Attitude Formation Through Exploration: Valence Asymmetries

Russell H. Fazio
Ohio State University

J. Richard Eiser
University of Sheffield

Natalie J. Shook
Ohio State University

The formation of attitudes toward novel objects was examined as a function of exploratory behavior. An initial experiment, in which participants played a computer game that required them to learn which stimuli, when sampled, produced favorable or unfavorable outcomes, demonstrated learning, attitude formation, and generalization to novel objects. The findings also revealed 2 interesting valence asymmetries: a learning asymmetry involving better learning for negatively valenced than positively valenced objects and a generalization asymmetry involving stronger generalization as a function of negative than of positive attitudes. Findings from 4 experiments led to an explanation of the learning asymmetry in terms of information gain being contingent on approach behavior and related the generalization asymmetry to a negativity bias that weighs resemblance to a known negative more heavily than resemblance to a positive.

Despite the central role that the attitude concept has occupied in social psychology, relatively little attention has focused on the processes involved in attitude formation per se. The field as a whole has attended more to questions regarding attitude change, attitude structure and function, and influences of attitudes on judgments and behavior than it has to attitude formation (Eagly & Chaiken, 1993; J. M. Olson & Zanna, 1993). For example, in discussing future directions in the study of attitudes, Eagly and Chaiken (1993) referred to the field's current "lack of attention to the developmental issue of how attitudes are formed and become strong" as a "serious omission and limitation" (p. 681). In fact, very little research in experimental social psychology has examined the formation of attitudes toward novel objects—ones for which individuals have no relevant a priori knowledge. The most notable exception may be work concerning Zajonc's (1968) mere exposure hypothesis, in which Chinese ideographs and Turkish words often have been used as stimuli.

This is not to say that the question of how attitudes form has been ignored. To the contrary, a number of mechanisms have been implicated as means of attitude formation. These include not only the consequences of mere exposure (for recent theoretical perspectives on the phenomenon, see Winkielman, Schwarz, Fazendeiro, & Reber, 2003; Zajonc, 2001) but also the conditioning of attitudes (for a recent review, see De Houwer, Thomas, & Baeyens, 2001) and the inferring of attitudes from observation of one's own behavior and the conditions under which the behavior occurred (e.g., Bem, 1972; Fazio, 1987). In general, such formation mechanisms are examples of attitude formation on the basis of what Fazio and Zanna (1981) referred to as *direct experience,* which can be contrasted with the development of attitudes through *indirect experience,* that is, on the basis of information that one receives from others about a given attitude object. In the latter case, individuals may form their attitudes vicariously, whether it be through general socialization (e.g., Newcomb, 1943), inferential reasoning about the communicated attributes of the object (e.g., Fishbein, 1963; Fishbein & Middlestadt, 1995; McGuire, 1960), or consideration of the value with which others regard the object (e.g., Heider, 1946, 1958).

What is lacking in the literature, however, is systematic research examining how attitudes toward novel objects develop over time

Russell H. Fazio and Natalie J. Shook, Department of Psychology, Ohio State University; J. Richard Eiser, Psychology Department, University of Sheffield, Sheffield, England.

This research was supported by National Institute of Mental Health (NIMH) Senior Scientist Award MH01646 and NIMH Grant MH38832 to Russell H. Fazio and by Economic and Social Research Council, United Kingdom Grant R000223077 to J. Richard Eiser. Experiments 3 and 4 were part of a master's thesis conducted by Natalie J. Shook under the supervision of Russell H. Fazio. We thank Rennae Feggins for her many contributions to the initial pilot research with which this line of work began, Roxanne Flowers for her assistance with data collection and analysis, and Michael Bailey for developing and testing the software that was used in the experiments. We also thank Ed Hirt and Nira Liberman for their helpful commentary as the research program progressed.

Correspondence concerning this article should be addressed to Russell H. Fazio, Department of Psychology, Ohio State University, 1885 Neil Avenue, Columbus, OH 43210-1222. E-mail: fazio.11@osu.edu

Journal of Personality and Social Psychology, 2004, Vol. 87, No. 3, 293–311

0022-3514/04/$12.00 DOI: 10.1037/0022-3514.87.3.293

as a function of individuals' own exploratory behavior. Certainly, people are sometimes provided with information about novel objects by social agents of one sort or another. Certainly, people are sensitive to covariations that they observe between the presence of a given object and the presence of related positive or negative cues—the essence of associative learning through classical conditioning (De Houwer et al., 2001; M. A. Olson & Fazio, 2001, 2002). However, sometimes the social world is constructed differently. Whether people receive any information about the value of an object can depend on their willingness to sample the object in the first place. So, for example, they can come to anticipate oysters as a culinary delight because friends or food critics recommend them. They can anticipate positivity because oysters tend to be served in smart French restaurants. On the other hand, they might anticipate negativity from a sense that oysters appear slimy and disgusting. Ultimately, however, if people are to learn whether they like or dislike oysters, they need to be willing to sample them and experience for themselves the sensory feedback that oysters produce.

Thus, our primary concern in the present research is with the situation in which people have the opportunity to explore novel objects in their environment—ones for which they have no relevant a priori knowledge—and to develop attitudes toward those objects on the basis of the consequences of this exploratory behavior. The individuals are not basing their opinions on past experiences or beliefs, because the objects are truly novel. Nor are they receiving information from others about the objects or being subjected to any socialization pressures. Instead, they are simply choosing whether to interact with a given object or not. Our focus is on the development of evaluative associations as a consequence of learning the outcomes that accrue from interaction with the object. Our aim is to illuminate and test some fundamental principles regarding the development of such evaluative associations. In other words, our interest is in the dynamic interplay between the decision to interact with an object; the experience of whatever outcome such interaction produces; and, hence, the opportunity to learn the value of the object and update the hypothesis that governed the initial choice.

There is no question that attitudes can develop on the basis of the positivity or negativity of one's experiences with the attitude objects. Classic attitude formation studies have established this very well (e.g., Insko, 1965; Zanna, Kiesler, & Pilkonis, 1970). In addition, the outcomes that are likely to accrue from interaction with the object form a central component of the generally well-accepted expectancy–value model of attitudes (e.g., Fishbein, 1963; for a recent review, see Ajzen, 2001). Nevertheless, very little is known about the associative learning process per se or about how exploratory behavior might affect and be affected by such learning.

Essentially, we are considering people's tendency to categorize the objects in their social world into good and bad. However, very important differences can be identified between classification or category development research as typically conducted by cognitive scientists and attitude formation via exploratory behavior. When cognitive scientists study category development, they typically provide participants with feedback on each and every trial (e.g., Kruschke, 1992; McClelland & Rumelhart, 1988)—for example, the participant is asked, "Is *X* an *A* or a *B*?" He or she responds and then is told whether the response is correct or incorrect. In contrast, when developing attitudes as a consequence of the positive or negative outcomes that objects provide, individuals experience the outcome only if they have chosen to approach the object. Hence, feedback and the opportunity to learn from it are available only if one chooses to approach. There is no feedback, and thus there can be no learning, if one chooses to avoid the object. As shall become evident, this asymmetry between approach and avoidance behavior is one of the fundamental principles that the present research serves to highlight.

Although not previously considered within the context of attitude formation, the distinction between approach and avoidance behavior is of course a classic issue in research on reinforcement learning and instrumental conditioning. In reinforcement learning, the outcomes received by the learner are no longer under the complete control of the experimenter (as in Pavlovian conditioning or passive associative learning) but are contingent on the learner's own behavior. In order to obtain rewards or avoid punishment, the learner must discover which actions produce which outcomes, and this requires trying out a novel response (e.g., pressing a lever, investigating a new arm of a maze)—in short, *exploration* (Sutton, 1992; Sutton & Barto, 1998). Where such exploratory behavior produces reward, it will tend to be repeated, and this tendency will be strengthened by continued positive reinforcement. However, if it produces punishment, the action will not be repeated and a form of avoidance response will be substituted. In either case, however, the decision to perform the initial exploratory behavior rests with the learner.

The BeanFest Game

We suspect that the development of attitudes toward novel objects has not been investigated extensively largely because of the sheer difficulty of overcoming various operational hurdles. Identifying or constructing stimuli with which individuals have no relevant a priori history surely is part of the problem. Equally important, however, is gaining control over the outcomes that interaction with any given object will produce. The present research accomplished these aims by having participants play a computer game that involved their spending time in a virtual world. Because the world that participants were asked to occupy consisted only of beans, we affectionately refer to the game as "BeanFest."

In the variation of BeanFest with which this program of research begins, the participant's goal is to survive in the virtual world in which he or she has been placed. Survival is represented by one's current energy level, which can range from 0 to 100. If one's energy level reaches 0, then one dies. The problem the participant faces is that energy diminishes simply as a function of time. Time is represented in terms of trials. Each trial of the game has a fixed cost; that is, energy is depleted by one unit on each trial. The beans themselves vary in the outcomes they produce. When eaten, some provide energy, whereas others deplete energy. (In most of the experiments to be reported, the values used were $+10$ and -10.) If one is to survive in the BeanFest world, then one must learn which beans to eat and which to avoid.

Visually, the beans vary along two dimensions: shape and number of speckles (see Figure 1 for a few examples). Unbeknownst to the participants, the specific beans that are presented were selected from a population making up the 10×10 matrix shown in Figure 2. The beans, and their associated valences, were selected very carefully so as to avoid the creation of any linear relations. Neither the shape of the beans nor the number of speckles is correlated with valence. Thus, participants cannot be successful by learning a simple linear rule. Instead, they must come to associate different regions of the matrix—that is, different kinds of beans—with different outcomes.

Each trial of the game begins with the presentation of a target bean, centered in the upper half of the computer screen. The lower right quadrant of the screen always displays the participant's energy meter. Both a graphic and a numeric indicator of the current energy level are visible. The participant must decide whether to eat the bean (approach) or not (avoid) and presses one of two keys on a response box to indicate the decision. The lower left quadrant of the computer screen contains three rows of text that provide outcome information after each and every decision. Each trial involves a value of −1 in the first row of text: "ENERGY LOSS VIA TIME: −1." If the participant chooses not to eat on that particular trial, no value is listed for the second row, "EFFECT OF BEAN: __"; and the third row of text, "NET GAIN OR LOSS: __," simply shows a value of −1. Thus, feedback about the value of a given bean is contingent on a decision to sample the bean. If, and only if, the participant chooses to eat, a value of +10 or −10 appears as the effect of the bean, and the net gain or loss is +9 or −11. In all cases, the energy meter is updated after each decision. This information—the bean at the top of the screen and any effects of eating it at the bottom of the screen—remains visible for 5 s, providing the participant with ample study time to associate the outcome of eating with the specific bean. Then, the text values in the feedback box are erased and a new trial begins with the presentation of a new bean.

	Y1	Y2	Y3	Y4	Y5	Y6	Y7	Y8	Y9	Y10
X1	1	1	1			-1	-1	-1		
X2	1	1			-1	-1	-1			
X3						-1				
X4									1	
X5		-1						1	1	1
X6	-1	-1	-1						1	1
X7	-1	-1								
X8					1					
X9				1	1	1			-1	-1
X10			1	1	1			-1	-1	-1

Figure 2. Matrix of beans presented during the learning phase and their values. X = shape, from circular (1) to oval to oblong (10); Y = number of speckles, from 1 to 10; 1 = a bean with a value of +10; −1 = a bean with a value of −10.

Experiment 1

Our first investigation was exploratory in nature. The aim was simply to see if participants were able to develop attitudes that reflected reality, that is, to learn the valence associated with the beans. Because only a subset of the beans from the 10 × 10 matrix were presented during the course of the game, we also were in a position to ask whether any developed attitudes would generalize to novel stimuli, that is, never-before-seen beans, that bore varying degrees of similarity to the presented stimuli. This initial investigation uncovered intriguing phenomena regarding learning and generalization. The findings motivated us to adopt a more traditional hypothesis-testing approach in the subsequent experiments, which were aimed at illuminating and understanding the mechanisms responsible for the effects observed in the original experiment.

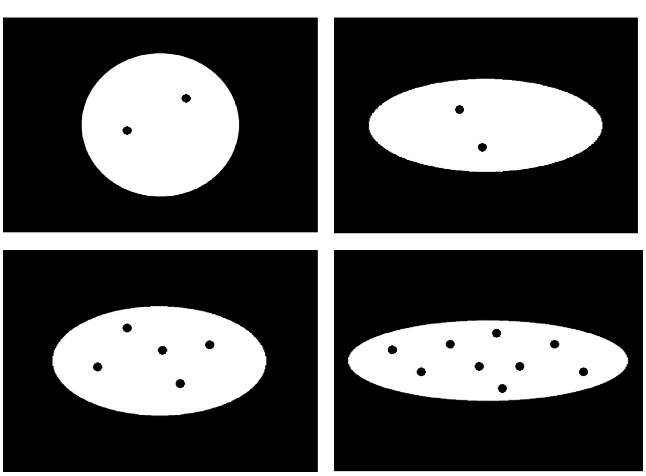

Figure 1. Example beans.

Method

Participants. Eighty-eight Indiana University undergraduates (44 of each gender) participated in the experiment in partial fulfillment of a psychology course requirement.[1]

Procedure. Participants read a lengthy description of the game before beginning. They were informed that their mission was to survive on a planet occupied only by various kinds of beans and that survival was represented by their maintaining an energy level above 0. Because pilot testing had revealed that some people had difficulty discerning the critical dimensions of the stimuli in the number of trials that were to be presented, these instructions explicitly informed the participants that the beans varied in shape ("from circular to oval to oblong") and in the extent to which they were speckled ("marked by anywhere from very few to some to many speckles"). The various parameters of the game, including the fact that beans had values of +10 or −10 and that each trial involved the loss of 1 energy unit simply as a function of time, were noted. In addition, the energy meter and text information that was to be displayed in the lower portion of the screen was illustrated and described. The instruction sheets also informed the participants of a payoff scheme aimed at increasing their motivation to do well. Their energy level at the end of the game would determine the monetary bonus that they would receive, 5 cents/energy unit. Thus, participants could earn as much as $5 if they ended with an energy level of 100. A reminder note attached to the monitor kept this information salient, showing 100 points = $5, 80 points = $4, 60 = $3, and so on.

Finally, the instructions indicated that the experiment would begin with a practice block of six trials. Participants were forewarned that the six beans to be presented would be just a few of the ones that they would see

[1] No effects of gender were observed in this or any of the subsequent experiments.

during the actual experiment but that the beans would have the same value as they would during the experiment. These trials were described as their first opportunity to begin to learn about some of the beans and to familiarize themselves with the energy meter and feedback displays. They were explicitly instructed that given these goals, they should respond "yes" on each of the six practice trials.

After the practice trials, the actual game began, with the participant's energy level meter set to an initial value of 100. The learning, or game, phase of the experiment consisted of three blocks of trials in which each of the 36 beans shown in Figure 2 was presented once. Within each block, the beans were presented randomly for each participant, with the exception of the first 12 trials of the very first block. The order of presentation of these trials, which involved 2 beans from each region of the matrix, was fixed across all participants. This was done to reduce the likelihood of premature "death" simply as a consequence of an unlucky run of negative beans.

This learning phase was followed by a test phase, during which the gamelike aspects of BeanFest no longer operated. No energy meter or feedback was displayed. Instead, a bean was presented, and participants were instructed to classify it as good (one "that you would eat, i.e., one that you believe has beneficial effects on your energy level") or bad (one "that you would not eat, i.e., one that you believe has harmful effects") by pressing one of two buttons on their response pads.[2] Participants responded to each of 52 different beans without receiving any feedback about the accuracy of their responses. Twenty-six of these beans, a minimum of 4 from each region, were beans that had been presented during the learning phase. The other 26 were beans that had not been presented earlier but were included for the purpose of examining the generalization of attitudes.

For half the participants, the matrix of bean values was exactly as displayed in Figure 2. For the other half, the very same beans were presented, but the values associated with the beans were reversed, that is, positive beans from the original matrix were assigned negative values, and negative beans were assigned positive values. This counterbalancing scheme ensured that the findings would not be confounded by the specific location of positive and negative regions in the matrix. This variable produced no theoretically relevant effects in any of the analyses. Hence, all means reported below collapse across the two matrices.

Results

The data are most easily presented in terms of a series of questions about participants' learning and attitude development.

Did participants learn? A very convenient way to assess learning is simply to examine the phi coefficient indexing the strength of the relation between the valence of the beans and a given participant's response to them. Mean phi coefficients across the three blocks of the learning phase were .14, .24, and .35, indicating improving performance over time, $F(2, 174) = 26.96$, $p < .001$. For the test phase of the experiment, the average phi coefficient relating the valence of the beans that had been presented during the game with the good–bad classification was .39, well above chance levels, $t(87) = 12.98$, $p < .001$. Moreover, signs of learning were very consistent across participants; fully 52.3% were characterized by test phase phi coefficients that exceeded the critical value for statistical significance. Thus, the participants showed very clear evidence of having learned.

Did learning vary as a function of valence? Examination of the proportion of beans that were correctly classified during the final test phase revealed an effect of bean valence. Although performance was well above the chance level of .50 for both the positive beans ($M = .61$), $t(87) = 4.91$, $p < .001$, and the negative beans ($M = .77$), $t(87) = 12.99$, $p < .001$, the positive beans were significantly less likely to be classified correctly than the negative beans, $t(87) = 5.45$, $p < .001$. Thus, an intriguing asymmetry in learning was apparent. Participants learned which beans to avoid better than they learned which beans to approach.

How did approach behavior change over time? Early in the game, participants approached the beans fairly indiscriminately (over 60% of the beans in Block 1 were eaten). This makes sense, for it is only through such approach behavior that information about a bean's value is gained. Such indiscriminate approach behavior declined across blocks. When approach behavior was examined as a function of bean valence and block, a significant interaction emerged, $F(2, 174) = 33.79$, $p < .001$. The relevant mean proportions are displayed in Figure 3. As is evident from the graph, all the behavior change was in terms of learning to stay away from negative beans. Approach behavior toward the positive beans was constant across blocks, $F(2, 174) = 1.60$, $p > .20$, whereas approaching negative beans declined sharply over time, $F(2, 174) = 30.84$, $p < .001$.

Did attitudes generalize? Beans that had not been presented during the learning phase of the experiment were included during the test phase to permit an examination of attitudes generalizing to novel stimuli. The novel beans varied in their degree of similarity to the presented beans. We compared the Euclidean distance between a novel bean's location in the 10×10 matrix and its nearest positive neighbor with the distance from its nearest negative neighbor to classify these generalization beans as closer to positive ($n = 10$), equidistant ($n = 6$), or closer to negative ($n = 10$). Responses to the generalization beans (scored as $+1$ and -1 for beans considered good vs. bad, respectively) were strongly affected by this proximity variable, $F(2, 174) = 18.99$, $p < .001$, such that novel beans closer to a negative ($M = -.43$) were more

[2] A few comments about this measure of attitude may prove useful. First, it should be noted that the measure is clearly a subjective measure of evaluation, not a behavioral index thereof. The game had been obviously concluded by the time the test phase began. Both the computer display announcing the end of the game and the experimenter's verbal instructions made that very clear. In addition, during the test phase, the bottom portion of the screen did not display the feedback and energy level information that had been visible during the game phase. Thus, participants were no longer choosing to eat or not eat a bean. Instead, they were judging the beans as either good or bad. Second, this evaluative classification of beans, on the basis of the outcomes that they are expected to produce, corresponds well conceptually with the expectancy–value framework of attitudes (e.g., Fishbein, 1963). Moreover, although our attitude measure is binary in nature, it also corresponds well with often-used semantic differential measures of attitude that anchor the rating scale with such labels as *good–bad* or *beneficial–harmful.* Finally, our choice of wording was very deliberate. It seemed much more appropriate to ask participants if they thought a given bean to be good or bad than to ask them a more affective question, such as whether they liked it or not. Consider, for example, people's attitudinal expressions regarding multivitamin pills. Despite most people's positive attitudes, it is rare for anyone to say that they "like" vitamin pills. Asking participants whether they liked or disliked a given bean would have seemed similarly awkward in light of standard communication practices (although we suspect that participants would have answered any such question simply by drawing an inference from their knowledge of the valence of the outcome produced by a given bean). Nevertheless, the evaluative association to a given bean, as well as the test phase measure, clearly qualify as attitudinal in nature.

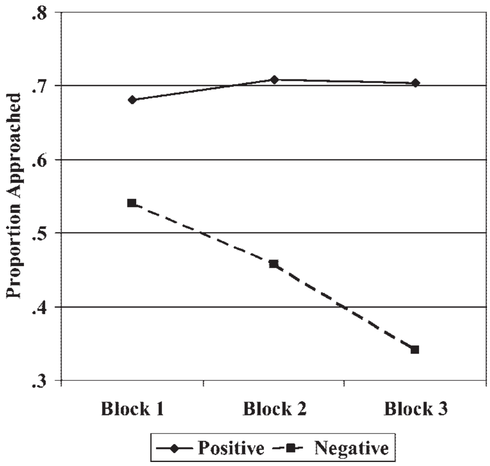

Figure 3. Mean proportion of beans approached as a function of block and valence.

likely to be classified as bad than were equidistant beans ($M = -.22$) or beans closer to a positive ($M = -.08$).

Is there a valence asymmetry in generalization? Intriguingly, the mean responses to the generalization beans implied a difference in the generalization of positive versus negative attitudes. Generalization beans that were equidistant between positive and negative neighbors were likely to be regarded as negative ($M = -.22$), $t(87) = 4.23$, $p < .001$. Moreover, the novel generalization beans as a whole were much more likely to be classified as negative than as positive ($M = -.25$), $t(87) = 7.22$, $p < .001$. In other words, resemblance to a known negative appears to have been weighted more heavily than resemblance to a known positive.

Of course, the generalization data are not at all independent of learning itself. So, any apparent asymmetry with respect to generalization may simply reflect differential learning of the originally presented positive and negative beans. Indeed, participants' mean responses to the novel generalization beans correlated strongly with the degree to which they displayed a learning asymmetry. Greater learning of the negative beans relative to the positives was strongly associated with a greater likelihood of considering the generalization beans to be negative, $r(86) = .71$, $p < .001$. However, the tendency for generalization beans to be considered negative (adjusted $M = -.12$) remained statistically reliable even after controlling for the learning asymmetry, $t(86) = 4.09$, $p < .001$.[3]

Discussion

The findings from this initial investigation proved very interesting. Participants showed clear evidence of having learned the valence of the outcomes to be expected for various types of beans. Moreover, the attitudes that developed then generalized to stimuli that had never been seen before. The degree to which novel beans resembled specific beans to which the participants had been exposed earlier influenced their evaluative responses.

More surprising were the valence asymmetries that emerged. In particular, participants displayed much better learning for the negative beans than for the positive beans. The negative attitudes also exerted a greater influence on generalization. Less similarity to a negatively valued object was required for a novel object to be considered negative than was required for a novel object to be considered positive.

The valence findings are reminiscent of the well-documented asymmetry in impression formation concerning the typically greater impact of negative information about a target on social judgments (see Skowronski & Carlston, 1989; Taylor, 1991). This bias is usually considered to stem from the typically greater diagnosticity of negative information and has been shown to reverse itself in trait domains for which positive information proves more diagnostic (Skowronski & Carlston, 1987). In contrast, the present findings do not involve any such potential difference in the diagnostic value of the positive versus negative outcomes produced by the various beans. The beans are equivalent in the sense of their increasing or decreasing energy level by 10 units. Thus, although related to these past findings, the asymmetry apparent in the associative learning of attitudes seems to be due to a different mechanism than cue diagnosticity.

Possibly more relevant to the current finding is the hypothesis that negative valence generally exerts a stronger influence than positive valence. For example, Cacioppo and his colleagues have proposed a "negativity bias" with respect to attitude formation (e.g., Cacioppo, Gardner, & Berntson, 1997), again suggesting that negative information carries more weight than positive information. Their research has involved the provision of equally extreme positive and negative information about target persons. In contrast, the present work involves not the provision of verbal information ascribing positive or negative information about a target, but positive or negative outcomes that actually occur directly to the participant as a function of interacting with the bean. In general,

[3] As an alternative way of examining the relationship between the generalization asymmetry and learning, we computed an index to represent the extent to which the generalization beans were classified more similarly to each participant's view of the beans that were actually positive or those that were actually negative. This index (A) had a possible range from $-.5$, meaning that the average response to the generalization beans was identical with that given to the bad beans, through $+.5$, meaning that the participant responded to the generalization beans in exactly the same way as to the good beans. The actual computation was $A = .5 - (P - G)/(P - N)$, where P = the mean response to positive beans, N = the mean response to negative beans, and G = the mean response to the generalization beans, all coded as $+1$ for approach and -1 for avoidance. The denominator ($P - N$) simply reflects the extent to which good beans are preferred to bad ones. Seventy-nine of the 88 participants achieved scores greater than zero, that is, displayed the expected preference for the positive beans, and hence were included in the analysis. Mean $A = -.09$, which differed significantly from zero, $t(78) = 2.54$, $p < .02$, indicating that the generalization beans were, on average, viewed more similarly to the way that the participants classified the negative beans than to the way they classified the positive beans. Thus, different analytical approaches converge in suggesting that the valence asymmetry observed with respect to generalization is an independent phenomenon, over and above the valence asymmetry in learning. Although not reported in the interest of brevity, similar analyses were conducted for each of the later experiments. In all cases, consideration of the asymmetry ratio scores (A) corroborated the inferences drawn on the basis of the analyses that examined the generalization asymmetry while controlling statistically for the learning asymmetry. The generalization asymmetry is apparent over and above the effects that the learning asymmetry itself has on generalization.

though, findings in the existing literature point to the greater dominance of negative outcomes. Indeed, lengthy review articles recently have been published concerning a variety of empirical findings suggesting that bad is stronger than good (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Rozin & Royzman, 2001).

Of course, such a general assertion is very difficult to confirm. We ourselves are somewhat skeptical of the generality of the negativity bias principle. In no small part, our hesitation stems from the inherent difficulty of unconfounding valence and its many typical concomitants. We note above Skowronski and Carlston's (1987, 1989) highlighting of the important role played by the diagnosticity of positive versus negative information. Analogously, recent research by Wentura, Rothermund, and Bak (2000) has observed serious limitations to the attention-grabbing power of negative social information that had been demonstrated by earlier research (e.g., Pratto & John, 1991). Whereas trait ascriptions to a target that signaled safety or threat for the perceiver (the target is described as, e.g., generous or aggressive) automatically triggered approach or avoidance tendencies, the same was not true for traits that would be less relevant to the perceiver and, hence, would imply less regarding the risk involved in any social interaction with the target (e.g., the target is described as intelligent or shy; for elaboration of this distinction, see Peeters & Czapinski, 1990). Wentura et al. (2000) suggested that any observed valence asymmetries may stem from stimuli that confound trait valence and relevance. Moreover, Roskos-Ewoldsen and Fazio (1992) did not find the degree to which attitude objects automatically attracted attention in the visual field to vary as a function of valence, although such attention attraction was affected by attitude accessibility. Such findings of equivalence, or even reversal, suggest that the postulated negativity bias may not be as general as it is sometimes portrayed.

Nevertheless, reference to some form of a negativity bias certainly merits serious consideration as a possible mechanism for the learning asymmetry that we observe in the BeanFest paradigm. The discussions of negativity dominance suggest that individuals may attend to and rehearse the events associated with negative outcomes more than those associated with positive outcomes. Such differential attention and rehearsal may have played a critical role in the participants' having learned the negative beans better than they learned the positives.

One might even argue that such greater attention to the negatives was fostered by an entire ensemble of contextual cues. In terms of the distinction offered by Higgins (1998) regarding promotion versus prevention focus, the BeanFest game—at least as instantiated in our initial experiment—is characterized by a very strong prevention-focus flavor. The game concerns survival; it emphasizes hunger and eating behavior; and, because of the fact that energy level cannot exceed 100, it includes a satiation component. Moreover, participants begin with an energy level of 100 and can only suffer "losses" with respect to this initial value. This framing may encourage participants to be concerned with security (i.e., to adopt a prevention focus), which may contribute to their greater attention to negatively valenced beans. Framing the game differently may diminish, or even reverse, the asymmetries that were observed in the earlier work. Instead of illustrating a general principle regarding the associative learning of evaluations, the valence asymmetry observed in the first experiment may be dependent on the learning context. It may be that these asymmetries can be moderated, or even reversed, by contextual factors.

Experiment 2

Examining this potential effect of context constituted one of the two aims of Experiment 2. In pursuit of this aim, we created a blander version of the BeanFest game. The presentation of the game was stripped of all its language regarding hunger, eating, energy, survival, and death. Instead, it focused only on points. We then manipulated the framing of the goal of the game, having it concern either gaining points or avoiding the loss of points (Kahneman & Tversky, 1988). Obviously, the question of interest was whether valence asymmetries involving stronger effects for the negative beans would be equally evident in both conditions or more apparent in the loss framing condition than in the gain framing condition.

The second aim underlying this experiment concerned the very mechanism by which the valence asymmetries occur. Why do people learn the negatives better than the positives? One possibility focuses on the potential negativity bias that we discussed earlier. Participants may simply attend to and rehearse the events associated with negative outcomes more than those associated with positive outcomes. Largely because of its blander language, the points version of the game permitted us to test this possibility in a very direct fashion. Given the format of the survival version of BeanFest, it did not make sense, within the context of the game, to provide people with information about beans that they chose to avoid. We were simulating a world in which one experienced outcomes only when the object was approached. In the points version of the game, we could include a condition in which feedback was provided on each and every trial instead of having feedback be contingent on approach behavior. That is, we could display to participants the effect the bean would have had if they had chosen it and not have that feedback violate the spirit of the game. That is exactly what we did in what we termed the "full feedback" condition. If the learning asymmetry is due to an attentional and rehearsal mechanism, it should be present in this condition as well as in the standard condition, in which feedback is contingent on approach behavior. In both cases, there should be greater attention to negative outcomes. The effect might be even larger as a consequence of directly experiencing any and all observed negative outcomes in the contingent feedback condition, but the negativity bias explanation calls for an effect of valence even in the full feedback condition.

The alternative possibility that we considered stems from the fundamental structural difference between approach and avoidance that we have been highlighting, that is, the idea that one can learn only from approach behavior. If the learning asymmetry were to occur only with contingent feedback, then it must be a consequence of this fundamental structural asymmetry. It was in order to shed light on the viability of an attentional mechanism versus this more structural possibility that the experiment included the manipulation of full versus contingent feedback.

Method

Participants. A total of 216 Indiana University undergraduates participated in the experiment in partial fulfillment of an introductory psychology course requirement. The data from 8 participants were excluded from the analysis because of computer malfunctions during the test phase of the experiment, experimenter error, or in one case the participant's responding with the same button press to each and every trial of the test phase. The final sample consisted of 83 men and 125 women. Participants were assigned randomly to one of the eight cells of the experimental design: Framing (gains vs. loss) × Feedback (contingent vs. full) × Matrix (original vs. reversed). As in the earlier experiment, all results are reported collapsed across the matrix variable.

Procedure. With a few exceptions largely concerning the change in terminology to points instead of energy level, the general procedure was much as in Experiment 1. The game was described as involving beans that varied in their point value. Participants were instructed that they would need to decide whether to accept (approach) the bean presented on any given trial or to reject (avoid) the bean. They were to indicate their responses by pressing either the "yes" or the "no" button on their response pads.

As before, the lower portion of the screen displayed relevant information for the participant. Current point status was presented both numerically and in the form of a "point bar" ranging from 0 to 100. Also included was a line labeled "DECISION: __" for which "YES" or "NO" was displayed once the participant responded. Listed on a second row was "EFFECT OF BEAN: __" on which the consequence of accepting a bean was displayed. As in the earlier experiment, beans had a value of +10 or −10. Unlike the earlier experiment, there was no loss of points simply as a function of time. We eliminated this feature of the game because it was one of the aspects of the original game that highlighted a focus on losses. We needed the new version to be neutral in this respect to permit us to manipulate the framing of the game in terms of gains or losses as symmetrically as possible.

Framing manipulation. The instructions for BeanFest were framed in terms of either gains or losses. Participants in the gains condition read,

> You will begin the game with 0 points. You should try to gain points by making good decisions about which beans to select. Your goal is to reach a point value of 100. In fact, 100 represents winning the game. If your point level ever reaches 100, you will immediately be notified of the fact that you have won. You will then restart the game with 0 points. You should try to win as many times as you can during the course of the experiment.

In contrast, those in the loss framing condition read the following:

> You will begin the game with 100 points. You should try not to lose points by making good decisions about which beans to select. Your goal is to avoid reaching a point value of 0. In fact, 0 represents losing the game. If your point level ever reaches 0, you will immediately be notified of the fact that you have lost. You will then restart the game with 100 points. You should try to avoid losing as much as you can during the course of the experiment.

In either case, that is, if participants in the gains framing condition reached 100 or if those in the losses framing condition reached 0, the game restarted as many times as the participants won or lost. With any restarted games, the point values assigned to the beans remained exactly the same as in the initial game—a fact that the participants' instruction sheets highlighted.

Feedback manipulation. Participants in both conditions were shown the effect of any bean that they selected on the information panel. A value of +10 or −10 appeared as the "EFFECT OF BEAN," and their point value was updated. When participants in the contingent feedback condition rejected a bean, no feedback was presented. The row concerning "EFFECT OF BEAN" remained blank, and the point value remained unchanged. In the full feedback condition, point values also were unaffected by decisions to avoid a bean. However, for these participants, any such "NO" responses were followed by a row of text that read "EFFECT BEAN WOULD HAVE HAD." Thus, although their point value was unaffected by any avoidance decisions, these participants did receive information of the valence of a bean on each and every trial, irrespective of their decisions. In the case of a "YES" response, they were shown the effect the bean had on their point total. In the case of a "NO" response, they were provided with hypothetical information regarding the effect the bean would have had if they had chosen it.

Results

Learning. As in Experiment 1, participants showed evidence of having learned substantially. The average phi coefficient relating responses during the test phase to the actual valence of the beans was .41, well above a chance performance level, $t(207) = 18.03$, $p < .001$. The phi coefficients did not vary as a function of either framing or feedback conditions ($F < 1$).

Examination of the proportion of positive versus negative beans classified correctly during the test phase did reveal evidence of a learning asymmetry. More negative beans ($M = .73$) were classified correctly than positive beans ($M = .68$), $F(1, 204) = 8.51$, $p = .004$. This learning asymmetry was not at all moderated by framing condition ($F < 1$). However, as shown in Figure 4, the asymmetry was moderated by feedback condition, $F(1, 204) = 8.12$, $p = .005$. It was strikingly large in the contingent feedback condition, $F(1, 103) = 14.60$, $p < .001$, but completely absent in the full feedback condition ($F < 1$).

Generalization. Responses to the generalization beans (scored as +1 and −1 for beans considered good vs. bad, respectively) were strongly affected by the valence of the proximal presented beans, $F(2, 408) = 42.76$, $p < .001$. Novel beans closer to a negative ($M = -.30$) were more likely to be classified as bad than were novel beans closer to a positive ($M = .04$); the mean for equidistant beans (−.14) fell between these other two levels of the proximity variable. This effect of proximity was not moderated to a significant degree by either the framing, $F(2, 408) = 2.07$, $p > .12$, or the feedback, $F(2, 408) = 2.00$, $p > .13$, manipulation.

Replicating the earlier experiment, the data provide strong evidence of a generalization asymmetry. The average response to the equidistant beans was significantly more negative than the ex-

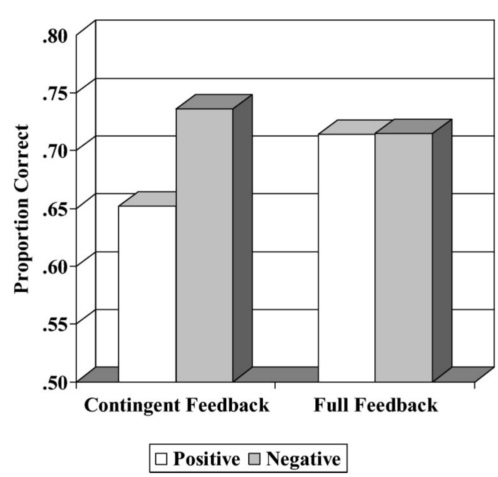

Figure 4. Mean proportion of beans correctly classified as a function of bean valence and feedback contingency.

pected value of zero ($M = -.14$), $t(207) = 4.94$, $p < .001$, as was the average response to the full set of novel beans ($M = -.13$), $t(207) = 6.93$, $p < .001$. As in Experiment 1, however, the learning and generalization asymmetries were strongly related to one another, as evidenced by the correlation involving the difference between test performance for positive versus negative beans and the average response to the novel beans, $r(206) = .63$, $p < .001$. Nevertheless, the tendency for generalization beans to be considered negative (adjusted $M = -.10$) remained statistically reliable even after controlling for the learning asymmetry, $t(206) = 6.45$, $p < .001$. Moreover, and unlike the learning asymmetry itself, this generalization asymmetry was unaffected by the feedback manipulation ($F < 1$).

Discussion

The findings from Experiment 2 are very informative, providing a substantial advance in our understanding of the valence asymmetries that emerged in Experiment 1. Not only were these asymmetries replicated, they were found to be very robust, occurring regardless of whether the BeanFest game was framed in terms of gains or losses. The fact that the learning asymmetry was observed even when the focus of the game concerned advancement, gains, and winning, that is, increasing one's points from the initial value of 0 to the winning value of 100, is very important. It reveals that the learning asymmetry in favor of negatively valenced objects does not stem from negative outcomes simply being more relevant to the concerns for security and prevention that characterized the original survival version of the game. Instead, it appears that the learning asymmetry is not dependent on gain or loss framing.

Equally informative were the results of the feedback manipulation. The learning asymmetry was very strongly affected by feedback contingency. Although the asymmetry was very evident when feedback was contingent on approach behavior, as in Experiment 1, it was absent when feedback was made available irrespective of the behavioral decision. The latter finding is contrary to negativity bias formulations postulating greater attention to and learning from negative outcomes than positive ones. Any such tendency should have produced some evidence of an asymmetry favoring the negative beans in this full feedback condition. During the 5-s time period that followed their decisions on each trial, participants certainly were free to examine and rehearse, as much as they might wish, the specific characteristics of the bean that produced (or would have produced) a negative outcome. Yet, if they engaged in any such greater attention and rehearsal for negative beans than for positive ones, it was insufficient to affect their learning. Their performance during the test phase was unrelated to valence. In effect, they displayed very good and equivalent learning of a two-category classification system.

Although it was necessary for feedback to be contingent on approach behavior for the learning asymmetry to emerge, the generalization asymmetry was not influenced by feedback contingency. In fact, evidence of stronger generalization to novel beans from negative ones than from positive ones was obtained in both the full and the contingent feedback conditions. This finding is very supportive of a specific form of the negative bias postulate. Negative information about an object appears to be weighted more heavily than positive information, just as many theorists have postulated (e.g., Baumeister et al., 2001; Cacioppo et al., 1997; Rozin & Royzman, 2001; Taylor, 1991). More specifically, less resemblance to a known negative is necessary for a novel object to be deemed negative than is required for a novel object to be considered positive.

Despite its clear relevance to the generalization findings, a negativity bias regarding greater attention to negative outcomes than to positive ones does not appear relevant to the learning process. If differential attention and rehearsal as a function of valence is not responsible for the learning asymmetry, then what is? The observation that feedback contingent on approach behavior is necessary for the asymmetry to emerge suggests that the learning asymmetry may itself be an outgrowth of approach behavior. The finding led us to entertain a more structural explanation for the learning asymmetry—the possibility that the differential learning of positively and negatively valued objects might stem from a fundamental asymmetry between approach and avoidance behavior. If information about an object's value is contingent on approaching the object, then one cannot learn from avoiding the object. No information is gained in such a case. The BeanFest game situation is relatively complex, requiring the accurate discrimination of fairly similar stimuli. As a result, participants are bound to develop some initial false beliefs about some of the beans. However, any such hypotheses, even tentative ones, regarding the valence of a given bean are likely to influence the approach–avoidance decision. Beliefs that a bean is positive will promote approach behavior—even if those beliefs are incorrect. Any such false beliefs will be disconfirmed as a result of experiencing the negative outcome that follows the approach behavior. The end result is learning that a negative bean, once thought to be positive, is indeed negative. However, the same is not true of false beliefs that a bean is negative. Any such beliefs encourage avoidance, and if a truly positive bean is avoided, one never learns of its positive value. Instead, the misconception that it is negative persists uncorrected. Thus, the learning asymmetry may itself stem from this fundamental structural difference between approach and avoidance behavior.

The next few experiments in this series were intended to test this structural explanation directly. Experiments 3 and 4 focused explicitly on the postulated mediating role of approach behavior in producing the learning asymmetry. The goal was to manipulate approach behavior experimentally to determine whether the extent of approach behavior does indeed determine the extent of any asymmetry in the learning of positives and negatives. In Experiment 3, promotion- and prevention-focused primes were used to encourage approach or avoidance behavior. In Experiment 4, parameters of the BeanFest game itself were manipulated to influence the likelihood of approach behavior.

Experiment 3

In Experiment 3, two conditions—one in which approach behavior was encouraged via promotion focus and one in which it was discouraged via prevention focus—were compared (Higgins, 1997). *Promotion focus* centers more on gains and the attainment of positive outcomes and produces behavioral tendencies involving approach strategies. Thus, promotion focus is more open to risk taking. On the other hand, *prevention focus* is more sensitive to losses and the avoidance of negative outcomes and tends to produce more conservative strategies geared toward safety and security (Crowe & Higgins, 1997; Foerster, Higgins, & Bianco, 2003).

The experiment used a two-pronged manipulation of promotion and prevention foci. First, the points version of the BeanFest game was framed as concerning gains or losses, just as in Experiment 2. Although this differential framing had no effects on learning in Experiment 2, it was included as part of the present experimental manipulation in order to set the stage appropriately for the second portion of the manipulation. What we added was an attempt to prime approach or avoidance behavior via the use of prevention- or promotion-focused mazes (Friedman & Foerster, 2001). In a series of experiments aimed at examining the influence of prevention and promotion foci on creativity, Friedman and Foerster (2001) found that promotion-focused mazes produced riskier and more novel responses than prevention-focused mazes on a word-fragment completion task. In our experiment, these mazes served as primes of approach or avoidance strategies. The promotion-focused mazes depicted a mouse, squirrel, or monkey that had to find its way through a maze to reach some food (cheese, acorn, and banana, respectively). So the animals were approaching a positive goal. This was intended to prime participants with an approach strategy. In the prevention-focused mazes, the animals' motivation to get through the maze involved the avoidance of a threat (a hawk, a fox, and a lion, respectively). These mazes were intended to produce an avoidance strategy. Participants completed one maze immediately before playing the BeanFest game and completed two more during the breaks between blocks of the game.

Thus, two conditions were created: a promotion focus condition, in which the game was framed as concerning gains, and participants completed mazes that implied approaching a positively valued goal, and a prevention focus condition, in which the game was framed as concerning the avoidance of losses, and participants completed mazes that implied avoiding a threat. We hypothesized that relative to those in the prevention focus condition, participants in the promotion focus condition would approach more beans during the course of the game. According to the structural account of the mechanism responsible for the learning asymmetry, this differential approach behavior should in turn influence the extent to which participants display a learning asymmetry. The asymmetry was expected to be reduced in the promotion focus condition.

Method

Participants. One hundred thirty-five Ohio State University undergraduate students enrolled in introductory psychology courses participated in this experiment for research credit. Participants were assigned randomly to one of the four cells of the experimental design: Focus (promotion vs. prevention) × Matrix (original vs. reversed). As before, all results are reported collapsed across the matrix variable. The data from 5 participants were excluded from the analyses, either because the data were not properly recorded by the computer or because the participant clearly was not engaged in the task as evidenced by a constant pattern of button presses throughout the test phase of the experiment. Sixty men and 70 women made up the final sample.

Procedure. On their arrival, participants were led to believe that the purpose of the study was to examine the effects of task switching on learning, which would be done by having them alternate between playing the BeanFest game, an associative learning task, and completing the mazes, a spatial learning task. As in the earlier experiments, participants read a lengthy description of the BeanFest game. The game was presented just as in Experiment 2, with the game being framed in terms of gains for approximately half the participants and in terms of losses for the other half.

Once the experimenter responded to any questions participants might have had about the BeanFest game instructions, the experimenter instructed the participants to turn over the packet of mazes and complete the first maze. When the participants had completed the maze, they were told that BeanFest would now begin, but when they came to the two rest periods in the game, they would complete the second and third maze. Thus, they would alternate between playing the game and completing the mazes. Participants worked on one maze between the first and second blocks of the BeanFest game and on the final maze between the second and third blocks. Participants assigned to the promotion focus condition, for whom the game had been framed in terms of gains, completed the mazes involving approach behavior (the animal finding its way to a desired food). For those in the prevention focus condition, the game had been framed as concerning losses, and they completed the mazes involving avoidance of a threat (the animal escaping from a predator).

The BeanFest game itself operated exactly as in the earlier experiments, with one exception. To obtain the maximal number of observations, the test phase was expanded. All 100 beans from the 10 × 10 matrix were presented, including the 36 beans that had been involved during the game phase and 64 novel generalization beans. Half of the beans of each type were presented in each of two blocks (i.e., 50 trials), with a rest period between the blocks.

Results

Approach behavior during the game. The focus manipulation was expected to influence the percentage of game trials on which the participants approached the beans. These data were analyzed in a 2 (focus) × 2 (matrix) × 3 (block) × 2 (valence) analysis of variance (ANOVA), with the latter two variables as repeated measures.

The analysis revealed a strong main effect of block, $F(1, 125) = 57.19$, $p < .001$, indicating that approach behavior decreased as the game progressed (*M*s = .70, .62, and .56 for Blocks 1–3, respectively). This main effect was qualified by a two-way Block × Valence interaction, $F(1, 125) = 24.56$, $p < .001$, such that the approaching of negative beans (*M*s = .65, .53, and .43) decreased, whereas the approaching of positive beans remained fairly constant (*M*s = .75, .72, and .68) as the game progressed. This interaction merely indicates that participants were attending to the game and appeared to be learning.

The predicted main effect of focus also was evident, $F(1, 126) = 28.15$, $p < .001$. Participants in the promotion focus condition ($M = .69$) approached significantly more than those in the prevention focus condition ($M = .56$). A Focus × Valence interaction, $F(1, 125) = 4.42$, $p < .05$, indicated that this more extensive approach behavior on the part of promotion-focused participants (*M*s = .76 and .61 for positive and negative beans, respectively), relative to the prevention-focused participants (*M*s = .67 and .46), was especially true for negative beans. Thus, approach behavior during the game was affected by the experimental manipulation.

Learning. As in the earlier experiments, learning was examined by considering performance during the test phase of the experiment. The average phi coefficient relating responses during the test phase to the actual valence of the beans was .31, well above a chance performance level, $t(129) = 14.24$, $p < .001$. The phi coefficients did not vary as a function of condition ($F < 1$).

Although the participants obviously learned, their learning of positive beans versus negatives did vary. Examination of the proportion of positive versus negative beans classified correctly during the test phase did reveal evidence of a learning asymmetry. More negative beans ($M = .71$) were classified correctly than positive beans ($M = .58$), $F(1, 126) = 38.57$, $p < .001$. However, this main effect was qualified by the predicted Valence × Focus interaction, $F(1, 126) = 4.49$, $p < .05$. The relevant means are presented in Figure 5. Although both groups exhibited the learning asymmetry in favor of negative beans, those in the promotion focus condition, $t(68) = 2.98$, $p < .01$, were characterized by a significantly smaller learning asymmetry than those in the prevention focus condition, $t(60) = 5.83$, $p < .001$.

Mediational analysis. To examine the role that approach behavior was postulated to have on the learning asymmetry, we performed a mediational analysis. First, the average level of approach behavior during the game correlated with the learning asymmetry ($r = -.49$, $p < .001$), such that the more one approached during the game, the smaller the learning asymmetry. Also, average approach correlated with focus ($r = .42$, $p < .001$), indicating that those in the promotion focus condition approached more during the game. Finally, the learning asymmetry significantly correlated with focus ($r = -.19$, $p < .05$), such that participants in the promotion focus condition exhibited a smaller learning asymmetry. These latter two outcomes of the focus manipulation had been apparent in the analyses reported above. However, the mediational analysis revealed that controlling for approach behavior resulted in the disappearance of any relation between focus conditions and the learning asymmetry ($r = .03$, *ns*). A Sobel test (MacKinnon & Dwyer, 1993; MacKinnon, Warsi, & Dwyer, 1995) indicated that the reduction in the magnitude of the relation was statistically significant ($z = 3.97$, $p < .001$). Thus, just as expected by the structural account, the learning asymmetry did depend on approach behavior while playing the game.

Generalization. The generalization data replicated the results observed in the earlier experiments. Responses to the novel beans (scored as +1 and −1 for beans considered good vs. bad, respectively) were strongly affected by the proximity factor, $F(2, 252) = 3.67$, $p = .027$. Novel beans more proximal to a negative ($M = -.23$) were more likely be classified as bad than were novel beans that were equidistant between a positive and a negative ($M = -.17$) or ones that were closer to a positive ($M = -.15$). The analysis also revealed a main effect of the focus variable, $F(1, 126) = 4.06$, $p < .05$, such that the average response to the novel beans was more negative in the prevention focus condition ($M = -.23$) than in the promotion focus condition ($M = -.13$). As these means suggest, the data provide strong evidence of a generalization asymmetry. The average response to the novel beans was significantly more negative than the expected value of zero ($M = -.18$), $t(129) = 7.12$, $p < .001$. The main effect of focus indicates that this asymmetry effect was larger in the prevention focus condition, but it was evident to a statistically reliable degree in both the prevention ($M = -.23$), $t(60) = 6.81$, $p < .001$, and promotion ($M = -.13$), $t(68) = 3.69$, $p < .001$, focus conditions.

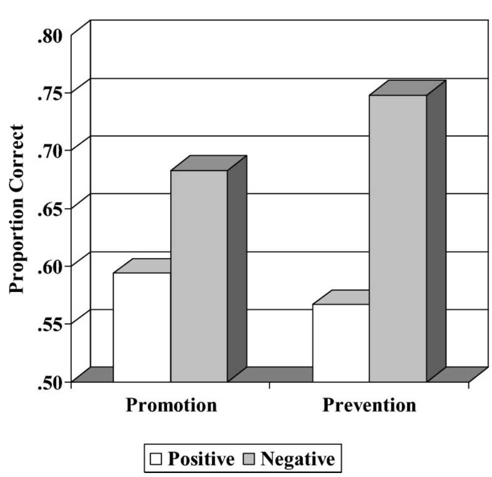

Figure 5. Mean proportion of beans correctly classified as a function of bean valence and promotion versus prevention focus.

As in the earlier experiments, a strong correlation existed between the learning and the generalization asymmetries, $r(129) = .76$, $p < .001$. However, as before, taking the learning asymmetry into account when considering the generalization asymmetry continued to yield clear evidence for the latter effect. The tendency for generalization beans to be considered negative (adjusted $M = -.11$) remained statistically reliable even after controlling for the learning asymmetry, $t(129) = 4.59$, $p < .001$. However, the effect of focus condition was no longer significant after controlling for learning, $t(128) = 1.41$, $p = .16$. The generalization asymmetry was evident in each condition to a statistically equivalent degree: For prevention (adjusted $M = -.17$), $t(60) = 4.94$, $p < .001$; for promotion (adjusted $M = -.07$), $t(68) = 1.91$, $p = .06$.

Discussion

The results of Experiment 3 confirm the predictions. The amount of approach behavior engaged in by the participants was influenced by the focus manipulation. Framing the game as concerning gains and priming the participants with promotion-focused mazes encouraged them to approach more during the game. This increased approach behavior also led to a reduced learning asymmetry, relative to what was apparent for participants with loss framing and prevention-focused mazes. Finally, the mediational analysis indicated that the effect of focus condition on the learning asymmetry was itself mediated by the differential approach behavior. In other words, participants in the promotion focus condition approached more beans during the game, and this increased approach behavior, in turn, led to a decreased difference in the learning of positive and negative beans.

Evidence of attitude generalization also was obtained. Resemblance to a known positive or negative affected responses to the novel stimuli. As in the previous research, generalization also involved an asymmetry. Less resemblance was required for a novel bean to be declared negative than for it to be declared positive. This greater weighting of resemblance to a negative was affected by the focus manipulation, but only indirectly. Participants in the prevention focus condition did display a stronger tendency to consider novel beans negative, but this seemingly differential generalization asymmetry was apparently due to the differential learning asymmetry produced by the focus manipulation. Having learned more of the negatives than the positives, prevention focus participants were more likely to consider novel beans to be negative. However, no effect of focus was apparent over and above this influence of the learning asymmetry. Experiment 3's major aim was to test the structural explanation for the learning asymmetry. The results provided clear support for the idea that the learning asymmetry ultimately stems from the fact that one does not gain any information about an object when one chooses to avoid it. This empirical support involved our use of the mazes as primes, essentially as agents external to the game that could encourage or discourage risky approach behavior. In the next experiment, we attempted to accomplish the same goal, but with manipulation internal to the BeanFest game itself. In this way, we hoped to provide converging evidence for the important mediating role played by approach behavior.

Experiment 4

In Experiment 4, a parameter of the BeanFest game itself was varied in order to prompt differential levels of approach behavior. The values assigned to the various beans were altered, such that beans varied not only in valence but also in value extremity. The matrices were created so that all positive beans had a point value of +10 and all negative beans had a point value of −2 in the extreme positive condition, and all positive beans had a point value of +2 and negative beans had a point value of −10 in the extreme negative condition. It was expected that participants would be more likely to engage in approach behavior during the course of the game in the former condition, that is, if all positive beans were extreme and mild beans negative. Given the bean values, the risk assumed by approaching a bean that was possibly negative would be small and the potential reward would be large. If, on the other hand, all positives were mild and all negatives were extreme, then the risks would be large and the potential rewards small. Thus, participants were expected to pursue less approach behavior in the extreme negative condition than in the extreme positive condition.

Method

Participants. One hundred forty-six Ohio State University students enrolled in introductory psychology courses participated in this experiment. Participants were assigned randomly to one of the four cells of the Extremity (extreme positive vs. extreme negative) × Matrix (original vs. reversed) experimental design. Data from 9 participants had to be excluded from the analyses either for technical reasons ($n = 7$) or because the participant was not engaged in the task ($n = 2$). The final sample consisted of 70 men and 67 women.

Procedure. With three exceptions, the procedure for Experiment 4 was the same as in Experiment 3. First, participants did not complete mazes; the experiment was presented as being concerned solely with associative learning. Second, the game description to which all participants were exposed attempted to frame the game in a balanced manner that did not focus exclusively on either gains or losses. That is, instead of focusing solely on either increasing points or avoiding the loss of points, this more balanced presentation made note of choosing "beans in such a way so as to gain points and avoid losing points" and defined both winning and losing ("reaching 100 represents winning the game, and reaching 0 represents losing the game"). In addition, participants began the game as well as any restarted games following a win or loss with 50 points rather than 0 or 100. Finally, instead of the beans having the values of either +10 or −10, the values of the beans during the game phase of the experiment varied by condition: +10 and −2 in the extreme positive condition and −10 and +2 in the extreme negative condition. These values were not explicitly noted in the game description that the participants read, but they were in operation during the six practice trials with which the game began and, hence, quickly became apparent to the participants.

Results

Approach behavior during the game. The percentage of game trials on which the participants approached the beans was examined via a 2 (extremity) × 2 (matrix) × 3 (block) × 2 (valence) ANOVA, with the latter two variables as repeated measures. A main effect of block, $F(1, 132) = 42.64$, $p < .001$, emerged, indicating that approach behavior decreased as the game progressed (*M*s = .75, .67, and .63 for Blocks 1–3, respectively). This main effect was qualified by a Block × Valence interaction, $F(1, 132) = 30.12$, $p < .001$, such that approaching negative beans (*M*s = .72, .60, and .52) decreased substantially over time whereas approaching positive beans did not (*M*s = .78, .74, and .74). Most important, however, the analysis revealed the expected main effect of the extremity variable, $F(1, 133) = 17.80$, $p < .001$. Participants in the extreme positive condition ($M = .73$), approached significantly more often than did those in the extreme negative condition ($M = .63$). Thus, whether positive beans or negative beans were the more extreme in value did affect approach behavior in the expected manner.

Learning. Once again, performance during the test phase of the experiment showed evidence of learning, with the average phi coefficient relating a given participant's responses to the actual valence of the beans being well above chance ($M = .32$), $t(136) = 15.19$, $p < .001$. The phi coefficients did not vary as a function of condition ($F < 1$).

An ANOVA on the proportion of positive versus negative beans classified correctly during the test phase revealed a valence main effect, $F(1, 133) = 48.33$, $p < .001$, indicative of the usual learning asymmetry in favor of the negative beans. However, this main effect was qualified by a Valence × Extremity interaction, $F(1, 133) = 10.30$, $p < .01$. Although both groups exhibited the learning asymmetry in favor of negative beans (see Figure 6), those in the extreme positive condition, $t(67) = 2.62$, $p < .02$, displayed a significantly smaller learning asymmetry than those in the extreme negative condition, $t(68) = 7.35$, $p < .001$. That is,

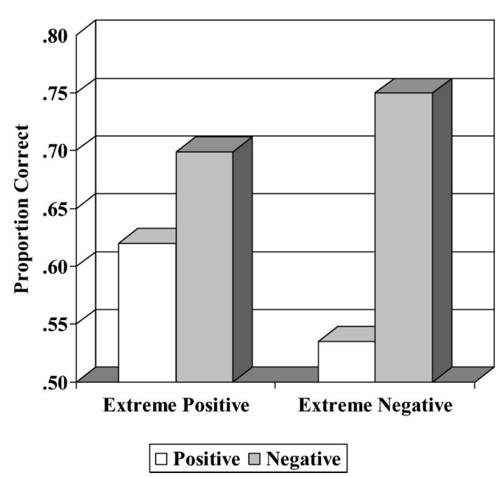

Figure 6. Mean proportion of beans correctly classified as a function of bean valence and matrix value.

those in the condition that approached more during the game phase showed a smaller difference in their learning of the positive and negative beans.

Mediational analysis. As in Experiment 3, a mediational analysis was conducted to examine the role of approach behavior in producing the learning asymmetry. The extent of approach behavior during the game correlated with the learning asymmetry ($r = -.61$, $p < .001$), such that the more often participants approached beans during the game, the smaller the learning asymmetry at test. Also, average approach correlated with the extremity manipulation ($r = .34$, $p < .001$), as did the learning asymmetry itself ($r = -.27$, $p < .01$). However, controlling for the extent of approach behavior reduces the correlation between extremity condition and the learning asymmetry ($r = -.08$), and does so to a statistically significant degree according to a Sobel test ($z = 3.77$, $p < .001$). Thus, the effect of condition on the learning asymmetry was mediated by the extent of approach behavior that the extremity manipulation produced during the game itself.

Generalization. The generalization findings closely resembled what was observed in Experiment 3. Proximity to known positives or negatives significantly influenced responses to the novel beans, $F(2, 266) = 12.82$, $p < .00$ (Ms $= -.26$, $-.17$, and $-.11$ for beans closer to a negative, equidistant, and closer to a positive, respectively). The means also suggested a tendency for more negative responses in the extreme negative condition ($M = -.22$) than in the extreme positive condition ($M = -.14$), but the main effect of the extremity manipulation did not attain a conventional level of significance, $F(1, 133) = 2.62$, $p = .11$. The average response to the novel beans was clearly negative and indicative of a generalization asymmetry across all participants ($M = -.18$), $t(136) = 7.17$, $p < .001$; within the extreme negative condition ($M = -.22$), $t(68) = 6.29$, $p < .001$; and within the extreme positive condition ($M = -.14$), $t(67) = 3.92$, $p < .001$. Moreover, even after controlling for the relation between the learning and the generalization asymmetries, $r(137) = .74$, $p < .001$, the tendency for generalization beans to be considered negative remained significant. Again, this held true across all participants (adjusted $M = -.13$), $t(136) = 5.05$, $p < .001$; within the extreme negative condition (adjusted $M = -.17$), $t(68) = 4.77$, $p < .001$; and within the extreme positive condition (adjusted $M = -.09$), $t(67) = 2.43$, $p < .05$.

Discussion

Experiment 4 provided evidence that a parameter of the game itself—the extremity of positive versus negative beans—could influence approach behavior during the play of the game and the resulting asymmetry in learning. Participants who encountered extreme positive beans and mild negative beans approached significantly more often than did participants for whom the negative beans were extreme and the positive mild. The increased approach behavior decreased the learning asymmetry that was apparent during the test phase, when participants were classifying the beans as good or bad. These findings provide a conceptual replication of the results of Experiment 3, in which differential approach behavior was primed via a task extrinsic to the game. Together, Experiments 3 and 4 provide strong converging evidence for the structural explanation of the learning asymmetry.

The generalization asymmetry also was very evident, as a phenomenon over and above the learning asymmetry. Moreover, and as in Experiment 3, the generalization asymmetry was not affected by the manipulation, at least not once we controlled for the learning asymmetry. Thus, two different manipulations that strongly influenced both approach behavior and the learning asymmetry did not exert any independent influence on the generalization asymmetry. The generalization asymmetry was statistically reliable and statistically equivalent within each condition of the two experiments. We discuss this at greater length in the General Discussion section. It is simply worth highlighting here that the absence of any effects of the manipulations on the generalization asymmetry and its presence in both conditions suggest that the generalization asymmetry is driven by different forces than the learning asymmetry. Despite their substantial covariation, the two forms of valence asymmetries are distinct from one another.

Experiment 5

In confirming predictions derived from the structural explanation for the learning asymmetry, Experiments 3 and 4 have highlighted the importance of approach behavior. More frequent approach decisions reduced the generally superior learning of the negative beans relative to the positive ones. By encouraging participants to be riskier than they otherwise might have been, the promotion focus of Experiment 3 and the extreme positive but mild negative values of Experiment 4 resulted in more sampling of more beans. This more extensive experiencing of the actual outcomes produced by the various beans resulted in more equivalent learning of positive and negative beans. Presumably, if approach behavior had been increased to the point that it approximated the full (noncontingent) feedback condition of Experiment 2, the learning asymmetry would have been not just reduced but actually eliminated.

Apparently, the riskier outlook that the manipulations successfully produced was sufficient to lower the participants' thresholds regarding the probability of a positive outcome deemed necessary to warrant approaching a given bean. As a result, beans whose positivity participants were unsure about were more likely to be sampled. In this way, suspicions that a positive bean might be negative were corrected. The greater willingness to approach the bean allowed the participants to discover the true positive value of the bean. Normally, suspicions that a positive bean is negative would not be subject to such correction, because the suspicion would lead to avoidance and, hence, maintenance of the false belief that the bean is negative. However, Experiments 3 and 4 focused only on generally encouraging (or discouraging) risk-taking behavior, and although the manipulations presumably exerted their influence through this sort of correction process, the specific mechanism was not tested directly. Experiment 5 was conducted to examine the likelihood of such correction and, hence, to provide further evidence illustrating the fundamental principle that learning occurs as a result of approach behavior but cannot follow from avoidance behavior.

To examine this fundamental tenet of the structural explanation, Experiment 5 focused on an intriguing implication of the present conceptual reasoning for the domain of prejudice. What would happen if participants were to develop invalid prejudices? That is, what would happen to learning if participants were provided with initial misleading information? According to our structural explanation, the answer depends on whether the invalid information leads one to avoid or to approach a given type of bean. If one develops an invalid, negative expectancy about a particular kind of bean (a negative prejudice), then one is likely to avoid the bean when it becomes available, which means that its true positive value is likely to go undiscovered. The same is not true of an invalid, positive expectancy. That preconception fosters approach behavior, and by approaching, one will eventually learn the true value of the object. Thus, the asymmetry regarding approach versus avoidance behavior itself suggests that it will prove more difficult to overcome invalid negative prejudices than to overcome invalid positive prejudices.

Experiment 5 used the BeanFest paradigm to simulate a real-world situation involving the social transmission of a mistaken belief and then examined how such culturally shared prejudices influence individuals' decision making and, ultimately, the likelihood of their overcoming invalid positive versus invalid negative prejudices. This was accomplished by having participants develop an initial attitude toward one particular kind of bean on the basis of what Fazio and Zanna (1981) have termed "indirect experience." After reading about the game, but prior to the beginning of play, participants were provided with some critical information. They were told that the focus of the BeanFest experiment was not on individual learning but on how people learn across generations. They were presumably later generation participants. They were also told that a group of first-generation participants in the game had provided written observations and suggestions and that they would have access to one such first-generation report. Under this guise, participants received a suggestion asserting that circular beans with few speckles were either good or bad. The other factor in the design is whether such beans actually are good or bad. The prediction that follows from our structural model is that participants would prove successful at overcoming invalid positive expectancies but not at overcoming invalid negative expectancies. By encouraging avoidance behavior, the invalid negative attitudes based on the indirect experience would never be subjected to an actual test and, hence, would persevere, despite their inaccuracy.

Method

Participants. Indiana University undergraduates participated in the experiment in partial fulfillment of an introductory psychology course requirement. The data from 1 participant, who pressed the same key on each test trial, were excluded from the analyses. The final sample was made up of 31 men and 44 women.

Procedure. The experiment involved the survival version of the BeanFest game, as in Experiment 1. After instructing participants to read the description of the game and responding to any questions, the experimenter informed the participants that primary focus of the experiment was on "how well people can learn across generations" of participants. They were what was being called the third generation of players. The experimenter maintained that in previous semesters, two different groups of participants, referred to as first- and second-generation game players, had participated in this experiment and offered observations about the game and suggestions about how to play to the next generation of players. She then indicated that she would present each participant with a set of folders, each from a different pair of first- and second-generation partners, and they would choose one folder at random.

Each folder contained two forms. The "First Generation Report Form" began with a heading that stated, "Please answer the following questions as best you can. Your answers to these questions will be used in later experiments to aid others in the game you just completed." A handwritten response to the first question, "How successful were you at the game?" provided an opportunity to establish the credibility of the first-generation partner. The response indicated that the player had been "very successful" and "ended up with an energy level near 100." The response to the second question, which asked for suggestions for the next generation of players regarding beans to eat versus avoid, provided the opportunity for the manipulation. Circular beans with few speckles were the focus of the suggestion and were described differently as a function of condition:

> It's very hard to describe which beans were good to eat and which were not. Over time, I just came to know which were which. About the only suggestion I can put into words concerns circle beans with few speckles. For some reason, I focused on them right away. They were good [bad]. I'd say eat them whenever you see them [stay away from them; avoid them whenever you see them]. I can't say anything about other types of beans or about which ones were bad [good] to eat. I didn't focus on any other type of bean.

The "Second Generation Report Form" corroborated the first-generation partner's suggestion.[4] The response to an initial question again indicated the person having succeeded at the game. The second question asked for a rating on a 0–10 scale of the accuracy of the first-generation partner's advice. The endpoint of the scale, 10 (*very accurate*), was circled, and in the space under the question was the handwritten comment: "My partner told me to try [to stay away from] circles with few speckles. I followed the advice. I think my partner's hint really helped me to do well." The third and final question displayed a circled scale point of +4 on a −5 (*very unhelpful*) to +5 (*very helpful*) rating of how helpful the advice had been, followed by the handwritten note, "I would've had a difficult time without the suggestion. It really started me off on the right track."

Thus, participants were led to believe that circles with few speckles were either good or bad. The other factor in the design involved the original or reverse matrix. Hence, circles with few speckles actually were either good (in the original matrix) or bad (in the reverse matrix).

The game proceeded just as in Experiment 1, with two exceptions. The initial practice trials, intended to acquaint the participants with the game displays, did not include a bean from the circles with few speckles region of the matrix. In addition, all five of the beans from this region were presented as part of the test phase of the experiment.

Results

Obviously, our focus is on how well participants learned the value of the circular beans with few speckles. The proportion of correct responses to these beans during the test trials was examined as a function of the initial suggestion the participants were provided and the actual valence of these beans. A 2 (told positive vs. told negative) × 2 (actually positive vs. actually negative) ANOVA revealed main effects of both the suggestion factor, $F(1,$

[4] In a conceptually similar experiment that preceded this (Feggins & Fazio, 1999), no Second Generation Report Form was used, and the first-generation partner's suggestion was worded more mildly. This manipulation proved ineffective. However, approximately half the participants did display evidence of having placed credence in the suggestion that they received, as operationalized by their having behaved consistently with the suggestion the first two times a relevant bean was presented. These "believers" displayed the same effects observed in the present experiment, that is, statistically reliable evidence of overcoming an invalid positive expectation more than an invalid negative expectancy. No such differences were apparent among those characterized as "nonbelievers."

71) = 16.27, $p < .001$, and the actual valence factor, $F(1, 71) = 27.29$, $p < .001$. More important, the predicted interaction was highly significant as well, $F(1, 71) = 18.65$, $p < .001$. The means are presented in Figure 7.

Performance was outstanding in each condition of the experiment except the told negative–actually positive cell, which differed significantly from each of the other three conditions (all *t*s > 5.93, $p < .001$). Forming a positive initial attitude on the basis of socially transmitted information resulted in excellent learning of the beans' actual value. Even when this initial attitude was invalid, its encouraging of approach behavior resulted in individuals discovering and learning that the true value of these beans was negative. Hence, when participants received a positive suggestion, they were equally successful in learning the beans' actual positive or negative value. The same was not true when participants received an invalid negative prejudice regarding the target beans. Participants were unable to overcome an initial prejudice of this sort. The negative attitude promoted avoidance, and the consequence of avoidance is that participants were relatively unlikely to discover that the beans were actually positive.

The differential likelihood of approach and avoidance behavior prompted by the initial suggestion was evident during the course of the game. We examined the proportion of participants who approached, during each of the 15 trials on which a target bean was presented, as a function of trial number, suggestion, and actual valence. The ANOVA revealed a significant three-way interaction, $F(14, 994) = 2.23$, $p = .006$. The data are presented in Figure 8. Avoidance behavior in the condition in which the target beans were actually negative was an interactive function of suggestion and trial, $F(14, 490) = 2.84$, $p < .001$. Those led to believe that these beans were negative avoided them continuously. However, those who had been provided with an invalid positive suggestion showed significant improvement across trials, $F(14, 252) = 8.79$, $p < .001$. The invalid positive expectancy resulted in substantially more participants approaching on the first trial, $t(35) = 3.69$, $p = .001$; the second trial, $t(35) = 3.09$, $p = .004$; and the third trial, $t(35) = 2.56$, $p = .015$. However, by the fourth trial, just as many of those given the invalid positive suggestion had learned to avoid as was true for those who had been provided a valid suggestion. In fact, by the third trial, the two groups of participants who had been told the target beans were positive began to display differential behavior as a function of actual valence, $t(36) = 1.68$, $p < .11$, and this difference proved statistically significant from the fifth trial on. These differential trends across trials produced a strong Trial × Actual Valence interaction among those participants who had received a positive suggestion, $F(14, 504) = 6.91$, $p < .001$.

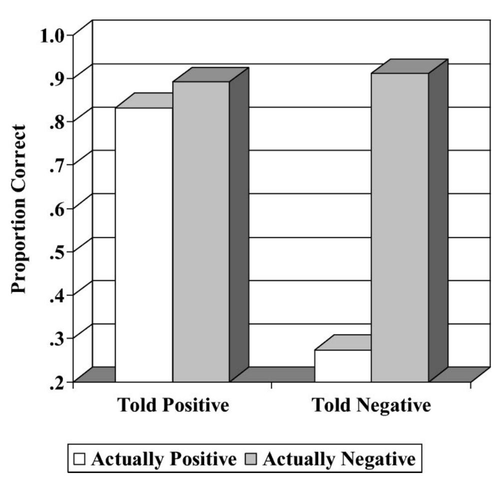

Figure 7. Mean proportion of beans correctly classified as a function of initial suggestion and actual valence.

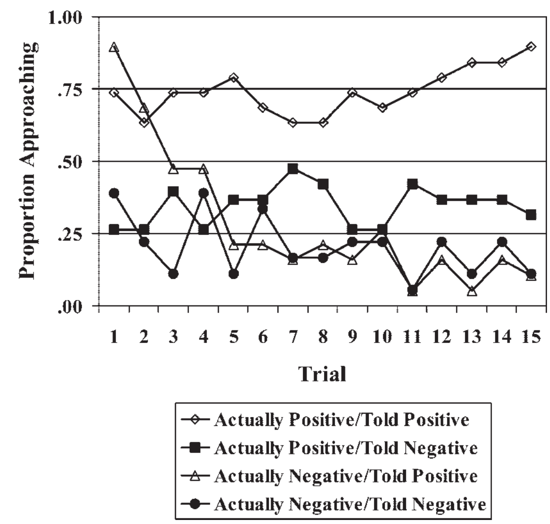

Figure 8. Proportion of participants approaching a target bean as a function of trial, initial suggestion, and actual valence.

In contrast, approach behavior in the condition in which the target beans actually were positive varied significantly only as a function of the suggestion factor, $F(1, 36) = 11.12$, $p = .002$, unmoderated by trial, $F(14, 504) = 1.34$, $p < .20$. Thus, those falsely led to believe that the targets were negative did not succeed in overcoming their initial prejudice. Significantly fewer of these participants approached on the first trial than did those who had been provided an accurate suggestion, $t(36) = 3.23$, $p = .003$, and this lower incidence of approach behavior remained essentially constant as the game progressed. Moreover, participants who received a negative suggestion were largely unaffected by the actual valence of the target bean, as evidenced by the absence of a significant main effect of the actual valence factor, $F(1, 35) = 1.91$, $p < .20$. Given that a minority of the participants who were told the target beans were negative nonetheless approached, over time actual valence did begin to matter somewhat, as indicated by the Trial × Actual Valence interaction, $F(14, 490) = 2.00$, $p = .016$. However, the three-way interaction reported earlier indicates that this occurred to a substantially smaller degree than was true among participants who received a positive suggestion.

To examine the hypothesized role of approach behavior more deeply, we conducted a mediational analysis among the participants for whom the target beans were positive. It is only within this actually positive condition that the suggestion manipulation influenced accuracy during the test phase (see Figure 7; $r = .62$, $p < .001$). According to our reasoning, the effect of the suggestion should be mediated by its impact on exploratory behavior during the game phase. The suggestion variable did correlate with the number of times participants approached the target beans ($r = .49$, $p = .002$), and approach behavior during the game correlated very highly with accuracy at test ($r = .89$, $p < .001$). The relation between the suggestion manipulation and accuracy remained statistically significant after controlling for the frequency of approach behavior ($r = .47$, $p < .002$), implying that some participants relied on the suggestion they had received when responding during the test phase irrespective of their behavior during the game. However, the correlation was reduced markedly in magnitude according to a Sobel test ($z = 3.16$, $p < .002$), indicating that accurately discerning the positive value of the target beans was mediated by the greater sampling promoted by the positive suggestion relative to the negative.[5]

Discussion

The experimental findings provide clear support for the predictions. Invalid positive expectancies encouraged approach behavior and, as a result, were overcome and corrected over time. As a result of their sampling, participants who falsely believed that the target beans were positive eventually learned that they actually were negative. Invalid negative prejudices, on the other hand, were not overcome. Because negative expectancies encourage avoidance behavior, they are less likely to be corrected by actual experience. Their consistent avoidance resulted in participants who falsely believed that the target beans were negative maintaining their inaccurate attitude.

These findings regarding the likelihood of sampling and subsequent correction corroborate the conceptual reasoning regarding the structural account. The inherent difference between approach and avoidance behavior—the former's provision of feedback and the latter's not—plays a major role in the learning of positive versus negative attitudes. Invalid positive attitudes are much more subject to experiential correction than are invalid negative attitudes. In the present experiment, such misconceptions were created experimentally through socially transmitted indirect experience. In the earlier experiments, misconceptions presumably would have developed as a result of the sheer difficulty of discriminating the various types of beans. Regardless of how they develop, false beliefs that an object is negative are less likely to be corrected than are false beliefs that it is positive. As a result, relatively more objects with a truly positive value go undetected, continuing to be avoided because of the invalid negative attitude. This, then, appears to be the mechanism responsible for the robust learning asymmetry that we have observed. Negatives are learned better than positives because misconceptions affect approach–avoidance behavior and, hence, are differentially subject to correction.

As suggested above, these findings also have important implications regarding prejudice. The experiment illustrates the problems that can occur when people form negative attitudes through indirect experience. In fact, the experiment can be viewed as a laboratory simulation of the cultural transmission of stereotypes and prejudice. Provided the socially transmitted information is accepted as credible (see Footnote 4), an initial attitude can develop and can influence subsequent approach–avoidance decisions. Negative attitudes, even if completely erroneous, can prompt avoidance behavior and, consequently, prevent one from discovering the truth and coming to appreciate a truly positive object or person. As a result, one needlessly forgoes opportunities for what would be positive interactions if one were only to give the object or person a chance.

General Discussion

The findings from this series of experiments have uncovered two important valence asymmetries illustrating the relatively greater power of negativity. First, individuals displayed a very robust proclivity to learn the negatively valenced beans better than the positively valenced beans. Second, negative attitudes more strongly influenced generalization than did positive attitudes.

The Learning Asymmetry

Let us consider the learning asymmetry first. The magnitude of this asymmetry was reduced substantially in settings that served to encourage approach behavior (Experiments 3 and 4), but it consistently remained evident. In fact, the learning asymmetry was successfully eliminated only when participants were provided with feedback regarding the valence of each bean they encountered irrespective of their approach–avoidance decisions (Experiment 2). Real-life situations that offer such noncontingent feedback may not be very common. Some teaching situations in which an individual is informed of the correctness of a response by an instructor, a coach, or some other supervisory agent are of this sort. More typically, however, learning about the validity of the outcome that one anticipates on interacting with some object will depend on having chosen to pursue the interaction. It is only through approach behavior that the guiding attitude receives either confirmation or disconfirmation. Avoidance results in no information gain.

It is this fundamental difference between approach and avoidance behavior that lies at the heart of the learning asymmetry. Because they encourage approach, positive attitudes ultimately yield gains in information. Positive attitudes that accurately reflect reality receive validation. Invalid positive attitudes are disconfirmed and hence subject to change when the object's negative value is experienced. As a result of these varying corrective processes, it is rare for an individual who has had repeated behavioral opportunities to wrongly view what is truly a negative object as positive.

Although invalid positive expectancies can be overcome, as illustrated by Experiment 5, invalid negative expectancies tend to persist. They encourage avoidance, which provides no information gain. Thus, even after repeated exploratory behavior, it is much more likely that an individual will hold erroneous attitudes about an object with positive value, that is, wrongly believe that it is negative, than wrongly believe that an object with negative value is positive. The asymmetry stems directly from the essential struc-

[5] As the game progresses, a decision to eat is more and more likely to reflect learning on the basis of prior sampling. However, the mediational outcome revealed by a Sobel test is statistically significant, even if one limits the index of approach behavior to the number of times participants approached during the first three trials ($z = 2.70$, $p < .01$) or the first two trials ($z = 2.29$, $p < .025$). The same is true if one considers only the very first trial ($z = 2.26$, $p < .025$). Thus, mediation is evident as a function of early game behavior, when a decision to approach all the more clearly reflects exploration instead of prior direct experience.

tural difference between approach and avoidance: The availability of feedback is contingent on approach behavior.

The structural mechanism underlying the learning asymmetry received support, not only in Experiment 5, in which expectations about a specific set of beans were manipulated, but also in the earlier experiments involving manipulations that served to encourage or discourage risky approach behavior more generally. The promotion focus of Experiment 3 increased approach behavior during the game and consequently attenuated the learning asymmetry. The same was true in Experiment 4, in which risk taking was encouraged by a point system that placed more extreme value on positive objects than on negative objects. When participants had much to gain and little to lose by approaching a bean about whose value they were uncertain, approach tendencies were generally enhanced, and this, in turn, led to a reduced asymmetry in learning.

Our arguments in favor of a structural account of the learning asymmetry should not be viewed as trivializing the phenomenon in any way. We are not at all implying that this effect is "merely" a product of the experimental situation. On the contrary, the principal structural constraint in our experiments—that feedback is contingent on approach—deliberately reflects the challenge faced in real life by anyone (or any creature) seeking to discover how to navigate safely and/or profitably through their natural and social environments.

In addition, the structural mechanism allows for substantial biases due to prejudice or social influence, as shown by Experiment 5. Insofar as they affect the likelihood of approach or avoidance, prejudicial attitudes are, or are not, updated through actual experience. As such, the present findings emphasize an important aspect of racial or ethnic prejudice: its promotion of avoidance behavior. There's no question that prejudice can involve outright hostility toward minority group members. However, prejudice also can lead individuals to anticipate less enjoyment interacting with a minority group member and, hence, prompt them to avoid entering situations involving such minority persons. Towles-Schwen and Fazio (2003) recently documented such a relation between automatically activated racial attitudes and expressed willingness to enter various situations involving interracial interaction. Moreover, Plant and Devine (2003) found that such expressions of anticipated anxiety about interacting with an African American partner predicted whether individuals actually returned for the interaction. Choosing whether or not to enter a situation is one of the ways in which people exert control over and construct their social worlds. Individuals navigate themselves toward situations and interaction partners that suit their personal preferences and characteristics (Ickes, Snyder, & Garcia, 1997; Snyder & Gangestad, 1982). The present findings illustrate what can happen when the relevant attitudes are unwarrantedly negative. The resulting avoidance behavior markedly reduces the chances of discovering the invalidity of the negative attitudes.

The persistence of beliefs that promote avoidance has an obvious parallel in the animal learning literature concerning conditioned fear. An animal that receives an electric shock in one part of a cage following a warning signal (e.g., light or tone) will quickly learn to move to a different part of the cage. It is crucial to note that such avoidance behavior will persist even though the animal receives no further shocks and so, on the face of it, is no longer being reinforced (Solomon & Wynne, 1954). Such findings have led to a long-standing debate over the extent to which avoidance might be reinforced by reduction in fear or attributed to the absence of new information contradicting the learner's acquired expectancy of punishment contingent on nonavoidance (e.g. Mineka, 1979; Seligman & Johnston, 1973; Tarpy, 1982). In either case, the learned avoidance behavior is very resistant to extinction. Moreover, as long as the organism continues to emit the avoidance response in the presence of the stimulus, the belief that the avoidance behavior prevents a negative event is not open to disconfirmation. Similar notions have been extended to the modeling of risk-averse decision making under uncertainty (March, 1996) and explanations of addictive behavior in humans (Baker, Piper, McCarthy, Majeskie, & Fiore, 2004; Solomon, 1980). The present findings illustrate the relevance of exploratory behavior to attitude formation and persistence.

Negativity Bias and the Generalization Asymmetry

The absence of a learning asymmetry when feedback was not contingent on approach behavior but was instead available following the presentation of each and every bean is inconsistent with the general idea that negativity is more powerful than positivity (Baumeister et al., 2001; Cacioppo et al., 1997; Rozin & Royzman, 2001; Taylor, 1991). The data suggest an important limitation to this general proposition. At least in situations in which the positivity and negativity are equated in terms of their extremity and diagnosticity, we find no evidence supporting the assertion that people learn more on the basis of negative than positive events. The data provide no indication of greater attention and rehearsal following the provision of feedback that a bean had negative value than when it was positive. Any such mechanism should have produced a learning asymmetry within the full feedback condition of Experiment 2.

However, a more specific form of the negativity dominance proposition received very strong and consistent support in our experiments. The attitudes that participants formed toward the beans that were presented during the game clearly generalized to novel beans. The proximity effects that were observed indicate that the similarity of a novel bean to a known positive or negative strongly influenced the valence that the novel bean was thought to assume. Novel beans more similar to a known positive were thought to be more positive than novel beans more similar to a known negative.

Resemblance to a known positive or negative did not affect generalization equally, however. A strong generalization asymmetry was apparent, even after controlling for the learning asymmetry, and even in the full feedback condition of Experiment 2 in which no learning asymmetry occurred. Resemblance to a known negative was weighted more heavily than resemblance to a known positive. So, for example, novel beans that in term of Euclidean distance in our matrix were equidistant between learned positives and learned negatives were themselves likely to be viewed as negative. This generalization asymmetry, then, represents a clear case of the negativity bias that various scientists have discussed (Baumeister et al., 2001; Cacioppo et al., 1997; Rozin & Royzman, 2001; Taylor, 1991). Assuming that they have some relevant information, people in effect weigh the negative features of a novel stimulus more than the positive features.

A Connectionist Approach

The present line of research has involved a companion project concerned with connectionist simulations of the BeanFest learning situation (Eiser, Fazio, Stafford, & Prescott, 2003). Connectionist models use "neural networks" of interconnected nodes or units, roughly analogous to neurons in a brain. Activation is transmitted between the nodes as a function of the strength of the connections between them, with such "connection weights" themselves being modified through "learning." A simple example of connectionist learning involves the network, when presented with a given input, generating an output that is compared with a target value. The discrepancy between the output and the target is then treated as an error message, and this leads to the connections being modified so as to reduce the output–target discrepancy following subsequent inputs. This process is repeated for a specified number of trials or until error falls below some criterion.

The network designed by Eiser et al. (2003) involved a series of inputs representing the shape of the beans and the number of speckles, a hidden layer that received activation from the input units, and an output unit that represented the judgment of the stimulus. The salient feature of the simulations was a comparison between a (standard) full feedback learning algorithm, whereby the connection weights were modified on every trial, and a contingent feedback algorithm, whereby the weights were modified only if the network generated an output above a specified threshold, analogous to eating a bean. When the network was presented with input patterns corresponding to the 36 beans presented in Figure 1 and trained using the full feedback algorithm, it achieved errorless discrimination between the good and bad beans. However, when the contingent feedback algorithm was used, the network still achieved near-perfect identification of the bad beans, but only 70% correct identification of the good beans (*M*s over 10 replications; see Eiser et al., 2003, Study 1, Table 1, p. 1228). In other words, a learning asymmetry was observed under contingent but not full feedback, just as in Experiment 2 reported here.

Like our human participants, the connectionist simulations also showed strong evidence of generalization. When the responses of the network were recorded to all remaining 64 possible patterns of the 10×10 matrix at the end of training, these novel beans tended to be categorized similarly to nearby beans presented during training. Furthermore, under contingent but not full feedback, the average response to these novel beans was significantly to the negative side of neutral. When rescaled to be comparable with the data presented here, for Study 1 of Eiser et al. (2003), the mean in the contingent feedback simulations was $-.17$, $t(9) = 2.85$, $p < .02$. In this same condition, there was also a strong correlation between the responses to the novel beans and the initial learning asymmetry, that is, the difference between the network's output to the good and bad beans ($r = .74$, $p < .02$). However, unlike the human data findings presented here, there was no evidence of a valence asymmetry in generalization after statistically controlling for the learning asymmetry (adjusted $M = -.05$, $t < 1$). Moreover, a ratio score analogous to those reported for the present experiments also achieved a mean that did not differ reliably from zero ($-.02$, $t < 1$), indicating that novel beans were equally likely to be classified as positive or negative beans.

This discrepancy between the simulation and human data underlines the importance of distinguishing between structural and attentional processes. The simulations incorporated the same structural features as the human experiments and the same dependency of learning on exploratory or approach behavior. The learning asymmetry follows directly from these structural features, and this in turn influences generalization such that more novel beans are "judged" to be bad. The question is whether there is any evidence of a generalization asymmetry over and above the effects of the learning asymmetry. In the simulations, the answer appears to be no—a finding that is less surprising when it is remembered that the simulations did not incorporate anything corresponding to the human bias toward attending more to negatively valenced information or treating it as more diagnostic. The fact that the connectionist model does not display a generalization asymmetry but does accord with all other aspects of the human performance suggests that individuals may enter the experiment with an already formed predisposition to weigh resemblance to a known negative more than resemblance to a positive, just as argued by those scientists who have postulated the existence of a negativity bias (Baumeister et al., 2001; Cacioppo et al., 1997; Rozin & Royzman, 2001; Taylor, 1991).

A Final Irony

In concluding, we wish to highlight a rather striking irony regarding attitude development as a function of exploratory behavior. By learning about the effects associated with the various objects, that is, by developing attitudes, people come to experience more positive events than negative events. In all the experiments, participants generally performed well. They maintained a reasonable energy level, accrued points, or maintained their allocated points. By having learned the value of at least some of the beans, participants came to experience more positive events than negative. However, the two valence asymmetries that our research has illuminated—one in learning and one in generalization—lead people to believe that their world consists of more negative objects than positive objects. The learning asymmetry indicates that positive objects are more likely to be mistaken as negative than negative objects are to be mistaken as positive. Given that generalization occurs as a function of the values one has come to associate with the various stimuli that have been presented, this learning asymmetry alone produces a tendency for novel stimuli to be classified as negative instead of positive. However, our analyses indicate that even over and above this consequence of differential learning, resemblance to a known negative will be weighed more heavily in judging a novel object than will resemblance to a known positive. Ironically then, people come to believe that it's a cruel world out there. However, they are quite capable of navigating the world so as to experience more positive than negative outcomes. Such is the functional value of attitudes.

References

Ajzen, I. (2001). Nature and operation of attitudes. *Annual Review of Psychology, 52,* 27–58.

Baker, T. B., Piper, M. E., McCarthy, D. E., Majeskie, M. R., & Fiore, M. C. (2004). Addiction motivation reformulated: An affective processing model of negative reinforcement. *Psychological Review, 111,* 33–51.

Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5,* 323–370.

Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances* *in experimental social psychology* (Vol. 6, pp. 1–62). New York: Academic Press.

Cacioppo, J. T., Gardner, W. L., & Berntson, G. G. (1997). Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review, 1,* 3–25.

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.

De Houwer, J., Thomas, S., & Baeyens, F. (2001). Associative learning of likes and dislikes: A review of 25 years of research on human evaluative conditioning. *Psychological Bulletin, 127,* 853–869.

Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes.* Fort Worth, TX: Harcourt Brace Jovanovich.

Eiser, J. R., Fazio, R. H., Stafford, T., & Prescott, T. J. (2003). Connectionist simulation of attitude learning: Asymmetries in the acquisition of positive and negative evaluations. *Personality and Social Psychology Bulletin, 29,* 1221–1235.

Fazio, R. H. (1987). Self-perception theory: A current perspective. In M. P. Zanna, J. M. Olson, & C. P. Herman (Eds.), *Ontario Symposium on Personality and Social Psychology: Vol. 5. Social influence* (pp. 129–150). Hillsdale, NJ: Erlbaum.

Fazio, R. H., & Zanna, M. P. (1981). Direct experience and attitude–behavior consistency. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 14, pp. 161–202). New York: Academic Press.

Feggins, L. R., & Fazio, R. H. (1999, April). *Forming attitudes through associative learning: The effects of initial prejudices.* Paper presented at the Annual Meeting of the Midwestern Psychological Association, Chicago.

Fishbein, M. (1963). An investigation of the relationships between beliefs about an object and the attitude toward that object. *Human Relations, 16,* 233–240.

Fishbein, M., & Middlestadt, S. E. (1995). Noncognitive effects on attitude formation and change: Fact or artifact. *Journal of Consumer Psychology, 4,* 181–202.

Foerster, J., Higgins, E. T., & Bianco, A. T. (2003). Speed/accuracy decisions in task performance: Built-in trade-off or separate strategic concerns? *Organizational Behavior and Human Decision Processes, 90,* 148–164.

Friedman, R. S., & Foerster, J. (2001). The effects of promotion and prevention cues on creativity. *Journal of Personality and Social Psychology, 81,* 1001–1013.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology, 21,* 107–112.

Heider, F. (1958). *The psychology of interpersonal relations.* New York: Wiley.

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

Higgins, E. T. (1998). Promotion and prevention: Regulatory focus as a motivational principle. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 30, pp. 1–46). New York: Academic Press.

Ickes, W., Snyder, M., & Garcia, S. (1997). Personality influences on the choice of situations. In R. Hogan, J. A. Johnson, & S. Briggs (Eds.), *Handbook of personality psychology* (pp. 165–195). San Diego, CA: Academic Press.

Insko, C. A. (1965). Verbal reinforcement of attitude. *Journal of Personality and Social Psychology, 2,* 621–623.

Kahneman, D., & Tversky, A. (1988). Prospect theory: An analysis of decision under risk. In P. Gärdenfors & N.-E. Sahlin (Eds.), *Decision, probability, and utility: Selected readings* (pp. 183–214). New York: Cambridge University Press.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99,* 22–44.

MacKinnon, D. P., & Dwyer, J. H. (1993). Estimating mediated effects in prevention studies. *Evaluation Review, 17,* 144–158.

MacKinnon, D. P., Warsi, G., & Dwyer, J. H. (1995). A simulation study of mediated effect measures. *Multivariate Behavioral Research, 30,* 41–62.

March, J. G. (1996). Learning to be risk averse. *Psychological Review, 103,* 309–319.

McClelland, J., & Rumelhart, D. (1988). *Explorations in parallel distributed processing.* Cambridge, MA: MIT Press.

McGuire, W. J. (1960). Cognitive consistency and attitude change. *Journal of Abnormal and Social Psychology, 60,* 345–353.

Mineka, S. (1979). The role of fear in theories of avoidance learning, flooding, and extinction. *Psychological Bulletin, 86,* 985–1010.

Newcomb, T. M. (1943). *Personality and social change: Attitude formation in a student community.* New York: Dryden.

Olson, J. M., & Zanna, M. P. (1993). Attitudes and attitude change. *Annual Review of Psychology, 44,* 117–154.

Olson, M. A., & Fazio, R. H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science, 12,* 413–417.

Olson, M. A., & Fazio, R. H. (2002). Implicit acquisition and manifestation of classically conditioned attitudes. *Social Cognition, 20,* 89–103.

Peeters, G., & Czapinski, J. (1990). Positive–negative asymmetry in evaluations: The distinction between affective and informational negativity effects. *European Review of Social Psychology, 1,* 33–60.

Plant, E. A., & Devine, P. G. (2003). The antecedents and implications of interracial anxiety. *Personality and Social Psychology Bulletin, 29,* 790–801.

Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology, 61,* 380–391.

Roskos-Ewoldsen, D. R., & Fazio, R. H. (1992). On the orienting value of attitudes: Attitude accessibility as a determinant of an object's attraction of visual attention. *Journal of Personality and Social Psychology, 63,* 198–211.

Rozin, P., & Royzman, E. B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review, 5,* 296–320.

Seligman, M. E., & Johnston, J. C. (1973). A cognitive theory of avoidance learning. In F. J. McGuigan & D. B. Lumsden (Eds.), *Contemporary approaches to conditioning and learning* (pp. 69–110). Oxford, England: Winston & Sons.

Skowronski, J. J., & Carlston, D. E. (1987). Social judgment and social memory: The role of cue diagnosticity in negativity, positivity, and extremity biases. *Journal of Personality and Social Psychology, 52,* 689–699.

Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin, 105,* 131–142.

Snyder, M., & Gangestad, S. (1982). Choosing social situations: Two investigations of self-monitoring processes. *Journal of Personality and Social Psychology, 43,* 123–135.

Solomon, R. L. (1980). The opponent-process theory of acquired motivation: The cost of pleasure and the benefits of pain. *American Psychologist, 35,* 691–712.

Solomon, R. L., & Wynne, L. C. (1954). Traumatic avoidance learning: The principles of anxiety conservation and partial irreversibility. *Psychological Review, 61,* 353–385.

Sutton, R. S. (1992). Introduction: The challenge of reinforcement learning. *Machine Learning, 8,* 225–227.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction.* Cambridge, MA: MIT Press.

Tarpy, R. M. (1982). *Principles of animal learning and motivation.* Glenview, IL: Scott, Foresman.

Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: The mobilization–minimization hypothesis. *Psychological Bulletin, 110,* 67–85.

Towles-Schwen, T., & Fazio, R. H. (2003). Choosing social situations: The relation between automatically activated racial attitudes and anticipated comfort interacting with African Americans. *Personality and Social Psychology Bulletin, 29,* 170–182.

Wentura, D., Rothermund, K., & Bak, P. (2000). Automatic vigilance: The attention-grabbing power of approach- and avoidance-related social information. *Journal of Personality and Social Psychology, 78,* 1024–1037.

Winkielman, P., Schwarz, N., Fazendeiro, T., & Reber, R. (2003). The hedonic marking of processing fluency: Implications for evaluative judgment. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 189–217). Mahwah, NJ: Erlbaum.

Zajonc, R. B. (1968). Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology Monograph, 9*(2, Pt. 2), 1–27.

Zajonc, R. B. (2001). Mere exposure: A gateway to the subliminal. *Current Directions in Psychological Science, 10,* 224–228.

Zanna, M. P., Kiesler, C. A., & Pilkonis, P. A. (1970). Positive and negative attitudinal affect established by classical conditioning. *Journal of Personality and Social Psychology, 14,* 321–328.

Received February 23, 2004
Revision received April 20, 2004
Accepted April 23, 2004 ■

New Editor Appointed for *Journal of Occupational Health Psychology*

The American Psychological Association announces the appointment of Lois E. Tetrick, PhD, as editor of *Journal of Occupational Health Psychology* for a 5-year term (2006–2010).

As of January 1, 2005, manuscripts should be submitted electronically via the journal's Manuscript Submission Portal (www.apa.org/journals/ocp.html). Authors who are unable to do so should correspond with the editor's office about alternatives:

Lois E. Tetrick, PhD
Incoming Editor, *JOHP*
George Mason University
Department of Psychology, MSN, 3F5
4400 University Drive, Fairfax, VA 22030

Manuscript submission patterns make the precise date of completion of the 2005 volume uncertain. The current editor, Julian Barling, PhD, will receive and consider manuscripts through December 31, 2004. Should the 2005 volume be completed before that date, manuscripts will be redirected to the new editor for consideration in the 2006 volume.