

The Irrationality of Categorical Perception

Stephen M. Fleming,^{1,2} Laurence T. Maloney,¹ and Nathaniel D. Daw¹

¹Department of Psychology and Center for Neural Science, New York University, New York, New York 10003, and ²Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, United Kingdom

Perception is often categorical: the perceptual system selects one interpretation of a stimulus even when evidence in favor of other interpretations is appreciable. Such categorization is potentially in conflict with normative decision theory, which mandates that the utility of various courses of action should depend on the probabilities of all possible states of the world, not just that of the one perceived. If these probabilities are lost as a result of categorization, choice will be suboptimal. Here we test for such irrationality in a task that requires human observers to combine perceptual evidence with the uncertain consequences of action. Observers made rapid pointing movements to targets on a touch screen, with rewards determined by perceptual and motor uncertainty. Across both visual and auditory decision tasks, observers consistently placed too much weight on perceptual uncertainty relative to action uncertainty. We show that this suboptimality can be explained as a consequence of categorical perception. Our findings indicate that normative decision making may be fundamentally constrained by the architecture of the perceptual system.

Introduction

Bayesian decision theory imparts a clear goal to the observer: actions should be chosen to maximize expected future utility given the current state of the world. This scheme mandates considering the utility a candidate action would achieve under each possible state of the world and averaging these utilities weighted by the probability of each state (Maloney, 2002; Bossaerts et al., 2008; Gershman and Daw, 2012). However, there is considerable evidence that perception is categorical. Kiani et al. (2008) reported evidence for an internal bound on evidence accumulation during the random-dot motion task (Newsome et al., 1989). Such a bound is associated with threshold levels of firing in the parietal and prefrontal cortices (Gold and Shadlen, 2007). The interpretation of this finding is that even when a behavioral response is not required, the stimulus is implicitly and prematurely categorized as reflecting one or another state of the world. More generally, categorical perception of ambiguous stimuli, as with rivalry, is subjectively apparent and empirically widespread (Harnard, 1987).

If the probabilities associated with possible states are lost through categorical perception, the organism cannot determine which action maximizes utility. Consider the decision tree laid out in Figure 1. At night or in poor visibility, aircraft pilots occasionally experience an illusion in which the plane is perceived as

being banked to one side when it is not. If the aircraft is categorically perceived as banked, the arm of the decision tree that considers the costs and benefits of the alternative state will be ignored. This can be potentially catastrophic if the unnecessary corrective action leads to a crash.

However, it is unknown whether categorical perception disrupts normative decision computations. There are several reasons for this lacuna in the literature. First, in tasks in which categorical perception has been reported, the utility function is most often implicit or fixed so as to preclude dissociating state inference from utility maximization: for example, observers may receive a reward for correctly reporting the world state and nothing otherwise. Second, whereas payoffs for different outcomes can be manipulated to affect decision criteria (Whiteley and Sahani, 2008; Feng et al., 2009; Fleming et al., 2010), this has mostly been examined in detection regimes in which the categorical aspects of perception are subtler (Kiani et al., 2008; Aly and Yonelinas, 2012). Finally, an important contributor to potential losses and gains from an action is the action itself (Trommershäuser et al., 2003a; Trommershäuser et al., 2003b; Körding and Wolpert, 2004; Trommershäuser et al., 2006; Landy et al., 2012), which in categorization and detection tasks is usually trivial (a button press or eye movement).

We hypothesized that the sequential organization of perception and action—both conceptually and in terms of the brain's gross functional organization—makes their combination a particularly likely setting in which categorical perception might adversely constrain decision making. To test this hypothesis, we manipulated state and action uncertainty in a random-dot motion task to determine whether categorical perception disrupts normative decision computation.

Materials and Methods

Subjects

Eighteen subjects in Experiment 1a (mean age 22.2, range 18–33, 12 females) and 7 subjects in Experiment 1b (mean age 22, range 19–29, 5

Received March 24, 2013; revised Aug. 22, 2013; accepted Sept. 24, 2013.

Author contributions: S.M.F., L.T.M., and N.D.D. designed research; S.M.F. performed research; S.M.F. and N.D.D. analyzed data; S.M.F., L.T.M., and N.D.D. wrote the paper.

This work was supported by the Wellcome Trust (Sir Henry Wellcome Fellowship WT096185 to S.M.F.), the National Eye Institute, National Institutes of Health (award EY019889 to L.T.M.), the McKnight Foundation (Scholar Award to N.D.D.), and the James S. McDonnell Foundation (Scholar Award to N.D.D.). We thank Adam Kepecs and Joshua Sanders for sharing code used to generate the auditory stimuli in Experiment 1b, Sophie Tam for assistance with data collection, and Michael Lee for advice on MCMC sampling.

Correspondence should be addressed to Stephen M. Fleming, Center for Neural Science, New York University, 6 Washington Place, New York, NY 10003. E-mail: fleming.sm@gmail.com.

DOI:10.1523/JNEUROSCI.1263-13.2013

Copyright © 2013 the authors 0270-6474/13/3319060-11\$15.00/0

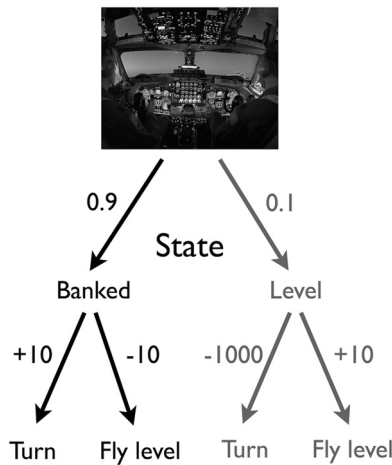


Figure 1. Schematic decision tree illustrating why categorical perception is at odds with normative decision theory. The example draws on the “graveyard spiral” illusion in which pilots under conditions of poor visibility may perceive their aircraft as banked when it is actually flying level. To decide which corrective action to take, an ideal observer computes the posterior probability of being in one or another state of the world (banked or flying level) and uses it to weight the expected utility of the actions available in each world state. Here, the plane is in fact flying level and trying to correct for an illusory bank will result in disaster. The ideal observer incorporates the small subjective probability of being in the alternate state and selects the best action (fly level) given the potentially catastrophic consequences of correcting the bank. A categorical perceiver, in contrast, settles on a single (illusory) state of the world, thus closing off the actions contingent on the nonperceived state (in gray). This will lead to disaster if, as in this example, there is a large cost associated with the alternate state.

females) gave written informed consent to participate. Twelve subjects (mean age 23.1, range 19–33, 6 females) gave written informed consent to participate in Experiment 2. Seven of these subjects had previously provided data in Experiment 1a. Subjects received US\$12 per hour plus a performance-related bonus as described below. The study was approved by the New York University Committee on Activities Involving Human Subjects.

Apparatus

Stimuli were displayed on a vertically mounted touchscreen LCD monitor (338 mm × 270 mm) in a dimly lit room. The monitor resolution was 1280 × 1024 pixels, with a refresh rate of 75 Hz. Subjects were positioned 45 cm away from the screen. The experiment was programmed in MATLAB (MathWorks) using Psychtoolbox (Brainard, 1997; Pelli, 1997).

Stimuli

Experiment 1. Stimuli were presented against a gray background. All trials began with a central white fixation point (0.2° diameter). All stimuli, including the fixation point, were displaced 5° below the horizontal meridian of the touchscreen. In pilot work, we found that this shift reduced anisotropy in the distribution of pointing end points by minimizing the vertical component of the movement.

Targets for the motor component of the task were circular gray disks presented at an eccentricity of 8° either side of fixation. In the calibration phase, only one target was presented; in the main experiment, two targets were presented. Targets used in the calibration phase were 0.5°, 1°, 1.5°, and 2° in diameter.

Experiment 1a. In Experiment 1a, stimuli consisted of random-dot kinematograms (RDKs). Each stimulus consisted of a field of random dots (2 × 2 pixels; 0.52 × 0.52 mm) contained in a 7° diameter aperture. Each set of dots lasted for 1 video frame and was replotted 3 frames later (Roitman and Shadlen, 2002). Each time the same set of dots was replotted, a subset determined by the percent coherence was offset from their original location in the direction of motion and the remaining dots were replotted randomly. Coherently moving dots moved at a speed of 8°/s and the number of dots in each frame were specified such that their density was 30 dots/deg²/s. Each RDK stimulus lasted for 800 ms (60

video frames). The percent coherences used in the calibration phase were 3%, 8%, 12%, 24%, 48%, or 100%. Motion direction (left or right) was randomized and independent of coherence.

Experiment 1b. In Experiment 1b, the sensory evidence was delivered in the auditory modality. Stimuli consisted of Poisson-distributed “click-trains” played binaurally through headphones. Each click was a 23 ms burst of white noise sampled at a rate of 44 kHz. The overall click rate was set to 200 Hz and each clicktrain lasted for 1 s. On each trial, the number of clicks played to each ear varied depending on the ratio of click rates for the correct to incorrect response. In the calibration phase, these ratios were drawn from the set [0.51 0.53 0.58 0.68 0.8 1.0]. The direction was randomized and independent of ratio. Subjects listened to the click stimuli over a pair of Bose OE2i headphones set to a comfortable volume.

Experiment 2. Stimuli were the same as in Experiment 1a.

Task and procedure

Calibration phase. Before performing the main task, each subject underwent a calibration session designed to quantify task-related state and action uncertainty.

To quantify state uncertainty, we asked subjects to perform a 2-alternative forced choice judgment as to whether the predominant motion direction or clicktrain was to the left or right. Stimulus strength and direction were randomized across trials. Subjects made their judgment using the left and right arrow keys on a computer keyboard after the offset of each stimulus. The response was unsped. After each response, auditory feedback was delivered as to whether the judgment was correct (high pitched tone) or incorrect (low pitched tone). The intertrial interval was 1 s. The calibration consisted of 3 blocks of 100 trials.

To quantify action uncertainty, we asked each subject to make a series of speeded pointing movements to a target of variable size that could appear either 8° to the left or right of the fixation point. The x and y coordinates of each target was adjusted on each trial by sampling from a uniform distribution of ±1° to discourage preplanning of movements. Target side and size were randomized across trials. Subjects could initiate each trial by pressing down the space bar with their dominant hand on a keyboard firmly fixed to the table. If the spacebar was still being held after 500 ms, a target was displayed for 750 ms. Subjects were required to reach out and hit the target with their index finger before 750 ms had elapsed. If a target was successfully acquired, it changed color from gray to red. If the subject released the spacebar before the target appeared, the trial was aborted and a message was displayed reminding subjects to continue holding the spacebar before target onset. If the touchscreen did not register a response within 750 ms, the message “Too slow!” was displayed for a time-out of 5 s. Calibration consisted of 3 blocks of 96 trials.

Experiment phase: Experiment 1. In the main experiment, subjects were instructed that they would receive a reward by successfully acquiring the target indicated by the direction of the stimulus. The ideal observer should take into account both the chances of being in each possible state (each possible motion direction) and the chances of successfully hitting the targets on the left or right (see below). For example, if the sensory evidence indicating left is borderline and the target on the left is much smaller than the target on the right, selecting the right target may lead to greater expected reward. If, however, the sensory evidence is compelling, the subject should resign himself to attempting to hit the smaller target.

The subject initiated each trial by depressing the spacebar, as in the motor calibration phase. After 500 ms, if the spacebar were still depressed, two targets differing in size appeared either side of fixation (Fig. 2C). After a further 300 ms, the RDK or clicktrain stimulus was presented. Subjects were instructed that they would not be able to release the spacebar until the end of the RDK or clicktrain. If the spacebar was released prematurely, an error message was displayed and the trial was aborted. At the end of the stimulus, the fixation point turned red and subjects were able to release the spacebar and point to one of the two targets. If they hit the correct target within 750 ms, it turned green and a high-pitched tone was heard. If the incorrect target was hit, it turned red and a low-pitched tone was heard. If anywhere else on the touchscreen was contacted or the time limit was exceeded, the targets remained gray and a low-pitched tone was heard. The intertrial interval was 1 s.

Subjects were instructed that they would have the potential to earn a bonus payment of up to \$12 by hitting the target indicated by the direction of motion. The instructions emphasized that a bonus could not be earned unless the target was hit and that to perform well on the task requires paying attention to both the stimulus and the sizes of the targets. The main experiment consisted of 3 blocks of 108 trials each with brief rest periods between blocks.

At the end of the experiment, six trials were randomly selected from the choice set. For each trial on which the subject was successful (i.e., hit the target in the direction of the stimulus), \$2 was added to their overall payment. Subjects received an average bonus of \$5.3 (range \$0–\$10) in Experiment 1a and \$6 (range \$4–\$8) in Experiment 1b.

Experiment phase: Experiment 2. Experiment 2 was identical to Experiment 1a, except that we elicited a probability judgment after target selection but before the pointing attempt. Calibration of state and action uncertainty was performed exactly as for Experiment 1a. In the main experiment, targets appeared for 300 ms, followed by the random-dot motion stimulus for 800 ms. Subjects were required to indicate, with a button press, which target they would prefer to hit (left or right). The reward schedule was the same as in Experiment 1. After target selection, a probability screen appeared, with probability indicated both by a circular pie chart and a vertical scale labeled from 0% to 100% in 10% increments. Using a trackball, observers could adjust the probability scale to their selected rating and confirm their choice by clicking one of the trackball buttons. There was no time limit on probability ratings.

On each trial, one of three probabilities was elicited, indicated by the question displayed on the probability screen. The occurrence of a particular question was random and could not be predicted during the judgment phase of the trial. “Probability dots?” asked subjects to rate the probability that their response was in the direction of motion regardless of whether they thought they could hit the chosen target. This question aimed to measure the subjective state probability. “Probability hit?” asked subjects to rate the probability that they would be able to hit their chosen target, regardless of whether they thought they had chosen the correct direction of motion. This question aimed to elicit the subjective action probability. “Probability overall?” asked subjects to estimate their overall probability of success. This question aimed to elicit overall subjective confidence.

After the rating, observers were asked to hit their chosen target. The procedure here was similar to the motor calibration phase—after pressing the spacebar for at least 500 ms, the subject was required to reach out and hit the chosen target with his or her index finger before the 750 ms had elapsed.

At the end of the experiment, six trials were randomly selected from the choice set. For each trial on which the subject was successful (i.e., hit the target in the direction of the stimulus), \$1 was added to the overall payment.

We incentivized probability ratings using the “Lottery Rule” (Karni, 2009). This rule is similar to the Becker-DeGroot-Marshak mechanism in behavioral economics (Becker et al., 1964) and provides incentives for the subject to truthfully reveal a subjective probability p . For each of the 6 selected trials, subjects drew a ball from a bag of bingo balls labeled 1–100. If $p > l_1$, the computer checked to see if the subject had correctly

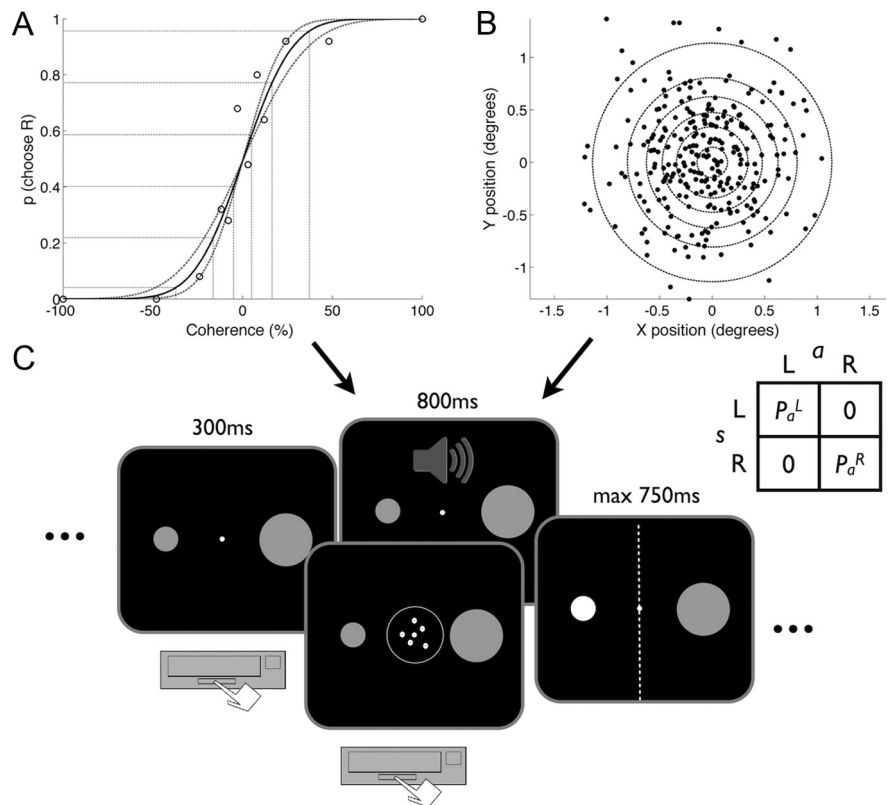


Figure 2. Calibration data and task setup. **A**, Example subject data from the sensory calibration stage. Open circles reflect the proportion of times rightward motion was indicated as a function of coherence. Lines indicate the fit of a cumulative Gaussian to these data, with 95% confidence intervals on the slope parameter. Dotted vertical lines indicate coherence values selected for use in the main experiment for this subject. **B**, Example subject data from the motor calibration stage. Dots reflect the x, y position of movement end points collected from several trials. Circles indicate the target sizes selected for use in the main experiment for this subject. **C**, Schematic outline of the task. Observers began the trial by depressing the spacebar, after which two targets appeared on the left or right of the screen. A random-dot kinematogram (Experiment 1a) or auditory stimulus (Experiment 1b) was then presented for 800 ms. Responses were made by a rapid pointing movement to one of the two targets. The inset tabulates the utility function for the task. Rewards are available for hitting the target in the direction of the world state (left or right), so the expected utility of each action is determined solely by the chances of hitting each target (P_a).

answered the relevant component of the task (selected direction for “Probability dots?,” target acquisition for “Probability hit?,” and overall success for “Probability overall?”). If the judgment was correct, an additional \$1 was won; if incorrect, nothing was won. If $p < l_1$, subjects drew a new bingo ball, l_2 . If $l_2 < l_1$, \$1 was won; if $l_1 < l_2$, nothing was won. The rule can be intuitively understood as follows. The higher the initial rating of p , the more likely the correctness of the decision will determine earnings. The lower the rating, the more likely earnings will be determined by chance (a second lottery). A particular rating value (e.g., 75%) thus reveals how subjects trade off a belief in an element of a decision being correct against a randomly determined reward. Before the experiment, we explained various possible outcomes to subjects, along with their intuitive interpretation, until they understood how different rating strategies affected their potential earnings.

Stimulus and target selection. Stimulus and target sets used in the main experiment were individually determined for each subject from online fits to the calibration data. Stimulus coherence levels θ were identified that led to “rightward” state probability values of [0.05 0.23 0.44 0.56 0.77 0.95] using probit regression (Fig. 2A). Rightward target sizes were chosen that corresponded to matching action probability values (P_a^R) of [0.05 0.23 0.44 0.56 0.77 0.95] by integrating a bivariate probability density as outlined below (note outliers were not excluded at this stage; Fig. 2B). For each condition, the target probability on the left (P_a^L) was set to $1 - P_a^R$. Stimuli and targets were randomized across trials and crossed to create a fully factorial design

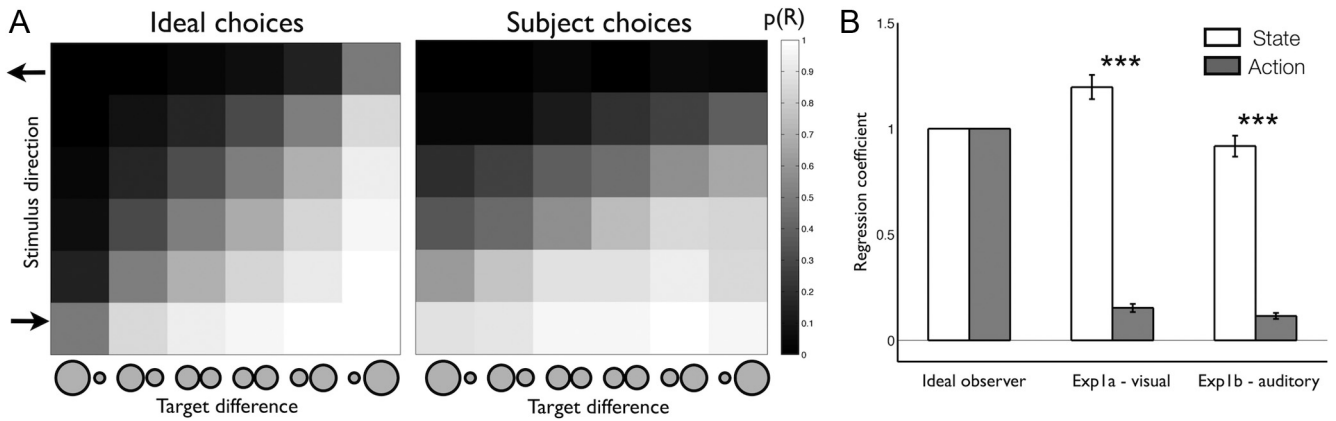


Figure 3. Experiment 1 results. **A**, Ideal observer model of choice probabilities and average subject data from Experiment 1a. The heat map reflects the predicted proportion of times the rightward target should be selected as a function of both stimulus coherence (y -axis) and target asymmetry (x -axis). **B**, State and action regression coefficients from a probit analysis of Experiments 1a and 1b plotted alongside the ideal observer predictions. The size of the coefficient represents the influence state and action probabilities had on subjects' choices. $***p < 0.0001$; error bars reflect SEM.

(Fig. 3A). The model fitting procedures outlined in subsequent sections were then performed offline to refine our estimates of state and action probabilities for subsequent analysis.

Computational models

Ideal observer model. The agent receives a fixed reward r (in this case, a fixed possibility of a reward) if the target associated with the particular state is successfully acquired. The utility function is tabulated in the inset to Figure 2C. Because r is constant throughout the experiment, the utility of each action is determined only by the posterior belief $P(s|X)$ (the “state” probability) and the probabilities of acquiring each target P_a^L and P_a^R (“action” probabilities):

$$\begin{aligned} U(a = R) &= P(s = R|X) \cdot P_a^R \\ U(a = L) &= P(s = L|X) \cdot P_a^L \end{aligned} \quad (1)$$

In other words, the optimal agent considers the chances of hitting each target and weighs these by the posterior probability of being in each state. The decision rule is to respond “right” if:

$$P(s = R|X) \cdot P_a^R > P(s = L|X) \cdot P_a^L \quad (2)$$

We can equivalently write this rule in terms of state and action log-odds ratios:

$$\log \frac{P(s = R|X)}{P(s = L|X)} + \log \frac{P_a^R}{P_a^L} > 0 \quad (3)$$

The state probability for motion direction has the following generative model. The true motion direction s gives rise to a noisy sample of this direction from one of two Gaussian distributions with means that depend linearly on motion coherence θ via free parameter k (Palmer et al., 2005). The variance of these distributions reflects the combined contributions of external (stimulus) and internal (observer perceptual) noise; we take it to be 1 without loss of generality because k is free to rescale. We can then write down the likelihood function for the observer's noisy sensory measurement X on each trial as follows:

$$P(X|s) \sim N(dk\theta, 1) \quad (4)$$

where d is an indicator variable that is +1 if motion is rightward and -1 if motion is leftward. Because leftward and rightward motion directions are equally likely (the prior is flat), the log posterior odds for belief in rightward motion (B ; the first component of Equation 3) is equal to the log-likelihood ratio:

$$B = \log \frac{P(s = R|X)}{P(s = L|X)} = \log \frac{P(X|s = R)}{P(X|s = L)} \quad (5)$$

$$\begin{aligned} &= \log \frac{e^{(k\theta+X)^2/2}}{e^{(k\theta-X)^2/2}} \\ &= 2k\theta X \end{aligned}$$

We use M to denote the action log-odds ratio (the second component of Eq. 3):

$$M = \log \frac{P_a^R}{P_a^L} \quad (6)$$

Therefore, the compound decision variable, D , is as follows:

$$D = B + M + \beta \quad (7)$$

The observer should respond “right” when $D > 0$ and “left” when $D < 0$. Free parameter β (which should be zero in the ideal case) accommodates any overall bias toward left or right regardless of experimental condition.

“Categorical” model. The model outlined above is optimal in that it maintains a posterior probability of being in one or another state to select the action with the highest expected utility. If, in contrast, the observer has an internal bound on evidence accumulation (Kiani et al., 2008), they may prematurely settle on one or other state of the world before combination with a utility function. We approximate bounded accumulation with a threshold ($\pm A$) on the log-posterior odds (Fig. 6A):

$$\begin{aligned} \text{If } |B| > A & \quad D = B + \beta \\ \text{else} & \quad D = B + M + \beta \end{aligned} \quad (8)$$

If A is very large, the categorical model becomes the ideal observer. As A becomes smaller, a greater proportion of trials are categorically perceived and action probabilities are ignored (Fig. 5B).

“Weighted” model. An alternative to the categorical model is that observers draw on both state and action information, but weight state information to a greater degree when making their decision. In this model, parameter w controls how state and action probabilities are combined:

$$D = wB + (2 - w)M + \beta \quad (9)$$

Data analysis

Estimating state and action uncertainty from calibration data. In what follows, we use i and j to index the target pair t and coherence θ for each

condition in the experimental design, respectively. In the state calibration phase, the decision variable (Equation 7) depends only on the state log-odds B and bias parameter β :

$$D_{calib} = B + \beta \quad (10)$$

The probability of answering “right” for coherence level θ_j is then:

$$\begin{aligned} P(a = R|\theta_j) &= P(D_{calib} > 0) \\ &= P(B > -\beta) \\ &= P(2k\theta_j X > -\beta) \text{ (From Eq. 5)} \\ &= P\left(X > -\frac{\beta}{2k\theta_j}\right) \\ &= \Phi\left(dk\theta_j + \frac{\beta}{2k\theta_j}\right) \text{ (From Eq. 4)} \end{aligned} \quad (11)$$

where $\Phi()$ is the standard cumulative Gaussian function.

For each subject, k and β were fit to the calibration data using Markov Chain Monte Carlo methods (see “Model fitting” below). The mean and variance of each subject’s k and β estimates were stored and used to specify prior distributions when fitting experiment phase data.

It is possible to underestimate k due to random lapses made by the subject at high stimulus strength (Wichmann and Hill, 2001). To address this possibility, we additionally refit k assuming that all errors made when the stimulus was unambiguous (i.e., when all moving dots or clicks were associated with one or other stimulus direction) could be attributed to lapses (Prins, 2012). The probability of a lapse, λ , was set to the sum of the lapse rate to rightward and leftward stimuli, leaving the probability that the response was instead based on the cumulative normal function as $1 - \lambda$ (where θ_1 is unambiguous motion to the left and θ_n is unambiguous motion to the right):

$$\begin{aligned} P(a = R|\theta_1) &= \frac{\lambda}{2}; P(a = R|\theta_n) = \left(1 - \frac{\lambda}{2}\right); \\ P(a = R|\theta_{1 < j < n}) &= (1 - \lambda) \cdot \Phi\left(dk\theta_j + \frac{\beta}{2k\theta_j}\right) + \frac{\lambda}{2} \end{aligned} \quad (12)$$

In the action calibration phase, we assume that the motor system imposes a probability density on the space of possible movement trajectories that could occur once a target is selected (Trommershäuser et al., 2003b). The center of each target, t_j , is at a given location (x, y) on the touchscreen. For each subject, the probability of landing at a particular location (x', y') having aimed at the center of target t_j is given by the multivariate normal density $P(x', y'|x, y) = N(x, y; \mu_a, \Sigma_a)$. The probability of hitting the target is then the integral of this density over the target area:

$$P_a(t_j) = \int_{t_j} P(x', y'|x, y) dx' dy' \quad (13)$$

To determine P_a for each subject, we computed the covariance Σ_a of the bivariate distribution of end points during calibration. Endpoints >3 SDs from the mean in either x or y were excluded as outliers. This covariance matrix was used to specify a bivariate normal probability density centered on the target location. All points on a 100×100 grid falling within each target circle were numerically integrated to approximate the solution to Equation 13. $\hat{M}(t_j)$ is then the estimated log-odds in support of a rightward action for each target pair:

$$\hat{M}(t_i) = \log \frac{P_a^R(t_i)}{P_a^L(t_i)} \quad (14)$$

Probit regression analysis. We visualized data from Experiment 1 by plotting the probability of choosing the rightward target (regardless of whether it was acquired) as a function of both target asymmetry and stimulus identity (Fig. 3A). Rightward or leftward choices were defined by whether the movement end point fell to the left or right of the vertical meridian of the touchscreen.

The ideal observer model (Equation 7), marginalizing over the inter-nal sample X (which is unknown to the experimenter), predicts that the probability of responding “right” is a cumulative Gaussian thresholded by the action log-odds M :

$$\begin{aligned} P(a = R|\theta_j, t_i) &= P(B + M + \beta > 0) \\ &= P(2k\theta_j X > -M - \beta) \text{ (From Eq. 5)} \\ &= P\left(X > -\frac{M + \beta}{2k\theta_j}\right) \\ &= \Phi\left(dk\theta_j + \frac{M + \beta}{2k\theta_j}\right) \text{ (From Eq. 4)} \end{aligned} \quad (15)$$

Given this identity between the ideal observer and Equation 15, the ideal observer makes predictions about effect sizes that can be tested using probit regression. Specifically, the estimated action log-odds \hat{M} and coherence θ_j can then be entered as predictors of subjects’ choices:

$$a \sim \Phi\left(\beta_s \cdot dk\theta_j + \beta_a \cdot \frac{\hat{M}}{2k\theta_j} + \beta_0\right) \quad (16)$$

where $a = 1$ for a rightward response and 0 otherwise, and \hat{M} and k are derived from calibration data as above. We estimated state (β_s) and action (β_a) regression coefficients (summary statistics) at the individual subject level and compared them at the group level using paired t tests (Holmes and Friston, 1998). If subjects are following the ideal observer model, $\beta_s = \beta_a = 1$. In contrast, if categorical perception is at work, we expect observers to sometimes disregard the possible rewards available from the unperceived state, leading in the aggregate to a weaker overall influence of action probabilities on behavior and $\beta_a < \beta_s$.

Model fitting. We used Markov chain Monte Carlo methods to sample from each candidate model using Gibbs sampling as implemented in JAGS (<http://mcmc-jags.sourceforge.net/>). Uninformative (high variance) prior distributions on k and β were used when fitting calibration data (after JAGS convention, variances are written as precisions, or the reciprocal of the variance):

$$\mu_k \sim N(0, 0.0001)$$

$$\mu_\beta \sim N(0, 0.0001)$$

The mean and variance of the posterior densities of k and β were then entered as priors when fitting models to the experimental phase data. Priors on w and A were specified as uniform distributions:

$$k \sim N\left(\mu_k, \frac{1}{\sigma_k^2}\right)$$

$$\beta \sim N\left(\mu_\beta, \frac{1}{\sigma_\beta^2}\right)$$

$$w \sim \text{unif}(0, 2)$$

$$A \sim \text{unif}(0, 5)$$

A is bounded above at 5, which is equivalent to a posterior belief in rightward or leftward motion of 0.99 or greater. We note that the categorical model coincides exactly with the ideal model in the limit as $A \rightarrow \infty$ when the flanking criteria in Figure 6A are pushed out past the edges. We simplify the model by bounding A a priori within the perceptually relevant range; the behavior at $A = 5$ is effectively identical to the ideal observer and further increases in A do not appreciably affect its behavior.

For each subject/model, JAGS sampling was run with 2000 adaptation steps, 20,000 burn-in samples and 250,000 effective samples. Traces for parameters k , β , w , and A were monitored and estimates of the posterior density for each parameter were calculated. Three chains for each parameter were run, each with different starting points, and visually checked for convergence. We additionally report Gelman and Rubin’s (1992) poten-

tial scale reduction statistic \hat{R} for all parameters. Large values of \hat{R} indicate convergence problems and values ~ 1 suggest convergence. Average \hat{R} (across parameters and subjects) was 1.001 and all values were < 1.1 , indicating good convergence.

To determine the ability of each model to account for subjects' choices, we extracted the maximum a posteriori parameter estimates from each subject's model fits and simulated 50,000 trials with these parameter settings. The predictions of each model are plotted alongside the data in Figure 6B. As a quantitative measure of model fit, we computed Deviance Information Criterion (DIC) scores after model convergence (Spiegelhalter et al., 2002). DIC provides a measure for how well each model fits the data while penalizing for model complexity (effective number of parameters). DIC is approximately equivalent to Akaike's Information Criterion assuming negligible prior information (Spiegelhalter et al., 2002). A lower DIC score indicates better model fit, with differences > 6 considered strong evidence (Burnham and Anderson, 1998; Spiegelhalter et al., 2002).

Ideal observer for confidence. Subjective confidence is often equated with the subjective probability of success (Kepecs and Mainen, 2012). We modeled an ideal observer for confidence as the probability of success from Equation 2:

$$P_c = \operatorname{argmax}[P(s = R|x) \cdot P_a^R P(s = L|x) \cdot P_a^L] \quad (17)$$

To examine how (objective) P_s and P_a contribute to subjective confidence ($P_{c,sub}$), we constructed the following regression model:

$$\log P_{c,sub} = \beta_a \cdot \log P_a^{chosen}(t_i) + \beta_s \cdot \log P_s^{chosen}(\theta_j) + \beta_0 \quad (18)$$

where $P_a^{chosen}(t_i)$ indicates the value of P_a associated with the chosen target and $P_s^{chosen}(\theta_j) = 1 - \Phi(dk\theta_j)$ if the leftward target is chosen and $\Phi(dk\theta_j)$ if the rightward target is chosen.

Results

We developed a task in which observers were first given a stochastic perceptual cue (a visual or auditory stimulus) that indicated which of two circular targets carried a reward (Fig. 2C). The stimulus duration was fixed and observers could only respond after observing the entire evidence stream. The duration chosen occupied a regime known to be affected by internal bounds on evidence accumulation in primates (Kiani et al., 2008). The reliability of the perceptual cue and the chances of hitting each target were varied independently. We tailored stimuli and targets for each subject based on performance in a calibration phase.

The ideal Bayesian observer chooses whichever target would most likely lead to reward, giving equal weight to state and action probabilities (Fig. 3A). In contrast, if categorical perception is at work, we expect observers to sometimes disregard the possible rewards available from the unperceived state, leading in the aggregate to a weaker overall influence of action probabilities on behavior. We test for categorical perception effects by comparing the fit of ideal observer models with and without an internal threshold on sensory evidence.

Choice probabilities

Examples of individual subject data from the state and action calibration stages are shown in Figure 2A, B. These calibration phases were used to select stimuli and targets for use in the main experiment (Fig. 2C; see Materials and Methods). We visualized data from Experiment 1a by plotting the probability of choosing the rightward target (regardless of whether it was acquired) as a function of both target asymmetry and stimulus coherence (Fig. 3A, right). Due to inherent state and action uncertainty, an ideal observer should choose the target with the greatest probability of success. The ideal observer model thus predicts that state and

Table 1. Mean regression coefficients for state and action probabilities

Experiment	Method to fit k	State β (SEM)	Action β (SEM)
Visual	Cumulative normal	1.20 (0.05)	0.15 (0.02)
	Cumulative normal UB*	0.95 (0.05)	0.19 (0.02)
	Cumulative normal + lapse	1.15 (0.07)	0.16 (0.02)
Auditory	Cumulative normal	0.92 (0.05)	0.12 (0.02)
	Cumulative normal UB*	0.73 (0.04)	0.12 (0.01)
	Cumulative normal + lapse	0.85 (0.06)	0.12 (0.02)

*UB indicates upper bound on 95% credible interval.

action probabilities will be weighed equally in the decision process, resulting in the diagonal tradeoff shown in Figure 3A, left. We observed that, although subjects showed sensitivity to both sources of uncertainty, there was a systematic skew in the average heat map due to a greater influence of state probabilities on decision making. This influence is qualitatively consistent with a categorical observer who ignores action probabilities on a subset of trials.

The underweighting of action probabilities can be quantified in a probit regression analysis of the influence of the two sources of uncertainty on choices. If subjects are following the ideal observer (Equation 6), $\beta_s = \beta_a = 1$. Both the state and action regression coefficients were significantly nonzero (one-sample t test, state $t_{(17)} = 20.9, p < 10^{-12}$; action $t_{(17)} = 8.0, p < 10^{-6}$). The state coefficient was individually significant in all 18 subjects; the action coefficient was individually significant in 16 of 18 subjects. The action regression coefficient was systematically lower than the state regression coefficient in all 18 subjects (Fig. 3B; paired-samples t test, $t_{(17)} = 16.4, p < 10^{-11}$). We also established that the interaction between state and action probabilities was not significant (one-sample t test, $t_{(17)} = 0.64, p = 0.53$). This result indicates that subjects place greater weight on state compared with action probabilities during decision making.

Why do subjects underweight action probabilities?

We sought to rule out several explanations for this effect. First, we considered that an underestimation of state probabilities during calibration could give rise to its apparent overweighting during the decision phase. We recalculated the regression parameters after deriving state probability values from either the upper 95% confidence bound on the psychometric function slope or a psychometric function incorporating a lapse rate (Wichmann and Hill, 2001). Neither method altered our conclusions (Table 1).

Second, we considered that subjects might underweight action probabilities due to inattention to the target information. Stimulus and target information were presented through the same modality (vision) in Experiment 1a, so it may have been difficult for subjects to attend to peripheral targets while processing centrally presented sensory information. In Experiment 1b, we therefore adapted our task to present sensory information in the auditory, rather than visual, modality (using a binaural click stimulus; see Materials and Methods) while maintaining visual presentation of action targets. The same pattern of results was obtained (Fig. 3B), with state and action regression coefficients both significantly above zero (one-sample t test, state $t_{(6)} = 18.5, p < 10^{-5}$; action $t_{(6)} = 7.9, p < 0.001$) and the state coefficient again greater than the action coefficient in all 7 subjects ($t_{(6)} = 14.7, p < 10^{-5}$).

Finally, we considered that subjects may distort subjective probabilities such as failing to appreciate the success rate associated with pointing at targets of different sizes. We performed an additional experiment (Experiment 2; see Materials and Methods

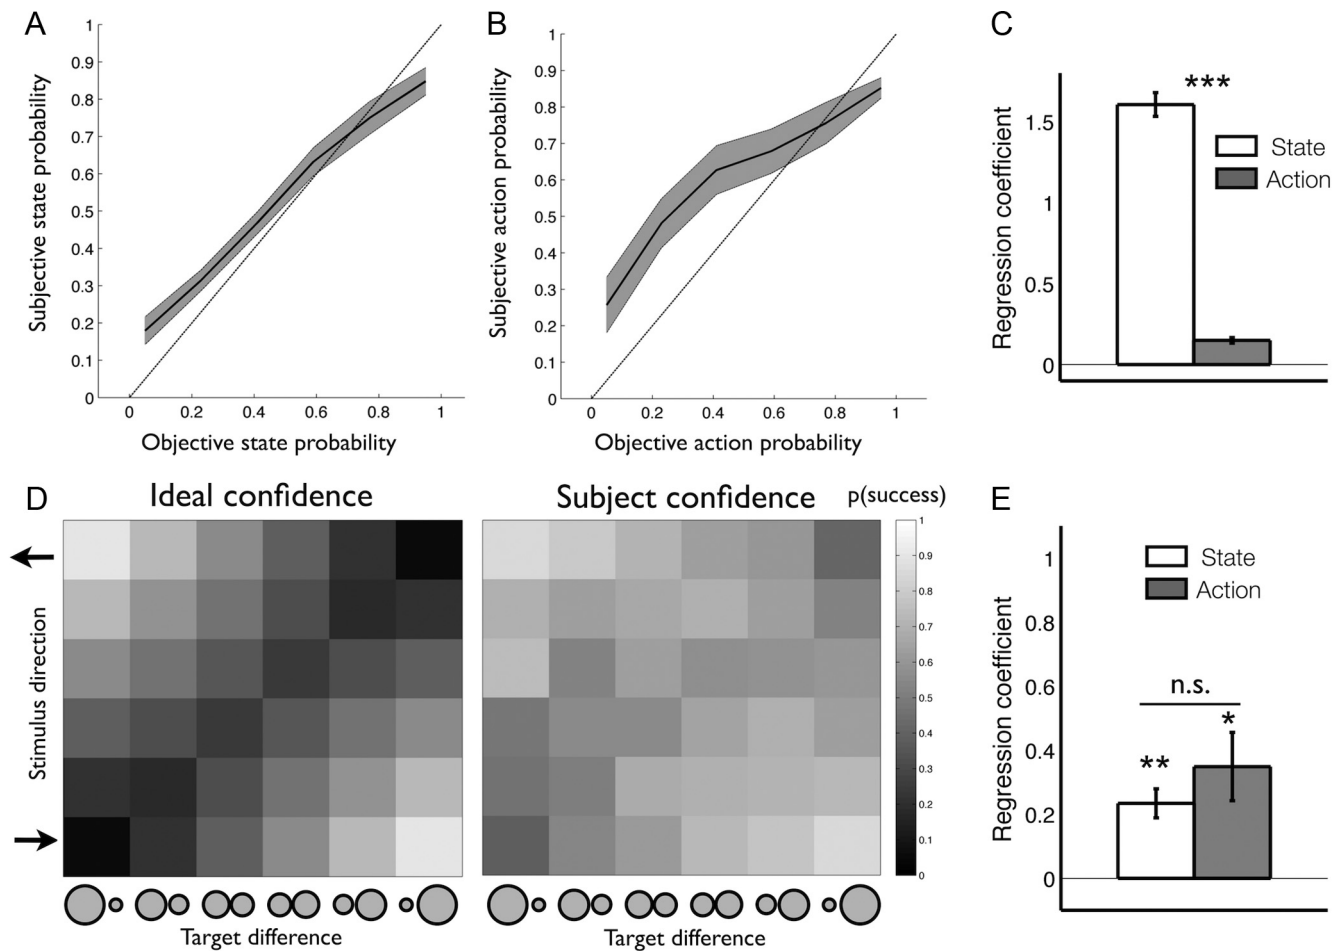


Figure 4. Experiment 2 results. **A**, Average subjective state probability in Experiment 2 plotted against objective state probabilities derived from the calibration stage. **B**, Average subjective probability of hitting chosen targets in Experiment 2 plotted against objective action probabilities derived from the calibration stage. **C**, State and action regression coefficients from a probit analysis of Experiment 1 (pooling over visual and auditory experiments), but now assuming the probability weighting functions derived from Experiment 2 (**A**, **B**). **D**, Ideal observer model of confidence (probability of success) and average subject data from Experiment 2. Colors reflect the predicted confidence rating as a function of both stimulus coherence (*y*-axis) and target asymmetry (*x*-axis). **E**, Regression coefficients for the influence of state and action probabilities on confidence judgments (Experiment 2). n.s., not significant; * $p < 0.01$; ** $p < 0.001$; *** $p < 0.0001$; error bars reflect SEM.

for details) in which we elicited subjective beliefs about success in various aspects of the task. Calibration curves for state and action probabilities are shown in Figure 4*A, B*. We then used these functions as proxies for state and action probabilities in a reanalysis of Experiments 1*a* and 1*b*. Despite the overconfidence shown for small targets in Experiment 2 (Fig. 4*B*), subjective probability distortion was not sufficient to explain the underweighting of action probabilities seen in subjects' behavior (Fig. 4*C*). In addition, in the elicited beliefs, both state and action probabilities contributed significantly and with equal weight to overall confidence in success (Fig. 3*D, E*; one-sample *t* test against zero, state $t_{(11)} = 5.41$, $p < 0.001$; action $t_{(11)} = 3.42$, $p < 0.01$; paired-samples *t* test $t_{(11)} = 1.23$, $p = 0.25$). In other words, subjects are well aware of the two limiting factors on their performance—state and action uncertainty—and, when reporting beliefs about them, they do not place greater weight on one or the other. These results indicate that the underweighting of action probability observed in Experiment 1 is not due to a lack of understanding of action-outcome uncertainty or of the structure or requirements of the task.

Computational model comparison

Instead, our results are consistent with a systematic constraint on the perceptual decision process that leads to an underweighting

of action probabilities. Such a constraint could result from a categorical commitment to one or other state of the world. In a signal detection model, this commitment can be represented as an additional bound that controls the transition from probabilistic to discrete state judgments (Fig. 5*A*). On this model, judgments become categorical on a subset of trials when the subjective evidence exceeds a threshold; the chance of this happening increases with stimulus strength. We next compared the fit of categorical and noncategorical observer models to subjects' choice probabilities in Experiment 1*a* (Fig. 5*B*; see Materials and Methods).

Examining the model fits, the categorical model was strongly preferred to the ideal observer model (Table 2; median DIC scores: ideal = 187.8, categorical = 51.8). A weighted model (equivalent to the probit regression) also provided a better fit than the ideal observer, as expected, but was not preferred to the categorical model (median DIC score = 59.1; difference in group-level DIC = 66.8). The median *A* parameter for the categorical model was 0.25 (Table 3; range 0–0.78), which corresponds to responding categorically when the state probability exceeds 0.56 (range 0.5–0.69).

The ability of each model to account for the pattern of subjects' choices is displayed in Figure 6. These plots show the pro-

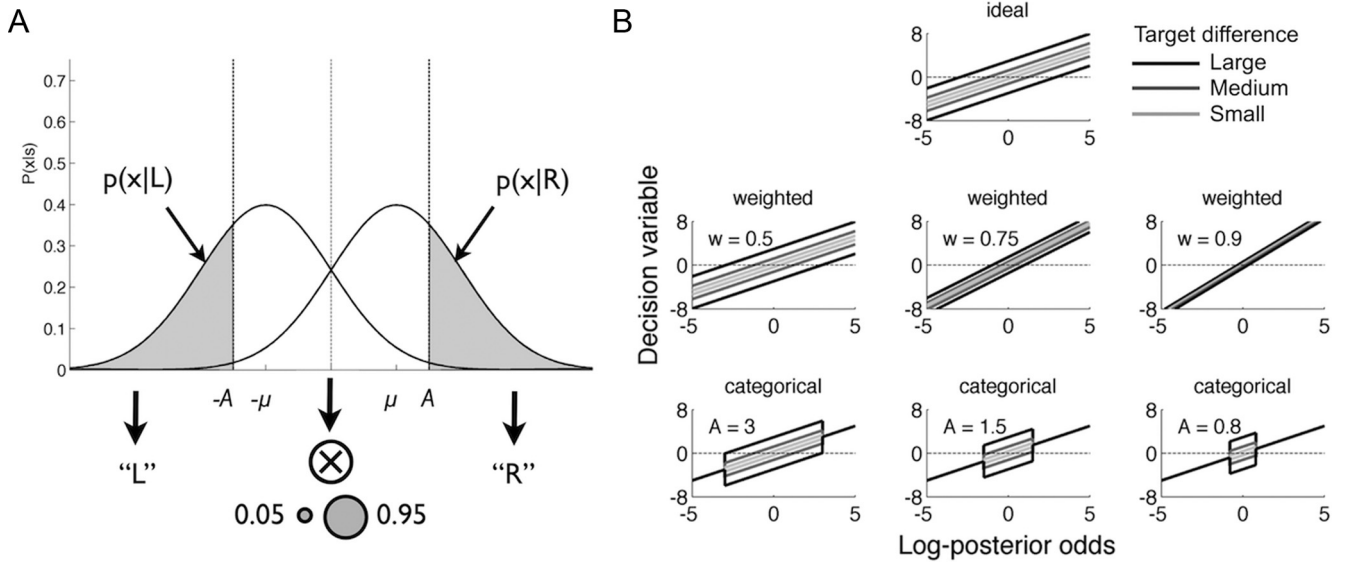


Figure 5. Computational models. **A**, Schematic of the “categorical observer” model described in Materials and Methods. Stimuli of varying strength provide evidence sampled from one of two Gaussian distributions depending on the motion direction. In the categorical model, samples from the posterior falling to the right of A or the left of $-A$ lead to a categorical belief in rightward or leftward states without further consideration of action-outcome uncertainty. See Materials and Methods for further details. **B**, Illustration of the decision variable D for each model variant (on the y -axis) as a function of the posterior log-odds B (on the x -axis). The shading of the lines indicate the difference in target size (small, medium, and large).

Table 2. Comparison of model fits to data

Model	Parameters	Median DIC	Summed DIC
Ideal	2	187.8	3321.8
Weighted	3	59.1	1007.2
Categorical	3	51.8	940.4

Table 3. Median maximum a posteriori parameter estimates \pm SD for each computational model

Model	k	β	w	A
Ideal	6.49 ± 3.16	0.06 ± 0.19	—	—
Weighted	5.07 ± 2.47	0.11 ± 0.16	1.76 ± 0.13	—
Categorical	5.14 ± 2.50	0.17 ± 0.13	—	0.25 ± 0.21

portion of times the model/subjects chose the bigger target in the pair as a function of motion coherence, target size, and whether the motion direction was congruent or incongruent with the bigger target. The ideal observer is unable to accommodate subjects’ avoidance of the bigger target (i.e., the underweighting of action probabilities). Both the weighted and categorical models accommodate this variable weighting, but only the categorical model captures the overlap of choice probabilities for medium and large target differences, particularly when coherence is low. In particular, the two models (weighting and categorical) differ in how they manage the tradeoff between the posterior belief B and the target probabilities M . For the categorical model, the range within which M exerts its effect is compressed, but because in this range B is also small, the model often goes with the bigger target regardless of the value of M (medium or large).

Discussion

Perception is often categorical (Harnard, 1987) and perceptual categorization can potentially affect decision making (Kiani et al., 2008; Aly and Yonelinas, 2012). Discrete categorization at the perceptual stage is at odds with the requirement that optimal decisions require evaluating the consequences of actions consid-

ering all possible states of the world—not just the state perceived. Neglecting alternate states as a result of categorization can lead to suboptimal selection of actions (Fig. 1).

We tested for this irrationality in a task that required the combination of sensory evidence with motor uncertainty. In a sensory discrimination task, observers responded with rapid pointing movements to targets on a touch screen, with potential rewards conditional both on the movement’s success and on variable perceptual information. We found that observers erred by placing more weight on perceptual state probabilities compared with action probabilities, even though an ideal observer would balance these factors symmetrically.

We showed that this pattern of results could be due to categorical perception distorting the computation of expected value, closing off actions that would be more beneficial. Harnessing computational model comparison, we found that a model in which categorical perception occurred on a subset of trials provided a good fit to observers’ choices. This finding was supported by a second experiment in which the sensory information was presented in a different modality (auditory) to the targets of action, indicating that visual inattention to action targets could not explain our results. Finally, we confirmed that our pattern of results was not due to a failure to appreciate action-outcome uncertainty. Subjects were sensitive, in explicit belief reports, to the probabilities of acquiring different targets, and gave equal weight to state and action probabilities when placing bets on their performance.

An internal bound on evidence accumulation provides a parsimonious explanation of these findings (Kiani et al., 2008). In a generalization of signal detection theory, observers accumulate information over time until a threshold is reached, at which point one or other state of the world is settled upon. A flattening of performance as stimulus duration is increased and a greater influence of early evidence on choice are both signatures of such an internal termination of accumulation (Vickers et al., 1971; Kiani et al., 2008). Conversely, other experiments reveal that if this threshold is not crossed, a probability of being in one or another




Figure 6. Comparison of ideal, weighted, and categorical model predictions (solid lines) against subjects' behavior (open circles). These plots show the proportion of times the model (subject) chose the bigger of the two targets as a function of difference in target size (small, medium, and large), coherence (low, medium, and high), and congruence (motion direction agrees or disagrees with bigger target location).

state is available to the observer to guide their decision making (Kiani and Shadlen, 2009). At a neural level, this probability is associated with graded changes in firing of posterior parietal cortex neurons (Yang and Shadlen, 2007; Kiani and Shadlen, 2009). In contrast, a bound crossing may be associated with an attractor-like process in which the system “commits” to one or another state of the world (Wong and Wang, 2006; Gold and Shadlen, 2007).

Consistent with these accounts, a categorical observer model faithfully combines state probabilities with action consequences when an internal threshold is not crossed, but disregards the utilities of alternative actions on a subset of trials that lead to categorical perception. We note that alternative explanations are possible. In particular, the weighted model we tested here also accommodates the behavioral finding of lowered sensitivity to action probabilities. However, we suggest that the categorical model provides a more satisfactory explanation of subjects' choice patterns. First, the categorical model accommodates qualitative features of the data that are not predicted by the weighted model, such as the near-equal effect of medium and large target differences when sensory evidence is low. Second, in explicit belief reports, subjects did not weight state probabilities differently from action probabilities. Finally, Bayesian model comparison indicates a small but significant advantage for the categorical model.

Thresholds on evidence are cornerstones of optimal decision rules and the role of evidence thresholding in optimal and suboptimal choice is somewhat subtle. For example, optimizing expected reward rate in the reaction time version of the random-dot motion task (in which observers can terminate the stimulus when they are ready to answer) prescribes responding when evidence reaches a threshold (Wald, 1945; Gold and Shadlen, 2007). However, as we have described, when a fixed amount of noisy evidence is available, prematurely ceasing to consider it or in any case treating the state probability distribution resulting from accumulated evidence as categorical can lead to inefficient choices. Finally, it is true that, given any particular loss function, the optimal decision rule that would arise from computing utility in expectation over the full state distribution can equivalently be reexpressed in terms of a threshold on the perceptual evidence: a criterion shift in signal detection terms (Green and Swets, 1966). However, in such a one-stage perception/decision system, determining the appropriate threshold depends on the loss

function. Therefore, in an arguably more realistic computational architecture in which perceptions are computed and utility-bearing actions are evaluated serially—as we suggest, in evaluating noisy movements conditional on noisy percepts—the threshold is instantiated at the movement stage and premature categorization at the perceptual stage can lead to errors.

Previous studies have demonstrated that observers can combine state uncertainty with loss information to guide decision making (Bohil and Maddox, 2001; Whiteley and Sahani, 2008; Feng et al., 2009; Fleming et al., 2010; Rorie et al., 2010; Summerfield and Koechlin, 2010). We analogously find that observers are sensitive to expected utilities (here, the probability of successful actions) in perceptual decision making. However, we additionally find that action consequences are underweighted. In signal detection terms, this result is consistent with empirically measured criterion shifts being closer to neutral than would be expected from an ideal observer model (Green and Swets, 1966; Healy and Kubovy, 1981; Maloney and Thomas, 1991; Bohil and Maddox, 2001). A commonly invoked *post hoc* explanation for this conservatism is that observers place weight on being accurate, rather than only on maximizing reward, leading to a competition between reward and accuracy maximization (Maddox, 2002; Maddox and Bohil, 2004). A categorical observer model offers a different explanation: it also attenuates overall bias effects, but does so not because observers place a premium on accuracy. Instead, the criterion undercompensates for the action probabilities on average because a subset of trials leads to categorical perception. In other words, an apparent insensitivity to payoff information can be explained by the sequential nature of perception and action together with categorical perception at the perceptual stage.

A categorical observer effectively exhibits nonlinearity in the computation of state probabilities (small state probabilities are underweighted and large probabilities are overweighted). Nonlinear probability weighting is a ubiquitous feature of decision under risk (for review, see Zhang and Maloney, 2012). However, the weighting function implied by the categorical perception model is opposite to that commonly observed in “decision from description” tasks, in which small probabilities are overweighted and large probabilities are underweighted (Kahneman and Tversky, 1979; Gonzalez and Wu, 1999). Categorical perception instead produces a pattern similar to that observed when probabilities are learned from experience (Hertwig et al., 2004).

However, whereas explanations of such underweighting in that domain relate to the statistics of decision makers' objective trial experience (e.g., rare events tend not to be included in an average subject's sample), our results suggest that similar nonlinear weighting may arise internally due to how experience is interpreted.

If categorical perception throws normative decision computations off course, why is it so common? One possibility is that categorical perception is a byproduct of a limited-resource computation of a posterior belief distribution (Gershman et al., 2012). Although approximate perceptual inference may be costly due to producing relatively poorer decisions, these inefficiencies may be balanced, on average, by significant savings in computational costs due to simplifying the problem. Therefore, such approximations are justified overall even though they can lead to irrationalities in individual decision scenarios. A second point is that categorical judgments are not, in themselves, problematic. To the contrary, decisions by definition require the production of discrete responses, such as whether it is safe to cross the road or whether to subscribe to pension plan A or B. However, if evolution has iteratively replicated or extended core mechanisms or circuits for simple sensory decisions to support higher-order decision making (Shadlen et al, 2008; Dehaene and Sigman, 2012), then categorization might occur inappropriately at intermediate stages before the final decision (Murphy and Ross, 2010). Our findings indicate that normative decision making may, in this way, be fundamentally constrained by the architecture of the perceptual system.

References

- Aly M, Yonelinas AP (2012) Bridging consciousness and cognition in memory and perception: evidence for both state and strength processes. *PLoS One*, 7:e30231. [CrossRef Medline](#)
- Becker GM, DeGroot MH, Marschak J (1964) Measuring utility by a single-response sequential method. *Behav Sci* 9:226–232. [CrossRef Medline](#)
- Bohil CJ, Maddox WT (2001) Category discriminability, base-rate, and payoff effects in perceptual categorization. *Percept Psychophys* 63:361–376. [CrossRef Medline](#)
- Bossaerts P, Preusschoff K, Hsu M (2008) The neurobiological foundations of valuation in human decision making under uncertainty. In: *Neuroeconomics: decision making and the brain* (Glimcher PW, Camerer CF, Fehr E, Poldrack RA, eds), pp 351–364. New York: Elsevier.
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433–436. [CrossRef Medline](#)
- Burnham KP, Anderson DR (1998) *Model selection and inference*. New York: Springer.
- Dehaene S, Sigman M (2012) From a single decision to a multi-step algorithm. *Curr Opin Neurobiol* 22:937–945. [CrossRef Medline](#)
- Feng S, Holmes P, Rorie A, Newsome WT (2009) Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Comput Biol* 5, e1000284. [CrossRef Medline](#)
- Fleming SM, Whiteley L, Hulme OJ, Sahani M, Dolan RJ (2010) Effects of category-specific costs on neural systems for perceptual decision-making. *J Neurophysiol* 103:3238–3247. [CrossRef Medline](#)
- Gelman A, Rubin DB (1992) Inference from iterative simulation using multiple sequences. *Statist Sci* 457–472.
- Gershman SJ, Vul E, Tenenbaum JB (2012) Multistability and perceptual inference. *Neural Comput* 24:1–24. [CrossRef Medline](#)
- Gershman SJ, Daw ND (2012) Perception, action and utility: the tangled skein. In: *Principles of brain dynamics: global state interactions* (Rabinovich MI, Friston KJ, Varona P, eds), pp 293–312. Cambridge: MIT.
- Gold J, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574. [CrossRef Medline](#)
- Gonzalez R, Wu G (1999) On the shape of the probability weighting function. *Cogn Psychol* 38:129–166. [CrossRef Medline](#)
- Green D, Swets J (1966) *Signal detection theory and psychophysics*. New York: Wiley.
- Harnard S (1987) *Categorical perception*. Cambridge UP.
- Healy AF, Kubovy M (1981) Probability matching and the formation of conservative decision rules in a numerical analog of signal detection. *J Exp Psychol Learn Mem Cogn* 7:344–354. [CrossRef](#)
- Hertwig R, Barron G, Weber EU, Erev I (2004) Decisions from experience and the effect of rare events in risky choice. *Psychol Sci* 15:534–539. [CrossRef Medline](#)
- Holmes A, Friston KJ (1998) Generalisability, random effects and population inference. *Neuroimage* 7:S754.
- Kahneman D, Tversky A (1979) Prospect theory: analysis of decision under risk. *Econometrica* 47:263–291. [CrossRef](#)
- Karni E (2009) A mechanism for eliciting probabilities. *Econometrica* 77:603–606. [CrossRef](#)
- Kepecs A, Mainen ZF (2012) A computational framework for the study of confidence in humans and animals. *Phil Trans R Soc B: Biol Sci* 367:1322–1337. [CrossRef Medline](#)
- Kiani R, Shadlen M (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324:759–764. [CrossRef Medline](#)
- Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci* 28:3017–3029. [CrossRef Medline](#)
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247. [CrossRef Medline](#)
- Landy MS, Trommershäuser J, Daw ND (2012) Dynamic estimation of task-relevant variance in movement under risk. *J Neurosci* 32:12702–12711. [CrossRef Medline](#)
- Maddox WT (2002) Toward a unified theory of decision criterion learning in perceptual categorization. *J Exp Anal Behav* 78:567–595. [CrossRef Medline](#)
- Maddox WT, Bohil CJ (2004) Probability matching, accuracy maximization, and a test of the optimal classifier's independence assumption in perceptual categorization. *Percept Psychophys* 66:104–118. [CrossRef Medline](#)
- Maloney LT (2002) Statistical decision theory and biological vision. In: *Perception and the physical world: psychological and philosophical issues in perception* (Heyer D, Mausfeld R, eds), pp 145–189. New York: Wiley.
- Maloney LT, Thomas EAC (1991) Distributional assumptions and observed conservatism in the theory of signal detectability. *Journal of Mathematical Psychology* 35:443–470. [CrossRef](#)
- Murphy GL, Ross BH (2010) Uncertainty in category-based induction: When do people integrate across categories? *J Exp Psychol Learn Mem Cogn* 36:263–276. [CrossRef Medline](#)
- Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature* 341:52–54. [CrossRef Medline](#)
- Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J Vis* 5:376–404.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442. [CrossRef Medline](#)
- Prins N (2012) The psychometric function: the lapse rate revisited. *J Vis* 12(6).
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489. [Medline](#)
- Rorie AE, Gao J, McClelland JL, Newsome WT (2010) Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS One* 5:e9308. [CrossRef Medline](#)
- Shadlen MN, Kiani R, Hanks TD (2008) An intentional framework. In: *Better than conscious? Decision-making, the human mind, and implications for institutions* (Engel C, Singer W, eds), pp 71–101. Cambridge: MIT.
- Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A (2002) Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society B* 64:583–639. [CrossRef](#)
- Summerfield C, Koehlin E (2010) Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. *Front Hum Neurosci* 4:208. [CrossRef Medline](#)
- Trommershäuser J, Maloney L, Landy M (2003) Statistical decision theory and trade-offs in the control of motor response. *Spat Vis* 16 3:255–275. [Medline](#)

- Trommershäuser J, Maloney LT, Landy MS (2003) Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A Opt Image Sci Vis* 20:1419–1433. [CrossRef Medline](#)
- Trommershäuser J, Landy MS, Maloney LT (2006) Humans rapidly estimate expected gain in movement planning. *Psychol Sci* 17:981–988. [CrossRef Medline](#)
- Vickers D, Caudrey D, Willson R (1971) Discriminating between the frequency of occurrence of two alternative events. *Acta Psychol* 35:151–172. [CrossRef](#)
- Wald A (1945) Sequential tests of statistical hypotheses. *Ann Math Statist* 16:117–186. [CrossRef](#)
- Whiteley L, Sahani M (2008) Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *J Vis* 8:2.1–2.15. [CrossRef Medline](#)
- Wichmann FA, Hill NJ (2001) The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept Psychophys* 63:1293–1313. [CrossRef Medline](#)
- Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 26:1314–1328. [CrossRef Medline](#)
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080. [CrossRef Medline](#)
- Zhang H, Maloney LT (2012) Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Front Neurosci* 6:1. [CrossRef Medline](#)