Journal of Experimental Psychology:
Learning, Memory, and Cognition
2003, Vol. 29, No. 6, 1174–1193

0278-7393/03/$12.00 DOI: 10.1037/0278-7393.29.6.1174

Linear Transformations of the Payoff Matrix and Decision Criterion Learning in Perceptual Categorization

W. Todd Maddox, Corey J. Bohil, and Jeffrey L. Dodd
University of Texas at Austin

The effects of payoff-matrix multiplication, payoff-matrix addition, the presence of long-run gains versus long-run losses, category discriminability, and base rate on decision criterion learning were examined in 2 perceptual categorization experiments. Observers were found to be sensitive to the effects of payoff-matrix multiplication (and category discriminability) on the steepness of the objective reward function in line with predictions from the flat-maxima hypothesis and contrary to the predictions from the payoff-variance hypothesis. Decision criterion learning was best in base-rate conditions, was worst when losses were associated with incorrect responding, and was intermediate when no losses were associated with incorrect responding. This performance profile was well captured by the competition between reward and accuracy (COBRA) hypothesis. A hybrid model framework that instantiates both the flat-maxima and COBRA hypotheses was necessary to account for the data from both experiments.

Throughout each day, living organisms are forced to decide between multiple courses of action on the basis of uncertain information (Ashby & Maddox, 1998). For example, an animal might categorize food as nutrient or poison on the basis of its color. An office manager might categorize applicants as good or poor fits for a particular job on the basis of a composite test score. A medical doctor might diagnose a patient as suffering from a heart attack or indigestion on the basis of the degree of chest pain. Finally, the president might decide to go to war or increase diplomacy on the basis of the perceived threat to national security. These binary categorization decisions vary in their survival value or importance. For example, the need to identify good applicants for a janitorial position might be less important than the need to identify good applicants for a chief executive officer position. Similarly, the need to identify highly nourishing food is more important in the winter when food is scarce than in the spring when food is plentiful. Because of the prevalence of categorization across species and its importance to survival in many cases, it is likely that the cognitive operations involved in making categorization decisions are relatively general. Thus it is reasonable to hypothesize that they should be described by a model that applies across a wide variety of situations, and that places a premium on successful categorization (Ashby & Maddox, 1998).

To fully understand the cognitive processes involved in categorization and to adequately model performance, one needs to identify the basic properties of everyday categorization problems and build them into the experimental categories used in the laboratory. Then one needs to identify the important factors that might vary across category-learning situations, manipulate these experimentally, and determine how well the model accounts for performance changes that occur as a result of these manipulations.

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

This research was supported in part by National Science Foundation Grant SBR-9796206 and National Institute of Mental Health Grant 5 R01 MH59196. We thank Lorilei Cardenas for help with data collection, Jerome R. Busemeyer and Ido Erev for several helpful discussions, and three anonymous reviewers for helpful comments on an earlier version of this article.

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

The Normal Distribution as a Model of Natural Categories

Although no set of properties characterizes all real-world categories, many real-world categories have the following four properties. First, the stimulus dimensions typically are continuous valued, as opposed to binary valued. For example, the degree of chest pain evidenced by a patient varies continuously. Second, most categories contain a large, often infinite, number of exemplars. For example, there are many degrees of perceived threat to national security that would lead to a decision to go to war. Third, many categories have a graded structure where the exemplars are symmetrically and unimodally distributed around some prototype, or at the very least there is good evidence that people make this assumption (e.g., Fried & Holyoak, 1984). Fourth, categories generally overlap, meaning that perfect performance is impossible. For example, no matter how advanced an animal's visual or olfactory system, there will be times when it categorizes a nutrient as poisonous or a poison as a nutrient. The normal distribution possesses these four properties. It is a continuous valued, unimodal, and symmetric distribution that contains an infinite number of exemplars that overlap with exemplars from other distributions. The studies outlined in this article used normally distributed overlapping categories that were composed of a large number of continuous valued stimuli within the framework of a perceptual categorization task (this is referred to as the *randomization technique;* Ashby & Gott, 1988).

Overview of the Current Research

A number of factors can vary across category-learning situations. Three are prevalent in the real world and are of particular importance to the present research. In addition, each of these has specific effects on the behavior of the optimal-classifier model of signal detection theory, a generalization of which forms the foundation of our modeling approach. These include category discriminability, category base rates (prior probabilities), and the costs and benefits associated with correct and incorrect categorization decisions. The mathematical details of the optimal classifier are outlined briefly below, but for now suffice it to say that the optimal-classifier model of signal detection theory has been successfully applied to category learning, as well as to a number of other areas of psychology (e.g., perception, recognition memory, and clinical diagnosis). When applied to category learning, the most common version of signal detection theory assumes that the categories are normally distributed. The optimal classifier uses information about the category representation (i.e., the category distribution and associated parameters) to compute the likelihood ratio of the stimulus for each category. The likelihood ratio is strongly affected by category overlap, which is directly related to the predictability or discriminability of the stimulus dimension for correct categorization. For example, in some cases the nutrient–poison distinction might be based on an easy red–green distinction whereas in other cases it might be based on a hard red–pink distinction. Throughout this article, the term *category discriminability* is used to refer to the standardized distance between category means, also called Category d' (Green & Swets, 1967). The more discriminable two categories are along a particular dimension, the larger the d'. The optimal classifier uses information about the category base rates and costs and benefits associated with correct and incorrect categorization decisions to compute the decision criterion, and compares the likelihood ratio with the decision criterion, selecting the response that maximizes long-run reward. For example, during a recession, many job applicants might be overqualified for the position resulting in more good applicants than poor applicants (a base-rate difference) leading the optimal classifier to set a decision criterion that results in more decisions to hire. Similarly, the benefits associated with correct categorization decisions might differ, and the costs associated with incorrect categorization decisions might differ (the costs and benefits make up the elements of the payoff matrix). For example, the benefit of correctly diagnosing a heart attack exceeds that of correctly diagnosing indigestion, insofar as the former might save the patient's life. Similarly, the cost of incorrectly diagnosing a heart attack patient as suffering from indigestion will exceed that of incorrectly diagnosing an indigestion sufferer as having a heart attack, because the former could lead to the patient's death. This situation would lead the optimal classifier to set a decision criterion that resulted in more heart attack diagnoses even though it might reduce overall accuracy. The costs and benefits are also affected by the survival value (or importance) of the categorization decision. For example, the benefit of correctly choosing war or diplomacy and the cost of an incorrect decision are less when dealing with a militarily weak country, and are much greater when dealing with a militarily strong country with nuclear weapons. Finally, the costs and benefits determine whether long-run reward maximization will lead to overall gains or overall losses. In many cases, optimal responding leads to long-run gains in reward (e.g., food foraged, lives saved), but in other cases optimal responding leads to long-run losses (e.g., gambling).

The aim of the current research is to examine the complex interplay among category discriminability, base-rate, and cost–benefit information on decision criterion learning. The main focus is on the effects of cost–benefit manipulations on decision criterion learning. In previous research, we compared situations in which there was no loss of reward associated with an incorrect response (a no-cost condition) with situations in which there was a loss of reward associated with an incorrect response (cost condition). We found more nearly optimal responding in the no-cost conditions relative to the cost conditions, and attributed the difference to a greater emphasis on accuracy maximization when losses in reward were associated with incorrect responding (e.g., Maddox & Bohil, 2000; Maddox & Dodd, 2001). The present research extends this previous work in two important ways: First, we manipulate systematically the importance of each categorization decision by comparing cases in which the magnitude of the costs and benefits are small with cases in which the magnitude of the costs and benefits are large. For example, in Condition 1 the observer might earn \$.02 for each correct Category A response, \$.00 for each correct Category B response, and −\$.01 (i.e., lose \$.01) for an incorrect response, but in Condition 2 they might earn \$.12 for each correct Category A response, \$.00 for each correct Category B response, and −\$.06 for an incorrect response. Notice that the costs and benefits from Condition 2 can be derived from the costs and benefits in Condition 1 by multiplying each value by a factor of 6, and that in both cases the cost–benefit ratio is 3:1 (see Equation 3). We refer to this as *payoff matrix multiplication* (PMM). Second, we manipulate systematically the nature of the long-run reward, specifically, whether there is a gain or loss associated with long-run reward. For example, Condition 1 might be as defined above, but in Condition 3 the observer might earn \$.01 for each correct Category A response, −\$.01 for each correct Category B response, and −\$.02 for an incorrect response. Importantly, there is a long-run gain associated with optimal performance in Condition 1, but a long-run loss associated with optimal performance in Condition 3. Notice also that the costs and benefits from Condition 3 can be derived from the costs and benefits in Condition 1 by subtracting 1 from each value, and again in both cases the cost–benefit ratio is 3:1 (see Equation 3). We refer to this as *payoff matrix addition.* Collectively, PMM and payoff-matrix addition are referred to as *linear transformations* of the payoff matrix.

Categorization and decision-making problems that differ in the importance of each decision and the direction (loss or gain) of the long-run reward are prevalent in the real world and are thus of empirical interest. In addition, as we will discuss shortly, these manipulations are also of theoretical importance because they provide novel tests of a recent theory of decision criterion learning developed in our lab (Maddox & Dodd, 2001) and provide a critical test of predictions from our theory with others in the literature. Specifically, the flat-maxima hypothesis that forms one cornerstone of Maddox and Dodd's (2001) theory of decision criterion learning predicts better decision criterion learning as the PMM factor increases. The payoff-variance hypothesis predicts better decision criterion learning as the variance among the payoff matrix entries decreases (Bereby-Meyer & Erev, 1998; Erev, 1998). Because increasing the PMM factor increases payoff variance, the two hypotheses make opposing predictions. Finally, this research is unique because it bridges the gap between traditional studies of categorization that focus on the processes involved in category structure learning, and studies in decision making that focus on the processes involved in base-rate and cost–benefit learning by allowing the study of either issue (or both) within a single unified theoretical framework.

Next, we outline briefly Maddox and Dodd's (2001) theory of decision criterion learning and generate predictions from the model for the payoff matrix, base-rate, and category-discriminability manipulations. We then present the results from Experiment 1, which provides a critical test of the flat-maxima and payoff variance hypotheses. We then present the results from Experiment 2, which examines an extended range of linear transformations, and includes a condition in which optimal responding results in long-run losses. We then conclude with some general comments.

A Theory of Decision Criterion Learning and a Hybrid Model Framework

Suppose there are two categories (A and B) in which membership is determined from some variable *x* with normally distributed means μ_A and μ_B, and the standard deviations σ_A and σ_B. For any given *x*, the optimal classifier computes the likelihood ratio, which is called the *optimal decision function,*

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

where f(x|i) is the likelihood of result *x* for category *i.* The optimal classifier has perfect knowledge of the base-rates and the costs and benefits that make up the elements of the payoff matrix and determine the value of the *optimal-decision criterion:*

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

where *P*(*i*) is the base-rate probability for category *i*, V_{aA} and V_{bB} are the benefits associated with correct responses, and V_{bA} and V_{aB} are the costs associated with incorrect responses (with lowercase letters denoting response and uppercase letters denoting categories). Notice that base-rate differences and cost–benefit differences can have the same effect on β_o. For example, if A is three times as common as B, and the payoff matrix is symmetric [i.e., if P(A) = 3P(B) and $(V_{aA} - V_{bA}) = (V_{bB} - V_{aB})$], referred to as a *3:1 base-rate condition,* or if the cost–benefit difference for A is three times larger than for B and the base-rates are equal (i.e., when $[V_{aA} - V_{bA}] = 3[V_{bB} - V_{aB}]$ and P[A] = P[B]) referred to as a *3:1 cost–benefit condition,* then $\beta_o = 3.0$. Notice also that linear transformations of the payoff matrix do not affect the value of the optimal decision criterion (see Equation 2). The optimal classifier uses $l_o(x)$ and β_o to construct the optimal decision rule:

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

Humans rarely behave optimally but appear to use the same strategy as the optimal classifier (e.g., Ashby & Maddox, 1990, 1992; Busemeyer & Myung, 1992; Erev, 1998). Ashby and colleagues proposed decision bound theory, which 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. 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 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 σ_{pi} equals σ_p). Criterial noise exists because there is trial-by-trial variability in the placement of the decision criterion—that is, 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). 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,

$$\text{If } l_o(x_{pi}) > \beta_o + e_c \text{ then respond } B\text{, otherwise respond } A. \quad (4)$$

Early applications of decision-bound theory to data collected from unequal base-rate and unequal cost–benefit conditions replicated two earlier findings; namely, that observers tend to use a decision criterion that was more conservative than the optimal-decision criterion (termed conservative cutoff placement), and that observers were more conservative in cost–benefit as opposed to base-rate conditions, even when the optimal-decision criterion was identical. These applications freely estimated a suboptimal decision criterion, β, from the data (i.e., β replaced β_o in Equation 4; Bohil & Maddox, 2001; Maddox, 1995; Maddox & Bohil, 1998a, 1998b; Maddox & Bohil, 2000). One weakness of this approach is that no mechanism was postulated or formalized to guide decision-criterion placement. Maddox and Dodd (2001) offered a formal theory of decision criterion learning that uses decision-bound theory as the basic modeling framework, but supplements the model by postulating psychologically meaningful mechanisms that 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. The theory proposes two mechanisms (described next) that determine decision-criterion placement.

Before outlining these two mechanisms, it is important to reiterate that the theory has its roots in signal detection theory, and thus makes use of parametric properties of stimulus distributions, likelihood ratios, likelihood-ratio-based decision criteria, and other related constructs. Even so, although these constructs help us to understand and characterize decision criterion learning behavior, we are not arguing that people (or nonhuman animals) possess explicit information about the stimulus distributions, likelihood ratios, or decision criteria. People gain information (likely implicitly[1]) about the categorization problem by experiencing stimuli and the reinforcing consequences of the decisions they

[1] The distinction between implicit and explicit memory and the importance of each system in category learning has received much attention in the past few years. The current thinking is that there are at least two category-learning systems and that one relies predominantly on explicit memory processes, whereas the other relies predominantly on implicit memory processes (Ashby & Ell, 2001, 2002; Ashby, Maddox, & Bohil, 2002; Erickson & Kruschke, 1998; Pickering, 1997; Smith, Patalano, & Jonides, 1998). The neurobiological basis of these systems is an area of active research (Filoteo, Maddox, & Davis, 2001a, 2001b; Knowlton, Mangels, & Squire, 1996; Maddox & Filoteo, 2001; see Ashby & Ell, 2001 for a review).

make (i.e., their reinforcement history). There are likely many mathematical systems (or algorithms) that can capture the behavioral profile of decision criterion learning and ours is only one. Even so, we have found that our approach provides insight into categorization behavior, and leads to specific testable predictions.

Flat-Maxima Hypothesis

The first mechanism is based on the flat-maxima hypothesis (Busemeyer & Myung, 1992; vonWinterfeld & Edwards, 1982). 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; Kubovy & Healy, 1977; Thomas, 1975; Thomas & Legge, 1970). To formalize this hypothesis one can construct the objective reward function that 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; vonWinterfeldt & Edwards, 1982). To generate an objective reward function, one chooses a value for the decision criterion and computes the long-run 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 1A plots expected reward as a function of the deviation between a hypothetical observer's decision criterion $\ln(\beta)$ and the optimal-decision criterion $\ln(\beta_o)$ standardized by category d'. This is referred to as $k - k_o = \ln(\beta)/d' - \ln(\beta_o)/d'$. Notice that for large deviations from the optimal decision criterion, the expected reward is small, and as the deviation from the optimal-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 $k - k_o$ 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 Tangent Lines 1, 2, and 3 in Figure 1A. 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 1 to 2 to 3). In other words, the change in the rate of reward or steepness declines as the decision criterion approaches the optimal-decision criterion. Figure 1B 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 1A are highlighted in Figure 1B. 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 small changes in the placement of the decision criterion near the optimal value will yield perceivable changes in reward. Flat objective reward functions, on the other hand, will lead to less optimal decision-criterion placement because small changes in the placement of the decision criterion near the optimal value will not yield perceivable changes in reward.

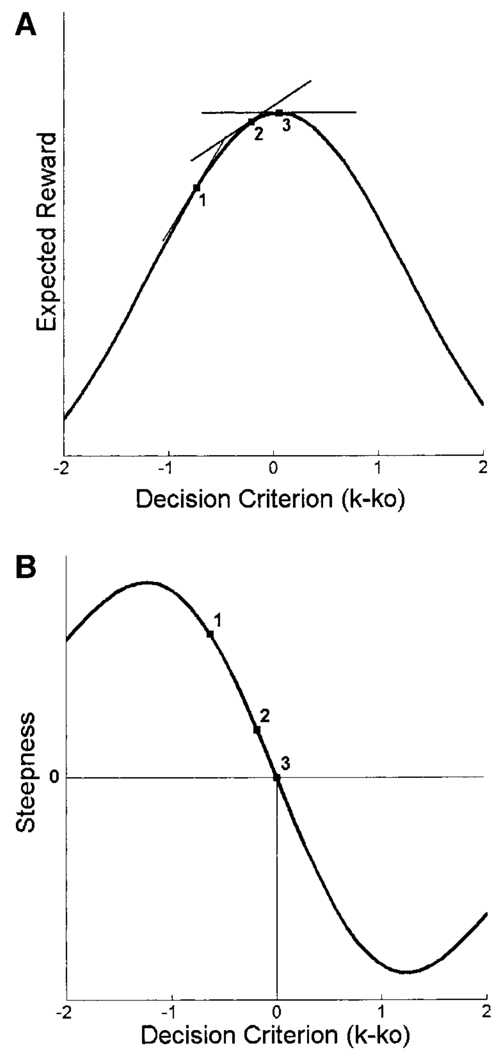

Figure 1. A: Expected reward as a function of the decision criterion (relative to the optimal decision criterion; i.e., $k - k_o$), called the objective reward function. The three lines are the tangent lines at Points 1, 2, and 3 on the objective reward function that denote the derivative or steepness of the objective reward function at each point. B: Steepness of the objective reward function from Panel A along with the three points highlighted in Panel A.

Figure 2A displays the objective reward functions for Category $d' = 1.0$, 2.2, and 3.2 for two 3:1 payoff matrices that are related via a PMM factor of 2 (these are the shallow/no-cost and steep/no-cost payoff matrices from Experiment 1; see Table 1). Figure 2B plots the relationship between the steepness for each objective reward function and $k - k_o$, with the tangent lines in Figure 2A corresponding to the $k - k_o$ values associated with the same fixed steepness value on all six objective-reward functions. The solid horizontal line on Figure 2B denotes the same fixed nonzero steepness value, and the vertical lines denote the associated $k - k_o$ values for each condition. Three comments are in order: First, $k - k_o$ is smaller for conditions in which the PMM factor is 2, implying that decision criterion learning should be better for payoff matrices associated with a larger payoff-matrix multiplication factor. In light of this fact, we refer to payoff matrices with large multipli-

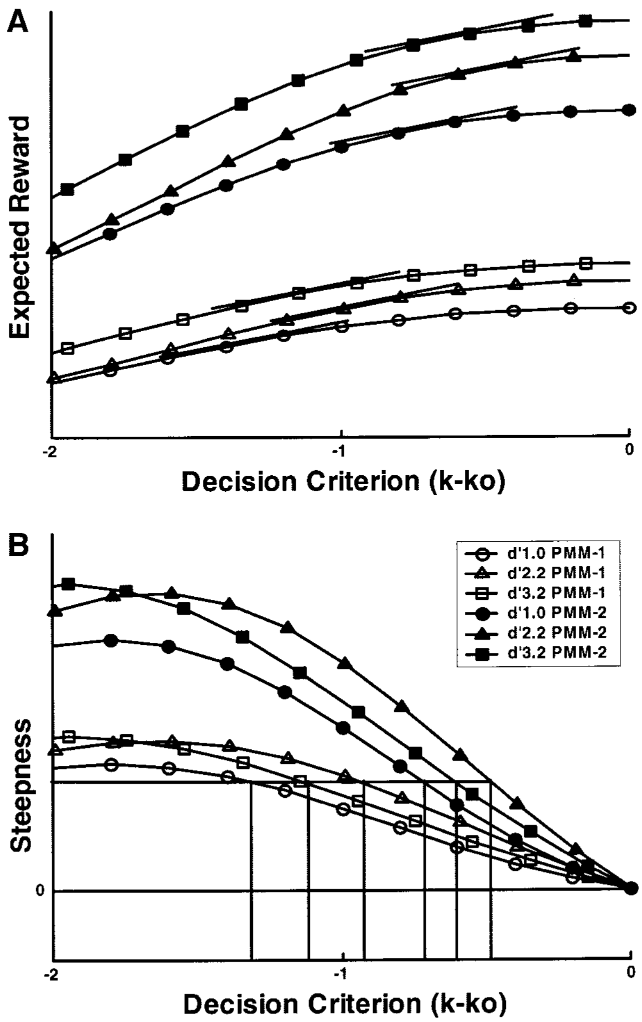

Figure 2. A: Objective reward functions for Payoff-Matrix Multiplication (PMM) Factors 1 and 2 for $d' = 1.0$, 2.2, and 3.2. The tangent lines correspond to the same steepness value on each function. B: Steepness of the objective reward functions from Panel A. The solid horizontal line intersects the tangent points from Panel A.

cation factors as *steep*, and those with small multiplication factors as *shallow*. Second, $k - k_o$ is smallest for $d' = 2.2$, is intermediate for $d' = 3.2$, and is largest for $d' = 1.0$. Third, although not shown in Figure 2, payoff matrix addition does not affect steepness of the objective reward function.

To summarize, if the flat-maxima hypothesis provides a good description of the observer's estimate of the reward-maximizing decision criterion then the following pattern of results should emerge. The observer's estimate of the reward-maximizing decision criterion should be (a) closer to optimal for large PMM factors than for small PMM, (b) closer to optimal for $d' = 2.2$ than for $d' = 1.0$ or 3.2, and (c) unaffected by payoff-matrix addition manipulations.

Because the flat-maxima hypothesis is based on the objective reward function, it applies only to learning of the reward-maximizing decision criterion. The observed decision criterion is assumed to be a weighted average of the reward- and accuracy-maximizing decision criteria. The competition between reward and maximization (COBRA) hypothesis (described in the following section) instantiates the weighting process.

COBRA

The second mechanism assumed to influence decision criterion placement is based on Maddox and Bohil's (1998a) COBRA and was developed to account for the finding that observers show more nearly optimal decision-criterion placement in unequal base-rate than in unequal cost–benefit 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. Figure 3A displays a 3:1 base-rate condition and Figure 3B displays a 3:1 cost–benefit condition. Expected reward is maximized in both cases by using the optimal reward-maximizing decision criterion, $k_{ro} = \ln(3)/d'$. In the 3:1 base-rate condition, the decision criterion that maximizes reward also maximizes accuracy so k_{ao} equals k_{ro}, and so reward and accuracy can be maximized simultaneously. However, in the 3:1 cost–benefit condition $k_{ao} = \ln(1)/d'$, which is different from k_{ro} and so reward- and accuracy-maximization cannot be achieved simultaneously. (When base-rates are equal, it is always the case that the accuracy-maximizing decision criterion, β_{ao}, equals 1.) An observer who places importance on both goals will use an intermediate decision criterion and will show more conservatism than in a 3:1 base-rate condition. To instantiate this hypothesis we assume a simple weighting function: $k = wk_a +$

Table 1
Category Base Rates and Costs–Benefits for Each of the Experimental Conditions From Experiments 1 and 2

	Base rate		Cost–Benefit			
Condition	P(A)	P(B)	VaA	VbA	VbB	VaB
			Experiment 1			
3:1 Base rate	.75	.25	2	0	2	0
Shallow						
Cost	.50	.50	2	−1	−1	0
No cost	.50	.50	3	0	0	1
Steep						
Cost	.50	.50	4	−2	−2	0
No cost	.50	.50	6	0	0	2
			Experiment 2			
Shallow						
Cost-LRG	.50	.50	2	−1	−1	0
Cost-LRL	.50	.50	1	−2	−2	−1
No cost(A)	.50	.50	3	0	0	1
No cost(B)	.50	.50	4	1	1	2
Steep						
Cost-LRG	.50	.50	12	−6	−6	0
Cost-LRL	.50	.50	6	−12	−12	−6
No cost(A)	.50	.50	18	0	0	6
No cost(B)	.50	.50	24	6	6	12

Note. (A) and (B) are versions of the condition. P(A) = the base-rate probability of Category A; P(B) = the base-rate probability of Category B; VaA = the value associated with an *A* response to a Category A stimulus; VbA = the value associated with a *B* response to a Category A stimulus; VbB = the value associated with a *B* response to a Category B stimulus; VaB = the value associated with an *A* response to a Category B stimulus; LRL = long-run loss; LRG = long-run gain.

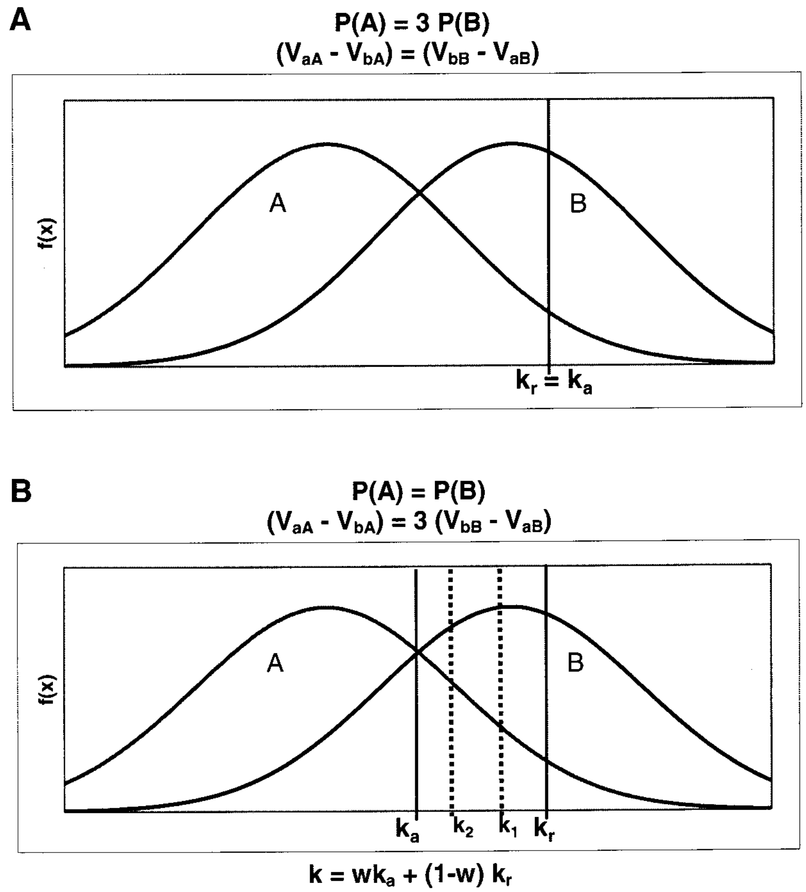

Figure 3. Schematic illustration of the COmpetition between reward and accuracy maximization (COBRA) hypothesis. The k_r decision criterion denotes the decision criterion that is being used by the observer in an attempt to maximize expected reward. The k_a decision criterion denotes the decision criterion that maximizes expected accuracy. The k_1 decision criterion denotes the decision criterion resulting from the COBRA hypothesis with the assumption that less importance or weight ($w < .5$) is being placed on accuracy maximization. The k_2 decision criterion denotes the decision criterion resulting from the COBRA hypothesis with the assumption that more importance or weight ($w > .5$) is being placed on reward maximization. P(A) = probability of a stimulus from Category A; P(B) = probability of a stimulus from Category B; V_{aA} = the value associated with an *A* response to a Category A stimulus; V_{bA} = the value associated with a *B* response to a Category A stimulus; V_{bB} = the value associated with a *B* response to a Category B stimulus; V_{aB} = the value associated with an *A* response to a Category B stimulus.

$(1 - w)k_r$, where w ($0 \leq w \leq 1$) denotes the weight placed on expected accuracy maximization.[2] In the 3:1 base-rate condition, this weighting function is irrelevant because $k_a = k_r$ and so all values of w will yield the same value for k. In the 3:1 cost–benefit condition, this weighting function results in an intermediate decision criterion. For example, in Figure 3B, k_1 denotes a case in which w is less than .5, and 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. The model assumes that the observer's reward-maximizing decision criterion (k_r) is determined by the steepness of the objective-reward function (see Figures 1 and 2), and the decision criterion used on each trial in condition i (k_i) is determined from a weighted combination of k_{ai} and k_{ri} as follows:

$$k_i = wk_{ai} + (1 - w)k_{ri} \tag{5}$$

To ensure that observers have adequate knowledge of the accuracy-maximizing decision criterion, k_a, we pretrain each observer on the category structures in the baseline condition (which is described in the *Method* section). The specific version of the hybrid model examined in each experiment is outlined in the Results and Theoretical Analyses section of each experiment.

[2] Although other weighting schemes are possible (see, e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Maddox & Estes, 1997), the current approach is simple to instantiate and has met with reasonable success (Maddox & Dodd, 2001).

Experiment 1

Maddox and Dodd (2001) examined decision criterion learning in 3:1 base-rate, 3:1 no-cost, and 3:1 cost conditions at three levels of d' (1.0, 2.2, and 3.2). (For related work from the judgment/decision-making literature see Busemeyer & Myung, 1992; Erev, 1998; Erev, Gopher, Itkin, & Greenshpan, 1995; Wallsten, Bender, & Li, 1999; Wallsten & González-Vallejo, 1994). They found support for the flat-maxima hypothesis prediction that decision criterion learning should be better for $d' = 2.2$ than for $d' = 1.0$ or 3.2. They found support for the COBRA prediction that decision criterion learning should be better for base rate than payoff conditions, and after applying the hybrid model found that the weight placed on accuracy was greater in cost than in no-cost conditions and that the importance of accuracy maximization increased when actual losses were associated with incorrect responding (i.e., in the cost conditions).

There are at least two problems with Maddox and Dodd's (2001) conclusion regarding the cost/no-cost comparison. First, the cost–no-cost manipulation was confounded with a difference in payoff variance. There was less variability in the no-cost condition than in the cost condition because the benefit of a correct *A* response was 3 points, the benefit of a correct *B* response was 1 point, and the cost of either incorrect response was 0 points in the no-cost condition, whereas the benefit of a correct *A* response was 4 points, the benefit of a correct *B* response was 0 points, and the cost of either incorrect response was −2 points in the cost condition. Erev and colleagues (Bereby-Meyer & Erev, 1998; Erev, 1998) suggest that decision criterion learning is slower for high payoff-variance matrices relative to low payoff-variance matrices, and this result is predicted a priori from their criterion reinforcement learning (CRL) model. Second, the objective reward function for the cost payoff matrix is steeper than the objective reward function for the no-cost payoff matrix because the cost payoff matrix results when 1 point is subtracted from all entries of the no-cost payoff matrix and the resulting values are multiplied by 2 (i.e., cost entries = 2[no-cost entries − 1]). In short, Maddox and Dodd's (2001) cost and no-cost payoff matrices differed in the presence or absence of costs, payoff variance, and objective reward-function steepness, with one payoff matrix being characterized by no costs, low payoff variance, and a shallow objective-reward function, and the other being characterized by costs, high payoff variance, and a steep objective-reward function.

Experiment 1 of the current study alleviates this confounding of factors by including four cost–benefit conditions constructed from the factorial combination of two levels of objective reward steepness (shallow and steep) with two levels of cost (no cost and cost). The payoff-matrix entries for the resulting four payoff matrices are displayed in Table 1 and are hereafter referred to as the *shallow/cost, steep/cost, shallow/no-cost,* and *steep/no-cost* conditions. The two shallow conditions are related via payoff-matrix addition, and have the same payoff variance. In addition, the two steep conditions are related via payoff-matrix addition, and have the same payoff variance. The two no-cost conditions are related via PMM, and the steep variant has a larger payoff variance. In addition, the two cost conditions are related via PMM, and the steep variant has a larger payoff variance. By including all four payoff matrices in the same experiment we are able to examine the effects of costs versus no costs while holding payoff variance constant, and we are able to examine the effects of payoff variance separately for cost and no-cost conditions. These manipulations also allow a critical test of the flat-maxima and payoff-variance hypotheses because the payoff variance hypothesis predicts worse decision criterion learning in the steep/cost and steep/no-cost conditions relative to the shallow/cost and shallow/no-cost conditions, respectively, and the flat-maxima hypothesis predicts the opposite.

Method

Observers. Six observers were recruited from the community of the University of Texas at Austin. All observers claimed to have 20/20 vision or vision corrected to 20/20. Each observer completed 18 sessions, each of which lasted approximately 30–40 min. Observers were paid on the basis of their accumulated points from the sessions, with a bonus going to the observer with the highest total points.

Stimuli and stimulus generation. The stimulus was a filled white rectangular bar (40 pixels wide) set flush upon a stationary base (60 pixels wide) that was centered on a computer monitor. The length of the bar from the far end to the stationary base varied from trial to trial. There were two categories, A and B, each defined by a specific univariate normal distribution. The distance of separation between the means of Categories A and B were 21, 45, and 67 pixels for $d' = 1.0$, 2.2, and 3.2 respectively. The standard deviation for the distribution of stimuli for Categories A and B was 21 pixels for all three levels of discriminability.

For the $d' = 1.0$ condition, two sets of 60 stimuli were generated. One set was used in all cases for which the base rates were equal (i.e., baseline and payoff-manipulation conditions), and the other set was used for the 3:1 base-rate condition. Both sets were generated by taking numerous random samples of Size 60 from the population, and by selecting the sample set that best matched the population's objective reward function. Stimuli for the $d' = 2.2$ and 3.2 conditions were generated by performing the appropriate transformations on the $d' = 1.0$ sample set.

Four measures were taken to discourage the transfer of information across conditions with different levels of d'. First, and most important, before the observer was allowed to begin each of the 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* (Equation 4) 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. The inclusion of 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. Second, the location of the equal-likelihood decision criterion in the stimulus space was varied across levels of d'. Third, the direction of the bar was changed across levels of d'. For all observers, the bar in the $d' = 1.0$ condition varied in length vertically above the stationary base, in $d' = 2.2$ the bar varied in length horizontally to the right of the stationary base, and in the $d' = 3.2$ condition the bar varied in length vertically below the stationary base. Finally, the labels for the disease categories were changed across different levels of d' conditions.

Each session in the experiment consisted of one block of baseline trials and five 60-trial blocks of training. Corrective feedback was provided on each trial (detailed in the *Procedure* section). The same 60 stimuli were presented once in every block, and the order of the presentation was randomized across blocks. The costs and benefits associated with *A* and *B* responses for each of the five base-rate–payoff conditions are displayed in Table 1. Table 2 displays the point totals, accuracy rates, and optimal decision-criterion values (β_o) for each of the 15 experimental conditions.

Procedure. Observers were informed at the beginning of each session that perfect performance was impossible. However, an optimal level of performance was specified as the goal (in the form of desired point totals). Participants were told they were performing in a simulated medical-diagnosis task, and the length of the bar on the computer monitor indicated the results of a hypothetical medical test. The medical test was designed to distinguish between two diseases, with such names as *burlosis* and *namitis* (these names changed across conditions [see the *Stimuli and stimulus generation* section] and are referred to simply as *Diseases A* and *B* throughout this article). Observers were informed that all of the "patients" they would be diagnosing in the experiment would have one of these two diseases, and each trial would represent a new patient for review. They were informed that the total points earned in the experiment would be converted to money that they would receive at the end of the experiment. As an additional incentive to maximize their points, they were informed that the participant with the highest total of points for the experiment would receive a monetary bonus. Finally, observers were told not to worry about the speed of responding, and the experimenter answered any questions the observers had about the procedure.

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. The observer's response was followed by 750 ms of feedback. Three lines of feedback were presented. The top line indicated the amount of points earned (or lost) for the response. The middle line indicated the potential earnings for a correct response on the trial. If an observer responded correctly, the first and second lines had equivalent values; if the response was incorrect, the second line indicated the amount they could have earned had they chosen the correct response. The third line indicated the sum points the observer had accumulated in that particular session. The feedback was followed by a 125-ms interval during which the computer monitor was blank. At the end of each block of 60 trials, observers were given a break. During the break, the monitor displayed the observer's accumulated points for the session, and the optimal sum points attainable (i.e., the sum points gained by the optimal classifier given the same set of stimuli).

The order of the three different d' conditions was counterbalanced. During the first session for each level of d', the observer completed five 60-trial blocks of baseline condition training. The baseline was completed to ensure that observers had accurate knowledge of the category structures before beginning conditions with base-rate or cost–benefit manipulations. The remaining five experimental conditions were completed over the next four days (so as to finish all sessions for a given level of discriminability in one week). During the second and third sessions, the observer completed the 3:1 base-rate and shallow/no-cost conditions in counterbalanced order. The steep/no-cost, shallow/cost and steep/cost conditions were completed in the remaining three sessions, again in counterbalanced order. Before each experimental session, the observer completed a minimum of 60 baseline trials and was required to meet the accuracy and model-based criterion described earlier before transferring to the experimental condition. To simplify analyses, all data were reorganized so that Category A always referred to the high base-rate or high payoff (i.e., biased) category.

Table 2
Points and Accuracy Predicted from the Optimal Decision Criterion (β_o) that Maximizes Long-Run Reward For Experiments 1 and 2

	$d' = 1.0$		$d' = 2.2$		$d' = 3.2$	
Condition	P	Acc	P	Acc	P	Acc
Experiment 1						
3:1 Base rate	83	69.2	103	85.9	113	94.5
Shallow						
Cost	33	[a]	46	[b]	55	[c]
No cost	93	[a]	106	[b]	115	[c]
Steep						
Cost	67	[a]	93	[b]	109	[c]
No cost	186	[a]	213	[b]	229	[c]
Experiment 2						
Shallow						
Cost-LRG	33	[a]				
Cost-LRL	−27	[a]				
No cost(A)	93	[a]				
No cost(B)	154	[a]				
Steep						
Cost-LRG	200	[a]				
Cost-LRL	−160	[a]				
No cost(A)	560	[a]				
No cost(B)	920	[a]				

Note. d' = category discriminability. (A) and (B) are versions of the condition. $\beta_o = 3$. Blank cells indicate nonapplicable data. P = points; Acc = accuracy. LRG = long-run gain; LRL = long-run loss.
[a] 61.0. [b] 82.9. [c] 93.5.

Results and Theoretical Analysis

Point Totals

Observers were instructed to maximize points, and their monetary compensation was directly tied to their point totals. To determine how category discriminability, and the base-rate/cost–benefit conditions influenced point totals, we computed the deviation from optimal points (observed points − optimal points)/(optimal points − points for 0% correct). We computed the deviation from optimal points for all 15 experimental conditions separately for each of the five blocks and 6 observers. These values (averaged across observer and block) are displayed in Table 3 by category discriminability and base-rate/cost–benefit condition, and were subjected to analyses of variance (ANOVA). The main effect of category discriminability was significant, $F(2, 10) = 43.95$, $p < .001$, with posthoc analyses indicating that decision-criterion placement was significantly closer to optimal for Category $d' = 2.2$ than for Category $d' = 1.0$, and for $d' = 3.2$ than for $d' = 1.0$, but was only marginally significantly different for $d' = 2.2$ and $d' = 3.2$ ($p = .092$). This pattern of results supports the flat-maxima hypothesis. The main effect of base-rate/payoff condition was nonsignificant, but the pattern of results is in line with previous research (Healy & Kubovy, 1981; Higgins, 1987; Kahneman & Tversky, 1979; Maddox & Dodd, 2001), and the predictions from COBRA. Specifically, decision criterion learning was closest to optimal in the base-rate conditions, relative to the cost–benefit conditions. The speculation that decision criterion learning should be closer to optimal in the no-cost condition relative to the cost condition was supported for the steep payoff matrices but not for the shallow payoff matrices. The main effects of block and all two- and three-way interactions were nonsignificant.

Table 3
Deviation from Optimal Points Averaged Across Blocks and Observers For Experiment 1

Condition	d' = 1.0	d' = 2.2	d' = 3.2	Average
Shallow				
Cost	−6.43	−1.36	−2.50	−3.43
No cost	−8.07	−1.62	−2.01	−3.90
Steep				
Cost	−7.39	−1.15	−1.74	−3.43
No cost	−5.68	−0.40	−3.17	−3.08
Cost (Average)	−6.91	−1.25	−2.12	−3.43
No cost (Average)	−6.88	−1.01	−2.59	−3.49
Base rate	−6.61	−0.93	−0.23	−2.59
Overall average	−6.84	−1.09	−1.93	−3.29

Note. d' = category discriminability.

Of central importance to the current study was the performance comparison between steep and shallow payoff matrices. Recall that the flat-maxima hypothesis predicts superior performance for payoff matrices with a steep objective-reward function over payoff matrices with a shallow objective-reward function, whereas the payoff-variance hypothesis predicts the opposite. With the base-rate data removed, an analysis was conducted to determine whether decision criterion learning was affected by objective-reward function steepness. The effect was nonsignificant. A careful examination of Table 3 suggests that the results are mixed. Support for the flat-maxima hypothesis over the payoff-variance hypothesis was observed in the d' = 1.0 no-cost condition, d' = 2.2 no-cost condition, d' = 2.2 cost condition, and the d' = 3.2 cost condition (i.e., the deviation from optimal points was smaller for the steep than for the shallow payoff matrix). Support for the payoff-variance hypothesis was observed in the d' = 1.0 cost, and d' = 3.2 no-cost conditions (i.e., the deviation from optimal points was smaller for the shallow than for the steep payoff matrix).

Model-Based Analyses

All of the models developed in this article are based on the decision-bound model outlined in Equation 4. Specifically, 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., $l_o(x_{pi})$) because each observer completed a number of baseline trials and met 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, COBRA hypothesis, or both. All model-based 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).

To determine whether the flat-maxima and COBRA hypotheses are important in accounting for each observer's data, we developed four base models. Each model makes different assumptions about the k_r and w values. The nested structure of the four base models is presented in Figure 4 with the arrow pointing to a more general model and models at the same level having the same number of free parameters. The number of free parameters (in addition to the noise parameters described earlier) are presented in parentheses with the value on the left denoting the number of parameters needed to fit the data from Experiment 1.

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(d') model assumes that the observer's reward-maximizing decision criterion (k_r) is determined by the steepness of the objective reward function. A single steepness parameter is estimated from the data that determines three distinct k_r values, one for each of the three Category d' conditions (see Figure 2B), but assumes that the same k_r is used for steep and shallow payoff

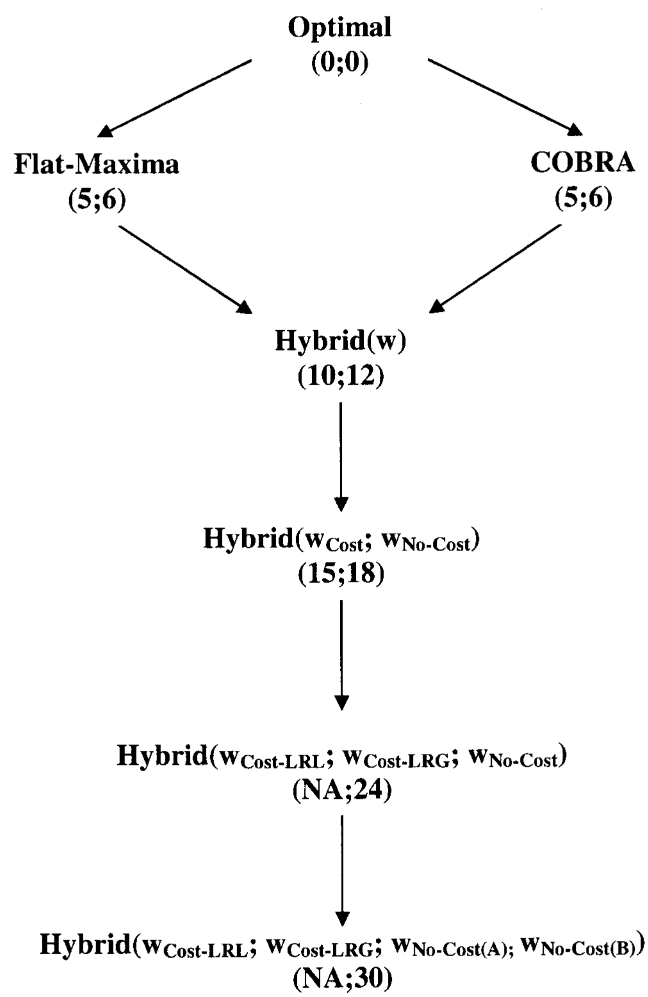

Figure 4. Nested relationship among the decision bound models applied simultaneously to the data from all 15 Experiment 1 conditions, and 8 Experiment 2 conditions. The left-most number in parentheses denotes the number of free parameters needed to apply the model to the data from Experiment 1, and the right-most number denotes the number of free parameters needed to apply the model to the data from Experiment 2. The arrows point to a more general model. NA = not applicable because the latter two models were not applied to the data from Experiment 1. COBRA = COmpetition between reward and accuracy maximization; LRL = long-run loss; LRG = long-run gain.

matrices. In other words, this version assumes that the observer is not sensitive to steepness differences caused by PMM. In the flat-maxima(d'/PMM) model, this single steepness parameter determines six distinct k_r values, two for each of the three Category d' conditions, where one applies for the steep payoff matrices and the other applies for the shallow payoff matrices (see Figure 2B). This version assumes that the observer is sensitive to steepness differences caused by PMM. Both flat-maxima models assume that there is no competition between accuracy and reward maximization (i.e., $w = 0$), and contain the optimal model as a special case in which the steepness value is equal to zero (i.e., $k_r = k_o$). Maddox and Dodd (2001; see also Bohil & Maddox, 2001) also examined a model that assumed no effect of category discriminability. This model was rejected for every observer in the two current experiments and will not be discussed further. 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 5 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. One version of the hybrid model instantiated the flat-maxima(d') assumptions, and another instantiated the flat-maxima(d'/PMM) assumptions. Both versions of the hybrid model contain two free parameters, and they include the previous three models as special cases. A more general version of the hybrid model was also applied to the data. The hybrid(w_{cost}; $w_{no\ cost}$) model is identical to the hybrid model except that one accuracy weight was estimated from the two cost conditions (i.e., shallow/cost and steep/cost), and a separate accuracy weight was estimated from the two no-cost conditions (i.e., shallow/no cost and steep/no cost). This model was developed to determine whether more weight is placed on accuracy maximization in cost conditions as compared with no-cost conditions when payoff variance (and objective reward-function steepness) is controlled.

Each of these five models were applied simultaneously to the data from all 15 experimental conditions and five blocks, but were applied separately to the data from each of the 6 observers. All models contained 15 noise parameters: one for each of the three d' levels in each of the five blocks. Each block consisted of 60 experimental trials, and each observer was required to respond *A* or *B* for each stimulus. Thus each model was fit to a total of 9,000 estimated response probabilities (60 trials $\times$ 2 response types [*A* or *B*] $\times$ 15 conditions $\times$ 5 blocks). Because the predicted probability of responding *B*, P(*B*), equals $1 - \text{P}(A)$ there were 4,500 degrees of freedom. Maximum-likelihood procedures (Ashby, 1992b; Wickens, 1982) were used to estimate the model parameters with the aim being to minimize the maximum-likelihood fit values ($-\ln L$). 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 $\alpha = .05$ (for a discussion of the complexities of model comparison see Myung, 2000; Pitt, Myung, & Zhang, 2002).

Sensitivity to objective reward function steepness. The first aim of the model-based analyses was to determine whether observers were sensitive to the steepness of the objective reward function as it relates to PMM and category discriminability, or just to category discriminability. To achieve this goal we compared the d' and d'/PMM versions of the flat-maxima hypothesis. Although the d' version assumes only two unique objective reward functions, whereas the d'/PMM version assumes four unique objective reward functions, both versions assume that the reward-maximizing decision criterion is determined from a single estimated steepness parameter. Thus both versions of the model have the same number of parameters, and the fit values can be compared directly. The maximum-likelihood fit values (averaged across observers) for all the relevant models are displayed in Table 4 (the smaller the value the better the fit) along with the percent of responses accounted for. First, we compared the d' version and d'/PMM versions of the hybrid(w_{cost}; $w_{no\ cost}$) model. This is the most general model and by definition provides the best account of the data. For all 6 observers the version of the d'/PMM version provided the better account of the data. The improvement in fit was small, but yielded a nearly 1% increase in percentage of responses accounted for. Second, we compared the d' and d'/PMM versions of the hybrid model. For 3 of the 6 observers, the d'/PMM version provided the better account of the data, but in this case the fit values and percent of responses accounted for statistic were essentially identical. Finally, we compared the d' and d'/PMM versions of the flat-maxima model. Recall that this model assumes that the observer does not place any weight on accuracy maximization. As we will see shortly, the flat-maxima model never provides the most parsimonious account of the data, and so these results should be interpreted with caution. Even so, the d'/PMM version provided a better account of the data from only 1 of the 6 observers.

Most parsimonious model by observer and block. Because the d'/PMM version of the flat-maxima hypothesis was generally supported in both versions of the hybrid model (6 of 6 observers for the hybrid [w_{cost}; $w_{no\ cost}$] model, and 3 of 6 observers for the hybrid model), we restrict attention to this version of the model. To determine the most parsimonious of the top four Figure 4 models and the hybrid(w_{cost}; $w_{no\ cost}$) 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 2 of 6 observers the flat-maxima model provided the better fit, and for the remaining 4 of 6 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

Table 4
Maximum Likelihood Fit Value and Percent of Responses Accounted For (Averaged Across Observers) in Experiment 1

Model	$-\ln L$	% Responses
Optimal	1215.83	85.20
Flat maxima		
d'	1065.24	87.59
d'/PMM	1089.37	87.31
COBRA	1057.33	87.75
Hybrid		
d'	1025.16	90.53
d'/PMM	1025.70	90.73
d'; w_{cost}; $w_{no\ cost}$	1009.12	91.59
d'/PMM; w_{cost}; $w_{no\ cost}$	1005.36	92.34

Note. $-\ln L$ = maximum likelihood fit value; d' = category discriminability; PMM = payoff-matrix multiplication; COBRA = COmpetition between reward and accuracy maximization; w = accuracy weight.

model depending on which of the two provided the better fit for a particular observer. The G^2 values ranged from 84.17 to 623.53. Because the critical value was based on $\alpha = .05(df = 5) = 11.07$, 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 5.88 to 86.57, with five of the six G^2 values falling above the critical value of 11.07, again assuming $\alpha = .05(5)$. For the 6th observer, the fit of the flat-maxima model was not significantly improved upon by the hybrid model. Fourth, we conducted likelihood-ratio tests comparing the fit of the flat-maxima with the fit of the hybrid(w_{cost}; $w_{no\ cost}$) model for this one observer. The G^2 value was 55.51 which falls above the critical value of 18.31(10) suggesting that the hybrid(w_{cost}; $w_{no\ cost}$) model provided the most parsimonious account of the data from this observer. Finally, we conducted likelihood-ratio tests comparing the fit of the hybrid model to the fit of the hybrid(w_{cost}; $w_{no\ cost}$) model for the remaining 5 observers. The G^2 values ranged from 5.82 to 106.35, with three of the five G^2 values falling above the critical value of 11.07, again assuming $\alpha = .05(5)$. Thus, the more general hybrid(w_{cost}; $w_{no\ cost}$) model provided the most parsimonious account of the data from 4 of the 6 observers, whereas the less general hybrid model provided the most parsimonious account of the data from the remaining 2 of 6 observers.

Notice that performance of all versions of the hybrid model, especially the hybrid(w_{cost}; $w_{no\ cost}$) model that assumed sensitivity to both d' and PMM, was quite good ranging from 90.53%–92.34% of responses accounted for. The fits of the less general models on the other hand were clearly worse accounting for 85.20%–87.75% of responses in the data. Taken together, these findings suggest 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, payoff matrix addition, PMM, and category discriminability are manipulated.

To determine how the observer's estimate of the reward-maximizing decision criterion and the weight placed on accuracy changed across blocks we examined the steepness and accuracy weight, w, parameters from the hybrid(w_{cost}; $w_{no\ cost}$) model. These values are displayed for the five blocks of trials averaged across observers in Figures 5A and 5B along with standard error bars. Several results stand out. First, a one-way ANOVA on the steepness values suggested a significant effect of block, $F(4, 20) = 7.04$, $p < .01$, that was characterized by a nearly monotonic decline in the steepness value, implying a nearly monotonic approach toward the reward-maximizing decision criterion across blocks. Second, a two-way ANOVA on the accuracy weight (w) values suggests no effect of cost/no-cost, block, or an interaction, although in every block of trials, the average weight placed on accuracy in the no-cost conditions was smaller than that for the cost conditions, suggesting that observers placed more importance on accuracy maximization when actual point losses were associated with incorrect responses.

The fact that the hybrid models accounted for a large percentage of the responses in the data suggests that the conclusions drawn from them have validity. Even so, we decided to compute the decision-criterion estimates derived from signal detection theory (i.e., $k = \ln\beta/d'$) and compared the pattern of results observed for

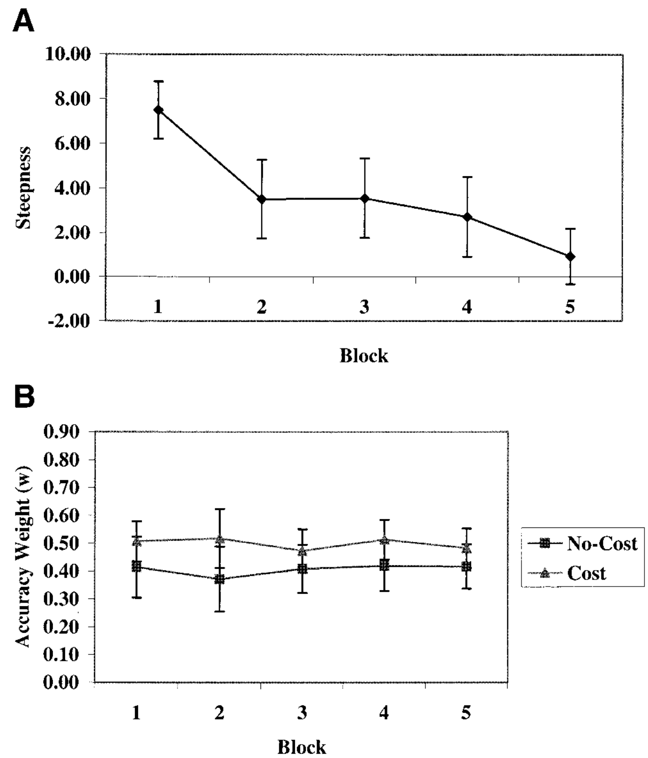

Figure 5. A: Steepness values and B: accuracy weight (w) values from the hybrid(d'/PMM; w_{cost}; $w_{no\ cost}$) model applied simultaneously to the 15 conditions and the five blocks of trials averaged across observers from Experiment 1. Standard error bars are included. PMM = payoff-matrix multiplication.

this measure with the predictions from the hybrid model. Specifically, we computed the decision criterion for all 15 experimental conditions separately for each of the five blocks and 6 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). Then we computed the deviation between the observer's decision criterion and the optimal decision criterion, or $k - k_o$. These values (averaged across observer and block) are displayed in Table 5 by category discriminability and base-rate/payoff condition, and were subjected to ANOVAs. (These values averaged across blocks separately for each observer are displayed in the appendix.) The main effect of category discriminability was significant, $F(2, 10) = 136.71$, $p < .001$, with posthoc analyses indicating that decision-criterion placement was significantly closer to optimal for Category $d' = 2.2$ than for Category $d' = 1.0$, and for $d' = 3.2$ than for $d' = 1.0$, but was not significantly different for $d' = 2.2$ and $d' = 3.2$. The superiority of $d' = 2.2$ and 3.2 over $d' = 1.0$ supports the flat-maxima hypothesis. The nonsignificant difference between $d' = 2.2$ and 3.2 neither supports nor provides evidence against the flat-maxima hypothesis. Even so, there does appear to be some evidence (at least on the basis of the means) to suggest superior decision criterion learning for $d' = 3.2$ over $d' = 2.2$, contrary to the predictions from the flat-maxima hypothesis. It is important to note though that the flat-maxima hypothesis predicts only a small difference in performance between $d' = 2.2$ and 3.2,

Table 5
Deviation from Optimal Decision Criterion ($k - k_o$) Averaged Across Blocks and Observers in Experiment 1

Condition	$d' = 1.0$	$d' = 2.2$	$d' = 3.2$	Average
Shallow				
Cost	−0.67	−0.32	−0.20	−0.40
No cost	−0.70	−0.26	−0.07	−0.34
Steep				
Cost	−0.77	−0.37	−0.13	−0.42
No cost	−0.56	−0.28	−0.19	−0.34
Cost (average)	−0.72	−0.35	−0.16	−0.41
No cost (average)	−0.03	−0.27	−0.13	−0.34
Base rate	−0.56	−0.21	−0.10	−0.29
Overall average	−0.65	−0.29	−0.14	−0.36

Note. Deviations from $k - k_o$ are based on signal-detection criterion estimates. d' = category discriminability.

whereas it predicts much larger performance differences between $d' = 1.0$ and 2.2 and $d' = 1.0$ and 3.2.

The main effect of the base-rate–cost–benefit conditions was nonsignificant, but the pattern of results is in line with previous research (Healy & Kubovy, 1981; Higgins, 1987; Kahneman & Tversky, 1979; Maddox & Dodd, 2001), and the predictions from COBRA. Specifically, decision criterion learning was closest to optimal in the base-rate conditions, was farthest from optimal in the cost conditions, and was intermediate in the no-cost conditions. The main effect of block was significant, $F(4, 20) = 4.75$, $p < .01$, and suggested a gradual shift toward the optimal-decision criterion. The Base-Rate–Cost–Benefit Condition × Block interaction was also significant, $F(16, 80) = 2.49$, $p < .01$, suggesting decision criterion learning in all conditions that was of a larger magnitude in the base-rate condition. All other two- and three-way interactions were nonsignificant.

The flat-maxima hypothesis predicts better decision criterion learning for steep than for shallow payoff matrices, whereas the payoff variance hypothesis predicts the opposite. From Table 5 note that the flat-maxima hypothesis was supported in the $d' = 1.0$, no-cost condition, and in the $d' = 3.2$, cost condition (i.e., $k - k_o$ was smaller for the steep than for the shallow payoff matrix), the payoff-variance hypothesis was supported in the $d' = 3.2$ no-cost condition, the $d' = 1.0$ cost condition, and to a lesser extent in the $d' = 2.2$ cost condition, and the results were equivocal in the $d' = 2.2$ no-cost condition.

Discussion

The principle goal of Experiment 1 was to provide a critical test of the flat-maxima and payoff-variance hypotheses by factorially combining two levels of PMM with a cost and a no-cost condition in such a way that payoff variance could be either controlled or studied directly. As outlined in the beginning of this article, PMM, at a more general level, can be thought of as increasing the significance or importance of each categorization decision because PMM increases the magnitude of the benefits associated with correct categorization and increases the magnitude of the costs associated with incorrect categorization. The flat-maxima hypothesis predicts that this increase in importance will lead to better decision criterion learning, whereas the payoff-variance hypothesis predicts that decision criterion learning will be worse. The results were mixed with some support for the flat-maxima hypothesis prediction that decision criterion learning should be better when the PMM factor was two (i.e., for steep relative to shallow payoff matrices), but also evidence for the payoff-variance hypothesis that makes the opposite prediction. In neither case were the decision criterion learning differences large. Strong support for the flat-maxima hypothesis as it relates to differences in category discriminability were found. Category-discriminability information is likely learned gradually and implicitly over many trials, whereas information about the payoff-matrix entries is likely learned quickly and explicitly over a few trials. Information learned explicitly is more likely available to conscious awareness, and thus is more easily attended or ignored. Although speculative, it may be the case that the PMM constant of two did not result in a big enough difference in payoff-matrix entries across steep and shallow conditions to warrant consistent attention by observers. There is some precedent for this speculation from the base-rate literature. It is well established that base-rate use is poor when learned explicitly (e.g., Kahneman & Tversky, 1973; Tversky & Kahneman, 1974, 1980), but is much better when learned implicitly (e.g., Estes, Campbell, Hatsopoulos, & Hurwitz, 1989; Maddox & Bohil, 1998a, 1998b). It is also possible that the difference in subjective utility across the shallow and steep payoff matrices was not large enough in Experiment 1. In light of these facts, we chose a larger PMM of six in Experiment 2. In addition, we included conditions for which there was a long-run gain or a long-run loss associated with optimal responding.

Experiment 2

The goals of Experiment 2 were three-fold: First, and foremost, we wished to continue our examination of the effects of PMM on decision criterion learning. In light of the equivocal point total and signal-detection decision-criterion results from Experiment 1, and the relatively small performance advantage for the d'/PMM version of the hybrid model over the d' version, we decided to increase the PMM factor to six in Experiment 2. The idea was to determine whether observers might show greater sensitivity to PMM when the steepness difference was increased. Second, we extended the range of payoff-matrix addition factors that we examined. We examined two no-cost conditions: The no-cost(A) condition was identical to that run in Experiment 1. The no-cost(B) condition was derived from the no-cost(A) condition by adding the constant 1 to each payoff-matrix entry. In the no-cost(B) condition the observer always gains points, even when they make an incorrect response. We also examined two cost conditions: The cost-LRG (long-run gain) condition was identical to that run in Experiment 1. The cost-LRL (long-run loss) was derived from the cost-LRG condition by subtracting the constant 1 from all payoff-matrix entries. Interestingly, this small change to the payoff-matrix entries leads to a qualitatively different situation in which even the optimal classifier will lose points over trials. Third, we again tested the validity of the hybrid model, and tested versions of the model that differed in their assumptions about the effects of cost–no cost and long-run reward (long-run gains or losses). To achieve these goals we had each observer complete eight 3:1 cost–benefit conditions constructed from the factorial combination of two levels of PMM (1 or 6) with four levels of payoff-matrix addition (yielding the no-cost(A), no-cost(B), cost-LRG, and cost-LRL conditions) for Category (discriminability) $d' = 1.0$.

Method

Observers. Eight observers were recruited 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 eight sessions, each of which lasted approximately 30–40 minutes. Observers were paid based on their accumulated points from the sessions, with a bonus going to the observer with the highest total points.

Stimuli and stimulus generation. The stimuli used in Experiment 1 ($d' = 1.0$ only) were used in Experiment 2. The costs and benefits associated with *A* and *B* responses for each of the eight 3:1 payoff conditions are displayed in Table 1. Table 2 displays the point totals, accuracy rates, and optimal decision-criterion values (β_o) for the same eight conditions.

Procedure. The procedures were identical to those used in Experiment 1, except that the order of the eight conditions was generated from a Latin square, and each observer completed six 60-trial blocks instead of five 60-trial blocks as in Experiment 1.

Results and Theoretical Analysis

The data from one observer was excluded from all subsequent analyses because of an error that occurred during data collection.

Point Totals

Following the approach taken in Experiment 1, we computed the deviation from optimal points, and subjected these values to a 2 (PMM factor) × 4 (payoff-matrix addition factor) × 6 (block) within-observer ANOVA. The deviation from optimal point values (averaged across observer and block) are displayed in Table 6. The main effect of PMM factor was nonsignificant but held ordinally at all four levels of payoff-matrix addition (see Table 6). The main effect of cost–benefit condition was marginally significant, $F(3, 18) = 2.59$, $p = .085$. Posthoc analyses suggested that the point totals in the no-cost(A) and no-cost(B) conditions were significantly closer to optimal than in the cost-LRG condition. The main effect of block was also significant, $F(5, 30) = 4.42$, $p < .01$, and suggested a gradual shift toward the optimal point total over blocks. There was a significant PMM × Block interaction, $F(5, 30) = 3.99$, $p < .01$, that suggested an approach toward the optimal point total for the steep payoff matrices but little learning for the shallow payoff matrices. To elaborate, during Block 1 the averaged decision from optimal points was −.08 and −.09 and during Block 6 was −.10 and −.05 for the shallow and steep payoff matrices, respectively. All other two-way and three-way interactions were nonsignificant.

Model-Based Analyses

The five Figure 4 models that were applied to the data from Experiment 1 were also applied to the data from the remaining 7 observers from Experiment 2. However, because there was no category-discriminability manipulation in Experiment 2, one variant of the flat-maxima hypothesis assumed that the observer was not sensitive to the effects of PMM on the steepness of the objective reward function (no PMM version), whereas the other variant assumed that the observer was sensitive to the effects of PMM on the steepness of the objective reward function (PMM version). Following the procedure used in Experiment 1, the models were applied simultaneously to the data from all eight experimental conditions and six blocks, but were applied separately to the data from each of the 7 observers. Two additional, more general, versions of the hybrid model (assuming both the no PMM and PMM variants of the flat-maxima hypothesis) were also examined (see Figure 4). The hybrid($w_{\text{cost-LRL}}$; $w_{\text{cost-LRG}}$; $w_{\text{no cost}}$) model estimated one accuracy weight from the two cost-LRL conditions (shallow/cost-LRL and steep/cost-LRL), a second accuracy weight from the two cost-LRG conditions (shallow/cost-LRG and steep/cost-LRG), and a third accuracy weight from the four no-cost conditions (i.e., shallow/no cost[A], shallow/no cost[B], steep/no cost[A], and steep/no cost[B]). This model was developed to instantiate the hypothesis that more weight is placed on accuracy maximization in cost conditions relative to no-cost conditions and to explore the possibility that the weight placed on accuracy might differ across cost conditions in which there was a long-run gain or a long-run loss. A hybrid($w_{\text{cost-LRL}}$; $w_{\text{cost-LRG}}$; $w_{\text{no cost(A)}}$; $w_{\text{no cost(B)}}$) model that estimated four separate accuracy weights was also applied to the data, but in no case did this model provide a significant improvement in fit over a more restricted model, and thus will not be discussed further.

All models contained six noise parameters, one for each of the six blocks of trials. Each model was fit to a total of 5,760 estimated response probabilities, 60 (trials) × 2 (response types: *A* or *B*) × 8 (conditions) × 6 (blocks), yielding 2,880 degrees of freedom. Maximum-likelihood procedures and likelihood-ratio (G^2) tests were again used to identify the most parsimonious model.

Sensitivity to objective reward function steepness. The first aim was to determine whether observers were sensitive to the PMM effect on the steepness of the objective reward function by comparing no PMM with PMM versions of the models. Because the two versions have the same number of parameters, the fit values can be compared directly. The maximum-likelihood fit values (averaged across observers) for all the relevant models are displayed in Table 7 (the smaller the value the better the fit) along with the percentage of responses accounted for each measure. First, we compared the no PMM version with the PMM version for all three versions of the hybrid model (i.e., hybrid, hybrid(w_{cost}; $w_{\text{no cost}}$), hybrid($w_{\text{cost-LRL}}$; $w_{\text{cost-LRG}}$; $w_{\text{no cost}}$)). The results were clear. For all three models, all 7 observers' data were better fit by the PMM than by the no PMM version of the model. For all three versions of the hybrid model, the improvement in fit was much larger than that observed in Experiment 1, and the improvement in percentage of responses accounted for was larger ranging from just under 1% to nearly 2%. Taken together, these results suggest that the larger payoff matrix multiplier used in Experiment 2 led to greater sensitivity to the PMM manipulation. We also compared the no PMM and PMM versions of the flat-maxima model. In line with the results from Experiment 1, the flat-maxima model never provides the most parsimonious account of the data, and so these results should be interpreted with caution. Even so, the PMM

Table 6
Deviation from Optimal Points Averaged Across Blocks and Observers in Experiment 2

Condition	No cost(A)	No cost(B)	Cost-LRG	Cost-LRL	Average
Shallow	−0.08	−0.10	−0.14	−0.12	−0.11
Steep	−0.07	−0.09	−0.13	−0.10	−0.10
Average	−0.08	−0.10	−0.14	−0.11	—

Note. (A) and (B) are versions of condition. LRG = long-run gain; LRL = long-run loss.

Table 7
Maximum Likelihood Fit Value and Percentage Responses Averaged Across Observers in Experiment 2

Model	−lnL	% Responses
Optimal	1,488.95	85.73
Flat maxima		
PMM	1,409.87	86.43
No PMM	1,348.83	86.99
COBRA	1,348.89	87.01
Hybrid		
PMM	1,330.19	88.83
No PMM	1,348.81	87.00
PMM; w_{cost}; $w_{no\ cost}$	1,313.13	90.13
No PMM; w_{cost}; $w_{no\ cost}$	1,322.58	89.35
PMM; $w_{cost\text{-}LRL}$; $w_{cost\text{-}LRG}$; $w_{no\ cost}$	1,307.07	91.01
No PMM; $w_{cost\text{-}LRL}$; $w_{cost\text{-}LRG}$; $w_{no\ cost}$	1,317.48	90.05

Note. −lnL = maximum likelihood fit value; PMM = payoff-matrix multiplication; COBRA = COmpetition between reward and accuracy maximization; LRL = long-run loss; LRG = long-run gain; w = accuracy weight.

version of the model never outperformed the no PMM version of the model.

Most parsimonious model by observer and block. Because the PMM version of the hybrid model was consistently superior to the no PMM version, we restrict attention to this version of the model. In addition, we followed the same procedure used in Experiment 1 to determine the most parsimonious model. 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 all 7 observers, the COBRA model provided the better fit. Second, likelihood-ratio tests comparing the fit of the optimal model to that from the COBRA model revealed G^2 values ranging from 119.15 to 624.29. On the basis of a critical value of 12.59 (assuming α = .05[6]) the COBRA model provided a significant improvement in fit over the optimal model for every observer. Third, we conducted likelihood-ratio tests comparing the fit of the COBRA model to the fit of the hybrid model. The G^2 values ranged from 7.17 to 124.05, with five of the seven G^2 values falling above the critical value of 12.59 (assuming α = .05[6]). Fourth, we conducted likelihood-ratio tests comparing the fit of the COBRA model with that of the hybrid(w_{cost}; $w_{no\ cost}$) model for the two observers best fit so far by the COBRA model. The G^2 values were 22.25 and 43.92, both of which are larger than the critical value of 21.03(12) suggesting that the hybrid(w_{cost}; $w_{no\ cost}$) model provided the most parsimonious account of the data from these two observers. Fifth, we conducted likelihood ratio tests comparing the fit of the hybrid model to the fit of the hybrid(w_{cost}; $w_{no\ cost}$) model for the remaining 5 observers. The G^2 values ranged from 16.35 to 66.30, all of which are larger than the critical value of 12.59(6). Finally, we compared the fit of the hybrid(w_{cost}; $w_{no\ cost}$) model with that from the hybrid($w_{cost\text{-}LRL}$; $w_{cost\text{-}LRG}$; $w_{no\ cost}$) model. The G^2 values ranged from 0.00 to 40.82. On the basis of a critical value of 12.59 (assuming α = .05[6]) the ($w_{cost\text{-}LRL}$; $w_{cost\text{-}LRG}$; $w_{no\ cost}$) model provided a significant improvement in fit over the hybrid(w_{cost}; $w_{no\ cost}$) model for 2 of the 7 observers, and did not provide a significant improvement in fit for the remaining 5 of 7 observers.

In line with the results from Experiment 1, the performance of the PMM variant for all three versions of the hybrid model was quite good ranging from 88.83%–91.01% of responses accounted for, whereas the fits of the less general optimal, flat-maxima, and COBRA models were worse accounting at most for 87% of the responses in the data. Taken together with the results from Experiment 1, these findings provide further evidence that both the flat-maxima and COBRA hypotheses are necessary to provide an adequate account of human decision criterion learning.

To determine how the observer's estimate of the reward-maximizing decision criterion and weight placed on accuracy changed across blocks, we examined the steepness and accuracy weight, w, parameters from the hybrid($w_{cost\text{-}LRL}$; $w_{cost\text{-}LRG}$; $w_{no\ cost}$) model. Although the hybrid(w_{cost}; $w_{no\ cost}$) model provided the more parsimonious account of the data from 5 of the 7 observers, we decided to focus on the more general hybrid($w_{cost\text{-}LRL}$; $w_{cost\text{-}LRG}$; $w_{no\ cost}$) model to provide some insight into the effects of long-run reward (gain vs. loss) on the weight placed on accuracy maximization. These values are displayed for the six blocks of trials averaged across observers in Figures 6A and 6B along with standard error bars. Several results stand out: First, a one-way ANOVA on the steepness values revealed a nonsignificant effect of block. Even so, the steepness value declined sharply from Block 1 to Block 2 and then stabilized.

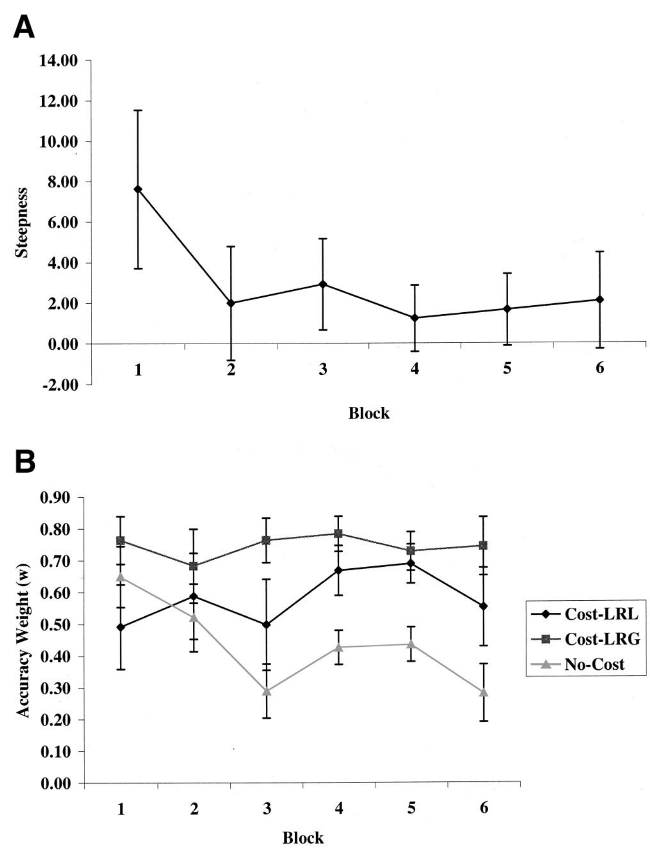

Figure 6. A: Steepness values and B: accuracy weight (w) values from the hybrid(PMM; $w_{cost\text{-}LRL}$; $w_{cost\text{-}LRG}$; $w_{no\ cost}$) model applied simultaneously to the eight conditions and six blocks of trials averaged across observers from Experiment 2. Standard error bars are included. PMM = payoff-matrix multiplication; LRL = long-run loss; LRG = long-run gain.

Second, a two-way ANOVA on the accuracy weight (w) values revealed a significant effect of condition, $F(2, 12) = 4.47, p < .05$. Bonferroni posthoc analyses indicated that the accuracy weight in the cost-LRG condition was significantly larger than the accuracy weight in the no-cost condition. The accuracy weight in the cost-LRL condition was intermediate between that from the cost-LRG and no-cost conditions, but did not differ significantly from either of these. The Condition × Block interaction was significant, $F(10, 60) = 2.48, p < .05$. Further analyses suggested that the interaction was primarily due to the gradual decline in the weight placed on accuracy over blocks in the no-cost condition relative to the fairly stable weight placed on accuracy across blocks in the two cost conditions.

Following the approach taken in Experiment 1, we computed the decision-criterion estimates derived from signal detection theory (i.e., $k = \ln\beta/d'$) and compared the pattern of results observed for this measure with the predictions from the hybrid model. The resulting $k - k_o$ values were subjected to a 2 (PMM factor) × 4 (payoff matrix addition factor) × 6 (block) within-observer ANOVA. The $k - k_o$ values (averaged across observer and block) are displayed in Table 8. The main effect of PMM factor was marginally significant, $F(1, 6) = 4.53, p = .077$, and supported the flat-maxima hypothesis prediction that steep payoff matrices yield better decision criterion learning (i.e., smaller $k - k_o$ values) than shallow payoff matrices. It is worth mentioning that the superiority for steep over shallow payoff matrices held at three of the four levels of payoff-matrix addition (see Table 8), and for the remaining condition [no cost(B)] performance was identical across steep and shallow payoff matrices. The main effect of block was also significant, $F(5, 30) = 9.62, p < .05$, and suggested a gradual shift toward the optimal decision criterion over blocks. The block and PMM main effects were qualified by a significant PMM × Block interaction, $F(5, 30) = 5.94, p < .05$, that suggested much faster decision criterion learning for the steep payoff matrices than for the shallow payoff matrices. To elaborate, during Block 1 the averaged decision-criterion deviation was −.75 and −.73, and during Block 6 it was −.71 and −.37 for the shallow and steep payoff matrices, respectively. Thus, for the shallow payoff matrices there was very little learning (a change of .04 from Block 1 to Block 6), whereas there was substantial learning for the steep payoff matrices (a change of .37 from Block 1 to Block 6). The payoff-matrix addition effect was nonsignificant. Even so, in line with the results from Experiment 1, we found better decision criterion learning in the no-cost conditions relative to the cost conditions. All other two-way and the three-way interactions were nonsignificant.

Table 8
Deviation from Optimal Decision Criterion ($k - k_o$) Averaged Across Blocks and Observers in Experiment 2

Condition	No cost(A)	No cost(B)	Cost-LRG	Cost-LRL	Average
Shallow	−0.71	−0.68	−0.92	−0.80	−0.78
Steep	−0.63	−0.68	−0.63	−0.70	−0.66
Average	−0.67	−0.68	−0.78	−0.75	—

Note. Deviations from $k - k_o$ are based on signal-detection decision criterion estimates. (A) and (B) are versions of condition. LRG = long-run gain; LRL = long-run loss.

Discussion

Unlike some of the results from Experiment 1, in Experiment 2 the model-based analyses and the signal-detection decision-criterion analyses converged. Both the model-based analyses and the detection-theory analyses suggested that decision criterion learning was consistently superior for the PMM factor of six relative to the PMM factor of one. This finding supports the flat-maxima hypothesis and provides evidence against the payoff-variance hypothesis. A more detailed comparison of these two hypotheses is reserved for the General Discussion section. The model-based analyses suggested that more weight was placed on accuracy in cost conditions relative to no-cost conditions, and this result was also supported by the signal-detection analyses. In addition, there was some evidence that more weight was placed on accuracy in cost conditions for which there was a long-run gain than in cost conditions for which there was a long-run loss. This pattern was supported in the signal detection analyses for the shallow payoff matrices but not for the steep matrices and so remains speculative. Even so, we comment on this result in the following section.

General Discussion

The present study extends our understanding of categorization and decision-making processes by examining the effects of linear transformations of the payoff matrix (Experiments 1 and 2), long-run losses versus long-run gains (Experiment 2), and category discriminability (Experiment 1) on decision criterion learning. Maddox and Dodd's (2001) hybrid model of decision criterion learning, which instantiates simultaneously the flat-maxima and COBRA hypotheses, provided a good description of the results. The flat-maxima hypothesis prediction that PMM leads to a steeper objective reward function and thus better decision criterion learning was generally supported in the data from Experiment 1, which used a small PMM factor of two, and was strongly supported by the data from Experiment 2, which used a larger PMM factor of six. These data argue against the payoff-variance hypothesis which predicts worse decision criterion learning for the steep payoff matrices because they are also associated with higher payoff variance. The COBRA prediction that biased cost benefits leads to worse decision criterion learning than biased base rates was supported by the data. In addition, the hypothesis that a loss of points for incorrect responding leads to worse decision criterion learning than no loss of points for incorrect responding was supported by the data. Importantly, and unlike Maddox and Dodd, this pattern was observed when payoff variance and the steepness of the objective reward function were held fixed. Finally, there was some evidence to suggest that decision criterion learning was better in cost conditions where there was a long-run loss associated with optimal responding than in cost conditions where there was a long-run gain associated with optimal responding. COBRA accounted for this pattern by assuming that the weight placed on accuracy maximization was smaller in long-run loss than in long-run gain conditions.

Information-Processing and the Hybrid Model Framework

The hybrid model has its roots in signal detection theory, and thus makes use of parametric properties of stimulus distributions, likelihood ratios, likelihood-ratio-based decision criteria, and other related constructs. Even so, we are not arguing that people possess explicit information about the stimulus distributions, likelihood ratios, or decision criteria. It is possible that they use some form of exemplar memory, prototype abstraction, or some other algorithm. Given the computational level of analysis (Marr, 1982) that we are working at, there are a number of different mathematical systems (or algorithms) that can capture the behavioral profile predicted by the hybrid model. Even so, the hybrid model approach has been very successful at capturing a wide range of category-learning results (see Maddox, 2002, for a review). The flat-maxima hypothesis suggests that learning of the reward-maximizing decision criterion is adaptive and incremental much like the assumptions of the CRL (see also Busemeyer and Myung's, 1992, hill climbing model). It suggests that people behave in accordance with the steepness of the objective reward function, and gradually adjust their reward-maximizing decision criterion toward the optimal value. The COBRA hypothesis suggests that the notion that observers attempt to maximize is complex. Specifically, COBRA suggests that at least two goals are important to observers. These include reward maximization, as instructed, and an emphasis on accuracy maximization. In fact, when costs and benefits are manipulated, the observer who wishes to maximize reward must sacrifice some measure of accuracy. This is analogous to the real-world situation in which the medical doctor biases his or her decision making in such a way that he or she is more willing to incorrectly diagnose a heart attack, than to incorrectly diagnose indigestion. This biasing decreases overall accuracy of responding, but allows the doctor to make certain to minimize the possibility that he or she erroneously diagnoses a heart attack victim as suffering from indigestion. This notion that performance results from a competition between goals (or systems) is gaining popularity in the category-learning literature with the growing body of research in support of multiple systems (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Erickson & Kruschke, 1998; Pickering, 1997; Smith, Patalano, & Jonides, 1998).

PMM: A Test of the Flat-Maxima and Payoff Variance Hypotheses

PMM (by a factor greater than one) increases the steepness of the objective reward function, and thus, on the basis of the flat-maxima hypothesis, should lead to better decision criterion learning. PMM also increases payoff variance. In fact, because PMM is a linear transformation, the variance of the payoff matrix increases by the square of the matrix-multiplication scalar. Erev and colleagues (Bereby-Meyer & Erev, 1998; Erev, 1998) show that (under certain conditions) increasing payoff variance leads to slower decision criterion learning. Erev and colleagues (Erev, 1998; Roth & Erev, 1995) developed a successful model of decision criterion learning called the CRL model that makes the same prediction. Thus, the flat-maxima hypothesis predicts better decision criterion learning for steep (and high variance) payoff matrices, whereas the payoff-variance hypothesis predicts better decision criterion learning for shallow (and low variance) payoff matrices. For a small PMM factor of two (Experiment 1), the results were inconclusive. Under some conditions, the flat-maxima hypothesis was supported, under others the payoff-variance hypothesis was supported, and under still other conditions neither was clearly supported. For a large PMM factor of six, on the other hand (Experiment 2), the results were clearly in line with the predictions from the flat-maxima hypothesis and not the payoff-variance hypothesis. Even so, it is important not to draw strong conclusions from these findings for two reasons: First, Erev and colleagues were not interested in providing a comprehensive test of the payoff-variance hypothesis. Rather, their aim was to compare the predictions from their CRL with Busemeyer and Myung's (1992) hill climbing model to evaluate each model's ability to account for decision criterion learning across two payoff matrices (see Erev, 1998, Table 3). From their analyses, they concluded that the hill climbing model predicted a performance difference because of differences in outcome rank ordering, whereas the CRL predicted a performance difference because of differences in payoff variance. Importantly, the two payoff matrices had different variances, but were characterized by identical objective reward functions, and thus had the same steepness. More work is needed to determine how the CRL is affected by payoff-variance differences that are the result of PMM. Erev (personal communication, September 15, 2001) speculates that payoff variance should not affect the predictions from the CRL, but he recognizes that more work is needed.

Second, there is a critical difference between the framework of the hybrid model and the CRL. Whereas the CRL model predicts a payoff-variance effect on the observable decision criterion, the flat-maxima hypothesis concerns the effects of PMM on the observer's expected reward decision criterion, k_r. This is only one mechanism of the hybrid model framework that helps determine the observable decision criterion. The other mechanism is instantiated by COBRA and assumes that the observable decision criterion, k, is derived from a weighted combination of k_r and the accuracy-maximizing decision criterion, k_a (see Equation 5). Thus, it is possible (and likely) that the effect of PMM on k_r is partially "canceled" by an increase in the weight being placed on accuracy maximization, thus yielding less effect of payoff-matrix steepness at the level of the observable decision criterion. In our view, the fact that these competing mechanisms can be made observable is a testament to the importance of using (a) a within-observer design in which several experimental variables are combined factorially and (b) a model-based approach in which the data from all experimental conditions are modeled simultaneously using a nested hierarchy of models.

Besides offering a critical test of the flat-maxima and payoff-variance hypotheses, the PMM manipulations also provide information about how the importance or seriousness of each category decision affects performance. As the PMM factor increases, the magnitude of both the benefits and costs increases. For example, the benefit of correctly diagnosing a heart attack or cancer and the cost of incorrectly diagnosing either are much greater in magnitude than the costs and benefits associated with correctly or incorrectly diagnosing indigestion or a cold. Although more work is clearly needed, the current research suggests that observers are sensitive to changes in the magnitude of the costs and benefits, and that the optimality of performance increases as the importance of each decision increases.

Payoff Matrix Multiplication Effects Across Experiments

Experiment 1 examined the effects of category discriminability and PMM on decision criterion learning, and tested the flat-maxima hypothesis prediction that both factors should influence decision-criterion placement. This prediction was generally supported by the model-based analyses, but the signal-detection decision-criterion results were equivocal. Experiment 2 focused on the PMM effect on decision criterion learning, and increased the PMM factor from two to six. In Experiment 2, the flat-maxima hypothesis was strongly supported by both the model-based analyses and the signal-detection decision-criterion results. We concluded that strong support for the flat-maxima hypothesis was obtained in Experiment 2 because of the increase in the PMM factor, but a second possibility is that the lack of a category-discriminability manipulation in Experiment 2 might explain the results. As a preliminary test of this hypothesis we reanalyzed some data from a recent study (Bohil & Maddox, in press) that examined the effects of different types of trial-by-trial feedback on decision criterion learning. A detailed discussion of the goals of Bohil and Maddox's (in press) study are beyond the scope of this article. Suffice it to say that two levels of Category d' (1.0 and 2.2) were combined factorially with two levels of PMM (multiplication by a factor of one or six) and a number of other factors related to the nature of the trial-by-trial feedback. Bohil and Maddox (in press) examined only versions of the flat-maxima and hybrid model that assumed sensitivity to both Category d' and PMM (i.e., the d'/PMM versions of the model). For the present purposes, we refit these models under the assumption that observers were sensitive to the Category d' manipulation but not to the PMM manipulation (i.e., the d' version of the model). Three versions of the hybrid model with either 1, 2, or 4 accuracy weight parameters were applied in Bohil and Maddox (in press), and the same three versions were refit under these assumptions. The results were clear, and mirrored those from Experiment 2. Specifically, for 6, 7, and 8 of the 8 observers the d'/PMM version of the model fit best for the hybrid model with 1, 2, and 4 accuracy weight parameters, respectively. In addition, and again in line with the results from Experiment 2, the d'/PMM version of the flat-maxima model never fit better than the d' version, but in no case did the flat-maxima model provide the most parsimonious account of the data. These results suggest that the support for the flat-maxima hypothesis as it relates to PMM manipulations in Experiment 2 is due to the larger PMM factor, and is not due to the lack of a category-discriminability manipulation. Of course, Bohil and Maddox (in press) included only two levels of category discriminability, and manipulated the nature of the trial-by-trial feedback, whereas Experiment 1 included three levels of category discriminability and did not manipulate trial-by-trial feedback, and so further work is needed.

Long-Run Gains Versus Long-Run Losses

Although studies in the decision-making literature have examined conditions in which the behavior of the optimal classifier is characterized by long-run losses (e.g., Barkan, Zohar, & Erev, 1998), all of our previous work has examined situations in which the behavior of the optimal classifier is characterized by long-run gains. Thus, it was of interest to compare cost conditions in Experiment 2 that were characterized by either long-run gains or long-run losses. Importantly, other relevant factors, such as payoff variance, steepness of the objective reward function, and the value of the optimal decision criterion were controlled. Because both conditions are characterized as cost conditions, and because the objective reward function steepness was controlled, no a priori predictions (except possibly the null prediction of no difference) were offered. Instead, we used the hybrid model as a data-analysis tool to explore possible effects of long-run losses versus long-run gains on decision criterion learning. Several patterns of results seemed reasonable. One possibility is that the presence of costs leads to greater weight placed on accuracy, but that the sign of the long-run reward (gain vs. loss) does not affect the magnitude of the accuracy weight. In this scenario the weight placed on accuracy should be identical across the cost-LRL and cost-LRG conditions. A second possibility is that the weight placed on accuracy is greater when there are losses associated with long-run responding, because the observer speculates (perhaps implicitly) that more accurate responding leads to long-run gains. In this scenario the weight placed on accuracy should be larger in the cost-LRL condition than in the cost-LRG condition. A third possibility is that the weight placed on accuracy is less when there are losses associated with long-run responding because the observer recognizes (again possibly implicitly) that he or she is losing points over trials and that no strategy will yield a long-run gain. This allows the observer to focus less on accuracy maximization leading to a smaller accuracy weight in the cost-LRL condition. The accuracy weights from the hybrid($w_{\text{cost-LRL}}$; $w_{\text{cost-LRG}}$; $w_{\text{no cost}}$) model suggest that greater weight is placed on accuracy-maximization in the cost-LRG condition relative to the cost-LRL condition providing evidence against the second possibility, and in support of the third possibility. That said, the fact that the hybrid(w_{cost}; $w_{\text{no cost}}$) model provided the best account of the data from 5 of the 7 observers provides evidence in support of the possibility that the sign of the long-run reward has no effect on performance. Although clearly more systematic work is needed on this important issue, the results are quite interesting and suggest that future research should focus on the distinction between long-run gains and long-run losses.

In conclusion, the present study suggests that observers are sensitive to the effects of category discriminability and PMM (at least when the PMM factor is large) on the steepness of the objective reward function, and adjust their reward-maximizing decision criterion in accordance with predictions from the flat-maxima hypothesis. Payoff matrix addition affects decision criterion learning in such a way that the presence of losses for incorrect responding leads to worse decision criterion learning than the absence of costs associated with incorrect responding. This pattern of results is accounted for by the COBRA hypothesis by assuming that losses for incorrect responding lead to a greater emphasis being placed on accuracy maximization. Preliminary data are presented that suggest that long-run gains associated with optimal responding (along with losses for incorrect responding) lead to a greater emphasis being placed on accuracy maximization than long-run losses associated with optimal responding (along with losses for incorrect responding). This collective pattern of results was well captured by a hybrid model of decision criterion learning that instantiates both the flat-maxima and COBRA hypotheses.

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., & Ell, S. W. (2001). The neurobiological basis of category learning. *Trends in Cognitive Sciences, 5,* 204–210.

Ashby, F. G., & Ell, S. W. (2002). Single versus multiple systems of learning and memory. In J. Wixted & H. Pashler (Eds.), *Stevens' handbook of experimental psychology: Vol. 4. Methodology in experimental psychology* (3rd ed., pp. 655–692). New York: Wiley.

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., & Maddox, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 598–612.

Ashby, F. G., & Maddox, W. T. (1992). Complex decision rules in categorization: Contrasting novice and experienced performance. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 50–71.

Ashby, F. G., & Maddox, W. T. (1998). Stimulus categorization. In M. H. Birnbaum (Ed.), *Measurement, Judgment, and Decision Making* (pp. 251–301). New York: Academic Press.

Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information-integration category learning. *Memory & Cognition, 30,* 666–667.

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.

Barkan, R., Zohar, D., & Erev, I. (1998). Accidents and decision making under uncertainty: A comparison of four models. *Organizational Behavior and Human Decision Processes, 74,* 118–144.

Bereby-Meyer, Y., & Erev, I. (1998). On learning to become a successful loser: A comparison of alternative abstraction of learning processes in the loss domain. *Journal of Mathematical Psychology, 42,* 266–286.

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. (2003). On the generality of objective versus optimal classifier feedback effects on decision criterion learning in perceptual categorization. *Memory & Cognition, 31,* 181–198.

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 and 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.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127,* 107–140.

Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological Bulletin, 53,* 134–140.

Estes, W. K., Campbell, J. A., Hatsopoulos, N., & Hurwitz, J. B. (1989). Base-rate effects in category learning: A comparison of parallel network and memory storage-retrieval models. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15,* 556–571.

Filoteo, J. V., Maddox, W. T., & Davis, J. D. (2001a). Quantitative modeling of category learning in amnesic patients. *Journal of the International Neuropsychological Society, 7,* 1–19.

Filoteo, J. V., Maddox, W. T., & Davis, J. D. (2001b). A possible role of the striatum in linear and nonlinear categorization rule learning: Evidence from patients with Huntington's disease. *Behavioral Neuroscience, 115,* 786–798.

Fried, L. S., & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 234–257.

Green, D. M., & Swets, J. A. (1967). *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.

Higgins, E. T. (1987). Self-discrepancy: A theory relating self and affect. *Psychological Review, 94,* 319–340.

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

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

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996, September 6). A neostriatal habit-learning system in humans. *Science, 273,* 1399–1402.

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.

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. (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*(1), 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.*

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. (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. (1997). Direct and indirect stimulus-frequency effects in recognition. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 539–559.

Maddox, W. T., & Filoteo, J. V. (2001). Striatal contribution to category learning: Quantitative modeling of simple linear and complex nonlinear rule learning in patients with Parkinson's disease. *Journal of the International Neuropsychological Society, 7,* 710–727.

Marr, D. (1982). *Vision.* New York: Freeman.

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

Pickering, A. D. (1997). New approaches to study of amnesic patients: What can a neurofunctional philosophy and neural network methods offer? *Memory, 5,* 255–300.

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.

Smith, E. E., Patalano, A., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition, 65,* 167–196.

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). San Francisco: 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, September). 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.* Hillsdale, NJ: Erlbaum.

von Winterfeldt, D., & Edwards, W. (1982). Costs and payoffs in perceptual research. *Psychological Bulletin, 91,* 609–622.

Wallsten, T. S., Bender, R. H., & Li, Y. (1999). Dissociating judgment from response processes in statement verification: The effects of experience on each component. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25,* 96–115.

Wallsten, T. S., & González-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.

Appendix

Table A1
Deviation From Optimal Decision Criterion ($k - k_o$) for Each of the Experimental Conditions and Observers (Averaged Across Blocks) From Experiment 1

	Observer					
Condition	1	2	3	4	5	6
			$d' = 1.0$			
Shallow						
No cost	−0.60	−0.45	−0.90	−0.86	−0.85	−0.56
Cost	−1.04	−0.70	−0.51	−0.74	−0.41	−0.59
Steep						
No cost	−1.01	0.20	−0.45	−0.97	−0.38	−0.78
Cost	−0.46	−0.76	−0.76	−1.03	−1.07	−0.55
Base rate	−0.62	−0.71	−0.77	0.11	−0.69	−0.66
			$d' = 2.2$			
Shallow						
No cost	−0.30	0.06	−0.30	−0.05	−0.62	−0.36
Cost	−0.65	−0.44	−0.12	−0.13	−0.16	−0.44
Steep						
No cost	−0.54	−0.03	−0.32	−0.27	−0.12	−0.37
Cost	−0.43	−0.18	−0.12	−0.64	−0.37	−0.49
Base rate	0.04	−0.09	−0.18	−0.58	−0.14	−0.29
			$d' = 3.2$			
Shallow						
No cost	−0.25	−0.08	−0.08	−0.19	0.11	0.06
Cost	−0.34	0.02	−0.01	−0.19	−0.22	−0.22
Steep						
No cost	−0.28	−0.21	−0.37	−0.13	−0.04	−0.31
Cost	−0.16	−0.01	0.06	−0.09	−0.18	−0.40
Base rate	−0.62	0.09	−0.11	0.34	0.03	−0.31

Note. d' = Category discriminability.

Table A2

Deviation From Optimal Decision Criterion ($k - k_o$) for Each of the Experimental Conditions and Observers (Averaged Across Blocks) From Experiment 2

	Observer						
Condition	1	2	3	4	5	6	7
Shallow							
Cost-LRL	−0.77	−0.88	−0.40	−0.82	−0.83	−1.00	−0.94
Cost-LRG	−0.87	−0.79	−1.06	−1.03	−1.06	−0.88	−0.79
No cost(A)	−0.81	−0.53	−0.87	−0.51	−0.45	−0.88	−0.94
No cost(B)	−0.87	−0.25	−0.36	−0.77	−0.76	−0.89	−0.89
Steep							
Cost-LRL	−0.90	−0.73	−0.46	−0.78	−0.48	−0.84	−0.70
Cost-LRG	−0.37	−0.87	−0.78	−0.75	−0.40	−0.39	−0.86
No cost(A)	−0.67	−0.75	−0.23	−0.42	−0.91	−0.68	−0.77
No cost(B)	−0.53	−0.91	−0.49	−0.40	−0.57	−0.91	−0.91

Note. (A) and (B) are versions of condition. LRL = long-run loss; LRG = long-run gain.

Received November 4, 2002
Revision received May 20, 2003
Accepted May 21, 2003 ■