Research Report

The Interaction of Payoff Structure and Regulatory Focus in Classification

Arthur B. Markman, Grant C. Baldwin, and W. Todd Maddox

University of Texas

ABSTRACT—*This report brings together research on motivation and learning by exploring how fit of regulatory focus affects people's ability to acquire new categories. Perceptual categories were learned by people with a promotion focus (a situationally determined sensitivity to gains) or a prevention focus (a sensitivity to losses). Classification performance was closest to optimal (as determined by models fit to individual subjects' data) when the regulatory focus matched the structure of the payoffs for the categories. Promotion-focus subjects performed best when the payoffs consisted of all gains. Prevention-focus subjects performed best when the payoffs consisted of all losses.*

This report brings together research on category learning and motivation. There has been significant progress examining how changes in the rewards associated with correct and incorrect responses affect people's ability to identify members of perceptual categories (Maddox, 2002; Maddox & Dodd, 2001). In a typical classification experiment, every correct response is rewarded (or punished) in the same way regardless of the category to which the stimulus is classified; similarly, every incorrect response is rewarded (or punished) in the same way regardless of the category to which the stimulus is classified. In this case, the optimal strategy is one that maximizes accuracy of responses. If the rewards or costs associated with correct and incorrect responses for different categories are asymmetric, however, then optimizing reward requires deviating from optimal accuracy.

For example, Figure 1 illustrates a simple unidimensional classification task. Each stimulus is a dot that appears in a particular location on a computer screen. As the figure shows, the categories have overlapping distributions. The bold solid line indicates the decision rule that yields optimal accuracy. If subjects respond "category A" to every stimulus to the left of this boundary and "category B" to every stimulus to the right, then they will maximize their accuracy of performance.

If one category (which we call the *high-payoff category*) has a higher payoff than the other, and the payoffs for incorrect responses are the same for both categories, then the decision rule that optimizes reward is shifted away from the center of the high-payoff category (see Fig. 1). In previous research varying the payoffs associated with categories, people moved their criterion away from optimal accuracy and toward the optimal-reward criterion, though their adjustments were often more conservative than would be optimal for the particular combination of rewards and costs facing them (Maddox & Bohil, 1998).

In the studies reported here, we used categories that were difficult to learn. The signal detection discriminability of the categories (d') was 1. Pilot research suggested that when the categories were easily discriminable, there were no influences of motivational manipulations on learning. It is not surprising that influences of motivational variables are most evident when the cognitive task is difficult.

As soon as rewards and costs for responses enter into classification learning, it becomes important to consider how motivational factors affect learning (Kruglanski et al., 2002). The motivation literature has made a broad distinction between approach stimuli, which are positive or valuable items, and avoidance stimuli, which are negative or undesirable items (e.g., Carver & Scheier, 1998; Lewin, 1935; Markman & Brendl, 2000). An important line of research has explored the possibility that motivation can be guided by a regulatory focus (Higgins, 1987, 1997). According to this view, people may have a *promotion focus*, in which case they are sensitive to the presence or absence of approach stimuli, or they may have a *prevention focus*, in which case they are sensitive to the presence or absence

Address correspondence to Arthur B. Markman, Department of Psychology, University of Texas, 1 University Station, A8000, Austin, TX 78712; e-mail: markman@psy.utexas.edu.

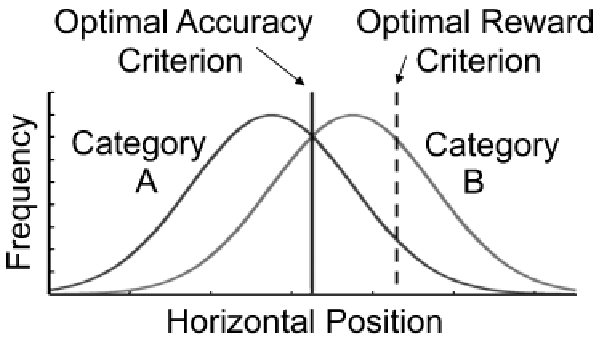

Fig. 1. Category distributions and optimal decision criteria. The two categories represented here are described by one relevant dimension (position of a dot on a computer screen). Category A is associated with a higher payoff or lower punishment than is category B. Consequently, maximizing reward requires selecting a decision criterion to the right of the criterion that maximizes accuracy.

of avoidance stimuli. Higgins and his colleagues demonstrated that the type of regulatory focus people have in a given task affects their cognitive processing and decision making (Crowe & Higgins, 1997; Shah & Higgins, 1997). For example, people who had to solve difficult anagrams were more likely to persevere at their task if they had a situationally induced promotion focus than if they had a prevention focus (Crowe & Higgins, 1997).

How should regulatory focus influence classification learning? Recent work in this area suggests that people prefer a fit between their regulatory focus and their environment (e.g., Higgins, 2000; Higgins, Chen Idson, Freitas, Spiegel, & Molden, 2003). For example, in one study, people with a chronic promotion focus valued a prize most when they focused on what they would gain from owning that prize, whereas people with a chronic prevention focus valued that same prize most when they focused on what they would give up by not having that prize.

Regulatory fit could influence category learning through the associated payoff matrix. Consider the three payoff matrices illustrated in Table 1. In all cases, correct responses yield a higher payoff (or lower punishment) for one category than the other, and incorrect responses are treated equally for the two categories. In the *mixed matrix*, subjects are rewarded with points for correct responses and penalized for incorrect responses. In the *gain matrix*, subjects receive points for both correct and incorrect responses, though they receive more points for a correct response than for an incorrect response. In the *loss matrix*, subjects lose points for both correct and incorrect responses, though they lose fewer points for correct responses than for incorrect responses.

TABLE 1

Payoff Matrices and Performance Criteria for the Three Payoff Conditions

	High-payoff category		Low-payoff category		
Matrix	Correct response	Incorrect response	Correct response	Incorrect response	Performance criterion
Mixed	200	−100	0	−100	3,700
Gain	400	100	300	100	33,700
Loss	−111	−411	−311	−411	−43,000

Note. Subjects started each task with 0 points.

These payoff matrices were all designed to have a signal detection decision criterion (β) of 3. Thus, the optimal classifier would use the same decision criterion across matrices. For this task, if the stimuli are dots presented on a computer screen, and the two categories are defined by location along an arbitrary dimension of 650 pixels, where the high-payoff category has a mean of 275 and a standard deviation of 100 and the low-payoff category has a mean of 375 and a standard deviation of 100, then the optimal accuracy criterion is at 325 pixels and the optimal reward criterion is at 434.5 pixels.

According to the regulatory-fit view, people's performance should be closest to optimal when their regulatory focus matches the structure of the payoff matrix. To assess this prediction, we had subjects perform three category-learning tasks, one with each of the matrices in Table 1. They did so under either a promotion focus or a prevention focus. We manipulated regulatory focus using a situational manipulation derived from previous experiments by Higgins and his colleagues (Higgins, 1997). In the promotion-focus condition, participants were told that in each block, they would receive an entry into a drawing to win $50 if their performance exceeded some criterion. In the prevention-focus condition, an entry ticket for a $50 raffle was displayed on the computer screen at the start of each block, and participants were told that they could keep the ticket unless their performance fell below a criterion, in which case they would lose that ticket. As in previous research, this manipulation was designed so that participants in the promotion- and prevention-focus conditions were in the same objective situation. However, the promotion-focus condition framed the goal as an approach state, and the prevention-focus condition framed the goal as an avoidance state.

Gains in points were accompanied by the sound of a ringing cash register, and losses by an unpleasant buzzer, to heighten the sense of gain and loss on each trial. Subjects' current total of points was shown on a bar on the right side of the screen, with the performance criterion clearly marked.

We fit a decision-bound model to each subject's performance on each task, to determine both the decision criterion that best fit the subject's data for that task and the degree of noise that characterized the subject's error in using that criterion. We expected that people with a promotion focus would have a criterion closer to optimal when they were learning categories with the gain matrix than when they were learning categories with the loss matrix. In contrast, we expected that people with a prevention focus would have a criterion closer to optimal when they were learning categories with a loss matrix than when they were learning categories with a gain matrix. We expected that promotion- and prevention-focus subjects would have similar performance on the task with the mixed matrix, which combined gains and losses.

METHOD

Subjects were 44 members of the University of Texas community (22 each in the promotion- and prevention-focus conditions). They were either paid \$6 for their time or given course credit. In addition, participants were all given the opportunity to receive entries into a drawing to win \$50. One drawing was held for every 10 subjects in the study. The data from 8 subjects (5 in the promotion-focus condition and 3 in the prevention-focus condition) were eliminated from analysis because these data were best fit by a model that had no decision criterion. That is, the performance of these subjects reflected random responding across the stimulus space.

The two primary independent variables were payoff matrix (mixed, gain, and loss), and regulatory focus (promotion and prevention focus). Payoff matrix was manipulated within subjects, and regulatory focus was manipulated between subjects.

Each stimulus was a single dot that appeared on a computer screen. In each block, one dimension (vertical or horizontal) was kept constant, and the other varied. The position of the dot along the dimension that varied was determined by the normal distribution describing the category from which the dot was drawn. The means and standard deviations of the distributions of the two categories were set up so that the d' for the categories was 1.

Participants performed three perceptual classification tasks in a session. They were told to determine which of two categories each dot belonged to, based on its position on the screen. Regulatory focus was manipulated before the first task. Subjects given a promotion focus were told that in each task, if their performance exceeded a certain criterion, they would receive one entry into a drawing to win \$50 for that task. Subjects given a prevention focus were told that for each task, they were going to receive an entry to win \$50, but that they would lose this entry if their performance fell below a particular criterion. Throughout the experiment, participants were reminded that receiving an entry into the drawing for a particular task was contingent on their performance.

Before each task, a different payoff matrix was introduced. Each task consisted of three blocks of 50 trials. To set a performance criterion, we determined the number of points that the optimal classifier would obtain. We then took the difference between this number of points and the number of points a participant would end up with if he or she scored 0% correct over the course of that task. The performance criterion was set at 80% of the value of this difference (see Table 1).

On each trial, participants saw a dot on a computer screen and pressed one button to respond "category A" and a second button to respond "category B." They were told the number of points they received for their response, as well as the maximum number of points they could have received on that trial (i.e., the number of points they would have received for a correct response on that trial). We did not give accuracy feedback directly, because we wanted subjects to focus on the goal of maximizing points rather than the goal of maximizing accuracy.

RESULTS

The data were analyzed by fitting a decision-criterion model to the data from each subject in each block. Only the middle 100 trials were used for these analyses. The first 25 trials were considered practice. The last 25 trials were noisy, because subjects started to reach criterion by that point.

The mean criterion for each matrix for promotion- and prevention-focus subjects is shown in Figure 2. A 2 (regulatory focus) × 3 (payoff matrix) analysis of variance on these criteria revealed a significant interaction between these factors, $F(2, 68) = 3.36$, $p < .05$, $\eta_p^2 = .09$. As shown in Figure 2, this interaction reflected the interaction predicted by the regulatory-fit view. For the gain matrix, subjects given a promotion focus had a criterion significantly closer to optimal than did subjects given a prevention focus, $t(34) = 2.11$, $p < .05$, $d = 0.72$. For the loss matrix, subjects given a prevention focus had a criterion marginally significantly closer to optimal than did subjects given a promotion focus, $t(34) = 1.84$, $p = .074$, $d = 0.63$. For the mixed matrix, the criterion subjects used did not differ significantly between the two regulatory-focus conditions, $t(34) = 0.65$, $p > .10$, $d = 0.22$.

To demonstrate how these models were based on participants' classification performance, we calculated two sets of correlations. First, we looked at the correlation between the deviation of a subject's criterion from the optimal reward criterion and the number of points that subject amassed in a block. One would expect that the further someone's criterion from the optimal reward criterion, the fewer points he or she would amass. Indeed,

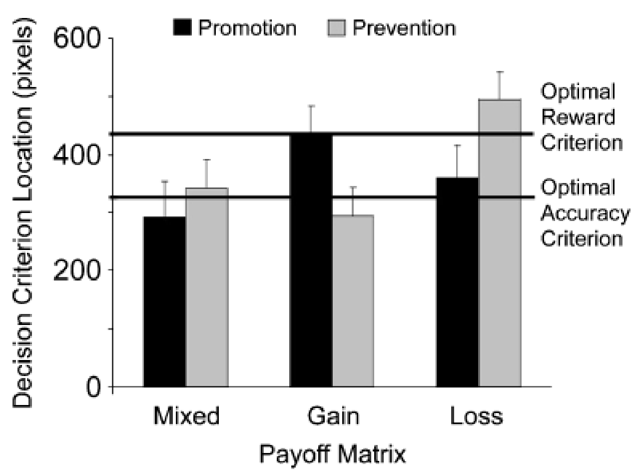

Fig. 2. Mean decision criteria of promotion- and prevention-focus subjects as a function of the payoff matrix. The optimal reward criterion and optimal accuracy criterion are also shown.

TABLE 2

Correlations Between Performance and Deviations From the Optimal Reward Criterion and the Optimal Accuracy Criterion

Regulatory focus	Payoff matrix: Mixed	Gain	Loss
Correlation between deviation from optimal reward and total points			
Promotion	−.95*	−.25	−.49*
Prevention	−.68*	−.82*	−.65*
Correlation between deviation from optimal accuracy and overall accuracy			
Promotion	−.89*	−.48*	−.47
Prevention	−.80*	−.69*	−.71*

*$p < .05$.

as Table 2 shows, all of the correlations were negative, and five of the six were statistically significant.

Second, we correlated the deviation of each subject's criterion from the optimal accuracy criterion with that subject's accuracy. One would expect that accuracy would decrease the further a subject's criterion was from the optimal accuracy criterion. Again, as Table 2 shows, the correlations were consistent with this interpretation; all of the correlations were negative, and five of the six were statistically significant (and the remaining correlation was marginally significant).

DISCUSSION

This perceptual classification task provided us with the opportunity to explore the influence of regulatory focus on perceptual classification learning. The data are consistent with a regulatory-fit hypothesis. People's performance is closest to optimal when their regulatory focus fits the payoff structure of the learning task. People with a promotion focus had a decision criterion closer to optimal than did people with a prevention focus when the payoff structure consisted of all gains. People with a prevention focus had a decision criterion closer to optimal than did people with a promotion focus when the payoff structure consisted of all losses. Performance in the two regulatory-focus conditions was roughly equivalent when the payoff structure had both gains and losses.

The results of this study suggest that combining research on motivation with research on classification will bear fruit in both domains. In classification learning, the emerging research on differences in payoffs for responses to different categories has not incorporated the distinction between approach and avoidance motivation. Clearly, positive and negative reward structures have an important influence on learning performance, and thus a better understanding of motivation will provide insight into classification learning.

On the motivational side, many of the previous demonstrations of effects of regulatory focus have been impressive, but have involved domains where it is hard to characterize the particular strategies that people are using. Thus, many of the existing studies support only broad conclusions, such as that a person may value an item if he or she interacts with that item using processes that are related to the person's current regulatory focus (e.g., Higgins, 2000). The classification methodology allows one to characterize the strategies that subjects use with a great deal of precision, and thus allows a fine-grained view into the way that regulatory focus influences performance.

Acknowledgments—This research was supported by National Institutes of Health (NIH) Grant 1 R21 DA015211-01A1 to A.B.M. and NIH Grant R01 MH59196 to W.T.M. The authors thank Scott Lauritzen for his help running subjects.

REFERENCES

Carver, C.S., & Scheier, M.F. (1998). *On the self-regulation of behavior.* New York: Cambridge University Press.

Crowe, E., & Higgins, E.T. (1997). Regulatory focus and strategic inclinations: Promotion and prevention in decision-making. *Organizational Behavior and Human Decision Processes*, *69*, 117–132.

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

Higgins, E.T. (1997). Beyond pleasure and pain. *American Psychologist*, *52*, 1280–1300.

Higgins, E.T. (2000). Making a good decision: Value from fit. *American Psychologist*, *55*, 1217–1230.

Higgins, E.T., Chen Idson, L., Freitas, A.L., Spiegel, S., & Molden, D.C. (2003). Transfer of value from fit. *Journal of Personality and Social Psychology*, *84*, 1140–1153.

Kruglanski, A.W., Shah, J.Y., Fishbach, A., Friedman, R., Chun, W.Y., & Sleeth-Keppler, D. (2002). A theory of goal systems. *Advances in Experimental Social Psychology*, *34*, 331–378.

Lewin, K. (1935). *A dynamic theory of personality.* New York: McGraw-Hill.

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

Maddox, W.T., & Dodd, J.L. (2001). On the relation between base-rate and cost-benefit learning in simulated medical diagnosis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*, 1367–1384.

Markman, A.B., & Brendl, C.M. (2000). The influence of goals on value and choice. In D.L. Medin (Ed.), *The psychology of learning and motivation* (Vol. 39, pp. 97–129). San Diego, CA: Academic Press.

Shah, J., & Higgins, E.T. (1997). Expectancy * value effects: Regulatory focus as determinant of magnitude and direction. *Journal of Personality and Social Psychology*, *73*, 447–458.

(RECEIVED 6/25/04; REVISION ACCEPTED 2/15/05; FINAL MATERIALS RECEIVED 2/28/05)