

© 2019 American Psychological Association 0735-7036/19/\$12.00

2019, Vol. 133, No. 3, 380-396 http://dx.doi.org/10.1037/com0000168

Probabilistic Inferential Decision-Making Under Time Pressure in Rhesus Macaques (Macaca mulatta)

Andrew C. Toader, Hrishikesh M. Rao, Minyoung Ryoo, Martin O. Bohlen, Jessi S. Cruger, and Hanna Oh-Descher Duke University

> Tobias Egner Duke University

Silvia Ferrari Cornell University

Jeff Beck Duke University School of Medicine

Marc A. Sommer Duke University

Decisions often involve the consideration of multiple cues, each of which may inform selection on the basis of learned probabilities. Our ability to use probabilistic inference for decisions is bounded by uncertainty and constraints such as time pressure. Previous work showed that when humans choose between visual objects in a multiple-cue, probabilistic task, they cope with time pressure by discounting the least informative cues, an example of satisficing or "good enough" decision-making. We tested two rhesus macaques (Macaca mulatta) on a similar task to assess their capacity for probabilistic inference and satisficing in comparison with humans. In each trial, a monkey viewed two compound stimuli consisting of four cue dimensions. Each dimension (e.g., color) had two possible states (e.g., red or blue) with different probabilistic weights. Selecting the stimulus with highest total weight yielded higher odds of receiving reward. Both monkeys learned the assigned weights at high accuracy. Under time pressure, both monkeys were less accurate as a result of decreased use of cue information. One monkey adopted the same satisficing strategy used by humans, ignoring the least informative cue dimension. Both monkeys, however, exhibited a strategy not reported for humans, a "group-the-best" strategy in which the top two cues were used similarly despite their different assigned weights. The results validate macaques as an animal model of probabilistic decision-making, establishing their capacity to discriminate between objects using at least four visual dimensions simultaneously. The time pressure data suggest caution, however, in using macaques as models of human satisficing.

Keywords: rhesus macaque, decisions, probabilistic inference, satisficing

Everyday decision-making may require the use of multiple, simultaneous cues with different diagnostic values. A person or animal choosing between pieces of fruit, for example, may assess their ripeness, symmetry, and abrasions, that is, the dimensions of color, shape, and texture. Probabilistically, color may be a better indicator of palatability than shape, which may in turn be more important than texture. In the human cognitive literature, learning the statistical properties of such cues is termed *probabilistic inference* (for reviews, see Fiser, Berkes, Orbán, & Lengyel, 2010 and Ma, 2012).

The ability to make decisions based on probabilistic inference of multiple, simultaneous cues, such as in our fruit example, has been

Hrishikesh M. Rao is now at MIT Lincoln Laboratory, Lexington, Massachusetts.

This document is copyrighted by the American Psychological Association or one of its allied publishers.

This article was published Online First February 25, 2019.

Andrew C. Toader, Hrishikesh M. Rao, Minyoung Ryoo, Martin O. Bohlen, and Jessi S. Cruger, Department of Biomedical Engineering, Duke University; Hanna Oh-Descher, Department of Psychology & Neuroscience and the Center for Cognitive Neuroscience, Duke University; Silvia Ferrari, Laboratory for Intelligent Systems and Control, Sibley School of Mechanical and Aerospace Engineering, Cornell University; Tobias Egner, Department of Psychology & Neuroscience and the Center for Cognitive Neuroscience, Duke University; Jeff Beck, Department of Neurobiology, Duke University School of Medicine; Marc A. Sommer, Department of Biomedical Engineering and Center for Cognitive Neuroscience, Duke University.

We thank C. Giattino and K. Rafie for preliminary data collection and Z. Abzug for assistance with experimental design. Funded by the Office of Naval Research (award N000141310561) and the National Science Foundation (IGERT Award DGE-1068871).

Correspondence concerning this article should be addressed to Marc A. Sommer, Department of Biomedical Engineering and Center for Cognitive Neuroscience, Duke University, 1427 CIEMAS, Box 90281, 101 Science Drive, Durham, NC 27708. E-mail: marc.sommer@duke .edu

381

studied extensively in humans (Gluck & Bower, 1988; Knowlton, Squire, & Gluck, 1994; Medin, Wattenmaker, & Hampson, 1987; Nelson, 1984). Much of that work focused on categorization, and recent studies extended the approach to macaques. Presented with stimuli consisting of multiple visual cues, some of which are strong indicators of category ("criterial cues"), it was shown that macaques infer and use the criterial cues to varying degrees depending on task design and cue dimensionality (Couchman, Coutinho, & Smith, 2010; Smith, Beran, Crossley, Boomer, & Ashby, 2010; Smith et al., 2012). Sigala and Logothetis (2002), similarly, trained macaques on a categorization task that required probabilistic inference of two out of four cue dimensions. Macaques can also infer the probabilities of visual cues that signify various forms of reinforcement (negative vs. positive: Amemori & Graybiel, 2012; risky vs. ambiguous: Hayden, Heilbronner, & Platt, 2010; flavor and amount: Raghuraman & Padoa-Schioppa, 2014) or cues that are revealed sequentially (Kira, Yang, & Shadlen, 2015; Yang & Shadlen, 2007). Furthermore, by using opt out and bail out options as a measure of confidence, it has been demonstrated that animals and humans have a probabilistic interpretation of their solutions to such tasks (Kiani & Shadlen, 2009; for a review of both humans and animals, see Kepecs & Mainen, 2014). The present study extends all this work to examine, in a highly quantified task framework, the extent to which macaques apply probabilistic inference in the naturalistic situation of deciding between two objects that each consist of several visual features, all of which have varying levels of learned importance.

Decisions based on multiple, simultaneous cues may need to be made under conditions of uncertainty or time pressure. In the example of choosing fruit, a fast decision may be necessary to beat a conspecific to the best piece. Faced with such challenges, humans switch to satisficing strategies to reach decisions that are "good enough"; although not necessarily optimal, they are satisfactory and sufficient for the task at hand (Gluck & Bower, 1988; Simon, 1955, 1956). Satisficing includes, for instance, the use of heuristics or "rules of thumb" (Gigerenzer & Gaissmaier, 2011; Rieskamp & Hoffrage, 2008). Human satisficing strategies during probabilistic decision-making were studied recently by Oh and colleagues (Oh et al., 2016) using tasks that involved four simultaneously presented dimensions of visual information. They found that, under time pressure, humans adopt a "drop-the-worst" heuristic by ignoring the least informative cue dimensions. This change in strategy is accompanied by a shift in brain activity from cortical to subcortical networks as indicated by functional MRI (Oh-Descher, Beck, Ferrari, Sommer, & Egner, 2017). Another common satisficing strategy is the "take-the-best" heuristic (Gigerenzer & Goldstein, 1996) in which subjects search through cues from the most to least informative and stop when they find the first cue that discriminates between the possible choices (Gigerenzer & Goldstein, 1996; Karelaia, 2006; Newell, Weston, & Shanks, 2003).

The neural basis of probabilistic inference and satisficing could be explored in more detail using rhesus macaques if we had a better understanding of their behavior. Therefore, the goal of the present study was to test two hypotheses:

Hypothesis 1: Rhesus macaques use probabilistic inference to choose between stimuli that are each defined by jointly presented, statistically informative visual dimensions.

Hypothesis 2: They adopt satisficing strategies when making those decisions under time pressure.

We trained macaques on the same decision-making task tested on humans (Oh et al., 2016) that involved visual stimuli composed of four, probabilistically weighted cue dimensions. Using psychometric and modeling approaches, we analyzed the animals' strategies for performing the task under low and high time pressure. Our results supported both hypotheses, with the caveat that satisficing strategies differ in some ways between rhesus macaques and humans.

Method

Animals and Apparatus

Two female rhesus monkeys (Macaca mulatta; Monkey A, 6 kg, and Monkey L, 8 kg) were tested in their cages using touchscreens (Elo 1537L 15" Intellitouch SER/USB, Elo Touch Solutions, Inc., Milpitas, CA) mounted onto the cage using custom built attachments. A tube above the screen delivered 1 cc of water reward for every correct trial. The animals received at least 20 cc/kg of water per day. If they did not earn this amount of water during a session, they were provided with the remaining water afterward. They also received supplemental fruit daily. Aside from water control, there were no other modifications to the monkeys' routine husbandry. No invasive procedures were required, animals were pair-housed, and environmental enrichment was provided. Experimental control, including stimulus and reward delivery, was performed using Psychtoolbox in MATLAB (www.mathworks.com). All procedures were approved by the Institutional Animal Care and Use Committee at Duke University (protocol A079-17-03).

Task Overview

We used the same probabilistic decision-making task used in Experiments 1 and 2 in the Oh et al. (2016) study except for slight modifications to facilitate monkey testing (e.g., four possible locations of stimuli instead of two, to discourage spatial habits and maintain attentiveness). Figure 1 illustrates the overall procedure. A monkey viewed stimuli on its in-cage touchscreen. Trials started automatically with a white cross appearing in the center of the gray background for 0.7 s. When the cross disappeared, two 2.5×2.5 cm stimuli appeared in two of four possible quadrants. Each stimulus was constructed of features that varied across four visual dimensions (Figure 1A). The two stimuli remained visible for a maximum duration termed the response window (details in next paragraph). Before the response window elapsed, the monkey had to select one of the stimuli by touching it (Figure 1B). If it did not, the trial was aborted. On each trial, two stimuli were drawn without replacement from the 16 possible stimuli (Figure 1C). Distributions of touches to the target locations were distinct (Figure 1D). As soon as a touch was detected, the stimuli disappeared. If the animal earned water reward, it was delivered immediately. After a 2-s intertrial interval, the central cross reappeared to start a new trial. Monkeys worked to satiation (~4 hr, several hundred trials/day).

In each testing session, a single response window was used that defined the time pressure. For "Low Pressure" (LP) sessions, the



Figure 1. Overview of the task and data collection. (A) The four cue dimensions and the cue states for each. (B) Monkey A using the in-cage touchscreen setup. (C) All 16 possible compound stimuli. (D) Heat map of responses around each of the four possible stimulus locations during a typical session. Solid lines indicate boundaries of stimulus display $(2.5 \times 2.5 \text{ cm})$, and dotted lines show the allowed response area $(4.4 \times 4.4 \text{ cm})$.

response window ranged from 1.75 to 2 s. For "High Pressure" (HP) sessions, the response window ranged from .35 to .50 s. We tested Hypothesis 1, which states macaques use probabilistic inference to choose between stimuli that are each defined by jointly presented, statistically informative visual dimensions, by analyzing data from LP sessions. We tested Hypothesis 2, which states that monkeys make these decisions under time pressure using satisficing strategies, by comparing data between HP and LP sessions.

Compound Visual Stimuli and Weighting

Each visual stimulus was a compound object consisting of four individual visual features, which we refer to as *cue dimensions*, each of which could take on one of two *cue states* (Figure 1A), as follows: *Color dimension*: A stimulus could have a blue or red background; *Shape dimension*: A stimulus could be a circle or a square; *Border dimension*: A stimulus could be outlined in white or black; and *Orientation dimension*: Lines within a stimulus could be horizontal or vertical. Given the combination of these cue dimensions and possible cue states, there were 16 unique compound stimuli (Figure 1C). On a given trial in which two different stimuli were presented, there were therefore 120 possible unique pairs of compound stimuli.

An optimal decision to select one stimulus over another required evaluation of evidence from the four cue dimensions and a comparison between the two different stimuli. Within a cue dimension (e.g., color), each of the two possible cue states (blue or red) had an *assigned weight*. The difference between these assigned weights defined the *net weight* or relative importance for a cue dimension. Table 1 shows all the assigned weights. In experiments with Monkey A, for example, "blue" had an assigned weight of .9, and "red" had an assigned weight of .1, giving the cue dimension "color" a net weight of .8. At the same time, the cue dimension

Table 1Visual Stimulus Parameters

		Assigned weights			
Cue dimension	Cue states	Monkey A	Monkey L		
Color	Blue	.9	.4		
	Red	.1	.6		
Shape	Circle	.8	.3		
1	Square	.2	.7		
Border	White	.7	.2		
	Black	.3	.8		
Orientation	Horizontal	.6	.1		
	Vertical	.4	.9		

Note. Each stimulus had four cue dimensions, each with two states. Weights were assigned as shown. Between monkeys, the ordering of importance of the cue dimensions was reversed, and, within each dimension, the ordering of the states was flipped.

"border" had a net weight of .4 because the assigned weights for white and black borders were .7 and .3, respectively. For Monkey A, therefore, color evidence was more informative than border evidence. Net weights were limited to four values, .2, .4, .6, or .8, but the cue dimensions corresponding to those net weights were reversed between the two monkeys, that is, in order of probabilistic importance, color > shape > border > orientation for Monkey A, but orientation > border > shape > color for Monkey L. Also, within each dimension, the relative importance of each cue state was flipped between the monkeys. In the color dimension, for instance, blue had a higher weight than red for Monkey A but a lower weight than red for Monkey L (Table 1; for details, see the Training section in the following text).

The assigned weights of each cue state were independent and additive, yielding a summed, *total weight* for each compound stimulus. For each pair of stimuli in a trial, the *sum of evidence* (SoE) in favor of one stimulus was the difference between its total weight and that of the other stimulus. This SoE was used to calculate the probability of receiving reward after touching a stimulus as follows:

$$P(\text{reward} \mid SoE) = \frac{1}{1 + 10^{-SoE}} \tag{1}$$

For each pair of stimuli, therefore, selecting the one with the positive SoE (greater total weight) resulted in higher, abovechance probability of reward; it was the *correct* decision for optimizing reward. Performance was calculated on the basis of making these correct decisions. Trials where the SoE was 0 were ignored when calculating accuracy. Note, monkeys were most likely to receive reward if they selected the stimulus with positive SoE, but because reward was probabilistic, selecting that stimulus might yield no reward, or selecting the other stimulus might be rewarded. This reward contingency encouraged the animals to make their decisions by combining the probabilistic information of the stimuli, rather than on the basis of stimulus-response contingencies. The latter would require memorization of 120 possible pairs of compound stimuli.

Training

To counterbalance our experimental design, the two monkeys learned cue dimensions with reversed order of importance and, within each cue dimension, opposite weighting of cue states (Table 1). For example, color was the most informative dimension for Monkey A (net weight = .8) but the least informative for Monkey L (net weight = .2). Within the color dimension, blue was weighted more than red for Monkey A but red was weighted more than blue for Monkey L. The purpose of this manipulation between monkeys was to control for any differences in visual saliency within and across dimensions. Performance therefore benefited from learning the probabilistic weights of the features regardless of whether some features seemed more salient than others.

Neither of the monkeys had previous experience with this set of visual cue dimensions. Both monkeys required several weeks of training to learn the task (see Learning Phase in Results). Response windows were set to 2 s during training, so that monkeys did not experience time pressure until the later testing sessions. Monkey A was the first one trained. Because this monkey was used to pilot the task and settle on appropriate spatial, temporal, and visual

dimension parameters, we do not have a learning curve for it. For the second one, Monkey L, we used constant parameters and documented the learning curve thoroughly. After a monkey reached asymptotic levels of performance, formal testing sessions commenced, one per day.

Subjective Weights

A major goal of our analysis was to infer the subjective importance of the various experimental parameters on each monkey's decision-making. The central analytic approach was to perform a logistic regression using variational Bayesian inference (Drugowitsch, 2013; Oh et al., 2016) to calculate subjective weights, the inferred values assigned by each monkey to different experimental parameters. Specifically, we analyzed the contribution of the four cue dimension net weights (.2, .4, .6, and .8), as well as the six possible spatial configurations of the stimulus pairs (up left-up right, up left-down left, up left-down right, up right-down left, up right-down right, and down left-down right). Although stimulus location was irrelevant to reward probability, monkeys exhibited varying degrees of preference for certain locations in each location pair, a form of spatial bias likely related to body posture and the arm used. After completing data collection, we calculated the average location preference order across all trials for both monkeys. To keep the direction of location preference consistent between monkeys, the stimulus in the preferred location on a given trial was treated as the positive response in the logistic regression. This was done by first labeling the stimuli in the more and less preferred locations as "Stimulus A" and "Stimulus B," respectively. Bayesian logistic regression was then performed to infer the posterior probability of choosing Stimulus A, given all of a monkey's selections on an experimental session. In this way, any spatial biases would be realized positively in the "A" direction for all location pairs and monkeys, independently of cue weight effects

Details of the Bayesian logistic regression are provided in Appendix 1. In brief, the analysis used the cue dimension net weights and target location pairs for each trial, in each session, to fit the monkey's choice responses. The analysis yielded subjective cue weights for each of the four cue dimensions and subjective spatial weights for each of the six spatial location pairs that best explained the monkey's responses. As a measure of decision noise and spatial bias, we calculated the square root of the sum of squares (SRSS) of the subjective cue weights and subjective spatial weights, respectively. To assess strength of the learned cue weight associations, it is useful to compare the SRSS of cue weights to that of an ideal observer. In related human literature, humans tend to match their choice probabilities to the true outcome probabilities (Oh et al., 2016; Vulkan, 2000). We therefore chose to compare monkey performance to an observer that makes decisions based on sampling from the true posterior, that is, on the observed probability of outcomes. This idealized decision process would yield SRSS = 1.095 (the norm of the net weights) and, if nothing else affected performance, an accuracy of approximately 79% (the expected probability of reward). Decisions based on inference of the underlying cue weights rather than reward probability would increase SRSS and accuracy above these levels.

To isolate relative differences in subjective cue weighting from overall task knowledge, subjective cue weights were then normalized to their SRSS. Finally, the weights were scaled by the precision of the regression estimates for each session and averaged for each cue to yield the weighted-mean normalized subjective cue weight \bar{w}_{c_i} and its uncertainty $\sigma_{\bar{w}_{c_i}}$ for that cue. Monkeys performing with variable degrees of decision noise but the correct relative valuation of cues would have the same normalized subjective cue weight (represented by a 0.2:0.4:0.6:0.8 ratio). For subjective spatial weights, SRSS measured spatial bias, the a priori influence of location before stimuli onset. Decisions made with no spatial bias would yield SRSS = 0. Because the evidence was spatially unbiased, spatially biased behavior would increase SRSS and decrease accuracy.

Decision Model Comparisons

To determine which experimental factors were being utilized for task performance, we searched a model space of 32 decision strategies that accounted for every combination of cue dimension and inclusion or exclusion of spatial location (similar to Oh et al., 2016; Figure 2). Specific models of interest include an "optimal compensatory strategy" of using information from all four cue dimensions to make decisions (Figure 2, Models 15 and 31), the "drop-the-worst" strategy used by humans under time pressure (Figure 2, Models 11 and 27; Oh et al., 2016), and the "take-thebest strategy" that searches through cue dimensions in descending order of their assigned weights and stops on the first cue that differentiates the two options (Figure 2, Models 16 and 32; Gigerenzer & Goldstein, 1996). The remaining 29 models comprised all possible combinations of cue dimensions crossed (starred models in Figure 2), or not, with the six spatial location pairings. The Bayesian regression described earlier (Drugowitsch, 2013) was performed for all 32 possible decision strategy models to calculate a lower bound on the marginal log probability of each model in each session. To characterize decision strategy across multiple sessions, these log model evidences were submitted to a Bayesian model selection procedure (Stephan, Penny, Daunizeau, Moran, & Friston, 2009). This approach fits a hierarchical model by treating the decision strategy models as random effects that could vary across sessions and estimates the protected exceedance probabilities, which represent the probability that each given model is more likely than others, corrected for the possibility that observed differences are due to chance (Rigoux, Stephan, Friston, & Daunizeau, 2014). The model with the highest protected exceedance probability was selected as the preferred decision strategy model for a given monkey and time pressure. Protected exceedance probabilities were calculated using the spm_BMS routine of the SPM12 software suite (http://www.fil.ion.ucl.ac.uk/ spm/software/spm12/).

Results

Learning Phase

Monkeys were initially trained on the LP version of the task. The learning curve, measured as a fraction of correct responses, is shown for Monkey L in Figure 3A. "Session 1" corresponds to the first day of training. The power law fit shows a rapid increase early in training as the animal began to recognize the probabilistic importance of the different visual stimuli, followed by a more gradual rise toward a steady state plateau. In Figure 3B-D, normalized subjective cue weights are plotted against net weights for three individual sessions (labeled with arrows in Figure 3A). Note that spatial location was included in these regressions but not depicted here. Early in training, this monkey heavily favored the third cue dimension, which had net weight of 0.6, even though the fourth dimension was more informative at net weight 0.8 (Figure 3B). In time, the monkey appeared to recognize the importance of the fourth cue dimension (Figure 3C), culminating in the correct ordering of cue dimensions during steady state performance (Figure 3D). In the learning curve of Figure 3A, black dots indicate sessions in which the monkey accomplished the correct ordering of cue dimensions. The first occurrence of 3 consecutive days of correct ordering (Session 95 in Figure 3A) defined the end of the learning phase.

Through the rest of this report, data from the learning phase were excluded. As an additional criterion, we excluded aborted trials and each trial that immediately followed a no-response aborted trial (n = 7,679 trials, 12.9% of all trials for Monkey A; n = 13,958 trials, 14.5% of all trials for Monkey L), to include only those trials performed at a similar state of alertness or engagement with the task and allow for an accurate measure of response time. Both monkeys were free to roam in their cage and generally performed the task in "bursts" of trials, so this criterion typically excluded the first in a string of completed trials. Experimental sessions were excluded if the number of trials performed was too low for the logistic regression to converge. The final data



Figure 2. Complete space of strategy models explored. Filled black circles indicate that the corresponding cue dimension is included in the regression for the given model. Gray circles under Models 16 and 32 are used to indicate that these models use a "take-the-best" strategy, where the highest valued cue dimension that differs between both stimuli is used to make a decision.



Figure 3. Performance of Monkey L during training. (A) Fraction correct across sessions, fit with a power law curve. Starting with first exposure to the task, performance in 116 sequential sessions are shown (except for six sessions omitted due to insufficient trials for regression convergence). Black circles indicate sessions in which subjective cue weights were properly ordered according to net weights (B-D). Examples of improvement in cue weight ordering during training. Normalized subjective cue weights are plotted against net weights for (B) Session 46, (C) Session 83, and (D) Session 116. In panel "A," the data points corresponding to these sessions are labeled *b*, *c*, and *d*, respectively. All training sessions were low pressure trials.

set consisted of 52,085 trials, with median of 396 (interquartile range 204 to 700) trials per session from Monkey A and 68,631 trials, with median of 895 (interquartile range 466 to 1,252) trials per session from Monkey L.

Steady-State Behavior at LP

Once trained, the monkeys continued to perform LP sessions (Monkey A: 56 sessions, 27,773 total trials; Monkey L: 51 sessions, 57,772 total trials). In these LP trials, Monkey A selected the correct stimulus at a rate of 86.0%, and Monkey L was correct at 86.5%, with mean response times of 399 ms and 634 ms, respectively. Across these trials, both monkeys accurately discriminated between features in all four visual cue dimensions by inferring the correct order of subjective cue (Figure 4). All subjective cue weights were significantly different in a monotonic rising progression as a function of net weights (statistics are in the Satisficing section, and numerical data are in Table 2 associated with that section). Recall that the assigned dimension order of importance and within-dimension cue feature weights were counterbalanced between Monkeys A and L, so the ordering was not a salience artifact in monkeys (e.g., if color happened to be easier to discern than orientation of the lines). Both monkeys had an average SRSS

of subjective cue weights much greater than 1.095, implying that they used probabilistic inference of cue weights to improve accuracy above the expected reward rate (see Method).

Performance Under HP

Under HP conditions, mean accuracies across all trials for Monkeys A and L were significantly decreased to 77.23% (Wilcoxon's rank sum, Cohen's d = -.230, z = 25.56, p < .0001) and 72.57% (Wilcoxon's rank sum, Cohen's d = -.387, z = 35.89, p < .0001), respectively. The mean response times decreased to 353 ms (*t* test, Cohen's d = -.354, t = 40.78, p < .0001) and 486 ms (*t* test, Cohen's d = -.515, t = 50.26, p < .0001) over 48 sessions (24,312 total trials) for Monkey A and 21 sessions (10,533 total trials) for Monkey L.

The degradation of performance in HP trials coincided with a decreased ability of the monkeys to utilize and discriminate between the cue features. Figure 5 shows the psychometric functions between the proportion of trials a stimulus was selected (ordinate) versus its SoE (abscissa). The greater the SoE of a stimulus (recall Equation 1), the more likely it was to be selected. The slope at the inflection point is a measure of the perceptual sensitivity to SoE as a factor in performance. In HP sessions (Figure 5, white circles), the



Figure 4. Performance for Monkey A (top) and Monkey L (bottom) at steady state. Left: mean normalized subjective cue weights (\bar{w}_{c_i}) across low pressure sessions. Error bars, smaller than the symbols, indicate uncertainty $(\sigma_{\bar{w}_{c_i}})$. Right: Mean (and standard error of the mean) of the SRSS of subjective cue weights across low pressure sessions. The high SRSS levels, much greater than 1.095 (dotted line), indicates the use of probabilistic inference for high accuracy rather than the matching of expected reward rates.

psychometric slopes clearly decreased relative to LP sessions (Figure 5, gray circles) for both monkeys. That is, under time pressure, performance was still driven by SoE, but with less sensitivity to it.

Multiple hypotheses could account for these effects of time pressure on performance. One possibility is that the monkeys occasionally made cue-independent decisions in HP trials due to lapses in attention (Wichmann & Hill, 2001). This predicts, however, that the psychometric curves would not asymptote to zero or one, and yet they do (Figure 5). Hence this "lapse" hypothesis was ruled out. Three other hypotheses and their predictions are as follows:

(1) Satisficing: Limitations in cognitive capacity under time pressure could cause subjects to change how they evaluate the cue dimensions. The least informative cue di-

Table	2
-------	---

Data Shown in Figure 6	, Plus	Effect Sizes	(Cohen's	d) for	High Pressure-	-Low Pressure	Difference at	Each Net	Weight
------------------------	--------	--------------	----------	--------	----------------	---------------	---------------	----------	--------

		Monl	key A			Monl	key L		
Net weights	.2	.4	.6	.8	.2	.4	.6	.8	
Low pressure condition		N = 56				N = 51			
M, \bar{w}_{c}	.0797	.3213	.5630	.7395	.2151	.4124	.5877	.6472	
Uncertainty, $\sigma_{\bar{w}}$.0040	.0038	.0037	.0037	.0029	.0030	.0030	.0030	
High pressure condition		N = 48				N = 21			
M, \overline{w}_{c}	.0740	.3087	.5980	.6861	.1984	.3652	.6027	.6045	
Uncertainty, $\sigma_{\bar{w}}$.0047	.0046	.0045	.0045	.0090	.0091	.0093	.0092	
Cohen's $d^{n_{c_i}}$	0867	1952	+.5487	8372	2451	6854	+.2165	6182	



Figure 5. Performance as quantified by the psychometric curves for selecting a stimulus as a function of evidence for it (sum of evidence [SoE]), across all low pressure (gray) and high pressure (white) trials. The shallower the slope at SoE = 0, the less sensitive the performance was to SoE.

mension could be ignored, for example, as found by Oh et al. (2016). Such satisficing strategies would predict changes in the strengths or ordering of the normalized subjective cue weights.

- (2) Decision noise: Time pressure may introduce more noise, or randomness, into the decision-making process. This would predict a decrease in the SRSS of subjective cue weights under time pressure, with no effect on the relative strengths or ordering of the normalized subjective cue weights.
- (3) Biased priors: When evidence is weak due to time pressure, subjects might rely more on a priori preferences. Here, we consider priors on stimulus locations, that is, the animals' spatial biases. Biases toward these priors under time pressure would predict increases in the SRSS of subjective spatial weights.

These three hypotheses are not mutually exclusive. In the following sections, we consider the evidence for each of them.

Satisficing

As a first step in determining whether the monkeys used satisficing strategies to change their relative weighting of cue dimensions under time pressure, we directly compared the individual subjective cue weights between LP and HP sessions (Figure 6). Specifically, the precision-weighted means of each (normalized) subjective cue weight was compared four times: between the cue dimensions within the LP or HP condition (three comparisons), and between the LP and HP condition for each cue dimension (one comparison), yielding a corrected criterion level of p < .0125. The null hypothesis-no effect of time pressure, that is, HP condition-predicts significant differences for the first three comparisons (clear ordering of the subjective cue weights) but no significant difference for the fourth comparison (no effect of the HP condition on subjective cue weights). In Figure 6, we label each result that deviated from this null hypothesis. For Monkey A, the first three comparisons (across cue dimensions) were all significant, as expected, in both the LP and HP conditions (t tests, all p <.0125). That is, subjective cue weights were ordered systematically



Figure 6. Mean normalized subjective cue weights (\bar{w}_{c_i}) for low pressure (Low Pressure [LP]; gray) and high pressure (High Pressure [HP]; white) sessions for each monkey. Arrows and labels show pairs of values that violated the null hypothesis in that they were significantly different between the LP and HP conditions (*) or not significantly different across dimensions (nsd). Error bars, smaller than the symbols, indicate the uncertainty $(\sigma_{\bar{w}_c})$. Note, the LP data are the same as in Figure 4, reproduced here for direct comparison with the HP data.

as a function of net weights. The fourth comparison (LP vs. HP) was not significant for the lower two dimensions (p > .0125) but was significant for the higher two dimensions. The subjective cue weight dropped in the HP condition for net weight = .8 (p = 4.00×10^{-8}) and rose for net weight = .6 ($p = 8.54 \times 10^{-15}$), bringing the two subjective weights closer together. For Monkey L, all comparisons across cue dimensions were significant in LP sessions (p < .0125). During HP sessions, however, the subjective cue weights for net weights .6 and .8 were not significantly different from each other (p = .89), as if the monkey treated them the same. This was due to a significant drop in subjective weighting for net weight .8 in HP sessions ($p = 1.74 \times 10^{-4}$). A similar drop was found for weight .4 ($p = 4.65 \times 10^{-5}$), but no HP versus LP difference was found for weights .2 and .6 (p > .0125). The numerical data from Figure 6 are listed in Table 2. The common finding between the two monkeys, therefore, was a shift in relative subjective cue weighting for the top cue dimensions, reducing (or in Monkey L, eliminating) the subjective difference between them. This "group-the-best" satisficing strategy would serve to reduce the dimensionality of the task from four sources of probabilistic information to three.

We then examined whether this "group-the-best" effect was present at the level of individual sessions. For each session, we calculated the difference in subjective cue weights for the .6 and .8 cue dimensions. Smaller differences would represent closer subjective weights and more of a grouping effect. We found that these differences were significantly smaller in the HP condition for both monkeys (Monkey A: means .153 in HP vs. .228 in LP, Cohen's d = .406, p = .041; Monkey L: means -.0245 in HP vs. .0655 in LP, Cohen's d = .779, p = .00370; t tests). To compare the effects for both monkeys directly, we performed a two-factor analysis of variance on the differences using monkey identity and time pressure condition as factors. Both main effects were significant, due to smaller overall differences for Monkey L (F = 43.16, p = 5.76×10^{-10}) and smaller overall differences in the HP condition (F = 10.08, p = .00178). The interaction was not significant, meaning that the test found no difference in HP versus LP effects between monkeys (F = .0899, p = .765). In sum, this session-bysession analysis confirmed that both monkeys used a "group-thebest" strategy under high time pressure.

When analyzing the same task in humans, Oh et al. (2016) applied Bayesian model selection analyses to infer how strategies changed in HP versus LP conditions. This approach evaluates strategies that involve using subsets of the visual cues (Figure 2). Under time pressure, humans consistently resorted to a "drop-theworst" satisficing strategy (Oh et al., 2016). We applied the same analysis to the data from our monkeys across all experimental sessions in the HP and LP conditions. Figure 7 illustrates the protected exceedance probabilities calculated for each decision model strategy. For Monkey A (Figure 7A), the maximum marginal likelihood strategy in the LP condition (gray bars) was Model 31, in which all four cue dimensions are used with spatial bias. Under time pressure (HP condition, white bars), however, Monkey A shifted to Model 27, the "drop-the-worst" strategy with spatial bias. Monkey L (Figure 7B) did not exhibit a significant shift in strategy that could be captured by these models. Its preferred strategy for LP and HP was Model 31. In sum, both animals used a "group-the-best" satisficing strategy for probabilistic decision-



Figure 7. Protected exceedance probabilities for all 32 strategy models during low pressure (gray) and high pressure (white) sessions for (A) Monkey A and (B) Monkey L.

making under time pressure, as described earlier, but Monkey A also used a "drop-the-worst" strategy as found for humans.

Decision Noise

To evaluate whether increases in decision noise affected performance in HP sessions, we calculated the average SRSS of the subjective cue weights for LP and HP sessions for both monkeys (Figure 8). In this pooled data, the average SRSS decreased significantly in HP sessions for Monkey L (Figure 8, right; mean SRSS = 1.65 for HP and 2.08 for LP, Cohen's d = -.909, Wilcoxon's rank sum, z = 2.21, p = .027) but not for Monkey A (Figure 8, left; mean SRSS = 2.22 for HP and 2.48 for LP, Cohen's d = -.354, Wilcoxon's rank sum, z = 1.96, p = .501). However, as will be shown in the Response Time Effects section, Monkey A did show effects of decision noise when controlling for response time. To a modest degree, therefore, increased decision noise was a factor in the monkeys' decision-making under time pressure.

Biased Priors

If monkeys were spatially unbiased, we would expect an even distribution of stimulus location selections. However, even in LP sessions, both monkeys showed consistent spatial preferences (Figure 9A, gray). They were biased toward selecting stimuli in the upper two quadrants of the screen and to the right for Monkey A or the left for Monkey L. This is counterproductive because stimulus locations are irrelevant to the probability of receiving reward. The reasons for these biases are unknown but may relate to individual handedness or postures while using the touchscreen. In HP sessions (Figure 9A, white), the innate spatial biases were accentuated. The resulting reduction in performance was evident upon recalculating the psychometric functions of Figure 5 with



Figure 8. Mean square root of the sum of squares (SRSS) of subjective cue weights across low pressure (LP; gray) and high pressure (HP; white) sessions for each monkey. The mean SRSS decreased significantly (p < .05) under time pressure for Monkey L (right) but not Monkey A (left). Error bars indicate standard error of the mean. Dotted line at 1.095 indicates the SRSS of assigned net weights. Note, the low pressure data are the same as in Figure 4, reproduced here for direct comparison with the high pressure data.

respect to selection of the stimulus in each pair that was at the more-preferred location (termed "Stimulus A"; Figure 9B). Not only did the psychometric slope decrease under HP, the inflection point exhibited a negative shift, indicating elevated decisions to preferentially located Stimulus A even when cue dimension evidence (SoE) was against it (<0). The contribution of prior biases to performance under time pressure is summarized by the SRSS of subjective spatial weights (Figure 9C), which increased significantly in HP sessions for both monkeys (Monkey A: means 2.59 in HP and 1.16 in LP, Cohen's d = 1.355, z = 5.77, $p = 7.75 \times 10^{-9}$; Monkey L: means 2.53 in HP and 0.988 in LP, Cohen's d = 1.798, z = 4.15, $p = 3.32 \times 10^{-5}$; Wilcoxon's rank sum tests). Hence, one way that the monkeys compensated under high time pressure was by exaggerating their biased priors for stimulus location.

Response Time Effects

One consequence of time pressure is shortened durations for viewing the stimuli. Having less time to see the stimuli could, by itself, worsen performance metrics such as accuracy, decision noise, and spatial bias. Viewing duration had an upper bound of the response window, but on any particular trial, it was truncated by the monkey's response time (latency between onset of the stimuli and touching the screen). Therefore, we compared data from LP and HP trials matched by response times, and thus viewing times, to see whether the effects of time pressure on accuracy, decision nose, and spatial bias were still present (Figure 10). For a list of all the numerical data from Figure 10, see Table A2 in Appendix 2. Specifically, we compared performance accuracy (Figure 10, top), SRSS of subjective cue weights (Figure 10, middle), and SRSS of subjective spatial weights (Figure 10, bottom) in LP versus HP trials within equal-sized response time bins (bootstrapping, 10,000 resamples with replacement, using p < p.0033 criterion to correct for the 15 LP-HP comparisons in each plot). The number of trials in each bin was set to the mean number of trials in LP and HP sessions for each monkey (Monkey A: 496 trials/bin for LP and 506 trials/bin for HP; Monkey L: 1132 trials/bin for LP and 502 trials/bin for HP). At many of the response time bins analyzed, both monkeys had lower accuracy,

lower SRSS of subjective cue weights, and higher SRSS of subjective spatial weights in HP compared with LP trials (Figure 10). Hence the effects found in HP trials were *not* just the equivalent of a data analysis that excluded slow trials. Furthermore, the effects of time pressure were not due to limited sensory exposure in HP trials but were consistent with changes in behavioral and cognitive strategies that allowed monkeys to adapt to the context of high time pressure during the session. Note that when the results are broken down by response time, the variability in SRSS of subjective cue weights for Monkey A explains the negative result in this animal's pooled data (Figure 8, left). The finer-grained analysis of Figure 10, middle, demonstrates that at nearly all response times, decision noise increased (i.e., SRSS of subjective cue weights decreased) similarly in HP trials for both monkeys.

Discussion

Two rhesus macaques were trained to perform discrimination decisions accurately by evaluating compound visual stimuli and probabilistic feedback under different levels of time pressure. A variational Bayesian regression was used to assess the monkeys' subjective weighting of each dimension of the compound stimuli as compared with the objective weight assignment. For LP sessions (response window of 1.75 s to 2 s), both monkeys performed the task accurately by correctly ranking the relative importance each cue dimension. This supported our first hypothesis, that rhesus macaques use probabilistic inference to choose between stimuli that are each defined by jointly presented, statistically informative visual dimensions. Importantly, the visual cues combined to form discrete objects much the same way that colocalized features define objects of interest in a monkey's environment (e.g., a piece of food or a face). This contrasts with previous work that used probabilistic cues to study categorization of single objects (Couchman et al., 2010; Sigala & Logothetis, 2002; Smith et al., 2010), reinforcement context (Amemori & Graybiel, 2012; Hayden et al., 2010; Raghuraman & Padoa-Schioppa, 2014), or the gradual accumulation of evidence (Kira et al., 2015; Yang & Shadlen, 2007). Our data contribute to a quantitative understanding of how macaques use probabilistic inference to compare and



Figure 9. Spatial effects during low pressure (LP; gray) and high pressure (HP; white) sessions for each monkey. (A) Overall fraction of trials in which a stimulus at each location was selected in LP and HP trials. UL = up left; UR = up right; DL = down left; DR = down right. (B) Psychometric curves for LP and HP trials for Stimulus A, the stimulus within each pair that was in the more-preferred location. (C) Square root of the sum of squares of subjective spatial weights in LP and HP trials.

choose between the types of visually complex objects they naturally encounter.

Under HP (response window \leq 500 ms), the animals still performed the task, but at lower accuracy. Several factors accounted for this change in performance under time pressure. First, decision noise increased, as indicated by a drop in the SRSS of subjective cue weights. Behaviorally, this translates to a higher degree of randomness in choices under pressure. Second, spatial biases increased, as indicated by a rise in the SRSS of subjective spatial weights. Even though spatial biases were always counterproductive (stimulus location was not a factor in correct decision-making), both monkeys exhibited a baseline level of bias even during low time pressure sessions. Under high time pressure, the same biases were present, but accentuated. Third, satisficing strategies emerged. Monkey A showed evidence of a "drop-the-worst" satisficing strategy under high time pressure as found in humans (Oh et al., 2016): The least important cue dimension was ignored in favor of more informative dimensions. Also found in this monkey, and even more so in Monkey L, was evidence for a "group-the-best" strategy, in that the subjective weights of the top two cue dimensions converged toward each other. Finally, all these effects were categorical, influenced by the context of performing decisions in low- versus high-pressure sessions. They were not simply artifacts of making responses more quickly as demonstrated by response time analyses. The results therefore supported our second hypothesis, that rhesus macaques adopt satisficing strategies when making probabilistic decisions under time pressure,



Figure 10. Response time effects for low pressure (gray) and high pressure (white) trials. Response time bins (*x* axis) are of equal size within a condition, centered at equal intervals. As a function of response time for Monkey A (left) and L (right), the graphs plot accuracy (top row), square root of the sum of squares of subjective cue weights (middle row), and square root of the sum of squares of subjective spatial weights (bottom row). Error bars are standard deviations. * p < .0033.

but showed, as well, that satisficing is only one of several factors influencing the behavior.

A limitation of our approach was that our visual cue dimensions were selected to be readily distinguishable but not to be of equal salience. It is possible that some cues, for example, orientation or border, were less visually salient to the macaques than others, for example, color or shape. These differences might have influenced performance. We mitigated this potential confound, however, by counterbalancing all of the visual stimulation between the two monkeys. It appears that this worked well because nearly all results were similar between the monkeys, indicating that they learned and used the probabilistic information in the cues despite any differences in salience. Specifically, even if some cues were less salient than others, both monkeys still sorted out the rankings of each cue dimension in both the LP condition (Figures 5, 7, 8, 9, and 10, gray data) and HP condition (Figures 5, 8, 9, and 10, white data). The only major difference between the monkeys was the use of a "drop-the-worst" satisficing strategy in Monkey A but not in Monkey L (Figure 7). It is possible that this difference was facilitated for Monkey A by assignment of the lowest probabilistic information to a cue (orientation) that may have been less salient than the others. We did not attempt to equalize salience between

the visual dimensions out of concern that we might influence the monkeys' judgments about the cues through repeated, preliminary testing to titrate salience levels, which might have required more exposure to some cues than others. Another concern was how often (or whether) we would need to interrupt data collection on the task to retest and potentially recalibrate salience levels. We therefore do not know, quantitatively, the extent to which visual salience contributed to the difference in decision strategies observed in HP conditions. It is not unusual for two monkeys in a study to have different levels of performance and apparently different strategies in a highly demanding task for reasons that are not definitively resolved (Couchman et al., 2010; Shields, Smith, Guttmannova, & Washburn, 2005). We do see it as advisable in follow-up work, however, to equalize salience across cue dimensions or to vary salience systematically as a trade-off with probabilistic information to understand how the two factors interact in macaque decision-making.

In the human literature, probabilistic inference has been studied with nuanced categorization paradigms, such as the weatherprediction task, that include manipulation of task statistics (Gluck & Bower, 1988; Knowlton et al., 1994). Studies on humans using the present task showed that they infer the probabilistic values of the different visual cues (Oh et al., 2016) and that prefrontal, parietal, and subcortical brain regions encode the inferential operations (Oh-Descher et al., 2017). Comparatively, our macaques exploited stimulus probabilities as humans did, but more slowly, requiring multiple sessions of training to learn the cue-probability associations (Figure 3). After they learned the associations, we did not manipulate task statistics because the goal was to achieve stable, not malleable, probabilistic decision-making to isolate the influence of time pressure. Paradigms more akin to classical probabilistic categorization, such as a weather-prediction task modified for macaques (Yang & Shadlen, 2007), were used previously to focus on probabilistic reasoning itself, its flexibility, and its neural correlates (Kira et al., 2015). Yang and Shadlen (2007, p. 1079) concluded that their results "... demonstrated a crude capacity for probabilistic inference in monkeys." Our main purpose was to build on this previous conclusion, using a more naturalistic task in which macaques simply chose between visual stimuli, as a way of studying the animals' satisficing strategies. Given that we did not manipulate task statistics, the specific type of probabilistic inference we documented might be more precisely termed "probabilistic preference," in that it was used to select the preferred stimulus (in terms of likelihood of reward). This alternative term, however, is already in use for other purposes in the literature (Vitelli, Sørensen, Crispino, Frigessi, & Arjas, 2018). All of these considerations should be kept in mind when comparing the present work to previous studies of probabilistic reasoning in macaques and humans.

Decision Noise and Spatial Bias

One of the nonsatisficing factors that influenced behavior under time pressure, decision noise, is well known to affect the selection of visual stimuli in macaques (reviewed in Gold & Shadlen, 2007), although it had not been evaluated previously in the conditions of our study, namely, probabilistic decision-making across multiple visual dimensions. Time pressure caused approximately a doubling of decision noise for both monkeys (halving of SRSS of subjective cue weights) across many response times (Figure 10, middle). For one monkey, a significant, smaller effect was evident even in its overall pooled data (Figure 8, right). The other nonsatisficing factor that was modulated by time pressure, spatial bias, manifested as an increased reliance on favored spatial locations. The specifics of this spatial bias may be partially explained by the in-cage experimental setup. The stimulus locations for each monkey's biases roughly correspond to both the handedness and height of each monkey. Decisions may have been made based on the effort required to make the appropriate reach, with further reaches judged as less valuable (Morel, Ulbrich, & Gail, 2017). Indeed, our behavioral results provide evidence that the monkeys applied a subjective value to the spatial location of onscreen stimuli that became more important to them under time constraints. Moreover, the time pressures used in this study may approach the physiological limit for a macaque to process the visual information and physically reach toward the stimulus. At high speeds, the use of a priori biases becomes especially important when executing a quick reach (Wolpert & Landy, 2012). Further studies may address this by either continuing to use touch responses but including variations of the spatial presentation of stimuli, or using eye movements to report decisions.

Satisficing

The "drop-the-worst" strategy observed in Monkey A is consistent with satisficing exhibited in the same task in humans (Oh et al., 2016). By selectively discounting less valuable information, the cognitive load for decisions is lowered. In addition, both monkeys showed a "group-the-best" strategy, reducing or eliminating the difference in subjective weights attributed to the top two stimulus dimensions under time pressure. Similar to a "take-thebest" strategy, the end result is to save time by focusing on the dimensions of information that are most valuable. Unlike a "takethe-best" strategy, however, this "group-the-best" strategy serves to categorize the two best dimensions together, evaluate this new category together with the lower two dimensions, and thus effect a form of dimensionality reduction. Grouping of the two best dimensions was manifested as a steady behavior exhibited across trials as the sessions unfolded. Human subjects are known to improve on "take-the-best" strategies by exploiting confirming evidence within trials (Karelaia, 2006; Newell et al., 2003). In our task, this would be revealed as selection based on a stimulus's top two cue dimensions (Strategy 5/21, without/with spatial regressors, in Figure 2), but we found no evidence for such a strategy in our monkeys (Figure 7).

The differences in type and degree of satisficing strategies between Monkeys A and L may suggest that satisficing in macaques may be more idiosyncratic than in humans, but such a conclusion has two caveats. First, our sample size (n = 2 monkeys) was much smaller than that tested in studies of human satisficing, for example, Oh et al. (2016; 46 humans). Second, humans also exhibit considerable intersubject variability in their satisficing behavior. On an individual basis, approximately 32% of participants studied by Oh et al. (2016) showed the "drop-theworst" strategy under time pressure, whereas 26% of individuals continued to use the full cue integration strategy and thus failed to satisfice at all (author Oh-Descher, personal observation). It is possible, of course, that the latter group may have learned to satisfice in the task if given more than one session to perform it. Individual use of a "group-the-best" strategy was not examined in humans (Oh et al., 2016). Satisficing is therefore variable in both species, and comparisons between them need to consider the strengths and limitations inherent to the study designs used for each species (e.g., few subjects in macaque studies and few testing sessions per subject in human studies).

Potential Neural Mechanisms

Studies of the neural basis of probabilistic inference and uncertain decision-making in humans have used functional MRI to identify several areas of interest (Behrens, Woolrich, Walton, & Rushworth, 2007; Gold & Shadlen, 2007; Knowlton, Mangels, & Squire, 1996; Oh-Descher et al., 2017; Poldrack et al., 2001; Shohamy, Myers, Kalanithi, & Gluck, 2008). However, these data are limited by both the temporal and spatial resolution of the blood-oxygen-level dependent signal. Animal models and invasive techniques allow for a finer grained analysis of the underlying neural mechanisms. Research into neural correlates of decisionmaking in macaques is extensive and has revealed that the subjective value of decisions can be tracked at the single neuron level in the lateral intraparietal area (Platt & Glimcher, 1999; Roitman & Shadlen, 2002; Rorie, Gao, McClelland, & Newsome, 2010) and prefrontal cortex (Kim & Shadlen, 1999). Furthermore, lateral intraparietal area neurons are able track the accumulation of evidence across sequentially presented visual cues (Kira et al., 2015; Yang & Shadlen, 2007). The present work demonstrates the feasibility of extending this research to neural studies of probabilistic visual discrimination under LP and HP. An important caveat, however, is that not all macaques react to high time pressure with the human "drop-the-worst" strategy. If the aim is to test that form of satisficing, monkeys should be prescreened to find those that exhibit it. Comparing neural activity between satisficing strategies would be informative too, and macaques appear to provide the opportunity to compare at least the "drop-the-worst" and "group-the-best" strategies, sometimes in the same subject (such as Monkey A in our study).

Conclusion

The present work explored the psychophysics of making decisions between objects based on probabilistic visual cues in rhesus macaques. We found similarities between macaque behavior and previously reported human behavior, in that both species can perform probabilistic discrimination between visual stimuli and can satisfice under time pressure, in some cases even using the same strategy ("drop-the-worst"). We found differences as well, likely reflecting a mix of real interspecies differences. Future work that compares macaques and humans with even more closely matched experiments, complemented by computational models that probe and compare the underlying algorithms used by each species, would help to address unresolved issues. Overall, our results provide support for using macaques to study the neural basis of probabilistic decision-making and satisficing under time pressure.

References

- Amemori, K., & Graybiel, A. M. (2012). Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nature Neuroscience*, 15, 776–785. http://dx.doi.org/10.1038/nn.3088
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10, 1214–1221. http://dx.doi.org/10.1038/nn1954
- Couchman, J. J., Coutinho, M. V., & Smith, J. D. (2010). Rules and resemblance: Their changing balance in the category learning of humans (*Homo sapiens*) and monkeys (*Macaca mulatta*). Journal of Experimental Psychology: Animal Behavior Processes, 36, 172–183. http://dx.doi .org/10.1037/a0016748
- Drugowitsch, J. (2013). Variational Bayesian inference for linear and logistic regression. Retrieved from https://arxiv.org/abs/1310.5438
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: From behavior to neural representations. *Trends* in Cognitive Sciences, 14, 119–130. http://dx.doi.org/10.1016/j.tics .2010.01.003
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. Annual Review of Psychology, 62, 451–482. http://dx.doi.org/10.1146/ annurev-psych-120709-145346
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650– 669. http://dx.doi.org/10.1037/0033-295X.103.4.650
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychol*ogy: General, 117, 227–247. http://dx.doi.org/10.1037/0096-3445.117.3 .227

- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574. http://dx.doi.org/10.1146/ annurev.neuro.29.051605.113038
- Hayden, B. Y., Heilbronner, S. R., & Platt, M. L. (2010). Ambiguity aversion in rhesus macaques. *Frontiers in Neuroscience*, 4, 166. http:// dx.doi.org/10.3389/fnins.2010.00166
- Karelaia, N. (2006). Thirst for confirmation in multi-attribute choice: Does search for consistency impair decision performance? *Organizational Behavior and Human Decision Processes*, 100, 128–143. http://dx.doi .org/10.1016/j.obhdp.2005.09.003
- Kepecs, A., & Mainen, Z. F. (2014). A computational framework for the study of confidence across species. In S. Fleming & C. Frith (Eds.), *The cognitive neuroscience of metacognition* (pp. 115–145). Berlin, Heidelberg: Springer. http://dx.doi.org/10.1007/978-3-642-45190-4_6
- Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 324, 759–764. http://dx.doi.org/10.1126/science.1169405
- Kim, J. N., & Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature Neuroscience*, 2, 176–185. http://dx.doi.org/10.1038/5739
- Kira, S., Yang, T., & Shadlen, M. N. (2015). A neural implementation of Wald's sequential probability ratio test. *Neuron*, 85, 861–873. http://dx .doi.org/10.1016/j.neuron.2015.01.007
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, 273, 1399–1402. http://dx.doi.org/ 10.1126/science.273.5280.1399
- Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning and Memory*, 1, 106–120.
- Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in Cognitive Sciences*, 16, 511–518. http://dx.doi.org/10.1016/j.tics.2012 .08.010
- Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family resemblance, conceptual cohesiveness, and category construction. *Cognitive Psychology*, 19, 242–279. http://dx.doi.org/10.1016/0010-0285(87)90012-0
- Morel, P., Ulbrich, P., & Gail, A. (2017). What makes a reach movement effortful? Physical effort discounting supports common minimization principles in decision making and motor control. *PLoS Biology*, 15, e2001323. http://dx.doi.org/10.1371/journal.pbio.2001323
- Nelson, D. G. K. (1984). The effect of intention on what concepts are acquired. *Journal of Verbal Learning and Verbal Behavior*, 23, 734– 759. http://dx.doi.org/10.1016/S0022-5371(84)90442-0
- Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone "takes-the-best". Organizational Behavior and Human Decision Processes, 91, 82–96. http://dx.doi.org/ 10.1016/S0749-5978(02)00525-3
- Oh, H., Beck, J. M., Zhu, P., Sommer, M. A., Ferrari, S., & Egner, T. (2016). Satisficing in split-second decision making is characterized by strategic cue discounting. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 42*, 1937–1956. http://dx.doi.org/10.1037/ xlm0000284
- Oh-Descher, H., Beck, J. M., Ferrari, S., Sommer, M. A., & Egner, T. (2017). Probabilistic inference under time pressure leads to a corticalto-subcortical shift in decision evidence integration. *NeuroImage*, *162*, 138–150. http://dx.doi.org/10.1016/j.neuroimage.2017.08.069
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400, 233–238. http://dx.doi.org/10 .1038/22268
- Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, 414, 546–550. http://dx.doi.org/10 .1038/35107080
- Raghuraman, A. P., & Padoa-Schioppa, C. (2014). Integration of multiple determinants in the neuronal computation of economic values. *The*

Journal of Neuroscience, 34, 11583–11603. http://dx.doi.org/10.1523/ JNEUROSCI.1235-14.2014

- Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. Acta Psychologica, 127, 258–276. http://dx.doi.org/10.1016/j.actpsy.2007.05.004
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies - revisited. *NeuroImage*, 84, 971–985. http://dx.doi.org/10.1016/j.neuroimage.2013.08.065
- Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, 22, 9475–9489. http://dx .doi.org/10.1523/JNEUROSCI.22-21-09475.2002
- Rorie, A. E., Gao, J., McClelland, J. L., & Newsome, W. T. (2010). Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS ONE*, *5*, e9308. http://dx.doi.org/10.1371/journal.pone .0009308
- Shields, W. E., Smith, J. D., Guttmannova, K., & Washburn, D. A. (2005). Confidence judgments by humans and rhesus monkeys. *Journal of General Psychology*, 132, 165–186.
- Shohamy, D., Myers, C. E., Kalanithi, J., & Gluck, M. A. (2008). Basal ganglia and dopamine contributions to probabilistic category learning. *Neuroscience and Biobehavioral Reviews*, 32, 219–236. http://dx.doi .org/10.1016/j.neubiorev.2007.07.008
- Sigala, N., & Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature*, 415, 318–320. http://dx.doi.org/10.1038/415318a
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118. http://dx.doi.org/10.2307/1884852

- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63, 129–138. http://dx.doi.org/10.1037/ h0042769
- Smith, J. D., Beran, M. J., Crossley, M. J., Boomer, J., & Ashby, F. G. (2010). Implicit and explicit category learning by macaques (*Macaca mulatta*) and humans (*Homo sapiens*). Journal of Experimental Psychology: Animal Behavior Processes, 36, 54–65. http://dx.doi.org/10.1037/ a0015892
- Smith, J. D., Berg, M. E., Cook, R. G., Murphy, M. S., Crossley, M. J., Boomer, J., . . . Grace, R. C. (2012). Implicit and explicit categorization: A tale of four species. *Neuroscience and Biobehavioral Reviews*, *36*, 2355–2369. http://dx.doi.org/10.1016/j.neubiorev.2012.09.003
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, 46, 1004–1017. http://dx.doi.org/10.1016/j.neuroimage.2009.03.025
- Vitelli, V., Sørensen, Ø., Crispino, M., Frigessi, A., & Arjas, E. (2018). Probabilistic preference learning with the Mallows rank model. *Journal* of Machine Learning Research, 18, 1–49.
- Vulkan, N. (2000). An economist's perspective on probability matching. Journal of Economic Surveys, 14, 101–118. http://dx.doi.org/10.1111/ 1467-6419.00106
- Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception and Psychophysics*, 63, 1293–1313. http://dx.doi.org/10.3758/BF03194544
- Wolpert, D. M., & Landy, M. S. (2012). Motor Control is decision-making. *Current Opinion in Neurobiology*, 22, 996–1003. http://dx.doi.org/10 .1016/j.conb.2012.05.003
- Yang, T., & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, 447, 1075–1080. http://dx.doi.org/10.1038/nature05852

395

Appendix 1

Details of the Bayesian Logistic Regression

Regressors for the effect of cue dimension net weights were coded into a T × 4 matrix, X_C , with columns for each dimension (1 through 4; color, shape, orientation, and border) and rows for each trial (up to T total trials per session). If $W_{Ai,j}$ and $W_{Bi,j}$ are the assigned weights of the cue state on trial *i* in cue dimension *j* for Stimuli A and B, respectively, then the elements of X_C are defined as follows: $X_{Ci,j} = 1$ if $W_{Ai,j} > W_{Bi,j}$, $X_{Ci,j} = -1$ if $W_{Ai,j} < W_{Bi,j}$, and $X_{Ci,j} = 0$ if $W_{Ai,j} = W_{Bi,j}$. The effect of spatial location was entered into the regression by assigning each of the six location pairs a unique row vector in a T × 6 matrix X_S . The row vectors contained five zeros and a single entry of '1' in a column unique to the stimulus location pair. This coding scheme allows location pairs to be treated as discrete categorical variables, while returning continuous numerical coefficients indicating the degree of spatial bias through regression.

Logistic regression was performed for each experimental session, fitting cue states and spatial locations in matrix $X = [X_c, X_s]$ to a T × 1 response vector. Each response was 1 if stimulus A was selected or -1 if Stimulus B was selected. The regression returned model parameters by fitting the following hierarchical model (Drugowitsch, 2013):

$$P(select A \mid X, w_{out}) = \frac{1}{1 + e^{-(Xw_{out})}}$$
(2)

$$P(w_{out} \mid \alpha) = \mathcal{N}(w_{out} \mid 0, a^{-1}I)$$
(3)

$$P(\alpha) = Gamma(\alpha | a_0, b_0)$$
(4)

where *I* is the 10 × 10 identity matrix. The regression returned estimated coefficients in 10 × 1 vectors w_{out} , and precisions in V_n^{-1} . The elements of w_{out} and V_n^{-1} each correspond to cue and spatial dimensions in the columns of X_C and X_S . The values of the hyper parameters a_0 and b_0 were set in accordance with an empirical Bayesian approach. First, w_{out} was calculated for each session using relatively uninformative hyperparameters ($a_0 = 0.01$, $b_0 = 0.0001$). Then all regressions were run again with $a_0 = 1$ and the value of b_0 estimated as the variance of w_{out} across all sessions individually for each monkey ($b_0 = 0.729$ for Monkey A and $b_0 = 0.382$ for Monkey L). This strategy selected an exponential conjugate prior on α with the appropriate mean.

To calculate subjective weights, each w_{out} was converted to exponential base 10 in vector w_n^* , (for experimental session *n*) to

remain consistent with the base used on each trial to calculate reward probability:

$$w_n^* = \log_{10} e^{w_{out}} \tag{5}$$

Subjective weights were then divided into sets corresponding to the four cue dimensions (*subjective cue weights* $w_{n,C}^*$) and six spatial location pairs (*subjective spatial weights* $w_{n,S}^*$). For each experimental session, the *square root of the sum of squares* (SRSS) of subjective cue weights and subjective spatial weights were calculated, respectively, as follows:

$$SRSS_{n,C} = |w_{n,C}^*| = \sqrt{\sum_{i=1}^4 w_{n,C_i}^*}$$
(6)

$$SRSS_{n,S} = |w_{n,S}^*| = \sqrt{\sum_{j=1}^{6} w_{n,S_j}^*}^2$$
(7)

To compare relative changes in subjective cue weights between HP and LP conditions, the weights in each session n and cue dimension i were normalized by the mean SRSS across all sessions both conditions as follows:

$$\hat{w}_{n,C_i} = \frac{w_{n,C_i}}{\frac{1}{N}\sum_{i} SRSS_{n,C}}$$
(8)

where *N* is the number of experimental sessions in the given condition. This normalization controls for the absolute magnitude of the subjective weights to isolate the relative importance between the weights. Finally, to account for variability in precision (V_{n,C_i}^{-1}) , the weighted mean of each normalized subjective cue weight, \bar{w}_{c_i} and its corresponding uncertainty $\sigma_{\bar{w}_c}$ were calculated as follows:

$$\overline{w}_{C_i} = \frac{\sum_{n}^{n} V_{n,C_i}^{-1} \hat{w}_{n,C_i}}{\sum_{n}^{n} V_{n,C_i}^{-1}}$$
(9)

$$\sigma_{\overline{w}_{C_i}} = \sqrt{\frac{1}{\sum_n V_{n,C_i}^{-1}}}$$
(10)

Because all the regressors were binary, \bar{w}_{c_i} provides a measure of the relative subjective importance of each cue dimension and spatial location pair.

(Appendices continue)

TOADER ET AL.

Appendix 2

Numerical Results of the Response Time Analyses

Table A2Data From Figure 10

Monkey A					Monkey L					
Low p	ressure	High p	pressure		Low pr	ressure	High p	ressure		
М	SD	M	SD	Cohen's d	М	SD	М	SD	Cohen's d	
				Fraction	Correct					
0.880	.015	0.731	.021	-8.204	.912	.013	0.676	.022	-14.626	
0.865	.016	0.755	.020	-6.189	.899	.014	0.693	.022	-12.245	
0.841	.017	0.759	.020	-4.478	.896	.014	0.756	.012	-8.616	
0.906	.014	0.748	.020	-9.295	.888	.015	0.722	.021	-9.803	
0.856	.016	0.782	.019	-4.203	.905	.014	0.738	.021	-10.406	
0.878	.015	0.762	.020	-6.517	.914	.013	0.730	.021	-11.694	
0.895	.014	0.758	.020	-7.944	.884	.015	0.772	.019	-6.860	
0.881	.015	0.775	.019	-6.115	.887	.015	0.779	.019	-6.725	
0.860	.016	0.802	.018	-3.340	.878	.015	0.785	.019	-5.622	
0.863	.016	0.810	.018	-3.139	.878	.015	0.778	.019	-5.972	
0.868	.016	0.778	.019	-5.129	.866	.016	0.799	.019	-3.984	
0.883	.015	0.832	.017	-3.181	.876	.015	0.791	.019	-5.111	
0.883	.015	0.826	.018	-3.464	.872	.016	0.800	.019	-4.380	
0.843	.017	0.840	.017	143	.853	.017	0.803	.018	-2.888	
0.852	.016	0.817	.018	-2.003	.869	.016	0.815	.018	-3.254	
			Square roo	ot of the sum of squ	ares of subjectiv	ve cue weights				
2.463	.282	1.334	.142	-5.066	2.007	.223	1.299	.169	-3.404	
2.057	.224	1.361	.147	-3.684	2.041	.182	0.916	.105	-6.922	
2.440	.226	1.304	.118	-6.328	1.976	.181	1.345	.134	-3.758	
2.652	.287	1.336	.129	-5.930	2.006	.194	1.102	.109	-5.246	
1.932	.185	1.643	.173	-1.618	2.511	.260	1.208	.121	-5.755	
2.672	.261	1.278	.118	-6.915	2.453	.255	1.041	.117	-6.371	
2.144	.216	1.324	.133	-4.586	2.040	.201	1.102	.104	-5.309	
2.249	.207	1.280	.121	-5.728	2.044	.189	1.230	.107	-4.856	
2.385	.220	1.577	.168	-4.143	2.062	.191	1.275	.115	-4.598	
2.178	.204	1.291	.126	-5.245	1.910	.172	1.083	.117	-5.265	
2.159	.188	1.265	.115	-5.755	1.963	.181	1.229	.119	-4.468	
2.339	.246	1.596	.142	-3.704	2.092	.191	1.432	.128	-3.785	
2.562	.295	1.461	.142	-4.772	2.029	.190	1.429	.126	-3.466	
2.187	.202	1.453	.149	-4.142	1.818	.176	1.255	.117	-3.516	
1.857	.179	1.364	.148	-2.997	1.788	.181	1.316	.126	-2.848	
			Square root	of the sum of squar	res of subjective	e spatial weigh	ts			
1.707	.248	2.583	.211	3.801	1.311	.225	3.636	.404	7.983	
1.498	.249	2.654	.277	4.389	0.972	.194	2.101	.219	5.581	
0.893	.189	1.942	.201	5.368	0.783	.170	2.633	.256	9.222	
0.837	.213	1.858	.199	4.957	0.598	.160	1.977	.201	7.945	
0.587	.131	2.328	.229	9.298	0.962	.200	2.100	.235	5.374	
1.012	.228	1.524	.172	2.533	0.749	.177	1.397	.170	3.711	
0.671	.148	1.720	.178	6.405	0.650	.154	0.985	.136	2.253	
0.709	.151	1.254	.153	3.589	0.560	.140	1.159	.144	4.252	
1.341	.190	0.825	.165	-2.900	0.602	.160	1.088	.152	3.079	
1.171	.179	0.983	.158	-1.111	0.598	.147	1.006	.149	2.765	
0.911	.159	0.809	.129	709	0.707	.145	1.000	.153	1.991	
1.122	.179	0.778	.147	-2.107	0.660	.145	1.272	.172	3.975	
1.439	.222	0.648	.134	-4.319	0.748	.151	1.271	.171	3.316	
1.439	.179	0.691	.141	-4.730	0.635	.148	1.087	.152	3.027	
0.964	.156	0.623	.138	-2.324	0.673	.143	1.025	.152	2.416	

Note. For each monkey and measure (e.g. fraction correct), entries are means, standard deviations, and effect sizes (Cohen's *d*). Rows in each section correspond to response time bins along the abscissas of the respective panels of Figure 10, from shorter (upper rows) to longer (lower rows) response times.

Received January 31, 2018

Revision received December 13, 2018

Accepted December 13, 2018