainty (42×42 possibilities). In each session, there were 20 presentations per unique combination of stimulus strength and sensory uncertainty. New dots were sampled for each repeat. Participants completed 5 one-hour sessions of 840 perceptual judgements and 420 confidence forced-choice judgements, resulting in a total of 4200 and 2100 judgements, respectively, per participant. During the experiment, no feedback was provided on the correctness of the perceptual decisions. Data from this experiment are available at <https://osf.io/k2nhq/>.

5.4 Stimulus Set Repetition

Due to a coding oversight, the experiment was conducted with the same random seed for every session. This, coupled with the practice of switching off the testing computer between sessions, led to many of the sessions displaying the exact same sequence of stimuli (Figure 2B). For example, eight participants saw the same identical sequence of dot clouds for all five sessions; three saw this sequence for four out of five sessions; and only one participant was given a unique stimulus set for each session. In the most frequent stimulus set, of the possible pairings of sensory uncertainty for the confidence comparison, the worst sampled combination had seven unique pairs of stimuli. If the interval order is ignored, this number increases to nine unique pairs of stimuli for two same-uncertainty pairings. While this random seed setting limited the richness of the collected dataset, we saw this as an opportunity to investigate the agreement of confidence reports. In our modelling, we were able to leverage these measurements of confidence agreement to perform predictive checks of the model fits.

6 Models

629

6.1 General modelling framework

630

We modelled our participants as Bayesian observers making a joint inference about the mean and precision (i.e., inverse spread) of the dot generating distribution from a noisy observation of the dots presented onscreen. Figure 6A shows the prior, likelihood, and posterior components for an example trial. It was important to include the inference about the precision of the generating distribution because uncertainty in the estimated mean depends on the precision. This can be seen in the triangular shape of the prior, likelihood, and posterior. If the distribution’s spread is small, precision is large, and the uncertainty in the distribution’s mean is low. To make a spatial judgement, the observer marginalises over all possible precision values to get the probabilistic representation shown in Figure 6B. From this they decide whether the evidence favours a distribution mean to the left or the right, and subsequently their confidence. The advantage of using the same general framework for models of all three metric types (Probability, Evidence-Strength, and Heuristic) was that conclusions drawn from the model fits were more likely to reflect the confidence computation than “nuisance” differences between different modelling frameworks.

643

Figure 6: **Elements of the decision models.** Depicted is the estimated joint probability of the dot-cloud generating distribution mean and precision for an example stimulus presentation of 2 dots. A) Centred Normal-gamma prior distribution (based on the stimulus statistics of the experiment), likelihood function based on the noisy dot observations, and the resulting normal-gamma posterior distribution. All three distributions have the same axes as the ones shown in the posterior panel. B) The marginal prior (grey), likelihood (orange), and posterior (blue) for estimating the generating distribution mean. An unbiased Type 1 decision criterion, k_1 , is also depicted. Inset: the displayed and perceived locations of the two sampled dots, with the dashed line indicating the screen midline.

6.2 Decision context

644

There are two categories of stimuli in the perceptual task: left ($C = L$) and right ($C = R$). For each interval, the observer reports their belief about stimulus category (i.e., $\hat{C} = L$ or $\hat{C} = R$), as their Type 1 response based on their noisy sensory measurements of the sampled dots as well as any prior beliefs about the underlying generative sensory process. For the Type 2 confidence judgement, the observer compares the probability of their Type 1 perceptual decisions being correct in Interval 1 and 2: $p(\hat{C}_1 = C_1)$ versus $p(\hat{C}_2 = C_2)$.

650

In a single interval, a dot sampling distribution, defined by its horizontal mean (μ_{cloud}) and spread (σ_{cloud}), is drawn from one of the respective categories with equal probability: $P(C_L) = P(C_R) = 0.5$. The exception being the case where $\mu = 0$, which favours neither category and we treat as coming from either category randomly. From the sampling distribution, N dots are independently drawn, represented by the vector of horizontal dot locations, $\mathbf{D} = d_1, d_2, \dots, d_N$. During the measurement process, additive sensory noise is applied per dot, $\epsilon \sim N(0, \sigma_{dot}^2)$, represented by the vector of noisy horizontal dot-location measurements, $\mathbf{X} = x_1, x_2, \dots, x_N$. We assume the observer knows N (either 2 or 5) but not σ_{cloud} of the sampling distribution on any given trial. Thus, the observer's Type 1 task is to infer the category of stimulus with uncertain sampling distribution mean and spread based on the observed dot samples.

660

6.3 The likelihood function

661

The observer must consider two sources of uncertainty when choosing a likelihood function, the noisy draws of N dots from the sampling distribution, affected by σ_{cloud}^2 , and the internal noise applied to each dot, σ_{dot}^2 . These two noise sources are additive, resulting in the combined precision per dot of

664

$$\tau_{comb} = \frac{1}{\sigma_{cloud}^2 + \sigma_{dot}^2} = \frac{1}{\tau_{cloud}^{-1} + \tau_{dot}^{-1}}. \quad (1)$$

Note that we have parameterised the variances in terms of precision: $\tau_{cloud} = \sigma_{cloud}^{-2}$ and $\tau_{dot} = \sigma_{dot}^{-2}$. The likelihood function for the observed dot cloud is

$$\begin{aligned} p(\mathbf{X} | \mu_{cloud}, \tau_{cloud}, \tau_{dot}) &= \prod_{i=1}^N p(x_i | \mu_{cloud}, \tau_{comb}) \\ &\propto \tau_{comb}^{\frac{N}{2}} \exp \left(\frac{-\tau_{comb}}{2} \sum_{i=1}^N (x_i - \mu_{cloud})^2 \right). \end{aligned} \quad (2)$$

This likelihood function is centred on the dot-cloud centroid, the average of all the dot locations, because each dot is considered to be an equally reliable cue to the location of the generating distribution. An example likelihood function is shown in Figure 6A.

6.4 The prior distribution

In the centred-prior variants, we assumed the observer correctly inferred the joint distribution of the sampling distribution mean and precision through experience with the task. Even though the mean and precision were discretised in our experiment (7 and 3 possible values, respectively), we assumed that the observer could not have such a detailed representation and thus we consider that the prior is a continuous distribution. The conjugate prior for our likelihood function (i.e., inferring a normal distribution with unknown mean and precision) is a normal-Gamma distribution (Bishop, 2006; Murphy, 2007):

$$\begin{aligned} p(\mu_{cloud}, \tau_{cloud}) &= \mathcal{NG}(\mu_{cloud}, \tau_{cloud} | \mu_0, \kappa_0, \alpha_0, \beta_0) \\ &= \mathcal{N}(\mu_{cloud} | \mu_0, (\kappa_0 \tau_{cloud})^{-1}) \mathcal{Ga}(\tau_{cloud} | \alpha_0, \beta_0). \end{aligned} \quad (3)$$

To find the parameters of the normal-Gamma prior, we calculated the values of $\mu_0, \kappa_0, \alpha_0$, and β_0 such that the marginal means and variances matched the true stimulus statistics from the experiment (derivation in the Supplementary Materials). This prior distribution, shown in Figure 6A, has $\mu_0 = 0, \kappa_0 = 0.68, \alpha_0 = 3.84$, and $\beta_0 = 13.48$. Note that κ_0 can be interpreted as a number of pseudo-observations, α_0 as related to degrees of freedom, and β_0 as related to the prior belief about pooled variance.

For the flat-prior variants, we applied a uniform distribution over all possible mean and spread values. The posterior distribution is thus equivalent to the normalised likelihood function. For the posterior distribution we detail next, we are referring to the centred-prior variant models.

6.5 The posterior distribution

Using Bayes' theorem, we can write the expression for the posterior that combines the prior beliefs about the joint probability of the distribution's mean and precision with the observed dot evidence. Because we have used the conjugate prior, the posterior is also a normal-Gamma distribution (deriva-

tion provided in the Supplementary Materials):

$$\begin{aligned} p(\mu_{cloud}, \tau_{cloud} | \mathbf{X}, \tau_{dot}) &\propto p(\mathbf{X} | \mu_{cloud}, \tau_{cloud}, \tau_{dot}) p(\mu_{cloud}, \tau_{cloud}) \\ &= \mathcal{NG}(\mu_{cloud}, \tau_{cloud} | \mu_p, \kappa_p, \alpha_p, \beta_p, \tau_{dot}), \end{aligned} \quad (4)$$

with the following posterior parameters:

$$\mu_p = \frac{\kappa_0 \tau_{cloud} \mu_0 + N \tau_{comb} \bar{x}}{\kappa_0 \tau_{cloud} + N \tau_{comb}}, \quad (5)$$

$$\kappa_p = \frac{\kappa_0 \tau_{cloud} + N \tau_{comb}}{\tau_{cloud}}, \quad (6)$$

$$\alpha_p = \alpha_0 + \frac{N}{2}, \quad (7)$$

and

$$\beta_p = \beta_0 + \frac{N}{2\tau_{cloud}} \left(\frac{\kappa_0 \tau_{cloud} \tau_{comb} (\bar{x} - \mu_0)^2}{\kappa_0 \tau_{cloud} + N \tau_{comb}} + \tau_{comb} s^2 + \log(\tau_{dot}^{-1} \tau_{cloud} + 1) \right). \quad (8)$$

Where $\bar{x}$ is the sample mean of the observed dots,

$$\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i, \quad (9)$$

and s^2 is the observed sample variance in maximum-likelihood terms,

$$s^2 = \frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2. \quad (10)$$

An example posterior distribution is shown in Figure 6C.

6.6 Taxonomy of confidence models

The seven base confidence models we considered differ in the level of access to the posterior distribution. They could access either 1) the full distribution, 2) only the mode, or 3) both mode and spread. Models that made use of the full distribution used Probability-metrics for computing confidence. Confidence computations that used the mode, with or without posterior spread, were DFC Evidence-Strength metrics. The Heuristic model combined separate estimates of stimulus strength and sensory uncertainty factors (i.e., number and spread of dots) in an idiosyncratic weighted sum. Model summaries are provided in Table 2 and model equations in Table 3.

We also considered two model variants when relevant, resulting in a total of twelve unique models. For the *centred-prior* variant, the observer uses the informative prior that matched the stimulus statistics (depicted in Figure 6). For the *flat-prior* variant, the observer uses a non-informative flat prior as if they had used the normalised likelihood function to make their perceptual and confidence decisions.

In regards to confidence interval bias and noise, as per its definition, the Ideal-Confidence-Observer model did not have either of these elements (Mamassian, 2020), but otherwise all confidence models had both. Confidence interval bias in the confidence forced-choice method is a preference for reporting Interval 2, for example (Mamassian, 2020; Mamassian and de Gardelle, 2021), and should not be confused with an over- or under-confidence bias. Confidence noise was dependent on the individual confidence model. Almost all models had additive Gaussian noise. However, for models that used raw probability values (i.e., the Basic-Probability models), we used Beta-distributed variability to keep the confidence decision variable between 0 and 1.

6.7 Ideal-Confidence-Observer model

The goal of the observer was to infer if the mean of the dot sampling distribution was left or right of centre. The ideal confidence observer would take into account the unknown precision of the sampling

Table 3: Equations for the seven base Type 2 Confidence models. The models consider the decision evidence, $\text{Ev}()$, in favour of the perceptual choice, r , given the dot measurements, X , and Type 1 decision criterion, k_1 . The DFC Evidence-Strength models take a point estimate of the posterior’s mode ($\hat{\mu}$), with or without scaling by the posterior’s spread ($\hat{\sigma}$), and compare its unsigned distance from k_1 . Confidence forced-choice judgements involve comparing the relative confidence evidence for the intervals (w_1 versus w_2), with an influence of a confidence interval bias, k_2 .

Model	Confidence Evidence	Decision Rule	Confidence Noise	Free Parameters
Ideal Conf. Observer	$w = \text{Ev}(r X, k_1)$	$w_2 > w_1$	none	none
Basic Probability	$w \sim \text{Beta}(\nu p, \nu(1-p)); p = \text{Ev}(r X, k_1)$	$w_2 + k_2 > w_1$	$0 < \nu < \infty$	ν, k_2
Probability Difference	$w = \text{Ev}(r_2 X_2, k_1) - \text{Ev}(r_1 X_1, k_1) + \epsilon$	$w + k_2 > 0$	$\epsilon \sim N(0, \sigma_{\text{conf}}^2)$	$\sigma_{\text{conf}}, k_2$
Log Probability Ratio	$w = \log\left(\frac{p}{1-p}\right) + \epsilon ; p = \text{Ev}(r X, k_1)$	$w_2 + k_2 > w_1$	$\epsilon \sim N(0, \sigma_{\text{conf}}^2)$	$\sigma_{\text{conf}}, k_2$
Unscaled Distance	$w = \hat{\mu} - k_1 + \epsilon $	$w_2 + k_2 > w_1$	$\epsilon \sim N(0, \sigma_{\text{conf}}^2)$	$\sigma_{\text{conf}}, k_2$
Scaled Distance	$w = \left \frac{\hat{\mu} - k_1}{\hat{\sigma}} + \epsilon\right $	$w_2 + k_2 > w_1$	$\epsilon \sim N(0, \sigma_{\text{conf}}^2)$	$\sigma_{\text{conf}}, k_2$
Heuristic	$w = \Delta \mu_c - k_1 + \beta_1 \Delta N - \beta_2 \Delta \frac{1}{\sigma_{\text{emp}}} + \epsilon$	$w + k_2 > 0$	$\epsilon \sim N(0, \sigma_{\text{conf}}^2)$	$\beta_1, \beta_2, \sigma_{\text{conf}}, k_2$

distribution by marginalising over all possible precision values

710

$$p(\mu_{cloud}|\mathbf{X}) = \int p(\mu_{cloud}, \tau_{cloud}|\mathbf{X}, \tau_{dot}) d\tau_{cloud}, \quad (11)$$

which results in the marginal posterior distribution shown in Figure 6B. The evidence for each category, given the observed measurements and centred stimulus prior, can be computed by

711

712

$$p(C = R|\mathbf{X}, k_1) = p(\mu_{cloud} > k_1|\mathbf{X}, k_1) = \int_{k_1}^{\infty} p(\mu_{cloud}|\mathbf{X}) d\mu_{cloud}, \quad (12)$$

where k_1 is the Type 1 discrimination boundary, and

713

$$p(C = L|\mathbf{X}, k_1) = 1 - p(C = R|\mathbf{X}, k_1). \quad (13)$$

For the Type 1 judgement, the inferred category of the stimulus, $\hat{C}$, is then determined by the relative probabilities of each category, with the observer selecting the most likely category. They select “right” if $p(C = R|\mathbf{X}, k_1) > p(C = L|\mathbf{X}, k_1)$, and “left” otherwise. The observer reports the more likely category with a keypress ($r = \hat{C}$), unless a lapse occurs, where the observer selects the unintended category, with the lapse rate of λ . As such,

714

715

716

717

718

$$p(\text{choose Right}) = \lambda + (1 - 2\lambda)p\left(p(C = R|\mathbf{X}, k_1) > p(C = L|\mathbf{X}, k_1)\right). \quad (14)$$

For the Type 2 confidence judgement, the Ideal Confidence Observer considers the relative strength of the evidence in two consecutive Type 1 judgements, selecting the interval with the response that is more likely to be correct as having higher confidence, where Interval 1 evidence, w_1 , is

719

720

721

$$w_1 = \text{Ev}(r_1|\mathbf{X}_1, k_1) = \max \left[p(C_1 = R|\mathbf{X}_1, k_1), p(C_1 = L|\mathbf{X}_1, k_1) \right] \quad (15)$$

and Interval 2 evidence, w_2 is

722

$$w_2 = \text{Ev}(r_2|\mathbf{X}_2, k_1) = \max \left[p(C_2 = R|\mathbf{X}_2, k_1), p(C_2 = L|\mathbf{X}_2, k_1) \right]. \quad (16)$$

The observer reports Interval 1 for the confidence judgement if $w_1 > w_2$ and Interval 2 if $w_2 > w_1$. When a lapse in the Type 1 report occurs, the observer selects the unintended response. We assume the participant is aware of the lapse and reflects this in their confidence report. Mathematically, the

723

724

725

max operation in the above equations is replaced by a min operation for instances of lapses. The confidence evidence computation and decision rule are shown in Table 3, along with those for all Type 2 confidence models we tested. As this is the Ideal-Confidence-Observer model, the observer does not incur any additional metacognitive noise in their computation of confidence, but they are still subject to Type 1 decision bias through parameter k_1 (Mamassian, 2020). Following the decision rule of the ideal confidence observer,

$$p(\text{choose } I_2) = \lambda + (1 - 2\lambda)p(w_2 > w_1). \quad (17)$$

We assumed the Type 2 lapse rate was identical to the Type 1 lapse rate, λ , and fixed it in the Type 2 model fits according to the best-fitting value from the Type 1 model fits. For both the Type 1 and 2 judgements, the choice probabilities were estimated by simulation. The Ideal-Confidence-Observer model had no free parameters after fixing the Type-1 parameters.

6.8 Probability-metric models

The remaining six Probability-metric models follow a similar logic to the Ideal Confidence Observer at the Type 1 and 2 levels, but included various forms of metacognitive noise (Table 3). We also considered flat- and centred-prior variants. The posterior distribution in Eq. 11 differs between the two variants, but otherwise the model computations are unchanged. In all of the Probability-metric models, we included a confidence interval bias term for interval preferences, which was implemented in the decision rule as $-\infty < k_2 < \infty$. Thus, the Probability-metric models had two free parameters: ν or σ_{conf} (i.e., confidence noise), and k_2 .

Basic-Probability models. The observer directly compares the evidence in favour of their choice, but these raw probability values have been corrupted by noise. As probabilities are constrained to the interval $[0, 1]$, we implemented a beta-noise model, with ν as a concentration parameter (larger values: less confidence noise). Effectively, ν is a number of “internal” observations, of which a certain fraction are consistent with the chosen category and the rest with the unchosen category, with proportions in line with the observer’s beliefs about being correct. Consequently, as $\nu \rightarrow \infty$, w becomes a Delta distribution at $\text{Ev}(r|\mathbf{X}, k_1)$.

Probability-Difference models. The decision evidence is first compared for the intervals and then late additive Gaussian noise is applied to their difference. The smaller this confidence noise, σ_{conf} ,

the closer the observer is to the ideal confidence observer.

753

Log-Probability-Ratio (LPR) models. The decision evidence of each interval is transformed onto a continuous scale of log probability and then Gaussian noise is applied before the intervals are compared.

754

755

756

6.9 DFC Evidence-Strength metric models

757

The Distance-From-Criterion (DFC) Evidence-Strength models rely on point estimates from the decision process. In both the flat- and centred-prior variants, the observer computes the mode of the posterior, $\hat{\mu}$. Using this point estimate in the Type 1 decision will lead to identical Type 1 choice behaviour as in Eq. 14, for the same prior variant. This is because a posterior mode right of the Type 1 criterion will always corresponds to $p(C = R|\mathbf{X}, k_1) > p(C = L|\mathbf{X}, k_1)$ for the normal-gamma posterior distribution in Eq. 4, and vice versa for a posterior mode to the left. However, the DFC Evidence-Strength models make different predictions for confidence. Similar to the Probability-metric models, we included a confidence interval bias, $-\infty < k_2 < \infty$. Thus, the DFC Evidence-Strength models also had two free parameters: σ_{conf} and k_2 .

758

759

760

761

762

763

764

765

766

Unscaled-Distance model. For confidence evidence, this model considers the distance of $\hat{\mu}$ from the Type 1 discrimination boundary, k_1 , with added Gaussian noise, $N(0, \sigma_{conf}^2)$, per interval. The further $\hat{\mu}$ is from the criterion, the stronger the evidence for the spatial discrimination judgement.

767

768

769

Scaled-Distance models. The observer also considers the decision uncertainty by assessing the spread of the marginal posterior distribution in the form of a second point estimate, $\hat{\sigma}$. The distance of the mode from k_1 is then computed units of standard deviation (i.e., signal-to-noise ratio; see Supplementary Materials for more details).

770

771

772

773

6.10 Heuristic model

774

The *Heuristic* model uses an estimate of stimulus strength and estimates of each separate factor affecting sensory uncertainty (i.e, the number of dots and dot spread) as inputs to the confidence computation, without any constraint on the relative weighting of these inputs on confidence. In effect, this model is the best-fitting model from the preliminary logistic regression analysis (see Supplementary Materials), but fit in accordance with the discrimination behaviour (Table 3). This involved combining the noisy observations of the sampled dots to compute the position predictor (i.e., the centroid DFC

775

776

777

778

779

780

$|\mu_c - k_1|$), the dot spread for the quality predictor (i.e., the empirical spread, σ_{emp}), and a count of the number of dots for the quantity predictor (N), as well as any confidence interval bias (k_2). To avoid equivalent best-fitting regressors, the position coefficient that was expected to be dominant was fixed at 1 in the model and the other two coefficients, β_1 (for the difference in the number of dots across intervals) and β_2 (for the difference in inverse dot spread), were free to vary. Confidence noise, σ_{conf} , was also free to vary. In total, the Heuristic model had 4 free parameters: β_1 , β_2 , σ_{conf} , and k_2 .

6.11 Model fitting, comparison, and validation

Models were fit in a two-step procedure using custom-written Matlab code (available at <https://osf.io/k2nhq/>). Type 1 parameters were fit first (σ_{dot} , k_1 , and λ) using the discrimination responses. Then the Type 2 parameters (ν or σ_{conf} , k_2 , and possibly the two β values) were estimated using the confidence forced-choice responses. Type 1 parameters were kept fixed at their best-fitting values for the Type 2 fits. Type 1 models were fit using a brute-force grid method, with response probabilities estimated by simulation of the observer. For the Type 2 model fits, only simulated sensory measurements consistent with the observer’s Type 1 responses were used for fitting to ensure the calculated response probabilities were conditional on the discrimination choice. We used Bayesian Adaptive Direct Search (BADS) with the BADS toolbox (Acerbi and Ma, 2017), for the Type 2 model fits. Models were then compared in terms of their corrected Akaike information criterion (AICc) scores (Akaike, 1974; Cavanaugh, 1997). Further details on the model fitting are provided in the Supplementary Materials. To confirm that the models are distinguishable for this experimental design, we performed a model recovery analysis. Using the MLE of the parameters, 10 data sets were simulated per observer per model.

Acknowledgements

We would like to thank Wei Ji Ma for helpful early discussions on data analysis and Damaris Beutel for her help in collecting the research data. This work was supported by NIH Grant EY08266 (MSL, SML) and National Science Foundation Collaborative Research in Computational Neuroscience Grant 1420262, an Anneliese Maier Award from the Alexander von Humboldt Foundation (PM, SML), a post-doctoral fellowship from the Fyssen Foundation (SML), as well as the French ANR grants ANR-18-CE28-0015-01 “VICONTE” and ANR-17-EURE-0017 “FrontCog” (PM, SML).

Author note

Individual author contributions are presented using the CRediT taxonomy (see <https://www.casrai.org/credit.html> for details). Conceptualisation: S.M.L., M.S.L. and P.M.; Data curation: S.M.L.; Formal analysis: S.M.L.; Funding acquisition: S.M.L., M.S.L. and P.M.; Investigation: S.M.L.; Methodology: S.M.L., M.S.L. and P.M.; Project administration: S.M.L., M.S.L. and P.M.; Resources: M.S.L. and P.M.; Software: S.M.L.; Supervision: M.S.L. and P.M.; Validation: S.M.L.; Visualisation: S.M.L., M.S.L. and P.M.; Writing – original draft: S.M.L.; Writing - review & editing: S.M.L., M.S.L. and P.M.

References

- Acerbi, L. and Ma, W. J. (2017). Practical Bayesian optimization for model fitting with Bayesian Adaptive Direct Search. *Advances in Neural Information Processing Systems*, 30:1834–1844.
- Adler, W. T. and Ma, W. J. (2018). Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLoS Computational Biology*, 14(11):e1006572.
- Ais, J., Zylberberg, A., Barttfeld, P., and Sigman, M. (2016). Individual consistency in the accuracy and distribution of confidence judgments. *Cognition*, 146:377–386.
- Aitchison, L., Bang, D., Bahrami, B., and Latham, P. E. (2015). Doubly bayesian analysis of confidence in perceptual decision-making. *PLoS computational biology*, 11(10):e1004519.
- Akaike, H. (1974). A New Look at the Statistical Model Identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.
- Balsdon, T., Wyart, V., and Mamassian, P. (2020). Confidence controls perceptual evidence accumulation. *Nature communications*, 11(1):1–11.
- Bang, J. W., Shekhar, M., and Rahnev, D. (2019). Sensory noise increases meta-cognitive efficiency. *Journal of Experimental Psychology: General*, 148(3):437–452.
- Barrett, A. B., Dienes, Z., and Seth, A. K. (2013). Measures of metacognition on signal-detection theoretic models. *Psychological Methods*, 18(4):535–552.
- Barthelmé, S. and Mamassian, P. (2010). Flexible mechanisms underlie the evaluation of visual confidence. *Proceedings of the National Academy of Sciences*, 107(48):20834–20839.
- Bertana, A., Chetverikov, A., van Bergen, R. S., Ling, S., and Jehee, J. F. (2021). Dual strategies in human confidence judgments. *Journal of vision*, 21(5):21–21.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer. New York, NY.

- Boldt, A., de Gardelle, V., and Yeung, N. (2017). The impact of evidence reliability on sensitivity and bias in decision confidence. *Journal of Experimental Psychology: Human Perception and Performance*, 43(8):1520–1531.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4):433–436.
- Burgess, A. E. and Colborne, B. (1988). Visual signal detection. iv. observer inconsistency. *JOSA A*, 5(4):617–627.
- Cavanaugh, J. E. (1997). Unifying the derivations for the Akaike and corrected Akaike information criteria. *Statistics & Probability Letters*, 33(2):201–208.
- Charles, L., Chardin, C., and Haggard, P. (2020). Evidence for metacognitive bias in perception of voluntary action. *Cognition*, 194:104041.
- Clarke, F. R., Birdsall, T. G., and Tanner, W. P. (1959). Two types of ROC curves and definitions of parameters. *The Journal of the Acoustical Society of America*, 31(5):629–630.
- De Gardelle, V. and Mamassian, P. (2015). Weighting mean and variability during confidence judgments. *PLoS One*, 10(3):e0120870.
- Denison, R. N., Adler, W. T., Carrasco, M., and Ma, W. J. (2018). Humans incorporate attention-dependent uncertainty into perceptual decisions and confidence. *Proceedings of the National Academy of Sciences*, 115(43):11090–11095.
- Fleming, S. M. and Daw, N. D. (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review*, 124(1):91–114.
- Fleming, S. M. and Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, 8:00443.
- Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J., and Rees, G. (2010). Relating introspective accuracy to individual differences in brain structure. *Science*, 329(5998):1541–1543.
- Fox, C. R. and Poldrack, R. A. (2009). Prospect theory and the brain. In Glimcher, P. W., Camerer, C. F., Fehr, E., and Poldrack, R. A., editors, *Neuroeconomics: Decision Making and the Brain*, pages 145–173. Elsevier, New York, NY.
- Frömer, R., Nassar, M. R., Bruckner, R., Stürmer, B., Sommer, W., and Yeung, N. (2021). Response-based outcome predictions and confidence regulate feedback processing and learning. *Elife*, 10:e62825.
- Galvin, S. J., Podd, J. V., Drga, V., and Whitmore, J. (2003). Type 2 tasks in the theory of signal detectability: Discrimination between correct and incorrect decisions. *Psychonomic Bulletin & Review*, 10(4):843–876.
- Gardner, J. L. (2019). Optimality and heuristics in perceptual neuroscience. *Nature neuroscience*, 22(4):514–523.

- Gigerenzer, G. and Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62:451–482. 873 874
- Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574. 875 876
- Gorea, A. and Sagi, D. (2000). Failure to handle more than one internal representation in visual detection tasks. *Proceedings of the National Academy of Sciences*, 97(22):12380–12384. 877 878
- Graziano, M. and Sigman, M. (2009). The spatial and temporal construction of confidence in the visual scene. *PLoS One*, 4(3):e4909. 879 880
- Hasan, B. A. S., Joosten, E., and Neri, P. (2012). Estimation of internal noise using double passes: Does it matter how the second pass is delivered? *Vision research*, 69:1–9. 881 882
- Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 4:263–291. 883 884
- Kiani, R., Corthell, L., and Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron*, 84(6):1329–1342. 885 886
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., and Broussard, C. (2007). What’s new in psychtoolbox-3. *Perception*, 36(14):1–16. 887 888
- Landy, M. S., Banks, M. S., and Knill, D. C. (2011). Ideal-observer models of cue integration. In *Sensory cue integration*, pages 5–29. Oxford University Press. 889 890
- Li, H. H. and Ma, W. J. (2020). Confidence reports in decision-making with multiple alternatives violate the bayesian confidence hypothesis. *Nature communications*, 11(1):1–11. 891 892
- Li, R. W., Klein, S. A., and Levi, D. M. (2006). The receptive field and internal noise for position acuity change with feature separation. *Journal of Vision*, 6(4):2–2. 893 894
- Lisi, M., Mongillo, G., Milne, G., Dekker, T., and Gorea, A. (2021). Discrete confidence levels revealed by sequential decisions. *Nature Human Behaviour*, 5(2):273–280. 895 896
- Locke, S. M., Gaffin-Cahn, E., Hosseinizadeh, N., Mamassian, P., and Landy, M. S. (2020a). Priors and payoffs in confidence judgments. *Attention, Perception, & Psychophysics*, 82(6):3158–3175. 897 898
- Locke, S. M., Mamassian, P., and Landy, M. S. (2020b). Performance monitoring for sensorimotor confidence: A visuomotor tracking study. *Cognition*, 205:104396. 899 900
- Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in cognitive sciences*, 16(10):511–518. 901 902
- Macmillan, N. A. and Creelman, C. D. (2005). *Detection Theory: A User’s Guide*. Lawrence Erlbaum Associates, Inc., Mahwah, NJ, 2nd edition. 903 904
- Mamassian, P. (2016). Visual confidence. *Annual Review of Vision Science*, 2(1):459–481. 905

- Mamassian, P. (2020). Confidence forced-choice and other metaperceptual tasks. *Perception*, 49(6):616–635.
- Mamassian, P. and de Gardelle, V. (2021). Modeling perceptual confidence and the confidence forced-choice paradigm. *Psychological Review*, Advance online publication:<https://doi.org/10.1037/rev0000312>.
- Maniscalco, B. and Lau, H. (2016). The signal processing architecture underlying subjective reports of sensory awareness. *Neuroscience of Consciousness*, 2016(1):1–17.
- Maniscalco, B. and Lau, H. C. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, 21:422–430.
- Maniscalco, B., Peters, M. A., and Lau, H. (2016). Heuristic use of perceptual evidence leads to dissociation between performance and metacognitive sensitivity. *Attention, Perception, and Psychophysics*, 78(3):923–937.
- Meyniel, F., Schlunegger, D., and Dehaene, S. (2015). The sense of confidence during probabilistic learning: A normative account. *PLoS Comput Biol*, 11(6):e1004305.
- Mole, C. D., Jersakova, R., Kountouriotis, G. K., Moulin, C. J. A., and Wilkie, R. M. (2018). Metacognitive judgements of perceptual-motor steering performance. *Quarterly Journal of Experimental Psychology*, 71(10):2223–2234.
- Murphy, K. P. (2007). Conjugate Bayesian analysis of the Gaussian distribution. Technical report: <https://www.cs.ubc.ca/~murphyk/Papers/bayesGauss.pdf>.
- Navajas, J., Hindocha, C., Foda, H., Keramati, M., Latham, P. E., and Bahrami, B. (2017). The idiosyncratic nature of confidence. *Nature human behaviour*, 1(11):810–818.
- Patel, D., Fleming, S. M., and Kilner, J. M. (2012). Inferring subjective states through the observation of actions. *Proceedings of the Royal Society B: Biological Sciences*, 279(1748):4853–4860.
- Peirce, C. S. and Jastrow, J. (1884). On small differences in sensation. *Memoirs of the National Academy of Science*, 3:73–83.
- Pelli, D. G. (1997). The Video Toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10:437–442.
- Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117(3):864–901.
- Pouget, A., Beck, J. M., Ma, W. J., and Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nature neuroscience*, 16(9):1170–1178.
- Pouget, A., Drugowitsch, J., and Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nature Neuroscience*, 19(3):366–374.
- Rahnev, D. (2021). A robust confidence–accuracy dissociation via criterion attraction. *Neuroscience of Consciousness*, 2021(1):niab039.

- Rahnev, D. and Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, 41:e223. 941 942
- Rahnev, D. et al. (2021). Consensus goals for the field of visual metacognition. *Psyarxiv*, <https://psyarxiv.com/z8v5x>. 943 944
- Rahnev, D., Koizumi, A., McCurdy, L. Y., D’Esposito, M., and Lau, H. (2015). Confidence leak in perceptual decision making. *Psychological science*, 26(11):1664–1680. 945 946
- Rahnev, D., Maniscalco, B., Graves, T., Huang, E., De Lange, F. P., and Lau, H. (2011). Attention induces conservative subjective biases in visual perception. *Nature Neuroscience*, 14(12):1513–1515. 947 948
- Sanders, J. I., Hangya, B., and Kepecs, A. (2016). Signatures of a statistical computation in the human sense of confidence. *Neuron*, 90(3):499–506. 949 950
- Shekhar, M. and Rahnev, D. (2021). The nature of metacognitive inefficiency in perceptual decision making. *Psychological Review*, 128(1):45–70. 951 952
- Spence, M. L., Dux, P. E., and Arnold, D. H. (2015). Computations underlying confidence in visual perception. *Journal of Experimental Psychology: Human Perception and Performance*, 42(5):671–682. 953 954 955
- Van den Berg, R., Zylberberg, A., Kiani, R., Shadlen, M. N., and Wolpert, D. M. (2016). Confidence is the bridge between multi-stage decisions. *Current Biology*, 26:3157–3168. 956 957
- Vickers, D. (1979). *Decision processes in visual perception*. Academic Press, New York, NY. 958
- Zhang, H., Daw, N. D., and Maloney, L. T. (2015). Human representation of visuo-motor uncertainty as mixtures of orthogonal basis distributions. *Nature neuroscience*, 18(8):1152–1158. 959 960
- Zhang, H. and Maloney, L. T. (2012). Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in neuroscience*, 6:1. 961 962
- Zylberberg, A., Barttfeld, P., and Sigman, M. (2012). The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience*, 6:79. 963 964
- Zylberberg, A., Fetsch, C. R., and Shadlen, M. N. (2016). The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *Elife*, 5:e17688. 965 966
- Zylberberg, A., Roelfsema, P. R., and Sigman, M. (2014). Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition*, 27:246–253. 967 968

Supplementary Information:

Suprathreshold perceptual decisions constrain models of confidence

Shannon M. Locke^a, Michael S. Landy^{b,c}, Pascal Mamassian^a

a) Laboratoire des Systèmes Perceptifs, Département d'Études Cognitives, École Normale Supérieure, PSL University, CNRS, 75005 Paris, France

b) Department of Psychology, New York University, New York, NY, United States

c) Center for Neural Science, New York University, New York, NY, United States

Contents

		1
1	Model fitting	2
1.1	Derivation of the posterior for the Bayesian ideal observer	2
1.2	Selecting the prior distribution parameters	7
1.3	Computing the standard deviation of the marginal posterior distribution for the Scaled-Distance models	8
1.4	Model-fitting procedure	8
1.5	Parameter recovery	12
2	Additional results	12
2.1	Preliminary logistic analysis to confirm that quantity and quality manipulations affected confidence	12
2.2	Results of model variants	14
2.3	Examining the Heuristic-model coefficients from the simulated datasets	15
2.4	Confidence agreement behaviour and model predictions	16

1 Model fitting

15

1.1 Derivation of the posterior for the Bayesian ideal observer

16

Here we show the derivation of the posterior distribution presented in the main paper. Note that for the sake of space, we are using $\mu = \mu_{cloud}$ for the mean of the dot-sampling distribution, $\tau = \tau_{cloud}$ the precision of this distribution, and $\tau_c = \tau_{comb}$ the combined precision.

17

18

19

Prior: The prior for the mean and precision of the dot-generation distribution can be defined as the following normal-Gamma distribution:

$$\begin{aligned} p(\mu, \tau) &= \mathcal{NG}(\mu, \tau | \mu_0, \kappa_0, \alpha_0, \beta_0) \\ &= \mathcal{N}(\mu | \mu_0, (\kappa_0 \tau)^{-1}) \mathcal{G}(\tau | \alpha_0, \beta_0) \\ &\propto \tau^{\frac{1}{2}} \exp\left(\frac{-\kappa_0 \tau}{2} (\mu - \mu_0)^2\right) \times \tau^{\alpha_0 - 1} \exp(-\tau \beta_0). \end{aligned} \quad (\text{S1})$$

The parameters μ_0 and κ_0 relate to the normal component, and α_0 and β_0 to the Gamma component.

20

Likelihood: The observer must consider two sources of uncertainty when choosing a likelihood function, the noisy draws of N dots from the dot-generation distribution and the measurement noise applied to each dot, σ_{dot}^2 (precision: τ_{dot}). These two noise sources are additive, resulting in the combined precision per dot of

21

22

23

24

$$\tau_c = \frac{1}{\sigma_{cloud}^2 + \sigma_{dot}^2} = \frac{1}{\tau^{-1} + \tau_{dot}^{-1}}. \quad (\text{S2})$$

The observer must consider the information provided by all dot locations $\mathbf{X}$ (that are the horizontal component of physical dot locations $\mathbf{D}$ corrupted by measurement noise), resulting in the following likelihood function:

$$\begin{aligned} p(\mathbf{X} | \mu, \tau, \tau_{dot}) &= \prod_{i=1}^N p(x_i | \mu, \tau_c) \\ &= \prod_{i=1}^N \frac{1}{(2\pi)^{\frac{1}{2}}} \tau_c^{\frac{1}{2}} \exp\left(\frac{-\tau_c}{2} (x_i - \mu)^2\right) \\ &\propto \tau_c^{\frac{N}{2}} \exp\left(\frac{-\tau_c}{2} \sum_{i=1}^N (x_i - \mu)^2\right), \end{aligned} \quad (\text{S3})$$

where τ_c is defined in Eq. S2 and we assume the observer has accurate knowledge of σ_{dot}^2 . Note that we have dropped the constant terms from the equation for simplicity as they can be applied by re-scaling

25

26

the posterior.

27

Posterior: As we have used a conjugate prior, the posterior is also a normal gamma of the form $\mathcal{NG}(\mu, \tau | \mu_p, \kappa_p, \alpha_p, \beta_p, \tau_{dot})$. However, to see this clearly, we decompose the posterior into three parts (A, B, and C).

$$\begin{aligned} p(\mu, \tau | \mathbf{X}, \tau_{dot}) &\propto \text{prior} \times \text{likelihood} \\ &\propto \tau^{\frac{1}{2}} \exp\left(\frac{-\kappa_0 \tau}{2} (\mu - \mu_0)^2\right) \tau^{\alpha_0 - 1} \exp(-\tau \beta_0) \tau_c^{\frac{N}{2}} \exp\left(\frac{-\tau_c}{2} \sum_{i=1}^N (x_i - \mu)^2\right) \\ &\propto A \times B \times C. \end{aligned} \tag{S4}$$

Part A is the non-exponent term of the normal component,

28

$$A = \tau^{\frac{1}{2}}. \tag{S5}$$

Part B is the non-exponentiated term of the Gamma component, which can be rearranged as follows

$$\begin{aligned} B &= \tau^{\alpha_0 - 1} \tau_c^{\frac{N}{2}} \\ &= \tau^{\alpha_0 - 1} \left(\frac{\tau}{\tau_{dot}^{-1} \tau + 1} \right)^{\frac{N}{2}} \\ &= \tau^{\alpha_0 + N/2 - 1} \exp\left(\frac{-N}{2} \log(\tau_{dot}^{-1} \tau + 1)\right). \end{aligned} \tag{S6}$$

The non-exponent part of B will be the non-exponent term of the Gamma component, and the remainder will be placed in Part C:

29

30

$$C = \exp\left(\frac{-\kappa_0 \tau}{2} (\mu - \mu_0)^2\right) \exp(-\tau \beta_0) \exp\left(\frac{-\tau_c}{2} \sum_{i=1}^N (x_i - \mu)^2\right). \tag{S7}$$

Thus, Part B becomes

31

$$B' = \tau^{\alpha_0 + N/2 - 1} = \tau^{\alpha_p - 1}, \tag{S8}$$

with

32

$$\alpha_p = \alpha_0 + N/2, \tag{S9}$$

and Part C, which now aggregates all the exponent terms, is

$$\begin{aligned}
C' &= \exp\left(\frac{-\kappa_0\tau}{2}(\mu - \mu_0)^2\right) \exp(-\tau\beta_0) \exp\left(\frac{-\tau_c}{2}\sum_{i=1}^N(x_i - \mu)^2\right) \exp\left(\frac{-N}{2}\log(\tau_{dot}^{-1}\tau + 1)\right) \\
&= \exp\left(-\frac{1}{2}\left[\kappa_0\tau(\mu - \mu_0)^2 + \tau_c\sum_{i=1}^N(x_i - \mu)^2\right] - \tau\beta_0 - \frac{N}{2}\log(\tau_{dot}^{-1}\tau + 1)\right) \\
&= \exp\left(-\frac{1}{2}D - \tau\beta_0 - \frac{N}{2}\log(\tau_{dot}^{-1}\tau + 1)\right), \tag{S10}
\end{aligned}$$

which will contribute to both the normal and the Gamma components. Part D of this equation can be simplified, by extracting Part E as follows

$$\begin{aligned}
D &= \kappa_0\tau(\mu - \mu_0)^2 + \tau_c\sum_{i=1}^N(x_i - \mu)^2 \\
&= \kappa_0\tau(\mu - \mu_0)^2 + N\tau_c(\mu - \bar{x})^2 + \tau_c\sum_{i=1}^N(x_i - \bar{x})^2 \\
&= E + \tau_c\sum_{i=1}^N(x_i - \bar{x})^2, \tag{S11}
\end{aligned}$$

given that

$$\begin{aligned}
\sum_{i=1}^N(x_i - \mu)^2 &= \sum_{i=1}^N[(x_i - \mu - \bar{x} + \bar{x})]^2 \\
&= \sum_{i=1}^N(\bar{x} - \mu)^2 + \sum_{i=1}^N(x_i - \bar{x})^2 - 2\sum_{i=1}^N(x_i - \bar{x})(\mu - \bar{x}) \\
&= N(\bar{x} - \mu)^2 + \sum_{i=1}^N(x_i - \bar{x})^2, \tag{S12}
\end{aligned}$$

and

$$\sum_{i=1}^N(x_i - \bar{x})(\mu - \bar{x}) = (\mu - \bar{x})\left(\sum_{i=1}^N x_i - N\bar{x}\right) = (\mu - \bar{x})(N\bar{x} - N\bar{x}) = 0. \tag{S13}$$

Afterwards, completing the square¹ and rearranging, we get

$$\begin{aligned}
E &= \kappa_0 \tau (\mu - \mu_0)^2 + N \tau_c (\mu - \bar{x})^2 \\
&= \kappa_0 \tau (\mu^2 - 2\mu\mu_0 + \mu_0^2) + N \tau_c (\mu^2 - 2\mu\bar{x} + \bar{x}^2) \\
&= \mu^2 (\kappa_0 \tau + N \tau_c) - 2\mu (\kappa_0 \tau \mu_0 + N \tau_c \bar{x}) + \kappa_0 \tau \mu_0^2 + N \tau_c \bar{x}^2 \\
&= (\kappa_0 \tau + N \tau_c) (\mu - \mu_p)^2 + \frac{(\kappa_0 \tau \mu_0 + N \tau_c \bar{x})^2}{\kappa_0 \tau + N \tau_c} \\
&= (\kappa_0 \tau + N \tau_c) (\mu - \mu_p)^2 + \frac{\kappa_0 \tau \mu_0^2 (\kappa_0 \tau + N \tau_c) + N \tau_c \bar{x}^2 (\kappa_0 \tau + N \tau_c) - (\kappa_0 \tau \mu_0 + N \tau_c \bar{x})^2}{\kappa_0 \tau + N \tau_c} \\
&= (\kappa_0 \tau + N \tau_c) (\mu - \mu_p)^2 + \frac{\kappa_0 N \tau \tau_c (\bar{x} - \mu_0)^2}{\kappa_0 \tau + N \tau_c}, \tag{S14}
\end{aligned}$$

where

$$\mu_p = \frac{\kappa_0 \tau \mu_0 + N \tau_c \bar{x}}{\kappa_0 \tau + N \tau_c}. \tag{S15}$$

Thus, Part D can be written as

34

$$D = (\kappa_0 \tau + N \tau_c) (\mu - \mu_p)^2 + \frac{\kappa_0 N \tau \tau_c (\bar{x} - \mu_0)^2}{\kappa_0 \tau + N \tau_c} + \tau_c \sum_{i=1}^N (x_i - \bar{x})^2, \tag{S16}$$

and we can now express Part C' in terms of two components, C_1 and C_2 ,

$$\begin{aligned}
C' &= \exp \left(-\frac{1}{2} D - \tau \beta_0 - \frac{N}{2} \log(\tau_{dot}^{-1} \tau + 1) \right) \\
&= \exp \left(-\frac{1}{2} (\kappa_0 \tau + N \tau_c) (\mu - \mu_p)^2 \right) \times \\
&\quad \exp \left(-\tau \beta_0 - \frac{\kappa_0 N \tau \tau_c (\bar{x} - \mu_0)^2}{2(\kappa_0 \tau + N \tau_c)} - \frac{\tau_c}{2} \sum_{i=1}^N (x_i - \bar{x})^2 - \frac{N}{2} \log(\tau_{dot}^{-1} \tau + 1) \right) \\
&= C_1 \times C_2, \tag{S17}
\end{aligned}$$

¹Completing the square: $ax^2 + bx + c = a(x + d)^2 + e$, where $d = b/2a$ and $e = c - b^2/4a$.

which contribute to the normal component and the Gamma component respectively. To see this, we must arrange Part C_1 as follows

$$\begin{aligned}
C_1 &= \exp \left(-\frac{1}{2}(\kappa_0\tau + N\tau_c)(\mu - \mu_p)^2 \right) \\
&= \exp \left(\frac{-\tau(\kappa_0\tau + N\tau_c)}{2\tau}(\mu - \mu_p)^2 \right) \\
&= \exp \left(\frac{-\kappa_p\tau}{2}(\mu - \mu_p)^2 \right),
\end{aligned} \tag{S18}$$

where

$$\kappa_p = \frac{\kappa_0\tau + N\tau_c}{\tau}. \tag{S19}$$

Similarly, for Part C_2 ,

$$\begin{aligned}
C_2 &= \exp \left(-\tau\beta_0 - \frac{\kappa_0 N\tau\tau_c(\bar{x} - \mu_0)^2}{2(\kappa_0\tau + N\tau_c)} - \frac{\tau_c}{2} \sum_{i=1}^N (x_i - \bar{x})^2 - \frac{N}{2} \log(\tau_{dot}^{-1}\tau + 1) \right) \\
&= \exp \left(-\tau \left[\beta_0 + \frac{\kappa_0 N\tau\tau_c(\bar{x} - \mu_0)^2}{2\tau(\kappa_0\tau + N\tau_c)} + \frac{\tau_c}{2\tau} \sum_{i=1}^N (x_i - \bar{x})^2 + \frac{N}{2\tau} \log(\tau_{dot}^{-1}\tau + 1) \right] \right) \\
&= \exp \left(-\tau \left[\beta_0 + \frac{N}{2\tau} \left(\frac{\kappa_0\tau\tau_c(\bar{x} - \mu_0)^2}{\kappa_0\tau + N\tau_c} + \tau_c s^2 + \log(\tau_{dot}^{-1}\tau + 1) \right) \right] \right) \\
&= \exp(-\tau\beta_p),
\end{aligned} \tag{S20}$$

where

$$\beta_p = \beta_0 + \frac{N}{2\tau} \left(\frac{\kappa_0\tau\tau_c(\bar{x} - \mu_0)^2}{\kappa_0\tau + N\tau_c} + \tau_c s^2 + \log(\tau_{dot}^{-1}\tau + 1) \right), \tag{S21}$$

and s^2 is the sample variance

$$s^2 = \frac{1}{N} \sum_{i=1}^N (x_i - \bar{x})^2. \tag{S22}$$

To see how this is a normal-Gamma we reassemble

$$\begin{aligned}
p(\mu, \tau | \mathbf{X}, \tau_{dot}) &\propto AC_1 \times B'C_2 \\
&\propto \tau^{\frac{1}{2}} \exp \left(\frac{-\kappa_p\tau}{2}(\mu - \mu_p)^2 \right) \times \tau^{\alpha_p-1} \exp(-\tau\beta_p) \\
&\propto \mathcal{NG}(\mu, \tau | \mu_p, \kappa_p, \alpha_p, \beta_p).
\end{aligned} \tag{S23}$$

Validation: We can check this derivation against the simpler case of $p(\mu, \tau | \mathbf{D})$, which has been documented previously in the literature (Murphy, 2007). The important distinction between these

derivations is that there is no additive noise per dot in the likelihood for $p(\mu, \tau | \mathbf{D})$, with the consequence that $\tau_c = \tau$. In this case, the parameters of the posterior simplify as follows:

$$\mu_p = \frac{\kappa_0 \tau \mu_0 + N \tau_c \bar{x}}{\kappa_0 \tau + N \tau_c} = \frac{\kappa_0 \mu_0 + N \bar{x}}{\kappa_0 + N}, \quad (\text{S24})$$

$$\kappa_p = \frac{\kappa_0 \tau + N \tau_c}{\tau} = \kappa_0 + N, \quad (\text{S25})$$

$$\alpha_p = \alpha_0 + N/2, \quad (\text{S26})$$

and

$$\begin{aligned} \beta_p &= \beta_0 + \frac{N}{2\tau} \left(\frac{\kappa_0 \tau \tau_c (\bar{x} - \mu_0)^2}{\kappa_0 \tau + N \tau_c} + \tau_c s^2 + \log(\tau_{dot}^{-1} \tau + 1) \right) \\ &= \beta_0 + \frac{\kappa_0 N (\bar{x} - \mu_0)^2}{2(\kappa_0 + N)} + \frac{N s^2}{2}. \end{aligned} \quad (\text{S27})$$

These parameter equations match those reported in the derivation of $p(\mu, \tau | \mathbf{D})$ in Murphy (2007).

1.2 Selecting the prior distribution parameters

We adjusted the parameters of the normal-Gamma prior ($\mu_0, \kappa_0, \alpha_0$, and β_0) to the true stimulus statistics of the dots mean μ and precision τ in the experiment. The possible centre locations of the dots were -4, -2, -1, 0, 1, 2, or 4 deg, and the possible standard deviations of the dots spreads were 1.5, 2, or 2.5 deg. Therefore, the statistics of the mean and precision were $E[\mu] = 0$, $\text{Var}[\mu] = 7$, $E[\tau] = 0.2848$, and $\text{Var}[\tau] = 0.0211$.

We matched the marginal distributions of the mean and precision variables to these statistics. The marginal distribution of the precision τ variable is the Gamma distribution with parameters (α_0, β_0) . The mean and variance of this distribution are

$$E[\tau] = \frac{\alpha_0}{\beta_0}, \quad (\text{S28})$$

$$\text{Var}[\tau] = \frac{\alpha_0}{\beta_0^2}. \quad (\text{S29})$$

The marginal distribution of the mean μ variable is the non-standardized Student's t-distribution with parameters $(\nu_0, \mu_0, \sigma_0^2) = (2\alpha_0, \mu_0, \beta_0/(\kappa_0 \alpha_0))$. The mean and variance of this distribution are

$$E[\mu] = \mu_0, \quad (\text{S30})$$

$$\text{Var}[\mu] = \frac{\sigma_0^2 \nu_0}{\nu_0 - 2} = \frac{\beta_0}{\kappa_0 \alpha_0} \frac{2\alpha_0}{2\alpha_0 - 2} = \frac{\beta_0}{\kappa_0(\alpha_0 - 1)}. \quad (\text{S31})$$

The match of these marginal statistics results in the following prior distribution parameters

$$\mu_0 = \text{E}[\mu] = 0, \quad (\text{S32})$$

$$\kappa_0 = \frac{\text{E}[\tau]}{\text{Var}[\mu](\text{E}[\tau]^2 - \text{Var}[\tau])} = 0.68, \quad (\text{S33})$$

$$\alpha_0 = \frac{\text{E}[\tau]^2}{\text{Var}[\tau]} = 3.84, \quad (\text{S34})$$

and

$$\beta_0 = \frac{\text{E}[\tau]}{\text{Var}[\tau]} = 13.48. \quad (\text{S35})$$

1.3 Computing the standard deviation of the marginal posterior distribution for the Scaled-Distance models

The Scaled-Distance models normalised the distance-From-Criterion (DFC) of the posterior mode, $\hat{\mu}$, by the standard deviation of the marginal posterior distribution, $\hat{\sigma}$, for each interval. We calculated the scaled DFC in the model fitting procedure using an inverse cumulative normal function approximation.

$$\text{DFC}_{\text{scaled}} = \left| \frac{\hat{\mu} - k_1}{\hat{\sigma}} + \epsilon \right| \approx |\Phi_{\text{inv}}(p(C = R | \mathbf{X}, k_1)) + \epsilon|. \quad (\text{S36})$$

Note the additive confidence noise, $\epsilon \sim N(0, \sigma_{\text{conf}}^2)$.

1.4 Model-fitting procedure

Type 1 models. Type 1 models were fit first, with σ_{dot} , k_1 , and λ as free parameters. Models were fit on a per-participant basis with custom MATLAB scripts that calculated the likelihood of the model for specific parameter combinations in a dense grid of parameter values (σ_{dot} : 0.1 to 2 deg in 50 log-spaced steps; k_1 : -1 to 1 deg in 20 steps; λ : 0.01 to 0.1 in 9 steps). Choice probabilities for each parameter combination were estimated by simulation. For each trial, 1000 observations were simulated, with choice probabilities estimated by computing the proportion of rightward choices made by the simulated observer. This involved constructing the prior, likelihood, and posterior distributions shown in Figure 6A of the main paper (grid resolution: μ : -10 to 10 in 200 steps; τ : 0.01 to 1.00 in 100 steps). After marginalising the posterior distribution across the τ dimension, the choice probabilities could be calculated using k_1 and λ . Because this was a computationally expensive simulation, the simulated

noisy observations and construction of the prior, likelihood, and posterior grids was done only once per level of σ_{dot} (for each trial and for each participant). Thus differences in choice probabilities from k_1 and λ were calculated using the same 1000 noisy simulations of the trial. Not taking this approach would have increased the time to fit a single Type 1 model for an observer from hours to days or weeks. Choice probabilities were then put into a Bernoulli model of binary choice behaviour with the participant's actual responses to calculate the likelihood of that specific parameter combination. Finally, the best-fitting parameters were extracted from the 3D parameter grid by finding combination that gave the largest the maximum likelihood value. Averages of the best-fitting parameter values are given in Table S1.

Type 2 models. The Type 2 models were fit using the same simulation-based approach combined with Bayesian Adaptive Direct Search (BADS; Acerbi and Ma, 2017) to find the best-fitting combination of parameters. The number of parameters in the Type 2 model fits depended on the model. The Ideal-Confidence-Observer model had no free parameters, the Heuristic model had 4 free parameters ($\beta_1, \beta_2, \sigma_{conf}, k_2$), and the remaining models had two free parameters (ν or σ_{conf}, k_2). The Type 1 parameters were fixed at their best-fitting values on a per-participant basis and the Type 1 model was matched to the Type 2 model in terms of whether the flat- or centred-prior variant was used because the location of a biased k_1 can differ between the variants.

The first step in Type 2 model fits was simulating the observer. For each participant, we simulated 3000 noisy observations of each trial, and computed the marginalised posterior distribution using the same μ - τ grid spacing as in the Type 1 model-fitting procedure. For each simulation, we computed $p(C = R|X, k_1)$ of Eq. 12 as well as the mode of the posterior distribution according to the prior variants of the Type 1 models. Each simulation was then coded as favouring a leftward or rightward Type 1 response. We sampled with replacement from the simulations that matched the Type 1 response given by the participant to ensure the simulated observer matched the actual observer in Type 1 choice behaviour. If all simulations were choice-consistent for a given trial we did not perform this sampling step. If no samples were choice-consistent, the trial was flagged as a likely lapse for later processing. Due to the mostly low difficulty of the task, the majority of samples usually favoured the same response and it was rare for the sampling to be drawn from a small number of simulations. This simulation procedure was performed only once per trial and observer and used for the fitting of all Type 2 models.

In the second step, we fit each of the Type 2 models with the set of choice-consistent simulations

using BADS. We considered the relative evidence for the Type 1 response given by the observer in each interval, according to the Type 2 model. Likely lapses were given the lowest confidence value possible with a tiny amount of jitter so the confidence report for comparing two lapses is entirely random. Thus, we were able to compute the choice probabilities for selecting “Interval 1” and “Interval 2” from the simulated observer. These were compared with the actual confidence forced-choice report of the participant in a Bernoulli model of binary choice behaviour. As choice probabilities were computed by simulation, we used the “uncertainty handling” option in the BADS package and drew 100 final samples to compute the estimate of the negative log-likelihood value of the model for the model comparison. The lower bound, upper bound, plausible lower bound, plausible upper bound, and starting location settings for each model can be found in the *fitType2.m* file available at <https://osf.io/k2nhq/>. Averages of the best-fitting parameter values per Type 2 model are given in Table S1.

Table S1: Summary of best-fitting parameters per model. Type 1 model fits were near identical, so parameter values were averaged across the four models for each observer before computing the parameter averages across observers. Type 2 models are detailed individually. Where relevant, the parameters from fits of the flat- and centred-prior variants are both reported.

Model	Parameter	Mean±SEM
All Type 1 (averaged)	σ_{dot}	$0.97^\circ \pm 0.04^\circ$
	k_1	$0.01^\circ \pm 0.03^\circ$
	λ	0.02 ± 0.002
Ideal Confidence Observer	none	none
Basic Probability	ν_{conf}	6.30 ± 1.24 ; 4.91 ± 0.90
	k_2	0.001 ± 0.001 ; 0.001 ± 0.001
Probability Difference	σ_{conf}	0.19 ± 0.03 ; 0.21 ± 0.03
	k_2	0.02 ± 0.01 ; 0.02 ± 0.01
Log Probability Ratio	σ_{conf}	2.29 ± 0.29 ; 2.15 ± 0.28
	k_2	0.11 ± 0.16 ; 0.09 ± 0.14
Unscaled Distance	σ_{conf}	$2.08^\circ \pm 0.18^\circ$; $1.42^\circ \pm 0.12^\circ$
	k_2	$0.04^\circ \pm 0.11^\circ$; $0.03^\circ \pm 0.08^\circ$
Scaled Distance	σ_{conf}	0.98 ± 0.11 sd; 0.94 ± 0.11 sd
	k_2	0.05 ± 0.07 sd; 0.04 ± 0.06 sd
Heuristic	σ_{conf}	1.12 ± 0.17
	k_2	0.13 ± 0.08
	β_1 (quantity)	0.68 ± 0.10
	β_2 (quality)	0.38 ± 0.09

Figure S1: Parameter recovery results for σ_{conf} (blue) and k_2 (red). Model and prior variant indicated by title. Data points are the mean values and error bars are ± 1 SD calculated from the fits of the 10 simulated data sets per observer per model.

1.5 Parameter recovery

To confirm that the model-fitting procedure was behaving correctly, we assessed whether refit models in the model-recovery analysis recovered the simulated parameters. Figure S1 shows that the confidence noise and bias terms were recovered well in almost all instances. The exception was the confidence-noise parameter for the Basic-Probability model, where the magnitude of the noise was often overestimated. We investigated this further and found that the Beta noise distribution was little changed for different confidence noise values for extreme probability values (i.e., close to 0 or 1). Thus, the design of the experiment likely affected the ability to fit this model.

2 Additional results

2.1 Preliminary logistic analysis to confirm that quantity and quality manipulations affected confidence

We examined whether the evidence quantity and quality manipulations affected confidence judgments. Observers could have ignored the sensory-uncertainty manipulations, and simply reported whichever interval had the more extreme leftward or rightward stimulus as more confident (the *Basic* model). Alternatively, the observer may modulate their confidence by information quantity and/or quality: increasing confidence if there were 5 dots or smaller spread, and decreasing it for 2 dots or larger spread. We considered these four possibilities (Basic, only quantity, only quality, and both) in a nested model comparison. It is reasonable to suspect that the observer inferred the mean and spread from the displayed dots, so we repeated the analysis for both the true mean and spread of the generating distribution and the empirical mean and spread values according to the dots displayed. If decision correctness is examined, observer reports better match the empirical mean of the dots (i.e., the centroid) as opposed to the mean of generating distribution (Figure S2A).

All models for the preliminary analysis were a logistic function of the probability of choosing Interval 2 as more confident,

$$p(\text{choose } I_2) = \lambda + (1 - 2\lambda)\text{Logit}(\theta), \quad (\text{S37})$$

with the predicted probabilities scaled according to the lapse rate, λ ($0.01 \leq \lambda \leq 0.15$). Here lapse rate is the rate of unintentionally giving the opposite decision than desired. The *Full* model had three predictors related to the distance of the stimuli from the screen centre, the number of dots, and the

dot spreads in each interval:

147

$$\theta_{\text{Full}} = \beta_0 + \beta_1 \Delta|\mu| + \beta_2 \Delta N_{\text{dots}} + \beta_3 \Delta \frac{1}{\sigma_{\text{cloud}}}. \quad (\text{S38})$$

The constant term, β_0 , is the confidence bias for choosing Interval 2 and was included in all models. 148
The relative distance of the stimuli was parameterised as the relative absolute distance of the dot 149
centroids from the screen centre, $\Delta|\mu|$, and was also included in all models. The relative number of 150
dots, ΔN_{dots} , and the inverse of the spreads, $\Delta 1/\sigma_{\text{cloud}}$, were optional predictors in the nested models. 151
The *Basic* model did not consider the evidence quantity or quality manipulations when computing 152
confidence, 153

$$\theta_{\text{Basic}} = \beta_0 + \beta_1 \Delta|\mu|. \quad (\text{S39})$$

Compared to the Basic model, the +*Quantity* model that included the dot-number predictor (β_2) and 154
the +*Quality* model that included the dot-spread predictor (β_3) provided a better fit to the data, 155
with the +Quantity model fitting better than the +Quality model (Figure S2B). This was true both 156
if the generating distribution spread was used or if the empirical spread was used (i.e., based on the 157
displayed dots). The empirical version did fit slightly better overall. 158

The models were fit with custom MATLAB scripts that calculated the maximum-likelihood esti- 159
mates of the parameters using gradient decent. Corrected Akaike information criterion (AICc) scores 160
(Akaike, 1974; Cavanaugh, 1997) were computed for model comparison from the log-likelihood values. 161

Figure S2: Preliminary results. A) Type 1 judgements. Task performance when decision correctness is scored against the true mean of the generating distribution (pink) versus the empirical mean of the dot cloud (purple). B) Type 2 Logistic regression results. Model fit is compared against the winning model: the Full model with empirical spread. Grey: models without the quality predictor. Pink: models with a quality predictor based on the true generating distribution spread. Purple: quality predictor based on the empirical spread. All error bars: $\pm\text{SEM}$.

AICc scores explicitly penalise the number of parameters in a model and correct for sample size. 162

We also repeated this nested analysis using a restricted dataset of only trial pairs where both 163
Type 1 judgements were correct. The purpose of this analysis was to reduce the probability that 164
lapse discrimination responses were included, because the logistic regression did not consider Type 1 165
performance in the confidence report. The pattern of model fits, however, did not change. 166

2.2 Results of model variants 167

The base model assumed that the likelihood function was calculated from the horizontal locations of 168
the noisy dot measurements. In the centred-prior variant, the prior distribution was a normal-Gamma 169
distribution matched to the stimulus statistics. These modelling choices were investigated by checking 170
whether the main experimental results were robust to changes in this general base model. 171

Uniform prior on the three precision levels. In this analysis, we changed the joint mean- 172
precision prior of the centred-prior to reflect that the observer had perfect knowledge of the three 173
levels of distribution spread used in the experiment. On an individual trial, the observer considered 174
the prior probability of the precision as uniform across these three levels. The prior probability of 175
the mean was matched to the stimulus statistics, $N(0, 7)$, for the centred-prior variant. The winning 176
model remained the Heuristic model, with the number best fit for this model unchanged (Figure S3A). 177
However, the number of observers best-fit by the Probability-Difference model increased by 2, and the 178
number of observers best-fit with a posterior variant doubled from 3 to 6. Note that the flat-prior 179
variant results are unchanged in this analysis, so changes of the best-fitting model were a result of the 180
centred-prior variants fitting better or worse. 181

Empirical spread from Euclidean distance. Here we included the vertical spread of the dots. 182
The vertical position of the generating distribution mean was always the half-height of the screen. Yet, 183
the distance of the dots from this half-height reference also give an indication of the true spread of the 184
circular-symmetric Gaussian generating distribution. The Heuristic model remained the overall winner 185
of the model comparison (Figure S3B). However, now the Scaled-Distance model has the slightly more 186
observers best fit compared to the Heuristic model. We do not consider this a significant departure 187
from the results reported in the main article. 188

Figure S3: Type 2 model-comparison results when the form of the general base model was changed. A) Models with a uniform prior on τ instead of the normal-Gamma prior. We assume here that the observer has perfect knowledge of the three levels of dot spread and implements a prior that gives equal weight to each of the three corresponding τ values. Bars: average relative AICc score. Markers: Individual participant results. Colour: flat-prior (orange) or centred-prior (blue) variant. Heuristic model has no prior. B) Models with the empirical spread computed from Euclidean distance instead of horizontal distance. Both the horizontal and vertical position of the dot are used in the standard deviation calculation. Error bars: $\pm SEM$.

2.3 Examining the Heuristic-model coefficients from the simulated datasets

To better understand the relationship between the best-fitting coefficients from the Heuristic model and the best-fitting model for each observer (Figure S4A), we examined the fits of the model simulations from the model-recovery analysis. In particular, there appeared to be some clustering of coefficient values based on the model type for the non-heuristic models. The model simulations, based on all observers, confirmed this clustering. It also showed a large spread in the coefficients for the Heuristic model simulations, consistent with the fits to the human data. A similar spread was

observed for the Ideal-Confidence-Observer model simulations (Figure S4C), and a small spread for the Basic-Probability models simulations. The simulated Unscaled-Distance model coefficients are near 0, consistent with these models having little to no influence of information quantity or quality on confidence, for the flat- and centred-prior variants respectively.

Figure S4: Comparing the Heuristic-model fit for the behaviour of observers against the Heuristic-model fit for the model simulations. A) Replotting of observer fits from Figure 5D. Each datapoint is a single observer, coloured according to the best-fitting model (note change in colour scheme). Prior variants are flat (F) or centred (C). B) Same as in A for the simulations of the selected models. Each datapoint is the fit from a single model simulation (10 simulations per observer; best-fitting parameters used). C) The best-fitting Heuristic-model coefficients for each of the models.

2.4 Confidence agreement behaviour and model predictions

The confidence agreement behaviour of all observers, and their corresponding model predictions according to their best-fitting parameters are shown in Figure S5. The pattern of results reveals the tendency of the observer to be more consistent than their best-fitting model would predict.

Figure S5: The 5-pass confidence agreement of all observers with the predicted confidence agreement of each model. Model predictions were calculated from 100 simulated datasets using the participant-specific best-fitting parameters. Error bars: ± 2 SD. For all observers, the Ideal-Confidence-Observer model had the highest confidence agreement, then the Basic-Probability model, then the remaining models. Red: confidence agreement of the observer. Green: predicted confidence agreement according to the best-fitting model for that observer. Black: predictions of the other models. Note that different observers had different maximum agreement counts depending on the number of repeated blocks (Figure 1B).

References

- Acerbi, L. and Ma, W. J. (2017). Practical Bayesian optimization for model fitting with Bayesian Adaptive Direct Search. *Advances in Neural Information Processing Systems*, 30:1834–1844.
- Akaike, H. (1974). A New Look at the Statistical Model Identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.
- Cavanaugh, J. E. (1997). Unifying the derivations for the Akaike and corrected Akaike information criteria. *Statistics & Probability Letters*, 33(2):201–208.
- Murphy, K. P. (2007). Conjugate Bayesian analysis of the Gaussian distribution. Technical report: <https://www.cs.ubc.ca/~murphyk/Papers/bayesGauss.pdf>.