Supporting Information

de Gardelle and Summerfield 10.1073/pnas.1104517108

SI Methods

Stimuli, Design, and Task. On each trial, eight color values (denoted C) and eight shape values (denoted S) were drawn randomly from normal distributions with prespecified means (mc, ms) and SDs (σc, σs). For each feature dimension, the mean could take one of four values (e.g., for color, from really red to slightly red to slightly blue to really blue), and the SD could take one of three values (low, medium, high). See Fig. S1 for a representation of the design.

The mean factor was then collapsed from four levels to two levels, by averaging across response categories: "slightly red" and "slightly blue" were considered together as "weak evidence" ("low mean") conditions, and "really red" and "really blue" were considered together as "strong evidence" ("high mean") conditions.

We parameterized the shape and color values (C and S) of individual elements between 0 and 1, and the means (mc, ms) and SDs (σc, σs) were controlled in a generic manner as

$$m = m_{\text{ref}} + m_{\text{categ}}\ (+1 \text{ or } -1, \text{ for left vs. right}) \times m_{\text{condition}}\ (1 \text{ or } 2, \text{ for low vs. high mean}) \times m_{\text{scaling}},$$
$$\sigma = \sigma_{\text{condition}}\ (1, 2, \text{ or } 3, \text{ for low, medium, and high SD}) \times \sigma_{\text{scaling}},$$

where m_{scaling} and σ_{scaling} are scaling values adjusting the spread of the conditions in the parameter space, and m_{ref} is the position in the parameter space (i.e., the interval [0, 1]) of the transition between left and right response categories. Specifically, we used $m_{\text{ref}} = 0.5$ for the red/blue task and the square/circle task (experiments 1–3), and we used $mc_{\text{ref}} = 0.75$ for the red/purple task (experiment 4a) and $mc_{\text{ref}} = 0.25$ for the purple/blue task (experiment 4b).

In the following, we denote by x_k the evidence value of the element k. That is, in the color task we have $x_k = C_k - m_{\text{ref}}$, and in the shape task we have $x_k = S_k - m_{\text{ref}}$ for each element. The evidence value is the position of the decision-relevant parameter relative to the decision boundary. It can be positive or negative, and for instance in the red/blue color task x_k takes positive values for redder elements and negative values for bluer elements, in the red-blue continuum presented in Fig. S1.

In experiment 1, four levels for mc, three levels for σc, and four levels for ms were manipulated orthogonally ($mc_{\text{scaling}} = 0.05$, $\sigma c_{\text{scaling}} = 0.05$, $ms_{\text{scaling}} = 0.05$). In each trial the value of σs was taken to be the same as the value of σc. In experiment 2, four levels for mc and three levels for σc were manipulated orthogonally ($mc_{\text{scaling}} = 0.025$, $\sigma c_{\text{scaling}} = 0.0333$). The shape parameters were kept constant ($ms = 0$, $\sigma s = 0.05$). In experiment 3, four levels for mc, three levels for σc, four levels for ms, and three levels for σs were manipulated orthogonally ($mc_{\text{scaling}} = 0.05$, $\sigma c_{\text{scaling}} = 0.0666$, $\mu s_{\text{scaling}} = 0.05$, $\sigma s_{\text{scaling}} = 0.0666$). In experiment 4, four levels for mc, three levels for σc, four levels for ms, and three levels for σs were manipulated orthogonally ($mc_{\text{scaling}} = 0.0375$, $\sigma c_{\text{scaling}} = 0.05$, $ms_{\text{scaling}} = 0.05$, $\sigma s_{\text{scaling}} = 0.0666$).

Data Analyses. The main analyses focused on the percentage of errors as well as response times (RTs) for correct trials. Then error rates and RTs were averaged for each subject and conditions ($2\mu \times 3\sigma = 6$ conditions) and submitted to between-subject ANOVAs. Three participants (in experiment 1) were excluded from the group analyses because they exhibited an average performance around chance level.

Additionally, we conducted logistic regression analyses to assess how participants weighted the different pieces of evidence (i.e., the eight values of the decision-relevant parameter) to make their choice. These analyses were carried out in two steps, by assessing the weight of evidence for each subject and then using t tests or ANOVAs at the group level. For each subject level, we estimated the coefficients for a generalized linear regression in which the decision-relevant values of the stimuli predicted the subject's responses. The responses were coded in a vector of dimensions (trials × 1) vector containing 0 for one category and 1 for the other. The predictors were coded as a vector of dimensions (trials × 9) matrix constituted by the decision-relevant values (i.e., the C values in the color task, the S values in the shape task) coded between 0 and 1 (consistently with the response codes) plus a constant term to estimate the intercept. Using probit or logistic models to estimate the beta coefficients produced equivalent results. Under the probit model, responses Y were taken to follow a binomial distribution with a probability P affected by the predictors in X, such that $P = \text{Phi}(X' \text{ beta})$ where Phi is the cumulative distribution function of the standard normal distribution. Under the logistic model, the betas are estimated such that $\log(P/(1-P)) = X' \text{ beta}$. This analysis was performed using the glmfit function in Matlab. Trials in which no response was given were not included in these analyses.

Crucially, the values of the eight elements were sorted before including them as predictors in the regression, so that the beta coefficients would correspond to the weighting of the evidence, respectively for each ranked element in the array.

For each participant we normalized the eight weights by dividing them by their root mean square. Importantly this normalization was neutral with respect to the finding that different elements might receive different weight of evidence. We used this normalization to downweight the contribution of individual subjects for which the estimation of the weights would be too unreliable (which might happen for a particular participant in a particular condition, if the number of errors in the trials is too small).

When comparing the treatment of inlying elements and outlying elements, we simply averaged the resulting normalized weights separately for elements ranked 1, 2, 7, and 8 (outlying) and for elements ranked 3, 4, 5, and 6 (inlying). This procedure was done subject by subject and we then used T-tests or ANOVA at the group level.

Computational Models. We simulated three random-walk models of decision making. These models were all based on the same "evidence accumulation up to a bound" principle. On each trial, a decision variable (DV) is formed by accumulating momentary evidence (A) corrupted by noise (a Gaussian random variable of variance c^2 and mean 0). A choice is made when this decision variable reaches a predefined boundary value (Z for one choice, $-Z$ for the other choice). Note that we used a simplified version of this type of model, in which the accumulation starts at zero, and the boundary conditions for the two choices are symmetrical. This is reasonable because choices in favor of one category are equally likely as choices in favor of the other category (i.e., there are as many "red" trials as "blue" trials), and participants do not show any bias toward one of the response categories.

Such a model can be formalized generically in the following way:

Start with no prior: $\text{DV}_0 = 0$.
Then for $t > 0$, the time in the number of simulation cycles, do
$\text{DV}_t = \text{DV}_{t-1} + A + N(0, c^2)$.

Then check whether the DV reaches one of the bounds at time t: If DV $\geq Z$ (resp. DV $\leq -Z$), then consider R_1 (resp. R_2) as the choice and t as the RT and stop the accumulation process.

For all models, our approach was to present the model with the same numerical input values as in the experimental data and adjust the numerical value of the bound and noise parameters separately for each subject to fit the subject's error rates across the 2 means × 3 variance conditions. Fitting was done by minimizing the mean square error over a large 2D search space for possible values for the noise deviance c and the bound Z. Then, once the best values for c and Z were found for each participant, the simulated RTs were rescaled from cycles to seconds by adjusting the mean and SD of the RTs in cycles to the participant's mean and SD, to allow for comparing simulated and real RTs within the same units and on the same graph. This rescaling of RTs was, critically, neutral to our hypothesis as it did not affect the relative pattern across conditions of simulated RTs and left the trial-by-trial choices of the model (and thus the weighting function) unchanged.

The models we simulated differed only in the way they increment the input values, that is, in the way the value of A was derived from the value of the elements x_k (we remind the reader that x_k denotes the position of the element k relative to the boundary between R_1 and R_2). We also denote by μ the mean of the evidence values of the sample and by σ the SD of these evidence values:

$$\mu = \frac{1}{n} \times \sum_{k=1}^{n} x_k \text{ and } \sigma = \sqrt{\frac{1}{n-1} \times \left(\sum_{k=1}^{n} (x_k - \mu)^2 \right)}.$$

In our three models, the increment A is defined in three different ways:

i) In the simple averaging model: $A = \mu$.
ii) In the SNR model: $A = \frac{\mu}{\sigma}$.
iii) In the LPR model: $A = \frac{1}{n} \times \sum_{k=1}^{n} \text{LPR}(x_k)$, where LPR($x$) relates the value x to the log of the posterior probability ratio between the two possible choice options (see details below).

LPR model. This model is an adaptation of the log-likelihood ratio (LLR) model, already described in previous studies, to our case of multiple elements (e.g., ref. 1). This model considers the parameter values not in their native space but in a probability space, which is shaped by the relation between the elements' values and the underlying state of the world. The idea is that if the value x in the parameter space is presented several times and if each presentation is paired with an indication about the underlying state of the world (R_1 or R_2), one can estimate the probabilities of R_1 or R_2 given the value x.

The LLR is an optimal decision variable for a bounded accumulation process, as it transforms it into a sequential probability ratio test (SPRT), which gives the fastest decision times for a given level of accuracy (2, 3). The estimation of the LLR can be done simply by computing the log posterior ratio (LPR) and seeing that in our case we have LPR = LLR.

To see why, let us write Bayes' rule for the posterior probabilities $p(R_1|x)$ and $p(R_2|x)$:

$$\frac{p(R_1|x)}{p(R_2|x)} = \frac{p(x|R_1) \times p(R_1)/p(x)}{p(x|R_2) \times p(R_2)/p(x)}.$$

Considering that R_1 and R_2 have equal prior probabilities [i.e., $P(R_1) = P(R_2)$], this simplifies as

$$\frac{p(R_1|x)}{p(R_2|x)} = \frac{p(x|R_1) \times p(R_1)}{p(x|R_2) \times p(R_2)} = \frac{p(x|R_1)}{p(x|R_2)}.$$

That is, the posterior ratio and the likelihood ratio are the same. Taking the log, we have

$$\text{LPR}(x) = \ln\left(\frac{p(R_1|x)}{p(R_2|x)}\right) = \ln\left(\frac{p(x|R_1)}{p(x|R_2)}\right) = \text{LLR}(x).$$

Additionally, given that R_1 and R_2 are the only two options [i.e., $P(R_1|x) + P(R_2|x) = 1$], we have

$$\text{LPR}(x) = \ln\left(\frac{p(R_1|x)}{1 - p(R_1|x)}\right).$$

Estimation of the LPR for each element. Thus, to simulate this model, we first estimated the LPR function of x (where x is a value in the decision space). To do so, we used a simple approach based on empirical probabilities (for a related approach, see supplementary material in ref. 1) to assess $p(R_1|x)$ and then used the equation above to convert this probability in a LPR. This procedure was done for each subject separately.

We derived the LPR function as follows. We considered all elements presented to the subject as associated with either R_1 or R_2 (according to the feedback). The goal is to determine a reasonable estimate of the posterior probability $p(R_1|x)$ for x varying along the parameter space. In each of 100 bins regularly spaced on the parameter space (after discarding 5% of extreme values for which the probability estimation involves very few data points) we computed $p(R_1|\text{bin})$ as the frequency of R_1 being the correct association for the elements falling in that bin [$p(R_1|\text{bin})$ = number of R_1/number of elements falling in that bin]. We fitted the resulting probabilities with a sigmoid function over the bins, to capture the shape of the probability profile. We chose the sigmoid function because it provided a better account than other profiles, notably because our design involved a mixture of 12 Gaussian distributions with different means and variances (indeed, theoretically, the inverse cumulative normal distribution would be the true profile in the case of only 2 symmetrical Gaussian distribution with equal variance). We finally applied the fitted function to all individual elements. Fig. S3 presents the LPR functions for a sample of individual subjects.

Optimality of the LPR and conditional independence assumptions. The sum of the log-likelihood ratios is the optimal decision variable to consider for a bounded accumulation model (2, 3). Because the likelihood ratios are equal to the posterior ratios (by Bayes' rule, when prior probabilities are equal, see above), the sum of the log of the posterior ratio is also the optimal decision variable. Additionally, because in our models the noise, bounds, and increments scale together (2), taking the summed or the averaged log ratio over the eight elements is equivalent. We favored the average to express all three models in similar forms.

Importantly, our estimation of the log posterior ratio described above is equal to the true log posterior ratio only under the conditions of conditional independence between the samples. To see this, let us write the true log posterior ratio of the full array:

$$\begin{aligned}\text{LPR}(x_1, x_2, \ldots, x_n) &= \ln\left(\frac{p(R_1|x_1, x_2, \ldots, x_n)}{p(R_2|x_1, x_2, \ldots, x_n)}\right) \\ &= \ln\left(\frac{p(x_1, x_2, \ldots, x_n|R_1)}{p(x_1, x_2, \ldots, x_n|R_2)}\right).\end{aligned}$$

When samples $x_1, \ldots, x_n$ are independent, then we have

$$p(x_1,x_2,\ldots,x_n|R_1) = p(x_1|R_1)\times p(x_2|R_2)\times \ldots \times p(x_n|R_n).$$

Incorporating this expression into the previous line, we have

$$\begin{aligned}\mathrm{LPR}(x_1,x_2,\ldots,x_n) &= \ldots = \ln\left(\frac{p(x_1,x_2,\ldots,x_n|R_1)}{p(x_1,x_2,\ldots,x_n|R_2)}\right)\\ &= \ln\left(\prod_{k=1}^{n}\left(\frac{p(x_k|R_1)}{p(x_k|R_2)}\right)\right).\end{aligned}$$

Rewriting the ln of the product as the sum of the ln, we see that the log posterior ratios for the individual elements appear in the expression

$$\mathrm{LPR}(x_1,x_2,\ldots,x_n) = \sum_{k=1}^{n}\ln\left(\frac{p(x_k|R_1)}{p(x_k|R_2)}\right) = \sum_{k=1}^{n}\mathrm{LPR}(x_k).$$

In this final line, we see that the decision variable used in the LPR model (i.e., the mean rather than the sum of the LPR over the elements) scales with the true log posterior ratio of the full array, if the samples are independent.

Thus, the decision variable we implemented is optimal only under this independence condition. This condition independence is not satisfied, as individual elements are taken from only one distribution (out of 12 possible distributions) in any trial. Thus even though they might be conditionally independent given the particular distribution presented on the current trial, they are not conditionally independent given the reward category. Consequently, the LPR model we implemented is technically not the "true" optimal model, but it rather represents the optimal process for an observer considering (inaccurately) the elements *as if* they were independent.

Stability over time. We note that our method to compute the LPR function uses all trials presented to the subject, thus using information available to the subject only at the end of the experiment, whereas the resulting LPR is used in all trials, including the ones at the very beginning of the experiment. This property seems problematic. However, we provide several arguments for why we think this method is reasonable. First, subjects do receive instructions at the beginning of the experiment and build very rapidly an understanding of the categories in the task, so implementing a mapping learned "from a blank slate" (i.e., a flat prior distribution) would be also unrealistic. Second, assessing the evolution of the LPR function over time (Fig. S3) seems to indicate that the mapping is stable quite soon after the experiment starts. Fig. S3 shows for a random sample of subjects the posterior probability function of the position of x in the parameter space, estimated using the ~1,000 trials of the whole experiment (red lines) or using only the first 300 trials of the experiment (blue lines). We also checked that the main results presented, and in particular the discarding of outliers, hold when only the first third of the experiment is considered. Given these points, we think that using a "stable" probability mapping computed over the whole experiment is a reasonable approximation for our simulations.

1. Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080.
2. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113:700–765.
3. Wald A, Wolfowitz J (1949) Bayes solutions of sequential decision problems. *Proc Natl Acad Sci USA* 35:99–102.

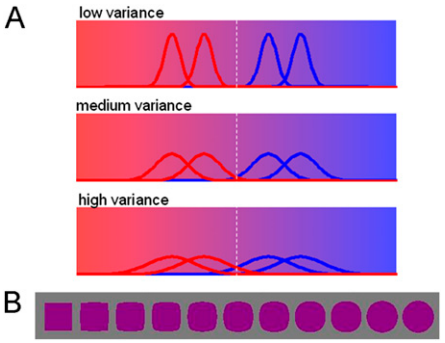

Fig. S1. (*A*) Schematic presentation of distribution of sensory information for the color manipulation. In each trial, 1 normal distribution was used to generate eight values. This source was one of the 12 possible distributions shown, corresponding to the 4 means × 3 SDs conditions of our design. Moving the mean away from categorical boundary (white dashed line) increased the evidence strength. The reliability of the evidence was manipulated in an orthogonal fashion, by using low (*Top*), medium (*Middle*), and high (*Bottom*) variance. (*B*) Parametric transition from a square to a circle using the "squircles" equations (see main text). Here 11 squircles are presented, from S = 0 (square) to S = 1 (circle), in step of 0.1 in the parameter space.

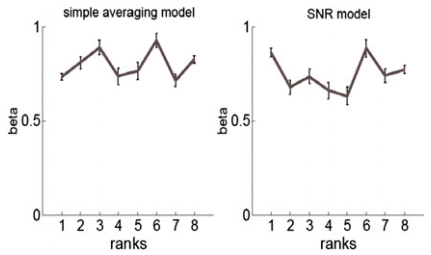

Fig. S2. Weighting profile across ranks for the simple averaging model (*Left*) and the SNR model (*Right*). In both cases, there was no effect of ranks on the beta coefficients [simple averaging model, $F(7,532) = 1.15$, $P = 0.29$; SNR model, $F(7,532) = 1.30$, $P = 0.26$], and in particular the beta coefficients were not different in inlying vs. outlying [simple averaging model, $t(77) < 1$, $P = 0.62$; SNR model, $t(77) < 1$, $P = 0.51$].

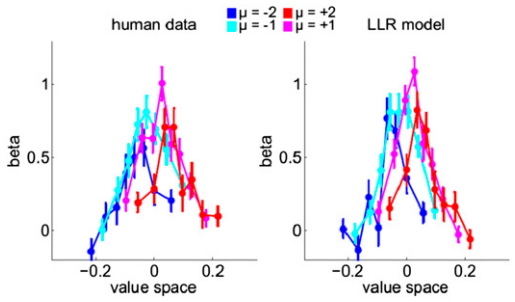

Fig. S3. Weighting profile across ranks for human observers (*Left*) and the LPR model (*Right*), as a function of the four possible categories of trial mean ($\mu = -2$, high mean evidence for left choices; $\mu = -1$, low mean evidence for left choices; $\mu = +1$, low mean evidence for right choices; $\mu = +2$, high mean evidence for right choices).

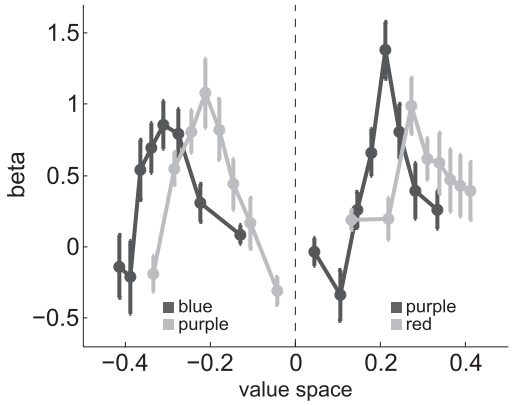

Fig. S4. Weighting profile across ranks in control experiment 4, with weights estimated separately for trials in which the evidence favored left (black) or right (gray) choices and separately for sessions in experiment 4 (*Left*, blue/purple task; *Right*, purple/red task). The *x*-axis positions correspond to the average value of the items (in the eight ranks), in the color space used in the previous experiments. The downweighting of outlying evidence values was observed in both portions of the color space [inlying vs. outlying: both $t_{(23)} > 4$, $P < 0.001$].

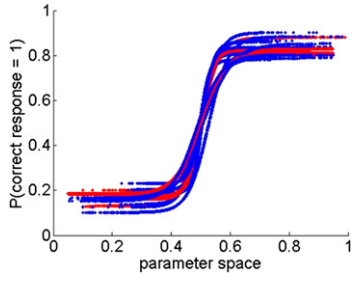

Fig. S5. Estimation of the probability function $P(R_1 \mid x)$ from the trials in the whole experiment (in red) or from only the first 300 trials (in blue), for several subjects in the study (each line is a subject).

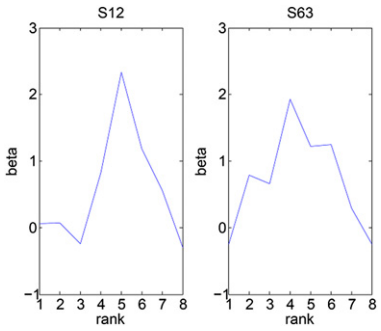

Fig. S6. Weighting profile across ranks, for two individual subjects in the study.