

# A test of the optimal classifier's independence assumption in perceptual categorization

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

Observers completed perceptual categorization tasks that included separate base-rate/payoff manipulations, corresponding simultaneous base-rate/payoff manipulations, and conflicting simultaneous base-rate/payoff manipulations. Performance (1) was closer to optimal for 2:1 than for 3:1 base-rate/payoff ratios and when base rates as opposed to payoffs were manipulated, and (2) was more in line with the predictions from the flat-maxima hypothesis than from the independence assumption of the optimal classifier in corresponding and conflicting simultaneous base-rate/payoff conditions. A hybrid model that instantiated simultaneously the flat-maxima and the competition between reward and accuracy maximization (COBRA) hypotheses was applied to the data. The hybrid model was superior to a model that incorporated the independence assumption, suggesting that violations of the independence assumption are to be expected and are well captured by the flat-maxima hypothesis without requiring any additional assumptions. The parameters indicated that observers' reward-maximizing decision criterion rapidly approaches the optimal value and that more weight is placed on accuracy maximization in separate and corresponding simultaneous base-rate/payoff conditions than in conflicting simultaneous base-rate/payoff conditions.

Categorization is fundamental to the survival of all organisms (Ashby & Maddox, 1998). For example, the deer must decide whether to continue feeding or to run when it detects a sound that could be leaves rustling or a human approaching. Similarly, the mountaineer must decide whether to continue climbing toward the summit or to retreat on the basis of a quick evaluation of weather conditions. Optimal categorization performance that maximizes long-run reward often requires the placement of a decision criterion along some relevant dimension. For example, a criterion might be set on the loudness of the leaves rustling, with loudness levels below the criterion leading to a human-absent decision and levels above the criterion leading to a human-present decision. The location of the optimal decision criterion is affected by the category base rates (i.e., the prior probability of each category) and the entries in the category payoff matrix (i.e., the costs and benefits associated with correct and incorrect categorization responses). Category base rates often differ across categorization situations. For example, during the off season, a human's presence has a low base-rate probability, but during hunting season the probability that a human is present increases, and so the deer might lower the loudness criterion associated with a human-present

decision. The costs and benefits associated with each categorization decision might also differ. We generally benefit when we make the correct decision, but the benefit of a correct human-present decision will be greater than the benefit of a correct human-absent decision, especially during hunting season, when the human might be armed. Similarly, there is often a cost associated with an incorrect decision; a more severe cost would be incurred for an incorrect human-absent decision than for an incorrect human-present decision, especially during hunting season. Under these conditions, the deer would be wise to lower the loudness criterion associated with a human-present decision to increase the number of correct human-present decisions, even though this would adversely affect the accuracy of human-absent decisions.<sup>1</sup>

Base rates and payoffs vary widely in real-world categories, and in many (likely most) cases, both differ simultaneously within the same categorization problem. Base rates and payoffs can bias the decision criterion in the same direction; for example, when the base-rate probability of humans is high and the benefit of a correct human-present response is high, the loudness criterion will likely be lowered. We refer to these as *corresponding* simultaneous base-rate/payoff conditions. In other cases, the base-rates and payoffs bias the decision criterion in different directions. For example, whereas the base-rate probability of cancer is relatively low, the benefit of a correct cancer diagnosis could save a life. In this case, the low base-rate probability biases the doctor toward a no-cancer diagnosis, but the high benefit of a correct cancer diagnosis biases the doctor toward cancer diagnoses. We refer to these as *conflicting* simultaneous base-rate/payoff conditions.

---

This research was supported in part by National Science Foundation Grant SBR-9796206 and Grant 5 R01 MH59196 from the National Institute of Mental Health, National Institutes of Health. We thank Neil Macmillan and three anonymous reviewers for helpful comments on an earlier version of this manuscript. We also thank Amy Kaderka for help with data collection. Correspondence should be addressed to W. T. Maddox, Department of Psychology, University of Texas, 1 University Station A8000, Austin, TX 78712 (e-mail: maddox@psy.utexas.edu).

Despite the prevalence of simultaneous base-rate/payoff differences in the environment, few empirical studies of decision criterion learning under simultaneous base-rate/payoff conditions have been undertaken, and even fewer psychologically motivated theories have been proposed. The overriding goal of this study was twofold. First, we examined human decision criterion learning when base-rates and payoffs were manipulated separately and compared performance with cases in which base-rates and payoffs were manipulated simultaneously across a range of base-rate/payoff ratios for both corresponding and conflicting base-rate/payoff conditions. Second, we generated and tested predictions from Maddox and Dodd's (2001) hybrid model of decision criterion placement for the separate and simultaneous base-rate/payoff conditions under investigation, and compared these with the independence assumption of the optimal classifier (described in detail below). Although human decision criterion learning rarely matches that of the optimal classifier, the optimal classifier provides a useful benchmark against which to compare human performance.

First, we briefly outline the optimal classifier and our modeling framework, both based on decision bound theory. Then we introduce Maddox and Dodd's (2001) theory of decision criterion learning, generating predictions from the model for several simultaneous base-rate/payoff conditions. Next, we outline the experiment, method, and results. Finally, we conclude with some general comments.

## THE OPTIMAL CLASSIFIER AND DECISION BOUND THEORY

### Optimal Classifier

The optimal classifier is a hypothetical device that maximizes long-run expected reward. Suppose a medical doctor must classify a patient into one of two disease categories, A or B, on the basis of Medical Test X,

whose outcomes for Diseases A and B are normally distributed as depicted in Figure 1. The optimal classifier has perfect knowledge of the form and parameters of each category distribution and records perfectly the test result, denoted  $x$ . This information is used to construct the *optimal decision function*, which is 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  $i$ . The optimal classifier has perfect knowledge of the category base rates and the payoffs associated with correct diagnoses. This information is used to construct the *optimal decision criterion*:

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

where  $P(A)$  and  $P(B)$  are the base-rate probabilities for Categories A and B, and  $V_{aA}$  and  $V_{bB}$  denote the payoffs associated with correct diagnoses.<sup>2</sup> (In the experiment outlined below, the cost of an incorrect response was set to zero, and so the error terms  $V_{bA}$  and  $V_{aB}$  were excluded from Equation 2.) The optimal classifier (e.g., Green & Swets, 1966) uses  $l_o(x)$  and  $\beta_o$  to construct the *optimal decision rule*:

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

Three points are in order. First, when  $P(A)(V_{aA}) = P(B)(V_{bB})$  then  $\beta_o = 1$  and the optimal classifier assigns the stimulus to the category with the highest likelihood. Second, base-rate and payoff manipulations have the same effect on the optimal decision criterion. For example, if Disease A is three times as common as Disease B—a 3:1 base-rate condition—or the payoff for Disease A is three times the payoff for Disease B—a 3:1 payoff condition [i.e., if  $P(A) = 3P(B)$  or  $(V_{aA}) = 3(V_{bB})$ ]—then  $\beta_o = 3.0$  (Figure 1). In this case, the optimal classifier will gener-




Figure 1. Hypothetical distributions for Categories A and B. The  $\beta_o = 1$  decision criterion denotes the criterion that is optimal when the base rates and payoffs are equal. This is also referred to as the equal likelihood decision criterion. The  $\beta_o = 3$  decision criterion denotes the criterion that is optimal when there is a 3:1 base-rate ratio or a 3:1 payoff ratio.

ate a Disease A diagnosis unless the likelihood of Disease B is at least three times larger than the likelihood of Disease A. Third, when base rates and payoffs are manipulated simultaneously, the optimal decision criterion can be derived from an independent combination of the separate base-rate and payoff decision criteria. This is seen more clearly if we use the alternative (but mathematically equivalent) formulation of Equation 2 where we apply the natural log to both sides, yielding

$$\ln\beta_0 = \ln[P(A)/P(B)] + \ln[(V_{aA}/V_{bB})]. \quad (4)$$

Notice that  $\ln\beta_0$  is determined completely by the sum of an independent base-rate and payoff term and is referred to as the independence assumption of the optimal classifier.

### Decision Bound Theory

The optimal classifier decision rule (Equation 3) has been rejected as a model of human performance, but performance often approaches that of the optimal classifier as the observer gains experience with the task. Ashby and colleagues argued that the observer attempts to respond using a strategy similar to that of the optimal classifier, but fails because of at least two sources of suboptimality in perceptual and cognitive processing (Ashby, 1992a; Ashby & Lee, 1991; Ashby & Maddox, 1993, 1994; Ashby & Townsend, 1986; Maddox & Ashby, 1993; see Wallsten & Gonzalez-Vallejo, 1994, for a similar approach). These include perceptual and criterial noise. Perceptual noise refers to trial-by-trial variability in the perceptual information associated with each stimulus. With one perceptual dimension the observer's percept of stimulus  $i$ , on any trial, is given by  $x_{pi} = x_i + e_p$ , where  $x_i$  is the observer's mean percept and  $e_p$  is a random variable denoting perceptual noise (we assume that  $\sigma_{pi} = \sigma_p$ ). Criterial noise refers to trial-by-trial variability in the observer's memory for the decision criterion. With criterial noise the decision criterion used on any trial is given by  $\beta_c = \beta + e_c$ , where  $\beta$  is the observer's average decision criterion and  $e_c$  is a random variable denoting criterial noise (assumed to be univariate normally distributed). Decision bound theory assumes that the observer attempts to use the same strategy as the optimal classifier, but with less success due to the effects of perceptual and criterial noise. Hence, the simplest decision bound model is the *optimal decision bound model*. The optimal decision bound model is identical to the optimal classifier (Equation 3) except that perceptual and criterial noise are incorporated into the decision rule. Specifically,

$$\text{If } I_0(x_{pi}) > \beta_0 + e_c \text{ then respond "B,"} \\ \text{otherwise respond "A."} \quad (5)$$

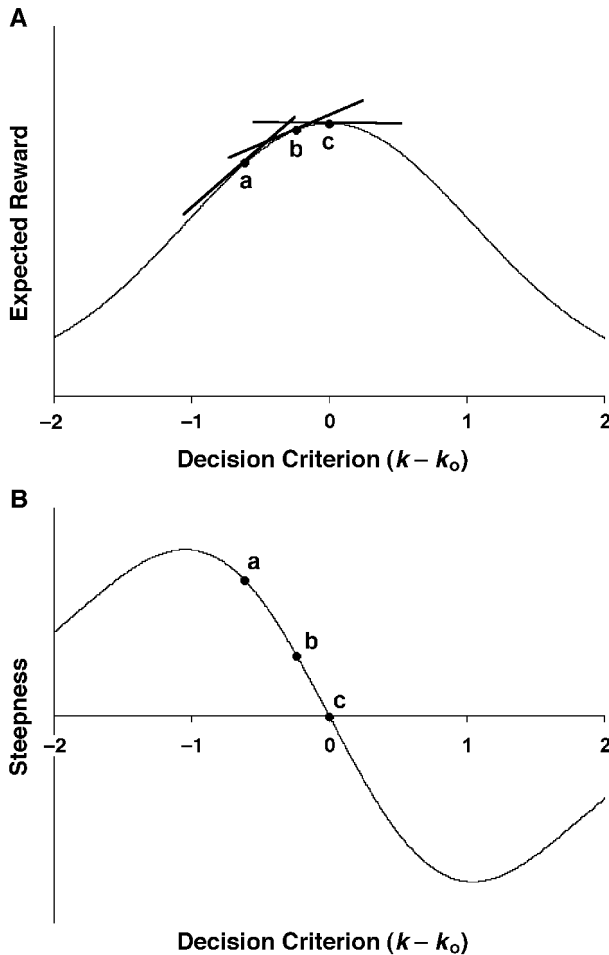
## A THEORY OF DECISION CRITERION LEARNING AND A HYBRID MODEL FRAMEWORK

Two robust findings emerged from a series of studies that examined decision criterion learning under unequal

base-rate or payoff conditions (e.g., Green & Swets, 1966; Healy & Kubovy, 1981; Kubovy & Healy, 1977; Lee & Janke, 1964, 1965; Lee & Zentall, 1966). First, observers tended to use a decision criterion that was more conservative than the optimal decision criterion. For example, if the base rates or payoffs were such that  $\beta_0 = 3$ , then observers tended to use a  $\beta$  between 1 and 3. This was termed *conservative cutoff placement*, because the decision criterion was not shifted far enough toward the optimal value. Second, observers' decision criterion estimates were closer to the optimal value when base rates as opposed to payoffs were manipulated, even when the optimal decision criterion was identical across base-rate and payoff conditions. Several explanations for these results have been offered in the literature (see Healy & Kubovy, 1981; Maddox & Bohil, 2000), but none has been generally accepted. Maddox and Dodd (2001; Maddox, 2002) offered a theory of decision criterion learning and a model-based instantiation that predicts these two results. The theory proposes two mechanisms (described next) that determine decision criterion placement.

### Flat-Maxima Hypothesis

The first mechanism is based on the *flat-maxima hypothesis* (Busemeyer & Myung, 1992; von Winterfeldt & Edwards, 1982) and was developed to account for the finding that observers tend to use a criterion that is more conservative than the optimal decision criterion when base rates or payoffs are manipulated (i.e., conservative cutoff placement). As suggested by many researchers, suppose that the observer adjusts the decision criterion based (at least in part) on the change in the rate of reward, with larger changes in rate being associated with faster, more nearly optimal, decision criterion learning (e.g., Busemeyer & Myung, 1992; Dusoier, 1980; Erev, 1998; Erev, Gopher, Itkin, & Greenspan, 1995; Kubovy & Healy, 1977; Roth & Erev, 1995; Thomas, 1975; Thomas & Legge, 1970). To formalize this hypothesis one can construct the *objective reward function*. The objective reward function plots objective expected reward on the y-axis and the decision criterion value on the x-axis (e.g., Busemeyer & Myung, 1992; Stevenson, Busemeyer, & Naylor, 1991; von Winterfeldt & Edwards, 1982). To generate an objective reward function one chooses a value for the decision criterion and computes the expected reward for that criterion value. This process is repeated over a range of criterion values. The expected reward is then plotted as a function of decision criterion value. Figure 2A plots expected reward as a function of the deviation between a hypothetical observer's decision criterion  $\ln(\beta)$  and the optimal decision criterion  $\ln(\beta_0)$  standardized by category  $d'$  where  $d'$  is defined as the standardized distance between the category means ( $d' = 2.2$  in Figure 1). This is referred to as  $k - k_0 = \ln(\beta)/d' - \ln(\beta_0)/d'$ .<sup>3</sup> Notice that for large deviations from the optimal decision criterion, the expected reward is small, and as the deviation from the op-



**Figure 2.** (A) Expected reward as a function of the decision criterion (relative to the optimal decision criterion; i.e.,  $k - k_0$ ), called the objective reward function for category discriminability,  $d' = 1.0$ . The three tangent lines at Points a, b, and c on the objective reward function denote the steepness (i.e., derivative) of the function at each point. (B) Steepness of the objective reward functions from Panel A along with the three points highlighted in Panel A.

timal decision criterion decreases, the expected reward increases. Notice also that when the deviation from optimal is zero (i.e., when the decision criterion is the optimal decision criterion), expected reward is maximized.

The derivative of the objective reward function at a specific  $k - k_0$  value determines the change in the rate of expected reward for that  $k - k_0$  value; the larger the change in the rate, the “steeper” the objective reward function at that point. Derivatives for three  $k - k_0$  values are indicated by the three tangent lines denoted as a, b, and c in Figure 2A. Notice that the slope of each tangent line, which corresponds to the derivative of the objective reward function at that point, decreases as the deviation from the optimal decision criterion decreases (i.e., as we go from Point a to b to c). In other words, the change in the rate of reward or steepness declines as the decision criterion approaches the optimal decision criterion. Fig-

ure 2B plots the relationship between the steepness of the objective reward function (i.e., the derivative at several  $k - k_0$  values) and  $k - k_0$ . The three derivatives denoted in Figure 2A are highlighted in Figure 2B. If the observer adjusts the decision criterion on the basis of the change in the rate of reward (or steepness), then steeper objective reward functions should be associated with more nearly optimal decision criterion values, because only a small range of decision criterion values around the optimal value have nearly zero derivatives (or small steepness values). Flat objective reward functions, on the other hand, will lead to less optimal decision criterion placement because a larger range of decision criterion values around the optimal value have derivatives near zero. Interestingly, nearly all previous work used tasks where the category  $d' = 1.0$ . The objective reward function is very shallow for  $d' = 1.0$ , which might partially explain the prevalence of conservative cutoff placement (Figure 2A).

Because the flat-maxima hypothesis is based on the objective reward function, it applies only to learning of the *reward-maximizing* decision criterion. As we will outline shortly, the observed decision criterion is assumed to be a weighted average of the reward- and accuracy-maximizing decision criteria. Although the flat-maxima hypothesis offers an explanation for the prevalence of conservative cutoff placement, it does not offer an explanation for the robust finding that the observed decision criterion is closer to optimal in base-rate, as compared with payoff, conditions.

### Competition Between Reward and Accuracy Maximization (COBRA) Hypothesis

The second mechanism assumed to influence decision criterion placement is based on Maddox and Bohil's (1998) COBRA hypothesis, which was developed to account for the finding that observers show more nearly optimal decision criterion placement in unequal base-rate than in unequal payoff conditions. COBRA postulates that observers attempt to maximize expected reward (consistent with instructions and monetary compensation contingencies), but they also place importance on accuracy maximization. Consider the univariate categorization problems depicted in Figure 3. Panel A displays a 3:1 base-rate condition and Panel B displays a 3:1 payoff condition. As suggested by Equation 2, expected reward is maximized in both cases by using the optimal reward-maximizing decision criterion,  $k_{r_0} = \ln(\beta_{r_0})/d' = \ln(3)/d'$ . Thus, an observer who attempts to maximize expected reward should use the same decision criterion in both conditions. In the 3:1 base-rate condition, the decision criterion that maximizes reward also maximizes accuracy, so  $k_{a_0} = k_{r_0}$ . However, in the 3:1 payoff condition the accuracy-maximizing decision criterion is different from the reward-maximizing decision criterion. Specifically, in the 3:1 payoff condition,  $k_{a_0} = \ln(\beta_{a_0})/d' = \ln(1)/d'$ . (When base rates are equal, it is always the case that the accuracy-maximizing decision criterion  $\beta_{a_0} = 1$ .) Thus, when base rates are manipulated, accuracy and re-

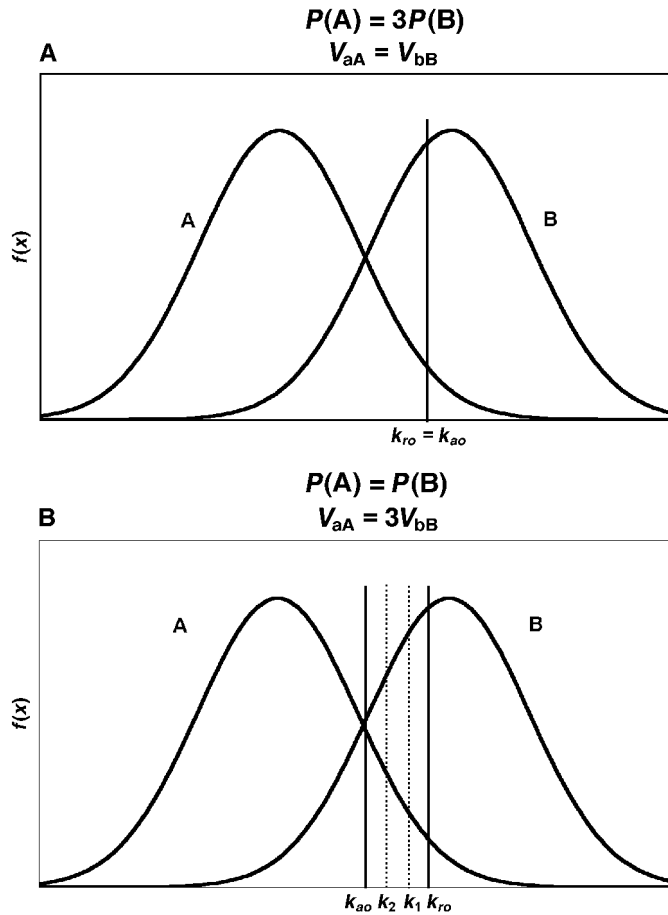


Figure 3. Schematic illustration of the competition between reward and accuracy maximization (COBRA) hypothesis. (See text for details.)

ward can be maximized simultaneously, since  $k_{ro} = k_{a0}$ , but when payoffs are manipulated, both goals *cannot* be achieved simultaneously since  $k_{ro} \neq k_{a0}$ . An observer who places importance (or weight) on both goals will use a decision criterion intermediate between the accuracy- and reward-maximizing decision criteria in the payoff condition, and thus will show more conservative cutoff placement in the payoff condition than in the base-rate condition. To instantiate this hypothesis 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 expected accuracy maximization. This weighting function results in a single decision criterion that is intermediate between that for accuracy maximization and that for reward maximization.<sup>4</sup> For example, in Figure 3B,  $k_1$  denotes a case in which  $w < .5$ , whereas  $k_2$  denotes a case in which  $w > .5$ .

**Flat-Maxima and Independence Predictions for Simultaneous Base-Rate/Payoff Conditions**

Recall from Equation 4 that  $\ln\beta_0$  for the optimal classifier is determined completely by the sum of an independent base rate and payoff term. This is referred to as

the independence assumption of the optimal classifier. Only a few studies have tested the independence assumption of the optimal classifier in human decision criterion learning (Bohil & Maddox, 2001; Maddox & Bohil, 1998, 2001; Stevenson et al., 1991). In an early test, Stevenson et al. plotted  $\ln(\beta)$  across base-rate conditions separately for each payoff condition using data collected by Healy and Kubovy (1981). If independence is satisfied, then the line plots should be linear and parallel. In analysis of variance (ANOVA) terms there would be *no* interaction. Stevenson et al. found generally linear, parallel line plots in support of the independence assumption. One weakness of this approach is that support for the independence assumption relies on the dangerous proposition of accepting the null hypothesis. In addition, when independence is violated, no specific alternative hypotheses that might explain potential violations of independence have been offered. In light of these facts, Maddox and Bohil (in press) developed a stronger test of the independence hypothesis through the application of a series of quantitative models—some that embodied the independence assumption of the optimal classifier and some that embodied the flat-maxima hypothesis—to

decision criterion learning data from a 3:1 base-rate, a 3:1 payoff, and a 3:1 base-rate/3:1 payoff condition (each was run at three levels of  $d'$ ). The flat-maxima hypothesis predicts that decision criterion placement in the 3:1 base-rate/3:1 payoff condition should differ systematically from that predicted from the independence assumption. The nature of this differential prediction is outlined next. In short, Maddox and Bohil (2001) found support for the models that embodied the flat-maxima hypothesis over the models that embodied the independence assumption.

Ten base-rate/payoff conditions were included in the present study. Four *separate* base-rate/payoff conditions were included: a 2:1 base-rate (2:1B), a 3:1 base-rate (3:1B), a 2:1 payoff (2:1P), and a 3:1 payoff (3:1P) condition. Four *corresponding* simultaneous base-rate/payoff conditions were included: a 2:1 base-rate/2:1 payoff (2:1B/2:1P), a 3:1 base-rate/3:1 payoff (3:1B/3:1P), a 2:1 base-rate/3:1 payoff (2:1B/3:1P), and a 3:1 base-rate/2:1 payoff (3:1B/2:1P) condition. Finally, two *conflicting* simultaneous base-rate/payoff conditions were included: a 1:2 base-rate/3:1 payoff (1:2B/3:1P) and a 3:1 base-rate/1:2 payoff (3:1B/1:2P) condition. To derive the flat-maxima hypothesis and independence assumption predictions we took the following approach. First, we generated the objective reward function and steepness functions for all 10 experimental conditions (on the basis of the method used to generate Figures 2A and 2B). Second, for a steepness value equal to zero, we derived the  $\ln(\beta)$  predicted in the four separate base-rate/payoff conditions. Third, for the same steepness value (zero in this case) we derived the  $\ln(\beta)$  predicted in the six simultaneous base-rate/payoff conditions. Because we used the same steepness to derive the separate *and* simultaneous base-rate/payoff conditions, these  $\ln(\beta)$  values are those predicted from the flat-maxima hypothesis. Fourth, we derived the  $\ln(\beta)$  values for the six simultaneous base-rate/payoff conditions from the independence assumption. Because a steepness value of zero is associated with the optimal decision criterion, under these conditions the flat-maxima hypothesis and the independence assumption make identical predictions. For example, in the 2:1B/2:1P condition, zero steepness is associated with the optimal decision criterion value  $\ln(\beta_o) = \ln(4) = 1.386$ . Similarly, in the 2:1B and 2:1P conditions, zero steepness is associated with the optimal decision criterion value  $\ln(\beta_o) = \ln(2) = .693$ . Assuming independence, the 2:1B/2:1P decision criterion value would be  $\ln(2) + \ln(2) = .693 + .693 = 1.386$ , which is identical to that predicted from the flat-maxima hypothesis.


The more interesting cases emerged when we repeated this process for a large range of positive steepness values. We selected only steepness values that resulted in  $\ln(\beta)$  values between the optimal value (steepness = 0) and  $\ln(\beta) = \ln(1) = 0$ . The results are summarized in Figure 4, with the solid line denoting the relationship between the flat-maxima and independence predictions and the broken line (included for comparative purposes)

denoting situations in which the two hypotheses make identical predictions. (Identical plots resulted for the 2:1B/3:1P and 3:1B/2:1P and for the 1:2B/3:1P and 3:1B/1:2P cases and so these were collapsed.) The  $x$ -axis of each panel displays criterion values derived from an independent combination of decision criteria from separate base-rate and payoff conditions. The  $y$ -axis denotes the simultaneous base-rate/payoff criterion values that share an objective reward function steepness with the separate base-rate and payoff criterion values (i.e., predicted from the flat-maxima hypothesis). If the independence assumption and the flat-maxima hypothesis predict the same simultaneous base-rate/payoff criterion values, then the solid line should fall on top of the dashed line. Two main points emerge from an examination of Figure 4. First, in the corresponding simultaneous base-rate/payoff conditions (Panels A–C), the decision criterion predicted from the flat-maxima hypothesis is always larger than that predicted from the independence assumption of the optimal classifier, and thus is closer to the optimal decision criterion (except when steepness is zero, in which case both the flat-maxima hypothesis and the independence assumption make identical predictions). Second, the opposite pattern emerges in the conflicting simultaneous base-rate/payoff conditions (Panel D). Specifically, the decision criterion predicted from the flat-maxima hypothesis is always smaller than that predicted from the independence assumption, and thus is farther from the optimal decision criterion (again, except when steepness is zero).

We tested these predictions in two ways. First, we compared the observed signal detection theory decision criterion estimates from each simultaneous base-rate/payoff condition with those predicted from the independence assumption. Second, we applied a series of quantitative models, based on Maddox and Dodd's (2001) hybrid model, to the data from each individual observer that instantiated either the flat-maxima hypothesis or the independence assumption in the simultaneous base-rate/payoff conditions. To anticipate, the results provide strong support for the flat-maxima hypothesis over the independence assumption and suggest that the hybrid model provides a good account of the data from all conditions, including the corresponding and conflicting simultaneous base-rate/payoff conditions, without the need to postulate any additional mechanisms. This hybrid model framework is outlined next.

### The Hybrid Model Framework

Maddox and Dodd (2001) developed a hybrid model of decision criterion learning that incorporated both the flat-maxima and COBRA hypotheses. Specifically, the model assumes that the decision criterion used by the observer to maximize expected reward ( $k_r$ ) is determined by the steepness of the objective reward function (Figure 2). A single steepness parameter is estimated from the data that determines a distinct decision criterion in every condition that has a unique objective reward function. The



**Figure 4.** Decision criterion  $[\ln(\beta)]$  predicted from the flat-maxima hypothesis plotted against the decision criterion  $[\ln(\beta)]$  predicted from the independence assumption of the optimal classifier for the six simultaneous base-rate/payoff conditions. (A) 2:1B/2:1P condition. (B) 3:1B/3:1P condition. (C) 2:1B/3:1P and 3:1B/2:1P conditions. (D) 1:2B/3:1P and 3:1B/1:2P conditions. (See text for details of procedure used to generate predicted decision criterion values.)

COBRA hypothesis is instantiated in the hybrid model by estimating the accuracy weight,  $w$ , from the data. To facilitate the development of each model, consider the following equation, which determines the decision criterion used by the observer on condition  $i$  trials ( $k_i$ ):

$$k_i = wk_{ai} + (1 - w)k_{ri}. \tag{6}$$

Importantly, when base rates are manipulated, the observer’s estimate of the reward-maximizing decision criterion, derived from the flat-maxima hypothesis, is also the best estimate of the accuracy-maximizing decision criterion, resulting not in competition but simply in use of the reward-maximizing decision criterion. When payoffs are manipulated, on the other hand, the reward- and accuracy-maximizing decision criteria differ from each another. By pretraining each observer on the category structures in a baseline condition (described shortly in the Method section), we essentially pretrain the accuracy-maximizing decision criterion. This criterion is then entered into the weighting function along with the observer’s estimate of the reward-maximizing decision

criterion to determine the criterion used on each trial in the experimental conditions.

All of the models developed in this article are based on the decision bound model in Equation 5. Specifically, each model includes one “noise” parameter that represents the sum of perceptual and criterial noise (Ashby, 1992a; Maddox & Ashby, 1993). Each model assumes that the observer has accurate knowledge of the category structures [i.e.,  $l_o(x_{pi})$ ]. To ensure that this was a reasonable assumption, each observer completed a number of baseline trials and was required to meet a stringent performance criterion (see Method section). Finally, each model allows for suboptimal decision criterion placement where the decision criterion is determined from the flat-maxima hypothesis, the COBRA hypothesis, or both, following Equation 6. To determine whether the flat-maxima and COBRA hypotheses are important in accounting for each observer’s data, we developed four models. Each model makes different assumptions about the  $k_r$  and  $w$  values used. The nested structure of the models is represented in Figure 5, with each arrow pointing to a more general model and

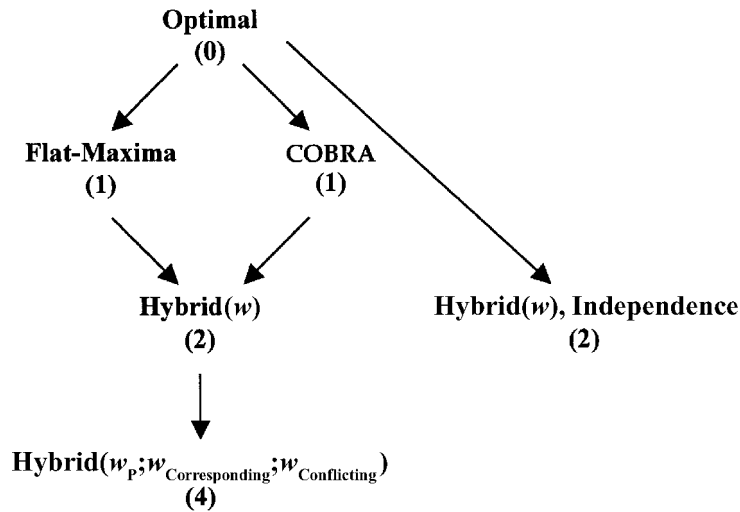


Figure 5. Nested relationship among the decision bound models applied simultaneously to the data from all experimental conditions. The number in parentheses denotes the number of free parameters in addition to the noise parameter. The arrows point to a more general model. (See text for details.)

models at the same level having the same number of free parameters. The number of free parameters (in addition to the noise parameter described above) is presented in parentheses. (The details of the model fitting procedure are outlined in the Results section.)

The *optimal model* instantiates neither the flat-maxima nor the COBRA hypotheses by assuming that the decision criterion used by the observer to maximize expected reward is the optimal decision criterion (i.e.,  $k_r = k_o$ ), and that there is no competition between reward and accuracy maximization (i.e.,  $w = 0$ ). The *flat-maxima model* instantiates the flat-maxima hypothesis, but not the COBRA hypothesis, by assuming that the decision criterion used by the observer to maximize expected reward ( $k_r$ ) is determined by the steepness of the objective reward function, and that there is no competition between reward and accuracy maximization (i.e.,  $w = 0$ ). A single steepness parameter is estimated from the data. This single steepness parameter determines the following six distinct  $k_r$  values: (1) 2:1B and 2:1P, (2) 3:1B and 3:1P, (3) 2:1B/2:1P, (4) 3:1B/3:1P, (5) 2:1B/3:1P and 3:1B/2:1P, and (6) 1:2B/3:1P and 3:1B/1:2P. Because  $k_r$  is determined from the steepness of the objective reward function, this model is constrained to predict that the  $k_r$  value in the corresponding conditions will be closer to optimal than those predicted from the independence assumption of the optimal classifier, and the  $k_r$  value in the conflicting conditions will be farther from optimal than those predicted from the independence assumption of the optimal classifier. Although not discussed up to this point, it is straightforward to show that the flat-maxima hypothesis also predicts more nearly optimal reward-maximizing decision criterion placement in 2:1 than in 3:1 base-rate or payoff conditions, and so one additional prediction is that the  $k_r$  value in the 2:1B and 2:1P conditions will be

closer to optimal than the  $k_r$  value in the 3:1B and 3:1P conditions. The flat-maxima model contains the optimal model as a special case. The *COBRA model* instantiates the COBRA hypothesis but not the flat-maxima hypothesis by assuming that  $k_r = k_o$ , while allowing for a competition between reward and accuracy maximization by estimating the Equation 6  $w$  parameter from the data. This model contains the optimal model as a special case. The hybrid( $w$ ) model instantiates both the flat-maxima and the COBRA hypotheses by assuming that  $k_r$  is determined by the steepness of the objective reward function, and that there is a competition between accuracy and reward maximization. This model contains the other three models as special cases. One more general version of the hybrid model was applied to the data. The hybrid( $w_p; w_{\text{Corresponding}}; w_{\text{Conflicting}}$ ) model estimated three accuracy weights. One was applied to the two separate payoff conditions (i.e., 2:1P and 3:1P), a second was applied to the four corresponding simultaneous base-rate/payoff conditions (i.e., 2:1B/2:1P, 3:1B/3:1P, 2:1B/3:1P, and 3:1B/2:1P), and a third was applied only to the two conflicting simultaneous base-rate/payoff conditions (i.e., 1:2B/3:1P and 3:1B/1:2P). This model contains the previously described models as special cases.

A variant of the hybrid( $w$ ) model that instantiated the independence assumption was also developed. The hybrid( $w$ )/independence model applied the hybrid model framework to determine the decision criteria in the separate base-rate and payoff conditions (i.e., 2:1B, 2:1P, 3:1B, and 3:1P) and then combined these independently to derive the decision criteria in the simultaneous base-rate/payoff conditions. Note that the hybrid( $w$ )/independence model has the same number of parameters as the hybrid( $w$ ) model and so the fits can be compared directly. Note also that the hybrid( $w$ )/independence model contains the op-

timal model, but no other model, as a special case. [An “independence” variant of the hybrid ( $w_P$ ;  $w_{\text{Corresponding}}$ ;  $w_{\text{Conflicting}}$ ) model cannot be developed because it contains accuracy weights associated specifically with the simultaneous base-rate/payoff conditions.] Another independence model was developed, but we reserve discussion of that model for the Results section.

## EXPERIMENT

This article reports the results of a categorization experiment that examined decision criterion learning in four separate (2:1B, 2:1P, 3:1B, and 3:1P), four corresponding simultaneous (2:1B/2:1P, 3:1B/3:1P, 2:1B/3:1P, and 3:1B/2:1P), and two conflicting simultaneous (1:2B/3:1P and 3:1B/1:2P) base-rate/payoff conditions. Including both corresponding and conflicting simultaneous base-rate/payoff conditions with 2:1 and 3:1 base-rate/payoff ratios leads to powerful tests of the flat-maxima hypothesis and the independence assumption predictions outlined in Figure 4, and to a first-ever test of the flat-maxima hypothesis prediction that decision criterion placement should be more nearly optimal in 2:1 than in 3:1 conditions. To examine decision criterion learning across separate and simultaneous base-rate/payoff conditions and to test the flat-maxima and independence assumption predictions, we had each observer complete several blocks of trials in all 10 base-rate/payoff conditions. First, we examined trends in decision criterion estimates from signal detection theory (and other performance measures), and followed these analyses with the application of a series of quantitative models applied to the data from all 10 experimental conditions *simultaneously*, although separately for each observer and block of trials. Each model instantiated a different set of hypotheses about the effects of base-rate/payoff condition and the validity of the flat-maxima hypothesis and independence assumption on decision criterion placement. All analyses were performed at the individual observer level because of concerns with modeling aggregate data (e.g., Ashby, Maddox, & Lee, 1994; Estes, 1956; Maddox, 1999; Maddox & Ashby, 1998; Smith & Minda, 1998).

### Method

**Observers.** Sixteen observers were solicited from the University of Texas community. All observers claimed to have 20/20 vision or vision corrected to 20/20. Each observer completed five approximately 60-min sessions. In each session the observer completed two base-rate/payoff conditions. The order of the 10 conditions was determined from a Latin square. Observers were paid on the basis of their performance in the task.

**Stimuli and stimulus generation.** The stimulus was a filled white rectangular bar that varied in height from trial to trial (40 pixels wide) set flush upon a stationary base (60 pixels wide). There were two categories of bar heights, A and B, each defined by a specific univariate normal distribution (Ashby & Gott, 1988). The separation between the Category A and B means was 45 pixels, with a standard deviation of 21 pixels for both Category A and B, yielding a category  $d' = 2.2$ . Four sets of 60 stimuli were generated prior to the experiment. One set contained 30 A and 30 B stimuli and was

used in all cases for which the base rates were equal—that is, the baseline (to be described shortly), the 2:1P, and the 3:1P conditions. The second contained 40 A and 20 B stimuli and was used in all cases with a 2:1 base rate—that is, the 2:1B, the 2:1B/2:1P, and the 2:1B/3:1P conditions. The third contained 45 A and 15 B stimuli and was used in all cases with a 3:1 base rate—that is, the 3:1B, the 3:1B/3:1P, the 3:1B/2:1P, and the 3:1B/1:2P conditions. The fourth contained 20 A and 40 B stimuli and was used in the 1:2B/3:1P condition. Each stimulus set was generated by taking numerous random samples (of Size 60) from the population and by selecting the sample whose objective reward function was most similar (on the basis of visual inspection) to that derived from the population. Two measures were taken to discourage information transfer across conditions. First, and most importantly, before being allowed to begin each of the 10 experimental conditions, the observer completed a minimum of 60 “baseline” trials in which the base rates were equal and the payoffs were equal. If the observer reached an accuracy-based performance criterion (no more than 2% below optimal), then two decision bound models were fit to the 60 trials of data (see Maddox & Bohil, 1998, for details). The *optimal decision criterion model* assumed that the observer used the optimal decision criterion (i.e.,  $\beta = 1$ ) in the presence of perceptual and criterial noise, whereas the *free decision criterion model* estimated the observer’s decision criterion from the data. Because the optimal decision criterion model is a special case of the free decision criterion model, likelihood ratio tests were used to determine whether the extra flexibility of the free decision criterion model provided a significant improvement in fit. If the free decision criterion model did *not* provide a significant improvement in fit over the optimal decision criterion model, then the observer was allowed to begin the experimental condition. If the free decision criterion model did provide a significant improvement in fit, then the observer completed 10 additional trials, and the same accuracy-based and model-based criteria were applied to the most recent 60 trials (i.e., Trials 11–70). This procedure continued until the observer reached the appropriate criterion. Including these baseline trials and this fairly conservative accuracy-based and model-based performance criterion ensured that each observer had accurate knowledge of the category structures before exposure to the base-rate or payoff manipulation, and minimized the possibility of within-observer carryover effects from one experimental condition to the next. As an additional safeguard, different category labels were used (e.g., “burlosis” and “namitis” in one condition and “coralgia” and “terragitis” in another condition) across the 10 experimental conditions.

Each experimental condition consisted of three 60-trial training blocks, followed by a 60-trial test block. During training, corrective feedback was provided on each trial (see details below). During test, no feedback was provided. The same 60 stimuli were presented (in random order) once in each training and test block. Table 1 displays the base rates, payoffs, point totals, accuracy rates,  $\beta$ , and  $\ln(\beta)$  values for the optimal classifier separately for each condition.

**Procedure.** 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) in each condition. Observers were told that they were participating in several hypothetical medical diagnosis tasks and that the length 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 as simply A and B. Observers were informed 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 condition. They were informed that these point totals would be converted into monetary values that they would receive at the end of the experiment. Observers were instructed to maximize points 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 Cat-

**Table 1**  
**Category Base Rates, Payoffs, Optimal Points (per 60 Trial Blocks), Accuracy,  $\beta_o$ ,**  
**and  $\ln(\beta_o)$  Values for the 10 Base-Rate/Payoff Conditions**

|           | Base Rates |        | Payoffs  |          | Points | Accuracy | $\beta_o$ | $\ln(\beta_o)$ |
|-----------|------------|--------|----------|----------|--------|----------|-----------|----------------|
|           | $P(A)$     | $P(B)$ | $V_{aA}$ | $V_{bB}$ |        |          |           |                |
| Baseline  | .50        | .50    | 6        | 6        | 308    | 85.5     | 1         | 0              |
| 2:1B      | .67        | .33    | 6        | 6        | 314    | 87.1     | 2         | 0.693          |
| 2:1P      | .50        | .50    | 8        | 4        | 314    | 84.7     | 2         | 0.693          |
| 3:1B      | .75        | .25    | 6        | 6        | 319    | 88.7     | 3         | 1.099          |
| 3:1P      | .50        | .50    | 9        | 3        | 319    | 82.9     | 3         | 1.099          |
| 2:1B/2:1P | .67        | .33    | 8        | 4        | 360    | 86.1     | 4         | 1.386          |
| 3:1B/3:1P | .75        | .25    | 9        | 3        | 421    | 86.7     | 9         | 2.197          |
| 2:1B/3:1P | .67        | .33    | 9        | 3        | 386    | 84.7     | 6         | 1.792          |
| 3:1B/2:1P | .75        | .25    | 8        | 4        | 386    | 87.8     | 6         | 1.792          |
| 1:2B/3:1P | .33        | .67    | 9        | 3        | 259    | 84.1     | 1.5       | 0.405          |
| 3:1B/1:2P | .75        | .25    | 4        | 8        | 259    | 87.6     | 1.5       | 0.405          |

egory B by pressing the appropriate button. During the training phase, the observer's response was followed by 750 msec of feedback. Five lines of feedback were presented. The top (first) line indicated the correct disease label on that trial. The second line indicated the points earned on that trial. The third line indicated the value of a correct response (i.e., the potential points) for that trial (regardless of the observer's response). The fourth line indicated the total points earned to that point in the experiment, and the fifth (last) line indicated the maximum total points (i.e., potential total points) that could have been earned to that point in the experiment (regardless of the observer's performance). The feedback was followed by a 125 msec intertrial interval (ITI) during which the screen was blank. In the test phase, the observer's response was followed by a 875-msec ITI during which the screen was blank. Observers were given a break after each block of trials. At each break, the observer's accumulated point total was displayed.

## RESULTS AND THEORETICAL ANALYSIS

Before we turn to the model-based analyses, a few words are in order regarding some basic trends in the signal detection theory decision criterion estimates (Green & Swets, 1966), point totals, and accuracy rates.

### Performance Trends

Because points, accuracy, and decision criterion estimates from signal detection theory are each measured on a different scale, we transformed each of these measures onto a common scale that allowed us to compare the observer's performance with that of the optimal classifier. Specifically, we computed the following three measures:

$$\begin{aligned} & \text{deviation from optimal decision criterion} \\ &= k - k_o \\ &= \frac{\ln(\beta)}{2.2} - \frac{\ln(\beta_o)}{2.2} \end{aligned} \quad (7)$$

$$\begin{aligned} & \text{deviation from optimal points} \\ &= \frac{\text{observed points} - \text{optimal points}}{\text{optimal points} - 0} \end{aligned} \quad (8)$$

$$\begin{aligned} & \text{deviation from optimal accuracy} \\ &= \frac{\text{observed accuracy} - \text{optimal accuracy}}{\text{optimal accuracy} - 0} \end{aligned} \quad (9)$$

Table 2 displays performance averaged across observers for the decision criterion, point, and accuracy deviation measures by base-rate/payoff condition and training/test block.<sup>5</sup> The flat-maxima hypothesis predicts that performance should be closer to optimal in the 2:1 than in the 3:1 condition. The COBRA hypothesis predicts that performance in the base-rate conditions should be closer to optimal than in the payoff conditions. As an initial test of these two hypotheses we conducted a 2 base-rate/payoff

**Table 2**  
**Decision Criterion, Point, and Accuracy Deviation Performance**  
**Measures (Averaged Across Observers) by Condition and Block**

| Condition | Measure   | Training Block |        |        | Test   |
|-----------|-----------|----------------|--------|--------|--------|
|           |           | 1              | 2      | 3      |        |
| 2:1B      | $k - k_o$ | -0.241         | -0.322 | -0.222 | -0.501 |
|           | Pt. Dev.  | -0.035         | -0.048 | -0.036 | -0.058 |
|           | Acc. Dev. | -0.035         | -0.048 | -0.036 | -0.058 |
| 2:1P      | $k - k_o$ | -0.176         | -0.242 | -0.309 | -0.303 |
|           | Pt. Dev.  | -0.027         | -0.066 | -0.062 | -0.027 |
|           | Acc. Dev. | -0.013         | -0.045 | -0.037 | 0.001  |
| 3:1B      | $k - k_o$ | -0.369         | -0.294 | -0.197 | -0.633 |
|           | Pt. Dev.  | -0.069         | -0.059 | -0.044 | -0.065 |
|           | Acc. Dev. | -0.069         | -0.059 | -0.044 | -0.065 |
| 3:1P      | $k - k_o$ | -0.337         | -0.415 | -0.443 | -0.568 |
|           | Pt. Dev.  | -0.028         | -0.086 | -0.057 | -0.086 |
|           | Acc. Dev. | 0.016          | -0.033 | 0.000  | -0.014 |
| 2:1B/2:1P | $k - k_o$ | -0.393         | -0.355 | -0.298 | -0.380 |
|           | Pt. Dev.  | -0.036         | -0.045 | -0.029 | -0.031 |
|           | Acc. Dev. | -0.014         | -0.026 | -0.014 | -0.004 |
| 3:1B/3:1P | $k - k_o$ | -0.695         | -0.443 | -0.478 | -0.797 |
|           | Pt. Dev.  | -0.060         | -0.028 | -0.027 | -0.096 |
|           | Acc. Dev. | -0.012         | 0.004  | 0.003  | -0.039 |
| 2:1B/3:1P | $k - k_o$ | -0.477         | -0.372 | -0.445 | -0.505 |
|           | Pt. Dev.  | -0.038         | -0.037 | -0.035 | -0.039 |
|           | Acc. Dev. | 0.002          | -0.015 | 0.005  | 0.016  |
| 3:1B/2:1P | $k - k_o$ | -0.541         | -0.412 | -0.386 | -0.642 |
|           | Pt. Dev.  | -0.043         | -0.029 | -0.022 | -0.082 |
|           | Acc. Dev. | -0.019         | -0.010 | -0.004 | -0.050 |
| 1:2B/3:1P | $k - k_o$ | -0.084         | -0.185 | -0.185 | -0.417 |
|           | Pt. Dev.  | -0.040         | -0.059 | -0.057 | -0.073 |
|           | Acc. Dev. | -0.032         | -0.032 | -0.028 | -0.015 |
| 3:1B/1:2P | $k - k_o$ | -0.211         | -0.048 | 0.061  | -0.160 |
|           | Pt. Dev.  | -0.025         | -0.039 | -0.026 | -0.031 |
|           | Acc. Dev. | -0.041         | -0.038 | -0.019 | -0.043 |

ratio (2:1 vs. 3:1) × 2 base-rate versus payoff condition × four blocks ANOVA on the decision criterion, point, and accuracy deviation scores from the four separate base-rate/payoff conditions.<sup>6</sup> The only interaction that reached significance consistently was the base-rate versus payoff condition × block interaction [ $F(3,45) = 9.18, 3.55, \text{ and } 3.98, MS_e = .028, .002, \text{ and } .002, p < .05$ , for the decision criterion, point, and accuracy measures, respectively]. The general pattern was for performance to become more optimal across the three training blocks for the base-rate condition and to become less optimal across the three training blocks for the payoff conditions. Importantly though, by the final block of training, the deviation from optimal decision criterion and points was smaller for the base-rate than for the payoff conditions, in support of the COBRA hypothesis. The deviation was smaller in the payoff condition for the accuracy measure, which is predicted from COBRA since the observers in the payoff condition place importance on accuracy maximization and do not sacrifice accuracy, as would be predicted from the optimal classifier. For both base-rate and payoff conditions, performance worsened during the test block. The main effect of base-rate/payoff ratio was significant for the decision criterion measure [ $F(1,15) = 12.35, MS_e = .071, p < .01$ ], was significant at the  $p = .07$  level for the point measure [ $F(1,15) = 3.77, MS_e = .005, p = .07$ ], and was nonsignificant for the accuracy measure. For both the decision criterion and point measures the deviation from optimal was smaller for the 2:1 than for the 3:1 conditions, in support of the flat-maxima hypothesis. The main effect of base-rate versus payoff condition was significant only for the accuracy measure [ $F(1,15) = 13.28, MS_e = .006, p < .01$ ], suggesting that performance was closer to optimal in the payoff than in the base-rate conditions, in line with COBRA.

**Simultaneous base-rate/payoff conditions: A comparison of the flat-maxima hypothesis with the independence assumption of the optimal classifier.** In the simultaneous base-rate/payoff conditions, the flat-maxima hypothesis predicts that decision criterion placement should be closer to optimal than would be predicted from the independence assumption in the corresponding base-rate/payoff conditions and farther from optimal in the

conflicting base-rate/payoff conditions. To provide an initial test of this prediction we compared the observed decision criterion estimates in the corresponding and conflicting simultaneous base-rate/payoff conditions with those generated from an independent combination of the separate base-rate/payoff condition decision criterion estimates. Specifically, we subtracted the independence predictions from the observed values. We then calculated the number of observers in each condition whose data supported the flat-maxima hypothesis prediction and conducted sign tests (assuming a null hypothesis probability of .50) on these data. In addition, we computed the average difference for each condition and conducted  $t$  tests (assuming a null hypothesis average of 0) on these data. These data are presented in Table 3. Several comments are in order. First, there was strong support for the flat-maxima hypothesis prediction in the corresponding base-rate/payoff conditions, with significant sign and  $t$ -test results for all four conditions collapsed across blocks, and in all but the first block of trials. Second, there was less support for the flat-maxima hypothesis prediction in the conflicting base-rate/payoff conditions, although the flat-maxima hypothesis was supported during the second training and during the test block for the 1:2B/3:1P condition. Finally, support for the flat-maxima hypothesis predictions was lacking during the first training block, but increased as the observers gained experience with the task and was quite large during the test block. These results provide good initial support for the flat-maxima hypothesis predictions in the simultaneous base-rate/payoff conditions and provide a foundation for the stronger model-based analyses that we turn to next.

**Model-Based Analyses**

Each of the Figure 5 models was applied simultaneously to the data from all 10 experimental conditions separately for each block and observer. Each block consisted of 60 experimental trials, and the observer was required to respond “A” or “B” for each stimulus. Thus each model was fit to a total of 600 independent response probabilities from each block (60 trials × 10 conditions). Maximum likelihood procedures (Ashby, 1992b; Wickens, 1982) were used to estimate the model

**Table 3**  
**Number of Observers/Total in Each Cell for Which the Observed Simultaneous Base-Rate/Payoff Condition Decision Criterion Minus the Independence Assumption Predicted Decision Criterion Was Positive (for the Corresponding) or Negative (for the Conflicting) by Condition and Block**

| Training Block | Condition |          |           |          |           |          |           |          |           |          |           |          | Total<br><i>N</i> |
|----------------|-----------|----------|-----------|----------|-----------|----------|-----------|----------|-----------|----------|-----------|----------|-------------------|
|                | 2:1B/2:1P |          | 3:1B/3:1P |          | 2:1B/3:1P |          | 3:1B/2:1P |          | 1:2B/3:1P |          | 3:1B/1:2P |          |                   |
|                | <i>N</i>  | <i>M</i> | <i>N</i>  | <i>M</i> | <i>N</i>  | <i>M</i> | <i>N</i>  | <i>M</i> | <i>N</i>  | <i>M</i> | <i>N</i>  | <i>M</i> |                   |
| 1              | 8/16      | 0.051    | 9/16      | 0.024    | 10/16     | 0.218    | 7/16      | 0.010    | 7/16      | 0.027    | 9/16      | -0.039   | 50/96             |
| 2              | 13/16**   | 0.452    | 9/16      | 0.574#   | 12/16*    | 0.787**  | 9/16      | 0.267    | 11/16#    | -0.198   | 7/16      | 0.008    | 61/96**           |
| 3              | 11/16#    | 0.501*   | 10/16     | 0.349    | 11/16#    | 0.474#   | 10/16     | 0.259    | 6/16      | 0.079    | 8/16      | -0.110   | 56/96*            |
| Test           | 12/16*    | 0.914    | 13/16**   | 0.869**  | 14/16**   | 1.214**  | 12/16*    | 0.634**  | 12/16*    | -0.755*  | 4/16      | 0.365    | 67/96**           |
| Total          | 44/64**   | 0.479**  | 41/64*    | 0.454*   | 47/64**   | 0.673**  | 38/64#    | 0.292*   | 36/64     | -0.212   | 28/64     | 0.056    | 234/384**         |

Note—The “mean” for the Total column is missing because it would require averaging data from conditions for which a positive difference is predicted with conditions for which a negative difference is predicted. Sign test assuming a null hypothesis probability = .50 or  $t$  test assuming a null hypothesis mean of 0 results: #.05 <  $p$  < .10; \* $p$  < .05; \*\* $p$  < .01.

parameters, and likelihood ratio ( $G^2$ ) tests were used to determine whether a more general model provided a significant improvement in fit (with  $\alpha = .05$ ) over a more restricted, nested model.<sup>7</sup> Maximum likelihood fit values for models with the same number of parameters were compared directly. Because the models were applied separately to each block of trials, we could identify the model with the fewest free parameters that could not be improved upon (statistically) by a more general model, referred to as the *most parsimonious model*, for each observer in each block. However, our interest was in identifying the most parsimonious model separately for each observer. Since decision criterion shifts across trials are likely large early in learning, but are much smaller later in learning and during the test block, we decided to determine this model from the final training block and test block. Thus, the  $G^2$  tests were based on the fit of each model summed over the last two blocks. (Using the cumulative fit across all four blocks did not change the overall pattern of results.) Even so, we do examine the parameter values separately for each block to better characterize decision criterion changes with experience.

**Identifying the most parsimonious model.** On the basis of the likelihood ratio ( $G^2$ ) tests of the maximum

likelihood fit values for the four “nonindependence” models (Figure 5), the hybrid( $w_p$ ;  $w_{\text{Corresponding}}$ ;  $w_{\text{Conflicting}}$ ) model provided the most parsimonious account of the data from 9 of the 16 observers, the hybrid( $w$ ) model from 3 of 16 observers, the flat-maxima model from 2 of 16 observers, and the COBRA model from 2 of 16 observers. The fact that the hybrid models were superior for 12 of 16 observers suggests that both hypotheses—the flat-maxima and COBRA hypotheses—are necessary to provide an adequate account of human decision criterion learning when base rates and payoffs are manipulated separately and simultaneously across a range of base-rate/payoff ratios.<sup>8</sup> In addition, it is important to note that the absolute fit of the hybrid( $w_p$ ;  $w_{\text{Corresponding}}$ ;  $w_{\text{Conflicting}}$ ) model was quite good, accounting for 88.3%–95.1% of the responses in the data from the 16 observers.

**Hybrid model steepness and accuracy weight parameter changes over blocks.** To determine how the observer's estimate of the reward-maximizing decision criterion changed across blocks and to determine the magnitude of the weight placed on accuracy and how it changed across blocks, we examined the steepness and accuracy weight,  $w$ , parameters from the hybrid( $w_p$ ;

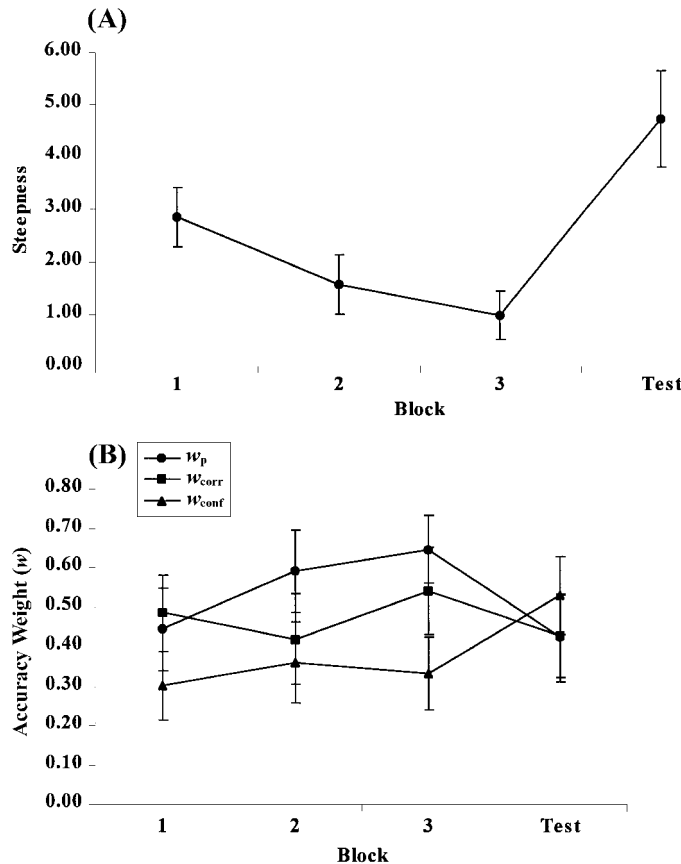


Figure 6. (A) Steepness values. (B) Accuracy weight,  $w$ , values from the hybrid( $w_p$ ;  $w_{\text{Corresponding}}$ ;  $w_{\text{Conflicting}}$ ) model for the three training blocks and the test block averaged across observers. Standard error bars are included.

$w_{\text{Corresponding}}$ ;  $w_{\text{Conflicting}}$ ) model. These values are displayed for the three training blocks and the test block averaged across observers in Figure 6. Several results stand out. First, a one-way ANOVA on the steepness values suggested a significant effect of block [ $F(3,45) = 10.06$ ,  $MS_e = 4.336$ ,  $p < .001$ ] that was characterized by a significant decline from Training Block 1 to Training Block 2 ( $p < .05$ ), a nonsignificant decline from Training Block 2 to Training Block 3, and a significant *increase* from the final training block to the test block ( $p < .05$ ). A focus on the training data suggests that the observer's reward-maximizing decision criterion approached the optimal value across the three training blocks. The increase from the final training to the test block was not completely unexpected; we have found this pattern in other studies that included only three 60-trial training blocks (e.g., Maddox & Bohil, in press). When more training is available (e.g., six 60-trial training blocks, as in Maddox & Dodd, 2001) we do not find this increase in the test block. Perhaps additional training leads to a more stable representation of the decision criterion that is less affected by the removal of trial-by-trial feedback during the test trials. Second, a one-way ANOVA on the accuracy weights,  $w$ , suggests no effect of training on the weight placed on accuracy. Instead, the weight placed on accuracy remained fairly constant, generally falling in the .30–.60 range—similar to the values found in earlier studies conducted in our lab. Despite the lack of significance, it is worth mentioning that the weight placed on accuracy (during the training blocks) was consistently largest in the separate payoff conditions, lowest in the conflicting simultaneous base-rate/payoff conditions, and intermediate in the corresponding simultaneous base-rate/payoff conditions. Speculatively, it appears that observers are able to overcome the need to maximize accuracy when they “realize” (perhaps implicitly) that the payoffs are biasing them toward a different response than the base rates.

**Comparison of the hybrid model with the hybrid/independence model.** With no additional post hoc mechanisms, the hybrid model captures performance in all the conditions, especially the corresponding and conflicting simultaneous base-rate/payoff conditions, for which the flat-maxima hypothesis makes different predictions from that of the independence assumption of the optimal classifier. Even so, the most rigorous comparison of the flat-maxima prediction with the independence assumption requires a direct model comparison of the hybrid model with a model that assumes an independent combination of base-rate and payoff information in the simultaneous base-rate/payoff conditions. To achieve this goal, we directly compared the fit of the hybrid( $w$ ) model with that of the hybrid( $w$ )/independence model (see Figure 5), since both have the same number of free parameters. The results were clear. For 9, 11, 12, and 14 of the 16 observers in Training Blocks 1–3 and the test block, respectively, the hybrid( $w$ ) model provided a better account of the data than the hybrid( $w$ )/independence

model. Based on a sign test with a null hypothesis probability of .50, the results for Blocks 1 and 2 were nonsignificant, and the results for Block 3 and the test block were significant ( $p < .05$ ).

One advantage of comparing the hybrid( $w$ ) and hybrid( $w$ )/independence models is that the models are identical in all respects except whether they apply the flat-maxima hypothesis or independence assumption to the simultaneous base-rate/payoff conditions. In the separate base-rate/payoff conditions, both models assume that the flat-maxima and COBRA hypotheses are operating. Even so, one might argue that this approach biases the results in favor of the flat-maxima hypothesis since the flat-maxima hypothesis is used to generate the decision criteria for the separate base-rate/payoff conditions. Another approach would be to develop a model that instantiates the independence assumption of the optimal classifier by freely estimating separate decision criterion values in the 2:1B, 2:1P, 3:1B, and 3:1P conditions and then combining these independently. Although this approach lacks an underlying theory for determining the four separate base-rate/payoff condition decision criteria, it does give the independence assumption the best chance of accounting for the data. We applied this five-parameter model (one noise and four decision criteria) to the data and compared it with the five-parameter hybrid( $w_p$ ;  $w_{\text{Corresponding}}$ ;  $w_{\text{Conflicting}}$ ) model. The hybrid( $w_p$ ;  $w_{\text{Corresponding}}$ ;  $w_{\text{Conflicting}}$ ) model provided a better account of the data for 3, 10, 10, and 7 of the 16 observers in Training Blocks 1–3 and the test block, respectively. Interestingly, whereas the independence model was superior during the first block of training, as the observer's gained experience with the task the hybrid model framework emerged as the superior model. The current and previous model-based comparisons suggest that the hybrid model that embodies the flat-maxima and COBRA hypotheses provides an excellent account of decision criterion learning across (1) different base-rate/payoff ratios, (2) separate base-rate/payoff manipulations, (3) corresponding simultaneous base-rate/payoff manipulations, and (4) conflicting simultaneous base-rate/payoff manipulations.

## GENERAL DISCUSSION

This article reports the results of an experiment that examined the effects of 2:1 and 3:1 base-rate/payoff ratios, four separate base-rate/payoff manipulations, four corresponding simultaneous base-rate/payoff manipulations, and two conflicting simultaneous base-rate/payoff manipulations on decision criterion learning in perceptual categorization. Each observer completed several blocks of trials in each of these 10 base-rate/payoff conditions. These conditions allowed a critical test of several predictions derived from the flat-maxima hypothesis and allowed a direct comparison with predictions derived from the independence assumption of the optimal classifier. These predictions were compared within the frame-

work of Maddox and Dodd's (2001; see also Maddox, in press) hybrid model, which instantiates both the flat-maxima hypothesis and the COBRA hypothesis.

The flat-maxima hypothesis suggests that the decision criterion used by the observer to maximize expected reward is determined by the steepness of the objective reward function (Figure 2), and makes the following three predictions. First, the decision criterion in the 2:1 condition should be closer to optimal than that from the 3:1 condition. Second, the decision criterion used in the corresponding simultaneous base-rate/payoff conditions should be closer to optimal than the decision criterion predicted from the independence assumption of the optimal classifier. Finally, the decision criterion used in the conflicting simultaneous base-rate/payoff conditions should be further from optimal than the decision criterion predicted from the independence assumption of the optimal classifier. COBRA predicts that observers place some importance (or weight) on reward *and* accuracy maximization. When both goals can be achieved simultaneously, as in unequal base-rate conditions (Figure 3A), there is effectively no competition, and the decision criterion approaches that used by the optimal classifier. When both goals cannot be achieved simultaneously, as in unequal payoff conditions, the competition leads to the use of a decision criterion intermediate between the optimal decision criterion and the decision criterion that maximizes expected accuracy. Thus, COBRA predicts more nearly optimal decision criterion placement in base-rate as compared with payoff conditions.

Maddox and Dodd's (2001) hybrid model provided good accounts of the data, providing support for the hypothesis that the observer's reward-maximizing decision criterion is a function of the steepness of the objective reward function in all base-rate/payoff conditions, even those in which base rates and payoffs are manipulated simultaneously and those in which the base rates and payoffs act in a corresponding or in a conflicting manner. This hypothesis argues against the validity of the independence assumption of the optimal classifier since the independence assumption consistently underpredicts decision criterion placement in corresponding simultaneous base-rate/payoff conditions and overpredicts decision criterion placement in conflicting simultaneous base-rate/payoff conditions. Additional support for this claim was obtained through a direct model-based comparison of these two hypotheses. In particular, the hybrid models that assumed that the decision criterion used in the simultaneous base-rate/payoff condition was determined from the steepness of the objective reward function (i.e., the flat-maxima hypothesis) were compared with analogous hybrid models that assumed that the flat-maxima and COBRA hypotheses determined the decision criterion values in the separate base-rate and payoff conditions, but that these values were combined independently in determining the decision criterion in the simultaneous base-rate/payoff conditions. In other words, a flat-maxima explanation for decision criterion place-

ment in the simultaneous base-rate/payoff conditions was compared directly with an independence explanation of decision criterion placement in these same conditions while holding all else equal. For 46 of the 64 comparisons (16 observers  $\times$  four blocks) the model that instantiated the flat-maxima explanation provided a better account of the data than the model that instantiated the independence explanation. A sign test with binomial probability = .5 indicates that this difference is highly significant ( $p < .001$ ). The hybrid model was also compared with a model that instantiated the independence assumption, but for which the decision criteria in the four separate base-rate/payoff conditions were freely estimated. Although the decision criterion values in the four separate base-rate/payoff conditions are not theoretically motivated, this model does give the independence assumption the most freedom to account for the data. When compared with the hybrid model, this independence model fared well. In 34 of 64 cases the independence model provided a superior account of the data, although this was due mainly to the superiority of the independence model during the first block of trials.

In conclusion, the present study extends our understanding of decision criterion learning across a range of base-rate/payoff ratios for separate and two qualitatively different types of simultaneous base-rate/payoff manipulations. Strong support for the flat-maxima and COBRA hypotheses was obtained. In addition, a flat-maxima interpretation of decision criterion learning in corresponding and conflicting simultaneous base-rate/payoff conditions provided a consistently better account of the data than an interpretation based on the independence assumption.

## 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., & 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., & LEE, W. W. (1991). Predicting similarity and categorization from identification. *Journal of Experimental Psychology: General*, **120**, 150-172.
- ASHBY, F. G., & MADDUX, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, **37**, 372-400.
- ASHBY, F. G., & MADDUX, W. T. (1994). A response time theory of separability and integrality in speeded classification. *Journal of Mathematical Psychology*, **38**, 423-466.
- ASHBY, F. G., & MADDUX, 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., MADDUX, 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., & 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.
- DUSOIR, A. (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., & GREENSPAN, 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.
- ESTES, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin*, **53**, 134-140.
- 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.
- KAHNEMAN, D., & TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47**, 263-291.
- KALISH, M. L., & KRUSCHKE, J. K. (1997). Decision boundaries in one dimensional categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 1362-1377.
- KRUSCHKE, J. K., & JOHANSEN, M. K. (1999). A model of probabilistic category learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **25**, 1083-1119.
- 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.
- LEE, W., & JANKE, M. (1964). Categorizing 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.
- LEE, W., & ZENTALL, T. R. (1966). Factorial effects in the categorization of externally distributed stimulus samples. *Perception & Psychophysics*, **1**, 120-124.
- MADDOX, W. T. (1999). On the danger of averaging across observers when comparing decision bound and generalized context models of categorization. *Perception & Psychophysics*, **61**, 354-374.
- MADDOX, W. T. (2002). Toward a unified theory of decision criterion learning in perceptual categorization. *Journal of the Experimental Analysis of Behavior*, **78**, 567-595.
- 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. (1998). 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. (2000). Costs and benefits in perceptual categorization. *Memory & Cognition*, **28**, 597-615.
- MADDOX, W. T., & BOHIL, C. J. (in press). A theoretical framework for understanding the effects of simultaneous base-rate and payoff manipulations on decision criterion learning in perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*.
- MADDOX, W. T., & DODD, J. L. (2001). On the relation between base-rate and cost-benefit learning in simulated medical diagnosis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 1367-1384.
- 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.
- 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.
- 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.
- 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.
- TVERSKY, A., & KAHNEMAN, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, **185**, 1124-1131.
- TVERSKY, A., & KAHNEMAN, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (pp. 49-72). Hillsdale, NJ: Erlbaum.
- TVERSKY, A., & KAHNEMAN, D. (1992). Prospect theory: An analysis of decision under risk. *Econometrica*, **47**, 276-287.
- 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.
- YATES, J. F. (1990). *Judgment and decision-making*. Englewood Cliffs, NJ: Prentice-Hall.

## NOTES

1. Other approaches have been offered that do not postulate the use of a decision criterion (e.g., Kalish & Kruschke, 1997). Most of this work focuses on cases in which base rates and payoffs are equal, although some work has examined unequal base-rate conditions (e.g., Kruschke & Johansen, 1999).

2. Before continuing, we need to discuss one important issue regarding the optimal decision criterion. The optimal decision criterion is constructed from "objective" or "true" payoff information. Much work suggests that people do not use the objective values, but rather base their decisions on subjective values that are directly related to the objective values (e.g., Kahneman & Tversky, 1979; Stevenson, Busemeyer, & Naylor, 1991; Tversky & Kahneman, 1974, 1980, 1992; Yates, 1990). Within the framework of decision theory, each of our  $V_{ij}$  terms should be converted into a subjective utility denoted  $u(V_{ij})$ , where  $u$  describes the functional relationship between the subjective and objective values. In the case of points converted to money, it is reasonable to assume that increasing value is approximately linearly associated with increasing utility.

3. In the experiment outlined below  $d' = 2.2$  in all conditions, so standardizing by  $d'$  is technically unnecessary. Even so, this standardized  $k - k_0$  measure facilitates comparison across objective reward functions in other studies and forms the basis of Maddox and Dodd's (2001) hybrid model. Thus, it will be used throughout this article.

4. Another possibility is that an intermediate decision criterion is not used, and instead the two decision criteria might 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). More work is needed to determine which of these approaches is more valid. In both approaches the  $w$  parameter would have a similar interpretation. In this alternative approach the  $w$  parameter would denote the probability that the decision criterion used on a particular trial was the accuracy-maximizing criterion. In the present

approach  $w$  denotes the weight placed on the accuracy-maximizing decision criterion in determining the criterion used on each trial. This approach is taken because it is easy to instantiate and has met with reasonable success (Maddox & Dodd, 2001).

5. It is worth mentioning that the decision criterion, point, and accuracy deviation measures are partially, albeit not perfectly, correlated. For example, the decision criterion measure is derived from the hit and false alarm rates and so is partially correlated with accuracy. In addition, accurate responding leads to more points (at least when base rates are manipulated) and so accuracy and point measures are partially correlated. We include all three measures for completeness, and do not expect perfect convergence in the conclusions drawn from each.

6. ANOVAs were also conducted on the  $d'$  estimates from signal detection theory. All effects were nonsignificant. These data are not discussed further.

7. These results should be interpreted with some caution since the maximum likelihood fit values assume independence in the data that is probably violated.

8. We also performed  $G^2$  tests on the fits of the Figure 5 models applied to the data from the final training block alone, and  $G^2$  tests on the

fits of the Figure 5 models applied to the data from the test block alone. The distribution of most parsimonious models and the superiority of the hybrid( $w_p; w_{\text{Corresponding}}; w_{\text{Conflicting}}$ ) model was similar to that obtained when the models were fit simultaneously to the data from the final training and test block. Specifically, for the final training block the hybrid( $w_p; w_{\text{Corresponding}}; w_{\text{Conflicting}}$ ) model provided the most parsimonious account of the data from 7 of the 16 observers, the hybrid( $w$ ) model from 3 of 16 observers, the flat-maxima model from 2 of 16 observers, and the COBRA model from 3 of 16 observers. For the test block the hybrid( $w_p; w_{\text{Corresponding}}; w_{\text{Conflicting}}$ ) model provided the most parsimonious account of the data from 7 of the 16 observers, the hybrid( $w$ ) model from 4 of 16 observers, the flat-maxima model from 2 of 16 observers, the COBRA model from 1 of 16 observers, and the optimal model from 2 of 16 observers.

(Manuscript received December 13, 2001;  
revision accepted for publication September 30, 2002.)