Emotion
2009, Vol. 9, No. 5, 599–608
© 2009 American Psychological Association
1528-3542/09/$12.00 DOI: 10.1037/a0016509

Integration of Social and Utilitarian Factors in Decision Making

Bruno B. Averbeck and Brad Duchaine
University London College

Decision making is influenced by social cues, but there is little understanding of how social information interacts with other cues that determine decisions. To address this quantitatively, participants were asked to learn which of two faces was associated with a higher probability of reward. They were repeatedly presented with two faces, each with a different, unknown probability of reward, and participants attempted to maximize gains by selecting the face that was most often rewarded. Both faces had the same identity, but one face had a happy expression and the other had either an angry or a sad expression. Ideal observer models predict that the facial expressions should not affect the decision-making process. Our results however showed that participants had a prior disposition to select the happy face when it was paired with the angry but not the sad face and overweighted the positive outcomes associated with happy faces and underweighted positive outcomes associated with either angry or sad faces. Nevertheless, participants also integrated the feedback information. As such, their decisions were a composite of social and utilitarian factors.

Keywords: ideal observer, decision making, social stimuli, fares

Decision making is often studied from the perspective of normative models that attempt to maximize utility. From this perspective decision making, as well as other activities requiring inference in the face of uncertainty, should be carried out in a statistically optimal way. For example, Bayes optimal theories have proven effective at predicting low-level behaviors, including sensory cue integration (Ernst & Banks, 2002; Jacobs, 1999; Knill & Saunders, 2003), motor control (Kording & Wolpert, 2006; Todorov & Jordan, 2002), and visual inference (Kersten, 1999; Knill & Richards, 1996; Poggio, Torre, & Koch, 1985), and models and data have shown that the brain can effectively represent and integrate uncertainty in these situations (Gold & Shadlen, 2001; Ma, Beck, Latham, & Pouget, 2006).

While these models can be effective at predicting behavior in simple sensorimotor situations, another line of research has also shown pervasive, systematic deviations of behavior from the predictions of normative models in higher level decision making (Gilovich, Griffin, & Kahneman, 2002; Kahneman, Slovic, & Tversky, 1982; Payne, Bettman, & Johnson, 1992; Tversky & Kahneman, 1986, but see Sher & McKenzie, 2008). For example, in the framing effect, participants often reverse their preference between a pair of decisions depending upon whether the options are described in terms of gains or losses even though the probabilities and outcomes are identical. In a specific example, when participants were asked to decide between treatment options to control the outbreak of a disease, their decision depended upon whether the options were stated in terms of the probability of people living (200 out of 600) or dying (400 out of 600) (Tversky & Kahneman, 1981, 1986). In related experiments, these effects were not eliminated by training and were similar in physicians and students (McNeil, Pauker, Sox, & Tversky, 1982). Other examples of deviations from normative models include ignoring prior information (base rate neglect; Bar-Hillel, 1990; Kahneman & Tversky, 1973), predicting that the combination of two events is more probable than either individual event (conjunction biases; Thuring & Jungermann, 1990; Tversky & Kahneman, 1983), and aversion to ambiguity (Ellsberg, 1961; Hsu, Bhatt, Adolphs, Tranel, & Camerer, 2005).

When behavior does deviate from the predictions of optimal models, it is often considered irrational (Gilovich et al., 2002; Sutherland, 1992). However, effective real-world decision making requires the integration of many sources of information including cues from the social context, and the construction of ideal observers in these situations is not always possible or meaningful. For example, how do we decide whether or not to trust a partner in a single-shot economic exchange (Scharlemann, Eckel, Kacelnik, & Wilson, 2001; Singer, Kiebel, Winston, Dolan, & Frith, 2004)? While such decisions are routinely made, little is known about which cues affect these decisions and how they interact with other decision-making factors. A few studies have examined the question of whether social factors affect decision making. These studies have shown that people were more likely to cooperate with smiling partners (Scharlemann et al., 2001), and thirsty participants shown a masked happy face drank more of a beverage and indicated a greater willingness to pay for the beverage when compared to the same condition in which they were shown an angry face (Winkielman, Berridge, & Wilbarger, 2005). In general, however, these studies have not quantified how social cues interact with utilitarian sources of information.

Bruno B. Averbeck, Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University London College, and Brad Duchaine, Institute of Cognitive Neuroscience, University London College.

This research was funded by a grant to B. D. from the ESRC RES-061-23-0051.

We thank the social club at the Institute of Cognitive Neuroscience for useful feedback on the project, as well as Laura Germine, Zara-Angela Abbas, and Simon Evans for help in collecting the data.

Correspondence concerning this article should be addressed to Bruno B. Averbeck, PhD, UCL Institute of Neurology, Sobell Department, Box 28, Queen Square, London WC1N 3BG. E-mail: b.averbeck@ion.ucl.ac.uk

To quantitatively characterize the effect of an important social cue, facial expression, we carried out a task that ostensibly had an optimal solution—deciding which of two faces within a block of trials was associated with a higher probability of reward. Both faces had the same identity, but one was smiling and the other was either angry (Experiment 1) or sad (Experiment 2). The participants' task was to pick the face that maximized their rewards. The task was made difficult, however, by the fact that one of the faces was associated with a positive outcome 40% of the time, and the other face with a positive outcome 60% of the time. In the results we examine the effect of the emotional expression on the participants' decision-making processes. Importantly, our model allows us to quantify the effect of the emotional cue (facial expression) relative to a utilitarian cue (the feedback). Furthermore, we can examine whether the effect of the social cue manifests only as a prior bias, which can be overcome by sufficient evidence, or whether the social cue also has an effect on how the evidence or reward feedback is accumulated over time. Our model, however, is not a proscriptive model that a priori assumes, for example, ecological behavior. Rather it is a descriptive model that gives us a tool for testing the hypothesis that the expressions influenced the subjects' decision-making processes. Finally, while we have focused on the effect of facial expression, we expect that other cues without social content could have a similar effect.

Methods

Task and Participants

Twenty-four participants (12 in Experiment 1: happy vs. angry, 12 in Experiment 2 happy vs. sad) performed a two-alternative forced choice decision-making task. We first collected data in Experiment 1 and then collected data on a separate set of participants in Experiment 2. Each subject did four blocks of 26 trials. In each block, there were two faces with the same identity but different expressions. In Experiment 1, one face had a happy expression and the other had an angry expression, and in Experiment 2 the expressions were happy and sad. Two different identities were also used in each experiment. The identities alternated across blocks, and we counterbalanced whether identity A or B was shown in the first block. The same identities were used in Experiments 1 and 2. In each block, one of the faces paid off 40% of the time when selected and the other paid off 60% of the time. Participants were instructed to make decisions to maximize their rewards. The order of high reward versus low reward associated with the happy and angry/sad face was also counterbalanced as much as possible across participants, such that some participants started with the happy face being rewarded most often, followed by the angry/sad face being rewarded most often, whereas other participants had the opposite order.

We chose 26 trials because an ideal observer is able to identify the most highly rewarded face correctly in 85% of blocks of this length with the probability values we used. Thus, there is sufficient evidence in most blocks to identify the correct face. This point is not actually relevant to our analysis, however, because we modeled the subjects' belief trial by trial, and therefore we do not rely on whether or not the subjects know which face "should" be correct in each block. Thus, if for some reason the sequence of rewards favors the face being rewarded stochastically less often, our analysis takes that into account.

Every subject was given the following instructions on the task: "On each trial in this task you will be presented with two faces. You will have to select one of the faces. Press 'z' to select the left face, '/' to select the right face. Your task is to try to figure out which face in each block has the highest probability of winning and pick that face as many times as possible. You will be told when the block switches, and at each switch the faces will be associated with new probabilities of winning."

Within an individual trial, the happy and angry faces were presented pseudorandomly on either the left or the right side of the screen (see Figure 1). Participants were given an unlimited duration to make their decision, and the faces were present until the participants responded. After the participants made their decision, they pressed one of two buttons to indicate whether they had chosen the left or the right face. The reaction time was the time from presentation of the faces until the response key was pressed. The chosen face was then presented at the center of the screen, and below it was text indicating whether they had "won" or "lost" in that trial, with a win worth 10 pence, and a loss worth nothing. A 5-kHz tone was played when they won, and a 2.5-kHz tone was played when they lost. In practice, all participants were paid the same amount, which was greater than their actual winnings, but they were not informed of this until after the experiment had ended.

Data Analysis

All data analysis was carried out in Matlab. Because the actual outcomes in the experiment were stochastic, it was possible for the face that had a lower probability of being rewarded in an individual block to actually be rewarded more often, especially over a short run of trials. Therefore we referenced all decisions of the participants to an ideal observer model in the first series of analyses. The ideal observer makes decisions using the reward feedback information, just as the subject. However, the ideal observer provides us with the best estimate of which face to pick, based upon the data, whereas the subject won't necessarily do so. In

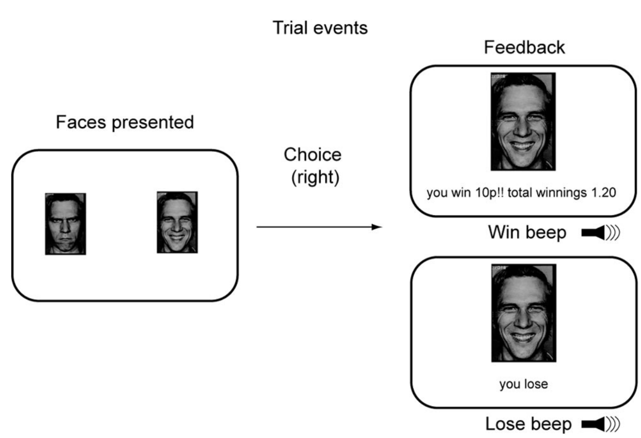

Figure 1. Trial events. Faces were presented on the left and right of screen. Participants then selected either the right or left face, following which they were given auditory and visual feedback about whether they had "won" or "lost." If they won, their total winnings were incremented 10 pence. Reprinted with permission from owner Paul Eckman.

subsequent analyses, we added parameters to the ideal observer so that it better modeled the decision-making processes of the participants. These parameters allowed us to test directly for particular stimulus effects on subject behavior. The ideal observer was modeled using a Bernoulli distribution for each of the two faces. The Bernoulli distribution simply allows us to keep track of the number of rewards the subject has received as a function of the number of times the subjects have picked each face, similar to keeping track of the number of times a coin toss has come up heads. While the ideal observer's estimates of which face are best are optimal given the feedback received, it is operating on the outcomes of the decisions of the subject, which might not be optimal. The likelihood function of the ideal observer was given by:

$$p(D|\theta_i) = \theta_i^{r_i}(1 - \theta_i)^{N_i - r_i} \quad (1)$$

Where θ_i is the probability that face i (either angry or happy) is rewarded, r_i is the number of times face i was rewarded, and N_i is the number of times face i was selected. The vector D represents the data, which in this case are the values of r and N. The probability that face i was more often rewarded than face j was given by:

$$p(\theta_i > \theta_j) = \int_0^1 p(\theta_i|D) \int_0^{\theta_i} p(\theta_j|D)\,d\theta_j\,d\theta_i. \quad (2)$$

We have here used the posterior, as we numerically normalized the distributions before carrying out the integral and assumed a flat prior on θ_i. The ideal choice or decision rule ($\hat{f}$) was then given by the face which was probably most highly rewarded, given by:

$$\begin{cases} p(\theta_i > \theta_j) > 0.5 & \hat{f} = i \\ p(\theta_i > \theta_j) < 0.5 & \hat{f} = j \end{cases} \quad (3)$$

In other words, the ideal choice was the face that would most likely be rewarded on the current trial. For the data analysis, when probabilities were tied each face was given half of a reward, which indicates that no information was gained.

This model makes several assumptions. Specifically, it assumes that the reward probabilities can take any values between 0 and 1, and that the probabilities for the two faces are independent. Both of these assumptions are consistent with the information given to the participants and the distributions from which the data was sampled during the experiment.

Ecological Observer

The ecological observer model contained three extra parameters. First, to model the limited memory capacity of the participants, we included a term that exponentially weighted past outcomes, with more recent outcomes being more heavily weighted. Second, we included a term that allowed for a differential weighting of the rewards for the happy and angry faces. This term allowed us to directly test the hypothesis that outcomes related to happy and angry or sad faces were differentially weighted. Finally, we also included a prior term that incorporated a prior disposition toward one of the faces. The exponential weighting was used to calculate a modified reward history, as

$$r_i^{\exp}(t) = \sum_{k=0}^{N-1} e^{-ak} r_i(t-k), \quad (4)$$

and correspondingly:

$$N^{\exp} = \sum_{k=0}^{N-1} e^{-ak}. \quad (5)$$

The variable a was a free parameter in the model. Next, the likelihood function was modified to include asymmetric reward effects. This was done by using an asymmetric reward outcome:

$$\begin{aligned} r_{happy}(t) &= r(t)(1+c) \\ r_{angry}(t) &= r(t)(1-c) \end{aligned} \quad (6)$$

Again, the variable c was fit as a free parameter that could be positive or negative, and as such, positive feedback associated with the happy face could be either over or under valued. The reward value, $r(t)$, entered in the analysis was 1 if the subject received a positive outcome for their decision (You Win!), and 0 if the subject received a negative outcome (You Lose). Thus, if the subject was rewarded, the model was updated by $1 \pm c$ depending upon the face, and if the subject was not rewarded, it was updated by 0. Finally, a Beta prior was used to model the prior disposition toward each face as:

$$p(\theta_i|\alpha_i,\beta_i) \propto \theta_i^{\alpha_i - 1}(1-\theta_i)^{\beta_i - 1}. \quad (7)$$

We constrained α_i and β_i to get a good model fit by reducing the four parameters to one degree of freedom. This was done by estimating a single parameter, b, and then computing α and β as:

$$\begin{aligned} \alpha_{happy} &= N_{prior}(0.5+b) \\ \alpha_{angry} &= N_{prior}(0.5-b) \\ \beta_{happy} &= N_{prior}(0.5-b) \\ \beta_{angry} &= N_{prior}(0.5+b) \end{aligned} \quad (8)$$

We set N_{prior} to 4. This set the strength of the prior and is equivalent to assuming participants had four trials prior experience with each face, with the number of positive and negative outcomes given by α and β respectively. Allowing it to float freely resulted in unrealistic values for b (i.e., values that were outside $+/-$ 0.5), and a Hessian (see below) that was not invertible, which indicated that the optimization algorithm was not finding the maximum likelihood estimates. Effectively, N_{prior} and b were correlated in the model, so we fixed one of them to eliminate this problem. The specific value of N_{prior} had minimal effect on the overall likelihood of the model.

The parameters of the model were fit by maximizing the likelihood of the parameters, given the data, where if we set $l = 0$ if $\hat{f} = i$ and $l = 1$ if $\hat{f} = j$, and collect the parameters of the model into a vector w, the likelihood was given by:

$$p(D|w) = \prod_{k=1}^{N} \left(p_k(\theta_i > \theta_j) l_k + (1 - p_k(\theta_i > \theta_j))(1 - l_k) \right) \quad (9)$$

We maximized the log of this likelihood using *fminsearch* in Matlab starting from initial values of zero for all parameters to minimize the probability of finding false positives when we did significance testing.

Significance Testing of Model Parameters

Significance testing of the model was done in two ways. First, we tested the significance of additional parameters by carrying out a likelihood ratio test with and without the parameter of interest included in the model (Papoulis, 1991). We always dropped one parameter while keeping all the other parameters, similar to a type III sum-of-squares F test, because parameters were not independent. Thus, we calculated

$$lr = 2\ln\frac{p(D|w)}{p(D|w_{/i})}, \tag{10}$$

Where w is the vector of parameters in the ecological observer, and the notation $w_{/i}$ indicates the model obtained by removing parameter i. Asymptotically, lr has a χ^2 distribution with 1 degree of freedom, and thus it can be used for significance testing. Additionally, we numerically calculated the Hessian, which is the matrix of second partial derivatives of the likelihood function (Bishop, 1995):

$$H = \nabla\nabla p(D|w). \tag{11}$$

The diagonal elements of the inverse of this matrix give the variance of each parameter, which we used to construct confidence bounds (Huber, 2006). In all cases, both methods gave the same answer for hypothesis testing.

We also fit the model to the data from individual participants. This allowed us to approximate a mixed model, by fitting parameters to individual participants and then doing hypothesis testing on the parameter distribution (Holmes & Friston, 1998). Model fitting to individual subjects was more strongly affected by local minima, thus we used several initial values for the optimization to see which gave the best fit.

Mutual Information

We also calculated the mutual information between the ideal observer model and various parameterizations of the ecological observer. Mutual information characterizes the degree to which two variables are related. If the two variables are independent, that is, if knowing the value of one variable tells you nothing about the value of the other variable, then mutual information would be 0. However, if knowing the value of one variable tells you something about the other variable, mutual information is positive. In a sense, mutual information generalizes the idea of correlations to non-Gaussian variables. In this case, we accumulated the confusion matrix of choices for each model, and then used that to calculate the mutual information. The confusion matrix is calculated by adding one to row i and column j each time Model 1 selects face i and Model 2 select face j. The mutual information between model k and model l in bits is given by:

$$I(m_k, m_l) = \sum_{i,j=1}^{2} p(i,j) \log_2 \frac{p(i,j)}{p(i)p(j)}. \tag{12}$$

Results

The task was challenging because the two faces were stochastically rewarded and the difference between the probability of reward for each face was small (0.6 vs. 0.4). Thus, we first focused on establishing whether or not participants were able to integrate reward feedback information across trials to determine which face they should choose, where the face that should have been chosen was the one that was most frequently rewarded. Because the rewards delivered in the task were stochastic, the face which should have been chosen, based upon the history of rewards, varied from trial to trial. It was possible to have a series of outcomes that favored the face that had a true reward probability of 0.4. Thus, in our initial analyses, the participants' behavior was compared to an ideal observer, which optimally integrated the outcomes across trials (see methods and Figure 2). The ideal observer behaved as if it knew the distribution from which the actual trial data was being drawn, and the information it used to make its decisions corresponded to the instructions that were given to our subjects (see Methods). This model kept a running count of the number of positive outcomes received for each face and the number of times each face had been selected. From this count, a distribution over the probability that each face would be rewarded was computed on a trial-by-trial basis (Figures 2a-c). For example, even if one flips a coin 10 times and it comes up heads 5 times, the true probability of heads could be 0.4. With finite data, one can only ever guess at the true probability and our model effectively hedges its bets by maintaining a distribution over possible underlying probabilities. The peak of this distribution corresponded to the maximum likelihood estimate of the frequency of rewards associated with each face, and was given by r_i/N_i, which is the number of times face i was rewarded, divided by the number of times it was picked. Additionally, the width of the distribution showed the model's estimate of how confident the participants were of their estimate of the reward probability. As additional evidence was gathered, the distribution became narrower. To gain some intuition for this feature, if a coin was flipped 10 times and it came up heads 3 times, the best estimate of the probability of obtaining heads at the next toss would be 0.3, but the true probability might be 0.5. However, if one flipped a coin 1,000 times and it came up heads 300 times, the chance that the true probability of heads was 0.5 would be much lower. Thus, we have a distribution over possible reward frequencies and these distributions get narrower as we collect more data, reflecting our increased confidence in the frequency estimates. From these distributions, the probability that the angry face was more likely to be rewarded than the happy face could also be calculated (see Methods equation 2 and Figures 2c and d). This tells us which face should be picked (i.e., the one that is more often rewarded) and our confidence that this is the correct choice. Thus, this model provided a reference against which each subject's behavior could be compared.

Angry Versus Happy Expressions

We began by analyzing the data from the experiment in which participants chose between a happy and an angry face. To see if participants were in fact integrating the feedback, we accumulated contingency tables, comparing the face choice predicted by the ideal observer model to the face actually selected by the participants on a trial-by-trial basis (see Table 1). This table, pooled across participants, was highly significant ($p < .0001$; $\chi^2 = 21.7$, $df = 1$, $N = 1{,}248$) and 3 of the 12 participants had individually significant contingency tables ($p < .05$). We found that the participants selected the same face as the ideal observer on average

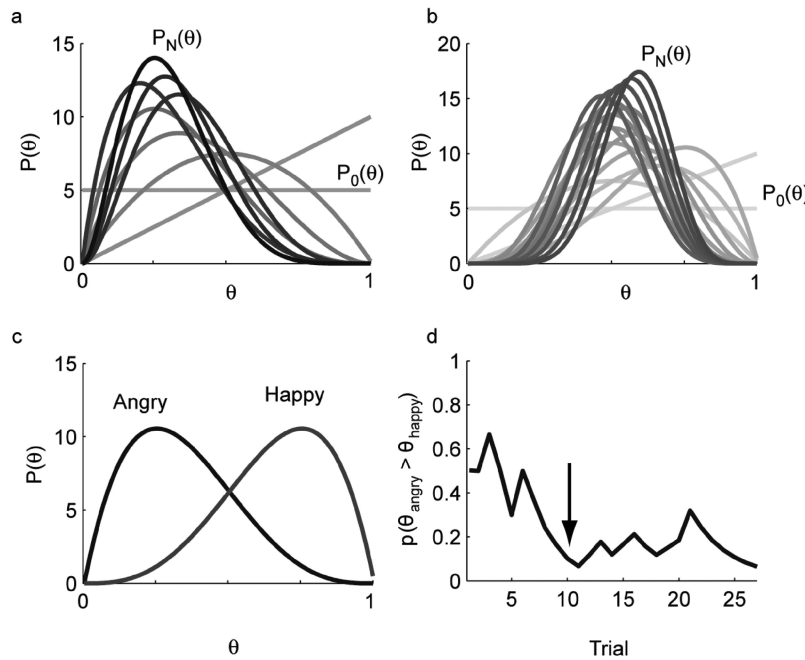

Figure 2. Trial-by-trial evolution of distributions and choice probability for an example block of data. a. Distribution of probability that the angry face will be rewarded derived from the ideal observer model. Trial evolution is encoded by increasing saturation of lines from 0 to *N*, where *N* is the number of trials in the current block. b. Same as panel a except distributions shown are for the happy face. c. Probability distributions of reward probability at $N = 10$. d. Probability that the angry face is more often rewarded than the happy face $p(\theta_{angry} > \theta_{happy})$ versus trial.

56% of the time. Further evidence supporting the assertion that participants were carrying out the task as instructed was given by a correlation between how strongly the model predicted the face that was selected and the reaction time of the participants on individual trials ($p < .0005$, $r = .14$). In this analysis the strength of the model's prediction was measured as the entropy, which measures the amount of variability. Entropy is maximum at a probability of 0.5 (no information about which target to pick) and then falls monotonically to zero when the probability is zero (which means pick the happy face with certainty) or 1 (which means pick the angry face with certainty). Thus, the more the evidence supported one face over the other and therefore the lower the entropy, the more quickly the participants responded (reaction time and entropy were z-transformed within participants before pooling for this analysis). We also examined whether or not the reaction time differed when selecting the happy versus the angry face and found that the RT was just significantly faster when selecting the angry face ($p < .05$, mean angry RT = 1.067 s, mean happy RT = 1.162 s). These results make it clear that the participants were able to carry out the task and extract information from the stochastic feedback.

Table 1
Contingency Table of Subject Selection vs. Selection of Ideal Observer for Angry vs. Happy Faces

		Ideal observer	
Choices	Expression	Angry	Happy
Participants	Angry	296	196
	Happy	353	403
		649	599

Next we began pursuing our primary hypothesis, which was that the facial expression would affect the decision-making process. We first approached this question by using the data in the contingency table (see Table 1) to calculate the probability that the participants chose the happy face when they should have chosen the angry face, $p(\text{happy}|\text{angry}) = 0.54$, and comparing it to the probability that the participants chose the angry face when they should have chosen the happy face, $p(\text{angry}|\text{happy}) = 0.33$, where the face that should have been chosen was given by the ideal observer model. These probabilities were in fact significantly different ($p < .0001$, likelihood ratio test, $df = 1$). Furthermore, 12 of the 12 participants showed a bias, such that $p(\text{happy}|\text{angry}) > p(\text{angry}|\text{happy})$, and the distribution of this bias across participants ($p(\text{happy}|\text{angry}) - p(\text{angry}|\text{happy}) = 0.22$, std $= 0.20$) was significantly different than zero ($p < .01$, *t* test, $n = 12$). Thus, there was a robust bias across participants to select the happy face, even when the evidence more strongly supported the angry face. This effect could also be seen by examining the frequency with which each face was picked, as a function of the strength of the evidence. This curve was biased toward happy faces, such that when the evidence equally supported the two faces, the happy face was being selected more than 60% of the time (see Figure 3). Thus, when we compared the performance of the

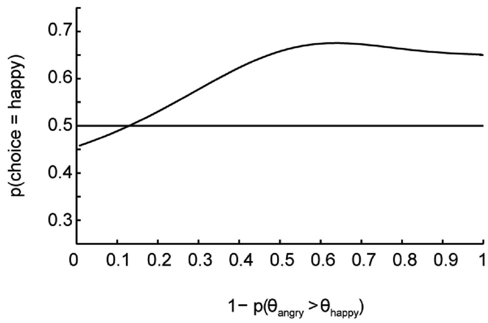

Figure 3. Fraction of times participants picked the happy face (*y*-axis) as a function of the evidence given by the ideal observer in favor of the happy face (*x*-axis) versus the angry face.

participants to an ideal observer, they had a marked bias to pick the happy face.

Ecological Model of Participants' Behavior

A drawback of the analyses in the previous section is that they do not tell us if participants simply had a prior bias to pick one face over the other, or if they in fact weighted the positive outcomes associated with the happy face more strongly than they weighted the positive outcomes associated with the angry face, where differential weighting of the outcomes would be a bias in evidence or reward accumulation. Either or both process could be responsible for the preference to choose the happy face. To address this, we tested a model that we refer to as the ecological observer (see Methods). The ecological observer takes into account a potential prior disposition toward one or the other face, as well as possible differential weighting of the reward feedback associated with the two faces. It also includes a factor that exponentially weights recent outcomes more than outcomes that occurred earlier in time because it is unlikely that participants would actually optimally integrate information from the beginning of the block. Importantly, however, this model does not presuppose ecological behavior, rather it is a tool that allows one to interpret the data and test our theory. We fit this model first to the data from all subjects pooled together and second, to each individual subject. Because the model fitting process is nonlinear, it is subject to local minima. In the pooled data, however, local minima are less likely. The individual subject analyses, however, are more conservative statistically.

Our first goal was to see if this model better predicted the choices of the participants than the ideal observer model, where prediction accuracy is measured with the likelihood function (see Methods equation 9). In linear regression the log of the likelihood function would be the log of the residual variance, but because our dependent variable is a binary decision and not a continuous variable, the log likelihood is the log of the model's prediction of each choice that the subjects made. If the model is predicting the subject's decisions well, it will predict the face the subject actually picked with a high probability. If the model predicts each decision of the subject with a probability of 0.5, it is guessing at what the subject will do, and it is not effective.

The ecological observer had three parameters. One that controlled the exponential forgetting (a), one that controlled the prior bias (b), and one that controlled the evidence bias or reward effect, which is the differential weighting of positive outcomes associated with a happy versus angry expressions (c). In the pooled analysis we found that all three parameters were statistically significant (mean [95% CI]; $a = 0.91$ [0.40 – 1.43], likelihood ratio = 75; $b = .19$ [0.12 – 0.26], likelihood ratio = 15.3; $c = 0.14$ [0.07 – 0.21], likelihood ratio = 5.4, log-likelihood of full model = 853.6, log-likelihood of the ideal observer = 971.0), although the differential weighting of positive outcomes was only weakly so. The exponential forgetting (a) and prior bias (b) terms had a synergistic interaction, such that dropping both had a larger effect than the sum of dropping either (likelihood ratio = 104). This occurs because the model is nonlinear, so dropping the two terms can be thought of as dropping two main effects and an interaction in a linear model, for example. Thus, the ecological observer accounted for the participants' behavior better than the ideal observer and, the expression of the face had an effect on the participants' decision-making processes, primarily mediated by a prior bias.

The parameters of the model accounted for the limited working memory capacity of the participants (Figure 4a) and accounted for the preference for the happy face (Figure 4b-c), which manifested as both a prior bias and a change in how positive outcomes were

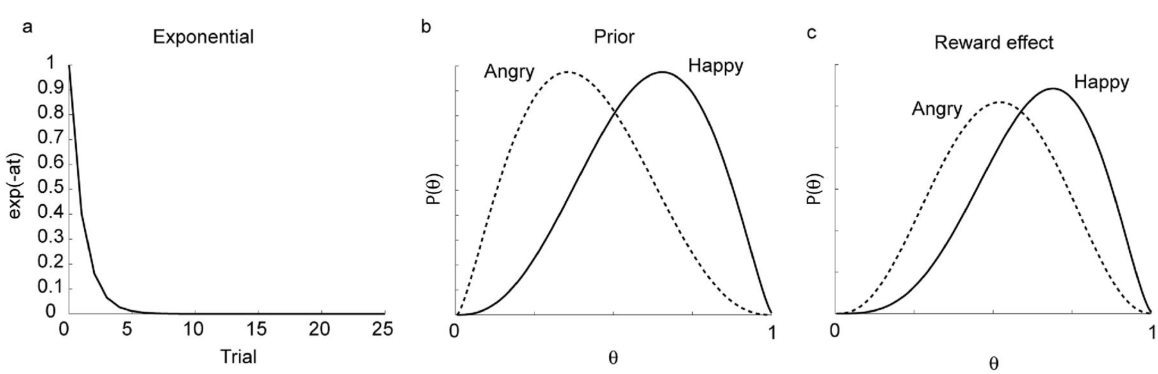

Figure 4. Illustration of ecological model. a. Exponential weighting function. b. Effects of prior bias on estimate of how often each face is rewarded. Distributions show subject estimates before any feedback has been received. c. Differential reward effect on estimate of how often each face has been rewarded. Results are shown after three positive outcomes for the happy face in five trials.

accumulated (reward effect). The prior bias affects decisions at the beginning of each block, but then its effects diminish, whereas the reward effect continues to have an impact throughout the block. In many cases, the ecological observer tracked the ideal observer relatively closely (Figure 5a), whereas in other cases it diverged from ideal behavior (Figure 5b). There are two salient features of the ecological observer that follow directly from the model. First, the ecological observer starts each block with a prior bias toward the happy target (Figure 5, Trial 1 at the left of the plot), and second, due to its exponential forgetting, it tends to have less stable behavior. This is because it does not average out the stochasticity in the feedback over all trials as the ideal observer does (e.g., Figure 5b). The reward effect cannot be seen as easily from these plots because it is a function of the feedback, which is not shown here. Thus, whereas the ideal observer often accurately identifies the target tat is being more often rewarded with a high degree of certainty, the ecological observer remains less certain (Figure 5b).

It is interesting to examine which features of the ecological observer model cause it to deviate most from the ideal observer model. Because the ecological observer, with all of its parameters set to 0, is equivalent to the ideal observer, we could set subsets of the parameters of the ecological observer to 0, and compare its predictions to the ideal observer (see Table 2). We compared the predictions of the two models by calculating the mutual information between their respective predictions. The mutual information (MI) is a measure of how often the predictions of one model match the predictions of the other model and has a maximum value of 1 bit because the decisions were binary. A value of 0 would indicate no predictive power. Parameters that cause the predictions of the ecological observer to deviate strongly from the ideal observer will strongly decrease the mutual information between the models. As can be seen from the table, the limited working memory has the largest effect (0.07 bits MI). However, the prior bias (0.40 bits) and the reward effect (0.68) also modified the predictions of the model. This analysis shows that the reward term has the smallest effect, followed by the prior bias, and then the limited working memory capacity. Combined, the prior and reward (0.40 bits) have the same effect as the prior term by itself, again showing that the reward effect has the smallest impact. It is possible for a parameter to significantly alter the likelihood of the model (i.e., be significant in a log-likelihood ratio test), without changing the actual prediction on any individual trials, because the prediction is found by thresholding the likelihood function (see equation 3 in methods). This is the case here, with the reward effect term.

Table 2
Mutual Information Between Ideal and Ecological Observers for Angry vs. Happy Faces. Single Parameter Indicates the MI Between the Ideal Observer and the Ecological Observer Model Containing Only the Corresponding Parameter. Dropped Parameter Indicates the Change in the MI by Dropping the Indicated Parameter From the Full Model

Parameter	Bits: Single parameter	Bits: Dropped parameter
a (exponential forgetting)	0.07	0.35
b (prior bias)	0.40	0.02
c (reward bias)	0.68	−0.01
b, c	0.40	0.02
Full model (a, b, c)	0.05	

Note. MI = mutual information.

We next examined the fit of the model to individual subjects by estimating the three parameters separately for each individual subject. We then carried out hypothesis testing by doing t tests on the parameter distributions across subjects. We found, similar to the pooled analysis, that the working memory (a) and prior bias terms (b) were significantly greater than zero ($p < .05$, t test, $n = 12$). However, the reward effect term did not reach significance. The reward effect term was marginally significant in the pooled analysis, and as the individual subject analysis is more conservative, it is not surprising that it did not reach significance here. For this analysis we also examined how much better the decisions of individual subjects were predicted by the full model than the ideal observer model. We found that the full model predicted the decisions of the subjects 65% of the time compared to 56% for the ideal observer as indicated above. The difference between the prediction of the ideal observer and the full model was statistically significant across subjects ($p < .05$, paired t test, $n = 12$).

Sad Versus Happy Expressions

We sought to extend the results of the experiment where we gave the participants a choice between angry and happy expres-

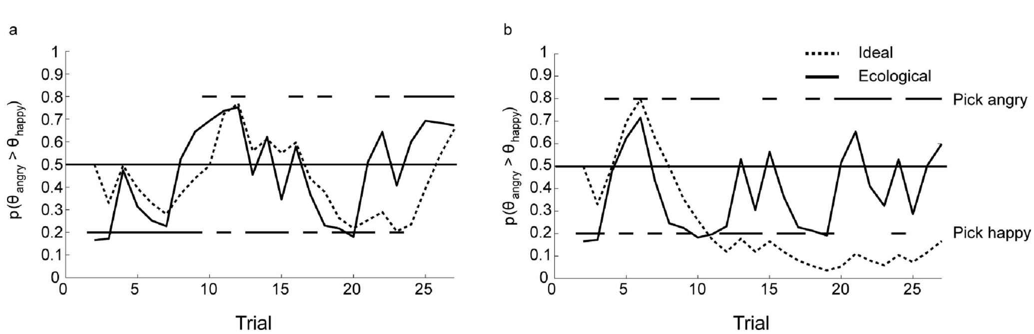

Figure 5. Two example blocks showing evolution of evidence under the ideal and ecological observer models. The intermittent lines in each plot at 0.2 and 0.8 indicate when either the happy (0.2) or angry (0.8) face was selected in the corresponding trial.

sions, by carrying out an experiment where they had a choice between happy and sad expressions. Analysis of the data from this experiment showed similar, although interestingly different, results. Specifically, the pooled contingency table (Table 3) was highly significant ($p < .0001$, $\chi^2 = 111.9$, $df = 1$, $N = 1{,}248$), and 7 of the 12 participants had individually significant tables ($p < .05$). Additionally, participants chose the same face as the ideal observer 66% of the time, and the distribution of this percentage across participants was significantly different than 50% ($p < .001$, t test, $n = 12$), The correlation between the entropy of the model and the subject's reaction time was also significant ($p < .01$, $r = .12$, $n = 1248$). Thus, participants were extracting information from the feedback about which face was most highly rewarded and picking that face more often.

Similar to what was found in the analysis of angry versus happy faces, when compared to the ideal observer, participants chose the happy expression when they should have chosen the sad expression more often than they chose the sad expression when they should have chosen the happy expression: $p(\text{happy}|\text{sad}) = 0.43$; $p(\text{sad}|\text{happy}) = 0.27$ ($p < .001$, likelihood ratio test, $df = 1$). This bias was seen in 10 of 12 participants, and the distribution across participants of probability differences was also significant: $p(\text{happy}|\text{sad}) - p(\text{sad}|\text{happy}) = 0.13$, std $= 0.18$ ($p = .03$, t test, $n = 12$). This difference was smaller than we found with the angry expressions, but because we did not run the experiment on the same participants, it is not clear whether or not the difference in effect size is due to the stimuli or the participants. A plot of evidence versus decisions also showed a bias toward happy faces (see Figure 6) although again this plot shows that the bias was not as strong for sad faces because the curve is shifted to the right when compared with Figure 3. However, participants that decided between happy versus sad faces also selected the happy face more often when the evidence more strongly supported the sad face (see Figure 6).

When we fit the ecological observer to the sad versus happy data, we found that the exponential forgetting and reward effect terms were significant, but that the prior bias was not significant (mean [95% CI] $a = 0.24$ [0.16 – 0.32], likelihood ratio = 46.3; $b = .03$ [-0.01 – 0.06], likelihood ratio = 0.7; $c = 0.15$ [0.10 – 0.20], likelihood ratio = 15.8, log-likelihood of full model = 766.2, likelihood of ideal observer = 830.5). Thus, in contrast to the angry faces, the sad faces elicited a reward effect, but not a prior bias toward the happy face, when it was compared to the sad face.

We also compared the mutual information between the ecological observer and the ideal observer for the sad versus happy expressions (see Table 4). We only made this comparison for the significant factors, exponential forgetting and reward. With only the exponential forgetting parameter in the model, the MI was 0.33, whereas the MI was 0.65 for the model with the reward effect term, and 0.38 for the full model (parameters a and c). In this case the exponential forgetting had a smaller effect on the MI, corresponding to the smaller value of the parameter in this data. This is also consistent with the fact that participants made the same choice as the ideal observer more often in this case than in the angry versus happy case.

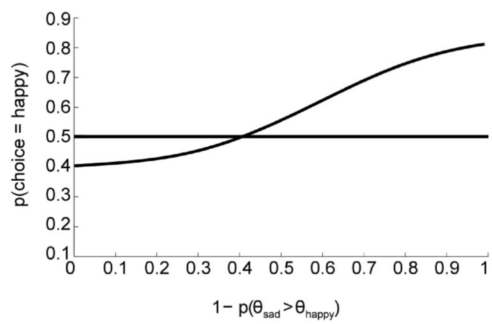

Figure 6. Fraction of times participants picked the happy face (y-axis) as a function of the evidence given by the ideal observer in favor of the happy face (x-axis) versus the sad face.

Finally, we fit the model to individual subjects for the sad versus happy data. In this case we found, again consistent with the pooled analysis, that the working memory and reward effect terms were significant ($p < .05$, one-tailed t test, $n = 12$). The prior bias term, was, however, not significant. The full model also predicted the decisions of individual subjects 71% of the time, and the difference between the fraction correct of the full and ideal observers subject-by-subject was significantly different than zero ($p < .05$, paired t test, $n = 12$).

Discussion

The results clearly demonstrate that the emotional expression on a face influenced the participants' decision making. Specifically, participants showed a prior bias toward smiling faces when compared to angry faces, and they also overestimated the number of positive outcomes associated with a smiling face and underestimated the number of positive outcomes associated with both angry and sad faces. There was asymmetry in these effects though as the largest effect for happy versus angry manifested in the prior bias and the largest effect for happy versus sad manifested in the reward effect. Additionally, the participants were not able to integrate information across an entire block of trials about which face was most often rewarded. Rather, their decisions were most

Table 3
Contingency Table of Subject Selection vs. Selection of Ideal Observer for Sad vs. Happy Faces

		Ideal observer	
Choices	Expression	Sad	Happy
Participants	Sad	322	186
	Happy	244	496
		566	682

Table 4
Mutual Information Between Ideal and Ecological Observers for Sad vs. Happy Faces

	Bits	
Parameter	Single parameter	Dropped parameter
a (exponential forgetting)	0.33	0.33
b (prior bias)	Not significant	Not significant
c (reward bias)	0.65	−0.08
Full model (a, c)	0.38	

heavily influenced by the previous two to five trials. Despite this limited integration, our data also shows that participants were able to extract information from the task about which face was more likely to be rewarded. As such, their decision-making process was integrating both social and utilitarian information. While we have found an effect of social cues, nonsocial cues may have similar effects.

In our task, a purely utilitarian or ideal observer should have ignored the expressions. However, the expressions had an effect on the participants' decisions. Often, when participants' performance is not consistent with ideal observers, it is considered irrational (Gilovich et al., 2002; Sutherland, 1992). A different perspective, however, suggests that such behavior is not irrational, but ecologically or boundedly rational, that is, that it is the result of phylogenetic or ontogenetic shaping of behavior to be reasonable given the design of the organism and the environmental context (Cosmides & Tooby, 2000; Gigerenzer, 2000; Simon, 1956). In social interactions, expressions are a critical indicator of the intention of an interactant. Smiles are often conceptualized as a signal of cooperative intent (Fridlund, 1995) and are associated with positive outcomes whereas expressions of anger and sadness are not. These considerations suggest that decision making should be influenced by facial expression information, so that individuals are more likely to engage in interactions with individuals expressing cooperative intent. This reasoning is consistent with a folk psychology view of expressions and social interactions, and our results indicate that expression information is included in decision making even when it is not relevant to the task. This finding is similar to other demonstrations that participants in a variety of situations are unable to de-couple decision making from its social context (Cosmides & Tooby, 2000; Haley & Fessler, 2005; Hoffman, McCabe, & Smith, 1996).

Effects of Faces on Reward and Communication Systems

Within our task, the impact of the emotional valence of the faces on decision making could have been driven by at least two processes; the faces could be functioning as primary reinforcers or as communication stimuli. If they were functioning as primary reinforcers, the participants may have been predisposed to select the happy faces, even when these faces were causing them to fail at the task, because they were balancing task failure with the positive affect associated with seeing the smiling face. There are several lines of evidence that suggest that the emotional valence of a face might function as a primary reinforcer. For example, happy expressions have been shown to function as primary rewards in behavioral studies (Matthews & Wells, 1999) and have also been shown to activate ventral medial prefrontal cortex (O'Doherty et al., 2003), an area which represents stimulus reward value. On the other hand, angry expressions may function as social communication stimuli indicating that the subject toward whom the angry expression is directed should suppress or change their behavior (Blair & Cipolotti, 2000). Consistent with this, neural activation to angry faces is centered in the anterior cingulate cortex (Blair, Morris, Frith, Perrett, & Dolan, 1999). Ultimately it may be difficult to disentangle the communication and primary reward functions of facial expression, because communicating happiness is likely rewarding, whereas communicating anger is likely not rewarding, and may even be punishing.

Complementary to studies which have shown activation in reward regions, facial attractiveness and expressions have also been shown to directly affect decision making. For example, in an experiment designed to test the hypothesis that smiling conveys an intention to cooperate by a partner, it was found that participants were more likely to trust smiling than nonsmiling partners (Scharlemann et al., 2001). Additionally, potential partners ranked as more attractive were more often trusted in a trust and reciprocity game. However, these attractive partners were also more aggressively punished than less attractive counterparts when they failed to reciprocate trust (Wilson & Eckel, 2006). Finally, when thirsty participants were shown happy faces subliminally, they drank more of a beverage and indicated a greater willingness to pay for the beverage when compared to the same condition in which they were shown an angry face (Winkielman et al., 2005).

Our results extend these findings in several important ways. First, our task combined with our analytical approach allowed us to separate the effects of the emotional expressions into a prior disposition, as well as a differential weighting of the feedback. Specifically, we show that facial expressions affect an explicit decision making process. When happy faces were contrasted with angry faces this manifested mostly as a prior bias, although there was a small reward bias and when happy faces were contrasted with sad faces this manifested mostly as a reward bias. Thus, in both bases, but most strongly in the happy versus sad case the effect of facial expression was not eliminated by feedback within the task. It is possible that extensive experience with the faces might ultimately eliminate these effects, but long blocks of trials would likely lead to inattention by participants, as the task is somewhat demanding.

It is not clear what is driving the difference between the effect of expression in Experiment 1 (happy vs. angry) and Experiment 2 (happy vs. sad), but it may be related to whether or not these expressions function as approach versus avoidance cues. Sadness may elicit approach, much like happiness, whereas anger may provoke avoidance (Marsh, Ambady, & Kleck, 2005). Furthermore, sad and angry expressions lead to responses in different brain regions (Blair et al., 1999); therefore, these expressions may be tapping different cognitive systems.

Our approach also allows us to quantify the amount by which behavior deviates from what would be expected of an ideal or utilitarian observer, as well as which factors most drive performance away from the ideal observer. Importantly, however, it seems likely that the deviation of participants' behavior from the ideal observer within our task is not irrational. Rather it may be ecologically rational. In most social situations, a subject who did not take into account facial expressions would behave suboptimally. Therefore, it is only within the restricted laboratory setting that an ideal observer model outperforms the participants.

Conclusion

We found that the emotional content of face stimuli affected participants' decisions in a task in which they were supposed to learn which of two faces was more rewarding. With angry faces, this effect manifested as both a prior bias toward the happy face as well as a small reward bias. With sad faces, the effect manifested significantly as a reward bias with no prior bias. While we have shown these effects using emotional expressions, we expect that other cues without social content could have a similar effect.

Participants in our task were also integrating the feedback information. Thus, they were combining social and utilitarian factors to arrive at their decision, a strategy that may reflect an ecologically rational approach. This opens up several interesting questions, including which cortical areas and anatomical pathways process social and utilitarian information, as well as whether or not a solid ecological foundation for these effects can be established. Our method also provides a flexible means of quantitatively comparing the effect of a wide variety of stimuli, both social and nonsocial, on decision making. Finally, while this study shows how social cues interact with utilitarian cues in this specific case, additional work will be necessary to understand how these findings generalize to different contexts.

References

Bar-Hillel, M. (1990). Back to base-rates. In R. M. Hogarth (Ed.), *Insights in decision making.* Chicago: University of Chicago Press.

Bishop, C. M. (1995). *Neural networks for pattern recognition (1st ed.).* Oxford: Oxford University Press.

Blair, R. J., & Cipolotti, L. (2000). Impaired social response reversal: A case of "acquired sociopathy." *Brain, 123,* 1122–1141.

Blair, R. J., Morris, J. S., Frith, C. D., Perrett, D. I., & Dolan, R. J. (1999). Dissociable neural responses to facial expressions of sadness and anger. *Brain, 122,* 883–893.

Cosmides, L., & Tooby, J. (2000). The cognitive neuroscience of social reasoning. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences.* Cambridge, MA: MIT Press.

Ellsberg, D. (1961). Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics, 75,* 643–669.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415,* 429–433.

Fridlund, A. J. (1995). *Human facial expression: An evolutionary view.* London: Academic Press.

Gigerenzer, G. (2000). *Adaptive thinking. Rationality in the real world.* Oxford: Oxford University Press.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment.* Cambridge, UK: Cambridge University Press.

Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Science, 5,* 10–16.

Haley, K. J., & Fessler, D. M. T. (2005). Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evolution and Human Behavior, 26,* 245–256.

Hoffman, E., McCabe, K., & Smith, V. L. (1996). Social distance and other-regarding behavior in dictator games. *The American Economic Review, 86,* 653–660.

Holmes, A. P., & Friston, K. J. (1998). *Generalisability, random effects and population inference. NeuroImage, 7,* s754.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science, 310,* 1680–1683.

Huber, D. E. (2006). Computer simulations of the ROUSE model: An analytic simulation technique and a comparison between the error variance-covariance and bootstrap methods for estimating parameter confidence. *Behavioral Research Methods, 38,* 557–568.

Jacobs, R. A. (1999). Optimal integration of texture and motion cues to depth. *Vision Research, 39, 3621*–3629.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge: Cambridge University Press.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychology Review, 80,* 237–251.

Kersten, D. (1999). High level vision as statistical inference. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences.* Cambridge: MIT Press.

Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference.* Cambridge: Cambridge University Press.

Knill, D. C., & Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research, 43,* 2539–2558.

Kording, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Science, 10,* 319–326.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience, 9,* 1432–1438.

Marsh, A. A., Ambady, N., & Kleck, R. E. (2005). The effects of fear and anger facial expressions on approach- and avoidance-related behaviors. *Emotion, 5,* 119–124.

Matthews, G., & Wells, A. (1999). The cognitive science of attention and emotion. In T. Dalgleish & M. J. Power (Eds.), *Handbook of cognition and emotion.* New York: Wiley.

McNeil, B. J., Pauker, S. G., Sox, H. C., Jr., & Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *New England Journal of Medicine, 306,* 1259–1262.

O'Doherty, J., Winston, J., Critchley, H., Perrett, D., Burt, D. M., & Dolan, R. J. (2003). Beauty in a smile: The role of medial orbitofrontal cortex in facial attractiveness. *Neuropsychologia, 41,* 147–155.

Papoulis, A. (1991). *Probability, random variables and stochastic processes* (3rd ed.). New York: McGraw-Hill Higher Education.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology, 43,* 87–131.

Poggio, T., Torre, V., & Koch, C. (1985). Computational vision and regularization theory. *Nature, 317,* 314–319.

Scharlemann, J. P. W., Eckel, C. C., Kacelnik, A., & Wilson, R. K. (2001). The value of a smile: Game theory with a human face. *Journal of Economic Psychology, 22,* 617–640.

Sher, S., & McKenzie, C. R. M. (2008). Framing effects and rationality. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 79–96). Oxford: Oxford University Press.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63,* 129–138.

Singer, T., Kiebel, S. J., Winston, J. S., Dolan, R. J., & Frith, C. D. (2004). Brain responses to the acquired moral status of faces. *Neuron, 41,* 653–662.

Sutherland, S. (1992). *Irrationality.* London: Pinter and Martin.

Thuring, M., & Jungermann, H. (1990). The conjunction fallacy: Causality vs. event probability. *Journal of Behavioral Decision Making, 3,* 61–74.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience, 5,* 1226–1235.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211,* 453–458.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgement. *Psychological Review, 90,* 293–315.

Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business, 59,* S251–S278.

Wilson, R. K., & Eckel, C. C. (2006). Judging a book by its cover: Beauty and expectations in the trust game. *Political Research Quarterly, 59,* 189–202.

Winkielman, P., Berridge, K. C., & Wilbarger, J. L. (2005). Unconscious affective reactions to masked happy versus angry faces influence consumption behavior and judgments of value. *Personality and Social Psychology Bulletin, 31,* 121–135.

Received March 10, 2008
Revision received April 16, 2009
Accepted April 17, 2009 ■