

Different types of uncertainty in multisensory perceptual decision making

Stacey Aston, Marko Nardini & Ulrik Beierholm

Dept of Psychology, Durham University, UK

In press, *Philosophical Transactions of the Royal Society B*, June 2023

Abstract

Efficient decision-making requires accounting for sources of uncertainty (noise, or variability). Many studies have shown how the nervous system is able to account for perceptual uncertainty (noise, variability) that arises from limitations in its own abilities to encode perceptual stimuli. However, many other sources of uncertainty exist, reflecting for example variability in the behaviour of other agents or physical processes. Here we review previous studies on decision making under uncertainty as a function of the different types of uncertainty that the nervous system encounters, showing that noise that is intrinsic to the perceptual system can often be accounted for near-optimally (i.e. not statistically different from optimally), whereas accounting for other types of uncertainty can be much more challenging. As an example, we present a study in which participants made decisions about multisensory stimuli with both intrinsic (perceptual) and extrinsic (environmental) uncertainty and show that the nervous system accounts for these differently when making decisions: they account for internal uncertainty but under-account for external. Human perceptual systems may be well equipped to account for intrinsic (perceptual) uncertainty because, in principle, they have access to this. Accounting for external uncertainty is more challenging because this uncertainty must be learned.

Keywords

Decision-making, perceptual, multisensory, uncertainty, noise

Background

Choosing the best course of action depends on evaluating the evidence in favour of different options. A problem for any decision-maker is the quality of the data available. Major limitations to data quality come from the *bias* (systematic error) and *uncertainty* (random error, or noise, or variability). An observer whose perception of the approach times of oncoming traffic is strongly biased, such that they systematically over-rate the available time to cross the road, is in danger. So is an observer whose perception is very uncertain: the high variability of their judgments means that, on any given occasion, there is a good chance of the percept deviating markedly from the true state.

An effective decision-maker should strive to minimise bias and uncertainty in the data on which they base their decisions. Here, we particularly consider the problem of uncertainty (noise). To preview our argument: 1. The data available for decision-making can include uncertainty (variability) for several reasons, and there is a key difference between intrinsic and extrinsic uncertainty. 2. Perceptual systems can be remarkably efficient at taking uncertainty into account, particularly via reliability-weighted combination of estimates, a noise-minimising strategy. 3. In contrast, people typically do not make such efficient decisions in tasks with uncertainty that is not perceptual. 4. Existing studies have explored ways in which perceptual and other tasks differ but have not clearly separated out the type of uncertainty alone as a factor. 5. We consider the possibility that “extrinsic” uncertainty, of the kind common to non-perceptual decision problems, is more difficult to account for than “intrinsic” uncertainty. To test this idea, we assess how people perform on a new spatial localisation task, in which all decisions are made within the perceptual domain, and all use the same stimuli, but we manipulate the degree of uncertainty from intrinsic (perceptual) vs extrinsic (environmental) sources. 6. Our finding is that perceptual judgments are indeed less effective in taking extrinsic uncertainty into account than intrinsic uncertainty. This illustrates that not all kinds of uncertainty are treated in the same way: internal uncertainty can in principle be accessible to perceptual systems, while external uncertainty needs to be learned.

1. Sources of uncertainty during perceptual decision making

Why are estimates uncertain (noisy, variable, limited in precision)? In the perceptual domain, limitations in sensory apparatus and computational imprecisions in the brain mean that sensory signals are subject to noise¹. In perception, this noise becomes particularly evident when fine discriminations are required: a person may struggle to correctly select the biggest of two similarly sized apples, or the shortest of two similarly long supermarket lines. This uncertainty greatly

increases when conditions for the sensory apparatus are degraded – for example, for a person choosing an apple in near darkness, or judging the supermarket line using eyes affected by progressive vision loss. We term this kind of uncertainty, related to variability (noise) in the sensory signal, *intrinsic* – in the sense that it is internal to the workings of a particular sensory system. We consider our example environmental (e.g. dark) or medical (e.g. vision loss) factors also to contribute to *intrinsic* uncertainty, as they lead to high variability (noise) in a particular sensory system. An ideal observer (or a different organism, or robot) equipped with better light-gathering apparatus would be subject to less noise during the nocturnal apple-selecting task¹. Intrinsic uncertainty reflects the variability (noise) within a particular sensory system, not variability (noise) in the environment.

In contrast, *extrinsic* uncertainty arises from variability in an environment, not in a sensory system. An example would be the variability in different lines' speeds at supermarket checkouts due to the differing behaviours of individual cashiers. Cashiers' speeds may vary as a function of their visuo-motor speed, their propensity to take time talking to customers, their familiarity with unusual items that otherwise need to be looked up, and so on. Excluding (in this example) reliable instantaneous perceptual cues to all the determinants of how fast a cashier is, the ideal observer could only establish relative speeds by observing the environment for some time. This process may *also* be limited by variability in the observer's perceptual processes – but crucially, even an ideal observer with negligible errors or computational limitations affecting these processes would still be subject to this externally determined uncertainty, and would need to spend time sampling the environment to build up an estimate of this source of variability - unless they had evolved, or been programmed to, have this information.

To sum up, we define intrinsic uncertainty as straightforwardly perceptual, in that it reflects the limitations of the perceptual process. Perception always comes with some uncertainty¹, but this uncertainty is particularly evident when making a fine discrimination and/or in degraded or sub-ideal conditions. In contrast, extrinsic uncertainty is outside of the perceptual process and reflects variability in the world itself. Examples include patterns of stochastic behaviour of agents or physical processes – e.g. the spread of droplets from a flow of water, the shooting accuracy of an archer, the likely waiting time for a bus. The statistics of these kinds of externally determined distributions can be learned, but the ideal observer needs to collect observations to learn them². Internal uncertainty

¹ Other properties of observers, such as computational limitations, can also distinguish them from ideal observers. For the present discussion, we focus just on the more elementary issue of the noise (variability) affecting observers' estimates.

² There are some interesting but limited cases in which evolution or (in artificial agents) programming could also provide an ideal observer with this information – mainly for simple physical processes, less obviously for

within a perceptual process itself could, in principle, be much more readily available to the perceiver.

2. Efficient mitigation of intrinsic uncertainty via combination of estimates during perceptual tasks

A useful strategy for reducing uncertainty is combining a perceptual estimate either with another estimate of the same property^{2,3}, or with prior knowledge of the statistics of the environment⁴. This has the effect of averaging-out random noise². In the perceptual domain, this process has been studied in *sensory cue combination* tasks, which measure people's abilities to reduce the uncertainty of a sensory estimate by combining it with other available estimates⁵. In a classic study⁶, people judged which of two bars was taller, using vision and/or touch. The ideal observer would obtain predictable reductions in their uncertainty given both estimates together vs either alone by reliability-weighted averaging – in which each estimate is weighted in inverse proportion to its uncertainty⁵. Participants followed this “optimal” strategy: they obtained the theoretically maximum uncertainty reduction when given the opportunity to combine cues, and they re-weighted (changed their relative reliance on) the visual cue to height as it was made more uncertain by addition of stereo noise⁶. Similarly, participants use prior statistical information to reduce their perceptual uncertainty, for example predicting visually noisy trajectories based on statistical distributions learned over the course of the experiment⁴. Observers often near-optimally combine noisy sensory estimates with each other^{6–9} and with prior distributions^{4,10–12}, although there are also cases of suboptimal combination^{13–15}. Perceptual cue combination is already evident in single neurons of early sensory areas during multisensory tasks¹⁶, and can also be detected using fMRI in humans, in sensory areas (even during passive viewing, in the case of combination of visual cues to depth)^{17–19} and across a whole cortical hierarchy up to the frontal lobe during perceptual decision tasks^{20–22}.

For this kind of reliability-weighting to work, the perceptual system must correctly account for its own uncertainty. For example, it must adjust the weighting given to a visual cue to bar height, vs a haptic one, in line with more or less stereo noise being added⁶. Similarly, during visual-auditory localisation⁷, the observer must adjust the weighting given to a visual cue to location, vs an auditory one, in line with progressive visual blurring. How uncertainty-weighted averaging is implemented at a cells-and-circuits level is still being established^{23–25}, but it is clear that uncertainty intrinsic to perceptual estimates can be represented implicitly by population responses, because elementary perceptual properties tend to be represented by populations of tuned cells. Thus, a very precise

specific cases – e.g. the waiting times on this bus route in particular, the shooting accuracy of this particular archer.

estimate of a line's orientation might strongly engage a relatively small population of cells tuned to a specific orientation, while an uncertain estimate of orientation, e.g. in a more blurred stimulus, would more weakly engage a broader population. The distribution of population responses can therefore reflect the level of intrinsic uncertainty. Consistent with this, widths of neuronal probability distributions decoded from participants' fMRI activity are predictive of their perceptual decisions^{26,27}.

In summary, perceptual tasks show abilities to account effectively for perceptual uncertainty. When sensory estimates are combined⁶⁻⁹, all the uncertainty is within the perceptual system itself (tasks combining sensory information and priors^{4,10-12} have the additional challenge of learning and representing the prior). It is not obvious that perceptual systems should necessarily accurately read out or represent the variabilities in their own estimates, but the evidence from cue combination tasks⁶⁻⁸ suggests that this must be the case. Considering how perceptual information is represented by neuronal populations²³⁻²⁷ also provides plausible mechanisms for representing and reading out this variability information within the perceptual system itself, without the need to additionally learn about this variability.

3. Less efficient mitigation of uncertainty during non-perceptual decision tasks

The literature on economic decision making under uncertainty stands in stark contrast to these findings from perceptual tasks, documenting many inefficient and seemingly sub-optimal patterns of decision making²⁸⁻³¹. By "classical" decision making, we mean decisions about verbally or symbolically presented options. One example of this type of problem is a lottery, such as "would you rather gain \$3000 for sure or have an 80% chance of gaining \$4000?"²⁹. A second is a decision informed by a *base rate* (similar to a *prior* in perceptual tasks), for example: how likely it is that a blue vs green taxi was responsible for a traffic accident, knowing that 85% of taxis in the city are green and 15% blue, and that the witness (who judged that it was blue) is only correct at making this discrimination 80% of the time? (the "taxicab problem"³⁰). A third kind is a multiple-factor combination judgment, such as deciding how toxic a fictitious bug is likely to be based on multiple predictive cues (e.g. leg length, colour)³¹.

Perhaps most pertinent to the questions asked in this study is the issue of weighting of experience vs externally provided statistics. A veteran physician will have accumulated years of experience of diagnosing common diseases, but may have no experience with very rare diseases. A possible problem is that they may rely too heavily on their experience, and under-value the potential of diseases not previously diagnosed, possibly due to the Representativeness heuristic³². (This is in

contrast to a more junior physician who might over-represent the possibility of diseases that are very rare, an example of Base rate neglect)³³. Phrased in terms of Bayesian inference, the experienced diagnostician might put too little weight on a likelihood of a disease that is verbally / symbolically presented to them, instead putting too much weight on their prior experience. Formally, this can also be used to describe the cause of stereotyping³⁴.

Tasks and situations like these are non-perceptual in the sense that the information (e.g. words, numbers, bug features, disease prevalences) is symbolic and not subject to relevant perceptual uncertainty. They can also differ from the perceptual tasks described above in the specific information integration problem (in the cue combination examples and in the taxi-cab and bug problems, multiplying likelihoods, probability distributions or probabilities; in the lottery example, combining posterior probability and value). Our aim here is to focus specifically on the source and representation of uncertainty, and to bring these kinds of problems closer, leading us to introduce a novel experimental task in which every aspect of the task except for the type of uncertainty is matched. In order to reach that point, we first review previous studies in which aspects of perceptual and non-perceptual decision tasks have been matched and compared in significant ways.

4. Task differences and uncertainty differences

Studies have directly compared people's abilities to deal with internal estimates of uncertainty on the one hand, and explicitly stated probability information on the other. Some studies measured participants' visuomotor precision at a speeded pointing task^{35,36}, allowing them to calibrate the visual stimuli so that the same verbal lottery could equivalently be presented as a visuo-motor one. In the verbal task, participants had to choose which of two numerically expressed probabilistic scenarios they would prefer (of the general type $x\%$ chance of winning A vs $y\%$ chance of winning B), while in the visuo-motor task, they made equivalent decisions by choosing which of two visual targets they would prefer to attempt to hit for their respective rewards. One study found different biases towards risk-seeking and different weighting of probability and value across these tasks³⁵; another found broadly comparable performance³⁶. This general approach compares intrinsic (visuo-motor) and extrinsic (stochastic) uncertainty, but there are other differences in the two task types: visuo-motor uncertainty is presented perceptually by a bar width, while extrinsic (stochastic) uncertainty is given via the classical numeric (symbolic) route. Therefore, differences in behaviour may depend in part on the origin and nature of the uncertainty (internal vs external) and in part on the manner of presentation (perceptual vs symbolic).

A study wholly within the perceptual domain manipulated an array of oriented gratings in two ways³⁷: by changes in the stimulus contrast, and changes in the variability of orientations across the different gratings in the array. Expressing this in our framework, reducing contrast increases intrinsic (perceptual) uncertainty, while increasing variability increases the difficulty of the information integration problem of correctly averaging together disparate estimates. This study did not have “extrinsic” uncertainty (nothing is stochastic: an ideal observer without limitations in perceptual precision or computational abilities could deal with both the contrast and the variability manipulations), but is related by the authors to cognitive decision-making in an interesting way. They found that participants took contrast-related (intrinsic) uncertainty near-optimally into account, but failed to do so for variability-related (integration-demanding) uncertainty, and were well described by a model blind to this source of noise. The authors propose that the “optimality gap” in perceptual vs classical cognitive decision tasks may be explained in part by a blindness to noise introduced during the integration process that combines multiple information sources. The argument is that even perceptually presented estimates (low-contrast grating orientations), when combined in a non-standard way for the sensory system (averaging multiple gratings in an array), become subject to cognitive-like decision errors. Presumably, the contrast is between this type of novel averaging and the highly familiar averaging during cue combination where multiple sensory cues are associated with a single redundant property, such as size⁶ or location⁷. This study interestingly bridges perceptual uncertainty and the integration/decision process, suggesting a privileged role for intrinsic (perceptual) uncertainty and a relative blindness to noise that affects a computational (averaging) process.

A recent study took a different approach to comparing how people presented with cognitive inference tasks approach the sort of precision weighting of estimates common in perceptual cue combination³⁸. Participants performed three different cognitive decision tasks that require combining prior information with new data, and a model-based analysis was used to categorise them according to the extent to which they followed different principles underlying optimal inference. Just over half of participants appreciated the need to consider the prior as well as the data, but only a quarter appreciated the need to weight them according to their uncertainties. While there is no direct comparison with the perceptual domains within this study, and perceptual studies rarely categorise individual subjects in this way (though see Ref³⁹), this very widespread blindness to uncertainty weighting during cognitive tasks is at odds with its usual strong implementation during perceptual tasks^{4,6-12}.

In summary, while perceptual decision making in the face of perceptual (intrinsic) uncertainty is often highly efficient in dealing with this uncertainty^{4,6-12}, this is not typically seen in classical or cognitive decision making^{28-31,35,38}. A crucial question is to what extent different systems and representations underlie decision making in these domains. However, these domains differ across many dimensions. Examples of ways in which studies have explored these differences include whether probabilistic information is presented perceptually or symbolically^{35,36}, how much uncertainty is at a perceptual vs at an integration stage³⁷, and which models best explain precision-weighting behaviour during cognitive tasks³⁸.

Behaviour during decision-making under uncertainty may differ for many of these and other reasons, but here we return to the key distinction introduced at the start: the source and nature of the uncertainty itself. It has been widely appreciated that internal uncertainty, due to variability within the perceptual system¹, and potentially amenable to being read out by the system itself²³⁻²⁷, has advantages compared with external uncertainty, which depends on stochastic events outside of the observer. It has also been appreciated that a privileged access to internal uncertainty may underlie many perceptual-cognitive task differences. However, tasks probing these differences tend not to allow for a clean interpretation of the role of type of uncertainty alone in judgments – since these also present the information in different ways^{35,36}, or do not compare these two types of uncertainty^{37,38}.

5. Explaining task differences: the present study

It may be that different systems and representations underlie decision making across perceptual and cognitive domains and that the relative engagement of these depends on a range of task, informational, and perceptual parameters. However, here we consider one crucial factor that has not been clearly distinguished and tested: the nature of the uncertainty involved, intrinsic vs extrinsic.

In the present study, we directly compare abilities to account for intrinsic vs extrinsic uncertainty during a perceptual decision-making task. Crucially, the task and even the stimuli are exactly matched across conditions, so that uncertainty is signalled in the same way (via spread of dots), and decisions and responses are made about the same property (location) and entered in the same way (clicking to indicate a screen location). We use a multisensory localisation task, in which visual (dot cloud) and auditory (white noise on a co-localised multi-speaker setup) provide redundant cues to the property to be estimated: location. We draw the dot clouds such that we calibrate the amount

of noise that is intrinsic (reflecting sensory imprecision at determining the centre of the cloud³) vs extrinsic (reflecting unreliability of the cloud centre as a guide to the target location). The colours of the dot clouds, together with a cover story (that they are guesses by two different players) are the cue to which levels of intrinsic and extrinsic noise underly a given trial. Since intrinsic uncertainty is available to the perceptual system itself, while extrinsic must be learned with experience, we predict that decision-making in the intrinsic case will show better following of signal reliabilities.

6. Preview of findings

We directly compared abilities to account for intrinsic vs extrinsic uncertainty during a perceptual decision-making task. Our main finding was near-optimal weighting of sensory estimates when uncertainty was only intrinsic, but mis-weighting of estimates when extrinsic uncertainty was added. Specifically, participants over-relied on an extrinsically uncertain signal, suggesting that they are to some degree “blind” to this source of uncertainty. This illustrates that not all kinds of uncertainty are treated in the same way: internal uncertainty may be accessible to the perceptual system, while external uncertainty needs to be learned.

Methods

Thirty participants (26 Female, min/median/max age 18/24/34 years, age information for one participant was lost) completed a series of trials where they used intrinsic+extrinsic visual cues, intrinsic-only visual cues, and intrinsic-only auditory cues to estimate the location of a hidden target on a projector screen.

Participants were recruited from a combination of undergraduate students in Psychology department and social media, and were compensated for their time (£10 per hour plus performance bonus). Only restrictions were corrected or uncorrected normal vision and normal hearing.

³ ³⁷We consider localising the centre of a dot-cloud to be an elementary spatial localisation task (similar task to finding the centre of a dim blob⁷). Dot clouds allow presentation of spatially-uncertain stimuli to be operationalised in a highly controllable way, and multisensory combination tasks using dot clouds are well fit by standard cue combination models^{20,40}.

Participants were fully briefed about the purpose of the experiment and gave informed consent, with ethical approval given by Durham Psychology department.

Participants were placed 140 cm in front of a projector screen (width 235cm, height 131cm), with all visual stimuli projected onto it (Optoma GT1080E). Behind the screen were placed 9 speakers (Visaton SC 5.9) at a separation of 2.5 visual degrees.

As a cover story, participants were told that they were taking part in a task similar to a fairground game, where they had to find a hidden object using visual cues of previous players, possibly together with the rustling sound made by the person placing the hidden object (see supplementary material for the full instructions). For the two types of visual stimuli (intrinsic-only and intrinsic+extrinsic, see below) participants were told the cues were guesses by two different individual players, “Rodger” and “June”, signalled by different coloured dots (blue or orange). The exact instructions given to participants are included in the Supplementary materials.

The experiment began with a visual cue calibration block where participants estimated the location of a hidden target using low or high variance intrinsic-only visual cues (4-dot clouds, each dot shown in succession for 100ms, generated by shifting and scaling the four dot centres drawn from a standard normal distribution so that the mean was exactly the true location, and the standard deviation (SD) was fixed at 0.05 or 0.2, respectively; stimulus parameters are defined as proportions of total screen width such that zero maps to the left of the screen and one to the right). There were 90 interleaved trials for each cue (10 repeats for 9 test locations; approximately evenly spaced points between 0.37 and 0.63). We interpret dot clouds generated in this way as being corrupted only by intrinsic noise as variability in location estimates made using these cues is due to noise that is intrinsic to the observer, such as sensory noise. Put differently, an ideal observer who is free from sensory noise, computational imprecision, memory imperfections, response noise, etc., could provide a perfect estimate on every trial (Figure 1A).

The purpose of the initial visual cue calibration block was to measure perceptual variance using each of the low and high variance intrinsic-only visual cues (σ_L^2 and σ_H^2 , respectively) so that the difference in variance using each of the two cues ($\sigma_E^2 = \sigma_H^2 - \sigma_L^2$) could be used as the variance of the extrinsic noise distribution. The extrinsic noise distribution was used to generate random shifts ($z_i \sim N(0, \sigma_E^2)$), where z_i is independently drawn for every appearance of each intrinsic+extrinsic sensory cue) of all dots in a cloud (same shift for each dot in a single cloud, or equivalently, only one z_i per trial) away from the true hidden location. For the remainder of the experiment, extrinsic noise generated in this way was only added to the low intrinsic-only cue to create a single intrinsic+extrinsic visual cue. Adding noise in this way, with the level of noise calibrated individually

for each participant, should lead to participants being equally variable using the intrinsic+extrinsic visual cue and the high intrinsic-only visual cue (referred to hereafter as the intrinsic-only visual cue). Indeed, that was our intention to allow for a direct comparison of the weight placed on each cue.

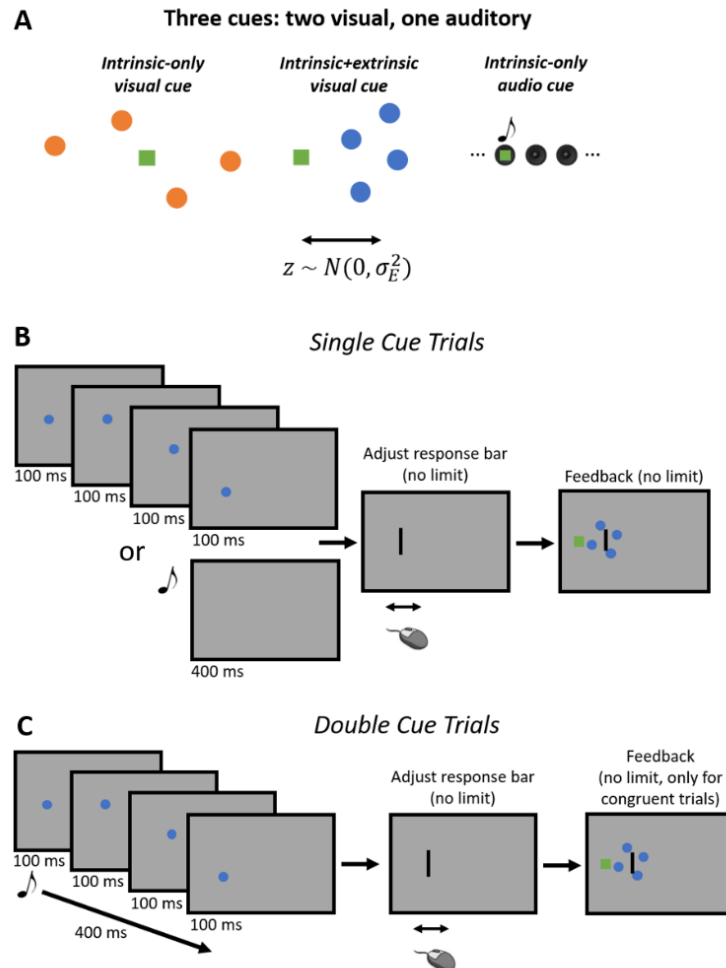


Figure 1: Figure 1: Schematics of experiment. A) Subjects were presented with three different types of cues: intrinsic-only visual, intrinsic+extrinsic visual, intrinsic-only auditory. The total variance of the intrinsic+extrinsic cue (blue) was matched to the intrinsic only visual cue (orange) by adding noise to the mean of the four dots. B) Subjects on some trials were presented with single cue stimuli of either the visual or auditory cue. Dots were presented sequentially, while the auditory cue was continuous for 400 ms. Feedback was provided for all single cue trials. C) For On other trials, subjects could be presented with two cues (visual and auditory). The cues were identical to the single cue trials, but feedback was only provided on trials where visual and auditory cues were congruent (i.e. from the same location)..

The visual cue calibration block was followed by the audio cue training block, where participants used an intrinsic-only auditory cue to estimate the location of the hidden object. The intrinsic-only auditory cue was a 400ms burst of white noise from one of nine speakers located behind the projector screen presenting the dot clouds. The locations corresponded to the locations used in the

calibration block and we tested each location 10 times for a total of 90 interleaved trials in this block. We say this auditory cue is only intrinsically uncertain as an ideal observer, who perfectly localises sounds, would choose the exact location of the hidden object on every trial. Thus, variability in the responses is intrinsic to the observer, again generated by factors such as sensory noise, computational imprecision, and response noise.

The audio cue training block was followed by the test block. In the test block of the experiment, participants used either the intrinsic+extrinsic visual cue, intrinsic-only visual cue, or an intrinsic-only auditory cue alone to estimate the hidden object location (single cue trials; Figure 1B), or one of the visual cues paired with the auditory cue (double cue trials; Figure 1C). In total, there were 405 single cue trials (Figure 1B) in the test block, made up of 15 repeats of each location used in the calibration block for each single cue (intrinsic+extrinsic visual cue, intrinsic-only visual cue, and intrinsic-only auditory cue). The single cue trials were interleaved with congruent and incongruent double cue trials (Figure 1C). The intrinsic-only and intrinsic+extrinsic trial types were randomly interleaved. In congruent double cue trials, each visual cue could be presented simultaneously with the auditory cue with both cues corresponding to the same hidden location. There were 270 congruent double cue trials made up of 15 repeats of each location used in the calibration block for each audio-visual cue pairing. Incongruent double cue trials were like the congruent cue trials, with each visual cue being presented simultaneously with the auditory cue, except that in these trials, rather than indicating the same location, these cues were in conflict, allowing us to estimate the amount of weight that participants placed on each visual cue when paired with the audio cue. There were 400 incongruent double cue trials, made up of 10 repeats of all combinations of intermittent calibration locations (0.37, 0.43, 0.5, 0.57, and 0.63) for each audio-visual cue pairing.

Due to an error in the experimental code for the incongruent conditions of the test block, only stimuli in the left hemisphere were presented. To ensure there were no left-right biases which could affect our results we compared variable error in Left vs Right hemisphere for the Intrinsic-only, as well as Intrinsic+External, congruent conditions and found no significant difference (Wilcoxon signed-rank test, $p=0.8555$, median difference $=-0.0037$ for congruent Intrinsic only, $p=0.8555$, median difference $=0.0011$ for congruent Intrinsic plus Extrinsic). Encouraged by this we continued with the originally planned analysis (a post-hoc analysis for the incongruent data that allowed the central tendency bias to be shifted away from the screen center found no differences, see Supplementary Material).

Participants received feedback on every trial in the visual cue calibration block, the audio cue training block, and on single cue and congruent double cue trials in the test block. Feedback

consisted of a green circle or square that indicated the location of the hidden object, presentation of all dots from the dot cloud simultaneously on any trial involving a visual cue, and a score out of 1000. Feedback showing the object location alongside all four dots was the means by which participants could, in principle, learn the magnitude of the extrinsic noise on intrinsic+extrinsic trials (noting that two different signals, shown by differently coloured dots, indicated two individual named ‘players’ – see participant instructions above and in Supplement).

Scores followed a squared error loss function so that the score reduced quadratically with increasing distance from the target until it reached zero. The function was scaled so that the function would reach 0 points at a distance from the target of 10 percent of the screen width (0.1). With this modification, participants only score if the absolute difference between the hidden and guessed locations is less than 0.1. This may seem small, but as the experiment was completed on a projector screen of width 190cm, this means participants only received points if their guess was within 19cm of the true location on the screen. The formula for calculating the score was:

$$s = 1000(1 - ((h - g)/0.1)^2),$$

where s is the score, h is the hidden location, and g is the guessed location.

Calculating perceptual variance

During the experiment, we calculated variance using each of the low and high variance intrinsic-only visual cues in the calibration block as the standard deviation of all errors (response – target location). These values (σ_L^2 and σ_H^2) were used to define the variance of the extrinsic noise individually for each participant ($\sigma_E^2 = \sigma_H^2 - \sigma_L^2$) that was added to the low intrinsic-only visual cue to create the intrinsic+extrinsic visual cue. This was calculated online in the experiment immediately following the calibration block so that test trials could be generated. As mentioned above, adding extrinsic noise in this way should lead to equal variance when estimating the location of the hidden object using either the intrinsic+extrinsic cue or the intrinsic-only (high noise) cue in the test block of the experiment.

Since collecting these data, we have become aware of an important issue with variance estimates from data with continuous responses, such as those here: a frequent finding of central tendency biases^{41–43}, where participants bias their responses towards the middle of the set of presented stimuli. Such biases lead to an under-estimation of variable error. We have recently described a method that recovers corrected estimates of variable error⁴⁴ in this situation. As described in the

Supplement and Ref. ⁴⁴, we apply this method to recovering correct variable errors in the present data set. Unfortunately, the variance calculations that were done online to calibrate stimuli within the experiment itself did not use this correction, and as a result we did not match variable error using the intrinsic-only and intrinsic+extrinsic visual cues in the test block as intended. Matching the conditions would allow us to have a simple non-model based comparison, and while the difference in variance makes the two conditions a little less comparable than intended, it does not prevent us from comparing the extent to which participants weighted stimuli appropriately in response to the two kinds of uncertainty using our model-based analysis. In short, given the corrected uni-modal variances it is straightforward to calculate the expected visual weighting and compare to the empirical measured weight.

Summary and predictions

In summary, we asked participants to localise a hidden object using visual (dot-cloud) and/or auditory (white noise) cues. The object's location was uncertain because of either only intrinsic noise (visual or auditory), or because of extrinsic as well as intrinsic noise. Extrinsic noise was implemented as an additional offset of cue dots vs the true location. Through feedback, participants had an opportunity to observe and learn about the level of this extrinsic noise.

Using trials with minor offsets in cue positions, we measured the relative weighting for (reliance on) visual vs auditory cues during perceptual judgments. We predicted that, in line with previous sensory cue combination studies⁶⁻⁹, participants would not weight the visual cue with only intrinsic noise differently to the optimal (reliability-weighted) prediction. We predicted that, in contrast, participants would overweight the intrinsic+extrinsic cue, as they would be less sensitive to the added extrinsic uncertainty.

All stimuli, data and code are available at <https://osf.io/6paq9/>

Results

Participants overweight the visual cue with extrinsic uncertainty

We used estimates of variable error using the intrinsic-only, intrinsic+extrinsic, and auditory cues alone in the test block (σ_I , σ_{IE} , and σ_A) to calculate the optimal, reliability based, weight that participants should place on each cue in a pair. The formulas we used to calculate the optimal weight to place on the intrinsic-only cue (w_I) and the intrinsic+extrinsic cue (w_{IE}) when either cue is paired with the auditory cue are given below.

$$w_I = \frac{\sigma_A^2}{\sigma_I^2 + \sigma_A^2} \quad w_{IE} = \frac{\sigma_A^2}{\sigma_{IE}^2 + \sigma_A^2}$$

This formulation is reliant on small discrepancies between visual and auditory stimuli (less than 10 visual degrees), for large discrepancies subjects might instead rely on a causal inference model^{40,45}.

The empirical weights were determined by modelling responses to the incongruent trials as

$$r = (1 - \hat{w}_P)(\hat{w}_V x_V + \hat{w}_A x_A) + 0.5\hat{w}_P + \epsilon,$$

where \hat{w}_P estimates the strength of the central tendency bias (for this cue pairing),

\hat{w}_V estimates the weight placed on the visual cue (the intrinsic-only or intrinsic+extrinsic cue: \hat{w}_I or \hat{w}_{IE}),

$\hat{w}_A = (1 - \hat{w}_V)$ estimates the weight placed on the auditory cue,

x_V is the centroid of the dot cloud (the same as the visual cue source location for the intrinsic-only cue but not for the intrinsic+extrinsic cue),

x_A is the source location of the auditory cue, and

$\epsilon \sim N(0, \sigma_n^2)$ is a noise term.

To allow for the fact that we assume subjects receive a noisy (uncertain) input, the perceptual cues had noise with variance of σ_V^2 and σ_A^2 , with $\hat{w}_V = \sigma_A^2 / (\sigma_A^2 + \sigma_V^2)$. A participant with a veridical estimate of their own relative perceptual uncertainties would thus correctly set their variances based on these. For a linear model this has no effect on the mean estimates, but can affect estimates for more complicated models.

We fit this model separately for each participant to all conflict trial responses for combinations of the intrinsic-only and auditory cue, and intrinsic+extrinsic and auditory cue separately. The model was fit using JAGS⁴⁶ via the MATLAB-to-JAGS interface matjags.m to estimate posterior probability distributions for the free parameters \hat{w}_P , \hat{w}_V , and σ_n .

We ran three independent chains, discarding the first 1000 samples of each chain as burn-in, and recording 4000 samples after the burn-in period, thinned by recording only every 5th sample. Both fitted weights (\hat{w}_P and \hat{w}_V) were initialised at 0.5 in all chains. The standard deviation of the noise (σ_n) was initialised at 0.01. The priors on w_V and w_P were uniform distributions between 0 and 1.

The prior on σ_n was a uniform distribution between 0.001 and 0.2. The final parameter estimates were taken as the mean of the expected values for each chain.

Figure 2A shows that the weight participants placed on the intrinsic-only cue relative to the auditory cue is positively correlated with the optimal prediction ($r = .536, p = .002$), suggesting participants weight the cues according to their reliabilities. Figure 2D shows that empirical and optimal weights do not differ significantly ($t(29) = -1.74, p = .092$), although the Bayes factor in favour of the null is approximately 1, suggesting both the null and alternative hypotheses are equally good explanations of the data ($BF_{01} = 1.34$).

Figure 2B shows that the weight participants placed on the intrinsic+extrinsic cue relative to the auditory cue is also positively correlated with the optimal prediction ($r = .524, p = .003$), again suggesting participants weight the cues according to their reliabilities. However, most points are above the identity line (empirical weight greater than optimal). Figure 2E shows that mean empirical and optimal weights differ significantly ($t(29) = -5.14, p < .001$), with participants overweighting the intrinsic+extrinsic visual cue. The Bayes factor in favour of the alternative suggests there is extreme evidence for a difference between empirical and optimal weights ($BF_{10} = 1291.3$).

We hypothesised that participants would overweight the intrinsic+extrinsic cue as they would be insensitive to the added extrinsic uncertainty. If participants are completely insensitive to the extrinsic uncertainty, they should weight the intrinsic+extrinsic cue according to the variability of their responses when using the low intrinsic-only cue measured in the calibration block. We call the use of this insensitive strategy the suboptimal intrinsic+extrinsic cue weight prediction. Figure 2C shows that empirical weights are correlated with the suboptimal insensitive predictions ($r = .482, p = .007$), but the position of most points below the identity line and Figure 2F show that empirical weights differed significantly from the suboptimal insensitive predictions ($t(29) = 4.16, p < .001$). The Bayes factor in favour of the alternative suggests there is extreme evidence for a difference between empirical and suboptimal weights ($BF_{10} = 110.7$). This suggests that participants are not completely insensitive to the extrinsic uncertainty, but only partially account for it.

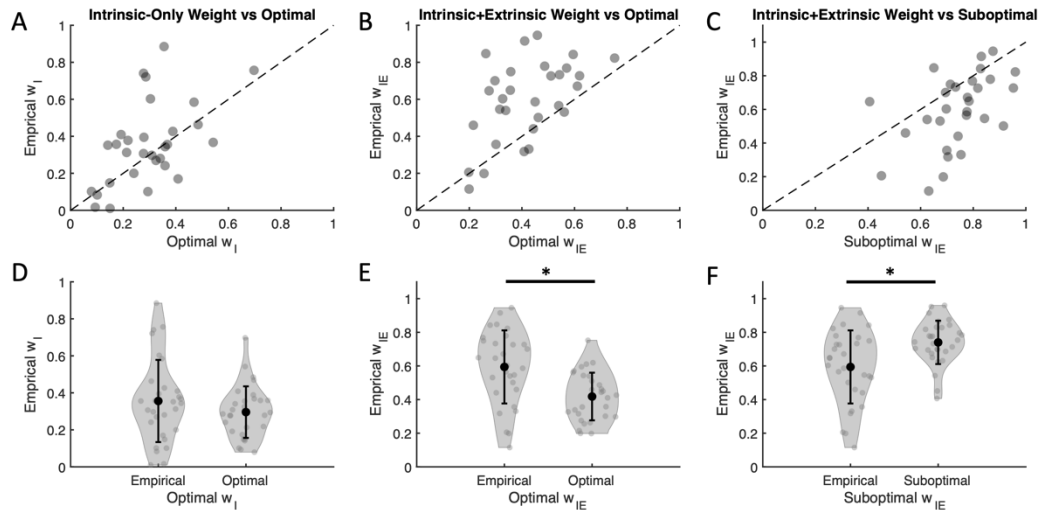


Figure 2: Comparison between visual weights for Intrinsic-only (A) and Intrinsic-extrinsic (B-C), empirically measured weights against either optimal weights (A,B) and suboptimal insensitive weights (C). While both Intrinsic-only and Intrinsic+Extrinsic empirical weights correlated with the optimal weights, the visual weights for Intrinsic-only were not significantly higher than optimal (D), whereas the weights for Intrinsic+Extrinsic were significantly higher than optimal (E). Intrinsic+Extrinsic visual weights were also significantly lower than weights from a suboptimal model that ignores the Extrinsic variability (F).

Individual differences

An advantage of running a Bayesian model with JAGS is that it provides uncertainties around the variable fits, unlike a Maximum Likelihood approach which just gives the best fitted parameter value. We can therefore also examine how well individual subject parameters were fitted, e.g. the Bayesian Confidence Intervals of parameters such as the visual weight. Using the JAGS output we calculated the Bayesian Confidence interval for the visual weight, finding that for the intrinsic-only data set for 18 out of 30 participants the optimal weight was outside their 95 percent confidence intervals, while for the intrinsic+extrinsic data set the proportion was 22 of 30 participants (see Fig. 3).

To explore further what participants did we also compared the empirically fitted weight values to those of a model participant who only used their best modality (whether visual or auditory). Only 2 of the participants in the intrinsic-only condition had a Bayesian confidence interval that included the possibility of just using their best modality, and none of the participants did in the intrinsic+extrinsic condition. For all the remaining participants (28 and 30) we could therefore rule out that they only used their best modality.

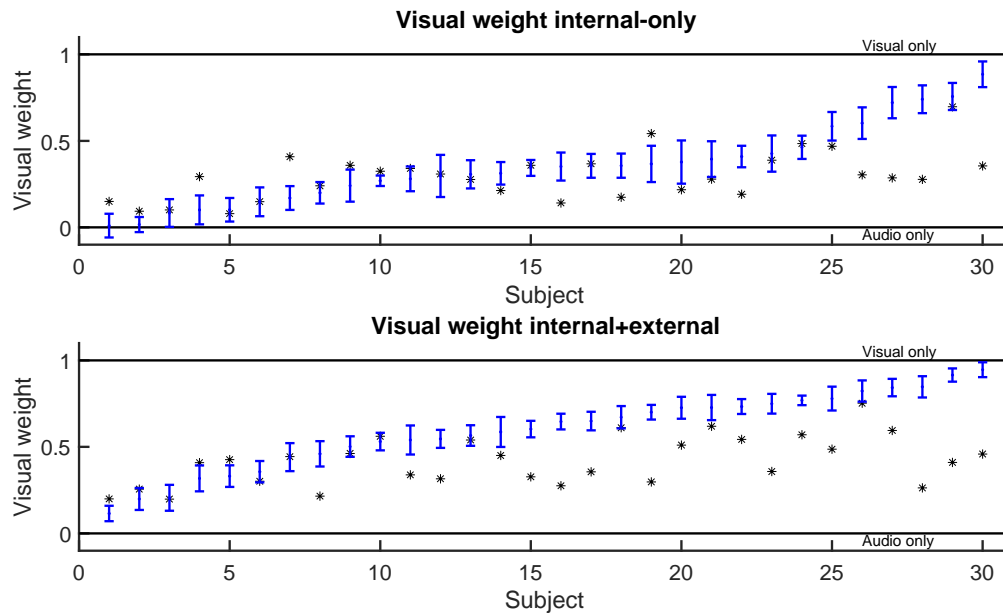


Figure 3: Individual visual weights, with 2σ standard deviation error bars, as reported from JAGS Bayesian fit to the intrinsic-only (above) and intrinsic+extrinsic (below) data sets. Gray asterisk indicated the optimal weighting. A value of 0 indicates all weight on audio, a value of 1 that all weight is on visual.

Switching

One possible reason that participants could attain a weight on the visual cue that would be indistinguishable from optimal would be to perform a random switching between reporting the noisy percept from the two modalities, at the right ratio but without actually combining them (also referred to as probability matching). Such behaviour has previously been found in perceptual cue integration experiments (see e.g. Ref. ⁴⁷). To rule out this idea we ran a variant of the model above that included an individually fitted probability of the participants performing switching instead of integration. We found a low average probability of subjects using a switching strategy of 0.195 for intrinsic-only and 0.212 for intrinsic+extrinsic. These values were not significantly different (ranksum test $p=0.631$), hence there was no evidence that the data sets differed based on propensity to use switching over integration. Overall, this implies that there is a small likelihood that participants were using a switching strategy and that it was more likely that were indeed combining stimuli based on the optimal weights for the intrinsic-only data set (but not for intrinsic+extrinsic).

Variable Error

An advantage of using Bayesian inference is the reduction in error, hence if we expected participants to be using optimal weighting there should be an advantage in error. However in practice these error reductions are only found when visual and auditory stimuli are calibrated to have well matched

error levels, otherwise the ideal observer predicted advantage may differ too little from the best single cue to be measurable in noisy responses^{48,49}. Unfortunately, as we did not foresee the issue of central tendency bias (see Methods above, Ref. ⁴⁴, and Supplement), we did not match variable error using the intrinsic-only and intrinsic+extrinsic visual cues in the test block as intended (see Supplement).

The Variable Error for the intrinsic-only data set (calculated for congruent audio and visual stimuli, unlike the calculation above for visual weight based on incongruent stimuli) did not show a significant improvement over the variable error for the single best modality for each participant (Wilcoxon signed-rank test, $p=0.371$, median=-0.002). Non-intuitively, there was however a small but significant advantage of using intrinsic+extrinsic visual together with audio ($p=0.032$, median=0.004), likely due to the better matching of the visual error with the audio error. The unintended higher uncertainty in the intrinsic only than the intrinsic+extrinsic condition meant that noise in this condition was less well matched to the audio noise, leading to a lower potential gain. See the Supplementary material for full details of this analysis.

Discussion

We introduced a distinction between intrinsic and extrinsic uncertainty during decision-making. While perceptual and other decision tasks can vary in many ways, a key difference is often the kind of uncertainty involved. Perceptual decision-making usually involves uncertainty that is intrinsic to the perceptual system itself, and to which it may therefore have good access²³⁻²⁷. In line with this, the intrinsic uncertainty in perceptual tasks is often efficiently mitigated by reliability-weighted combination⁶⁻⁹. In contrast, when uncertainty arises in processes external to the observer, it must be learned in order to be accounted for. While this difference may go some way towards accounting for differences in how people deal with uncertainty across decision tasks, those tasks typically have many other differences too. Therefore, in the present study, we developed a perceptual task to make this specific comparison directly. To do so, we manipulated levels of intrinsic vs extrinsic uncertainty while keeping the task and stimuli the same across these manipulations. We tested the extent to which decision-makers take uncertainty into account when using multiple sensory estimates for a localisation task, and tested the prediction that participants would follow the reliabilities of estimates better with intrinsic-only than with intrinsic+extrinsic uncertainty.

As predicted, we found near-optimal (or rather, not statistically different from optimal) weighting of sensory estimates when uncertainty was only intrinsic, but mis-weighting of estimates when extrinsic uncertainty was added. Specifically, participants over-relied on an extrinsically uncertain signal, suggesting that they are to some degree “blind” to this source of uncertainty. However, they relied on it more than the prediction for participants completely blind to the extrinsic uncertainty. That is, participants fully accounted for uncertainty when it was purely intrinsic, but only partially accounted for it when it was also extrinsic. These results are in line with the observation that intrinsic uncertainty within a perceptual process can in principle be read out from the estimate’s neural representation^{23–27}, while extrinsic uncertainty as reflected by the statistics of how an external process or agent behaves needs to be learned – either from an external information source^{29–31}, or (as here), by directly learning from experience.

These results also add to a growing literature on the description-experience gap⁵⁰, which shows how information that is not directly experienced, but based on abstract symbolic representation, is not correctly weighted within a Bayesian framework. This account of the issue of an experienced physician under-weighting the likelihood of a disease they have not personally encountered (see Introduction) would propose that they do not accurately represent the uncertainty of this information, as it was not experienced but merely described to them. In our task, the new information about external uncertainty was experienced, but was presented and learned in a manner different to internal noise. This provides a clue to which factors are important for the correct Bayesian inference: an insight is that not only whether the information is experienced, but the manner in which it is experienced, may be important.

Our task was designed to allow participants to learn about the extrinsic uncertainty. We explained the different trial types with a cover story that suggests a model for different kinds of uncertainty across conditions consistently denoted by differently coloured stimuli, and participants received feedback through which they could learn about the uncertainties of the two different cues. Our results showed that there was some learning of the extrinsic uncertainty: participants were not completely blind to it, as would be the case if they had not learned about it. An open question is the time scale over which observers might or might not perfectly learn and/or use this kind of uncertainty during decision-making: given much longer experience with stimuli like these, how much closer would performance come to optimal? Even with perfect learning of the external variability, would people still account for this differently to internal variability in their perceptual decisions? How does this generalise to other tasks and sources of external variability / uncertainty? This calls for longer studies, supported by ideal observer models of the process of learning about uncertainty based on the feedback given.

A key benefit of weighted combination of estimates is a reduction in variance. An unexpected result of the present study is that participants in the intrinsic uncertainty only condition did not reduce their variance, even though they used near-optimal weights. Our model-based analysis excludes the interpretation that they were switching rather than averaging estimates, so we believe that the most likely explanation for the absence of this finding is that our unplanned mis-matching of reliabilities across intrinsic and intrinsic+extrinsic uncertainty conditions. As with all cue combination studies, attending to this issue is important to reliably detect variance reductions^{48,49}. Using continuous response (as compared with alternative forced choice) methods, as here, has advantages for collecting highly informative measurements, but should in future be checked and corrected for central tendency biases in order to accurately recover measurements of variability⁴⁴.

Conclusions

Decision making under uncertainty varies in many ways across real-life settings and experimental tasks. We have highlighted one under-studied but crucial issue: the origin of the uncertainty itself, intrinsic or extrinsic. With a novel perceptual task designed to manipulate this factor alone, we show that while participants account well for intrinsic uncertainty, they under-account for extrinsic. The time scale for learning about this kind of uncertainty, and the generality of this finding for other tasks, settings, and sources of external uncertainty, are open questions for further research.

Acknowledgments

This work was supported by Leverhulme Trust Research Project Grant (RPG-2017-097) and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 820185). We are grateful to Artie Graham for help with testing.

References

1. Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nature Reviews Neuroscience* 2008 9:4 2008; 9: 292–303.
2. Clark JJ, Yuille AL. Data Fusion for Sensory Information Processing Systems. 1990.
3. Landy MS, Maloney LT, Johnston EB, et al. Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research* 1995; 35: 389–412.
4. Körding KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature* 2004; 427: 244–7.
5. Ernst MO. A Bayesian view on multimodal cue integration. *Human body perception from the inside out* 2006; 131: 105–131.
6. Ernst MO, Banks MS. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 2002; 415: 429–33.
7. Alais D, Burr D. The ventriloquist effect results from near-optimal bimodal integration. *Curr Biol* 2004; 14: 257–62.
8. Hillis JM, Watt SJ, Landy MS, et al. Slant from texture and disparity cues: optimal cue combination. *Journal of vision* 2004; 4: 967–992.
9. van Beers RJ, Sittig AC, Denier van der Gon JJ. How humans combine simultaneous proprioceptive and visual position information. *Exp Brain Res* 1996; 111: 253–61.
10. Bejjanki VR, Knill DC, Aslin RN. Learning and inference using complex generative models in a spatial localization task. 2016; 16: 1–13.
11. Berniker M, Voss M, Kording K. Learning Priors for Bayesian Computations in the Nervous System. *PLoS One* 2010; 5: e12686.
12. Tassinari H, Hudson TE, Landy MS. Combining Priors and Noisy Visual Cues in a Rapid Pointing Task. *Journal of Neuroscience* 2006; 26: 10154–10163.
13. Battaglia PW, Jacobs RA, Aslin RN. Bayesian integration of visual and auditory signals for spatial localization. *Journal of the Optical Society of America A* 2003; 20: 1391.

14. Rosas P, Wagemans J, Ernst MO, et al. Texture and haptic cues in slant discrimination: reliability-based cue weighting without statistically optimal cue combination. *Journal of the Optical Society of America A* 2005; 22: 801.
15. Rahnev D, Denison RN. Suboptimality in perceptual decision making. *Behavioral and Brain Sciences* 2018; 41: e223.
16. Gu Y, Angelaki DE, Deangelis GC. Neural correlates of multisensory cue integration in macaque MSTd. *Nature neuroscience* 2008; 11: 1201–1210.
17. Ban H, Preston TJ, Meeson A, et al. The integration of motion and disparity cues to depth in dorsal visual cortex. *Nature Neuroscience* 2012; 15: 636–643.
18. Murphy AP, Ban H, Welchman AE. Integration of texture and disparity cues to surface slant in dorsal visual cortex. *J Neurophysiol* 2013; 110: 190–203.
19. Dekker TM, Ban H, van der Velde B, et al. Late Development of Cue Integration Is Linked to Sensory Fusion in Cortex. *Curr Biol* 2015; 25: 2856–2861.
20. Rohe T, Noppeney U. Distinct computational principles govern multisensory integration in primary sensory and association cortices. *Current Biology* 2016; 26: 509–514.
21. Cao Y, Summerfield C, Park H, et al. Causal Inference in the Multisensory Brain. *Neuron* 2019; 102: 1076-1087.e8.
22. Noppeney U. Perceptual Inference, Learning, and Attention in a Multisensory World. *Annu Rev Neurosci* 2021; 44: 449–473.
23. Beck JM, Ma WJ, Kiani R, et al. Probabilistic population codes for Bayesian decision making. *Neuron* 2008; 60: 1142–52.
24. Ohshiro T, Angelaki DE, DeAngelis GC. A normalization model of multisensory integration. *Nat Neurosci* 2011; 14: 775–82.
25. Fetsch CR, Pouget A, Deangelis GC, et al. Neural correlates of reliability-based cue weighting during multisensory integration. *Nature Neuroscience* 2011 15:1 2011; 15: 146–154.
26. van Bergen RS, Ji Ma W, Pratte MS, et al. Sensory uncertainty decoded from visual cortex predicts behavior. *Nat Neurosci* 2015; 18: 1728–1730.
27. Bergen RSV, Jehee JFM. Probabilistic Representation in Human Visual Cortex Reflects Uncertainty in Serial Decisions. *Journal of Neuroscience* 2019; 39: 8164–8176.

28. Kahneman D, Slovic SP, Slovic P, et al. *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge university press, 1982.
29. Kahneman D, Tversky A. Prospect theory: An analysis of decision under risk. *Econometrica* 1979; 47: 263–292.
30. Tversky A, Kahneman D. Evidential impact of base rates. In: Kahneman D, Slovic P, Tversky A (eds) *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press, 1982, pp. 153–160.
31. Juslin P, Karlsson L, Olsson H. Information integration in multiple cue judgment: A division of labor hypothesis. *Cognition* 2008; 106: 259–298.
32. Whelehan DF, Conlon KC, Ridgway PF. Medicine and heuristics: cognitive biases and medical decision-making. *Irish Journal of Medical Science (1971 -)* 2020; 189: 1477–1484.
33. Kahneman DFS. Representativeness Revisited: Attribute Substitution in Intuitive Judgment. In: Gilovich T, Griffin D, Kahneman D (eds) *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge, UK: Cambridge University Press, 2002, pp. 49–81.
34. Benjamin DJ. Errors in probabilistic reasoning and judgment biases. 2019, pp. 69–186.
35. Wu S-W, Delgado MR, Maloney LT. Economic decision-making compared with an equivalent motor task. *Proceedings of the National Academy of Sciences of the United States of America* 2009; 106: 6088–6093.
36. Jarvstad A, Hahn U, Rushton SK, et al. Perceptuo-motor, cognitive, and description-based decision-making seem equally good. *Proceedings of the National Academy of Sciences* 2013; 110: 16271–16276.
37. Herce Castañón S, Moran R, Ding J, et al. Human noise blindness drives suboptimal cognitive inference. *Nature Communications* 2019 10:1 2019; 10: 1–11.
38. Stengård E, Juslin P, van den Berg R. How deep is your bayesianism?—Peeling the layers of the intuitive Bayesian. *Decision* 2022; 9: 321–346.
39. Negen J, Chere B, Bird L-A, et al. Sensory cue combination in children under 10 years of age. *Cognition*; 193. Epub ahead of print 2019. DOI: 10.1016/j.cognition.2019.104014.
40. Körding KP, Beierholm U, Ma WJ, et al. Causal inference in multisensory perception. *PLoS ONE* 2007; 2: e943-

41. Hollingworth HL. The central tendency of judgment. *The Journal of Philosophy, Psychology and Scientific Methods* 1910; 7: 461–469.
42. Huttenlocher J, Hedges L v., Vevea JL. Why do categories affect stimulus judgment? *J Exp Psychol Gen* 2000; 129: 220–241.
43. Olkkonen M, McCarthy PF, Allred SR. The central tendency bias in color perception: Effects of internal and external noise. *J Vis* 2014; 14: 5–5.
44. Aston S, Negen J, Nardini M, et al. Central tendency biases must be accounted for to consistently capture Bayesian cue combination in continuous response data. *Behav Res Methods* 2022; 54: 508–521.
45. Shams L, Beierholm U. Bayesian causal inference: A unifying neuroscience theory. *Neurosci Biobehav Rev* 2022; 137: 104619.
46. Plummer M. JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling. In: *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*. March, 20–22, 2003.
47. Wozny DR, Beierholm UR, Shams L. Probability Matching as a Computational Strategy Used in Perception. *PLoS Comput Biol* 2010; 6: e1000871.
48. Scarfe P. Experimentally disambiguating models of sensory cue integration. *J Vis* 2022; 22: 5.
49. Rohde M, van Dam LCJ, Ernst M. Statistically Optimal Multisensory Cue Integration: A Practical Tutorial. *Multisens Res* 2016; 29: 279–317.
50. Schulze C, Hertwig R. A description–experience gap in statistical intuitions: Of smart babies, risk-savvy chimps, intuitive statisticians, and stupid grown-ups. *Cognition* 2021; 210: 104580.

Different types of uncertainty in multisensory perceptual decision making

Stacey Aston, Marko Nardini & Ulrik Beierholm

Dept of Psychology, Durham University, UK

Supplement:

Supplemental Methods: Accounting for Central Tendency Bias

Supplemental Methods: Instructions to Participants

Supplemental Results: Variable Error

Supplemental Results: Non-central central bias

Supplemental Methods: Accounting for Central Tendency Bias

Figure S1A plots measured variable error (the standard deviation, or square root of the variance) using the intrinsic+extrinsic cue, σ_{L+E}^2 , from the test block against the measured variable error using the intrinsic-only cue during the test block. The dashed line in the figure is the identity line. Most points fall below this line, indicating that measured variable error using the intrinsic+extrinsic cue was often lower than when using the intrinsic-only cue and this difference was significant ($z = 4.56$, $p < .001$, two-tailed Wilcoxon signed-rank test with normal approximation).

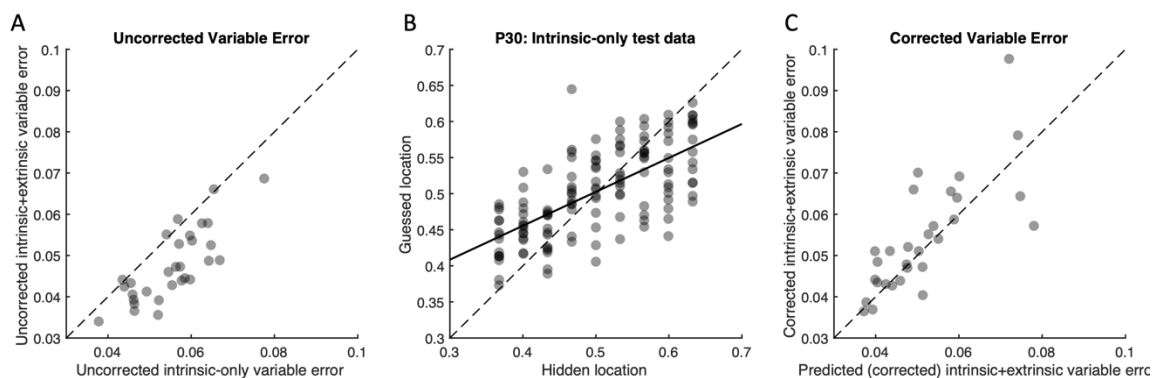


Figure S1: **A.** Individual participants' intrinsic-only vs intrinsic+extrinsic variable errors, uncorrected. **B.** One participant's guesses vs object locations, with a regression diverging from the identity line and indicative of a central tendency bias. **C.** Individual participants' intrinsic-only vs intrinsic+extrinsic variable errors, once corrected.

The mismatch between measured variable error using the intrinsic-only and intrinsic+extrinsic cues is caused by a feature of the continuous response data that, unless properly accounted for, can lead to underestimates of variable error. That feature is a **central tendency bias**¹⁻³ illustrated for a single participant's responses using the intrinsic-only visual cue in the test block in Figure S1B. The dashed line in Figure S1B is the identity line. Clearly, the mean response for each of the nine unique test locations is biased towards the middle of the range, with the magnitude of the bias increasing as the test location moves further from the centre. This is a classic central tendency bias, where estimates are pulled towards the mean of the stimulus distribution. We have recently described a method elsewhere that recovers corrected estimates of variable error⁴. The method is designed to account for central biases in continuous responses. To calculate measures of variability according to the method, we first regress responses for each trial type on the true hidden object locations and calculate the standard deviation of the residuals. If the slope of the fitted regression line is significantly less than one, the standard deviation of the residuals is divided by the fitted slope of the

regression line to correct for a central bias. Importantly, if there is no evidence of a central bias (the slope is not significantly less than one), no correction is performed.

These data demonstrate the appropriateness of calculating a corrected variable error by showing that the corrected estimate for the intrinsic+extrinsic cue does not differ from a prediction based on the corrected estimate for the low intrinsic-only cue (during the calibration block) and the amount of extrinsic noise that was added ($\sigma_{L+E}^2 \approx \sigma_L^2 + \sigma_E^2$ using corrected variable error estimates; $z = -1.9, p = .058$, two tailed Wilcoxon signed-rank test with normal approximation; Figure S1C). This is strong evidence that our recently proposed analysis method recovers a meaningful quantity from the data. Note that the calculation of the expected variable error with extrinsic noise does not depend on any assumptions about behaviour but is based on statistical theory. The fact that corrected variable errors are a better match to the predictions than the uncorrected variable errors suggests that it is the corrected variable errors that offer a better representation of how perceptual location estimates vary before any additional biases corrupt the resulting response. We used this method to generate the variable errors that are analysed in the main paper results section. Unfortunately, as we did not foresee the need for this analysis before we ran the experiment, this means that we did not match variable error using the intrinsic-only and intrinsic+extrinsic visual cues in the test block as intended.

Across 240 calculations (30 participants and eight trial types) we performed the correction of variable error 175 times. The number of corrections and mean strengths of the (corrected) central bias for each trial type were: low intrinsic-only (calibration) $\beta = 0.13$, high-intrinsic-only (calibration) $\beta = 0.25$, audio (training) $\beta = 0.1$, intrinsic-only (test) $\beta = 0.38$, intrinsic+extrinsic (test) $\beta = 0.21$, audio (test) $\beta = 0.18$, intrinsic-only and audio (test) $\beta = 0.15$, intrinsic+extrinsic and audio (test) $\beta = 0.12$.

Supplemental Methods: Instructions to Participants

The initial instructions for participants were:

'Hello! Thank you for agreeing to take part in this experiment. In this task you will always be trying to locate an object behind the screen as if playing a fairground game where you are trying to find a hidden object. Sometimes your clue to where the object is hidden will be the presentation of 4 dots in quick succession. You can think of these as the guesses of a previous player that were close to the correct answer but not quite right. The best way to use this clue is to guess that the location of the object coincides with the average horizontal location of the dots (the centroid). Here is an example.'

Instructions for calibration:

'Let's start with some practice using this visual clue. Remember, you will see 4 dots presented on the screen in quick succession. After the last dot has been presented, a black line will appear. Using the mouse, position the black line to where you think the dots were, on average, horizontally across the screen. Once you are happy with the placement of the line, click the left mouse button to save that response and see the true location (green square). You will gain points after each trial depending on how well you did on that trial. Points are converted to a cash reward at the end of the task. The trials will start immediately after you press enter so be ready!'

Instructions for auditory clue:

'Great! That's plenty of practice with the visual clue. Sometimes, rather than the visual clue to the hidden objects location, you will get an auditory clue. The auditory clue can be thought of as a rustling sound made by the person hiding the object that comes from behind the screen in the same location as the object is hidden. The best way to use this clue is to guess that the object is in the location that you heard the sound coming from. Here is an example.'

Followed by:

'Let's do some practice using this auditory clue. Remember, on each of these trials you will hear a noise from behind the screen. Once you have heard the noise a black line will appear. Position this

line to where to think the noise came from using the mouse. Like before, a left click issues your response and you will see feedback.'

The instructions given for the test phase of the experiment were:

'Great! You are now well practiced with both clues and we are ready to start playing the game. When playing the game you will get two types visual clues and an auditory clue. The two types of visual clues are previous guesses from different individuals and this will be signalled by the dot colour and the person's name (Rodger - blue dots, June - orange dots) will be displayed at the top of the screen. The auditory clue is always a sound made by the person hiding the object. On most trials, you will get only one of these clues. However, on some trials, you will get one of the visual clues and auditory clue together. Feedback will not be given on all trials, including the amount of points gained, but we will still be tallying them up in the background.'

Supplemental Results: Variable Error

Figure S2A shows that the measured reduction in variable error (best-both) when combining the *intrinsic-only* visual cue with the auditory cue is not significantly different to zero ($z = -0.89, p = .371$, two-tailed Wilcoxon signed-rank with normal approximation), significantly less than the optimal or best possible reduction from taking the reliability-weighted average ($z = -4.08, p < .001$), significantly less than a prediction tailored ($\sigma^2 = w_V^2 \sigma_V^2 + w_A^2 \sigma_A^2$) to the measured weights for each individual ($z = -3.05, p = .002$), and significantly greater than if participants were switching ($\sigma^2 = \widehat{w}_V \sigma_V^2 + \widehat{w}_A \sigma_A^2$) which cue they used from trial to trial ($z = 4.43, p < .001$). In sum, while participants did not show a benefit in precision when using the cues in combination, they did not do worse than with their best single cue either, implying that they were not switching between the cues. When combining the *intrinsic+extrinsic* and auditory cues, the measured gain was significantly greater than zero ($z = 2.15, p = .032$), significantly less than the optimal gain ($z = -4.17, p < .001$), not significantly different to the tailored prediction ($z = -1.43, p = .153$), and significantly greater than the switching prediction ($z = 4.62, p < .001$; Figure S2B).

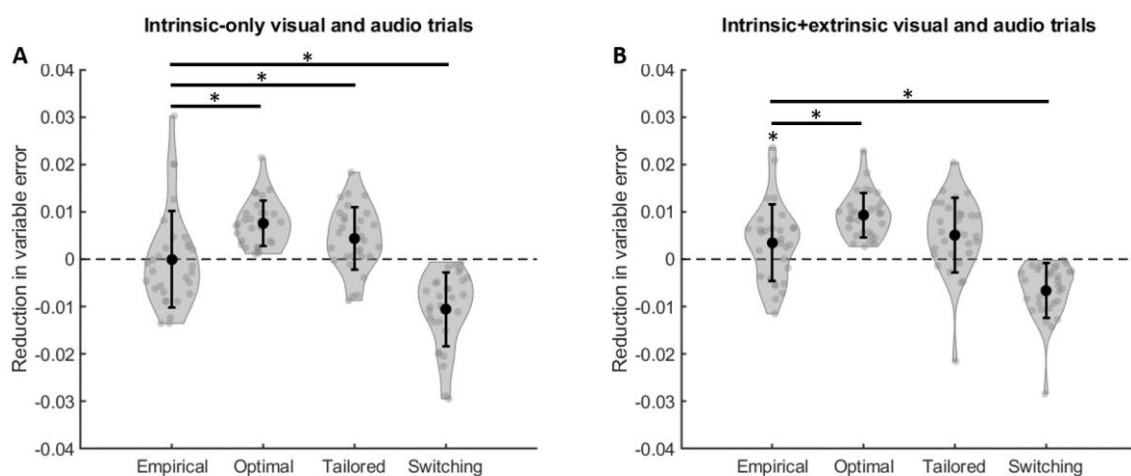


Figure S2: Reduction in variable error in the experiment, for congruent data. A) *Intrinsic-only* data set, empirical reduction (relative to best modality), the optimal reduction given the two modalities, reduction given the estimated weights and reduction for a model that in each trial either uses audio or visual modality only. B) Same for *Intrinsic+extrinsic* data set.

While it seems surprising that the reduction in error is not significant for intrinsic-only (where one would expect it to be) and is significant for intrinsic-extrinsic (where one would not expect it to be) it is worth keeping in mind that any benefit is only to be expected to be possible to detect if the

two perceptual modalities are evenly matched for variable error. If visual stimulus is very uncertain it will tend to have a large error, and there will be little benefit to auditory localisation in combining them.

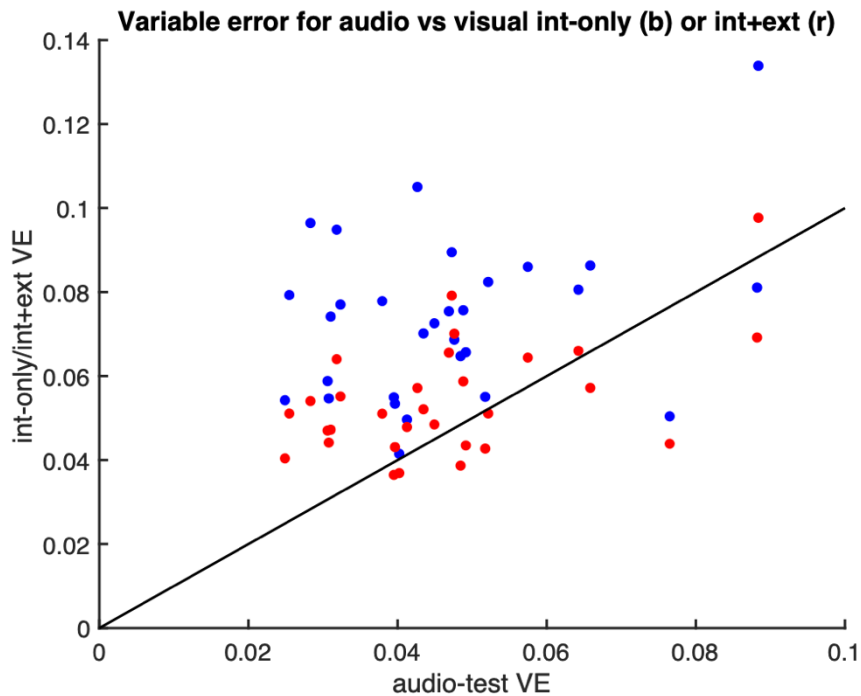


Figure S3: Variable error for audio-test (horizontal axis) against the visual int-only (blue) or int-ext (red). Points close to the diagonal line would have well matched variable errors.

However Fig. S3 shows that, because of the issues discussed in the Methods and Supplemental Methods, the intrinsic+extrinsic variable error is better matched with audio than intrinsic-only is, making it difficult to test for any multisensory advantage of combination. In line with this, the optimal prediction in Fig. S2 shows that there is more benefit available in the Intrinsic+Extrinsic (Mean=0.0093) than the Intrinsic-only condition (Mean=0.0076). This difference was significant ($t(29)=3.16, p<.005$).

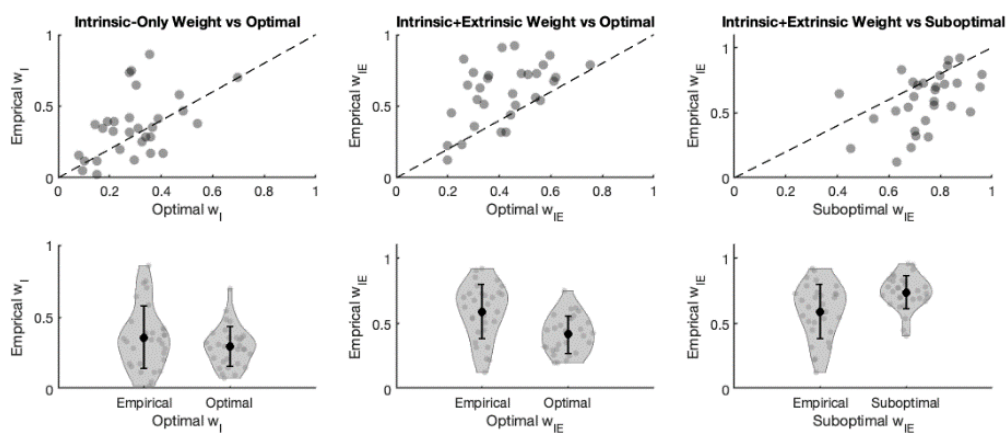
Supplemental Results: Non-central central bias

While the incongruent data were erroneously presented on the left side, congruent data were presented bilaterally. As subjects would find it difficult to tell which stimuli were congruent and incongruent, we would expect a small effect.

Nevertheless, we ran a version of the analysis for the incongruent data that allowed the central tendency bias to be shifted away from the screen center, to model the possibility that a central prior would be shifted away.

We would expect that presentation of stimuli only on the left side could potentially create a central tendency bias shifted towards the left side of the screen, but instead found a small, albeit significant bias towards the right side (intrinsic-only: mean 0.542, sign-rank test $p < 0.01$, intrinsic-only: mean 0.558, sign-rank test $p < 0.01$. Center of the screen is 0.5.)

Re-running the main the analysis using this alternative calculation of the central tendency bias, we again found a non- significant difference between empirical and optimal weights for Intrinsic-only ($t(29) = -1.79$, $p = 0.084$), and a significant difference between empirical and intrinsic+extrinsic data ($t(29) = -5.19$, $p < 0.001$). See figure below.



Overall, this implies no effect of the (unfortunate) presentation on the left side of the screen for the incongruent stimuli.

References

1. Hollingworth HL. The central tendency of judgment. *The Journal of Philosophy, Psychology and Scientific Methods* 1910; 7: 461–469.
2. Huttenlocher J, Hedges L v., Vevea JL. Why do categories affect stimulus judgment? *J Exp Psychol Gen* 2000; 129: 220–241.
3. Olkkonen M, McCarthy PF, Allred SR. The central tendency bias in color perception: Effects of internal and external noise. *J Vis* 2014; 14: 5–5.
4. Aston S, Negen J, Nardini M, et al. Central tendency biases must be accounted for to consistently capture Bayesian cue combination in continuous response data. *Behav Res Methods* 2022; 54: 508–521.