

Feedback effects on cost–benefit learning in perceptual categorization

W. TODD MADDOX and COREY J. BOHIL
University of Texas, Austin, Texas

Two experiments were conducted in which the effects of different feedback displays on decision criterion learning were examined in a perceptual categorization task with unequal cost–benefits. In Experiment 1, immediate versus delayed feedback was combined factorially with objective versus optimal classifier feedback. Immediate versus delayed feedback had no effect. Performance improved significantly over blocks with optimal classifier feedback and remained relatively stable with objective feedback. Experiment 2 used a within-subjects design that allowed a test of model-based instantiations of the flat-maxima (von Winterfeldt & Edwards, 1982) and competition between reward and accuracy (Maddox & Bohil, 1998a) hypotheses in isolation and of a hybrid model that incorporated assumptions from both hypotheses. The model-based analyses indicated that the flat-maxima model provided a good description of early learning but that the assumptions of the hybrid model were necessary to account for later learning. An examination of the hybrid model parameters indicated that the emphasis placed on accuracy maximization generally declined with experience for optimal classifier feedback but remained high, and fairly constant for objective classifier feedback. Implications for cost–benefit training are discussed.

Every day, we make hundreds of decisions on the basis of uncertain information. For example, the medical doctor must decide from an x-ray whether a patient has a tumor or not, the football referee must decide whether the defender interfered with the receiver, or the traveler on a two-lane road must decide whether there is enough time to pass the cars in front. These are categorization problems, because every pattern of light and dark on an x-ray, every pass play, and every road condition is unique, yet there are only two possible decisions—“tumor” or “no tumor,” “interference” or “no interference,” and “pass” or “do not pass.” In each case, we can be correct or incorrect, but perfect performance is impossible. Categorization problems of this sort differ along a number of factors. This article focuses on two: category discriminability and the costs and benefits associated with each categorization decision.¹ Category discriminability is defined as the standardized distance between category means, also called d' (Green & Swets, 1966; Macmillan & Creelman, 1991; Maddox & Dodd, in press). For example, a patient’s degree of chest pain might be a good discriminator between heart attack and no heart attack (high category discriminability) but might be a poor discriminator between can-

cer and no cancer (low category discriminability). Costs and benefits are the rewards and punishments accrued for correct and incorrect categorization responses. For example, the medical doctor generally benefits when making the *correct* diagnosis, but the benefit of a correct tumor diagnosis might be greater than the benefit of a correct no-tumor diagnosis. Similarly, there is generally a cost associated with an *incorrect* decision; a more severe cost might be incurred if a cancer patient is misdiagnosed as having no tumor than if a healthy individual is misdiagnosed as having a tumor.

When presented with a categorization problem of this sort, the optimal classifier (a hypothetical device) maximizes long-run reward by setting a decision criterion along the stimulus dimension and giving one categorization response for dimensional values below the criterion and the other categorization response for dimensional values above the criterion. The location of the optimal decision criterion is determined from the costs and benefits. A doctor who behaves like the optimal classifier would be biased toward making tumor diagnoses, because this strategy will save more lives, even though this bias will lead to more incorrect diagnoses in the long run. This doctor is more willing to misdiagnose a healthy person as having a tumor than he or she is to misdiagnose a cancer victim as healthy. Importantly, the magnitude of the sacrifice in accuracy associated with reward maximization is affected by category discriminability, with low category discriminability leading to a large sacrifice in accuracy and high category discriminability leading to a smaller sacrifice in accuracy.

This research was supported in part by National Science Foundation Grant SBR-9796206 and NIH Grant R01 MH59196. We thank Steven Sloman and one anonymous reviewer for helpful comments on an earlier version of this manuscript. We also thank Lorilei Cardenas for help with data collection. Correspondence concerning this article should be addressed to W. T. Maddox, Department of Psychology, Mezes Hall 330, University of Texas, Austin, TX 78712 (e-mail: maddox@psy.utexas.edu).

To date, a number of studies have examined decision criterion learning for biased cost–benefit ratios (e.g., Bohil & Maddox, 2001; Busemeyer & Myung, 1992; Busemeyer & Rappaport, 1988; Erev, 1998; Maddox & Bohil, 1998a, 1998b, 2000; Stevenson, Busemeyer, & Naylor, 1991; for related work, see Erev, Wallsten, & Budescu, 1994; Wallsten & Gonzalez-Vallejo, 1994). The most common result is that human observers use a decision criterion that is intermediate between the optimal reward-maximizing criterion, and the unbiased accuracy-maximization criterion, even though they are instructed (and paid) to maximize reward. Maddox and Bohil (1998a) offered a competition between reward and accuracy maximization (COBRA) hypothesis to account for this pattern of results. The idea is that observers attempt to maximize long-run reward, as instructed, but also place some importance on the accuracy of their responses. One reason that observers might place some importance on accuracy is that they generally receive local (trial-by-trial) information regarding the accuracy of their responding. For example, following the categorization response, the feedback display often includes the reward accrued on that trial and the reward accrued for a correct response (Bohil & Maddox, 2001; Maddox & Bohil, 1998a, Experiment 2, 2000). For correct responses these values are the same, and for incorrect responses these values are different. Notice that a display of this sort includes *corrective feedback* regarding reward and accuracy maximization.

This article reports the results of two simulated medical diagnosis categorization experiments with a biased cost–benefit ratio. In Experiment 1, different category discriminabilities were combined factorially with different types of corrective feedback in an attempt to identify feedback that would enhance the local information regarding reward maximization and reduce the local information regarding accuracy maximization, leading to more nearly optimal decision criterion learning. Each observer completed several blocks of trials in only one of the experimental conditions. This allowed an examination of the time course of decision criterion learning under various conditions. Experiment 2 included a subset of the Experiment 1 conditions, but in Experiment 2 each observer completed several blocks of trials in *all* the experimental conditions. The within-subjects design allowed us to develop and test a series of quantitative models, each of which was applied *simultaneously* to the data from all the experimental conditions separately for each observer and block of trials (Bohil & Maddox, 2001; Maddox & Ashby, 1993; Maddox & Bohil, 1998a, 2000; Maddox & Dodd, in press). Each model instantiated a different hypothesis regarding the effect of category discriminability, feedback, and cost–benefits on decision criterion learning. Two specific hypotheses will be outlined shortly, and the details of these models will be left for the Results section of Experiment 2 and the Theoretical Analyses section. Before we turn to the experiments, a brief review of the optimal classifier and these two hypotheses is in order.

The optimal classifier is a hypothetical device that maximizes long-run reward and is willing to sacrifice long-

run accuracy to attain this goal (e.g., Green & Swets, 1966). More formally, consider the situation facing a medical doctor who must classify a patient into one of two disease categories, A or B. Suppose the patient is given medical test X, which is diagnostic of the two diseases. In addition, suppose that the outcomes of test X for diseases A and B are normally distributed, with means μ_A and μ_B and standard deviation σ_A and σ_B , as depicted in Figure 1A. The optimal classifier records perfectly the test result, denoted x . The optimal classifier has perfect knowledge of the distribution of test results for each disease category (i.e., the form and parameters of the distribution). This information is used to construct the *optimal decision function*, which is computed from the likelihood ratio of the two category distributions,

$$l_o(x) = f(x|B)/f(x|A), \quad (1)$$

where $f(x|i)$ denotes the likelihood of test result x given disease category i . If test result x is more likely to result from disease B than from disease A, the likelihood ratio will be greater than one. If test result x is more likely to result from disease A than from disease B, the likelihood ratio will be less than one. Figures 1A and 1B depict categorization problems for two discriminabilities, $d' = 1.0$ and $d' = 2.2$, respectively. The optimal classifier has perfect knowledge of the costs associated with incorrect diagnoses and the benefits associated with correct diagnoses. This information is used to construct the *optimal decision criterion*

$$\beta_o = (V_{aA} - V_{bA})/(V_{bB} - V_{aB}), \quad (2)$$

where V_{aA} and V_{bB} denote the benefits associated with correct diagnoses and V_{bA} and V_{aB} denote the costs associated with incorrect diagnoses (lowercase subscripts denote responses, and uppercase subscripts denote categories). The optimal classifier (e.g., Green & Swets, 1966) uses $l_o(x)$ and β_o to construct the *optimal decision rule*:

$$\begin{aligned} &\text{If } l_o(x) > \beta_o, \text{ then respond "B";} \\ &\text{otherwise, respond "A."} \end{aligned} \quad (3)$$

When the cost–benefit differences are equal across categories (i.e., when $V_{aA} - V_{bA} = V_{bB} - V_{aB}$), then $\beta_o = 1$, and the optimal classifier assigns the stimulus to the category with the highest likelihood, simultaneously maximizing long-run accuracy and long-run reward. When the cost–benefit differences are unequal—for example, when $(V_{aA} - V_{bA}) = 3(V_{bB} - V_{aB})$, then $\beta_o = 3.0$, and the optimal classifier will generate a disease A diagnosis unless the likelihood of disease B is at least three times larger than the likelihood of disease A. Under these conditions, the optimal classifier will sacrifice long-run accuracy, which is obtained by using the $\beta = 1.0$ decision criterion, to maximize long-run reward by using $\beta_o = 3.0$.

Category Discriminability and the Flat-Maxima Hypothesis

When the cost–benefit differences are unequal, the optimal classifier sacrifices some degree of accuracy to




Figure 1. Hypothetical distributions for Categories A and B when (A) $d' = 1.0$, and (B) $d' = 2.2$. The $\beta = 1$ and $\beta_0 = 3$ values denote the accuracy-maximizing (equal likelihood) and reward-maximizing (optimal) decision criteria, respectively, for a 3:1 cost-benefit condition.

maximize long-run reward. The magnitude of this sacrifice in accuracy declines as category discriminability, d' , increases. For example, the optimal classifier must sacrifice 8% accuracy when $d' = 1.0$, but only 3% when $d' = 2.2$ (see Table 1). In addition, category discriminability has other important theoretical implications for decision criterion learning. As has been suggested by many researchers, suppose that the observer adjusts his or her decision criterion (at least in part) on the basis of the rate of change in reward, with larger values being associated with faster decision criterion learning (e.g., Busemeyer & Myung, 1992; Dusoir, 1980; Erev, 1998; Erev, Gopher, Itkin, & Greenspan, 1995; Kubovy & Healy, 1977; Roth & Erev, 1995; Thomas & Legge, 1970). To formalize this hypothesis, one can construct the objective reward function (ORF) that plots the objective expected (long-run) reward as a function of the decision criterion placement (e.g., Busemeyer & Myung, 1992; Stevenson et al., 1991; von Winterfeldt & Edwards, 1982). Figure 2A displays

the ORF for $d' = 1.0$ and $d' = 2.2$ for a 3:1 cost-benefit ratio. Specifically, Figure 2A plots expected reward as a function of the deviation between the observer's decision criterion (β) and the optimal decision criterion (β_0) standardized by d' . This $k - k_0$ measure, $\ln(\beta)/d' - \ln(\beta_0)/d' = \ln(\beta/\beta_0)/d'$, is the ratio of the actual and optimal decision criterion standardized by d' . The derivative of the ORF denotes the rate of change in expected reward; the larger rate of change, the steeper the ORF at that point. Figure 2B shows the relationship between the steepness of each ORF and $k - k_0$. The horizontal line on Figure 2B denotes a fixed steepness value, and the vertical lines denote the associated $k - k_0$ values for each d' . Note that the observer's decision criterion, k , is always closer to the optimal value, k_0 , for the steeper ORF— $d' = 2.2$, in this case. Thus, if the observer adjusts his or her decision criterion on the basis of the rate of change in reward (or steepness of the ORF), steeper ORFs should be associated with more nearly optimal decision criterion.

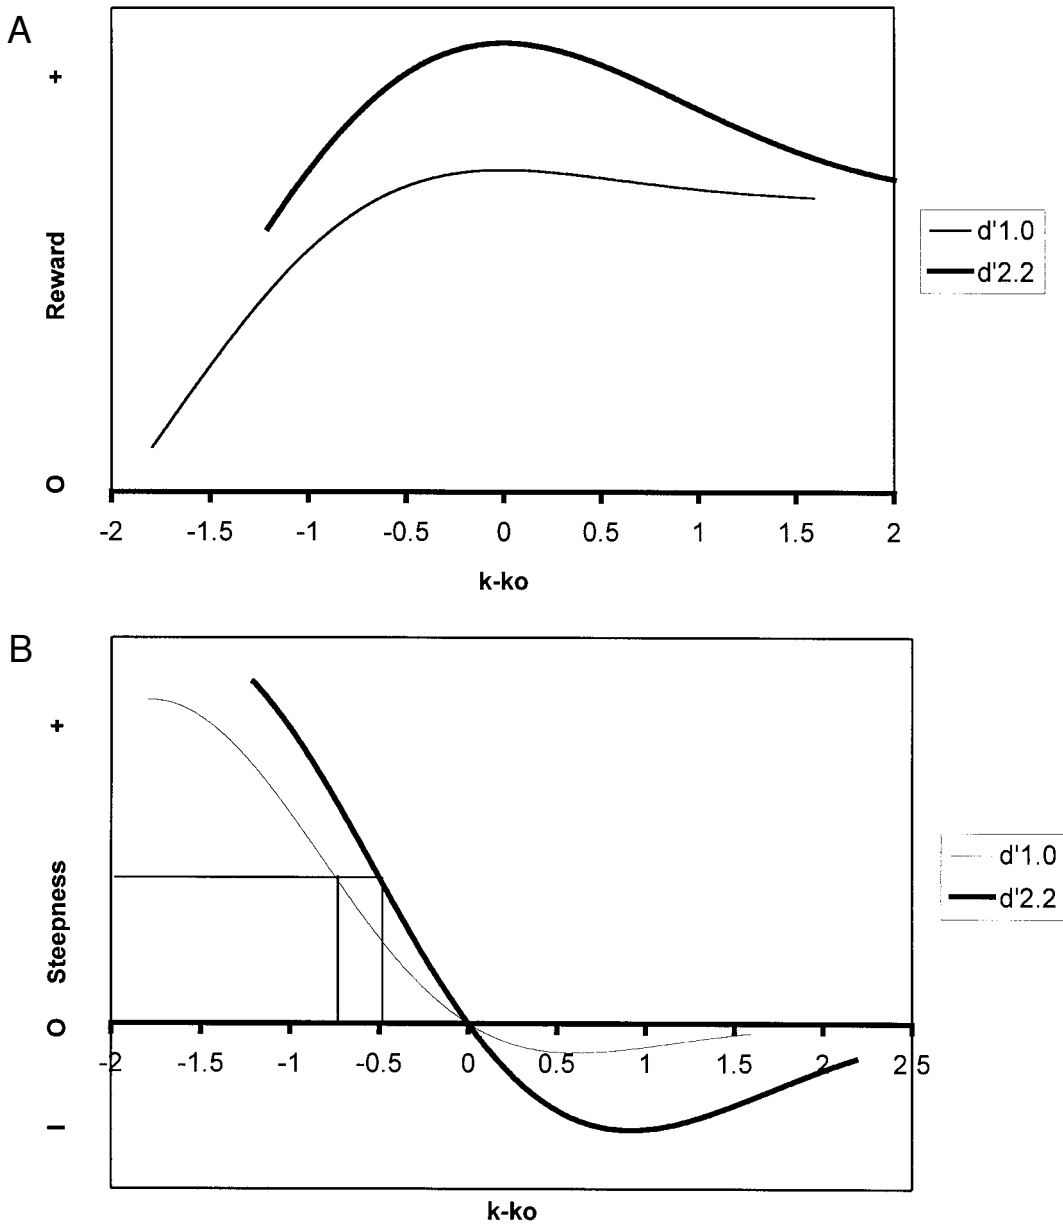


Figure 2. (A) Expected reward as a function of the deviation from the optimal decision criterion (i.e., $k - k_0$), called the objective reward function (ORF) for category discriminability, $d' = 1.0$ and 2.2 . (B) Steepness of the ORF (from panel A) for category discriminability, $d' = 1.0$ and 2.2 . The two vertical lines denote the $k - k_0$ values for a fixed value of steepness. Notice that the deviation from the optimal decision criterion is smallest for $d' = 2.2$ and largest for $d' = 1.0$.

von Winterfeld and Edwards (1982; see also Kubovy & Healy, 1980) referred to this as the *flat-maxima* hypothesis.

Unequal Cost–Benefit Differences and the Competition Between Reward and Accuracy Maximization Hypothesis of Decision Criterion Learning

A ubiquitous finding in the categorization literature is that the observer’s decision criterion is generally intermediate between the optimal reward-maximizing crite-

ri- on and the accuracy-maximizing criterion when the cost–benefit differences are unequal. Maddox and Bohil (1998a) offered the COBRA hypothesis, which assumes that the observer attempts to maximize long-run reward, as instructed, but also places some importance on the accuracy of his or her responses. For example, consider the two-category problem depicted in Figure 3. In this case, the cost–benefit difference is unequal. Notice that the decision criterion that maximizes expected reward (k_r) is different from the decision criterion that maximizes ex-

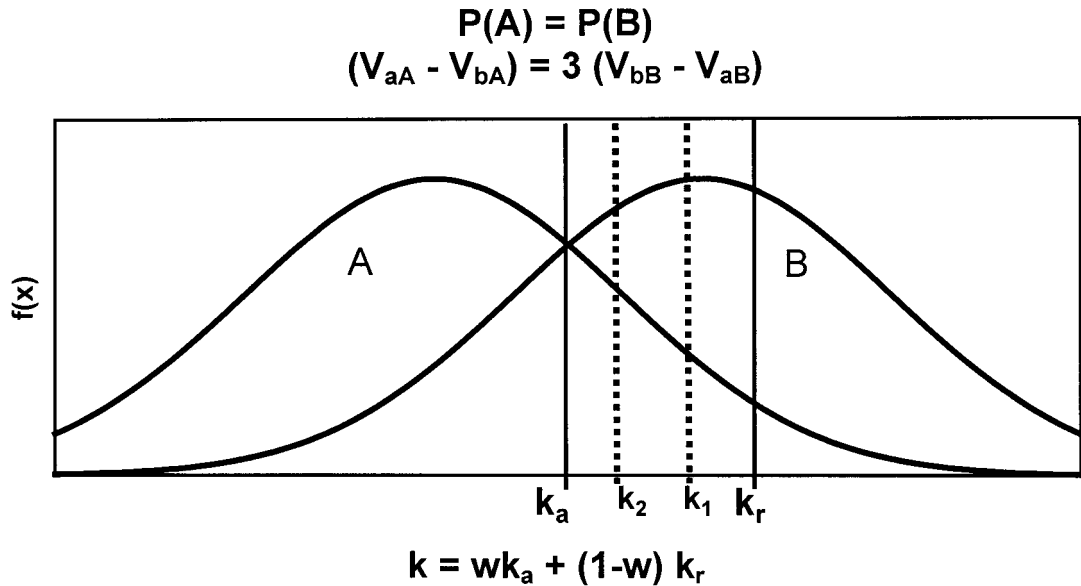


Figure 3. Schematic illustration of the competition between reward and accuracy (COBRA) hypothesis. The k_r decision criterion denotes the decision criterion that is being used by the observer in an attempt to maximize expected reward. The k_a decision criterion denotes the decision criterion that maximizes expected accuracy. The k_1 decision criterion denotes the decision criterion resulting from the COBRA hypothesis with the assumption that less importance (or weight; $w < .5$) is being placed on accuracy maximization. The k_2 decision criterion denotes the decision criterion resulting from the COBRA hypothesis with the assumption that more importance (or weight; $w > .5$) is being placed on reward maximization.

pected accuracy (k_a). Thus, an observer who places importance (or weight) on both goals will use a suboptimal decision criterion. To instantiate this hypothesis conceptually and within the framework of a mathematical model, we assume a simple weighting function, $k = wk_a + (1 - w)k_r$, where w ($0 \leq w \leq 1$) denotes the weight placed on accuracy maximization. This weighting function results in a single decision criterion that is intermediate between that for accuracy maximization and that for reward maximization.² In Figure 3, k_1 denotes a case in which $w < .5$, whereas k_2 denotes a case in which $w > .5$.

A major focus of the present study was to determine whether the flat-maxima and COBRA hypotheses were important in accounting for performance differences across feedback conditions. Specifically, model-based instantiations of each hypothesis in isolation were compared with a hybrid model that instantiated important assumptions from both hypotheses with respect to their ability to account for decision criterion learning across category discriminability and feedback conditions.

EXPERIMENT 1

In Experiment 1, we examined the optimality of cost-benefit learning for different types of feedback displays. Two factors were examined. First, we manipulated the number of trials on which feedback was provided. In the *immediate feedback* condition, feedback was provided on each trial. This is the approach taken in most studies of cost-benefit learning (e.g., Healy & Kubovy, 1981;

Maddox, 1995; Maddox & Bohil, 1998a, 1998b). The immediate feedback condition was contrasted with a *delayed feedback* condition, in which feedback was provided on every fifth trial. Importantly, the delayed feedback was based on aggregate performance across all five trials (i.e., the total reward earned across all five trials) and was not based only on performance for the fifth trial. This was important because we wanted to ensure that the amount of quantitative information was identical in both conditions. Delay procedures like this have been used occasionally to study such diverse topics as concept formation (Schroth, 1995, 1997), memory retention (Rankin & Trepper, 1978; Sassenrath, 1975), and training procedures (Duker, Hensgens, & Venderbosch, 1995; Reid & Parsons, 1996) and to test models of decision making (Bussemeyer & Myung, 1992).

Second, we manipulated the nature of the corrective feedback. In the *objective classifier* condition, we provided information about the reward that could have been earned on that trial (or on the last five trials) had the objectively correct response been given. Again, this is the approach taken in most studies of cost-benefit learning (e.g., Healy & Kubovy, 1981; Maddox, 1995; Maddox & Bohil, 1998a, 1998b). In the *optimal classifier* condition, we provided information about the reward that was earned by the optimal classifier on that trial (or on the last five trials). Importantly, on a certain proportion of trials, the optimal classifier will respond incorrectly and will earn no reward. However, the objective classifier, by definition, will never respond incorrectly and thus will always earn

Table 1
Points and Accuracy (per 120-Trial Block) for the
Optimal Classifier That Maximizes Long-Run Reward

| Condition | $d' = 1.0$ | | $d' = 2.2$ | |
|------------------|------------|----------|------------|----------|
| | Points | Accuracy | Points | Accuracy |
| Baseline | 166 | 69% | 206 | 86% |
| 3:1 Cost-Benefit | 186 | 61% | 212 | 83% |

a reward. Optimal classifier feedback should help the observers realize that they need not attempt to respond accurately on every trial. Combining the *delay* factor with the *nature of the feedback* factor yielded four feedback conditions: *immediate/objective classifier*, *immediate/optimal classifier*, *delay/objective classifier*, and *delay/optimal classifier*.

Each observer completed a single session of categorization with a 3:1 cost-benefit ratio. Each observer participated in only one of the eight experimental conditions. Each session consisted of three distinct phases: baseline, training, and transfer. In the baseline phase, the observer completed a minimum of 120 categorization trials in which no cost-benefit manipulation was present. If the observer reached a predetermined performance criterion in 120 trials, the training phase began. Otherwise, additional baseline trials were presented until the criterion was met. The baseline condition was included to ensure that the observers had accurate knowledge of the category distributions prior to any cost-benefit manipulations. The training phase consisted of three 120-trial blocks of categorization with a 3:1 cost-benefit ratio. The feedback manipulations were instantiated during the training phase. The transfer phase consisted of a single 120-trial block of categorization, in which no feedback was provided on any trial.

Method

Observers. One hundred and forty undergraduate students at the University of Texas, Austin participated in this experiment in partial fulfillment of a class requirement. There were 18 observers in each of the four $d' = 1.0$ conditions and 17 observers in each of the four $d' = 2.2$ conditions.

Stimuli and stimulus generation. The stimulus was a filled white rectangular bar (30 pixels wide) set flush upon a stationary base (40 pixels wide) that was centered on the computer monitor. The height of the bar varied from trial to trial. There were two categories, A and B, each defined by a specific univariate normal distribution (Ashby & Gott, 1988). The separation between the Category A and the Category B means were 42 and 91 pixels for $d' = 1.0$ and 2.2, respectively. The Category A and Category B standard deviation was 42 for both d' levels. For each level of d' , 120 unique stimuli were constructed by taking 60 random samples from each category distribution. The orders were randomized separately for each of the four blocks of trials.

Procedure. The observers were told that perfect performance was impossible. However, an optimal level of performance was specified as a goal (in the form of desired point totals). The observers were told that they were participating in a hypothetical medical diagnosis task and that the height of the bar represented the results of a particular medical test. The test was designed to distinguish between two diseases, such as "burlosis" and "namitis," hereafter referred to simply as "A" and "B." The observers were in-

formed that they would receive the medical test result for a new patient on each trial and that their goal was to maximize points in each session. The observers were told not to worry about speed of responding and that the stimulus would remain on the screen until they responded. Before being exposed to the cost-benefit manipulation, each observer completed a minimum of 120 baseline trials, in which the cost of an incorrect response was zero (i.e., $V_{aB} = V_{bA} = 0$) and the benefit of a correct response was 2 points (i.e., $V_{aA} = V_{bB} = 2$). If the observer reached an accuracy-based performance criterion (2% below optimal), he or she was allowed to begin the experimental condition. If the observer did not reach criterion, he or she continued in the baseline condition until criterion was reached. Once the observer reached criterion, he or she was allowed to begin the experimental condition. In the experimental condition, the benefit of a correct "A" response was 3 points (i.e., $V_{aA} = 3$), the benefit of a correct "B" response was 1 point (i.e., $V_{bB} = 1$), and the cost of an incorrect response was zero (i.e., $V_{aB} = V_{bA} = 0$). Table 1 displays the point totals and accuracy rates for the optimal classifier for each d' condition.

During each trial of the training phase, a stimulus was presented and remained on the screen until the observer responded. On feedback trials, the feedback remained on the screen for 750 msec. The feedback was followed by a 125-msec intertrial interval, during which the screen was blank. On no-feedback trials, the screen remained blank for 875 msec. In the immediate feedback condition, feedback was provided on each trial. In the delayed feedback condition, feedback was provided following every fifth trial, and the feedback was based on aggregate performance across all five trials. In all the conditions, the feedback display contained four lines of information. The top line presented the points earned on the previous trial (in the immediate feedback condition) or on the previous five trials (in the delayed feedback condition). In the objective classifier condition, the second line presented the number of points that could have been earned for a correct response on the previous trial (in the immediate feedback condition) or for correct responses on the previous five trials (in the delayed feedback condition). In the optimal classifier condition, the second line presented the number of points that were earned by the optimal classifier on the previous trial (in the immediate feedback condition) or on the previous five trials (in the delayed feedback condition). The third line presented a running cumulative point total for the observer. The fourth line presented a running cumulative total for either the objective or the optimal classifier, depending on condition. Sample

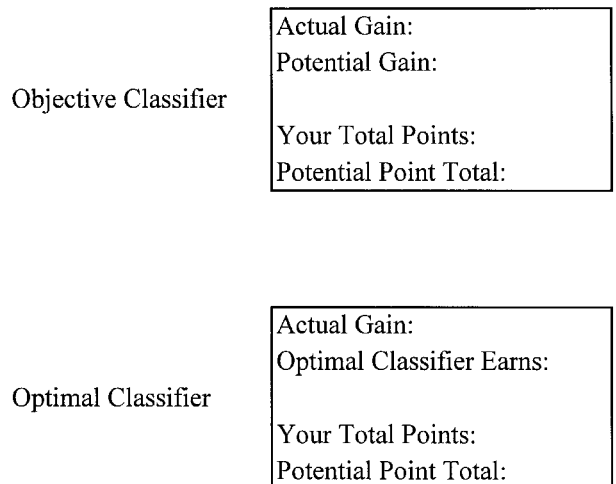


Figure 4. Hypothetical feedback displays for the objective classifier and optimal classifier conditions.

Table 2
Analysis of Variance Results
for Experiment 1 for the Three Performance Indices

| Source | % of Optimal Points | % of Optimal Accuracy | Deviation from Optimal Decision Criterion |
|-----------------------------------|---------------------|-----------------------|---|
| d' | * | * | * |
| Delay | n.s. | n.s. | n.s. |
| Nature of feedback | n.s. | n.s. | n.s. |
| Block | * | * | * |
| $d' \times$ delay | n.s. | n.s. | n.s. |
| $d' \times$ nature of feedback | n.s. | n.s. | n.s. |
| $d' \times$ block | .075 | n.s. | n.s. |
| Delay \times nature of feedback | n.s. | n.s. | n.s. |
| Delay \times block | n.s. | n.s. | n.s. |
| Nature of feedback \times block | .087 | n.s. | * |

* $p < .05$.

feedback displays for the objective and the optimal classifier conditions are depicted in the top and bottom panels of Figure 4, respectively.

Results

All analyses were performed on the three blocks of training data and on the transfer data. In comparing human performance with that of the optimal classifier, a number of performance indices can be examined. In this article, we took a *converging operations* approach (Garner, 1974). Specifically, we examined several different, but related, performance indices and focused on patterns of results that converged across indices. Because the optimal classifier attempts to maximize long-run reward, it was of interest to compare each observer's reward with that of the optimal classifier. In our experiment, reward came in the form of points that were earned on each trial, depending on the observer's response (see Table 1). These point totals accrued across trials, and at the end of the experiment, the observer received monetary payment that was proportional to their point totals. Because the optimal point totals differed across d' conditions, it was important to standardize the observer's point totals across conditions. Thus, we computed the percentage of optimal points as follows:

% of optimal points

$$= 100(\text{observed point total}/\text{optimal point total}).$$

An observer whose point total was less than that predicted by the optimal classifier would yield a percentage of optimal point value that was less than 100, whereas performance that was superior to that of the optimal classifier would yield a percentage of optimal point value that was greater than 100.

Another measure of interest was to compare each observer's accuracy rate with that predicted by the optimal classifier. Again, we standardized and computed the percentage of optimal accuracy as follows:

% of optimal accuracy

$$= 100(\text{observed accuracy}/\text{optimal accuracy}).$$

An observer whose accuracy rate was less than that predicted by the optimal classifier would yield a percentage of optimal accuracy value that was less than 100, whereas an observer whose accuracy rate was greater than that predicted by the optimal classifier would yield a percentage of optimal accuracy value that was greater than 100. This measure would be especially informative, because the optimal classifier sacrifices accuracy in order to maximize long-run reward. An observer who placed some importance on accuracy maximization would likely perform more accurately than the optimal classifier.

In addition, we examined the observer's decision criterion estimate, $k = \ln(\beta)/d'$ from signal detection theory (Green & Swets, 1966; Macmillan & Creelman, 1991; von Winterfeldt & Edwards, 1982). Because we were interested in comparing the observer's decision criterion with that of the optimal classifier, we computed the deviation from the optimal decision criterion as follows:

$$\text{deviation from optimal decision criterion} = k - k_0.$$

For each of the three measures, we start with a 2 (d') \times 2 (delay) \times 2 (nature of the feedback) \times 4 (blocks) mixed design analysis of variance (ANOVA) in which only the block factor was within-subjects.³ Table 2 summarizes the ANOVA results for each of the three performance measures for the main effects and two-way interactions. All the three-way interactions and the four-way interaction were nonsignificant for all three measures. All effects that were statistically significant at $p < .05$ are denoted with an asterisk (*). Marginal significance levels (i.e., $.05 < p < .10$) are denoted with the exact significance level. Nonsignificant effects (i.e., $p > .10$) are denoted with n.s. for nonsignificant. The main effects for category discriminability and block were significant for all three measures, and the nature of feedback \times block interaction was significant for two of the three measures. All other effects were significant for one or no performance index.

Figures 5–7 display the main effects of category discriminability and block and the nature of feedback \times block interaction, respectively, for each of the three per-

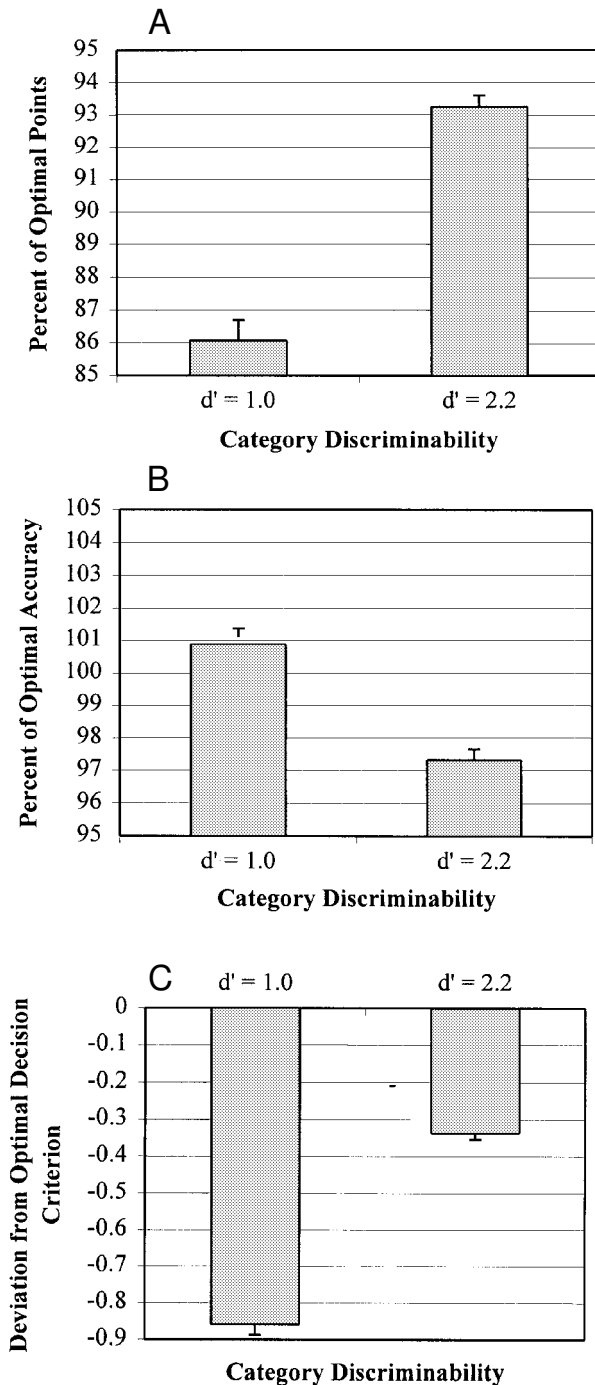


Figure 5. (A) Percentage of optimal points, (B) percentage of optimal accuracy, and (C) deviation from optimal decision criterion, averaged across observers, for the two category discriminabilities from Experiment 1. Standard error bars are included.

formance measures. Several comments are in order. First, note that performance was closer to optimal for $d' = 2.2$ than for $d' = 1.0$ for the point and decision criterion measures, as would be predicted from the flat-maxima hy-

pothesis. In addition, note that the percentage of optimal accuracy was larger for $d' = 1.0$ than for $d' = 2.2$ and was, in fact, greater than 100% (see Figure 5). Recall that the sacrifice in accuracy necessary to maximize long-run reward is larger for $d' = 1.0$ than for $d' = 2.2$. On the basis of COBRA, we hypothesized earlier that this larger accuracy sacrifice might make observers place more weight on accuracy maximization in the $d' = 1.0$ condition, thus leading to less optimal performance. This hypothesis is supported by the data. Second, note that performance generally improved over blocks for the three performance measures (see Figure 6). Finally, note that the nature of the feedback effect increased across blocks for all performance measures (Figure 7) but that the effect was largest for the point and decision criterion measures. Specifically, performance was closer to optimal for optimal classifier feedback, as compared with objective classifier feedback, and the magnitude of this effect increased across blocks. Post hoc *t* tests were conducted separately for each block to determine when the feedback effect was significant. For both the point and the decision criterion measures, the effect was nonsignificant during Training Blocks 1 and 2 ($p > .05$) and was significant during transfer Block 4 ($p < .05$). During Block 3, the effect was marginal for the point measure ($p = .061$) and was nonsignificant for the decision criterion measure ($p > .05$).

Discussion

In Experiment 1, two modifications to the feedback display were combined factorially with two levels of category discriminability in order to study their effects on decision criterion learning when the cost-benefit difference was unequal. Performance improved with experience, and higher category discriminability led to better learning. These findings converge with those from a number of published studies (Bohil & Maddox, 2001; Busemeyer & Myung, 1992; Healy & Kubovy, 1981; Maddox, 1995; Maddox & Bohil, 1998a, 1998b, 2000). Although the delay and the nature of the feedback manipulations did not reach statistical significance, the condition means suggested that delayed feedback and optimal feedback improved performance (for delayed vs. immediate feedback, the percentages of optimal points was 90.0% vs. 89.3%; the percentage of optimal accuracy was 99.4% vs. 98.8%, and the deviation from optimal decision criterion was $-.58$ vs. $-.62$; for optimal vs. objective feedback, the percentage of optimal points was 90.6% vs. 88.7%; the percentage of optimal accuracy, 99.5% vs. 98.7%, and the deviation from optimal decision criterion was $-.56$ vs. $-.64$). The feedback manipulation did interact with the level of experience with the task, revealing no feedback effect early in learning, but a large optimal feedback advantage late in learning (see Figure 7).

These data represent an important starting point. They suggest that performance can be improved when the observer is trained relative to the optimal classifier, as opposed to being trained relative to the objectively correct

EXPERIMENT 2




Figure 6. (A) Percentage of optimal points, (B) percentage of optimal accuracy, and (C) deviation from optimal decision criterion, averaged across observers, for the four blocks of trials from Experiment 1. Standard error bars are included.

response. Even so, these data are not rich enough to allow an examination of the psychological processes that might account for this feedback effect. This determination requires a model-based approach, which is the focus of Experiment 2.

Experiment 2 used a within-subjects design to rigorously examine the effects of category discriminability and feedback manipulations on cost-benefit learning. Because the delay manipulation did not reach statistical significance in Experiment 1, it was excluded from Experiment 2, and immediate feedback was provided. Thus, each observer completed four experimental conditions constructed from the factorial combination of two d' levels (1.0 and 2.2) with two feedback types (objective classifier and optimal classifier).

Method

Observers. Eight observers were solicited from the University of Texas community. Each observer completed four approximately 30-min sessions. The observers were paid on the basis of their day-to-day performance in the task.

Stimuli and stimulus generation. The stimuli and category structure were identical to those in Experiment 1.

Procedure. The procedures were identical to those in Experiment 1, with the following exceptions. First, a Latin-square design was used to counterbalance the order of the four conditions across observers. Second, different disease labels were used for each condition and were assigned randomly across conditions.

Results

For completeness, we briefly outline ANOVA results for the same performance measures as those used in Experiment 1. A 2 (d') \times 2 (feedback type) \times 4 (block) within-subjects ANOVA was performed separately on each of the three performance measures.⁴ Table 3 summarizes the ANOVA results for each of the three performance measures. The main effects of feedback type and block were significant (or marginally significant) for two of the three measures. In addition, the $d' \times$ feedback type interaction was significant for two of the three measures.

Figures 8–10 display the main effects of feedback type and block, and the $d' \times$ feedback type interaction, respectively, for each of the three performance measures. The results can be summarized as follows. First, performance was closer to optimal for the optimal classifier feedback than for the objective classifier feedback for all three measures, although the effect was much smaller for the accuracy measure (see Figure 8). Second, performance became more nearly optimal across blocks for the point and decision criterion measures but stayed nearly constant for the accuracy measure (see Figure 9). Finally, the $d' \times$ feedback interaction was due to a much larger effect of objective versus optimal classifier feedback on performance for the $d' = 1.0$ condition than for the $d' = 2.2$ condition. Post hoc analyses revealed that optimal classifier feedback yielded significantly better performance than objective classifier feedback for the $d' = 1.0$ condition ($p < .05$ for both the point and the decision criterion measures), but not for the $d' = 2.2$ condition ($p > .05$ for both measures).

Interestingly, the d' effect on decision criterion placement was nonsignificant. Even so, the decision criterion was closer to optimal for $d' = 2.2$ than for $d' = 1.0$ (av-

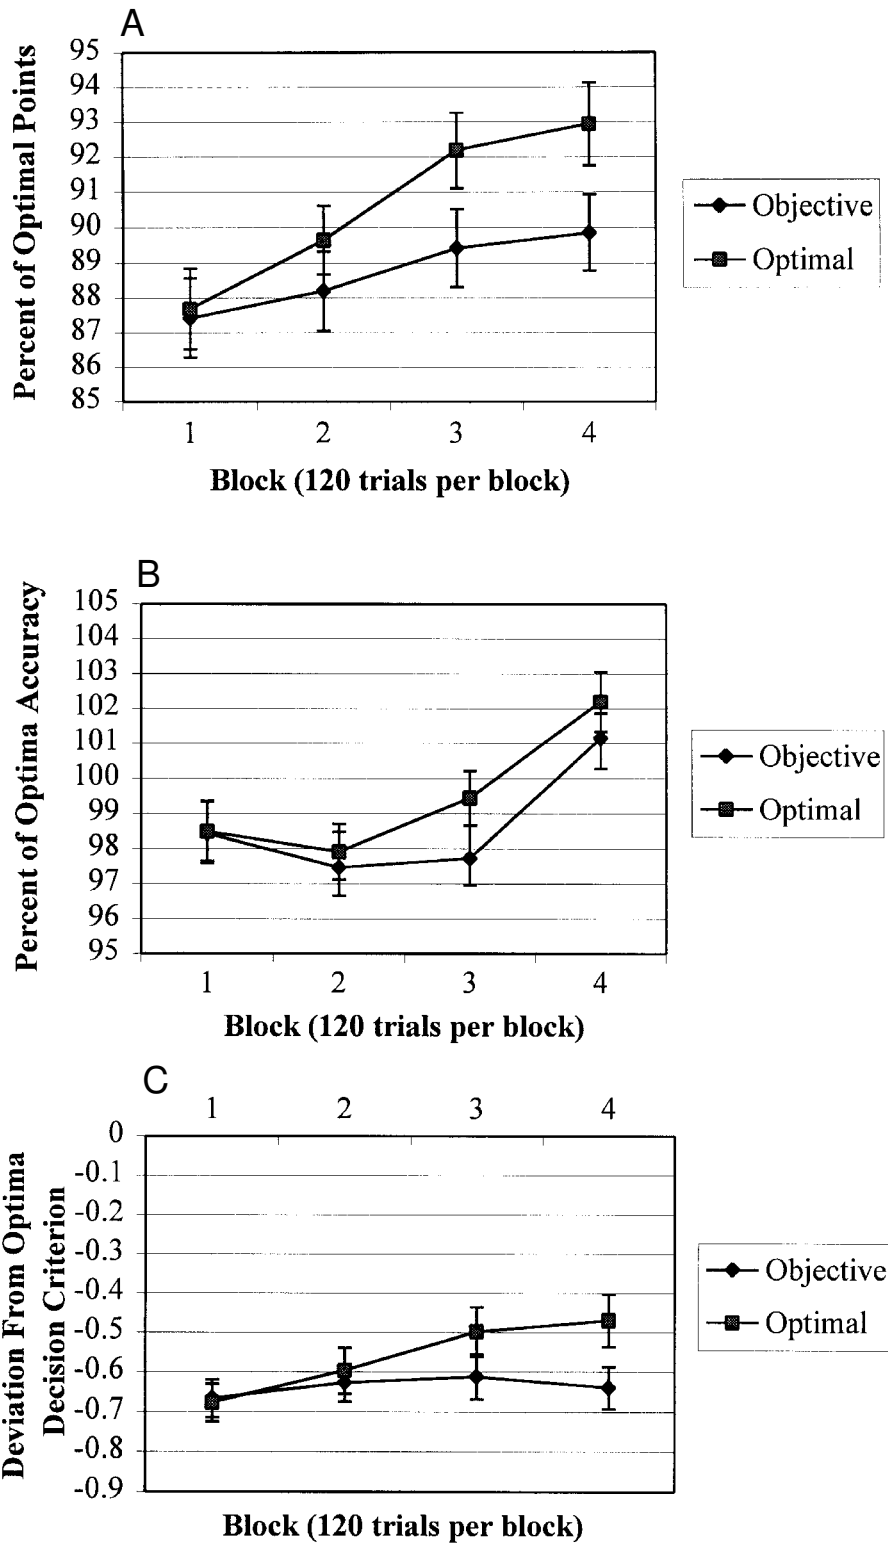


Figure 7. (A) Percentage of optimal points, (B) percentage of optimal accuracy, and (C) deviation from optimal decision criterion, averaged across observers, for the objective and optimal classifier feedback conditions by block from Experiment 1. Standard error bars are included.

Table 3
Analysis of Variance Results
for Experiment 2 for the Three Performance Indices

| Source | % of Optimal Points | % of Optimal Accuracy | Deviation From Optimal Decision Criterion |
|-----------------------------------|------------------------|--------------------------|---|
| d' | n.s. | .063 | n.s. |
| Nature of feedback | * | n.s. | * |
| Block | .060 | n.s. | .057 |
| $d' \times$ nature of feedback | * | n.s. | * |
| $d' \times$ block | .054 | n.s. | n.s. |
| Nature of feedback \times block | n.s. | * | n.s. |

* $p < .05$.

erage deviation from optimal decision criterion = $-.252$, and $-.344$ for $d' = 2.2$, and 1.0 , respectively), and more important, the same relation held at the individual observer level in 42 of 64 cases (8 observers \times 2 feedback types \times 4 blocks), providing initial support for the flat-maxima hypothesis.

THEORETICAL ANALYSES

To this point, we have focused on ANOVA-based results. Although fruitful, this approach is limited for at least two reasons. First, because an ANOVA focuses on “average” performance, performance profiles at the individual-observer level cannot be examined. Second, because an ANOVA determines only whether means are statistically equal, it does not allow one to test rigorously a number of important hypotheses. The focus of this section is on the development and application of a series of decision bound models. We begin with a brief overview of decision bound theory that provides the underlying framework for our modeling endeavor. The theory is described in detail in numerous articles (e.g., Ashby, 1992a; Ashby & Perin, 1988; Ashby & Townsend, 1986; Maddox & Ashby, 1993; Maddox & Bohil, 1998a, 1998b). We then outline a number of hypotheses to be tested, instantiate each within the framework of a decision bound model, and summarize the results of the model-based analyses. All analyses were performed at the individual-observer level, because of concerns with modeling aggregate data (e.g., Ashby, Maddox, & Lee, 1994; Estes, 1950; Maddox, 1999; Maddox & Ashby, 1998; Smith & Minda, 1998).

Decision Bound Theory

Decision bound theory assumes that the observer attempts to respond optimally but is unable to because two suboptimalities, perceptual noise and criterial noise, are inherent in all humans (and other organisms). Perceptual noise exists because there is trial-by-trial variability in the perceptual information associated with each stimulus. Criterial noise exists because there is trial-by-trial variability in the observer’s memory for the decision criterion. Because perceptual and criterial noise exist, the human observer cannot attain the level of performance reached by the optimal classifier (i.e., cannot maximize long-run reward). Even so, decision bound theory assumes that

the observer attempts to use the same *strategy* as the optimal classifier, but with less success owing to the effects of perceptual and criterial noise. Besides perceptual and criterial noise, other suboptimalities might exist. For example, suboptimalities might exist in category distribution knowledge. All of the models tested in this article assume that the observer had knowledge of the category structures. This was an important assumption because our interest is in studying observers’ decision criterion learning, and not potential suboptimalities in category distribution knowledge. To ensure that this was a reasonable assumption, the first block of trials consisted of baseline trials, in which no cost–benefit manipulation was present. The observers had to reach a rigid performance criterion during this block before shifting to the experimental trials, thus ensuring that the observer did have accurate knowledge of the category distributions (see the Method sections).

Suboptimalities might also exist in knowledge of the cost–benefits and, thus, in the placement of the decision criterion, k . All of the models tested in this section allow for suboptimalities in the decision criterion placement. The models differ only in the types of constraints imposed on these suboptimalities across category discriminability and feedback conditions. Each model was applied simultaneously to the data from all four experimental conditions, separately by block and observer. During training and transfer, each block consisted of 120 experimental trials, and the observer was required to respond “A” or “B” for each stimulus. Since there were four conditions, each model was fit to a total of 960 estimated response probabilities from each training block (120 trials \times 2 response types [“A” or “B”] \times 4 conditions). The model yielded predicted probabilities of responding “A” and “B” by solving the following equations:

$$P(R_A | x_i) = P[h(x_{pi}) < k | x_i]$$

and

$$P(R_B | x_i) = 1 - P[h(x_{pi}) < k | x_i],$$

where x_i is the bar height for stimulus i , x_{pi} is the perceptual effect for stimulus i under the assumption that normally distributed perceptual noise exists, h is the decision function, and k is the decision criterion (see Equation 4). Decision bound theory also postulates noise in

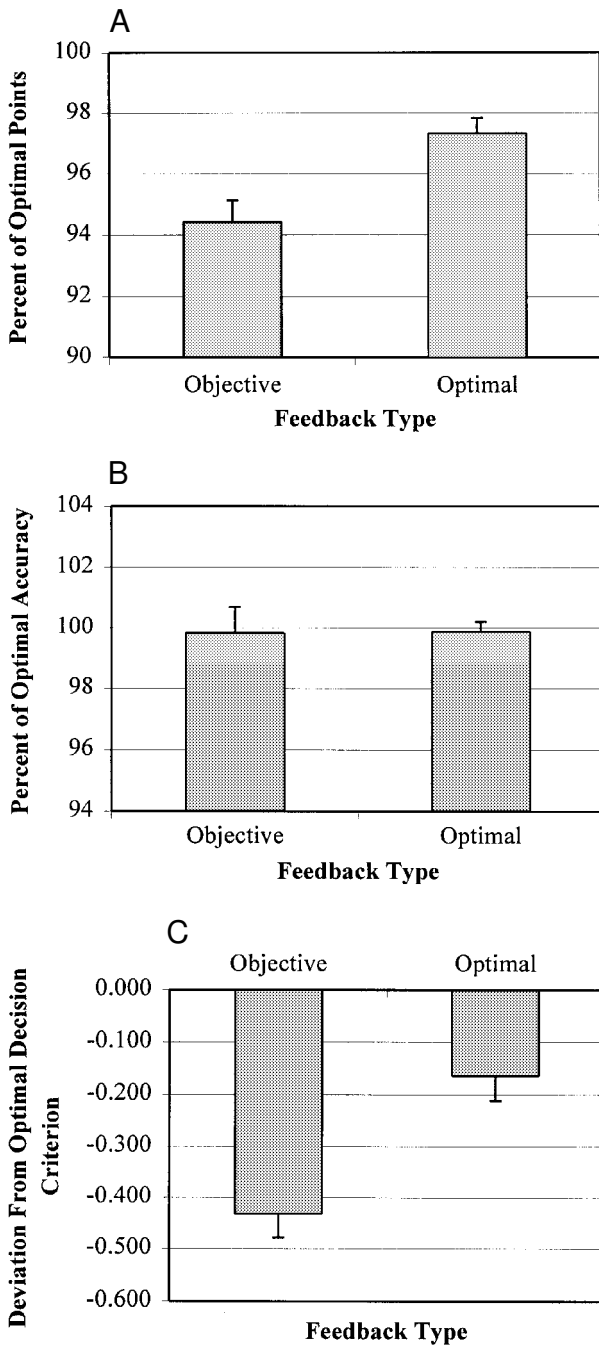


Figure 8. (A) Percentage of optimal points, (B) percentage of optimal accuracy, and (C) deviation from optimal decision criterion, averaged across observers, for the objective and optimal classifier feedback conditions from Experiment 2. Standard error bars are included.

the decision process—that is, criterial noise. However, in the present application, criterial noise is nonidentifiable with perceptual noise. All of the models outlined below included two *noise* parameters, one for each level of d' . The noise parameter represented the sum of perceptual and criterial noise (Ashby, 1992a; Maddox & Ashby, 1993).

The goal of the model-based analyses was to shed some light on the psychological processes that led to differences in decision criterion placement across the d' and feedback conditions. In particular, we were interested in determining whether variants of the flat-maxima and COBRA hypotheses could account for the data. To facilitate the development of each model, consider the following equation that determines the decision criterion used by the observer (k):

$$k = wk_a + (1 - w)k_r, \tag{4}$$

where k_a is the decision criterion that maximizes expected accuracy (i.e., the equal likelihood decision criterion), k_r is the decision criterion used by the observer to maximize expected reward, and w is the importance (or weight) given to expected accuracy maximization (see Figure 3). We began by developing four models, each of which makes different assumptions about the k_r and w values. The nested structure of the models is presented in Figure 11. The number of free parameters (in addition to the two noise parameters described above) is presented in parentheses. The arrows point to the more general model. Models at the same level have the same number of free parameters.

The *flat-maxima(stp) model* instantiates the flat-maxima hypothesis, but not the COBRA hypothesis. Specifically, it assumes that the decision criterion used by the observer to maximize expected reward (k_r) is determined by the steepness of the ORF. A single steepness parameter (stp) is estimated from the data. This single steepness parameter determines a distinct decision criterion value for each of the d' conditions (see Figure 2B). Because the decision criterion values are determined from the ORF, this model is constrained to predict that the decision criterion will be closest to optimal in the $d' = 2.2$ condition, relative to the $d' = 1.0$ condition. In addition, this model assumes that there is no effect of optimal versus objective feedback, thus assuming the same decision criterion in both cases. Finally, this model assumes that there is no competition between accuracy and reward maximization (i.e., $w = 0$).

The *COBRA(w) model* instantiates the COBRA hypothesis, but not the flat-maxima hypothesis. Specifically, it assumes that the decision criterion used by the observer to maximize expected reward is the optimal decision criterion (i.e., $k_r = k_o$) but allows for a competition between reward and accuracy maximization by estimating the Equation 4 w parameter from the data. This model assumes no effect of optimal versus objective feedback and estimates a single w parameter.

The *flat-maxima($stp_{opt};stp_{obj}$) model* generalizes the flat-maxima(stp) model by assuming that the decision criterion used by the observer to maximize expected reward (k_r) is different in the optimal and the objective feedback conditions. This model is constrained to predict that the decision criterion will be closest to optimal in the $d' = 2.2$ condition, relative to the $d' = 1.0$ condition, but it allows for differences across feedback conditions by estimating

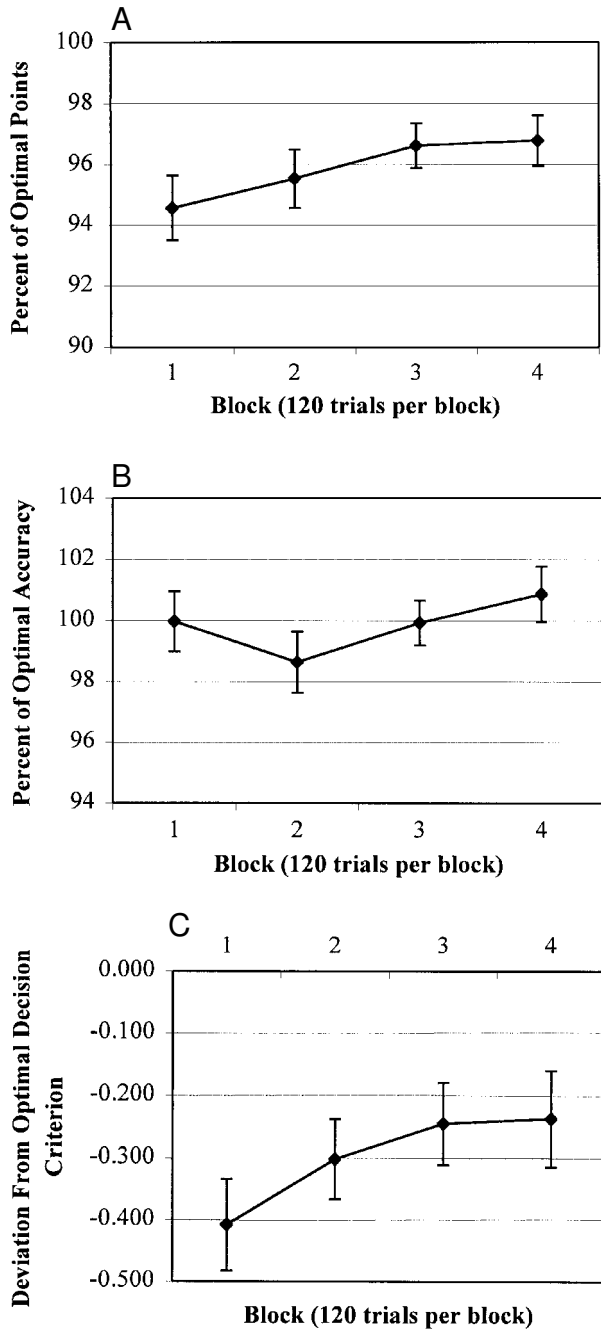


Figure 9. (A) Percentage of optimal points, (B) percentage of optimal accuracy, and (C) deviation from optimal decision criterion, averaged across observers, for the four blocks of trials from Experiment 2. Standard error bars are included.

two steepness parameters (stp_{opt} ; stp_{obj}). This model assumes no competition between reward and accuracy (i.e., $w = 0$) and contains the flat-maxima(stp) model as a special case.

The *COBRA*(w_{opt} ; w_{obj}) model generalizes the *COBRA*(w) model by assuming that the weight placed on accuracy maximization is different in the optimal (w_{opt})

and the objective (w_{obj}) feedback conditions. It does not instantiate the flat-maxima hypothesis and, instead, assumes that the decision criterion used by the observer to maximize expected reward is the optimal decision criterion (i.e., $k_r = k_o$). This model contains the *COBRA*(w) model as a special case.

Each of these models has problems from a psychological standpoint. For example, the flat-maxima(stp) and *COBRA*(w) models both assumed no feedback effect, when one was observed. The *COBRA*(w_{opt} ; w_{obj}) model has the potential to account for the feedback effect by assuming different accuracy weights across feedback conditions, but the model has no way to account for the d' effect. The flat-maxima(stp_{opt} ; stp_{obj}) model appears most reasonable on the surface, since it predicts more nearly optimal decision criterion placement for $d' = 2.2$ relative to $d' = 1.0$, as was observed, and it has the potential to account for the feedback effect by assuming that the nature of the feedback affects the steepness associated with the stopping point on the ORF. Even so, the flat-maxima hypothesis derives from the nature of the objective reward function, which is best thought of as a hypothesis about how the observer learns the decision criterion that maximizes long-run reward, k_r . It is unclear why the location of the stopping point on the ORF would be affected by the nature of the feedback.

In light of these problems, we developed and tested two additional *hybrid* models that instantiated simultaneously the important components of both the flat-maxima and the *COBRA* hypotheses. Both hybrid models instantiated the flat-maxima hypothesis by estimating a single steepness parameter (stp) that determined the decision criterion used by the observer to maximize expected reward (k_r). The two models differed only in their assumptions regarding *COBRA*.

The *hybrid*(stp ; w) model allows for a competition between reward and accuracy maximization by estimating the Equation 4 w parameter from the data. However, this model assumes that there is no effect of optimal versus objective feedback on the value of the w parameter. Note that this model contains both the flat-maxima(stp) model and the *COBRA*(w) model as a special case (see Figure 11).

The *hybrid*(stp ; w_{opt} ; w_{obj}) model generalizes the *hybrid*(stp ; w) model by estimating separate accuracy weights for the optimal and the objective feedback conditions. This model contains the *hybrid*(stp ; w), flat-maxima(stp), *COBRA*(w), and *COBRA*(w_{opt} ; w_{obj}) models as special cases (see Figure 11).

The model parameters were estimated using maximum likelihood (Ashby, 1992b; Wickens, 1982) and the goodness-of-fit statistic was

$$AIC = 2r - 2\ln L,$$

where r is the number of free parameters and L is the likelihood of the model given the data (Akaike, 1974; Takane & Shibayama, 1992). The AIC statistic penalizes a model for extra free parameters in such a way that the

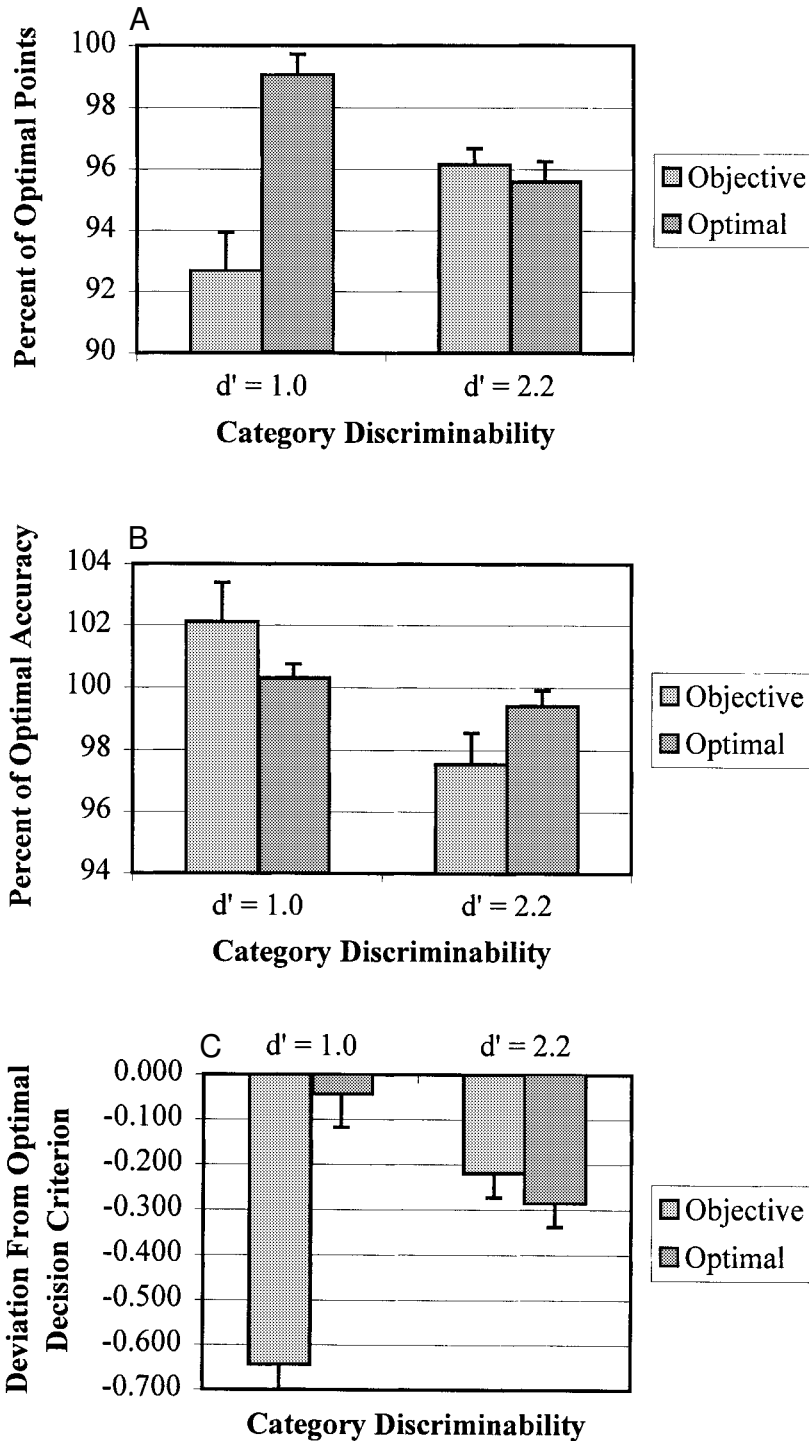


Figure 10. (A) Percentage of optimal points, (B) percentage of optimal accuracy, and (C) deviation from optimal decision criterion, averaged across observers, for the objective and optimal classifier feedback conditions by category discriminability from Experiment 2. Standard error bars are included.

smaller the AIC, the closer a model is to the “true model,” regardless of the number of free parameters. Thus, to find the best model among a given set of competitors, one simply computes an AIC value for each model and chooses

the model associated with the smallest AIC value. Table 4 presents the best-fitting model by observer and block. Several comments are in order. First, during the first two blocks of trials, the flat-maxima models provided the best




Figure 11. Nested relationship among the decision bound models applied simultaneously to the data from all the experimental conditions. The arrows point to a more general model (see the text for details). Note: All models assume two noise parameters.

account of the data in 9 of 16 (8 observers \times 2 blocks) cases. Both the hybrid and the COBRA models performed poorly, accounting best for 4 and 3 data sets, respectively. Second, during the last two blocks of trials the hybrid models dominated, accounting for 12 of 16 cases. Importantly, this was due mainly to the superiority of the Hybrid(stp; w_{opt} ; w_{obj}) model, which provided the best account for 4 of 8 observers during Block 3 and 6 of 8 observers during the final block. Third, a common pattern within observers (e.g., Observers 1, 3, 7, and 8) was for the flat-maxima(stp_{opt}; stp_{obj}) model to perform best early in learning and for the hybrid(stp; w_{opt} ; w_{obj}) model to perform best late in learning. Finally, the COBRA models generally performed poorly. Taken together, the model-based analyses suggest that the flat-maxima model in isolation can account reasonably well for early learning. However, a hybrid model that assumes that the reward-maximizing decision criterion is driven by the flat-maxima hypothesis and that COBRA accounts for the differential weight placed on accuracy maximization in objective and optimal classifier feedback conditions was necessary to account for decision criterion learning with experience.

Because the models were applied separately to each block of trials, we can examine the parameter values to determine how they changed as the observers gained experience with the task. Of particular interest is to examine the accuracy weights from the hybrid(stp; w_{opt} ; w_{obj}) model. Figure 12 plots the w_{opt} and w_{obj} values for each block averaged across observers. A two-way ANOVA was conducted on the accuracy weights. The main effect of feedback condition was significant [$F(1,7) = 11.535$, $p < .05$], whereas the main effect of block and the feedback \times block interaction were both nonsignificant ($p > .05$). Although the interaction was nonsignificant, we decided to conduct separate one-way ANOVAs on the accuracy weights for the two feedback conditions. In the objective classifier condition, the block effect was nonsignificant ($p > .05$), suggesting that the weight allo-

cated to accuracy maximization was relatively constant across learning, although the general trend was toward increasing weights. In the optimal classifier condition, on the other hand, the block effect was marginally significant [$F(3,21) = 2.775$, $p = .067$]. Post hoc analyses revealed that the weight declined significantly from Blocks 1 to 2, then remained relatively constant.

GENERAL DISCUSSION

This article reports the results of two experiments that examined the effects of different types of feedback on cost–benefit learning in a perceptual categorization task. A major focus was to identify feedback conditions that would increase the salience of information important for reward maximization and decrease the salience of information important for accuracy maximization. Because some measure of accuracy must be sacrificed to maximize reward when costs and benefits are unequal, feedback conditions that meet these requirements should lead to superior performance. In Experiment 1, two types of feedback manipulations, immediate/delay and objective/optimal classifier, were combined factorially with two levels of category discriminability ($d' = 1.0$ or 2.2). In the immediate feedback conditions, feedback was provided on each trial, whereas feedback was provided following every fifth trial in the delayed feedback conditions. In the objective classifier conditions, the feedback contained information regarding the number of points earned for an “objectively” correct response. In the optimal classifier conditions, the feedback contained information regarding the number of points earned by the optimal classifier. Importantly, the optimal classifier makes errors but makes more errors for the low cost–benefit category, thus using a decision strategy that maximizes long-run reward. Each observer completed several blocks of trials in one of the eight resulting conditions. Accuracy rates, point totals, and the theoretically motivated decision crite-

Table 4
Best-Fitting Model by Observer and Block

| Observer | Block | | | |
|----------|---|---|---|---|
| | 1 | 2 | 3 | 4 |
| 1 | flat-maxima (stp _{opt} ; stp _{obj}) | flat-maxima (stp) | hybrid (stp; w) | hybrid (stp; w _{opt} ; w _{obj}) |
| 2 | hybrid (stp; w) | COBRA (w) | hybrid (stp; w _{opt} ; w _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) |
| 3 | flat-maxima (stp _{opt} ; stp _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) | hybrid (stp; w) | hybrid (stp; w _{opt} ; w _{obj}) |
| 4 | hybrid (stp; w _{opt} ; w _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) |
| 5 | COBRA (w) | COBRA (w) | flat-maxima (stp _{opt} ; stp _{obj}) | flat-maxima (stp _{opt} ; stp _{obj}) |
| 6 | flat-maxima (stp) | flat-maxima (stp) | flat-maxima (stp) | COBRA (w _{opt} ; w _{obj}) |
| 7 | flat-maxima (stp _{opt} ; stp _{obj}) | flat-maxima (stp _{opt} ; stp _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) |
| 8 | flat-maxima (stp _{opt} ; stp _{obj}) | flat-maxima (stp _{opt} ; stp _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) | hybrid (stp; w _{opt} ; w _{obj}) |

tion index from signal detection theory (Green & Swets, 1966; Macmillan & Creelman, 1991) were examined with an ANOVA. Performance improved across blocks and was superior for more discriminable categories. Performance was generally superior in the delayed feedback conditions and in the optimal classifier conditions, but neither effect reached statistical significance. However, the nature of the feedback (objective vs. optimal classifier) interacted with the level of experience. Specifically, whereas performance was nearly identical for the objective and the optimal classifier feedback conditions during the first block of trials, as the observer gained experience with the

task, performance diverged for the two conditions, with superior performance being shown in the optimal classifier conditions (see Figure 7).

In Experiment 2, the nature of the feedback (objective vs. optimal classifier) was combined factorially with two levels of category discriminability ($d' = 1.0$ or 2.2), and each observer completed several blocks of trials in all four conditions. In line with the results from Experiment 1, performance improved over blocks and was superior when feedback was based on the optimal classifier. In addition, the feedback effect was much larger for the low category discriminability condition than for the high category dis-




Figure 12. Weight allocated to accuracy maximization for the objective and optimal classifier conditions averaged across observer from the hybrid (stp; w_{opt}; w_{obj}) model separately for each block of trials. Standard error bars are included.

criminability condition. This is reasonable, since the magnitude of the sacrifice in accuracy necessary to maximize reward increases as category discriminability decreases. Because a within-subjects design was utilized, a series of models, derived from decision bound theory, were applied simultaneously to the data from all four conditions separately for each observer and block. These model-based analyses allowed a fine-grained analysis of each observer's performance and allowed us to instantiate and test several hypotheses of theoretical importance. Two hypotheses were of particular importance. The first hypothesis, the flat-maxima hypothesis, suggests that the observer adjusts his or her decision criterion on the basis of the change in the rate of reward of the objective reward function. Because the ORF is steeper for $d' = 2.2$ than for $d' = 1.0$, the flat-maxima hypothesis predicts that the resulting decision criterion will be closer to the optimal value for $d' = 2.2$ than for $d' = 1.0$. The second hypothesis, the COBRA maximization hypothesis, suggests that observers place some importance (or weight) on reward and accuracy maximization. When both goals cannot be achieved simultaneously, as in unequal cost-benefit conditions, the competition leads to the use of a suboptimal decision criterion. Model-based variants of the flat-maxima and COBRA hypotheses, in isolation, and a hybrid model that instantiated important assumptions from both hypotheses were applied to the data in an attempt to account for decision criterion learning across d' and feedback conditions. The flat-maxima model provided a good account of performance early in learning, but a hybrid model that assumed that the reward-maximizing decision criterion was driven by the flat-maxima hypothesis and that the feedback effect was on the weight placed on accuracy maximization in COBRA was necessary to account for decision criterion learning with experience. The accuracy weights from the hybrid model indicated that the observer placed less importance on accuracy maximization when the feedback was based on the optimal classifier, as opposed to the objective classifier. Although more tentative, the importance placed on accuracy maximization remained fairly constant (or increased slightly) with learning in the objective feedback condition but showed a large decline early, followed by relative stability later in learning, in the optimal feedback condition (Figure 12). If we extrapolate beyond this single experimental session, the results support the general conclusion that traditional feedback (i.e., objective feedback) may never lead to optimal (or nearly optimal) levels of responding, whereas optimal classifier feedback yields excellent performance after only one session and may yield optimal (or nearly optimal) performance with only a few hundred additional trials.

This work has important implications for many types of real-world perceptual categorization or decision-making problems, such as medical diagnosis. This work suggests that training should not focus on correct trial-by-trial responding, something that even the optimal classifier can-

not attain, but rather should focus on training the novice categorizer to emulate the optimal classifier. In this way, the student can learn which types of error are less damaging to overall reward maximization and which are more damaging.

In conclusion, the present study provides useful information about the types of feedback that lead to superior cost-benefit learning. The results suggest that feedback based on the optimal classifier leads to better performance by decreasing the weight placed on accuracy maximization.

REFERENCES

- AKAIKE, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, **19**, 716-723.
- ASHBY, F. G. (1992a). Multidimensional models of categorization. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 449-484). Hillsdale, NJ: Erlbaum.
- ASHBY, F. G. (1992b). Multivariate probability distributions. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 1-34). Hillsdale, NJ: Erlbaum.
- ASHBY, F. G., ALFONSO-REESE, L. A., TURKEN, A. U., & WALDRON, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, **105**, 442-481.
- ASHBY, F. G., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.
- ASHBY, F. G., MADDOX, W. T., & LEE, W. W. (1994). On the dangers of averaging across subjects when using multidimensional scaling or the similarity-choice model. *Psychological Science*, **5**, 144-150.
- ASHBY, F. G., & PERRIN, N. A. (1988). Toward a unified theory of similarity and recognition. *Psychological Review*, **95**, 124-150.
- ASHBY, F. G., & TOWNSEND, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, **93**, 154-179.
- BOHIL, C. J., & MADDOX, W. T. (2001). Category discriminability, base-rate, and payoff effects in perceptual categorization. *Perception & Psychophysics*, **63**, 361-376.
- BUSEMEYER, J. R., & MYUNG, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General*, **121**, 177-194.
- BUSEMEYER, J. R., & RAPPAPORT, A. (1988). Psychological models of deferred decision making. *Journal of Mathematical Psychology*, **32**, 91-134.
- DUKER, P. C., HENSGENS, Y., & VENDERBOSCH, S. (1995). Effectiveness of delayed feedback on the accuracy of teaching communicative features to individuals with severe mental retardation. *Research in Developmental Disabilities*, **16**, 479-488.
- DUSOIR, A. E. (1980). Some evidence on additive learning models. *Perception & Psychophysics*, **27**, 163-175.
- EREV, I. (1998). Signal detection by human observers: A cutoff reinforcement learning model of categorization decisions under uncertainty. *Psychological Review*, **105**, 280-298.
- EREV, I., GOPHER, D., ITKIN, R., & GREENSHAN, Y. (1995). Toward a generalization of signal detection theory to N -person games: The example of two-person safety problem. *Journal of Mathematical Psychology*, **39**, 360-375.
- EREV, I., WALLSTEN, T. S., & BUDESCU, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, **101**, 519-527.
- ESTES, W. K. (1950). The problem of inference from curves based on group data. *Psychological Bulletin*, **53**, 134-140.
- GARNER, W. R. (1974). *The processing of information and structure*. New York: Wiley.
- GREEN, D. M., & SWETS, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- HEALY, A. F., & KUBOVY, M. (1981). Probability matching and the for-

- dation of conservative decision rules in a numerical analog of signal detection. *Journal of Experimental Psychology: Human Learning & Memory*, **7**, 344-354.
- KUBOVY, M., & HEALY, A. F. (1977). The decision rule in probabilistic categorization: What it is and how it is learned. *Journal of Experimental Psychology: General*, **106**, 427-466.
- KUBOVY, M., & HEALY, A. F. (1980). Process models of probabilistic categorization. In T. S. Wallsten (Ed.), *Cognitive processes in choice and decision behavior* (pp. 239-262). Hillsdale, NJ: Erlbaum.
- MACMILLAN, N. A., & CREELMAN, C. D. (1991). *Detection theory: A user's guide*. New York: Cambridge University Press.
- MADDOX, W. T. (1995). Base-rate effects in multidimensional perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 288-301.
- MADDOX, W. T. (1999). On the dangers of averaging across observers when comparing decision bound models and generalized context models of categorization. *Perception & Psychophysics*, **61**, 354-374.
- MADDOX, W. T., & ASHBY, F. G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics*, **53**, 49-70.
- MADDOX, W. T., & ASHBY, F. G. (1998). Selective attention and the formation of linear decision boundaries: Comment on McKinley and Nosofsky (1996). *Journal of Experimental Psychology: Human Perception & Performance*, **24**, 301-321.
- MADDOX, W. T., & BOHIL, C. J. (1998a). Base-rate and payoff effects in multidimensional perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 1459-1482.
- MADDOX, W. T., & BOHIL, C. J. (1998b). Overestimation of base-rate differences in complex perceptual categories. *Perception & Psychophysics*, **60**, 575-592.
- MADDOX, W. T., & BOHIL, C. J. (2000). Costs and benefits in perceptual categorization. *Memory & Cognition*, **28**, 597-615.
- MADDOX, W. T., & DODD, J. L. (in press). On the relation between base-rate and cost-benefit learning in simulated medical diagnosis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.
- MADDOX, W. T., & ESTES, W. K. (1996, August). *A dual process model of category learning*. Paper presented at the 31st Annual Meeting of the Society for Mathematical Psychology, University of North Carolina, Chapel Hill.
- RANKIN, R. J., & TREPPER, T. (1978). Retention and delay of feedback in a computer-assisted instructional task. *Journal of Experimental Education*, **46**, 67-70.
- REID, D. H., & PARSONS, M. B. (1996). A comparison of staff acceptability of immediate versus delayed verbal feedback in staff training. *Journal of Organizational Behavior Management*, **16**, 35-47.
- ROTH, A. E., & EREV, I. (1995). Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games & Economic Behavior*, **3**, 3-24.
- SASSENATH, J. M. (1975). Theory and results on feedback and retention. *Journal of Educational Psychology*, **67**, 894-899.
- SCHROTH, M. L. (1995). Variable delay of feedback procedures and subsequent concept formation transfer. *Journal of General Psychology*, **122**, 393-399.
- SCHROTH, M. L. (1997). The effects of different training conditions on transfer in concept formation. *Journal of General Psychology*, **124**, 157-165.
- SMITH, J. D., & MINDA, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 1411-1436.
- STEVENSON, M. K., BUSEMEYER, J. R., & NAYLOR, J. C. (1991). Judgment and decision-making theory. In M. D. Dunnette & L. M. Hough (Eds.), *Handbook of industrial and organizational psychology* (2nd ed., Vol. 1, pp. 283-374). Palo Alto, CA: Consulting Psychologists Press.
- TAKANE, Y., & SHIBAYAMA, T. (1992). Structures in stimulus identification data. In F. G. Ashby (Ed.), *Probabilistic multidimensional models of perception and cognition* (pp. 335-362). Hillsdale, NJ: Erlbaum.
- THOMAS, E. A., & LEGGE, D. (1970). Probability matching as a basis for detection and recognition decisions. *Psychological Review*, **77**, 65-72.
- VON WINTERFELDT, D., & EDWARDS, W. (1982). Costs and payoffs in perceptual research. *Psychological Bulletin*, **91**, 609-622.
- WALLSTEN, T. S., & GONZALEZ-VALLEJO, C. (1994). Statement verification: A stochastic model of judgment and response. *Psychological Review*, **101**, 490-504.
- WICKENS, T. D. (1982). *Models for behavior: Stochastic processes in psychology*. San Francisco: Freeman.

NOTES

1. Categorization problems also differ in the prior probabilities (or base rates) of each category. Base rates were equal in the present study but have been studied extensively in the literature (Green & Swets, 1966; Healy & Kubovy, 1981; Maddox & Bohil, 1998a, 1998b, 2000).
2. Other weighting schemes are possible. For example, instead of generating an intermediate decision criterion, it is possible that the two decision criteria compete on each trial for the opportunity to generate the categorization response (for related proposals, see Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Maddox & Estes, 1996). A rigorous comparison of these alternatives is beyond the scope of this article.
3. For completeness we also estimated d' from signal detection theory and submitted these data to a $2 \times 2 \times 2 \times 4$ ANOVA. The only significant effect was the main effect of category discriminability. These data will not be discussed further.
4. As in Experiment 1, the d' values from signal detection theory were submitted to a $2 \times 2 \times 4$ ANOVA. The only significant effect was the main effect of category discriminability. These data will not be discussed further.

(Manuscript received June 22, 2000;
revision accepted for publication February 7, 2001.)