

# Category discriminability, base-rate, and payoff effects in perceptual categorization

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

The optimality of perceptual categorization performance under manipulations of category discriminability (i.e.,  $d'$  level), base rates, and payoffs was examined. Base-rate and payoff manipulations across two category discriminabilities allowed a test of the hypothesis that the steepness of the objective reward function affects performance (i.e., the *flat-maxima* hypothesis), as well as the hypothesis that observers combine base-rate and payoff information independently. Performance was (1) closer to optimal for the steeper objective reward function, in line with the flat-maxima hypothesis, (2) closer to optimal in base-rate conditions than in payoff conditions, and (3) in partial support of the hypothesis that base-rate and payoff knowledge is combined independently. Implications for current theories of base-rate and payoff learning are discussed.

Every day we are faced with situations in which we must choose a course of action on the basis of uncertain (or imperfect) predictors. For example, we might decide to bring or not bring a heavy jacket to work on the basis of a quick feel of the early morning temperature. This is a categorization problem because there are many possible early morning temperatures that one might observe, but only two courses of action. If we decide to bring a heavy jacket and it is cold that day, or if we decide not to bring a heavy jacket and it is warm, we have made a correct decision. If we decide to bring a heavy jacket and it is warm that day, or if we decide not to bring a heavy jacket and it is cold, we have made an incorrect decision. Performance in these types of categorization tasks is dependent on at least three factors: the predictability of the stimulus information for correct categorization, the prior probabilities (or base rates) of the categories, and the costs and benefits (or payoffs) associated with each categorization decision.

Information sources differ in their predictability for correct categorization. For example, in Central Texas the early morning temperature is only moderately predictive of the weather later that day, whereas in Southern California the early morning temperature is a much better predictor. Throughout this article, we use the term *category discriminability* to refer to the expected diagnosticity of the information source. Categories with high discriminability lead to more accurate performance, whereas categories with low discriminability lead to less accurate per-

formance. The category base rates also affect performance. For example, if a cold front has set in, cold days will be more likely than warm days. Base rates can bias observers' decisions by shifting their decision criteria along the stimulus dimension. For example, during a cold front, one might be more likely to bring a heavy jacket even if it is moderately warm in the morning. Finally, the payoffs associated with each categorization decision will affect performance. We generally benefit when we make the "correct" decision, but the benefit for the two types of correct decisions might differ. For example, we might benefit more when we bring a heavy jacket on a cold day, by staving off the flu, than if we do not bring a heavy jacket on a warm day. The greater benefit of staving off the flu might bias one's decision in such a way that one is more likely to bring a heavy jacket even if it is warm in the early morning. (The costs of the two types of incorrect decisions might also differ, but the present study focuses only on the effects of unequal benefits.)

The perceptual and cognitive processes involved in solving categorization problems of this sort have been studied extensively (Ashby, 1992a; Busemeyer & Myung, 1992; Green & Swets, 1966; Healy & Kubovy, 1981; Koehler, 1996; Macmillan & Creelman, 1991; Maddox, 1995; Maddox & Ashby, 1993; Maddox & Bohil, 1998a, 1998b, 2000; Stevenson, Busemeyer, & Naylor, 1991; von Winterfeldt & Edwards, 1982). A fruitful approach has been to compare human performance with that of the optimal classifier—a hypothetical device whose categorization decisions maximize long-run reward. This approach is useful because the effects of category discriminability, base rates, and payoffs are identifiable in the optimal classifier's performance and thus can be investigated systematically. Several studies (to be reviewed shortly) have examined base-rate and payoff sensitivity in perceptual categorization tasks, although most of these have examined either base-rate or payoff manipulations (Busemeyer & Myung, 1992; Lee & Janke, 1964, 1965;

---

This research was supported in part by National Science Foundation Grant SBR-9796206 and NIH Grant R01 MH59196. We thank Jerome Busemeyer, Michael Kubovy, and Thomas Wallsten for their helpful comments on an earlier draft of this manuscript. We also thank Greg Ashby, Jeffrey Dodd, Wilson Geisler, and Arthur Markman for many interesting discussions that influenced this work. Special thanks go to Candace Blair and Greg Smith for help with the 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).

Maddox, 1995; Maddox & Bohil, 1998b) and have not manipulated both within the same experimental framework (however, see Green & Swets, 1966; Healy & Kubovy, 1981; Maddox & Bohil, 1998a). To our knowledge, no studies have examined the effects of category discriminability on base-rate and payoff sensitivity in perceptual categorization (however, see Kubovy & Healy, 1980).

### THE OPTIMAL CLASSIFIER

Consider an experiment involving two categories, A and B, whose exemplars vary continuously along a dimension,  $x$ . When presented with a stimulus, the optimal classifier perfectly encodes the dimensional value. In other words, for a fixed physical input the optimal classifier will show no variability in the perceptual representation (although there may be random variation in the physical input to the perceptual system from a fixed stimulus; Geisler, 1989). The optimal classifier has complete knowledge of the distribution characteristics (i.e., the form and parameters) of each category. It uses this information to generate the *optimal decision function*,

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

where  $f(x|i)$  denotes the likelihood of perceptual effect  $x$ , given a stimulus presentation from Category  $i$ . If this likelihood ratio is greater than one, the perceptual effect is more likely to have resulted from Category B.

The optimal classifier also has perfect knowledge of category base rates and the benefits associated with each response. The optimal classifier integrates this information to derive the *optimal decision criterion*,

$$\beta_o = [P(A)/P(B)] \times [(V_{aA} - V_{bA})/(V_{bB} - V_{aB})], \quad (2)$$

where  $P(A)$  and  $P(B)$  denote the category base rates, and  $V_{aA}$  and  $V_{bB}$  denote the benefit associated with correct responses (lowercase letters denote responses, uppercase letters denote category membership), and  $V_{bA}$  and  $V_{aB}$  denote the costs of incorrect responding. In the present research, all costs are zero, so these last two (cost) terms can be dropped.

The optimal classifier (e.g., Green & Swets, 1966) uses  $l_o(x)$  and  $\beta_o$  to formulate 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 base rates and benefits are such that responding is unbiased, then  $\beta_o = 1$ . When the base rates, benefits, or some combination of the two favors Category A,  $\beta_o > 1$ . Conversely, when the base rates, benefits, or some combination of the two favors Category B,  $\beta_o < 1$ . The partition that separates the "A" and "B" response regions [where  $l_o(x) = \beta_o$ ] is called the *optimal decision bound* (Ashby, 1992a; Maddox & Ashby, 1993).

### BASE-RATE AND PAYOFF EFFECTS IN PROBABILISTIC CATEGORIZATION

A seminal study of base-rate and payoff learning was conducted by Healy and Kubovy (1981; see also Green & Swets, 1966). In the study, two normally distributed categories were constructed with category  $d' = 1$ . A schematic illustration of two normally distributed categories with  $d' = 1$  is depicted in Figure 1 (for now, ignore the two decision criteria,  $\beta$  and  $\beta_o$ , and the shaded area denoted "CCP"). The stimuli were five-digit numbers representing the heights of individuals from a hypothetical population, and observers were instructed to categorize a presented stimulus as being the height of either a man or a woman. A within-observers design was used, and each observer completed a number of experimental conditions in which category base rates, payoffs, or both were manipulated. (Like the present study, only benefits were manipulated while holding costs fixed at zero.) Three results stand out. First, across base rate alone, payoff alone, and simultaneous base-rate/payoff manipulations, observers exhibited a pattern of responding known as *conservative cutoff placement (CCP)*. Specifically, they used decision criteria that fell somewhere between the optimal decision criterion and the equal likelihood ( $\beta = 1$ ) decision criterion. An example in which the optimal decision criterion,  $\beta_o$ , equals 3 is depicted in Figure 1. The shaded region denotes the set of decision criteria that result in CCP. Second, the observers' decision criteria were consistently closer to the optimal decision criterion when base rates, as opposed to payoffs, were manipulated, even though the optimal decision criterion was identical in the two cases. Finally, Stevenson et al. (1991) reanalyzed the Healy and Kubovy (1981) data and found some support for the notion that observers combine base-rate and payoff information independently, as is predicted by the optimal classifier (see Equation 2).

Several explanations for observers' suboptimal responding have been offered (e.g., Healy & Kubovy, 1981; Kubovy, 1977; Kubovy & Healy, 1980; Maloney & Thomas, 1991; Thomas & Legge, 1970). One hypothesis that is of




Figure 1. Decision criterion for the classifier that maximizes long-run reward ( $\beta_o = 3$ ) and the equal likelihood ( $\beta = 1$ ) decision criterion. The shaded area denotes the region for which conservative cutoff placement (CCP) results.

particular interest states that the *steepness* of the objective reward function (ORF) influences the observer's ability to learn the optimal decision criterion (e.g., Busemeyer & Myung, 1992; Dusoir, 1980; Kubovy & Healy, 1977; Thomas, 1975). The ORF plots the predicted reward for different values of the decision criterion, defined in signal detection theory as follows:  $\ln\beta/d'$  (Busemeyer & Myung, 1992; Stevenson et al., 1991; von Winterfeldt & Edwards, 1982). More specifically, the reward for a given criterion value is determined by  $V_{aA}P(A)P_{cA} + V_{bB}P(B)P_{cB}$ , where  $P_{cA}$  and  $P_{cB}$  represent the proportion correct from Categories A and B, respectively, for that criterion value. A large body of theoretical work has shown that the steepness (or change in reward as a function of changes in the decision criterion) of the ORF is strongly affected by the category discriminability, or  $d'$  (Busemeyer & Myung, 1992; Stevenson et al., 1991; von Winterfeldt & Edwards, 1982). The ORF for two category discriminabilities,  $d' = 1$  and  $d' = 2.16$ , are displayed in Figure 2. Notice that the function is relatively flat for  $d' = 1$  and is much steeper for  $d' = 2.16$ . In fact, the function is steepest for  $d'$  values near 2.16. If observers adjust their decision criterion on the basis of the rewards that they accrue in the task (e.g., Dusoir, 1980; Kubovy & Healy, 1977; Thomas, 1975), it is likely that observers will stop adjusting their criterion when they reach the flat portion of the reward function. Because the ORF is steeper for  $d' = 2.16$  than for  $d' = 1$ , the observer's decision criterion and performance should be closer to optimal when  $d' = 2.16$  than when  $d' = 1$ . To date, this *flat-maxima* hypothesis has not been tested empirically. A major goal of the present study was to test rigorously the flat-maxima hypothesis.

The flat-maxima hypothesis provides a testable explanation for the poor performance when  $d' = 1$ , but it does not offer an explanation for the fact that base-rate performance is generally superior to payoff performance. Maddox and Bohil (1998a) offered a competition between reward and accuracy maximization (COBRA) hypothe-

sis that predicts 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 responding. When payoffs are manipulated and base rates are equal, both goals cannot be achieved simultaneously, because the decision rule that maximizes accuracy is different from the decision rule that maximizes reward. Suppose the payoffs are such that  $\beta_0 = 3$ . In this case, accuracy is maximized when  $\beta = 1$ , but reward is maximized when  $\beta_0 = 3$  (see Figure 1). An observer who places importance on both goals will likely use an intermediate value of  $\beta$ . Suppose the base rates are such that  $\beta_0 = 3$ . In this case, accuracy *and* reward are maximized when  $\beta_0 = 3$  (see Figure 1). Thus, the observer's decision criterion should be closer to optimal when base rates, as opposed to payoffs, are unequal.<sup>1</sup>

A final issue to be addressed is the independence assumption of the optimal classifier. Specifically, the optimal classifier combines independently base-rate and payoff information when constructing the optimal decision criterion. Stevenson et al. (1991) tested the independence assumption in Healy and Kubovy's (1981) categorization data where  $d' = 1$ . Maddox and Bohil (1998a, Experiment 2) tested the independence assumption in data from a categorization problem where  $d' = 2.16$ . In both cases, the independence assumption received support. Unfortunately, a number of procedural differences exist between these studies that make a direct comparison difficult. The present study manipulated category  $d'$  within the same experimental framework, which provided a stronger test of the independence assumption.

## EXPERIMENT

The goals of the present research were many. First, we examined the effects of category discriminability,  $d'$ , on categorization performance when base rates, payoffs, or both were manipulated within a fixed experimental framework. Second, we assessed whether any discriminability

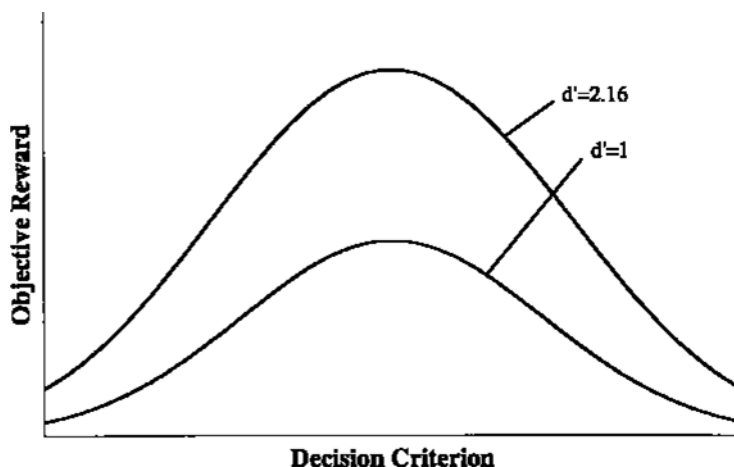


Figure 2. Objective reward functions associated with  $d' = 1$  and  $d' = 2.16$ .

effects matched the pattern of results predicted by the flat-maxima hypothesis. Third, we tested rigorously the hypothesis that observers combine base-rate and payoff information independently, as is predicted by the optimal classifier. Fourth, we examined the time course of learning by having observers complete several blocks of trials in each experimental condition and by analyzing the block-by-block data. To achieve these goals, we examined a number of summary statistics, including accuracy rates, point totals, and decision criterion values derived from hit and false alarm rates. In addition, we applied a series of decision bound models, each of which instantiated a specific hypothesis regarding base-rate and payoff learning and the potential effects of category discriminability.

As in two recent studies (Maddox & Bohil, 1998a, 2000) a within-observers design was utilized. A within-observers design is advantageous for at least two reasons. First, it minimizes the effects of within-group variance. Second, it provides a richer set of data that allows rigorous tests of a number of hypotheses that could not be tested on data from between-observers designs (e.g., the independence assumption of the optimal classifier). Five base-rate and payoff conditions were combined factorially with two category discriminabilities for a total of 10 experimental conditions. The 5 base-rate/payoff conditions were the same as those used by Maddox and Bohil (1998a, Experiment 2). These included a 3:1 base-rate-only condition, a 3:1 payoff-only condition, a 3:1 base-rate/3:1 payoff condition, a 3:1 base-rate/1:3 payoff condition, and a 1:3 base-rate/3:1 payoff condition. For brevity, these conditions are referred to throughout the remainder of the article as follows: 3:1B/1:1P (1:1P denotes the unbiased payoff values), 1:1B/3:1P (1:1B denotes unbiased base-rate ratio), 3:1B/3:1P, 3:1B/1:3P, and 1:3B/3:1P, respectively. We refer to the latter three conditions collectively as *combination conditions*, because in these conditions, we manipulate simultaneously base-rate and payoff ratios.

## Method

### Observers

Eight observers were solicited from the University of Texas community. The observers were paid on the basis of their day-to-day performance in the task. All the observers claimed to have 20/20 vision or vision corrected to 20/20.

### Stimuli and Stimulus Generation

The stimulus was a dot that varied in vertical (or horizontal) location from trial to trial, presented within a  $500 \times 500$  pixel square that was centered on a  $1,024 \times 768$  resolution computer screen. In conditions in which the dot varied in vertical (horizontal) position from trial to trial, the horizontal (vertical) position of the dot remained fixed at the center of the screen. For each observer, one category  $d'$  condition always used a dot that varied in vertical position, and the other  $d'$  condition always used a dot that varied in horizontal position. The  $d'$ 's with horizontal/vertical pairings were counterbalanced across observers.

Each of the two categories, A and B, was defined by a specific univariate normal distribution. For the horizontally varying dot stimuli in the  $d' = 1$  condition, the Category A mean was set 21 pixels

to the left of the center of the screen, and the Category B mean was set 21 pixels to the right of the center of the screen. For both categories, the standard deviation was set to 42 pixels. For the horizontally varying dot stimuli in the  $d' = 2.16$  condition, the Category A mean was set 45 pixels to the left of center, and the Category B mean was set 45 pixels to the right of center. Again, the standard deviation was set to 42 pixels. For the vertically varying dot stimuli, the parameter values were the same, except that the means were displaced vertically, instead of horizontally, from the center of the screen.

Prior to the experiment, three sets of 120 stimuli were sampled from the  $d' = 1$  category distributions, and three sets of 120 stimuli were sampled from the  $d' = 2.16$  category distributions. One set had equal numbers of A and B stimuli and were used in the baseline and 1:1B/3:1P conditions. The second set contained three times as many A as B stimuli, and the third set contained three times as many B as A stimuli. These were used in the appropriate unequal base-rate conditions. Each session in the experiment consisted of five blocks of 120 trials. In each block of trials, every stimulus from the appropriate 120-trial stimulus set was presented once. The presentation orders were randomized across blocks. The first four blocks (480 trials) made up the training phase. Corrective feedback was provided following each trial of the training phase (see details below). The final block of 120 trials comprised the transfer phase, during which none of the trials was followed by feedback.

In conditions in which the payoffs were equal, the payoff associated with a correct response was \$0.02. In the biased payoff conditions, the value of a correct high-payoff category response was \$0.03, and the value of a correct low-payoff category response was \$0.01. In all the conditions, the value of an incorrect response was \$0.00. Table 1 displays the optimal points, accuracy rates, and  $\beta$  values for the optimal classifier that maximizes long-run reward separately for each condition.

### Procedure

The observers were told that perfect performance was impossible. The observers were instructed to maximize the amount they earned and not to worry about speed of responding. A typical trial proceeded as follows. A stimulus was presented on the screen and remained until a response was made. The observer's task was to classify the presented stimulus as a member of Category A or Category B by pressing the appropriate button. During the training phase, the observer's response was followed, after a 250-msec delay, by 750 msec of feedback. Three lines of feedback were presented. The top line indicated the amount of money the observer earned for the response. The next line indicated the potential earnings for a correct response on each trial (i.e., following an incorrect response, this line indicated the amount that could have been earned had the observer chosen the correct response). The third line indicated the amount of money that the observer had accumulated up to that point in the session. The feedback was followed by a 125-msec intertrial interval, during which the screen was blank. In the transfer phase, corrective feedback was absent, and the observer was instructed to use the same decision criterion that he or she had been using during the final block of training. During these trials, the observer's response was

**Table 1**  
Optimal Points, Accuracy and  $\beta$  values

	$d' = 1$		$d' = 2.16$		$\beta_o$
	Points	Accuracy	Points	Accuracy	
Baseline	166	69.15	206	85.94	1
3:1B/1:1P	187	77.75	213	88.66	3
1:1B/3:1P	187	60.99	213	82.93	3
3:1B/3:1P	270	75.87	281	86.70	9
3:1B/1:3P	124	69.15	155	85.94	1
1:3B/3:1P	124	69.15	155	85.94	1

followed by a 1,125-msec intertrial interval, during which the screen was blank (i.e., the same length of time between stimulus presentations as in the training phase). The observers were given periodic breaks throughout the daily sessions. At each break, the observer's accumulated earnings were displayed, along with the total earnings possible (i.e., the optimal earnings) for that condition.

Half of the observers completed the  $d' = 1$  conditions first, and half completed the  $d' = 2.16$  conditions first. During the first session of each  $d'$  condition, each observer completed five blocks in the baseline condition in which base rates and payoffs were equal. The baseline condition was completed first, to ensure that each observer had accurate knowledge of the category structures before being exposed to the base-rate and/or payoff manipulations. The five experimental conditions were then completed, one session per day, on subsequent days, starting with the 3:1B/1:1P or the 1:1B/3:1P condition (counterbalanced between Days 2 and 3). The third experimental condition was always the 3:1B/3:1P condition, and the fourth and fifth conditions were either 3:1B/1:3P or 1:3B/3:1P (again, counterbalanced between days). Once an observer completed the first  $d'$  condition, they repeated the sequence (with different orders of the 3:1B/1:1P and 1:1B/3:1P conditions, as well as the 3:1B/1:3P and 1:3B/3:1P conditions) with the other  $d'$  value and with the orientation of the stimulus variation rotated 90° (i.e., from vertical to horizontal, or vice versa).

Before each experimental session, the observer completed a minimum of 60 baseline trials. If the observer reached an accuracy-based criterion of not more than 2% below optimal (i.e., 67% correct in  $d' = 1$  or 84% correct in  $d' = 2.16$ ), he or she was allowed to begin the experimental condition. If the observer did not reach criterion in 60 trials, he or she continued in the baseline condition until the criterion was reached. Once criterion was reached, he or she was allowed to begin that day's experimental condition. Including these baseline trials prior to each experimental condition helped ensure that each observer had accurate knowledge of the category structures before exposure to the base-rate and/or payoff manipulation and minimized the possibility of carryover effects from one experimental condition to the next.

## Results

This section is organized as follows. First, we evaluate performance in the baseline conditions to determine whether each observer had accurate knowledge of the category structures before being exposed to the base-rate and payoff manipulations. Second, we summarize the trends in the data by performing analyses of variance (ANOVAs) on a number of performance measures, such as accuracy rates, point totals, and decision criterion estimates from signal detection theory (Green & Swets, 1966). Finally, we applied a series of nested decision bound models simultaneously to the data from all 10 experimental conditions separately for each observer. These model-based analyses allowed more rigorous tests of the flat-maxima and *independence* hypotheses than are possible with ANOVAs.

### Baseline Condition

The purpose of the baseline condition was to ensure that the observers had a good understanding of the category distributions prior to any manipulation of base rates and payoffs. A number of analyses were conducted to determine whether this goal was achieved. Points and accuracy showed a general increase across the four training blocks, gradually approaching the optimal values. During the transfer block, average accuracy rates were 67%

and 85% for the  $d' = 1$  and  $d' = 2.16$  conditions, respectively. These are very close to the optimal values of 69% and 86%, respectively. Average point totals were 162 and 205 points for  $d' = 1$  and  $d' = 2.16$ , respectively. Again these are very similar to the optimal values of 166 and 206, respectively. Finally, the average  $\beta$  values (derived from hit and false alarm rates) were 1.039 and 1.092 for the  $d' = 1$  and  $d' = 2.16$  conditions, respectively. The optimal  $\beta$  value in both conditions was 1. Taken together, the point totals, accuracy rates, and  $\beta$  values suggest that the observers had accurate knowledge of the category distributions and were using a decision criterion that was very nearly optimal by the end of the baseline session.

### Experimental Conditions

Performance was evaluated for each of the four training blocks and the final transfer block. For completeness and to facilitate future model testing, hit and false alarm rates relative to Category A for each observer by block and experimental condition are presented in Appendices A and B. Because the point totals and accuracy rates predicted by the optimal classifier differed across conditions, we computed the percentage of optimal points and the percentage of optimal accuracy for each observer, as well as the deviation from the optimal decision criterion. We performed separate  $2(d') \times 5$  (base-rate/payoff condition)  $\times 5$  (block) within-observers ANOVAs on the percentage of optimal points, percentage of optimal accuracy, and deviations from optimal decision criterion.

The left-hand column of Figure 3 depicts the performance differences between  $d'$  conditions on the basis of percentage of optimal points, percentage of optimal accuracy, and deviation from optimal decision criterion. Notice that performance is closer to optimal for all three measures when  $d' = 2.16$  than when  $d' = 1$ . Specifically, the percentage of optimal points and percentage of optimal accuracy measures are closer to 100%, and the deviation from optimal decision criterion is smaller for  $d' = 2.16$  than for  $d' = 1$ . The main effect of  $d'$  reached significance for all three performance measures [percentage of optimal points,  $F(1,7) = 10.96, p < .05$ ; percentage of optimal accuracy,  $F(1,7) = 19.94, p < .005$ ; deviation from optimal decision criterion,  $F(1,7) = 146.82, p < .001$ ].

The right-hand column of Figure 3 depicts the performance differences between base-rate/payoff conditions on the basis of percentage of optimal points, percentage of optimal accuracy, and deviation from optimal decision criterion. The main effect of base-rate/payoff condition reached significance for the percentage of optimal accuracy measure [ $F(4,28) = 16.09, p < .001$ ] and the deviation from optimal decision criterion [ $F(4,28) = 55.47, p < .001$ ] but was nonsignificant for the percentage of optimal points [ $F(4,28) = 1.63, p > .05$ ]. Post hoc tests revealed that percentage of optimal accuracy was significantly closer to optimal in the 3:1B/1:1P condition than in the 1:1B/3:1P condition. Note, also, that percentage of optimal accuracy was greater than 100% in the 1:1B/3:1P condition. This result makes sense if the goals of

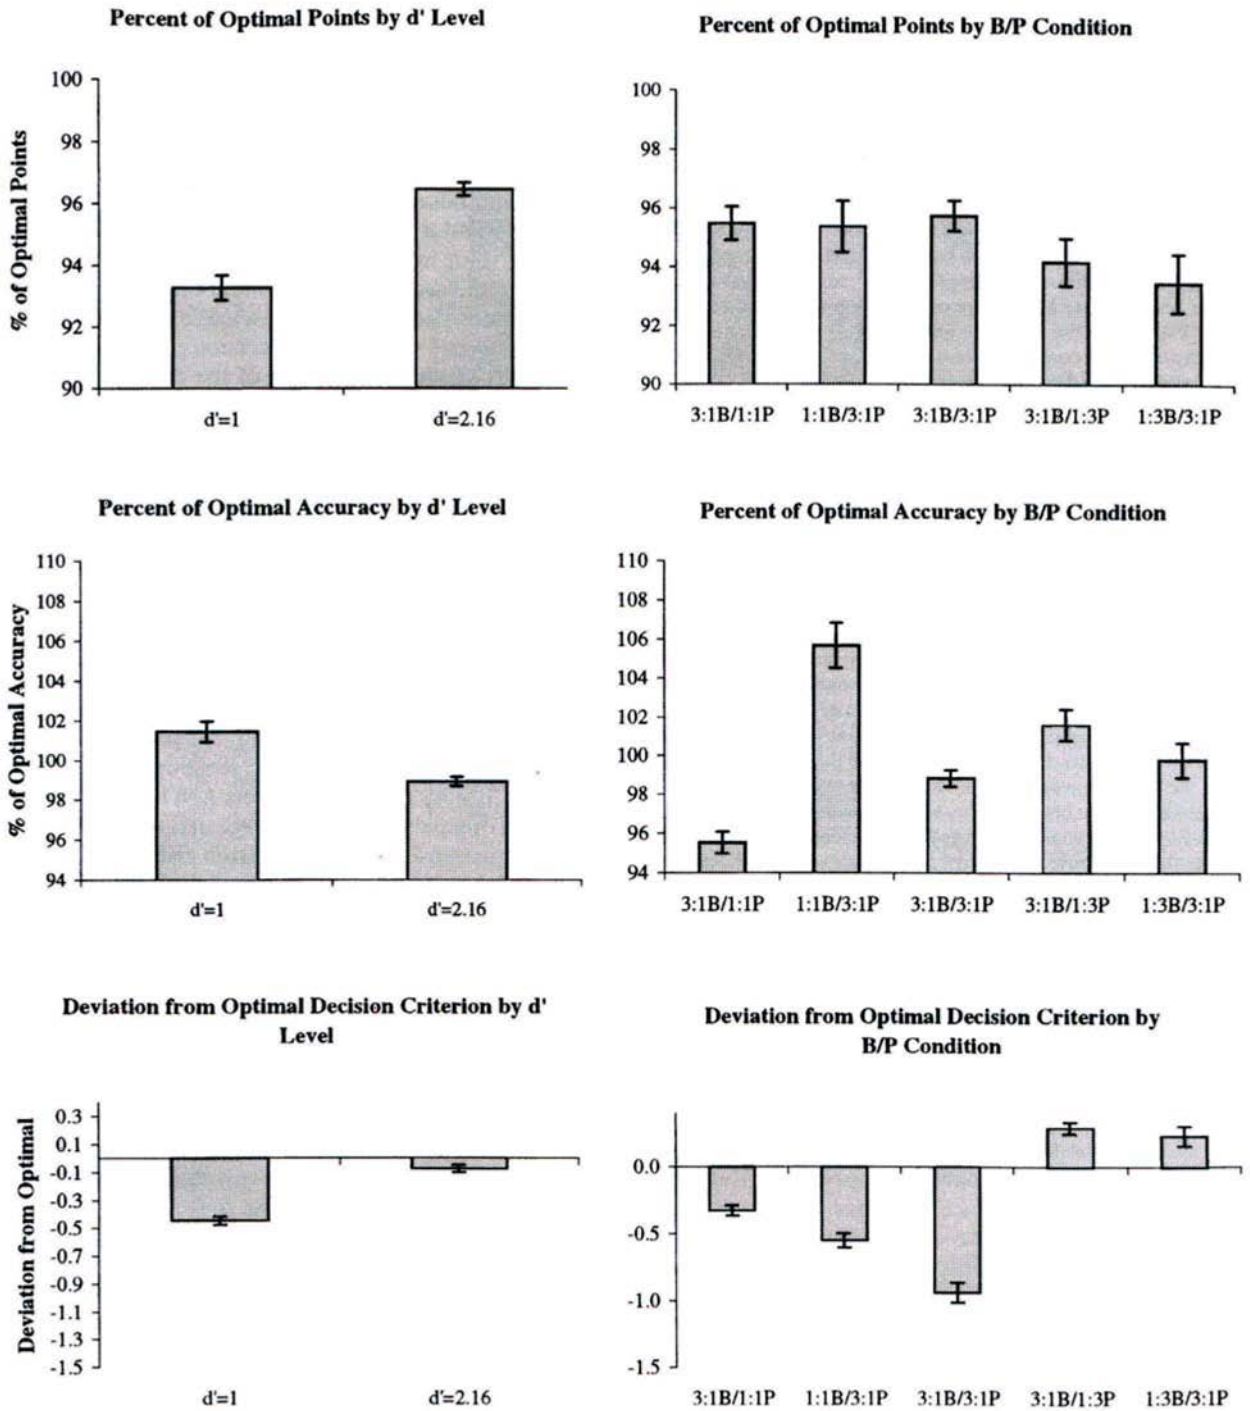


Figure 3. Average percentage of optimal points, percentage of optimal accuracy, and deviation from optimal decision criterion for each  $d'$  level (left-hand column) and for each base-rate/payoff condition (right-hand column).

reward maximization and accuracy maximization compete in the 1:1B/3:1P condition, but not in the 3:1B/1:1P condition. Percentage of optimal accuracy was significantly closer to optimal in the three combination conditions than in either the 3:1B/1:1P condition or the 1:1B/3:1P condition. Also, post hoc tests indicated that deviations from the optimal decision criterion was sig-

nificantly smaller in the 3:1B/1:1P condition than in the 1:1B/3:1P condition and that deviation was largest in the 3:1B/3:1P condition. The smallest deviation from the optimal decision criterion was found in the 3:1B/1:3P and 1:3B/3:1P conditions.

The main effect of block reached significance only for the deviation from optimal decision criterion measure

[ $F(4,28) = 5.43, p < .005$ ] and was nonsignificant for the other two measures. Although not monotonic across blocks, the deviation from the optimal decision criterion tended to decrease across blocks.

Figure 4 displays the percentage of optimal points, percentage of optimal accuracy, and deviation from optimal decision criterion by  $d'$  for each of the five base-rate/payoff conditions. The interaction between  $d'$  and base-rate/payoff condition reached significance for all three measures [percentage of optimal points,  $F(4,28) = 3.41, p < .05$ ; percentage of optimal accuracy,  $F(4,28) = 11.24, p < .001$ ; deviation from optimal decision criterion,  $F(4,28) = 22.54, p < .001$ ]. To evaluate the interaction more fully, we conducted separate  $2(d') \times 5$  (block) ANOVAs for each of the five base-rate/payoff conditions. The results indicated a significant  $d'$  effect for all three performance measures in the 3:1B/1:1P [percentage of optimal points,  $F(1,7) = 8.83, p < .05$ ; percentage of optimal accuracy,  $F(1,7) = 7.98, p < .05$ ; deviation from optimal decision criterion,  $F(1,7) = 24.29, p < .005$ ] and 1:1B/3:1P [percentage of optimal points,  $F(1,7) = 8.01, p < .05$ ; percentage of optimal accuracy,  $F(1,7) = 21.00, p < .005$ ; deviation from optimal decision criterion,  $F(1,7) = 57.29, p < .001$ ] conditions. For both conditions and for all three measures, performance was closer to optimal for  $d' = 2.16$  than for  $d' = 1$ . The  $d'$  effect was significant for two of three measures in the 3:1B/3:1P condition [percentage of optimal points,  $F(1,7) = 24.26, p < .005$ ; deviation from optimal decision criterion,  $F(1,7) = 73.60, p < .001$ ] and significant for only one of three measures in the 3:1B/1:3P [percentage of optimal accuracy,  $F(1,7) = 20.49, p < .005$ ] and 1:3B/3:1P [percentage of optimal points,  $F(1,7) = 15.44, p < .01$ ] conditions. Regardless of significance, performance was closer to optimal in all  $d' = 2.16$  conditions, on all three performance measures, than in the corresponding  $d' = 1$  conditions, with only a single exception (percentage of optimal accuracy, 1:3B/3:1P condition). In terms of percentage of optimal points, the interaction was driven primarily by the 1:1B/3:1P condition. When  $d' = 1$ , the percentage of optimal points in the 1:1B/3:1P condition was generally lower than in the other conditions (with one exception), but when  $d' = 2.16$ , the percentage of optimal points was higher than in any other conditions (see Figure 4). With respect to percentage of optimal accuracy, values were higher (or about equal) when  $d' = 1$  than when  $d' = 2.16$  for all but the 3:1B/1:1P condition. Finally, for deviation from optimal decision criterion, deviations were negative (i.e., criterion values were smaller than optimal) in the 3:1B/1:1P, 1:1B/3:1P, and 3:1B/3:1P conditions and differed substantially across  $d'$  values, whereas deviations were positive in the 3:1B/1:3P and 1:3B/3:1P conditions and were similar across  $d'$  values.

Finally, the interaction between block and condition reached significance for each measure [percentage of optimal points,  $F(16,112) = 3.07, p < .001$ ; percentage of optimal accuracy,  $F(16,112) = 3.02, p < .001$ ; deviation from optimal decision criterion,  $F(16,112) = 1.97, p < .05$ ]. The two factor ANOVA results ( $d' \times$  block,

conducted separately for each base-rate/payoff condition) indicate a significant block effect in the 3:1B/1:1P condition [percentage of optimal points,  $F(4,28) = 5.24, p < .005$ ; percentage of optimal accuracy,  $F(4,28) = 5.24, p < .005$ ; deviation from optimal decision criterion,  $F(4,28) = 5.28, p < .005$ ] and also in the 1:1B/3:1P condition [for two of three measures: percentage of optimal points,  $F(4,28) = 3.77, p < .05$ ; deviation from optimal decision criterion,  $F(4,28) = 3.18, p < .05$ ]. Post hoc tests revealed that there was a general performance improvement across blocks in these two conditions. The block effect was significant only for a single measure in the 3:1B/3:1P [percentage of optimal points,  $F(4,28) = 4.58, p < .01$ ] and 1:3B/3:1P [percentage of optimal accuracy,  $F(4,28) = 3.12, p < .05$ ] conditions and never reached significance in the 3:1B/1:3P condition.

**Flat-maxima hypothesis.** The flat-maxima hypothesis predicts that performance should be closer to optimal in the  $d' = 2.16$  condition than in the  $d' = 1$  condition. The ANOVA results provide strong support for this hypothesis. Percentage of optimal points, percentage of optimal accuracy, and the deviation from optimal decision criterion measures all suggested that performance was closer to optimal in the  $d' = 2.16$  condition (see Figures 3 and 4). This was true across training blocks, as well as for the transfer block.

**Competition between reward and accuracy: 3:1B/1:1P versus 1:1B/3:1P conditions.** The COBRA hypothesis predicts that performance in the unequal base-rate condition (3:1B/1:1P) should be closer to optimal than performance in the unequal payoff condition (1:1B/3:1P). Specifically, percentage of optimal points should be closer to 100%, and the decision criterion values should be closer to optimal for the unequal base-rate condition than for the unequal payoff condition. Percentage of optimal points and decision criterion values were, on average, closer to optimal in the 3:1B/1:1P condition than in the 1:1B/3:1P condition. In addition, the COBRA prediction that observed accuracy in the payoff condition should be higher than that predicted by the optimal classifier was supported in the data. This was especially prominent in the  $d' = 1, 1:1B/3:1P$  condition.

**A test of the independence of base-rate and payoff information hypothesis.** Equation 2 illustrates the way in which the optimal classifier combines knowledge of base rates and payoffs. An alternative formulation applies the natural log to both sides of Equation 2, yielding  $\ln \beta_o = \ln[P(A)/P(B)] + \ln[(V_{AA}/V_{BB})]$ . Notice that  $\ln \beta_o$  is determined completely by the sum of independent base-rate and payoff terms. In ANOVA terms, the interaction of base rates and payoffs should be nonsignificant. Following Stevenson et al. (1991), Figure 5 plots the natural log of the average  $\beta$  values for the baseline, 3:1B/1:1P, 1:1B/3:1P, and 3:1B/3:1P conditions. The top and bottom panels present data for the  $d' = 1$  and  $d' = 2.16$  conditions, respectively.

The lines are slightly nonparallel for  $d' = 1$  but are quite parallel for  $d' = 2.16$ . ANOVAs were conducted separately on the data from the two panels. In both cases, the




Figure 4. Average percentage of optimal points, percentage of optimal accuracy, and deviation from optimal decision criterion by  $d'$  level and base-rate/payoff condition.

interaction term was nonsignificant, providing support for an independent combination of base-rate and payoff information. The 3:1B/1:3P and the 1:3B/3:1P conditions were not included in these ANOVAs because 1:3B/1:1P and 1:1B/1:3P conditions would also be required, and they were not included in the study. However, the model-based analyses reported below do allow us to use the 3:1B/1:3P and the 1:3B/3:1P conditions in a test of independence without requiring the 1:3B/1:1P and 1:1B/1:3P conditions.

The analyses reported above, although informative for evaluating the hypotheses of interest, face certain limitations and are far from exhaustive. First, the ANOVA results are based on aggregate data and thus might mask potential individual differences (Maddox, 1999). Second, although providing some evidence in support of the flat-maxima hypothesis, they do not allow a rigorous test of a *strong* version of the flat-maxima hypothesis (to be described shortly). Finally, in testing the independence assumption, we were forced to omit the data from two of



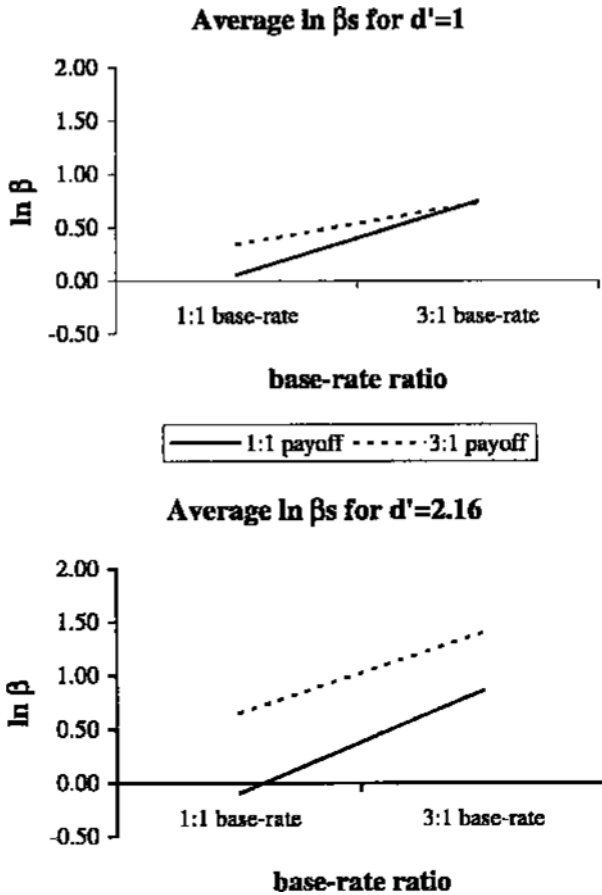


Figure 5. Average  $\ln \beta$  values from the baseline, 3:1B/1:1P, 1:1B/3:1P, and 3:1B/3:1P conditions for each  $d'$  level.

our five base-rate/payoff conditions. Each of these shortcomings was remedied by using the quantitative model-based approach described in the next section.

**Modeling Analyses Based on Decision Bound Theory**

The principle axiom of decision bound theory is that human performance rarely equals that of the optimal classifier because of two noise sources that are inherent in all biological organisms: perceptual noise and criterial noise (Ashby & Townsend, 1986; Maddox & Ashby, 1993). Perceptual noise results from trial-by-trial fluctuations in the observer’s percept of a given stimulus. Because of perceptual noise, the observer’s percept of stimulus  $i$ , on any trial, is given by  $x_{pi} = x_i + e_p$ , where  $e_p$  is a random variable that represents the effect of perceptual noise. Criterial noise results from trial-by-trial fluctuations in the observer’s memory for the decision criterion. Because of criterial noise, the observer’s decision criterion, on any trial  $i$ , is given by  $\beta_i = \beta + e_c$ , where  $\beta$  denotes the observer’s averaged decision criterion stored in memory, and  $e_c$  is a random variable that represents the effect of criterial noise (assumed to be univariate nor-

mally distributed). With the addition of perceptual and criterial noise, the Equation 3 decision rule becomes

$$\begin{aligned} &\text{If } l_o(x_{pi}) > \beta_o + e_c, \text{ then respond "B";} \\ &\text{otherwise, respond "A"} \end{aligned} \tag{4}$$

and is referred to as the *optimal decision bound model* (OPT).<sup>2</sup>

Two classes of suboptimality can be explored. One possibility is that the observer uses a suboptimal decision function,  $l(x_{pi})$ , rather than  $l_o(x_{pi})$ . The baseline trials were included specifically to alleviate this possibility. The analyses of the baseline data indicate that the observers had a good understanding of the category distributions prior to manipulations of base rates and payoffs and thus were likely using the optimal decision function. All of the models examined in this section assume that the observer used the optimal decision function. A second possibility is that the observer uses a suboptimal decision criterion. We will use a series of models that assume a suboptimal decision criterion to examine more rigorously the flat-maxima and independence of base-rate and payoff hypotheses. Specifically, we turn now to a series of nested decision bound models, each of which was applied simultaneously to the data from each of the 10 experimental conditions (2 category  $d' \times 5$  base-rate/payoff conditions), separately for each observer.<sup>3</sup> Because the block effects were negligible, we focus only on the transfer block. The first section addresses the flat-maxima hypothesis, and the second section addresses the independence hypothesis.

**Flat-maxima hypothesis.** The ANOVA results provide initial support for the flat-maxima hypothesis because performance (i.e., percentage of optimal accuracy, percentage of optimal points, and decision criterion estimates) was closer to optimal for  $d' = 2.16$  than for  $d' = 1$ . However, the ANOVA results do not test the stronger version of the flat-maxima hypothesis that states that the observer should stop adjusting their decision criterion once he or she reaches a point on the ORF at which changes in the decision criterion placement lead to little (if any) change in performance. Because the steepness (or slope) of the ORF determines how changes in the decision criterion will affect performance, the flat-maxima hypothesis can be instantiated by assuming that the observer adjusts their decision criterion until they reach a fixed steepness on the ORF. Since the ORF for  $d' = 2.16$  is steeper than the ORF for  $d' = 1$ , this model predicts that the observer will use a decision criterion that is closer to optimal for  $d' = 2.16$  than for  $d' = 1$  (see Figure 2). To test rigorously this hypothesis, we compared two models. The *flat-maxima model* estimated a single steepness value that determined the decision criterion to be used for both the  $d' = 1$  and the  $d' = 2.16$  conditions. Although the steepness parameter was “yoked” across  $d'$ , a separate steepness parameter was estimated in each of the five experimental conditions, yielding a total of five steepness parameters. This model was compared to the *fixed*

*decision criterion model*, which assumed that the observer used the same decision criterion in both  $d'$  conditions. In other words, this model assumed that there was no effect of category discriminability. A separate decision criterion parameter was estimated in each of the five experimental conditions (five parameters). For 6 of 8 observers, the flat-maxima model provided a better fit of the data than did the fixed decision criterion model and accounted (on average) for 93.75% of the variance in the data from each observer. This suggests that category discriminability affects decision criterion placement and supports the hypothesis that the decision criterion placement is determined by the steepness of the ORF.

#### **Independence of base-rate and payoff information.**

In this section, we provide a rigorous model-based test of the optimal classifier assumption that base-rate and payoff information are combined independently. Unlike the ANOVA results that excluded data from the 3:1B/1:3P and 1:3B/3:1P conditions, the model-based test of the independence assumption utilized the data from all five experimental conditions. The model-based analyses were conducted at the level of the individual observer and were not based on aggregate data. Two sets of models were developed. The first set was developed to test the independence assumption in the data from only one of the combination conditions—the 3:1B/3:1P condition. This allows a direct comparison with our previous tests of independence. The second set was more extensive and tested the independence assumption in all three combination conditions. Both sets of models also estimated parameters for the 3:1B/1:1P and 1:1B/3:1P conditions. Each model was fit to the appropriate experimental conditions separately for each  $d'$  level.

*3:1B/1:1P, 1:1B/3:1P, and 3:1B/3:1P conditions.* As a test of the independence assumption, we applied two models simultaneously to the data from the 3:1B/1:1P, 1:1B/3:1P, and 3:1B/3:1P conditions separately for each observer and  $d'$  condition. The *independence model* assumed that the decision criterion used in the 3:1B/3:1P condition was determined from an independent combination of the decision criterion estimates from the 3:1B/1:1P and 1:1B/3:1P conditions. Thus, only two suboptimal decision criterion values were estimated (one for the 3:1B/1:1P condition, and one for the 1:1B/3:1P condition). The *nonindependence model* allowed the decision criterion for the 3:1B/3:1P condition to also be a free parameter, thus violating independence. Notice that the independence model is a special case of the non-independence model and  $G^2$  tests can be performed. In 11 of 16 cases (2  $d'$  levels  $\times$  8 observers), the nonindependence model did *not* provide a significant improvement in fit over the independence model. This was due primarily to the  $d' = 2.16$  condition, in which 7 of 8 observers showed independence. These results converge generally with the ANOVA results outlined in Figure 5. When base rates and payoffs were manipulated simultaneously in the same direction, the observers had a tendency to use a decision criterion corresponding to an independent combi-

nation of these information sources, as was suggested by the ANOVA results.

*All experimental conditions.* In this section, we present a series of models that test the independence assumption of the optimal classifier in all three combination conditions. Each of the models was fit separately to the data from each category  $d'$  level. The assumptions of the models and their nested structure are outlined in Figure 6 (the arrows point to a more constrained model).

The simplest model tested contained two decision criterion parameters. The *suboptimal base-rate, suboptimal payoff, independence model* assumed that the observer combined base-rate and payoff information independently whenever base rates and payoffs were manipulated simultaneously. However, this model assumed a suboptimal estimate of *both* the base-rate and the payoff information. Three models with three decision criterion parameters were tested. Each of these models assumed independence in two of the three combination conditions. For example, the *suboptimal base-rate, suboptimal payoff, high/high non-independence model* was identical to the suboptimal base-rate, suboptimal payoff, independence model, except that independence is violated in the 3:1B/3:1P condition. The *suboptimal base-rate, suboptimal payoff, high/low nonindependence model* assumed that independence was violated in the 3:1B/1:3P condition, and the *suboptimal base-rate, suboptimal payoff, low/high nonindependence model* assumed that independence was violated in the 1:3B/3:1P condition. The most general model contained five free decision criterion parameters. The *suboptimal base-rate, suboptimal payoff, nonindependence model* assumed suboptimal estimates of both base-rate and payoff information and assumed that the observer did not combine base-rate and payoff information independently when base rates and payoffs were manipulated simultaneously. The decision criterion in each of the five conditions was a free parameter. It is important to be clear that all of the models that assumed both independence and suboptimality (in base-rate or payoff knowledge) predicted a suboptimal decision criterion in the 3:1B/3:1P, 3:1B/1:3P, and 1:3B/3:1P conditions. If the base-rate or payoff estimate is inaccurate but is combined independently, then suboptimality is predicted in the simultaneous base-rate/payoff manipulation conditions.

The most parsimonious model and predicted decision criteria by condition and observer are summarized in Table 2. In addition, when independence (i.e., Equation 2) was violated, the Equation 2 decision criterion is included in parentheses. Several interesting results emerge. First, the estimated 3:1B/1:1P and 1:1B/3:1P decision criteria derived from hit and false alarm rates and from the simultaneous fits are similar [ $r(30) = .86, p < .01$ ], suggesting that both approaches converge upon the same general conclusion. Second, independence was violated in about half of the cases—12 of 24 cases (8 observers  $\times$  3 combination conditions) when  $d' = 1$  and 11 of 24 cases when  $d' = 2.16$ . This result makes it difficult to support

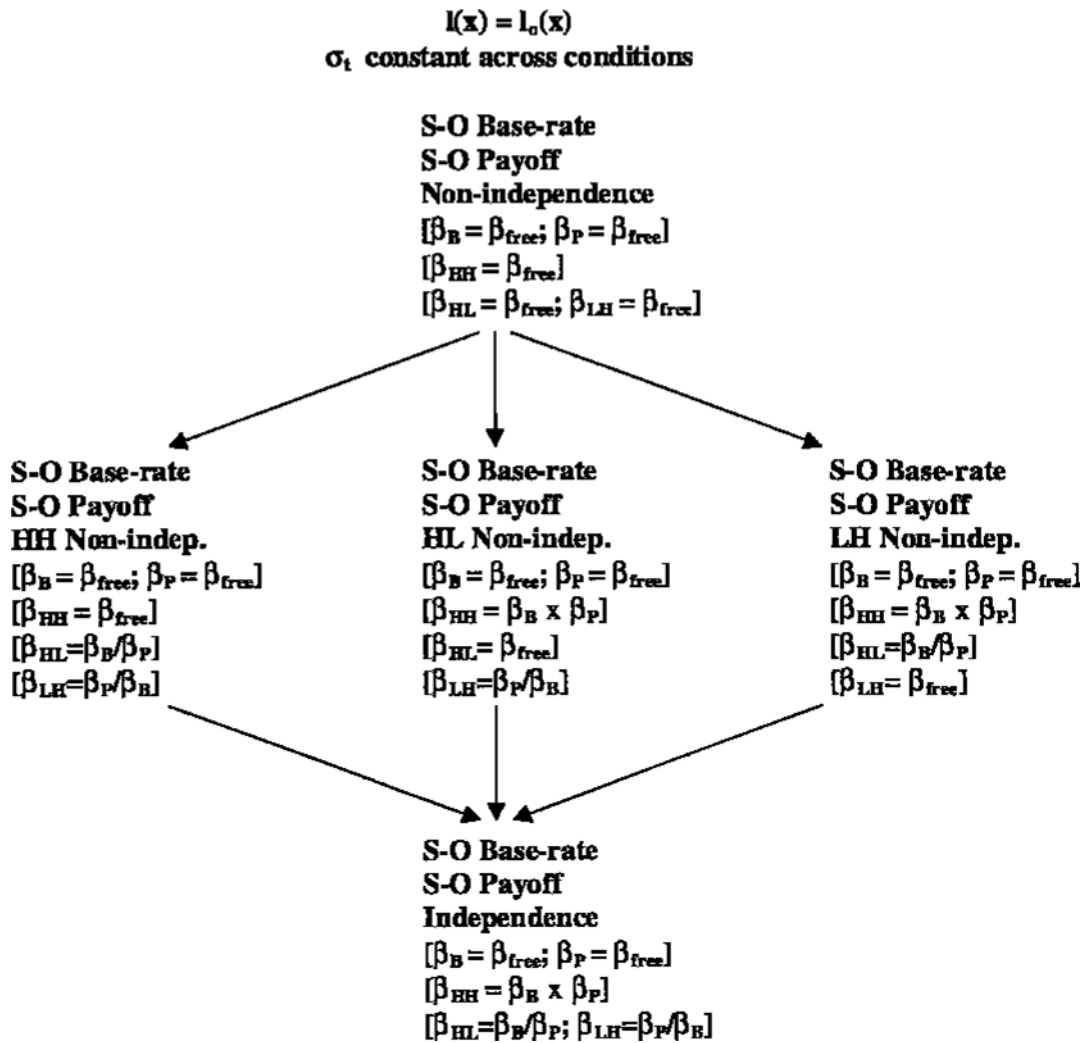


Figure 6. Nested relationship among the decision bound models applied to all three combination conditions within each level of  $d'$ . Each arrow points to a more constrained model. S-O, suboptimal; HH, high/high condition (or 3:1B/3:1P, see the text for an explanation); HL, high/low condition (or 3:1B/1:3P); LH, low/high condition (or 1:3B/3:1P).

an argument either for or against independent combination of base-rate and payoff knowledge. Notice that one observer violated the independence assumption in every combination condition. For the remaining observers, violations were not so systematic. In addition, no particular combination condition appears more likely to produce violations of independence than the others. Violations appeared approximately half the time for each combination condition.

The results of fitting the simultaneous models to all the experimental conditions (within a single  $d'$  level) appear to be somewhat at odds with earlier results reported by Maddox and Bohil (1998a, Experiment 2). In the earlier study, observers violated independence in only 25% of the cases. A direct comparison of the results, however, is difficult because of methodological differences between the studies. In the earlier study, 4 observers completed three combination conditions, all with a  $d'$  level of 2.16,

thus generating 12 cases for testing independence. In the present study, there were 48 possible cases to examine independence. Also, Maddox and Bohil's (1998a) earlier study utilized bivariate category structures, whereas in the present study, stimuli varied along only a single dimension. It is well established that observers require far more training to learn bivariate categories than univariate category structures. However, in both studies, the observers received enough training on the category structures to ensure a good understanding of the category distributions prior to beginning the experimental conditions. Because category learning and base-rate/payoff learning were separated in each study, and because the base-rate and payoff ratios were identical across experiments, it is reasonable to assume that the dimensionality of the categories was not the major cause of the discrepant results.

Also, in Maddox and Bohil's (1998a, Experiment 2) earlier study, the observers received a great deal more

**Table 2**  
**Most Parsimonious Model and  $\beta$  Estimates**  
**From the Simultaneous Fit Models Within Each Level of  $d'$**

Observer	Most Parsimonious Model	3:1B/1:1P ( $\beta_b$ )	1:1B/3:1P ( $\beta_p$ )	3:1B/3:1P ( $\beta_{HH}$ )	3:1B/1:3P ( $\beta_{HL}$ )	1:3B/3:1P ( $\beta_{LH}$ )
$d' = 1$						
1	HH	1.84	.67	2.82 (1.24)	2.73	.37
2	NI	5.51	1.69	3.21 (9.31)	5.53 (3.26)	.56(.31)
3	I	2.13	1.30	2.76	1.64	.61
4	NI	2.68	1.37	7.25 (3.66)	1.58 (1.96)	.83 (.51)
5	HH	2.78	1.74	2.93 (4.84)	1.59	.63
6	LH	7.50	8.85	66.39	.85	.94 (1.18)
7	NI	2.11	1.36	3.49 (2.87)	1.16 (1.56)	.99 (.64)
8	I	2.07	1.47	3.05	1.41	.71
Optimal $\beta$		3.00	3.00	9.00	1.00	1.00
$d' = 2.16$						
1	HL	4.46	.84	3.77	2.64 (5.28)	.19
2	NI	2.92	3.42	6.19 (9.98)	2.36 (.85)	.10 (1.17)
3	LH	2.11	1.53	3.21	1.38	.19 (.72)
4	I	3.58	2.66	9.52	1.35	.74
5	LH	1.97	2.22	4.39	.89	.56 (1.13)
6	NI	15.41	2.80	27.69 (43.16)	1.11 (5.5)	1.32 (.18)
7	HL	2.62	2.31	6.05	2.26 (1.14)	.88
8	HH	1.00	1.15	2.77 (1.16)	.87	1.15
Optimal $\beta$		3.00	3.00	9.00	1.00	1.00

Note—When independence was violated, the Equation 2 decision criterion is included in parentheses. HH, suboptimal base-rate, suboptimal payoff, high–high nonindependence model; NI, suboptimal base-rate, suboptimal payoff, nonindependence model; I, optimal base-rate, optimal payoff, independence model; LH, suboptimal base-rate, suboptimal payoff, low–high nonindependence model; HL, suboptimal base-rate, suboptimal payoff, high–low nonindependence model.

experience in the base-rate-alone and payoff-alone conditions (four 540-trial sessions in each experimental condition) than in the present study. In the earlier study, observers appeared to learn the base-rate and payoff ratios quickly (often in a single session) and reached asymptotic performance early on in each experimental condition. The observers then continued receiving training on the base-rate or payoff ratio until they had completed four sessions in that condition. Although speculative, greater experience with the base-rate and payoff ratios could explain the fact that violations of independence were less prevalent in Maddox and Bohil's (1998a) earlier study.

## GENERAL DISCUSSION

In this study, the differential effects of category discriminability on base-rate and payoff sensitivity within a fixed experimental framework were examined. The category discriminabilities were such that the flat-maxima hypothesis predicted less optimal performance in the  $d' = 1$  condition than in the  $d' = 2.16$  condition. Base-rate-alone, payoff-alone, and simultaneous base-rate/payoff conditions were included to test the assumption of the optimal classifier that base-rate information and payoff information are combined independently. These issues were examined through analyses of point totals, accuracy rates, and decision criterion values derived from hit and false alarm rates, as well as through the application of a series of decision bound models.

The results can be summarized as follows. First, on the basis of points, accuracy, and decision criterion estimates derived from hit and false alarm rates and from the model-based analyses, we found support for the flat-maxima hypothesis. Specifically, performance was closer to optimal in the  $d' = 2.16$  condition, for which the ORF was steepest, than in the  $d' = 1$  condition, for which the ORF was flatter. This is an important finding because much of the previous research utilized categories with  $d' = 1$ . The performance suboptimalities found in these earlier studies were taken as evidence that humans are poor at utilizing base-rate and payoff information. The present results suggest that the common use of small values of  $d'$  might have exaggerated the degree of suboptimality in previous studies.

Second, as was found in previous studies, performance was closer to optimal in the base-rate conditions than in the payoff conditions. This difference was more pronounced in the  $d' = 1$  condition than in the  $d' = 2.16$  condition. This  $d'$  effect makes sense if considered in light of both the flat-maxima hypothesis and the COBRA hypothesis. The COBRA hypothesis asserts that observers wish to maximize both reward and accuracy and that, with asymmetric category payoffs, overall accuracy must be sacrificed in order to maximize reward. This tradeoff between reward and accuracy maximization does not exist in the base-rate conditions. When  $d' = 1$  in payoff conditions, adjusting the decision criterion in the direction that maximizes reward results in a greater accuracy

loss than when  $d' = 2.16$  (see Table 1). Thus, when  $d' = 1$ , an observer who places a great deal of emphasis on accuracy might be more averse to selecting a reward-maximizing criterion than when  $d' = 2.16$ . This would amplify the performance difference between base-rate and payoff conditions when  $d' = 1$ . When  $d' = 2.16$ , the observer can emphasize reward maximization without a great loss in terms of accuracy.

Finally, we found mixed support for the notion that observers combine base-rate and payoff information independently. Independence was supported in the 3:1B/3:1P condition, based on an ANOVA on the criterion values derived from hit and false alarm rates, as well as decision criterion values estimated by the simultaneous model fits. However, when more rigorous simultaneous models were applied to all three combination conditions, the evidence for independence came into question. Specifically, independence was observed in about half of the experimental conditions, and in most cases, each observer violated independence in at least one experimental condition. Clearly, analyzing the entire data set for each individual observer constitutes the strongest test of the independence assumption, and future attempts to clarify this issue should be careful to follow a similar approach.

## FUTURE RESEARCH

### Flat-Maxima Interpretation of the Results

The fact that performance was closer to optimal when  $d' = 2.16$  supports the flat-maxima hypothesis. However, an alternative explanation is that performance approaches the optimal level as optimal point totals and accuracy rates increase. Because optimal accuracy and optimal points are higher in the  $d' = 2.16$  than in the  $d' = 1$  condition, this hypothesis predicts more nearly optimal performance in  $d' = 2.16$ , which was observed. Although this hypothesis accounts for the ordinal trends in the data, recall that the model-based analyses supported a strong version of the flat-maxima hypothesis that made quantitative performance predictions. Specifically, the model-based analyses supported a version of the flat-maxima hypothesis that postulated that the observer's decision criterion in the  $d' = 1$  and  $d' = 2.16$  conditions could be determined from a single *steepness* parameter derived from the objective reward functions. Although these more rigorous model-based analyses tend to support the flat-maxima hypothesis over the proposed alternative, clearly more work is needed.

One possibility would be to add a third, larger  $d'$  level to the experiment. For example, suppose that the  $d'$  values were 1, 2, and 3. In this case, optimal accuracy and optimal point totals increase as  $d'$  increases. Thus, the hypothesis that performance approaches the optimal level as optimal point totals and accuracy rates increase would predict that performance should be closer to optimal when  $d' = 3$  and would be furthest from optimal when  $d' = 1$ . The flat-maxima hypothesis, on the other hand, makes a different prediction. It is straightforward to show that the ORF is steepest for  $d' = 2$  and is flattest

for  $d' = 1$ . Thus, the flat-maxima hypothesis predicts performance that is closest to optimal for  $d' = 2$  (not  $d' = 3$ , as predicted from the alternative hypothesis) and that is furthest from optimal for  $d' = 1$ .

### The Influence of Trial-by-Trial Feedback

The results of this study suggest strongly that category discriminability affects the optimality of human performance. However, there are other factors that might also play a role in the optimality of human performance. One such possibility concerns the corrective feedback that observers are commonly presented with during experiments of this type. When observers are presented with trial-by-trial feedback that indicates their gain for the previous trial, they also receive information about the accuracy of their responses. Thus, researchers may unwittingly induce observers to attend to accuracy, which can result in suboptimal performance. It might be worthwhile to experiment with different types of feedback to assess their effect on base-rate and payoff learning. One possibility might be to present feedback only after the observer has completed a number of trials. In this way, the observer would get feedback for maximizing reward without receiving a constant reminder of his or her accuracy performance. Also, it might be useful to present observers with different types of feedback about their potential earnings (e.g., the amount that they could possibly earn). For example, observers could be presented with feedback regarding potential earnings that is based on the performance of the optimal classifier. In other words, the potential earnings feedback could reflect use of the optimal decision criterion (performance that is possible for them to nearly achieve), instead of the usual objective feedback (which represents a level of performance that is impossible to achieve, even for the optimal classifier).

### Learning Models

The two studies just described (using additional  $d'$  values and varying feedback type) could answer some important questions about the way humans utilize their knowledge of category base rates and payoffs. Clearly, however, there are a vast number of avenues for future research. For example, a number of learning models, such as Busemeyer and Myung's (1992) hill-climbing model, Erev's (1998; see also Erev, Wallsten, & Budescu, 1994) cutoff reinforcement learning model, and Wallsten and Gonzalez-Vallejo's (1994) stochastic judgment model, promise to shed new light on these issues. The prospect for many future studies of base-rate and payoff use is in keeping with the importance of understanding how humans utilize these sources of information.

## REFERENCES

- 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., & LEE, W. W. (1991). Predicting similarity and categorization from identification. *Journal of Experimental Psychology: General*, **120**, 150-172.
- ASHBY, F. G., & LEE, W. W. (1993). Perceptual variability as a fundamental axiom of perceptual science. In S. C. Masin (Ed.), *Foundations of perceptual theory* (pp. 369-397). New York: North-Holland.
- ASHBY, F. G., & MADDOX, W. T. (1998). Stimulus categorization. In M. H. Birnbaum (Ed.), *Measurement, judgment, and decision making* (pp. 251-301). New York: Academic Press.
- ASHBY, F. G., & TOWNSEND, J. T. (1986). Varieties of perceptual independence. *Psychological Review*, **93**, 154-179.
- BUSEMEYER, J. P., & 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.
- 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., WALLSTEN, T. S., & BUDESCU, D. V. (1994). Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review*, **101**, 519-527.
- GEISLER, W. (1989). Sequential ideal-observer analysis of visual discrimination. *Psychological Review*, **96**, 267-314.
- 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 formation of conservative decision rules in a numerical analog of signal detection. *Journal of Experimental Psychology: Human Learning & Memory*, **7**, 344-354.
- KOEHLER, J. J. (1996). The base-rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral & Brain Sciences*, **19**, 1-53.
- KUBOVY, M. (1977). A possible basis for conservatism in signal detection and probabilistic categorization tasks. *Perception & Psychophysics*, **22**, 277-281.
- 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.
- LEE, W., & JANKE, M. (1964). Categorizing three externally distributed stimulus samples for three continua. *Journal of Experimental Psychology*, **68**, 376-382.
- LEE, W., & JANKE, M. (1965). Categorizing externally distributed stimulus samples for unequal molar probabilities. *Psychological Reports*, **17**, 79-90.
- 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. (1996). Perceptual separability, decisional separability, and the identification-speeded classification relationship. *Journal of Experimental Psychology: Human Perception & Performance*, **22**, 795-817.
- 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., & BOGDANOV, S. V. (2000). On the relation between decision rules and perceptual representation in multidimensional perceptual categorization. *Perception & Psychophysics*, **62**, 984-997.
- 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., & 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.
- MALONEY, L. T., & THOMAS, E. A. C. (1991). Distributional assumptions and observed conservatism in the theory of signal detectability. *Journal of Mathematical Psychology*, **35**, 443-470.
- 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.
- THOMAS, E. A. C. (1975). Criterion adjustment and probability matching. *Perception & Psychophysics*, **18**, 158-162.
- THOMAS, E. A. C., & 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. There are several ways in which the weighting hypothesis might be instantiated. For example, it is possible that observers store two different decision criteria—one for accuracy maximization and one for reward maximization. On each trial, the two decision criteria might compete with one another for the opportunity to generate the categorization response. This competition could be modulated by a weighting function that emphasizes one goal more than the other (similar proposals have been offered by Ashby, Alfonso-Reese, Turken, & Waldron, 1998, and Maddox & Estes, 1996). Another possibility is that the weighting function results in a single decision criterion that is intermediate between that for accuracy maximization and that for reward maximization. A third possibility is that observers are attempting to maximize expected utility (e.g., Stevenson et al., 1991). A rigorous comparison of these possibilities is beyond the scope of this article. For ease of exposition, we assume that a single decision criterion is generated.

2. Throughout this article, it is assumed that perceptual noise is normally distributed and that the perceptual variance  $\sigma_{pi} = \sigma_p$ . Under these assumptions, the perceptual variance for each stimulus is identical. Clearly, in many cases these assumptions are incorrect (e.g., Ashby & Lee, 1991, 1993; Ashby & Maddox, 1998; Maddox & Ashby, 1996, 1998; Maddox & Bogdanov, 2000). However, with high-contrast, response-terminated displays and fairly simple stimuli, as in the present study, this is often a reasonable assumption. In addition, perceptual and criterial noise are nonidentifiable in the models outlined in this paper; thus, only the sum of perceptual plus criterial noise is estimated (Ashby, 1992a).

3. Two models are considered *nested* when one model can be derived from the other by setting some parameters of the more general model to constants. With nested models, likelihood ratio tests can be used to determine the most parsimonious model for each observer (Ashby, 1992b; Wickens, 1982). The most parsimonious model is defined as the model having the fewest free parameters that is not "significantly" improved on by a more general model (i.e., a model with more free parameters). Likelihood ratio tests should be interpreted with caution, because the test assumes independent observations that might be violated in our data.

**APPENDIX A**  
**Hit Rate (Relative to Category A) for**  
**Each Experimental Condition by Observer and Block**

Observer	Condition	$d' = 1$					$d' = 2.16$				
		1	2	3	4	5	1	2	3	4	5
1	3:1B/1:1P	.72	.84	.82	.83	.83	.92	.96	.98	.96	.93
	1:1B/3:1P	.55	.57	.58	.58	.58	.88	.88	.87	.85	.85
	3:1B/3:1P	.83	.89	.91	.91	.92	.94	.94	.96	.94	.94
	3:1B/1:3P	.92	.94	.93	.92	.91	.94	.94	.92	.90	.92
	1:3B/3:1P	.30	.27	.37	.13	.27	.73	.80	.80	.77	.70
2	3:1B/1:1P	.86	.93	.92	.93	.97	.83	.89	.92	.88	.88
	1:1B/3:1P	.82	.78	.67	.73	.85	.77	.87	.80	.90	.93
	3:1B/3:1P	.87	.88	.86	.87	.93	.92	.90	.93	.92	.97
	3:1B/1:3P	.86	.80	.86	.83	.98	.89	.87	.90	.90	.90
	1:3B/3:1P	.60	.10	.40	.27	.40	.73	.63	.63	.60	.50
3	3:1B/1:1P	.88	.83	.89	.83	.83	.88	.94	.94	.97	.88
	1:1B/3:1P	.83	.88	.85	.87	.85	.88	.95	.88	.95	.93
	3:1B/3:1P	.88	.88	.93	.93	.90	.96	.93	.92	.93	.92
	3:1B/1:3P	.84	.83	.90	.82	.82	.91	.91	.93	.90	.89
	1:3B/3:1P	.43	.53	.47	.40	.43	.73	.70	.73	.67	.70
4	3:1B/1:1P	.87	.93	.86	.88	.89	.90	.94	.93	.94	.91
	1:1B/3:1P	.87	.85	.77	.75	.80	.93	.92	.93	.93	.92
	3:1B/3:1P	.96	.98	.99	.97	.98	.94	1.00	.98	.97	.97
	3:1B/1:3P	.72	.80	.76	.74	.78	.87	.92	.91	.91	.90
	1:3B/3:1P	.57	.47	.47	.53	.60	.87	.83	.87	.83	.83
5	3:1B/1:1P	.72	.77	.83	.89	.89	.87	.91	.93	.91	.90
	1:1B/3:1P	.85	.90	.85	.82	.87	.93	.93	.90	.92	.92
	3:1B/3:1P	.90	.89	.91	.94	.91	.96	.99	.98	.97	.92
	3:1B/1:3P	.74	.70	.82	.80	.80	.87	.86	.87	.86	.83
	1:3B/3:1P	.53	.50	.50	.53	.47	.73	.87	.83	.90	.80
6	3:1B/1:1P	.89	.90	.96	.94	.98	.90	.97	.99	.96	.99
	1:1B/3:1P	.78	.88	.73	.77	.98	.87	.92	.85	.92	.93
	3:1B/3:1P	.93	.96	.96	.97	.99	.97	.97	.99	.99	.99
	3:1B/1:3P	.86	.71	.80	.80	.63	.81	.86	.87	.82	.84
	1:3B/3:1P	.80	.47	.43	.30	.70	.77	.87	.73	.73	.83
7	3:1B/1:1P	.77	.88	.86	.89	.86	.81	.84	.92	.90	.90
	1:1B/3:1P	.68	.73	.78	.77	.82	.93	.95	.92	.93	.93
	3:1B/3:1P	.87	.92	.90	.94	.92	.90	.91	.94	.89	.97
	3:1B/1:3P	.72	.67	.64	.72	.67	.82	.89	.84	.89	.91
	1:3B/3:1P	.57	.70	.70	.67	.67	.80	.83	.80	.83	.80
8	3:1B/1:1P	.72	.83	.87	.81	.90	.84	.86	.87	.84	.86
	1:1B/3:1P	.80	.78	.77	.75	.80	.88	.87	.88	.87	.88
	3:1B/3:1P	.86	.92	.92	.88	.91	.87	.88	.96	.90	.91
	3:1B/1:3P	.78	.79	.70	.72	.77	.84	.83	.82	.79	.82
	1:3B/3:1P	.70	.70	.67	.63	.53	.83	.83	.87	.83	.83
Average	3:1B/1:1P	.80	.87	.88	.88	.89	.87	.91	.94	.92	.91
	1:1B/3:1P	.77	.80	.75	.75	.82	.89	.91	.88	.91	.91
	3:1B/3:1P	.89	.91	.92	.93	.93	.93	.94	.96	.94	.95
	3:1B/1:3P	.81	.78	.80	.80	.79	.87	.88	.88	.87	.88
	1:3B/3:1P	.56	.47	.50	.43	.51	.78	.80	.78	.77	.75

(Continued on next page)

**APPENDIX B**  
**False Alarm Rate (Relative to Category A)**  
**for Each Experimental Condition by Observer and Block**

Observer	Condition	$d' = 1$					$d' = 2.16$				
		1	2	3	4	5	1	2	3	4	5
1	3:1B/1:1P	.57	.57	.60	.43	.60	.43	.27	.40	.40	.30
	1:1B/3:1P	.17	.17	.20	.23	.22	.13	.13	.15	.10	.13
	3:1B/3:1P	.70	.63	.63	.67	.70	.33	.27	.40	.37	.27
	3:1B/1:3P	.73	.77	.67	.77	.67	.27	.33	.27	.27	.23
	1:3B/3:1P	.09	.09	.12	.04	.06	.09	.10	.11	.07	.02
2	3:1B/1:1P	.73	.80	.73	.73	.87	.27	.27	.20	.27	.33
	1:1B/3:1P	.43	.55	.32	.37	.48	.17	.18	.13	.22	.35
	3:1B/3:1P	.77	.60	.57	.67	.70	.37	.37	.40	.43	.33
	3:1B/1:3P	.67	.53	.43	.57	.83	.30	.27	.30	.27	.27
3	1:3B/3:1P	.28	.19	.16	.16	.17	.03	.06	.03	.04	.06
	3:1B/1:1P	.60	.50	.73	.60	.63	.37	.33	.27	.30	.30
	1:1B/3:1P	.62	.55	.63	.55	.43	.27	.25	.27	.28	.22
	3:1B/3:1P	.63	.67	.70	.63	.70	.33	.27	.30	.30	.27
4	3:1B/1:3P	.60	.63	.63	.67	.60	.40	.27	.23	.23	.23
	1:3B/3:1P	.17	.24	.21	.19	.18	.12	.04	.06	.04	.03
	3:1B/1:1P	.73	.77	.73	.73	.70	.27	.30	.43	.43	.30
	1:1B/3:1P	.50	.50	.43	.42	.42	.28	.27	.27	.28	.25
	3:1B/3:1P	.83	.87	.83	.90	.90	.40	.47	.50	.53	.50
5	3:1B/1:3P	.23	.47	.33	.43	.57	.27	.23	.30	.40	.20
	1:3B/3:1P	.31	.23	.14	.24	.24	.18	.30	.24	.19	.17
	3:1B/1:1P	.30	.57	.57	.67	.70	.30	.27	.30	.30	.23
	1:1B/3:1P	.37	.52	.53	.47	.50	.17	.22	.20	.27	.25
	3:1B/3:1P	.77	.77	.83	.70	.70	.53	.60	.43	.47	.37
6	3:1B/1:3P	.33	.33	.53	.50	.60	.20	.23	.27	.23	.17
	1:3B/3:1P	.17	.29	.27	.21	.22	.12	.19	.20	.17	.10
	3:1B/1:1P	.70	.83	.77	.67	.87	.23	.43	.33	.33	.53
	1:1B/3:1P	.38	.65	.65	.68	.97	.15	.28	.27	.30	.30
	3:1B/3:1P	.77	.90	.90	.90	1.00	.33	.53	.57	.50	.67
7	3:1B/1:3P	.53	.40	.40	.37	.33	.13	.23	.20	.13	.23
	1:3B/3:1P	.36	.37	.38	.17	.29	.16	.14	.16	.13	.17
	3:1B/1:1P	.47	.53	.70	.67	.63	.23	.27	.23	.23	.23
	1:1B/3:1P	.37	.40	.35	.37	.38	.17	.27	.27	.33	.30
	3:1B/3:1P	.73	.77	.77	.90	.77	.40	.23	.33	.30	.33
8	3:1B/1:3P	.43	.30	.20	.47	.53	.20	.23	.23	.27	.23
	1:3B/3:1P	.27	.38	.31	.34	.33	.16	.09	.16	.12	.10
	3:1B/1:1P	.43	.57	.53	.70	.60	.20	.20	.23	.20	.20
	1:1B/3:1P	.38	.43	.43	.38	.37	.15	.15	.17	.17	.13
	3:1B/3:1P	.77	.73	.80	.80	.77	.37	.27	.37	.30	.30
Average	3:1B/1:3P	.40	.40	.47	.33	.53	.20	.27	.23	.17	.23
	1:3B/3:1P	.33	.38	.32	.38	.28	.16	.14	.16	.14	.19
	3:1B/1:1P	.57	.64	.67	.65	.70	.29	.29	.30	.31	.30
	1:1B/3:1P	.40	.47	.44	.43	.47	.19	.22	.21	.24	.24
	3:1B/3:1P	.75	.74	.75	.77	.78	.38	.38	.41	.40	.38
Average	3:1B/1:3P	.49	.48	.46	.51	.58	.25	.26	.25	.25	.23
	1:3B/3:1P	.25	.27	.24	.22	.22	.13	.13	.14	.11	.10

(Manuscript received July 7, 1999;  
revision accepted for publication August 10, 2000.)