

Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes

Louise Whiteley and Maneesh Sahani
Gatsby Computational Neuroscience Unit,
University College London,
London, UK.
www.gatsby.ucl.ac.uk

To appear in *Journal of Vision*

ABSTRACT

Perception is an ‘inverse problem’, in which the state of the world must be inferred from the sensory neural activity that results. However, this inference is both ill-posed (Helmholtz, 1856; Marr, 1982) and corrupted by noise (Green & Swets, 1989), requiring the brain to compute perceptual beliefs under conditions of uncertainty. Here we show that human observers performing a simple visual choice task under an externally imposed loss function approach the optimal strategy, as defined by Bayesian probability and decision theory (Cox, 1961; Berger, 1985). In concert with earlier work, this suggests that observers possess a model of their internal uncertainty, and can utilise this model in the neural computations that underlie their behaviour (Knill & Pouget, 2004). In our

experiment, optimal behaviour requires that observers integrate the loss function with an estimate of their internal uncertainty, rather than simply requiring that they use a modal estimate of the uncertain stimulus. Crucially, they approach optimal behaviour even when denied the opportunity to learn adaptive decision strategies based on immediate feedback. Our data thus support the idea that flexible representations of uncertainty are pre-existing, widespread, and can be propagated to decision making areas of the brain.

Keywords: Bayesian, early sensory noise, ideal observer

INTRODUCTION

The idea that perception should be viewed as unconscious inference dates back to Al Hazen's 11th century treatise on optics, and is still fundamental to vision science today. The challenge is to understand how such inference takes place in situations that are both ill-posed (Helmholtz, 1856; Marr, 1982) and noisy (Green & Swets, 1989). Considerable recent evidence suggests that in such situations human inference can closely approach the performance of a Bayes-optimal observer (the probabilistic inferential analogue to the ideal observer of signal detection theory). These demonstrations have largely focused on sensory cue-combination tasks, both crossmodal (Ernst & Banks, 2002; Deneve, Latham, & Pouget, 2001) and unimodal (see, e.g. Jacobs, 1999; Hillis, Watt, Landy, & Banks, 2004; Knill & Saunders, 2003; Landy & Kojima, 2001), and on the effects of motor and sensory uncertainties on motor planning (Trommershauser, Maloney, & Landy, 2003; Saunders & Knill, 2004; Kording & Wolpert, 2004; Saunders & Knill, 2005; Trommershauser, Gepshtein, Maloney, Landy, & Banks, 2005; Tassinari, Hudson, & Landy, 2006). There is much less evidence for Bayesian optimality in perceptual estimation tasks for a single, visual quantity (but see Weiss, Simoncelli, & Adelson, 2002; Schwartz, Sejnowski, & Dayan, 2005; Stocker & Simoncelli, 2006; Landy, Goutcher, Trommershauser, & Mamassian, 2007).

The present study was designed to probe what might be the simplest context in which Bayesian

optimality can be demonstrated, and to rule out other explanations for apparently Bayes optimal behaviour. The task was a simple unimodal visual Vernier or ‘offset’ discrimination (Westheimer, 1979), in which observers had to answer whether one stimulus was offset to the left or right of another. The stimulus set was fixed, and thus any perceptual uncertainty was due to sensory noise and subsequent processing. This enabled us to ask whether even for the simplest visual task the brain has an estimate of internal uncertainty available to guide behaviour (see also Schwartz et al., 2005). There are two key issues with this kind of behavioural Bayesian optimality experiment: (a) what optimality implies about which aspect of the sensory noise distribution is represented, and (b) whether we can conclude that observers do not mimic optimal behaviour via another strategy which does not require the representation of uncertainty. We attempt to address both these issues in the present study.

To probe the observers’ model of internal uncertainty we imposed an asymmetric loss function on their Vernier judgments. In the face of such an asymmetry, the Bayes-optimal observer combines information about their internal uncertainty, at least in the form of a likelihood ratio, with knowledge about the loss function in order to arrive at an optimal decision (Cox, 1961; Berger, 1985). In any such behavioural task information from a potentially uncertain distribution is used in the computation of a decision, but here we show that information about the *degree* of uncertainty is combined with the loss function, rather than simply an optimal estimate of the stimulus. This design provides an alternative to the cue-combination approach, allowing us to assess whether observers use information about uncertainty over a single visual quantity.

Similar external loss functions have long been used in the psychophysics literature to modify response criteria, and thus explore the receiver operating characteristic (ROC) curve (Green & Swets, 1989). In these earlier experiments, however, observers were provided with trial-by-trial feedback about their performance, potentially allowing them to adopt a feedback-driven threshold-adaptation strategy that mimics the Bayes optimal strategy without requiring an explicit model of internal uncertainty. Crucially, in our experiment, observers received only periodic cumulative feedback, thus ruling out the use of such simple adaptive threshold schemes and strengthening the conclusions that can be drawn

from behavioural results.

METHODS

Observers

Four participants (2 male, 2 female) took part in the experiment. They had a mean age of 25.5 years, were all right-handed, and had normal or corrected-to-normal vision. Three were entirely naive with respect to the aims of, and theory behind, the experiment; and the fourth (observer 4) was an author (LW).

Stimulus and Equipment

The stimulus consisted of a pair of vertically arranged Gabor-like patches, in which a sinusoidal grating with a spatial frequency of 0.03 cycles/mm (0.21 cycles/° of visual angle), was multiplied by a Gaussian envelope with a characteristic width (2σ) of 29.9 mm, truncated at a full width of 100 mm (14.4° of visual angle) in the horizontal direction, and a rectangular envelope with a width of 10.3 mm (1.48° of visual angle) in the vertical (see Fig. 1a). The pixel intensity in the two patches ranged from 0 to 255 (black to white) against a grey background of intensity 128. The separation of the two patches was 5.67 mm (0.813° of visual angle), and the stimulus appeared with the centre of the upper patch located 66.7 mm (9.56° of visual angle) either to the left or right of the fixation cross in a pseudorandomised order. On each trial the entire lower patch (both envelope and grating) was displaced relative to the upper patch by one of 20 pseudorandomised values, ranging from -15 to +15 pixels (positive numbers indicating offsets to the right). Each pixel corresponded to 0.333 mm (0.0478° of visual angle).

Before each block of main trials, observers were given a short block of practice trials in which the

stimulus duration was 500ms. In main trials the stimulus duration was 160ms, which is shorter than the latency for initiating a saccade (Carpenter, 1988; Hodgson, 2002), and was chosen to avoid fixation of the stimuli. There was a randomized delay period of 750-1250 ms between the time of response and the time of presentation of the next stimulus.

The experimental program was written in MATLAB (The Mathworks Inc, Natick, MA), using the Psychophysics Toolbox extensions (Pelli, 1997; Brainard, 1997).

Procedure

Observers sat at a table in front of a computer screen, and placed their head on a chin rest such that the perpendicular distance from their eyes to the screen was 400 mm. During the experiment, observers fixated a central cross on the screen, and were asked to make simple Vernier judgments (Westheimer, 1979), reporting whether the lower of the two Gabor patches was offset to the right or left of the upper one. The task is depicted schematically in Fig. 1a. Responses were given using a computer keyboard.

We imposed an asymmetric loss function to probe the observers' representation of uncertainty. On each trial observers were awarded points for a correct answer ('rewards'), or had points taken away for an incorrect answer ('costs'). Observers were instructed, and given an incentive, to maximise the cumulative number of points scored during the experiment. The loss function varied between blocks of trials — the rewards for correctly answering 'right' (\mathcal{R}_r) or 'left' (\mathcal{R}_l) were constant and equal, but the cost for incorrectly answering 'right' (\mathcal{C}_r) could be different from that for incorrectly answering 'left' (\mathcal{C}_l). A similar asymmetric penalty approach has been used to probe uncertainty in recent studies of motor planning (Trommershauser et al., 2003; Trommershauser et al., 2005). When the penalty for answering 'left' incorrectly is greater, a reasonable strategy would be to answer 'right' more often when uncertain of the answer. This would result in a psychometric curve shifted in the direction of the higher penalty, yielding a higher overall score (see Fig. 2a).

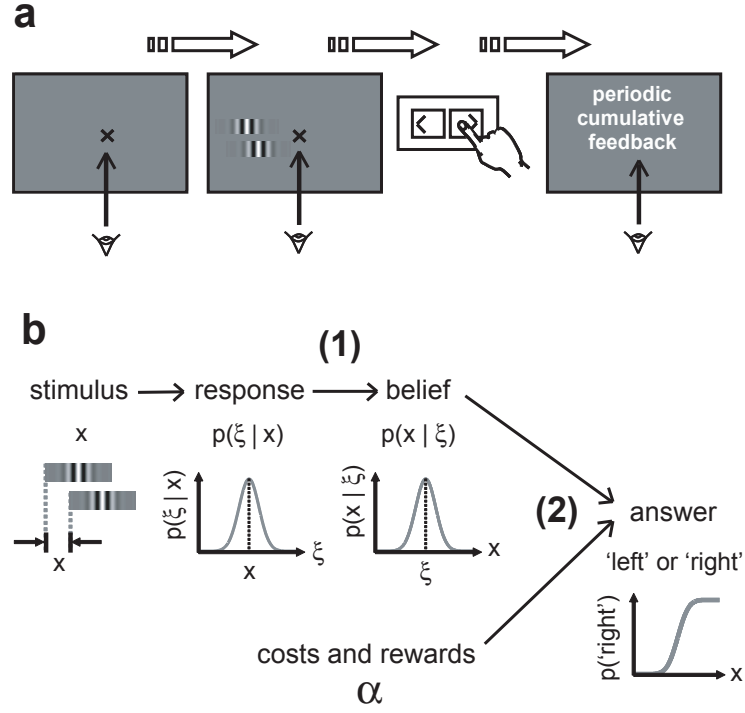


Figure 1: **Experimental design.** **a**, On each trial a stimulus consisting of two vertically arranged Gabor patches was briefly flashed, and the observer pressed a key to say whether the lower patch was offset to the left or right of the upper patch. Participants were asked to maximize their score, with varying numbers of points being awarded for a correct answer (‘reward’) and deducted for an incorrect answer (‘cost’). Participants received only periodic feedback about their performance, in the form of a cumulative score every 15 trials. **b**, Schematic of the quantities and transformations involved in the construction of the Bayes-optimal observer. The stimulus x produces a stochastic neural response. The observer transforms this neural response into a belief distribution (see (1)), and then uses this belief to decide which answer to give in the face of varying costs and penalties (see (2)). The Bayes-optimal observer specifies the optimal forms of transformations (1) and (2), thereby providing a behavioural benchmark of optimality.

We used five different sets of costs and rewards, listed in Table 1. The final column in this table shows, qualitatively, the relative shift we would expect for each set of costs and rewards according to the strategy just described. In the Results section we describe an optimal Bayes observer analysis which confirms, and quantifies, the optimality of this ‘curve shifting’ strategy.

The five different sets of costs and rewards were presented in a counterbalanced, pseudorandomised

Table 1: Values of α corresponding to costs and rewards

$\mathcal{R}_r = \mathcal{R}_l$	\mathcal{C}_l	\mathcal{C}_r	α	curve shift
+20	-10	-50	0.3	→
+20	-20	-40	0.4	→
+20	-30	-30	0.5	0
+20	-40	-20	0.6	←
+20	-50	-10	0.7	←←

block design. This was repeated in two experimental sessions on separate days, which provided a further test of our hypothesis — the level of observers’ internal uncertainty might be expected to differ in the two sessions, due to perceptual learning, consolidation or extrinsic factors. If observers were to behave optimally in two sessions with *different* noise distributions, this would support the claim that they carry a flexible representation of current internal uncertainty.

Each block consisted of a short practice session (60 trials), and a main session (260 trials). In the main session feedback on performance was provided only every 15 trials, when observers were shown the total score they obtained in the last 15 trials, as well as their cumulative total in the block so far. The sparseness of this feedback made it unlikely that observers could learn an optimal internal threshold by an adaptive threshold adjustment strategy. Control analyses (see Results) support this view. In the practice session, observers received trial-by-trial feedback to familiarize them with the cost values for that block, and encourage consistent performance in the main trials. However, the practice stimuli were presented for 500ms rather than 160ms, which made the task much easier. As the effective internal noise should therefore be different for the practice stimuli, feedback in the practice session could not be used to adaptively learn a response threshold relevant to the main-block trials. In addition, the easier stimulus meant that there were very few trials on which observers failed to give the correct answer, implying that there should have been very little uncertainty in their belief. This further limits the likelihood that they could use the practice session to test alternative strategies for dealing with the loss function.

An instruction screen appeared at the beginning of each block, and after each feedback screen, reminding observers of the task and costs for that block. After the experiment was finished, observers were debriefed using a simple questionnaire. Participants were paid according to ethical guidelines, with a score-related bonus in gift vouchers to motivate concentrated performance and encourage observers to try to maximise their total score.

RESULTS

Bayes Optimal Observer Analysis

The final column of Table 1 shows the relative shift of the psychometric function we might expect for different settings of the costs and rewards, under an intuitive strategy for maximizing score in which observers shift their psychometric function in the direction with the higher penalty. A quantitative Bayesian observer analysis was used to confirm and quantify the optimality of this strategy. Fig. 1b depicts the quantities involved in this analysis. The visual stimulus, with a Vernier offset x , evokes a stochastic neural response, on the basis of which the observer constructs an internal belief about the value of the stimulus offset (step (1) in Fig. 1b). This belief is then used to guide a decision about the appropriate response (step (2)).

An individual observer's responses to repeated presentations of the same visual stimulus may vary. We assume that this variability arises from at least two separate sources of noise. The first source perturbs the observer's sensory estimate of the Vernier offset by a random additive increment. This creates a noise distribution centred on the stimulus, the uncertainty due to which is reflected in the observer's belief. The second source affects the observer's decision directly, in a way that is independent of the value of the stimulus offset. We may think of this as 'decision noise', or as the result of motor errors or of lapses in attention. We do not expect this source of variation to affect the observer's internal belief about the value of the offset, and so it is neglected in the theoretical development below.

When modelling experimental responses, however, we introduce a separate ‘lapse rate’ parameter, so that these stimulus-independent errors do not corrupt our estimate of the stimulus-centred noise. Note that we do not distinguish between stimulus-centred *sensory* noise, and any stimulus- or estimate-centred *decision* noise which might, for instance, arise as sensory information is integrated with the loss function. Our definition of Bayesian optimality in decisions thus refers to all stimulus-centred variation. In concert with earlier analyses, we do however assume that the majority of this variation is ‘sensory noise’, and so treat and refer to it as such.

In keeping with the standard psychophysical treatments of sensory noise, our model assumes that the internal estimate of the Vernier offset, ξ , is normally distributed with constant variance σ^2 around the true stimulus offset: $p(\xi|x) = \mathcal{N}(\xi; x, \sigma^2)$ (Thurstone, 1927; Green & Swets, 1989). We test this assumption below, showing that the psychometric curves were all well fit by cumulative normal functions, with a constant slope parameter for each observer in each session. However, our observers each displayed a systematic bias in their responses; this will be addressed later.

In the Bayesian view, the observer’s belief about the Vernier offset x is not limited to a single estimated value ξ . Instead, ξ parameterises a belief distribution over all possible values of x that are consistent with the sensory evidence. The optimal form for this belief distribution is given by Bayes rule:

$$p(x|\xi) = \frac{p(\xi|x) \cdot p(x)}{p(\xi)}. \quad (1)$$

We assume that the prior belief about x is uniform, which implies that this optimal belief will also be Gaussian, with the same variance as the sensory noise distribution, and mean given by ξ : $p(x|\xi) = \mathcal{N}(x; \xi, \sigma^2)$ (we might also have assumed a broad zero-centred Gaussian prior, although then the variance of the posterior belief would have been slightly smaller than that of the sensory noise, for which there was no evidence in the data). In fact, if observers are able to learn the true distribution of x , their prior (and therefore posterior) belief should take the form of a series of delta functions located

at each discrete offset value used. In addition, for extreme values of x , the stochastic response ξ may fall outside the range of possible values, distorting the posterior. However, variability in decisions at the extremes was minimal, so that any divergence from normality at those points would have little impact on estimates of sensory variance. And within the central range, where decisions did vary, the values of the stimulus offset used were very closely spaced and we saw no evidence that observers were aware of the discretization.

The observer must base his or her response on the belief distribution (step (2) in Fig. 1b), and Bayesian decision theory gives the optimal form of this response (see Berger, 1985; Maloney, 2002; Yuille & Bulthoff, 1996). The observer should answer ‘right’ if and only if, on the basis of his or her belief, the expected (mean) reward (Γ) for answering ‘right’ is greater than the expected reward for answering ‘left’: *i.e.* if $\Gamma[\text{‘right’}] > \Gamma[\text{‘left’}]$. In this simple case, the expected reward is obtained by adding together the product of the probability of the answer being correct times the corresponding reward, and the probability of the answer being wrong times the corresponding cost. These two probabilities express the degree of the participant’s belief that the lower patch fell to the right or left of the upper patch, and are given by the areas under the belief distribution $p(x|\xi)$ that fall to the right and to the left of 0 respectively.

$$P(\text{answer ‘right’}) = P(\Gamma[\text{‘right’}] > \Gamma[\text{‘left’}]), \quad (2)$$

$$\Gamma[\text{‘right’}] = \int_0^{\infty} p(x|\xi) \cdot \mathcal{R}_r \, dx + \int_{-\infty}^0 p(x|\xi) \cdot \mathcal{C}_r \, dx, \quad (3)$$

$$\Gamma[\text{‘left’}] = \int_0^{\infty} p(x|\xi) \cdot \mathcal{C}_l \, dx + \int_{-\infty}^0 p(x|\xi) \cdot \mathcal{R}_l \, dx. \quad (4)$$

With some rearrangement, and combination of the integrals, we arrive at an expression in which the observer’s belief that the Vernier displacement was to the left is compared to a single quantity, α , that includes all the cost and reward terms. The values of α for each set of costs and rewards used in our

experiment are given in Table 1.

$$A_r(x, \alpha) = P\left(\int_{-\infty}^0 p(x'|\xi) dx' < \alpha\right); \quad \alpha \equiv \frac{\mathcal{C}_l - \mathcal{R}_r}{\mathcal{C}_l - \mathcal{R}_r + \mathcal{C}_r - \mathcal{R}_l}, \quad (5)$$

where we have introduced the notation $A_r(x, \alpha)$ for the probability that the Bayesian observer answers ‘right’, given a stimulus offset x and cost structure α , and x' is a dummy variable of integration over the observer’s belief.

It is useful at this point to introduce a normal density function which has the same width as $p(\xi|x)$ and $p(x|\xi)$, but zero mean: $f_\sigma(\zeta) = \exp(-\zeta^2/2\sigma^2)/\sqrt{2\pi\sigma^2}$. Thus $p(\xi|x) = f_\sigma(\xi - x)$ and $p(x|\xi) = f_\sigma(x - \xi)$, and the corresponding cumulative function is $\Phi_\sigma(\zeta) = \int_{-\infty}^{\zeta} f_\sigma(\zeta') d\zeta'$. Then,

$$A_r(x, \alpha) = P\left(\int_{-\infty}^0 f_\sigma(x' - \xi) dx' < \alpha\right) \quad (6)$$

$$= P\left(\int_{-\infty}^{-\xi} f_\sigma(\zeta) d\zeta < \alpha\right) \quad [\text{where } \zeta = x' - \xi] \quad (7)$$

$$= P(\Phi_\sigma(-\xi) < \alpha) \quad (8)$$

$$= P(\xi > -\Phi_\sigma^{-1}(\alpha)). \quad (9)$$

The probability that ξ is greater than $-\Phi_\sigma^{-1}(\alpha)$ can then be found by integrating the assumed sensory noise distribution:

$$A_r(x, \alpha) = \int_{-\Phi_\sigma^{-1}(\alpha)}^{\infty} p(\xi|x) d\xi. \quad (10)$$

If we again insert Φ_σ (and exploit its symmetry) we obtain the following easily computed expression for the optimal probability with which the observer should answer ‘right’ for a given value of α and x ;

$$A_r(x, \alpha) = \int_{-\Phi_\sigma^{-1}(\alpha)}^{\infty} f(\xi - x) d\xi \quad (11)$$

$$= \int_{-\infty}^{x + \Phi_\sigma^{-1}(\alpha)} f(\zeta') d\zeta' \quad [\text{where } \zeta' = x - \xi] \quad (12)$$

$$= \Phi_\sigma(x + \Phi_\sigma^{-1}(\alpha)). \quad (13)$$

The only unknown quantity in Equation 13 is the standard deviation, σ , of the zero-mean cumulative Gaussian distribution Φ_σ . This plays two roles in our Bayesian analysis; it is both the width of the sensory noise distribution, and, under the assumed uniform prior, the width of the consequent belief distribution. Under the symmetric cost condition ($\alpha = 0.5$) the observer’s decision reflects only whether the mean of his or her belief lies to the left or right of 0 (according to sensory noise), and is independent of the width of the belief distribution. Thus, following the standard psychometric approach, we estimate the variance of the noise by fitting a psychometric function based on a cumulative Gaussian to the behavioural data, with the slope of the function providing an estimate of σ .

The analysis of the Bayesian observer expressed in Equation 13 makes two predictions: as α changes, the psychometric curves (1) retain the same cumulative-normal shape, with the same width parameter, and (2) translate by an amount $\Phi_\sigma^{-1}(\alpha)$. Fig. 2b shows an example of the psychometric function fit to the data for one observer in one session, and shifts with changing α are clearly visible.

The fitting procedure, and the methods used to test these predictions, are detailed below. Briefly, we first verified that the shape of the psychometric function did not change with α by Bayesian model selection. We then tested the agreement of the observed curve translations with those predicted by the optimal Bayesian analysis. We fit psychometric functions to the data and measured the centre μ of

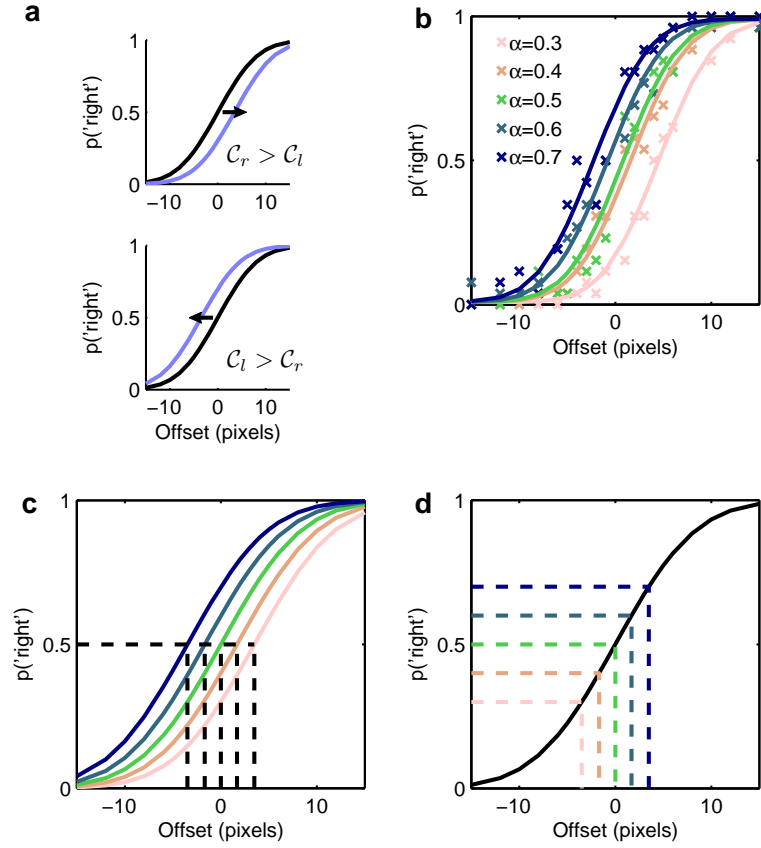


Figure 2: **Evaluating behavioural optimality.** **a**, Illustrates the qualitative prediction for optimization of the loss function - observers should give the answer with the lower penalty when uncertain, resulting in a shift of the psychometric curve in the direction of the higher cost. **b**, Example data from one observer in the five different α conditions. Crosses show data points, and the smooth lines show psychometric functions fit to the data, with the slope constrained to be the same for each α condition **c**, Illustrates the procedure for measuring observed shifts, once psychometric functions have been fit to the data from the five α conditions. **d**, Illustrates the procedure of taking the inverse value of the psychometric function at the values of α used in the experiment. The optimal shift between two psychometric curves is then given by the difference between the two corresponding inverse values.

each, *i.e.* the value of x at which the fitted psychometric curve gave equal probabilities of each answer, given by the mean of the underlying Gaussian (see Fig. 2c). We then used the maximal slope of the psychometric functions as a measure of σ , and inverted Equation 13 to recover the predicted *optimal*

values of the centre μ_j^* for each cost asymmetry value α_j (see Fig. 2d).

$$\begin{aligned}
0.5 &= \Phi_\sigma(\mu_j^* + \Phi_\sigma^{-1}(\alpha_j)), \\
\mu_j^* &= \Phi_\sigma^{-1}(0.5) - \Phi_\sigma^{-1}(\alpha_j), \\
\mu_j^* &= -\Phi_\sigma^{-1}(\alpha_j).
\end{aligned} \tag{14}$$

Fitting the Psychometric Function

The pattern of observers' responses was modelled by a cumulative normal psychometric function incorporating a random *lapse* term (see, e.g., Wichmann & Hill, 2001), and binomially distributed response counts. We used 20 different true offsets x_i , and 5 different cost asymmetries α_j , with N_{ij} trials in each condition. The number of trials n_{ij} in which observers answer 'right' for stimulus offset x_i and cost distribution α_j , is assumed to be drawn from a binomial distribution:

$$P(n_{ij}) = \binom{N_{ij}}{n_{ij}} p_{ij}^{n_{ij}} (1 - p_{ij})^{N_{ij} - n_{ij}}. \tag{15}$$

In the absence of lapses, the optimal probability p_{ij} should be given by $A_r(x_i, \alpha_j)$ in Equation 13, which has a cumulative normal form. To fit the data, we therefore also assumed an underlying cumulative Gaussian shape, parameterised in terms of the standard error function, such that the parameters μ_j and ρ_j gave the centre and maximal slope, respectively, of the curve under the j th value of α .

$$p_{ij}^{no\ lapse} = \frac{1 + \operatorname{erf}(\sqrt{\pi} \cdot \rho_j \cdot (x_i - \mu_j))}{2}; \quad \operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt. \tag{16}$$

However, it is likely that observers occasionally make errors due to stimulus-independent (but possibly cost-structure-dependent) sources such as ‘decision noise’, motor response errors, or moments of inattention (Green & Swets, 1989; Wichmann & Hill, 2001). In this case they might give either answer with equal probability, effectively setting p_{ij} in such cases to $\frac{1}{2}$ rather than the value given above. We took the probability of such an event occurring in any trial to be ϵ_j (the ‘lapse rate’ parameter referred to above), leaving the probability that the response was instead based on the cumulative Gaussian function as $1 - \epsilon_j$;

$$p_{ij} = (1 - \epsilon_j) \cdot \left[\frac{1 + \operatorname{erf}(\sqrt{\pi} \cdot \rho_j \cdot (x_i - \mu_j))}{2} \right] + \epsilon_j \cdot \frac{1}{2}. \quad (17)$$

There are thus three parameters, all of which potentially depend on α : the centre μ_j and slope ρ_j of the cumulative Gaussian, and the random error or lapse rate ϵ_j . An estimate of the slope parameter ρ_j provides an estimate of the width of the underlying Gaussian, according to

$$\sigma = \frac{1}{\sqrt{2\pi\rho}}. \quad (18)$$

When fitting the model to the data, we used Bayesian model comparison to determine whether the slope and lapse rate parameters should be shared between different α conditions, or fit separately (see below).

Shape of the Psychometric Curves

The Bayesian analysis predicts that, as the loss function varies, the psychometric curve will shift, but will retain both the cumulative Gaussian shape, as well as the same maximal slope. We tested both of these predictions.

To ask whether the cumulative Gaussian model with allowance for lapses (Equation 17) was appropriate for the data at all values of α , we examined the residual error between the measured response data

and the best fit psychometric curve. Fig. 3 shows the deviance residuals (McCullagh & Nelder, 1989; Wichmann & Hill, 2001) for all four participants, for each of the two sessions. The deviance residual is used to measure discrepancies in terms of the underlying likelihood model; in effect, it rescales the error by the locally predicted variance. Based on the total deviance, the cumulative normal model could not be rejected by a degrees-of-freedom-adjusted χ^2 -test, nor by a Monte-Carlo-based exact-binomial test (Wichmann & Hill, 2001) ($p > 0.3$ and $p > 0.8$ respectively, after correcting for multiple tests; in neither case could the distribution of p -values over the multiple tests be distinguished from uniform; Kolmogorov-Smirnov test, $p > 0.05$).

There is also no systematic trend evident in Fig. 3 to suggest that the *shape* of the psychometric function was inappropriate for any value of α , for any observer. This was confirmed using a runs test for randomness of the sign of the residuals, by which the hypothesis that the scatter of residuals was random could not be rejected ($p > 0.7$ after multiple-test correction, p -values uniform by K-S test, $p > 0.05$).

The second prediction was that the slope of the psychometric curve is also independent of the value of α ; that is, the parameters ρ_j in Equation 17 are, in fact, all the same. Visual inspection of the data supported this assumption (see Fig. 2b for an example). To assess this quantitatively, we employed Bayesian model selection to ask which of the various models with either shared or varying slope and lapse parameters was more probable, given the data. Details of this procedure are given in Appendix A.1, along with results in Table 2, and the fitted parameter values for the best model in Table 3. As predicted, for all observers, the best model had a single slope parameter for all α conditions within each session. However, lapse rates did vary with α for two observers, whilst being constant for the other two. The lapse rate captures decision noise, motor errors, and moments of inattention, and it seems reasonable that, whilst the internal uncertainty is the same for each α value, such random lapses might depend on the costs.

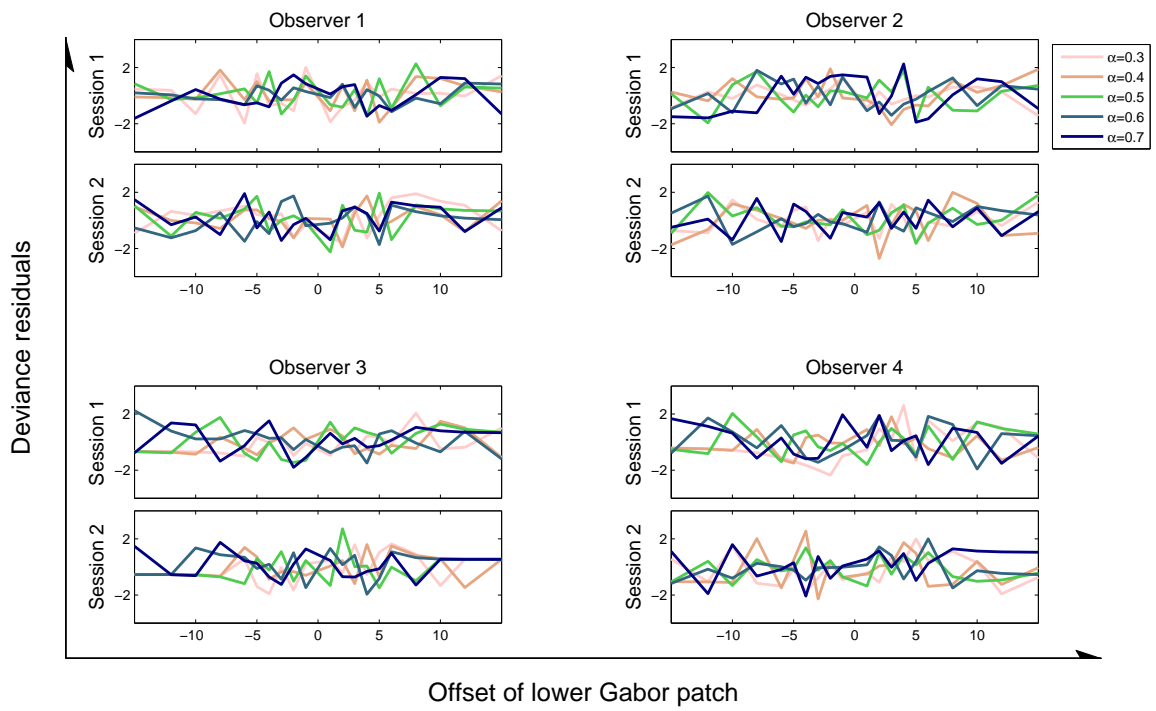


Figure 3: **Deviance residuals between model and data.** Deviance residuals between the model fitted to the behavioural data and the data points, for each observer in each of the two sessions.

Optimal and Observed Shifts

Consideration of the various models thus showed that the behaviour of each observer in each session was best modeled by a family of curves of the same shape and slope, but with centres depending on α . We next asked whether the observed shifts in the curve centres were aligned with the predictions of the ideal Bayesian observer model.

In at least one regard, observers were not optimal. The Bayesian prediction for the curve centre in the $\alpha = 0.5$ condition is always 0. However, for all observers, the curve centres for the $\alpha = 0.5$ condition were non-zero. Two observers showed a rightward bias in both sessions, and two showed a leftward bias in both sessions (see Table 3 in Appendix A.1), and we found no evidence that the bias was absent in the asymmetric penalty conditions. A similar directional bias has been reported widely in psychophysical studies (Green & Swets, 1989). In the analysis below we treat the directional bias as a constant constraint on observers' computations, and attempt to separate this form of non-optimality from the novel question of whether observers were able to integrate correctly the loss function with an estimate of internal uncertainty. Thus we compute shifts as *relative to the biased centre for $\alpha = 0.5$* , yielding 'predicted relative shifts' for the other four α conditions;

$$\Delta\mu_j^* = \mu_j^* - \mu_{0.5} = \Phi_\sigma^{-1}(\alpha_j) - \Phi_\sigma^{-1}(0.5) = \Phi_\sigma^{-1}(\alpha_j). \quad (19)$$

The comparison between observed and predicted relative shifts for the two sessions is shown in Fig. 4a and b. Note that both observed and predicted shifts derive from the same set of data, as the predicted shifts are based on an estimate of internal uncertainty derived from the slope of the psychometric curve. Thus estimation errors due to limited sampling may be correlated, and independent error bars for the two quantities cannot be drawn. Instead, we employed a bootstrap procedure to estimate the covariance of the errors in the two derived quantities, shown by the ellipses in Fig. 4a and b. A linear

fit to the observed shifts was computed by minimizing weighted squared error in the plane with respect to these covariances, and is also shown in Fig. 4a and b.

There is a strong qualitative match between measured and predicted relative shifts. Observers shift their psychometric curves in the right direction, and by an amount that is proportional to the size of the penalty asymmetry. This is in contrast to the simple strategy verbally reported by all observers, which was to give the answer with the lower penalty whenever they were unsure. Interference from this cognitive strategy might help to explain the non-linearity of the observed shift plot in Fig. 4a and b — the two leftward and two rightward shifts are more similar than in the quantitatively optimal scenario. However, if observers' behaviour was dictated by this simple strategy, we would expect the shifts to be of the same magnitude regardless of the size of the penalty asymmetry. In fact the shifts are significantly larger ($p < 0.01$ under a 2-tailed paired-samples t-test) for the greater cost asymmetries.

It has been observed previously that observers are reluctant to behave optimally when this entails an extreme bias in their responses (Green & Swets, 1989). To try and avoid such psychological effects we chose asymmetries which demanded a relatively small shift in the psychometric curve. In general, observers tended to *over* compensate for the penalty asymmetry (see Fig. 4a,b), but for most sessions a *smaller* over-compensation was seen for larger α values, as would be predicted by such an effect. This could also have contributed to the non-linearity in the observed shifts plot.

Optimality of achieved score

Although all observers showed the predicted pattern of shifts, the quantitative match was not exact. This is perhaps unsurprising given the requirement to integrate implicit knowledge of internal uncertainty with high level cognitive instructions. This raises the issue of which behavioural measure should be used to statistically test for optimality. In the present study, observers are asked to maximize their *point score*, not to work out what the optimal shift of the curve should be. It is possible that the function relating curve shift to total score is relatively flat in the region of the maximum score obtainable,

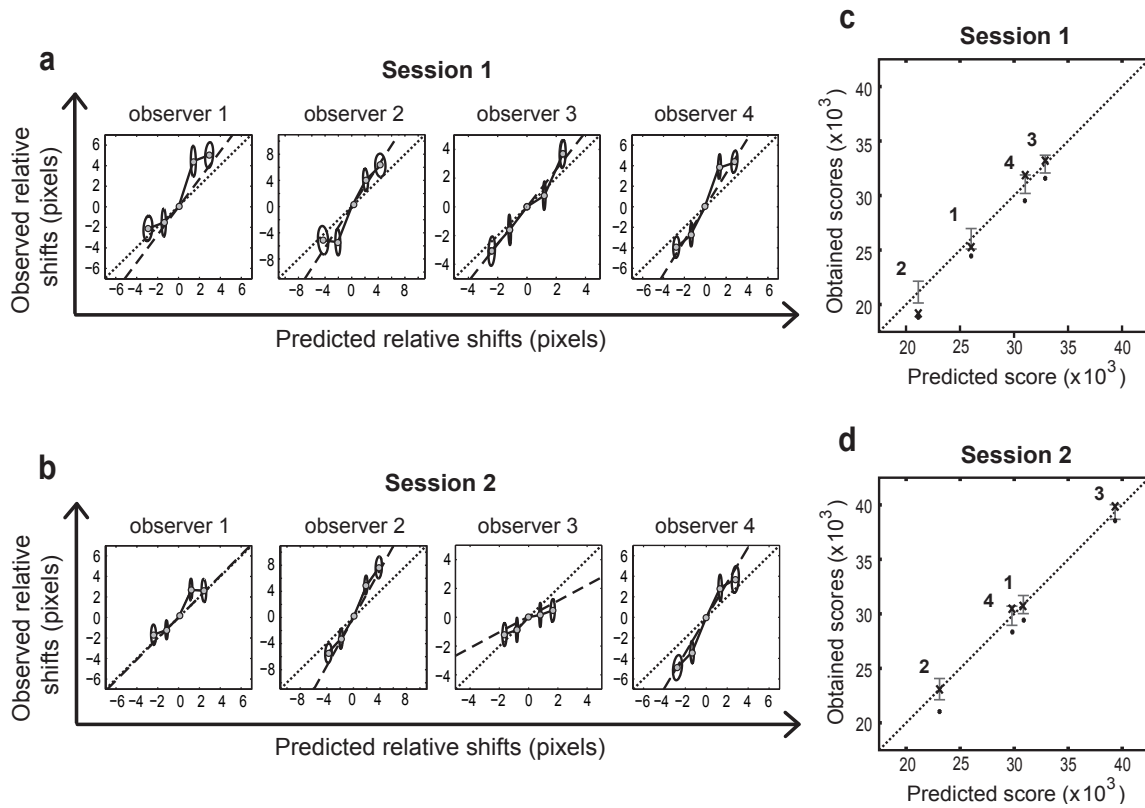


Figure 4: **Comparison of predicted and observed behaviour.** **a, b,** Predicted and observed relative shifts in the centres of the psychometric curves, for the first (**a**) and second (**b**) sessions. If performance is quantitatively optimal (up to a constant bias), the data points (grey circles) should lie on the dotted identity line. The ellipses show the 2σ covariance expected due to sampling errors, and the dashed line is a linear fit to the data points, computed by minimizing the weighted squared error in the plane with respect to these covariances. All observers showed the predicted pattern of shifts, but were not quantitatively exact. **c, d,** The mean and variance of the score that would be obtained if each observer behaved optimally given the directional bias was calculated for the first (**c**) and second (**d**) sessions (see Appendix A.2). Crosses plot predicted against observed score, with observers numbered as in **a** and **b**. The identity line again represents optimal performance (given the directional bias), and the vertical bars show one standard deviation from the mean. All points are within this range except for one observer in the first session. Filled circles show the mean score expected if observers failed to shift the centre of their psychometric curves from the biased centre of the curve for $\alpha = 0.5$, and all such points lie outside the predicted range.

such that there is little benefit from an exact quantitative match to the predicted curve shifts.

Our model of the psychometric curve was a composite function based on an underlying cumulative

Gaussian, which gives the probability of answering ‘right’ for a particular value of the stimulus x and α (see Equations 15 & 17). We used this model to compute the mean and standard deviation of the score observers would have obtained had they shifted their psychometric curves by the optimal amount from the biased mean of the curve for the $\alpha = 0.5$ condition (see Appendix A.2 for details). We were then able to ask whether the true score was statistically distinguishable from this predicted value.

Scores for all observers fell within one standard deviation of the predicted score in the second session (Fig. 4d), as did all but one in the first session (Fig. 4c). The failure of this one observer to obtain a score in this range in the first session could be due to cognitive interference or motivational issues. To test the sensitivity of total score as a measure of optimality, we computed the mean score that would have been obtained had participants failed to shift their psychometric curves from the biased central point. All such points lay outside the predicted range, as shown in Fig. 4c and d. This analysis suggests that, whilst not quantitatively optimal, the observed shifts were sufficient to obtain a score well within the predicted range.

Changes in performance

As described above, each observer participated in two experimental sessions on different days. We predicted that the level of observers’ internal uncertainty might have differed between the two sessions, due to perceptual learning, consolidation or extrinsic factors. If true, this would provide a further test of the hypothesis that observers’ behaviour is driven by internal beliefs that accurately reflect their sensory noise. If that sensory noise were to change, their beliefs, and thus their behaviour under the asymmetric loss function, should change concomitantly.

We first established that the level of observers’ sensory noise did, in fact, appear to change, as would be reflected by a change in the slope of the psychometric curve. We used Bayesian model selection (see Appendix A.1) to quantitatively compare models with the same slope in the two sessions to models in which the slope could differ. Table 2 shows that the model with different slopes in the two sessions was

overwhelmingly preferred in all cases, although, *within* each session, the model with the same slope for different loss functions was still the most probable. Thus, despite the apparent change in sensory noise, the basic prediction that the shape of the psychometric curve is unaffected by the loss function is confirmed.

In general, the slope of the psychometric curve was steeper in the second session, and observers' behaviour altered in accordance with the predictions of the Bayesian analysis. This can be seen as a trend towards smaller shifts in the second session (compare Fig. 4a and b), and towards higher scores (compare Fig. 4c and d). In particular, the three observers whose scores were in the predicted range in both sessions maintained this performance in the face of a clear change in apparent sensory noise. Furthermore, had the three observers who showed substantial changes in accuracy between the two sessions retained the same relative shifts in the second session as in the first, their expected scores would have fallen outside the optimal ranges shown. This suggests that observers were indeed adopting an efficient strategy, taking into account both the level of uncertainty and the external loss function.

As discussed in the Methods, we did not attempt to distinguish between stimulus-centred *sensory* noise, and any stimulus-centred *decision* noise not modeled by the stimulus-independent lapse-rate parameter. However, we assumed that the majority of this stimulus-centred variation was in fact due to sensory noise, and treated it as such. If this assumption is incorrect, the measured slope may incorporate stimulus-centred 'decision' noise associated with integrating the loss function with the true uncertainty, and thus an increase in slope might reflect an improvement in task performance rather than a change in internal uncertainty. However, inspecting Fig. 4a and b shows that for only one observer did the slope of the linear fit to performance (the dashed line) get *closer* to the identity line in the second session, supporting the assertion that the internal uncertainty was changing, rather than the ability to perform the task. Indeed, the observer whose fit improved was the same observer who obtained a score outside the predicted range in the first session, and it seems possible that she *did* change her strategy.

Controls for feedback

In order to use an ideal observer analysis to conclude that observers represent and compute with the relevant uncertainty, it is crucial to rule out alternative strategies for obtaining near optimal performance that do not require such knowledge. In the present task, it is possible that trial-by-trial feedback, had we provided it, would have allowed observers to incrementally adjust a internal threshold (perhaps in proportion to the size of the penalty) until their payoff was optimized. This could have led to psychometric curves that looked very much like those we predict from the analysis above. Indeed, classic Psychophysical studies have used a similar paradigm with trial-by-trial feedback to demonstrate this kind of ‘optimal’ criterion selection (Green & Swets, 1989).

Previous studies of uncertainty that have used trial-by-trial feedback have dealt with a similar potential confound by looking for evidence of incremental threshold adjustment in the data (Trommershauser et al., 2003; Trommershauser et al., 2005). The alternative strategy, that of withholding feedback, was adopted by Kording and Wolpert (2004) in a sensorimotor task, although without any asymmetry in costs. Here, we chose to provide only occasional (every 15 trials) cumulative feedback during the testing blocks (see Methods). This provided motivation, but did not allow observers to behave optimally via trial-by-trial threshold adjustment.

However, even such scarce feedback does provide some limited information about sensory noise, so we performed control analyses to confirm that the magnitude of the cumulative feedback had no measurable effect on behaviour. First, we fit a psychometric curve to all data which followed ‘good’ feedback (i.e. a cumulative score for the preceding 15 trials which fell above the 75th percentile), and to all data which followed ‘bad’ feedback (i.e. a cumulative score for the preceding 15 trials which fell below the 25th percentile). There were no feedback-related trends in the data (data not shown). To test for effects that might have been lost in averaging in this technique, we then examined whether ‘good’ feedback reinforced the direction of any changes in threshold, and whether ‘bad’ feedback reversed the direction of any such changes. If observers were modifying their behaviour in this way, we would

expect a positive correlation in threshold changes following ‘good’ feedback, and a negative correlation following ‘bad’ feedback. However, we observed only a slight negative correlation in both cases (data not shown).

DISCUSSION

Uncertainty inescapably affects almost all domains of brain function, arising due to variability in external processes, due to the under-constrained nature of many problems of perceptual inference and motor planning, from variability in motor execution, and due to noise in sensory processing. A long-standing and fundamental question in neuroscience is whether, and if so how, the brain takes account of this uncertainty in the course of perception, decision making, action and learning.

Our results address one aspect of this question. We show that observers possess an internal model of the visual processing uncertainty that affects their Vernier judgments, and that they use this model to guide their decisions. Crucially, observers’ decisions are sensitive to uncertainty even when they do not receive significant feedback about their accuracy or score. This indicates that the uncertainty-sensitive decision strategy is not learnt during the experiment itself, but is instead based on a pre-existing, implicit model of current internal uncertainty, that is presumably available at all times. In addition, observers’ decisions, and thus their models of internal uncertainty, track the changes in uncertainty which are associated with varying levels of sensory noise in different experimental sessions. Taken together, these observations suggest that the processing of uncertainty is a fundamental aspect of sensory computation. Furthermore, in our experiment observers must combine knowledge of this uncertainty, rather than simply a modal estimate of the stimulus, with an externally imposed loss function to perform well. Our results therefore also indicate that information about early sensory uncertainty, at least in the form of a two-alternative likelihood ratio between the models for left and right displacement, is preserved and made available to decision areas. This implies that the relevant information is propagated across multiple cortical layers.

Our new results join a growing body of work showing how uncertainty in its various forms affects behaviour (see introduction and Knill & Pouget, 2004, for a review). In the task reported here we aimed to show a ‘minimal context’ in which Bayesian optimality can be demonstrated, and in which alternative strategies that don’t require the representation of uncertainty can be ruled out. We therefore used a purely visual stimulus, and asked observers to make a simple Vernier offset judgement (Westheimer, 1979). We chose a Vernier task, as the requisite sensory processing is most likely to occur early in the visual pathway, perhaps principally in the relatively well-understood primary visual cortex. Combined with the existing evidence for optimality in crossmodal, motor, and visual cue combination experiments, our results support the claim that uncertainty is processed throughout the brain, even for simple, low-level visual quantities. Furthermore, such a task is a strong candidate for future integration with physiological data and with theoretical work concerning how neuronal populations represent and compute with uncertainty (Pouget, Dayan, & Zemel, 2003; Knill & Pouget, 2004).

Another key property of our task is that the set of visual stimuli was fixed, so that all stimulus-related uncertainty arose exclusively from visual processing, corresponding to ‘internal noise’ in psychophysical experiments (Green & Swets, 1989). This is in contrast to many previous studies of sensory uncertainty, where variability was driven by external manipulations, such as the random placement of dots or the addition of corrupting noise (but see Stocker & Simoncelli, 2006). Using a fixed stimulus set thus strengthens the conclusion that the mechanisms exposed are fundamental to sensory processing, rather than being limited to strategies for dealing with uncertainty in the external world.

There are many technical issues with Bayesian optimality experiments, which can obscure optimal behaviour and limit conclusions about the underlying computation. For example, with monetary loss functions observers often demonstrate over compensation, a reluctance to make extreme shifts, and failure to keep track of the current loss function (Green & Swets, 1989; Landy et al., 2007). We therefore used values of α that demanded relatively small shifts, and used practice trials and a blocked design with regular reminders of the current cost values.

Recent work using a similar loss function approach in an unspeeded visual orientation estimation task (Landy et al., 2007) found evidence for optimality, but also for a variety of suboptimal strategies. In psychological ‘betting’ paradigms failures of probabilistic reasoning characterise human behaviour, and it may be that unspeeded adjustment tasks are more vulnerable to such influences. Our task is very simple, and requires a forced-choice categorization response that observers performed very quickly. It seems likely that we therefore avoid some of these cognitive effects, making it easier to model observers’ use of information about their own uncertainty. However, as discussed above we do see evidence for over-compensation, and for relatively smaller over-compensation for larger cost asymmetries.

In our analysis we did not attempt to distinguish between stimulus-centred *sensory* noise, and any stimulus-centred *decision* noise not modeled by the lapse-rate parameter. However, we assumed that the majority of this stimulus-centred variation was due to sensory noise. This assumption was supported by examining the results across sessions — the ability of observers to choose optimally in the face of asymmetric costs did not change as their ability on the task, measured by the slope of the psychometric function, did (see Fig. 4a and b). Using external manipulations to produce randomly intermixed uncertainty levels on each trial would allow us to separate more explicitly any stimulus-centred decision noise from uncertainty due to sensory processing. However, Landy et al. (2007) found more suboptimality when levels of uncertainty were randomly intermixed rather than blocked as in our experiment, and it is not clear *why* this should be the case. It could reveal limits on the ability to perform online Bayesian processing, or alternatively arise from corrupting cognitive or psychological factors. In the present study we were not interested in trying to delineate these factors, and so used a blocked design. In addition, we wanted to retain the property of all stimulus uncertainty arising internally. Future investigations of the effect of task and experimental design on optimality are important to pull apart optimal sensory processing from cognitive reasoning effects.

In the present study we aimed to demonstrate minimal conditions for Bayes optimal behaviour under uncertainty. We showed that observers approach the quantitatively optimal strategy given a directional bias, and score within the predicted range, in a simple unimodal visual task that requires

them to integrate a model of their internal uncertainty with an external loss function. The assumptions of the model, and the predictions that arise from them, were tested, and we took care to rule out alternative strategies for achieving the observed behaviour. Our results therefore support the assertion that the processing of uncertainty is a fundamental aspect of sensory computation, and can be used to inform subsequent decision-making processes.

APPENDICES

A.1 Bayesian Model Comparison

A different psychometric function was obtained for each observer, for each value of α , in each session. In principle, each such function might have entirely different parameters. Our expectation, however, based on the Bayesian observer analysis, is that the slope parameter should remain constant as α varies within one session, although it may change between sessions. We used Bayesian model comparison to determine which group of shared parameters was best supported by our data, by evaluating an approximation to the marginal likelihood or ‘evidence’ for a number of different models. This approach to choosing an appropriate model originates with Jeffreys (Jeffreys, 1939), and incorporates an Occam’s razor-type penalty for models with more parameters (Gull, 1988; Kass & Raftery, 1995; Mackay, 2004). In the absence of prior bias towards any particular model, the marginal likelihoods are proportional to the probabilities of each model being the one from which the data arose.

For each observer we fit models that shared parameters between sessions (Supplementary Table 2, lower half), and models that had independent parameters for each session (Supplementary Table 2, upper half). In the latter case the total log evidence was the sum of the log evidences obtained for each session. For each case, we fit models with individual ρ_j and ϵ_j parameters for each α , models with the values of ρ_j and ϵ_j restricted to have the same value for all α , and models with one parameter

restricted while the other was allowed to vary. The centres μ_j always varied with α , as all data sets showed very clear shifts.

A gradient ascent procedure was used to find the most probable or *maximum a posteriori* (MAP) parameter values, given the data, under a non-informative prior. As the exact evidence could not be calculated, these MAP parameters for each model were then used to compute a Laplace approximation to the marginal likelihood in each case. The Laplace approximation results from taking the first three terms of a Taylor expansion about the MAP parameters (Mackay, 2004). In the equations below, D refers to the data, m to the model, θ to the vector of all parameters, θ^* to the MAP parameters, and d to the number of parameters in the model. The matrix A is the Hessian of the log posterior, i.e. the matrix of second partial derivatives of $\log P(\theta|D, m)$ with respect to θ , evaluated at θ^* .

$$\log P(D|m) = \log \int d\theta P(D, \theta | m), \quad (20)$$

$$\log P(D|m) \approx \log P(D|\theta^*, m) + \log P(\theta^*|m) + \frac{d}{2} \log 2\pi - \frac{1}{2} \log |A|. \quad (21)$$

The values of the Laplace approximation for each of the four models are shown in Table 2. The highest evidence, corresponding to the ‘best model’, is highlighted in bold for each observer and each case. In accordance with our assumption, the best model had a single ρ parameter for all α , whether or not this slope was the same or different between sessions. A single lapse parameter was best for two observers and separate lapse parameters for the other two. As mentioned above, the lapse rate is incorporated in the model to account for decision noise, motor errors, and moments of inattention, and it seems reasonable that, whilst the internal uncertainty is the same for each α value, such random lapses might vary.

In addition, the model with different parameters for the two sessions was always preferred, suggesting that observers’ sensory noise changed between the two experimental sessions. The values of the

Summed Laplace approximation for individual session models				
	single ρ		separate ρ	
observer	single ϵ	separate ϵ	single ϵ	separate ϵ
1	1995	1997	1970	1975
2	1602	1638	1613	1619
3	2351	2338	2346	2335
4	2061	2033	2055	2041
Laplace approximation for pooled session model				
	single ρ		separate ρ	
observer	single ϵ	separate ϵ	single ϵ	separate ϵ
1	1851	1841	1827	1830
2	1433	1487	1446	1441
3	2200	2185	2197	2181
4	1944	1948	1941	1929

Table 2: Results of Bayesian Model Selection: Laplace approximation to marginal log likelihood for each of four models for each observer. Bold text shows the model with the highest log likelihood for each participant. It should be noted that each unit difference in log likelihood corresponds to an e -fold ratio of model probabilities.

parameters fit to the best model for each observer, in each session, are given in Table 3.

A.2 Computation of Optimal Score Range

To assess whether observers' performance was significantly different from optimal, we computed the mean and standard deviation of the scores that they would have obtained under the optimal strategy, given their apparent internal uncertainties and their observed biases. The total score ('reward') for one session is obtained by summing the scores for each α value;

$$\mathcal{R}_{total} = \sum_j \mathcal{R}_j. \quad (22)$$

The score for each α value is given by the sum, over the different possible stimulus offsets x_i , of the number of trials on which the observer answers 'right' and 'left' correctly and incorrectly, multiplied by

		Session 1				Session 2			
observer	α	μ (pixels)	ρ (1/pixels)	σ (pixels)	ϵ (prob.)	μ (pixels)	ρ (1/pixels)	σ (pixels)	ϵ (prob.)
1	0.7	-4.0	}0.072	}5.6	0.052	-3.8	}0.086	}4.6	0.031
	0.6	-3.5			0.025	-3.3			0.00002
	0.5	-1.9			0.13	-2.1			0.017
	0.4	2.5			0.16	0.55			0.031
	0.3	3.1			0.045	0.63			0.072
2	0.7	-7.2	}0.049	}8.2	0.019	-7.3	}0.055	}7.2	0.013
	0.6	-7.5			0.0019	-5.1			0.00012
	0.5	-2.1			0.11	-1.8			0.10
	0.4	2.0			0.021	3.1			0.10
	0.3	4.2			0.0098	5.8			0.016
3	0.7	-2.3	}0.087	}4.6	}0.017	-0.92	}0.13	}3.1	}0.011
	0.6	-0.79				-0.60			
	0.5	0.82				0.28			
	0.4	1.6				0.44			
	0.3	4.5				0.76			
4	0.7	-3.5	}0.075	}5.3	}0.0069	-3.6	}0.079	}5.1	}0.042
	0.6	-2.3				-2.1			
	0.5	0.45				1.3			
	0.4	4.3				4.2			
	0.3	4.8				5.0			

Table 3: Results of Model Fitting to Experimental Data: centre (μ), slope (ρ), and lapse (ϵ) parameters for each observer in each α condition and each session (values given to 2 significant figures). For each observer there is a separate μ for each α condition, representing the centre of the psychometric function in pixels. However, there is only a single ρ for all α conditions, representing the fact that the observer's internal uncertainty is the same regardless of the value of α . The Gaussian standard deviation in pixels corresponding to these values of ρ is given in the next column. Two observers have a single ϵ for all α , and two have separate ϵ for each α . These constraints on parameter values were determined via Bayesian model comparison (see Table 2)

the appropriate reward or cost parameter. Using the same definitions of N_{ij} and n_{ij} as in Equation 15, this is

$$\mathcal{R}_j = \sum_{i | x_i < 0} (\mathcal{R}_l (N_{ij} - n_{ij}) + \mathcal{C}_r n_{ij}) + \sum_{i | x_i > 0} (\mathcal{R}_r n_{ij} + \mathcal{C}_l (N_{ij} - n_{ij})). \quad (23)$$

Under our model, n_{ij} is binomially distributed with mean $N_{ij} p_{ij}$, where p_{ij} is given by the psychometric function (Equation 17). To obtain the expected score under the optimal strategy (constrained by the observed bias), we evaluated Equation 16 for each offset, using the measured value of ρ , but using the optimal relative value of μ_j obtained by adding the optimal relative shift to the observed bias in the symmetric condition (*i.e.*, $\mu_{0.5} + \Delta\mu_j^*$). Calling these optimal relative values p_{ij}^* , the expected score under the constrained optimal strategy is,

$$\langle \mathcal{R}_{total} \rangle = \sum_j \langle \mathcal{R}_j \rangle, \quad (24)$$

with

$$\langle \mathcal{R}_j \rangle = \sum_{i | x_i < 0} N_{ij} (\mathcal{R}_l (1 - p_{ij}^*) + \mathcal{C}_r p_{ij}^*) + \sum_{i | x_i > 0} N_{ij} (\mathcal{R}_r p_{ij}^* + \mathcal{C}_l (1 - p_{ij}^*)). \quad (25)$$

To compare the measured scores to this value, we must also calculate the variance in the score that is to be expected as decisions vary due to sensory noise. As the score in each condition is independent, this variance is given by

$$\text{Var}(\mathcal{R}_{total}) = \sum_j \langle \mathcal{R}_j^2 \rangle - \langle \mathcal{R}_j \rangle^2, \quad (26)$$

The second moment above can also be computed in closed form, using the expression for \mathcal{R}_j in Equation 23, the binomial mean as before, and the binomial second moments:

$$\langle n_{ij} n_{i'j} \rangle = N_{ij} N_{i'j} p_{ij}^* p_{i'j}^* + \delta_{ii'} N_{ij} p_{ij}^* (1 - p_{ij}^*), \quad (27)$$

where $\delta_{ii'}$ is the Kronecker delta. We obtain

$$\langle \mathcal{R}_j^2 \rangle = \langle \mathcal{R}_j \rangle^2 + (\mathcal{C}_r - \mathcal{R}_l)^2 \sum_{i | x_i < 0} N_{ij} p_{ij}^* (1 - p_{ij}^*) + (\mathcal{R}_r - \mathcal{C}_l)^2 \sum_{i | x_i > 0} N_{ij} p_{ij}^* (1 - p_{ij}^*), \quad (28)$$

and so the expected variance in score is

$$\text{Var}(\mathcal{R}_{total}) = \sum_j \left((\mathcal{C}_r - \mathcal{R}_l)^2 \sum_{i | x_i < 0} N_{ij} p_{ij}^* (1 - p_{ij}^*) + (\mathcal{R}_r - \mathcal{C}_l)^2 \sum_{i | x_i > 0} N_{ij} p_{ij}^* (1 - p_{ij}^*) \right). \quad (29)$$

The corresponding standard deviation is shown in Figure 4.

Acknowledgements We thank P. Dayan, J. Solomon and M. Landy for comments. This work was supported by grants from the Wellcome Trust (to L.W.) and the Gatsby Foundation (to M.S.)

Author Information The authors declare no competing financial interests. Correspondence and requests for material should be addressed to L.W. (email: louisew@gatsby.ucl.ac.uk) or M.S. (email: maneesh@gatsby.ucl.ac.uk).

References

- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. New York: Springer Verlag.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436.

- Carpenter, R. H. S. (1988). *The movements of the eyes*. London: Pion Ltd.
- Cox, R. T. (1961). *The algebra of probable inference*. Baltimore, MD: Johns Hopkins University Press.
- Deneve, S., Latham, P. E., & Pouget, A. (2001). Efficient computation and cue integration with noisy population codes. *Nature Neuroscience*, *4*(8), 826–831.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.
- Green, D. M., & Swets, J. A. (1989). *Signal detection theory and psychophysics*. Los Altos, CA: Peninsula Publishing.
- Gull, S. F. (1988). Bayesian inductive inference and maximum entropy. In G. Erickson, & C. R. Smith (Eds.), *Maximum-entropy and Bayesian methods in science and engineering: Volume 1: Foundations*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Helmholtz, H. L. F. (1856). *Treatise on physiological optics*. Thoemmes Continuum.
- Hillis, J., Watt, S., Landy, M., & Banks, M. (2004). Slant from texture and disparity cues: optimal cue combination. *Journal of Vision*, *4*, 967–992.
- Hodgson, T. L. (2002). The location marker effect - saccadic latency increases with target eccentricity. *Experimental Brain Research*, *145*(4), 539–542.
- Jacobs, R. A. (1999). Optimal integration of texture and motion cues to depth. *Vision Research*, *39*(21), 3621–3629.
- Jeffreys, H. (1939). *Theory of probability*. Oxford: Oxford University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773–795.
- Knill, D. C., & Pouget, A. (2004). The bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719.
- Knill, D. C., & Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research*, *43*(24), 2539–2558.

- Kording, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.
- Landy, M., Goutcher, R., Trommershauser, J., & Mamassian, P. (2007). Visual estimation under risk. *Journal of Vision*, *7*(6), 1–15.
- Landy, M. S., & Kojima, H. (2001). Ideal cue combination for localizing texture-defined edges. *Journal of the Optical Society of America a-Optics Image Science and Vision*, *18*, 2307–2320.
- Mackay, D. J. C. (2004). *Information theory, inference, and learning algorithms*. Cambridge: Cambridge University Press.
- Maloney, L. T. (2002). Statistical decision theory and biological vision. In D. Heyer, & R. Mausfeld (Eds.), *Perception and the physical world: Psychological and philosophical issues in perception*. New York: Wiley.
- Marr, D. (1982). *Vision*. San Francisco, CA: W. H. Freeman and Company.
- McCullagh, P., & Nelder, J. (1989). *Generalised linear models*. London: Chapman Hall.
- Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.
- Pouget, A., Dayan, P., & Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience*, *26*, 381–410.
- Saunders, J. A., & Knill, D. C. (2004). Visual feedback control of hand movements. *Journal of Neuroscience*, *24*(13), 3223–3234.
- Saunders, J. A., & Knill, D. C. (2005). Humans use continuous visual feedback from the hand to control both the direction and distance of pointing movements. *Experimental Brain Research*, *162*(4), 458–473.
- Schwartz, O., Sejnowski, T. J., & Dayan, P. (2005). A bayesian framework for tilt perception and confidence. *Advances in Neural Information Processing Systems*, *17*.

- Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, *9*(4), 578–585.
- Tassinari, H., Hudson, T., & Landy, M. (2006). Combining priors and noisy visual cues in a rapid pointing task. *Journal of Neuroscience*, *26*(40), 10154–10163.
- Thurstone, L. L. (1927). A law of comparative judgement. *Psychological Review*, *34*, 273–286.
- Trommershauser, J., Gepshtein, S., Maloney, L. T., Landy, M. S., & Banks, M. S. (2005). Optimal compensation for changes in task-relevant movement variability. *Journal of Neuroscience*, *25*(31), 7169–7178.
- Trommershauser, J., Maloney, L. T., & Landy, M. (2003). Statistical decision theory and trade-offs in the control of motor response. *Spatial Vision*, *16*(3-4), 255–275.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, *5*(6), 598–604.
- Westheimer, G. (1979). Spatial sense of the eye – Proctor lecture. *Investigative Ophthalmology and Visual Science*, *18*(9), 893–912.
- Wichmann, F., & Hill, N. (2001). The psychometric function: I. fitting, sampling, and goodness of fit. *Perception and Psychophysics*, *63*(8), 1293–1313.
- Yuille, A., & Bulthoff, H. (1996). Bayesian decision theory and psychophysics. In D. Knill, & W. Richards (Eds.), *Perception as bayesian inference*. Cambridge, UK: Cambridge University Press.