Journal of Experimental Psychology:
Learning, Memory, and Cognition
2003, Vol. 29, No. 2, 307–320

0278-7393/03/$12.00 DOI: 10.1037/0278-7393.29.2.307

A Theoretical Framework for Understanding the Effects of Simultaneous Base-Rate and Payoff Manipulations on Decision Criterion Learning in Perceptual Categorization

W. Todd Maddox and Corey J. Bohil
University of Texas at Austin

Observers completed perceptual categorization tasks in which base rates and payoffs were manipulated separately or simultaneously across a range of category discriminabilities. Decision criterion estimates from the simultaneous base-rate/payoff conditions were closer to optimal than those predicted from the independence assumption, in line with predictions from the flat-maxima hypothesis. A hybrid model that instantiated the flat-maxima and competition between reward and accuracy maximization hypotheses was applied to the data as well as used in a reanalysis of C. J. Bohil and W. J. Maddox's (2001) study. 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.

Each day people are faced with situations in which they must choose a course of action based on uncertain information. For example, the student must decide whether to bring or not bring an umbrella to class on the basis of a quick look at the darkness of the cloud cover. Similarly, the expedition doctor camped out near the summit of Denali must decide whether a climber's complaint of difficulty breathing is a sign of pulmonary edema or exhaustion, with a pulmonary edema diagnosis leading to immediate retreat. These are both categorization problems because in each case, there are many possible information states (e.g., darkness of the cloud cover or degrees of difficulty breathing) but only two potential decisions (bring or do not bring an umbrella or diagnose or do not diagnose edema).

Optimal categorization performance that maximizes long-run reward often requires the placement of a decision criterion along some relevant dimension. For example, a criterion must be set on the climber's breathing difficulty, with breathing levels below the criterion leading to an exhaustion diagnosis and levels above the criterion leading to an edema diagnosis. 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, at low altitudes, edema has a low base-rate probability, but as the climber gains altitude, the probability that edema will develop increases, and so the expedition doctor might lower the breathing criterion associated with an edema diagnosis. 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 edema diagnosis will be greater than the benefit of a correct exhaustion diagnosis, especially at high altitude. Similarly, there is often a cost associated with an incorrect decision; a more severe cost would be incurred for an incorrect exhaustion diagnosis than for an incorrect edema diagnosis. Under these conditions, the expedition doctor would be wise to lower the breathing criterion associated with an edema diagnosis to increase the number of correct edema diagnoses, even though this would adversely affect the accuracy of exhaustion diagnoses.

Base rates and payoffs vary widely in real-world categories, and in many (likely most) cases, both differ simultaneously within the same categorization problem. Despite this fact, 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 is twofold. First, we examine human decision criterion learning when base rates and payoffs are manipulated separately and compare performance with cases in which base rates and payoffs are manipulated simultaneously. Second, we generate and test predictions from Maddox and Dodd's (2001) hybrid model of decision criterion placement for the separate and simultaneous base-rate/payoff conditions under investigation. Throughout this article, we use the behavior of the optimal classifier as a benchmark against which to compare human performance. Although human decision criterion learning rarely matches that of the optimal classifier, this benchmark provides a foundation for our investigation.

The Optimal Classifier and Decision Bound Theory

Optimal Classifier

The optimal classifier is a hypothetical device that maximizes long-run expected reward. Consider the situation facing a medical

---

W. Todd Maddox and Corey J. Bohil, Department of Psychology, University of Texas at Austin.

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 Tom Wallsten and three anonymous reviewers for helpful comments on an earlier version of this article and Amy Kaderka for help with data collection. We also thank David Ing for generating predictions from J. R. Busemeyer and I. J. Myung's (1992) hill-climbing model.

Correspondence concerning this article should be addressed to W. Todd Maddox, Department of Psychology, University of Texas, 1 University Station A8000, Austin, Texas 78712. E-mail: maddox@psy.utexas.edu

doctor who must classify a patient into one of two disease categories, A or B. Suppose the patient is given medical test X whose outcomes for Diseases A and B are normally distributed as depicted in Figure 1A. The optimal classifier has perfect knowledge of the form and parameters of each category distribution and records perfectly the test result, denoted as *x.* This information is used to construct the *optimal decision function,* which is the likelihood ratio of the two category distributions,

$$l_o(x) = f(x|\mathrm{B})/f(x|\mathrm{A}), \tag{1}$$

where $f(x|i)$ denotes the likelihood of test result x given disease category i. Test results more likely to come from Disease B will yield a likelihood ratio greater than 1, and test results more likely to come from Disease A will yield a likelihood ratio less than 1.

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,*

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

where $P(\mathrm{A})$ and $P(\mathrm{B})$ are the base-rate probabilities for Categories A and B, V_{aA} and V_{bB} denote the payoffs associated with correct diagnoses, and V_{bA} and V_{aB} denote the costs associated with incorrect diagnoses.[1] (In the experiment outlined below, the cost of an incorrect response was set at zero, and so these terms were excluded from Equation 2.) The optimal classifier (e.g., Green & Swets, 1966) uses $l_o(x)$ and β_o to construct the *optimal decision rule:*

$$\text{If } l_o(x) > \beta_o, \text{ then respond “B”; otherwise respond “A.”} \tag{3}$$

Three points are in order. First, if $P(\mathrm{A})V_{\mathrm{aA}} = P(\mathrm{B})V_{\mathrm{bB}}$, 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 $P(\mathrm{A}) = 3P(\mathrm{B})$ or $V_{\mathrm{aA}} = 3V_{\mathrm{bB}}$], then $\beta_o = 3.0$ (see Figure 1). In this case, the optimal classifier will generate a Disease A diagnosis unless the likelihood of Disease B is at least three times greater 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, in which we apply the natural log to both sides, yielding

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

Notice that $\ln\beta_o$ 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 (Stevenson, Busemeyer, & Naylor, 1991).

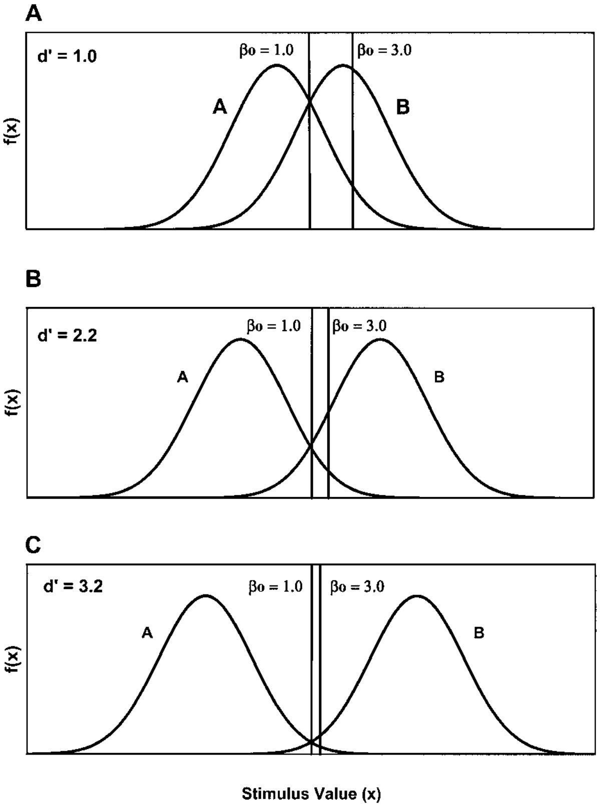

Figure 1. Hypothetical distributions for Categories A and B when category discriminability (d') is equal to 1.0 (A), 2.2 (B), and 3.2 (C). The $\beta_o = 1.0$ decision criterion denotes the criterion that is optimal when the base rates are equal and the payoff matrix is symmetric. This is also referred to as the *equal likelihood decision criterion.* The $\beta_o = 3.0$ decision criterion denotes the criterion that is optimal when there is a 3:1 base-rate ratio or when the payoff matrix is asymmetric with a 3:1 payoff ratio.

Decision Bound Theory

The optimal classifier decision rule (Equation 3) has been rejected as a model of human performance, but in many cases, performance approaches that of the optimal classifier as the observer gains experience with the task. Ashby and colleagues (Ashby, 1992a; Ashby & Lee, 1993; Ashby & Maddox, 1993, 1994; Ashby & Townsend, 1986; Maddox & Ashby, 1993) argued that the observer attempts to respond using a strategy similar to that of the optimal classifier but fails because of various suboptimalities in perceptual and cognitive processing. They proposed a series of decision bound models to test specific hypotheses about the locus of performance suboptimalities. Two suboptimalities inherent in humans and other organisms are perceptual and criterial noise. *Perceptual noise* exists because there is trial-by-trial variability in the perceptual information associated with each stimulus. Assuming a single perceptual dimension is relevant, the observer's percept of stimulus i, on any trial, is x_{pi}, which is equal

---

[1] The optimal decision criterion is constructed from the objective payoff information, whereas people 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 use, denoted as $u(V_{\mathrm{iJ}})$, 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 use.

to $x_i + e_p$, where x_i is the observer's mean percept and e_p is a random variable that represents the effect of perceptual noise. (We assume that e_p is normally distributed and that $\sigma_{pi} = \sigma_p$.) At the cognitive level, there is trial-by-trial variability in the placement of the decision criterion (termed *criterial noise*). Because of criterial noise, the decision criterion used on any trial is β_c, which is equal to $\beta + e_c$, where β is the observer's average decision criterion and e_c is a random variable that represents the effects of criterial noise. (We assume that e_c is normally distributed with a standard deviation σ_c.) Because perceptual and criterial noise exist, the observer cannot attain the level of performance reached by the optimal classifier (i.e., cannot maximize the expected reward). Even so, decision bound theory assumes that the observer attempts to use the same strategy as the optimal classifier but with less success, because of 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. That is,

$$\text{if } l_o(x_{pi}) > \beta_o + e_c, \text{ then respond "B"; otherwise respond "A."} \quad (5)$$

A Theory of Decision Criterion Learning and a Hybrid Model Framework

In several early applications of decision bound theory to data collected from unequal base-rate and unequal payoff conditions, the optimal decision bound model (Equation 5), which assumed that the observer used the optimal decision criterion (β_o), was compared with a suboptimal model that assumed that the observer used a suboptimal decision criterion (β) that was freely estimated from the data (Bohil & Maddox, 2001; Maddox, 1995; Maddox & Bohil, 1998a, 1998b; Maddox & Bohil, 2000). Two findings emerged from this work that were observed in numerous earlier studies (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 $\beta_o = 3$, then observers tended to use a β 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.

Although fruitful, one weakness of the decision bound theoretic approach to decision criterion learning is that no mechanism was postulated or formalized to guide decision criterion placement. Rather, the decision criterion, β, was freely estimated from the data. As a step toward alleviating this problem, Maddox and Dodd (2001) offered a formal theory of decision criterion learning. The idea was to use decision bound theory as the basic modeling framework but to supplement the model by postulating psychologically meaningful mechanisms that would guide decision criterion placement. In other words, instead of simply estimating the decision criterion value directly from the data, Maddox and Dodd outlined and formalized processes that might determine the decision criterion value. Maddox and Dodd used the two findings outlined above as a starting point and offered a theory of decision criterion learning and a model-based instantiation that predicts these two results. The theory proposes two mechanisms 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). Suppose that the observer adjusts the decision criterion on the basis of (at least in part) 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; Dusoir, 1980; Erev, 1998; Erev, Gopher, Itkin, & Greenshpan, 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 et al., 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 the decision criterion value. Figure 2A displays the objective reward function for a case in which category discriminability or d' is equal to 1.0 (where d' is defined as the standardized distance between the category means) and $\beta_o = 1$. That is, Figure 2A plots expected reward as a function of the deviation between a hypothetical observer's decision criterion (β) and the optimal decision criterion (β_o) standardized by category d'. This is denoted as $k - k_o$, which is equal to $\ln(\beta)/d' - \ln(\beta_o)/d'$.

The derivative of the objective reward function at a specific $k - k_o$ value determines the change in the rate of expected reward for that $k - k_o$ value; the larger the change in the rate, the steeper the objective reward function at that point. Derivatives for three $k - k_o$ values are denoted by the three tangent lines labeled 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 point *b* to point *c*). In other words, the steepness declines as the decision criterion approaches the optimal decision criterion. Figure 2B plots the relationship between the steepness of the objective reward function (i.e., the derivative at several $k - k_o$ values) and $k - k_o$. 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 of the objective reward function), 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.

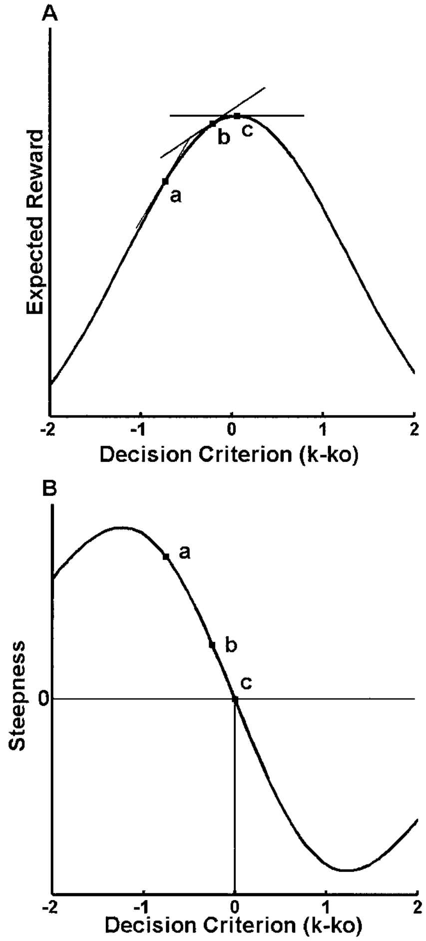

Figure 2. A: Expected reward as a function of the decision criterion (relative to the optimal decision criterion; i.e., $k - k_o$), called the *objective reward function,* for category discriminability, d', equal to 1.0. The three tangent lines at Points *a, b,* and *c* on the objective reward function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

Category discriminability predictions. Maddox and Dodd (2001; see also Bohil & Maddox, 2001) provided a test of the flat-maxima hypothesis as it applies to category discriminability by examining human decision criterion placement for category d' = 1.0, 2.2, and 3.2 (see Figure 1). Figure 3A displays the objective reward functions for these three d' values for a 3:1 base-rate or 3:1 payoff condition, and Figure 3B displays the associated steepness functions. The tangent lines in Figure 3A labeled 1, 2, and 3 correspond to the $k - k_o$ values associated with the same steepness value for $d' = 1.0$, 2.2, and 3.2, respectively. The horizontal line in Figure 3B denotes the associated steepness value, and the vertical lines denote the associated $k - k_o$ values for each category d'. Notice that for this fixed steepness, the deviation between the decision criterion and the optimal value, $k - k_o$, differs systematically across category d' conditions in such a way that the decision criterion (k) is closest to the optimal value (k_o) when category d' = 2.2, is farthest from optimal when $d' = 1.0$, and is intermediate when $d' = 3.2$. Thus, the flat-maxima hypothesis predicts that performance should be closest to optimal when category $d' = 2.2$, farthest from optimal when category $d' = 1.0$, and intermediate when category $d' = 3.2$. Maddox and Dodd's (2001) data supported this prediction.

Separate and simultaneous base-rate/payoff manipulation predictions. Only a few studies have examined decision criterion learning when base rates and payoffs are manipulated simultaneously within the same experimental context (Healy & Kubovy, 1981; Maddox & Bohil, 1998a; see also Stevenson et al., 1991), and only one of these has used more than one level of category discriminability (Bohil & Maddox, 2001). For example, Healy and Kubovy (1981) had observers complete several numerical decision tasks with $d' = 1.0$ in which base rates and payoffs were manipulated separately and simultaneously within observers. Stevenson et al. (1991) reanalyzed these data and compared human decision criterion learning with predictions from the optimal classifier. Recall from Equation 4 that $\ln\beta_o$ is determined completely by the sum of an independent base-rate and payoff term. In analysis of

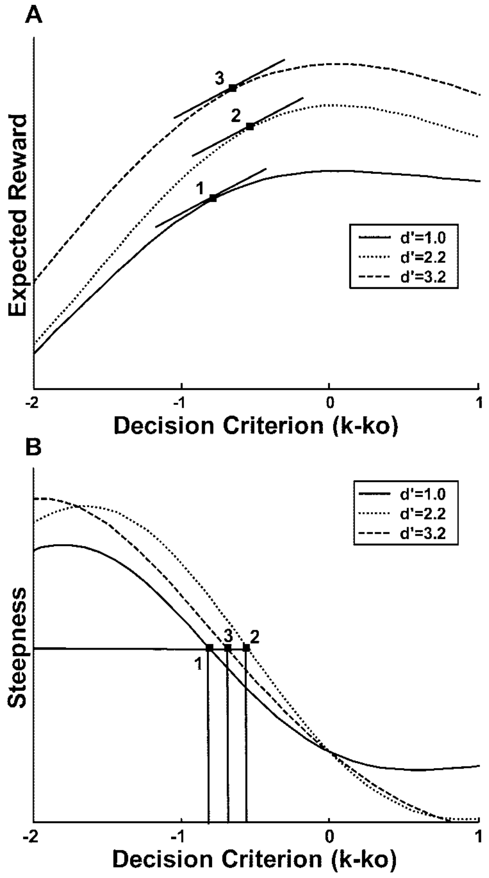

Figure 3. A: Expected reward as a function of the decision criterion (relative to the optimal decision criterion; i.e., $k - k_o$), called the *objective reward function,* for category discriminability, d', equal to 1.0, 2.2, and 3.2. The three tangent lines, one for each d' level, have the same slope and thus reflect the same derivative or steepness. B: Steepness of the objective reward functions from A along with the three points highlighted in A. The three points have identical steepness as denoted by the horizontal line; the vertical lines denote the associated $k - k_o$ values for each category d'. The corresponding decision criterion values are closest to the optimal value for $d' = 2.2$, farthest from optimal for $d' = 1.0$, and intermediate for $d' = 3.2$.

variance (ANOVA) terms, if this *independence assumption* of the optimal classifier holds in human decision criterion learning, then the interaction of base rates and payoffs should be nonsignificant, and a line plot should be characterized by parallel, linear trends. Stevenson et al. did not have the raw data to submit to an ANOVA, but they did generate the relevant plots (Stevenson et al., 1991, Figure 12) and found approximate linearity, in support of the independence hypothesis. Unfortunately, this approach is based on aggregate data, thus ignoring the possibility that individual differences might exist (Maddox, 1999). More important, support for the independence assumption relies on the dangerous proposition of accepting the null hypothesis. Finally, no theoretically motivated alternatives to the independence assumption were offered or tested.

In this section we derive predictions from the flat-maxima hypothesis for situations in which base rates and payoffs are manipulated separately or simultaneously. Figure 4A displays the objective reward functions for the three base-rate/payoff conditions examined in the experiment below with $d' = 1.0$: a 3:1 base-rate condition (hereafter referred to as the 3:1B condition) in which one category is presented three times as often as the other category; a 3:1 payoff condition (hereafter referred to as the 3:1P condition) in which the benefit of a correct response for one category is three times the benefit of a correct response for the other category (the cost of an incorrect response for both categories is zero); and a 3:1 base-rate/3:1 payoff condition (i.e., high base-rate/high payoff; hereafter referred to as the HH condition) in which the benefit of a correct response for the high base-rate category is three times the benefit of a correct response for the low base-rate category. It is important to notice that the objective reward functions for the 3:1B and 3:1P conditions are identical. Figure 4B displays the associated steepness functions. The tangent lines in Figure 4A labeled *BP* and *HH* correspond to the $k - k_o$ values associated with the same steepness value for the 3:1B and 3:1P conditions and the HH condition, respectively. The horizontal line in Figure 4B denotes the associated steepness value, and the vertical lines denote the associated $k - k_o$ values for each condition.

Of more central interest is a comparison of the flat-maxima hypothesis prediction for the HH condition with the associated prediction from the independence assumption of the optimal classifier. Notice that $k - k_o = -0.50$ in the 3:1B and 3:1P condition. Because $d' = 1.0$, $k_o = \ln(3)$, and thus $k = 0.60$. In deriving the independence assumption prediction, we assumed that the flat-maxima hypothesis determined the decision criterion for the 3:1B and 3:1P conditions but that the flat-maxima hypothesis was not used directly to determine the HH decision criterion. Instead, we combined the 3:1B and 3:1P decision criteria (derived from the flat-maxima hypothesis) using the independence assumption outlined in Equation 4 to derive the HH decision criterion and completely ignored the objective reward function for the HH condition. Applying Equation 4, the HH decision criterion predicted from the independence assumption equals $0.60 + 0.60$ or 1.20. Converted to $k - k_o$, this equals $1.20 - \ln(9)$ or -1.0. This $k - k_o$ value is denoted by the asterisk on the *x*-axis in Figures 4A and 4B. Notice that the $k - k_o$ value for the HH condition predicted directly from the flat-maxima hypothesis is -0.60, which is closer to optimal than that predicted from the independence assumption.

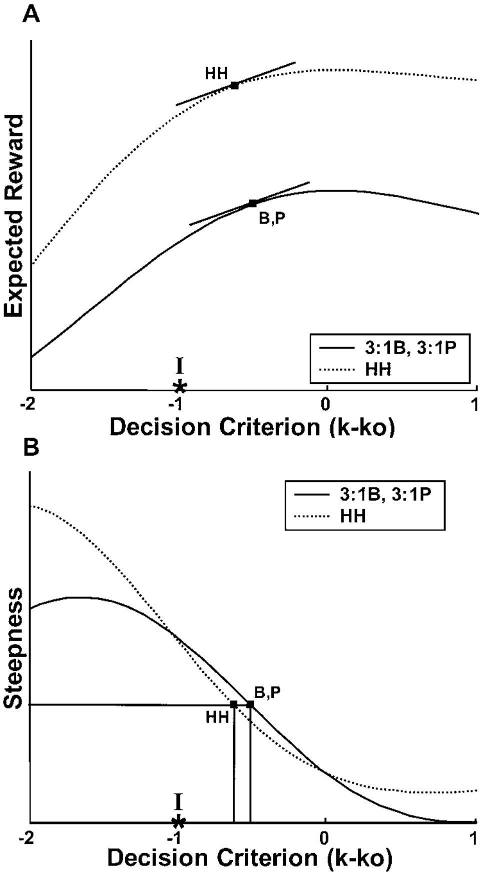

Figure 4. A: Expected reward as a function of the decision criterion (relative to the optimal decision criterion; i.e., $k - k_o$), called the *objective reward function,* for a 3:1 base-rate (3:1B), a 3:1 payoff (3:1P), and a high base-rate, high payoff (HH) condition assuming $d' = 1.0$. The tangent lines, one for the 3:1B and 3:1P conditions and the other for the HH condition, have the same slope and thus reflect the same derivative or steepness. B: Steepness of the objective reward functions from A along with the two points highlighted in A. The two points have identical steepness as denoted by the horizontal line; the vertical lines denote the associated $k - k_o$ values for each condition. The decision criterion ($k - k_o$) predicted from the independence assumption is denoted by the asterisk on the *x*-axis. Notice that the independence assumption predicts greater suboptimality in decision criterion placement in the HH condition than is predicted from the flat-maxima hypothesis. I = independence.

To determine whether the flat-maxima hypothesis always predicts more nearly optimal decision criterion placement in the HH condition than does the independence assumption, we repeated this process across a large number of steepness values [i.e., steepness values ranging from the optimal value of 0 to the steepness associated with $\beta = 1$ or equivalently $\ln(\beta) = 0$] and across the three relevant d' conditions. The results are plotted (in the natural log–natural log space) in Figures 5A, 5B, and 5C for $d' = 1$, 2.2, and 3.2, respectively. If the independence assumption and the flat-maxima hypothesis make the same prediction, then the solid line should fall on top of the dashed line. The results are clear. In every case, except for a small range of $d' = 1.0$ criterion values that are close to optimal (and were rarely observed in the empirical data), the solid line falls above the broken line, suggesting that the HH decision criterion predicted from the flat-maxima hypothesis is more nearly optimal than the HH decision criterion predicted from the independence assumption. [Of course, the predictions converge

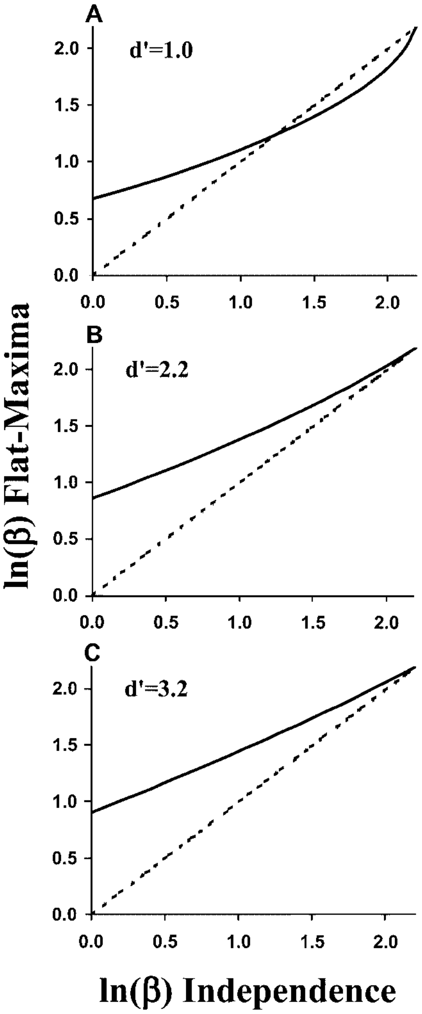

Figure 5. Decision criterion [ln(β)] predicted from the flat-maxima hypothesis plotted against the decision criterion [ln(β)] predicted from the independence assumption of the optimal classifier for the HH condition for $d' = 1.0$ (A), $d' = 2.2$ (B), and $d' = 3.2$ (C). In each panel, the solid line denotes the relationship between the flat-maxima and independence predictions; the broken line is included for comparative purposes and denotes situations in which the two hypotheses make identical predictions. The *x*-axis of each panel displays the criterion values derived from the independence assumption, and the *y*-axis displays the criterion values predicted from the flat-maxima hypothesis.

at ln(β_o).] We test this prediction through a series of model-based analyses and through an examination of decision criterion estimates derived from signal detection theory. As described later, the results provide strong support for the flat-maxima hypothesis over the independence assumption.

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

Competition Between Reward and Accuracy (COBRA) Maximization Hypothesis

The second mechanism assumed to influence decision criterion placement is based on Maddox and Bohil's (1998a) COBRA maximization hypothesis and 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. The COBRA maximization hypothesis 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 6: A displays a 3:1B condition and B displays a 3:1P condition. As suggested by Equation 2, expected reward is maximized in both cases by using the optimal reward-maximizing decision criterion, $k_{ro} = \ln(3)/d'$. Thus, an observer who attempts to maximize expected reward should use the same decision criterion in both conditions. In the 3:1B condition, the decision criterion that maximizes reward also maximizes accuracy, so $k_{ao} = k_{ro}$. Thus, there is effectively no competition between reward and accuracy because a single decision criterion maximizes both simultaneously. However, in the 3:1P condition the accuracy-maximizing decision criterion is different from the reward-maximizing decision criterion. In particular, in the 3:1P condition, $k_{ao} = \ln(1)/d'$. (When base rates are equal, it is always the case that the accuracy-maximizing decision criterion, β_{ao}, is equal to 1.) Thus, when payoffs are manipulated, reward and accuracy maximization cannot be achieved simultaneously because k_{ro} does not equal k_{ao}. 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 3:1P condition and thus will show more conservative cutoff placement in the 3:1P condition than in the 3:1B 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.[2] In the 3:1B condition, this weighting function is essentially irrelevant because $k_a = k_r$, and so all values of w will yield the same value for k. In the 3:1P condition, on the other hand, this weighting function results in a single decision criterion that is intermediate between that for accuracy maximization and that for reward maximization. For example, in Figure 6B, k_1 denotes a case in which w is less than .5, whereas k_2 denotes a case in which w is greater than .5.

Framework for a Hybrid Model

Maddox and Dodd (2001) developed a hybrid model of decision criterion learning that incorporated both the flat-maxima and COBRA hypotheses. In particular, the model assumes that the deci-

---

[2] Other weighting schemes are possible. For example, instead of generating an intermediate decision criterion, the two decision criteria could 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). The current approach is simple to instantiate and has met with reasonable success (Maddox & Dodd, 2001).

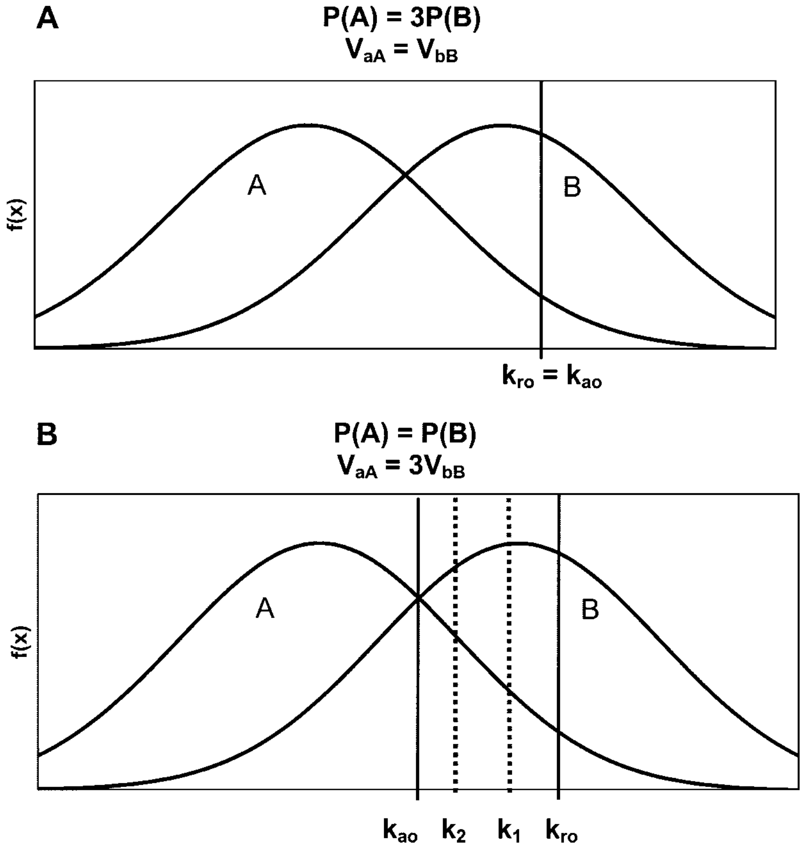

Figure 6. Schematic illustration of the competition between reward and accuracy (COBRA) maximization hypothesis. Solid vertical lines denote the reward- or accuracy-maximizing decision criteria. Dashed vertical lines denote decision criteria that might result when some weight is placed on reward- and accuracy-maximization.

sion criterion used by the observer to maximize expected reward (k_r) is determined by the steepness of the objective reward function (see Figures 2–4). 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 COBRA hypothesis is instantiated in the hybrid model by estimating the accuracy weight, *w*, from the data. A way to facilitate the development of each model is to consider the following equation, which determines the decision criterion used by the observer on condition *i* trials (k_i):

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

It is important to note that 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; so $k_i = k_{ri}$, with no need to apply the accuracy weight, *w*. When payoffs are manipulated, on the other hand, the reward- and accuracy-maximizing decision criteria differ from one another. Fortunately, by pretraining each observer on the category structures in the 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.

All of the models developed in this article are based on the decision bound model outlined in Equation 5. In particular, each model includes a separate noise parameter for each level of d' that represents the sum of perceptual and criterial noise (assumed to be normally distributed; 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, we had each observer complete a number of baseline trials and we required that they meet a stringent performance criterion (see the *Method* section). Finally, each model allows for suboptimal decision criterion placement in that the decision criterion is determined from the flat-maxima hypothesis, the COBRA hypothesis, or both.

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. The nested structure of the models is presented in Figure 7, with the arrows pointing to a more general model and with models at the same level having the same number of free parameters. The number of free parameters (in addition to the noise parameters described above) is presented in parentheses and

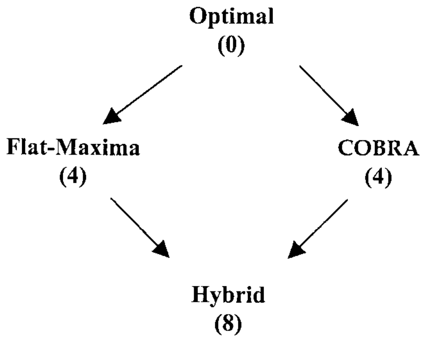

Figure 7. Nested relationship among the decision bound models applied simultaneously to the data from all experimental conditions and four blocks of trials. The number in parentheses denotes the number of free parameters in addition to the noise parameters. The arrows point to a more general model. COBRA = competition between reward and accuracy.

will be discussed, along with the details of the model fitting procedure, in the *Theoretical Analysis* 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 and determines six distinct decision criterion values, two for each of the three category d' conditions. One determines the decision criterion in the 3:1B and 3:1P conditions (which are the same) and the other determines the decision criterion in the HH condition, as depicted in Figure 4. Because the decision criterion values are determined from the objective reward function, this model is constrained to predict that the decision criterion will be closest to optimal in the category $d' = 2.2$ condition, farthest from optimal in the category $d' = 1.0$ condition, and intermediate in the category $d' = 3.2$ condition, as depicted in Figure 3. It is also constrained to predict an HH decision criterion that is closer to optimal than that predicted from the independence assumption, as depicted in Figures 4 and 5. This 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 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. A more general version of the hybrid model was applied to the data that assumed one accuracy weight, w, for the HH conditions and a second accuracy weight for the 3:1P conditions. This model performed poorly and is not discussed further.

Experiment

This article reports the results of a categorization experiment in which three category discriminabilities ($d' = 1.0$, 2.2, and 3.2) were combined factorially with three base-rate/payoff conditions (3:1B, 3:1P, and HH), for a total of nine experimental conditions. Every observer completed three 60-trial blocks of training followed by a 60-trial transfer block in which trial-by-trial feedback was withheld in each of the nine experimental conditions. The goals of this study were many. First and foremost, we wanted to examine decision criterion learning in separate and simultaneous base-rate/payoff conditions across a wide range of category discriminabilities. To achieve this goal, we had each of 12 observers participate in all nine categorization conditions outlined above. Second, we used these data to determine whether predictions from the flat-maxima hypothesis (Figures 4 and 5) were supported in the data from simultaneous base-rate/payoff conditions across levels of category discriminability. The focus of the data analysis was on performance of a series of quantitative models applied simultaneously to the data from all nine experimental conditions, separately for each observer and for each block of trials. Each model instantiated a different set of hypotheses about the effects of category discriminability, base-rate/payoff condition, and the validity of the flat-maxima hypothesis or the 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. Twelve observers were solicited from the University of Texas at Austin community. All observers claimed to have 20/20 vision or vision corrected to 20/20. Each observer completed three approximately 60-min sessions. In each session the observer completed all three base-rate/payoff conditions for a single level of category discriminability. Observers were paid based on 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) and was 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 were 21, 45, and 67 pixels for category $d' = 1.0$, 2.2, and 3.2, respectively. The standard deviation for Categories A and B was 21 pixels for all three category d' levels.

For the category $d' = 1.0$ condition, two sets of 60 stimuli were generated. One set was used in all cases for which the base rates were equal—that is, in the baseline conditions (to be described shortly) and in the 3:1P conditions—and the other was used in the 3:1B and HH conditions. Each set was generated by taking numerous random samples of size 60 from the population and by selecting the sample that best matched the population objective reward function. Stimuli for the category $d' = 2.2$ and category $d' = 3.2$ conditions were generated from the category $d' = 1.0$ samples by applying the appropriate transformation. Two measures were taken to discourage information transfer across category d' conditions. First and most important, before the observer was allowed to begin each of the nine 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, 1998a, 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 nine experimental conditions.

Each experimental condition consisted of three 60-trial training blocks, followed by a 60-trial transfer block. During training, corrective feedback was provided on each trial (see details below). During transfer, no feedback was provided. The same 60 stimuli were presented (in random order) once in each training and transfer block. Table 1 displays the base rates and payoffs for the 3:1B, 3:1P, and HH conditions. Table 2 displays the point totals, accuracy rates, and β values for the optimal classifier separately for each condition. In the baseline condition, $\beta_o = 1$; in the 3:1B and 3:1P conditions, $\beta_o = 3$; and in the HH condition, $\beta_o = 9$.

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 Category B by pressing the appropriate button. During the training phase, the observer's response was followed by 750 ms of feedback. Three lines of feedback were presented. The top line indicated the number of points the observer earned for the response. The next line indicated the potential point earnings for a correct response on each trial (i.e., if an observer responded incorrectly, this line indicated the amount that could have been earned had they chosen the correct response). The third line indicated the number of points that the observer had accumulated up to that point in the session. The feedback was followed by a 125-ms intertrial interval during which the screen was blank. In the transfer phase, the observer's response was followed by a 875-ms intertrial interval 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.

Table 1
Category Base Rates and Payoffs for the Three Base-Rate/Payoff Conditions

	Base rates		Payoffs	
Condition	P(A)	P(B)	V_{aA}	V_{bB}
3:1B	.75	.25	2	2
3:1P	.50	.50	3	1
HH	.75	.25	3	1

Note. P(A) and P(B) denote the base-rate probabilities for hypothetical disease categories A and B. V_{aA} and V_{bB} denote the payoffs associated with correct diagnoses of hypothetical diseases. 3:1B = 3:1 base-rate condition; 3:1P = 3:1 payoff condition; HH = high base-rate, high payoff condition.

Theoretical Analysis

Each of the models in Figure 7 was applied simultaneously to the data from all nine experimental conditions and four blocks but was applied separately to the data from each of the 12 observers. All models contained 12 noise parameters: one for each of the three d' levels in each of the four blocks. Besides the noise parameters, the optimal model contained zero parameters, the flat-maxima and COBRA models contained four parameters (one steepness or one accuracy weight for each block), and the hybrid model contained eight parameters (one steepness and one accuracy weight for each block). 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 4,320 estimated response probabilities from each block [60 trials × 2 response types (A or B) × 9 conditions × 4 blocks]. Because the predicted probability of responding *B,* P(*B*), is equal to 1 − P(*A*), there were 2,160 degrees of freedom in each block. Maximum likelihood procedures (Ashby, 1992b; Wickens, 1982) were used to estimate the model parameters, with the aim being to minimize $-\ln L$. The most parsimonious model was defined as the model with the fewest free parameters for which a more general model did not provide a statistically significant improvement in fit on the basis of likelihood ratio (G^2) tests with α equal to .05 (for a discussion of the complexities of model comparison, see Myung, 2000; Pitt, Myung, & Zhang, 2002).

To determine the most parsimonious model, we took the following steps. First, we compared the maximum likelihood values for the flat-maxima and COBRA models directly to determine which provided the superior account of the data. For 5 of 12 observers, the flat-maxima model provided the better fit; for the remaining 7 observers, the COBRA model provided the better fit. Second, we conducted likelihood ratio tests comparing the fit of the optimal model to the fit of either the flat-maxima model or the COBRA model depending on which of the two provided the better fit for a particular observer. The G^2 values ranged from 10.84 to 287.75. Because the critical value based on $\alpha = .05$ with four degrees of freedom is equal to 9.49, the flat-maxima and COBRA models provided a significant improvement in fit over the optimal model for every observer. Third, we conducted likelihood ratio tests comparing the fit of either the flat-maxima or COBRA model (whichever fit better) to the fit of the hybrid model. The G^2 values ranged from 0.08 to 120.30, with 9 of the 12 G^2 values falling above the critical value of 9.49, again assuming $\alpha = .05$ with four degrees of freedom. Thus, the hybrid model provided the most parsimonious account of the data for 9 of the 12 observers. For the remaining 3 observers, the flat-maxima model was superior in two cases, and the COBRA model was superior in one case. Table 3 displays the maximum-likelihood fit values (the smaller the value the better the fit) along with the percentages of responses accounted for (averaged across observers) for the four models in

Table 2
Optimal Points (per 60-Trial Block), Accuracy, and β Values

	Category $d' = 1$		Category $d' = 2.2$		Category $d' = 3.2$		
Condition	Points	Accuracy	Points	Accuracy	Points	Accuracy	β_o
Baseline	83	69.2	103	85.9	113	94.5	1
3:1B	93	77.8	106	88.7	115	95.4	3
3:1P	93	61.0	106	82.9	115	93.5	3
HH	135	75.9	140	86.7	146	94.7	9

Note. 3:1B = 3:1 base-rate condition; 3:1P = 3:1 payoff condition; HH = high base-rate, high payoff condition.

Figure 7, as well as for three additional models that will be described shortly. Notice that the performance of the hybrid model was quite good, accounting for over 90% of the responses in the data. This finding suggests that both hypotheses incorporated into the hybrid model—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 category discriminabilities.

To determine how an 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 model. These values for the three training blocks and the transfer block averaged across observers are displayed in Figures 8A and 8B along with standard error bars. Several results stand out. First, a one-way ANOVA on the steepness values suggested a significant effect of block, $F(3, 33) = 5.12$, $p < .01$, which was characterized by a marginally significant decline from Training Block 1 to Training Block 2 ($p = .09$), a significant decline from Training Block 2 to Training Block 3 ($p < .05$), and a significant increase from Training Block 3 to the transfer 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 Training Block 3 to the transfer block was unexpected. Maddox and Dodd (2001) did not find an increase in the steepness parameter from Training Block 3 to the transfer block, but their study included six 60-trial training blocks instead of three as in the current study. Perhaps the additional training in Maddox and Dodd's study led to a more stable representation of the decision criterion that was less affected by the removal of trial-by-trial feedback during the transfer trials. Second, a one-way ANOVA on the accuracy weight (*w*) values suggests no effect of training on the weight placed on accuracy. Instead, the weight placed on accuracy remained fairly constant, generally falling in the .40–.50 range. These values closely approximate those of Maddox and Dodd's, which ranged from .40–.60.

With no additional post hoc mechanisms, the hybrid model captures performance in all the conditions, including the simultaneous base-rate/payoff conditions. Recall that the flat-maxima hypothesis predicts that the decision criterion in the simultaneous base-rate/payoff condition will be closer to optimal than that predicted from the independence assumption (see Figure 5). Although the good fits of the hybrid model provide initial support for

Table 3
Maximum Likelihood Fit Values (−lnL) and Percentages of Responses Accounted for (Averaged Across Observers) for Several Models Applied Simultaneously to the Data

Model	$-\ln L$	% accounted for
Optimal	630.59	85.97
Flat maxima	587.84	87.97
COBRA	593.40	87.87
Hybrid	567.36	90.10
Hybrid independence	577.62	88.73
DB independence	532.21	89.67
DB nonindependence	506.27	90.34

Note. COBRA = competition between reward and accuracy; DB = decision bound.

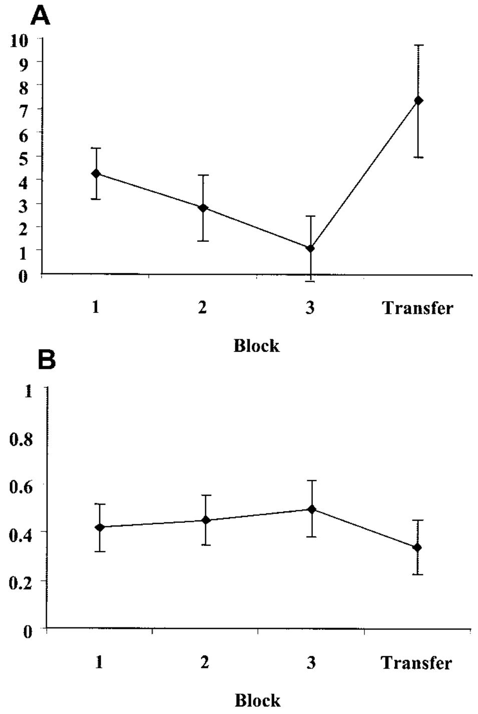

Figure 8. Steepness values (A) and accuracy weight (*w*) values (B) from the hybrid model for the three training blocks and for the transfer block averaged across observers. Error bars represent standard errors of the means.

this prediction, a more rigorous test requires a direct comparison of the hybrid model with a model that assumes an independent combination of base-rate and payoff information in the HH condition. To achieve this goal, we developed a new model called the *hybrid-independence model.* Like the hybrid model, the hybrid-independence model assumes that the reward-maximizing decision criterion in the 3:1B and 3:1P conditions is derived directly from the flat-maxima hypothesis. However, unlike the hybrid model, which assumes that the HH decision criterion is also derived directly from the flat-maxima hypothesis, the hybrid-independence model computes the HH decision criterion from an independent combination of the 3:1B and 3:1P criterion values (Equations 2 and 4). Fortunately, by assuming an independent combination of the 3:1B and 3:1P decision criteria determines the HH decision criterion, no additional parameters are needed; the hybrid model and the hybrid-independence model have the same number of free parameters and their fit values can be compared directly. (The fit value ($-\ln L$) and the percentage of variance accounted for from the hybrid-independence model [averaged across observers] are displayed in Table 3.) The results were clear. For 10 of the 12 observers, the hybrid model provided a better account of the data than did the hybrid-independence model. This effect is significant on the basis of a sign test with a null hypothesis probability of .50 ($p < .05$).

As a second test of the independence assumption, we abandoned the hybrid model framework and went back to the original decision bound model approach. We constructed two nested models. The simpler *decision bound independence* (*DB independence*) *model* freely estimated three 3:1B decision criteria (one for each of the three d' levels) and three 3:1P decision criteria (one for each of the three d' levels), then combined these independently to derive the three HH decision criteria (one for each of the three d' levels). The more general *decision bound nonindependence* (*DB nonindependence*) *model* generalized the DB independence model by freely estimating the three HH decision criteria. Because these models were nested, we performed G^2 tests. The G^2 values ranged from 13.65 to 153.72, with 9 of the 12 G^2 values falling above the critical value of 21.03, assuming $\alpha = .05$ with 12 degrees of freedom. Thus, the DB nonindependence model provided a more parsimonious account of the data for 9 of the 12 observers. In addition, 71% of the time (out of 144 cases; 3 $d' \times 4$ blocks × 12 observers) the freely estimated decision criterion in the HH condition was closer to the optimal decision criterion than that predicted from an independent combination of the freely estimated 3:1B and 3:1P decision criteria, in support of the flat-maxima hypothesis.

As one final test of the flat-maxima prediction, that the decision criterion in the HH condition will be closer to the optimal value than that predicted from the independence assumption, we examined decision criterion estimates derived from signal detection theory. In particular, we computed the decision criterion for all nine experimental conditions separately for each of the four training blocks and 12 observers from the hit rate (defined as a correct high base-rate or high payoff category response) and false alarm rate (defined as an incorrect high base-rate or high payoff category response). We then counted the number of times that the observed HH decision criterion was larger than the HH decision criterion predicted from an independent combination of the observed 3:1B and 3:1P decision criteria. These data are summarized in Table 4, along with sign tests that were conducted assuming a null hypothesis probability of .50.[3] Despite the low power of this test with the relatively small sample size, in 5 of the 12 cases (3 $d' \times 4$ blocks) the null hypothesis was rejected on the basis of a liberal significance level of $\alpha = .10$, and in one additional case the null hypothesis was rejected on the basis of the more traditional significance level of $\alpha = .05$. Perhaps more important, in each of the 12 cases the flat-maxima hypothesis prediction was supported more often than not. Along with the model-based analyses presented above, these data provide strong support for the flat-maxima hypothesis over the independence assumption of the optimal classifier and suggest a theoretically motivated explanation for decision criterion placement in simultaneous base-rate/payoff conditions.

Table 4

Number of Observers (out of 12) for Which the Observed HH Decision Criterion Was Larger Than That Predicted From an Independent Combination of the Observed 3:1B and 3:1P Decision Criteria

	Category discriminability (d')		
Block	1	2.2	3.2
1	9*	9*	7
2	7	7	7
3	7	9*	7
Transfer	10**	9*	9*

Note. HH = high base-rate, high payoff condition; 3:1B = 3:1 base-rate condition; 3:1P = 3:1 payoff condition.
* $p < .10$. ** $p < .05$.

As a further test of the hybrid model framework as it applies to separate and simultaneous base-rate/payoff decision criterion learning, in the next section we undertake a brief reanalysis of data collected by Bohil and Maddox (2001).

Bohil and Maddox (2001): A Reanalysis and Application of the Hybrid Model Framework

Bohil and Maddox (2001) had 8 observers participate in each of 10 categorization conditions constructed from the factorial combination of five base-rate/payoff conditions with two levels of category d' (1.0 and 2.2). Each observer participated in all experimental conditions and completed four 120-trial blocks of training followed by a single 120-trial block of transfer (in which no feedback was provided). Most relevant to the current analyses are the following six conditions: 3:1B, 3:1P, and HH conditions (analogous to the three conditions of the present study), each conducted at $d' = 1.0$ and 2.2. The four models in Figure 7 were applied simultaneously to the data from all six conditions and five blocks

---

[3] In the 3:1B and HH conditions, there were only 20 low base-rate category stimuli presented in each block, and so there was the possibility that no false alarms would be observed. In these cases we arbitrarily assumed that the false alarm rate was 1⁄20 or 5%. This occurred only in three cases: twice for $d' = 3.2$ in Training Block 3 and once for $d' = 3.2$ during the transfer block. In all three of these cases, we found support for the flat-maxima hypothesis over the independence assumption. With these cases excluded, we still found support for the flat-maxima hypothesis for 5 of 10 observers in Training Block 3 with $d' = 3.2$ and for 8 of 11 observers in the transfer block with $d' = 3.2$.

of trials but were applied separately to the data from each of the 8 observers. Following the approach outlined above, first we compared the maximum likelihood values for the flat-maxima and COBRA models directly to determine which provided the superior account of the data. For 5 of 8 observers, the flat-maxima model provided the better fit; for the remaining 3 observers, the COBRA model provided the better fit. Second, we conducted likelihood ratio tests comparing the fit of the optimal model to the fit of either the flat-maxima model or the COBRA model, depending on which of the two provided the better fit for a particular observer. The G^2 values ranged from 107.79 to 753.36. Because the critical value on the basis of $\alpha = .05$ with five degrees of freedom is equal to 11.07, the flat-maxima and COBRA models provided a significant improvement in fit over the optimal model for all 8 observers. Third, we conducted likelihood ratio tests comparing the fit of either the flat-maxima or COBRA model (whichever fit better) to the fit of the hybrid model. The G^2 values ranged from 9.03 to 68.81, with six of the eight G^2 values falling above the critical value of 11.07 (assuming $\alpha = .05$ with five degrees of freedom). Thus, the hybrid model provided the most parsimonious account of the data for 6 of the 8 observers, and the COBRA model provided the most parsimonious account of the data for the remaining 2 observers. Table 5 displays the maximum likelihood fit values along with the percentages of responses accounted for (averaged across observers) for the four models in Figure 7. Notice that the performance of the hybrid model was quite good, accounting for over 92% of the responses in the data. The dominance and good fit of the hybrid model mirrors that from the experiment reported above.

Following the approach outlined above, we examined the parameter values from the hybrid model to determine how they changed as the observers gained experience with the task. The steepness values (averaged across observers) were 5.67, 3.28, 0.95, 1.92, and 2.15 for Training Blocks 1–4 and the transfer block, respectively. A one-way ANOVA was conducted on the steepness values, revealing a significant effect of block, $F(4, 28) = 5.83$, $p < .01$, which was characterized by significant declines in the steepness values from Training Blocks 1 to 2 and 2 to 3 that stabilized from Training Blocks 3 to 4 and from Training Block 4 to the transfer block. As in the experiment outlined above, this suggests that the observers' reward-maximizing decision criterion approached the optimal value across the first 200 or 300 training trials. It is interesting that during the transfer block, the steepness was very close to that observed in the Training Block 4. Bohil and Maddox's (2001) observers completed four 120-trial blocks before transfer as compared with the three 60-trial blocks in the experiment described above. These results suggest that additional training makes the decision criterion more stable and thus less affected by the removal of trial-by-trial feedback. The accuracy weights (averaged across observers) were .37, .40, .55, .54, and .48 for Training Blocks 1–4 and the transfer block, respectively. A one-way ANOVA revealed a nonsignificant block effect. In line with the results from the experiment outlined above, the accuracy weights fell within an approximate .40–.50 range. Finally, and again in line with the results from the experiment outlined above, the fit of the hybrid model was quite good, accounting for 92.7% of the responses in the data averaged across the 8 observers.

Next we compared the hybrid and hybrid-independence models. Again, the models contained the same number of parameters and thus the fit values can be compared directly. The maximum likelihood fit value and percentage of variance accounted for are displayed in Table 5. In line with the results outlined above, the hybrid model was consistently superior to the hybrid-independence model, providing a better account of the data for 7 of 8 observers. We also applied the DB independence and DB nonindependence models to the data from each observer. Because these models were nested, we performed G^2 tests. The G^2 values ranged from 15.59 to 135.19, with seven of the eight G^2 values falling above the critical value of 18.31, assuming $\alpha = .05$ with 10 degrees of freedom. Thus, the DB nonindependence model provided a more parsimonious account of the data for 7 of the 8 observers. In addition, 59% of the time (out of 80 cases; 2 $d' \times 5$ blocks $\times$ 8 observers) the freely estimated decision criterion in the HH condition was closer to the optimal decision criterion than that predicted from an independent combination of the freely estimated 3:1B and 3:1P decision criteria, in support of the flat-maxima hypothesis.

Finally, we examined the decision criterion estimates derived from signal detection theory and counted the number of observers for which the observed HH decision criterion was larger than the HH decision criterion derived from an independent combination of the 3:1B and 3:1P decision criteria. (It is worth mentioning that we did not observe any cases in which the false alarm rate was zero.) These data are presented in Table 6. Because of the small sample size, we do not include sign test results. Even so, there is strong support for the flat-maxima hypothesis over the independence assumption, especially for $d' = 1.0$.

Table 5

Maximum Likelihood Fit Values ($-\ln L$) and Percentages of Responses Accounted for (Averaged Across Observers) for Several Models Applied Simultaneously to the Data From the Reanalysis of Bohil and Maddox (2001)

Model	$-\ln L$	% accounted for
Optimal	842.77	88.27
Flat maxima	699.50	92.03
COBRA	691.09	91.97
Hybrid	651.28	92.65
Hybrid independence	665.98	91.65
DB independence	615.11	92.12
DB nonindependence	586.36	93.59

Note. COBRA = competition between reward and accuracy; DB = decision bound.

General Discussion

The present study extends our understanding of decision criterion learning across a range of category discriminabilities to situations in which base rates and payoffs are manipulated separately or simultaneously within a condition. Strong support for the flat-maxima and COBRA hypotheses was obtained through model-based analyses of individual observer performance. In addition, a flat-maxima interpretation of decision criterion learning in simultaneous base-rate/payoff conditions, which predicts that decision criterion learning will be more nearly optimal than predicted from the independence assumption of the optimal classifier, provided a consistently better account of the data than an interpretation based on the independence assumption. This superiority held in applica-

Table 6

Number of Observers (out of 8) for Which the Observed HH Decision Criterion Was Larger Than That Predicted From an Independent Combination of the Observed 3:1B and 3:1P Decision Criteria From the Reanalysis of Bohil and Maddox (2001)

	Category discriminability (d')	
Block	1.0	2.2
1	5	7
2	5	2
3	5	4
4	5	3
Transfer	4	5

Note. HH = high base-rate, high payoff condition; 3:1B = 3:1 base-rate condition; 3:1P = 3:1 payoff condition.

tions to the current data and in a reanalysis of data collected by Bohil and Maddox (2001).

These results suggest that the flat-maxima hypothesis, which has successfully accounted for category discriminability effects, might provide a straightforward account of decision criterion placement when base rates and payoffs are manipulated simultaneously. Although promising, more work is clearly needed. For example, the current study examined only situations in which the base-rate and payoff biases were in the same direction. In other words, both the base-rate ratio and the payoff ratio biased the observer toward the same categorization response. It will be important to examine situations in which the base rates bias the observer toward one categorization response whereas the payoffs bias the observer toward another categorization response. In addition, base-rate and payoff ratios other than 3:1 need to be examined. Bohil and Maddox (2002) have begun to examine these issues by including 2:1 and 3:1 ratios and by comparing corresponding base-rate/payoff conditions, in which the base rates and payoffs bias the observer toward the same response, with conflicting base-rate/payoff conditions, in which the base rates and payoffs bias the observer toward different responses.

Comparisons with extant decision criterion learning models are also in order. Two popular and successful models of decision criterion learning are Busemeyer and Myung's (1992) hill-climbing model (HC) and Erev's (1998) criterion reinforcement learning model (CRL; see also Wallsten & Gonzalez-Vallejo's 1994 stochastic judgment model). As an initial test of these models, we generated predicted decision criterion values from each model in the 3:1B, 3:1P, and HH conditions using the parameter values outlined in the original articles (i.e., Busemeyer & Myung, 1992; Erev, 1998). Both models predict better decision criterion learning in the 3:1B condition than in the 3:1P condition, in line with the large body of empirical data. In addition, both models predict that the HH decision criterion will be farther from optimal (i.e., more suboptimal) than the HH decision criterion predicted from the independence assumption. Recall that the flat-maxima hypothesis predicts that the HH decision criterion will be closer to optimal (i.e., less suboptimal) than the HH decision criterion predicted from the independence assumption, and the empirical data support this prediction. Although these results support the flat-maxima hypothesis over the HC and CRL models, it is important to note that the simulations were run using the base version of the HC and CRL models and the original parameter settings. It is likely that the assumptions made in those original articles, although reasonable for those applications, may not be reasonable for the current application. Clearly more work is needed to provide a rigorous comparison of the various models.

In summary, the present study suggests that decision criterion placement in simultaneous base-rate/payoff conditions is more nearly optimal than that predicted from the independence assumption of the optimal classifier. This pattern of results is predicted from the flat-maxima hypothesis, and a model-based instantiation of the flat-maxima hypothesis provides a good account of the data, consistently outperforming model-based instantiations of 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, and Cognition, 14,* 33–53.

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. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology, 37,* 372–400.

Ashby, F. G., & Maddox, W. T. (1994). A response time theory of perceptual separability and perceptual integrality in speeded classification. *Journal of Mathematical Psychology, 33,* 423–466.

Ashby, F. G., Maddox, W. T., & Lee, W. W. (1994). On the dangers of averaging across subjects when using multidimensional scaling or the similarity-choice model. *Psychological Science, 5,* 144–150.

Ashby, F. G., & 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.

Bohil, C. J., & Maddox, W. T. (2002). [A test of the optimal classifiers independence assumption.] Unpublished raw data.

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., & Greenshpan, 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 and Memory, 7,* 344–354.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291.

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. (1995). Base-rate effects in multidimensional perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 288–301.

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., & 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 and Performance, 24,* 301–321.

Maddox, W. T., & Bohil, C. J. (1998a). Base-rate and payoff effects in multidimensional perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1459–1482.

Maddox, W. T., & Bohil, C. J. (1998b). Overestimation of base-rate differences in complex perceptual categories. *Perception & Psychophysics, 60,* 575–592.

Maddox, W. T., & Bohil, C. J. (2000). Costs and benefits in perceptual categorization. *Memory & Cognition, 28,* 597–615.

Maddox, W. T., & Dodd, J. L. (2001). On the relation between base-rate and cost–benefit learning in simulated medical diagnosis. *Journal of Experimental Psychology: Learning, Memory, and 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.

Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology, 44,* 190–204.

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109,* 472–491.

Roth, A. E., & Erev, I. (1995). Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games and 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, and 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, Vol. 1* (2nd ed., pp. 283–374).

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.*

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.

Received October 15, 2001
Revision received September 18, 2002
Accepted October 2, 2002 ■