European Journal of Social Psychology
Eur. J. Soc. Psychol. **37**, 1046–1056 (2007)
Published online 25 January 2007 in Wiley InterScience
(www.interscience.wiley.com) **DOI**: 10.1002/ejsp.419

Attitude learning through exploration: Advice and strategy appraisals

J. RICHARD EISER[1]*, NATALIE J. SHOOK[2] AND RUSSELL H. FAZIO[2]
[1]*University of Sheffield, UK*
[2]*Ohio State University, USA*

Abstract

Processes of attitude learning were investigated through a game requiring discrimination between good and bad objects, where feedback about object valence (involving gain or loss) is contingent on approach. Previous research demonstrates a preponderance of false-negative errors, with some good objects ('learning asymmetry') and most novel objects ('generalization asymmetry') being judged as bad, but provides no direct evidence concerning how participants appraise alternative strategies and their own performance. To compare alternative strategies, participants received advice, supposedly from a previous participant, that most objects were bad and should be avoided, or good and should be approached. Learning and generalization asymmetries were replicated, especially among participants who followed the former (risk-averse) advice. Additionally, participants' evaluations of their own game strategy were inversely related to amount of negative feedback (the number of bad objects approached), but unrelated to positive feedback (from good objects approached), pointing to the salience of negative information for self-appraisals. Copyright © 2007 John Wiley & Sons, Ltd.

A surprisingly under-researched question is how attitude formation depends on experience gained through exploration of one's social environment. Much theory views attitudes as object-evaluation associations (Fazio, 1995). However, research on how such associations are learnt has generally manipulated participants' *passive* experience, as when attitude objects are paired with stimuli of pre-established valence (De Houwer, Thomas, & Baeyens, 2001; Olson & Fazio, 2001). Our present research adopts a different focus by considering how attitudes are acquired through experience that depends on individuals' *active* decisions, that is, through experiential or reinforcement learning (e.g. Sutton & Barto, 1998).

Such learning is especially relevant when learners have to sample their (natural or laboratory) environment to identify the sources of good and/or bad outcomes. Expectancies about good and bad outcomes guide approach and avoidance. However, whereas approach behaviour (exploration)

*Correspondence to: J. Richard Eiser, Centre for Research in Social Attitudes, Department of Psychology, University of Sheffield, Sheffield S10 2TP, UK. E-mail: j.r.eiser@shef.ac.uk

Received 4 October 2005
Accepted 1 December 2006

Copyright © 2007 John Wiley & Sons, Ltd.

provides learners with feedback that confirms or contradicts their (positive) expectancies, avoidance behaviour typically does not. In the context of animal learning, this produces the classical finding that conditioned avoidance resists extinction (Solomon & Wynne, 1954). Within social psychology, it may underlie the tendency for negative attitudes and prejudices to persist. By leading to avoidance of contact with the negatively valued object, activity or group, negative attitudes are rarely subjected to test.

Fazio, Eiser and Shook (2004) devised a paradigm—the 'BeanFest' game—for investigating how individuals might develop attitudes on the basis of feedback from interaction with novel objects. Participants are introduced to a computer game involving a 'virtual world' containing a variety of good and bad objects or 'beans'. Survival, or success, in this world requires discriminating good beans (which provide 'energy') from bad beans (which cost energy) on the basis of visual cues (shape and number of speckles). However, such discriminations can be learnt only by 'eating', i.e. sampling the beans. Feedback about the value of a given bean is contingent upon approach behaviour. The main finding from a series of experiments is that participants are poorer at identifying good than bad beans—an effect referred to as the *learning asymmetry*. This is because participants only discover the valence of the beans they actually 'eat'. If they avoid a bean, thinking it to be bad, they never learn whether they were right or wrong and never receive corrective feedback if their negative expectancies were unfounded. As a consequence, a proportion of good beans tend to be avoided and assumed to be bad. A further effect is a *generalization asymmetry*, whereby novel beans are seen as more likely to be bad (see also Shook, Fazio, & Eiser, 2007).

Formally, BeanFest constitutes a case of risky decision-making. Participants cannot discover the valence of different beans without approaching them, and hence risking negative outcomes whenever beans turn out to be bad. Indeed, although Fazio et al. (2004) manipulated a number of factors designed to influence the likelihood of approach behaviour, the only way they found to eliminate the learning asymmetry was through providing 'full feedback' (see Fazio et al., 2004, Experiment 2). In this condition, participants were told the true valence of the beans on every trial of the game regardless of whether they chose to 'eat' or avoid a given bean, thus removing the risk intrinsic to the process of learning through exploration. This fundamental difference between full and contingent feedback is further supported by connectionist simulation of the BeanFest paradigm (Eiser, Fazio, Stafford, & Prescott, 2003).

In the standard conditions where feedback is contingent on approach, however, participants need to take some risks, at least early on, so that they can identify sufficient good beans. Thereafter a risk-averse strategy (involving less approach) is adequate for survival, even though it leads to incomplete learning. Previous work, however, is silent on how participants appraise the risks involved in exploration, how they choose between strategies involving greater or lesser risk, and how they evaluate their game outcomes in the light of the strategies they adopt. We do not even know whether participants actually perceive the game as one in which risk-aversion is the most appropriate or efficient strategy, or even whether approach as such is seen as risky. In short, although the Fazio et al. data show *how* positive and negative attitudes may be acquired, they provide less direct information about *why* participants adopt particular strategies, or how they appraise the outcomes these strategies produce.

We address these issues in the present study through explicitly recommending participants to adopt either a risk-accepting or a risk-averse strategy, observing how far they adhere to this recommendation, and relating such adherence to their learning outcomes *and* to their assessments of such strategies and their own performance. By recommending a specific strategy, we can focus the self-ratings on a known, identifiable approach. Recommendations were in the form of a written note, supposedly from a previous player. This procedure followed that used in Experiment 5 of Fazio et al. (2004), with two important modifications. First, participants received advice that the whole set of beans (rather than a subcategory) were predominantly good, and to be approached, or bad, and to be avoided. Second,

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

whereas Fazio et al. added a note supposedly from a second previous player, corroborating the first, we removed this corroboration so as deliberately to weaken the manipulation and allow more room for individual variation in levels of adherence to such advice, and of evaluation of the recommended strategies.

Note that the information provided about the valence of the beans is misleading in both conditions, since there are actually an equal number of good and bad beans. Nonetheless, previous findings on the learning and generalization asymmetries imply that participants typically infer a preponderance of bad beans. Hence, we predict that the risk-averse strategy should be evaluated more positively overall. Self-appraisals, however, should depend more on whether participants adhere to the advice given. Since taking advice from others could be assumed to reduce risk and uncertainty, adherence should predict more positive self-evaluations overall, but this is likely to depend on the nature of the advice given. In particular, those who adhere more to advice to be risk-averse (compared to those adhering to a risk-accepting strategy) will engage in less sampling and hence receive less direct experience to suggest they are following a wrong strategy. Hence, adherence should predict more positive self-evaluations following risk-averse advice.

'Direct experience', however, divides into two distinct categories: positive feedback from approaching good beans, and negative feedback from approaching bad beans. Other research (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Rozin & Roysman, 2001) indicates that negative information has greater weight than positive. In this context, this implies that participants' appraisals of alternative strategies and their own performance should be more strongly influenced by the amount of negative feedback (resulting from false positive decisions to approach bad beans) than by the amount of positive reinforcement for approaching good beans. This should hold true regardless of the advice provided, since in both conditions negative feedback contradicts positive expectations about the valence of the specific bean approached. This in turn should weaken beliefs that most beans are good, and hence a preference for risk-acceptance, or strengthen beliefs that most beans are bad, and hence a preference for risk-aversion.

Although the extra impact of negative feedback should hold true regardless of the specific strategy that was recommended, other effects of the manipulation on participants' appraisals should be apparent. These, however, should be a function not simply of the advice participants had been given but also how consistently they adhered to such advice. (As noted, the advice manipulation was set at a relatively weak level precisely so as to allow for individual differences in the amount of adherence.) In particular, participants who adhere to a recommendation to be risk-averse will engage in less direct testing of such advice and so continue to believe what they are told. This in turn should lead not only to a greater learning asymmetry but also to more positive assessment of the advice and one's own efficacy. By contrast, those who adhere to a recommendation to be risk-accepting will sample more beans, including some bad ones, leading not only to a lesser learning asymmetry but also to negative feedback that will reduce their confidence in the recommended strategy. This reasoning leads to the prediction of a stronger correlation between adherence and perceptions of efficacy in the risk-averse condition than in the risk-accepting condition.

METHOD

Participants

Ninety Ohio State University undergraduate students (53 females and 37 males) enrolled in introductory psychology courses participated in this experiment in return for course credit. At most,

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

four participants were present for each session. Within any given session, participants were randomly assigned to a condition.

Design

The design of the study was a 2 (Matrix) $\times$ 2 (Advice) between participants factorial. Half the participants were presented with beans where the relationship between valence and appearance was as shown in Figure 1, and half where the original values of the beans were reversed (this factor had no effect). Before starting the game, participants received one of two kinds of advice, supposedly from an earlier player. In the 'eat few' advice condition, participants were told that 'Most of the beans seemed to be bad ones. So if in doubt, don't eat'. In the 'eat lots' advice condition, participants were told 'Most of the beans seemed to be good ones. So if in doubt, eat'.

Procedure

Participants started with an energy level of 50 units and had to prevent this dropping to zero ('death'). Each good bean eaten increased their energy by 10, each bad bean eaten decreased it by 10. They also lost one energy unit per trial, to represent the need to find some food to avoid starvation. Participants were told that the purpose of the experiment was to determine how people learn from one generation to

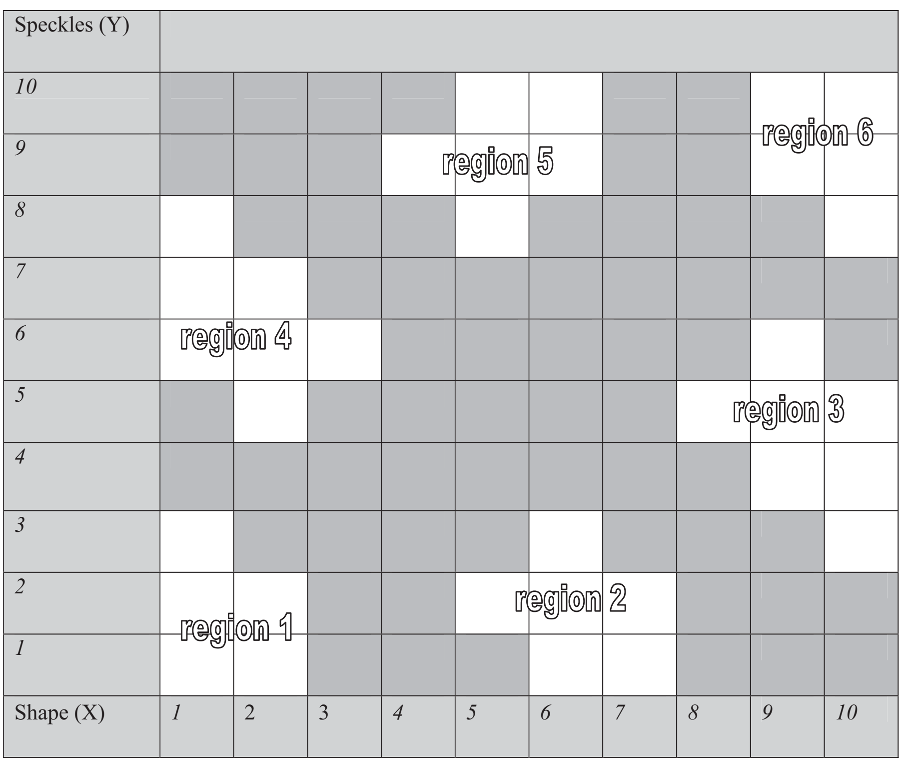

Figure 1. Matrix of beans showing good (light) and bad (dark) regions

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

the next and so they would receive advice from a previous generation player. They then chose a folder from a pile of six containing such advice. In reality, each folder contained the same advice for a given condition.

The game followed the Fazio et al. (2004) procedure, starting with a practice block of six beans (to which participants responded YES throughout) followed by the learning phase, consisting of three blocks of trials in which each of the 36 beans shown in Figure 1 was presented once. This was followed by a test phase, during which no feedback was displayed. Participants were presented, in a random order, with stimuli corresponding to *all* 100 cells of the matrix and were instructed to classify each 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').

Finally, participants completed a short questionnaire in which they rated (a) 'a strategy of eating lots of different beans and avoiding very few beans', (b) 'a strategy of avoiding lots of different beans and eating very few beans' and (c) ' your strategy for playing the game'. Each strategy was rated on a 7-point scale (1 = not at all, 4 = somewhat, 7 = extremely) in terms of 10 adjectives. The first two of these (effective, intelligent) were designed to provide a simple evaluation of the respective strategies and were combined to yield two scores: *advice-efficacy* (average rating for the recommended strategy minus its opposite, i.e. 'eat few' minus 'eat lots' for participants advised to 'eat few' and 'eat lots' minus 'eat few' for those advised to 'eat lots') and *own-efficacy* (average rating of own strategy). The remainder were designed to provide ratings of perceived riskiness or cautiousness, while unconfounding denotative descriptions of riskiness from evaluative connotations of approval or disapproval (see Eiser & Mower White, 1975, for a fuller discussion of the distinction between denotation and value connotation and implications for social judgement). These ratings were combined for each strategy by reverse-scoring the four scales denoting caution (careful, prudent, cowardly, timid) and then subtracting their average from the average of the other four scales denoting riskiness (dangerous, reckless, ambitious, bold). These scores provided, firstly, a manipulation check to see if the 'eat lots' strategy would be described as riskier than the 'eat few' strategy (it was: Ms = 2.35, −1.79, t (88) = 12.49, $p < 0.001$) and, secondly, a measure of how risky participants considered their own strategy to be (*own-riskiness*).

RESULTS

Since there were no significant differences between the responses of females and male participants on relevant dependent variables, analyses are reported for the combined sample without consideration of gender.

Analysis of Game Behaviour and Learning

Various scores were derived to assess aspects of participants' performance (see Table 1 for means for each condition). The proportion of beans approached (i.e. 'eaten') in the first block of 36 trials (*initial approach)* was only marginally lower for those advised to 'eat few' than 'eat lots', t (88) = 1.77, $p = 0.08$. *P good approached* and *P bad approached* represented the proportion of approaches to good and bad beans respectively in blocks 2 and 3 combined. A 2 × 2 (Advice × Valence: good vs. bad beans) ANOVA with repeated measures on the second factor revealed a highly significant effect of Valence, F (1,87) = 86.91, $p < 0.001$, $\eta^2 = 0.50$, with more approach occurring to good than bad beans, but no effect of Advice ($F < 1$).

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

Table 1. Learning as a function of advice

Advice	'Eat few'	'Eat lots'
Initial approach	0.68	0.73
P good approached	0.72	0.76
P bad approached	0.54	0.57
P good-judged good	0.56	0.58
P badjudged good	0.30	0.35
Learning asymmetry	0.14	0.07
Generalization	−0.22	−0.13

From participants' estimates of bean values derived from the test phase after learning was completed (where participants received no feedback), we calculated the proportions of good and bad beans from the original training set of 36 *judged* to be good (*P good-judged good*, *P bad-judged good*). These are considerably lower than the comparable behavioural measures based on approach in blocks 2 and 3 of the learning phase. Overall, participants were only slightly better than chance at identifying good beans, whereas they failed to identify bad beans correctly (i.e. made false-positive errors) on only about one-third of the trials. Again, the only significant effect is Valence, $F\ (1,88) = 86.76$, $p < 0.001$, $\eta^2 = 0.50$. For subsequent analyses, these data were combined into a *learning asymmetry* score, defined as (1—*P bad-judged good)* minus *P good-judged good*.

Generalization was assessed from participants' estimates of the remaining 64 beans not previously presented (coded as bad = −1, good = 1). This score was significantly biased toward the negative, overall $M = -0.17$, $t\ (89) = 6.97$, $p < 0.001$, but Advice only had a marginal effect, $t\ (88) = 1.76$, $p = 0.08$.

The Moderating Role of Adherence to Advice

Although Advice was unreliable in these analyses, it was not expected to influence attitude formation except insofar as it guided exploratory behaviour and hence the feedback participants received. In other words, the effects of advice should be moderated by the extent of participants' adherence to the recommendations, with advice having a stronger influence on individuals who scored higher on adherence. To control for individual variation in levels of adherence, we therefore converted participants' *initial approach* (block 1) scores to standard normal deviates, taking each condition separately; in the 'eat few' condition, we then reversed their sign. The resulting *adherence* score reflected how much any individual, compared to others in the *same* condition, ate fewer beans if advised to eat few, or ate more beans if advised to eat lots.

Hierarchical multiple regressions were conducted to predict learning asymmetry and generalization from Advice (coded 'Eat few' = 0, 'Eat lots' = 1) and Adherence (step 1) and the Advice × Adherence interaction (step 2). For *learning asymmetry*, the expected interaction was significant, $B = -0.13$, $t\ (86) = 2.30$, $p = 0.02$. Simple effects analysis (Aiken & West, 1991), revealed the effect of Advice to be significant at a level of adherence one standard deviation above the mean, $t(86) = 2.76$, $p = 0.007$, but not at all at one standard deviation below the mean, $t < 1$. Thus, just as predicted, the learning asymmetry was lowest—in fact, absent—for those individuals who were advised to 'eat lots' *and* adhered to this advice (see Figure 2).

For *generalization*, the interaction was again significant, $B = 0.16$, $t\ (86) = 2.88$, $p = 0.005$. Advice was significant at an adherence score one standard deviation above the mean, $t\ (86) = 3.39$, $p < 0.001$,

Copyright © 2007 John Wiley & Sons, Ltd.
DOI: 10.1002/ejsp

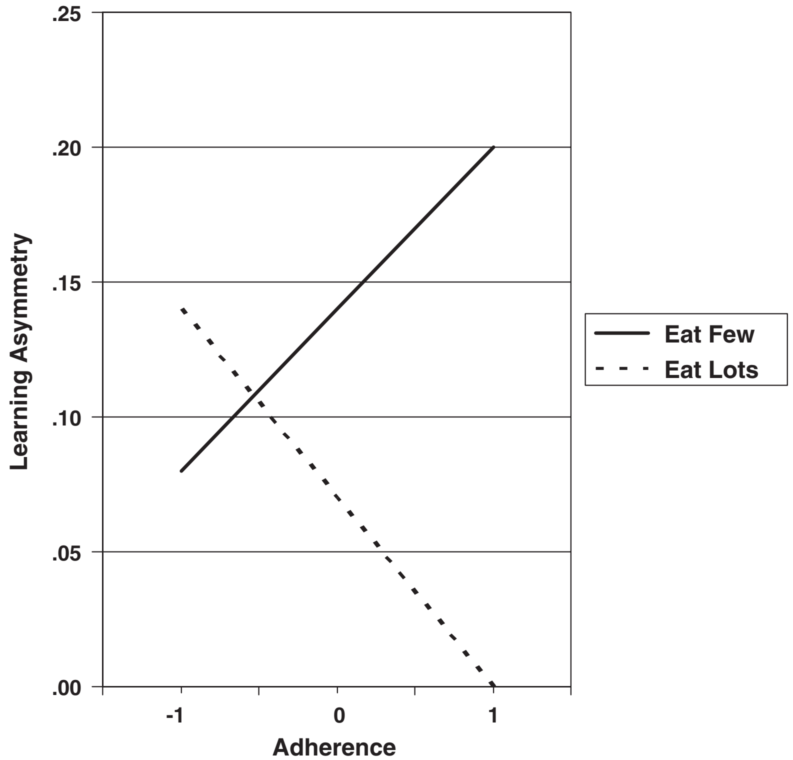

Figure 2. Learning asymmetry as a function of Advice and Adherence

but not at a score one standard deviation below the mean, t (86) = 1.03, $p > 0.30$. As predicted, those participants who adhered to advice to 'eat few' assumed that the novel beans were more likely to be bad, whereas those adhering to advice to 'eat lots' showed no such asymmetry in their predictions (see Figure 3).

Analysis of Questionnaire Ratings

We next considered participants' questionnaire ratings of the different strategies. *Advice-efficacy* was significantly higher for those advised to 'eat few' rather than to 'eat lots', Ms = 1.54, −0.55, t (86) = 3.61, $p < 0.001$, implying an overall preference for a risk-averse strategy. *Own-efficacy* (Ms = 4.13, 4.05, t (86) = 0.09) did not differ as a function of Advice. The *own-riskiness* measure showed that, overall, participants felt their adopted strategy had been somewhat risky, $M = 0.65$, t (87) = 3.72, $p < 0.001$, but these scores were unaffected by Advice (Ms = 0.73, 0.57, t (86) = 0.59).

Regression analyses were then performed to test the effects of Advice, Adherence and the interaction. For *advice-efficacy*, a strong effect of the intercept, t (86) = 3.99, $p < 0.001$, indicated an overall preference for the strategy consistent with the advice initially provided to participants over its opposite. The effect of Advice, $B = -2.09$, t (86) = 3.70, $p < 0.001$, confirmed the overall preference for risk aversion just noted, whereas the effect of Adherence, $B = 0.68$, t (86) = 2.61, $p = 0.01$, indicated that those who adhered more to the advice they had been given rated the recommended strategy more positively (relative to its opposite). There was no interaction, t (86) = 0.60.

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

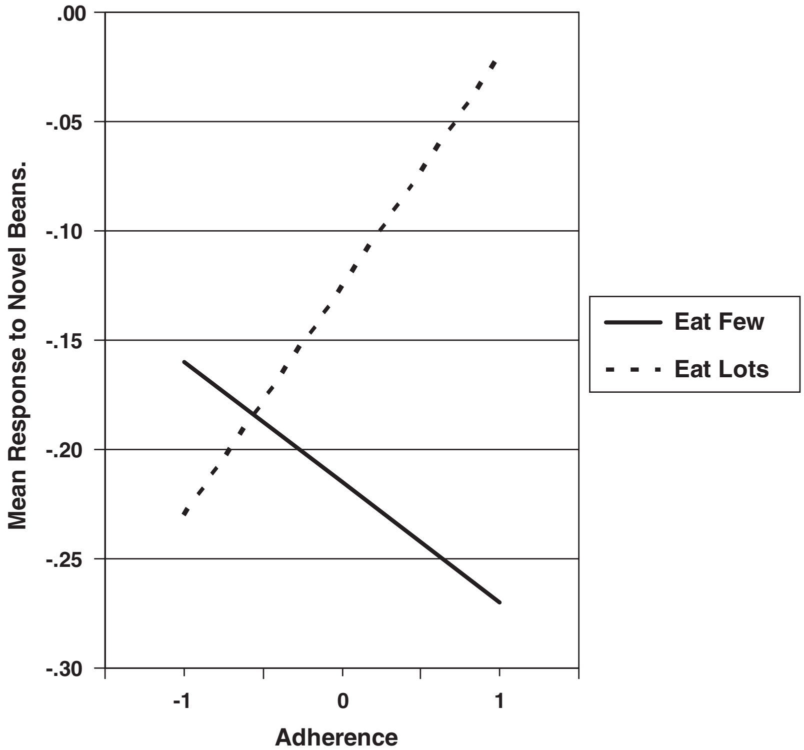

Figure 3. Mean estimates of novel beans as a function of Advice and Adherence

For *own-efficacy*, the effect of Advice was nonsignificant ($t < 1$), but there were significant effects for Adherence, $B = 0.62$, $t\ (84) = 3.94$, $p < 0.001$, and the interaction, $B = -0.99$, $t\ (84) = 2.30$, $p = 0.02$. For *own-riskiness*, there was, likewise, a significant effect for Adherence, $B = -0.54$, $t\ (84) = 3.36$, $p < 0.001$, but not Advice, $t < 1$. The interaction was marginal, $B = 0.77$, $t\ (84) = 1.74$, $p = 0.09$.

Regarding these last two interactions, we predicted that the relationship between adherence and appraisals should be stronger in the 'eat few' than 'eat lots' condition. That this was so can be seen in Table 2 which presents, separately for the two advice conditions, the correlations between advice-efficacy, own-efficacy and own-riskiness, and five scores derived from individuals' game performance: adherence, P good approached, P bad approached, learning asymmetry and generalization, Adherence significantly predicted higher own-efficacy and lower own-riskiness following 'eat few', but not 'eat lots' advice. The values of r differ significantly between the two conditions for both own-efficacy ($z = 3.36$, $p < 0.001$) and own-riskiness ($z = 2.54$, $p = 0.01$). The correlation between adherence and advice-efficacy did not differ between the two conditions ($z = 1.30$).

Note also that the correlation between advice-efficacy and learning is significantly positive in the 'eat few' condition, but marginally negative in the 'eat lots' condition. The difference between the two correlations is highly significant, $z = 3.30$, $p < 0.001$, and conceptually parallels what was observed earlier with respect to our behavioural adherence measure (see Figure 2), but this time in terms of participants' *perceptions* of the efficacy of the advice they received. The more participants viewed the 'eat few' advice as efficacious, the greater the learning asymmetry, whereas the more the 'eat lots' advice was viewed as efficacious, the less the learning asymmetry.

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

Table 2. Correlations between questionnaire and performance measures as a function of advice

	'Eat few'			'Eat lots'		
Advice	Advice-efficacy	Own-efficacy	Own-riskiness	Advice-efficacy	Own-efficacy	Own-riskiness
Adherence	0.35*	0.54***	−0.44**	0.08	−0.14	0.09
P good approached	−0.30*	−0.02	0.22	0.17	−0.06	0.07
P bad approached	−0.40**	−0.58***	0.68***	0.33*	−0.62***	0.26^
Learning asymmetry	0.39**	0.22	−0.33*	−0.30^	0.22	0.06
Generalization	−0.43**	−0.21	0.23	0.11	−0.17	0.11

^$p < 0.10$; *$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

Table 2 reveals another striking result. Within both conditions, own-efficacy is strongly negatively correlated with the proportion of bad beans approached, but unrelated to the proportion of good beans approached. Across all participants, the respective correlations were −0.60 and −0.04, a highly significant difference, $z = 7.10$, $p < 0.001$. Own-riskiness was likewise more strongly associated with P bad approached, but significantly so only in the 'eat few' condition, $z = 2.83$, $p < 0.01$. Thus, it was the *negative* feedback from approaching *bad* beans that was the main influence on assessments of one's own efficacy and, to a lesser extent, own-riskiness. In contrast, participants' successes, i.e. their *positive* outcomes from approaching *good* beans, had no effect on their self-evaluations. This finding of an asymmetry between the effectiveness of positive and negative feedback on self-appraisals confirms our prediction and constitutes the most important novel finding of this study.

DISCUSSION

Our study extends recent research on attitude learning through exploration. We replicate the basic findings of Fazio et al. (2004), by confirming that bad objects are learnt more thoroughly than good objects (learning asymmetry) and that estimates of the valence of novel objects tend to be biased toward the negative (generalization asymmetry). Beyond this, we attempted to influence the attitude learning process by providing advice, supposedly from a previous player, concerning the predominant valence of the stimulus set, and hence the appropriateness of a more risk-averse or risk-accepting search strategy ('eat few' vs. 'eat lots').

This manipulation was adapted from that used in Experiment 5 of Fazio et al. (2004), where participants were told that a specific subset of objects (in fact those in region 1 or 6, see Figure 1) were good or bad, and displayed a predicted tendency to correct false positive advice (that bad beans were good) but not, to the same extent, false negative advice (that good beans were bad). Perhaps because the advice provided in our present study was less specific (applying to the object set as a whole, rather than targeted to those in a single region of the matrix), and almost certainly because it was set at a weaker level (by removing the corroboration from a second player), our findings failed to show reliable main effects of Advice. However, this weakening of the manipulation was deliberate, since our concern was not with simply demonstrating that participants would change their search behaviour if forcibly instructed to do so, but with relating learning performance and appraisals to individual differences in the extent to which participants accepted advice to follow alternative strategies.

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

Our findings show that participants who were advised to 'eat few' *and acted upon such advice* showed greater learning asymmetry and generalization asymmetry effects, as predicted, than those who followed advice to 'eat lots'. These data point to how direct experience from individuals' own exploration may combine with indirect experience, in the form of information from other people about the valence of attitude objects. Whereas previous research has pointed to the greater accessibility of attitudes based on direct experience (Fazio, 1995; Fazio, Chen, McDonel, & Sherman, 1982), our data suggest that indirect experience may have greater influence where it leads individuals to modify their strategies of information search and exploratory behaviour.

Our findings are distinctive in terms of relating participants' subjective appraisals of different search strategies to their behaviour and learning outcomes. Participants' evaluations of the relative efficacy of the strategy recommended to them (advice-efficacy) tended to be significantly positive rather than negative, and more so if they had adhered to it. Their assessments of their own performance were quite veridical, substantially reflecting their direct experiences with the game beans. Interestingly, however, judgements of own-efficacy and own-riskiness were not based on successes and failures to an equivalent degree. The frequency with which negative outcomes were experienced was far more influential than the frequency with which positive outcomes were experienced. Poor decisions to approach a bad bean were weighted more heavily than good decisions to approach a good bean (despite their symmetrical effects on game points lost or gained). Even though, by the end of the game, participants were experiencing more positive than negative outcomes, it was the frequency of those negative outcomes that drove self-assessments. Like the generalization asymmetry observed in this and previous studies, this self-assessment asymmetry appears to be yet another form of a negativity bias in which perceivers are more influenced by negative information than positive information (Baumeister et al., 2001; Rozin & Roysman, 2001).

ACKNOWLEDGEMENTS

This research was supported by Grant R000223077 from the Economic and Social Research Council, United Kingdom to the J. Richard Eiser and by Grant MH38832 from the National Institute of Mental Health to the Russell H. Fazio.

REFERENCES

Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions*. Thousand Oaks, CA: Sage.

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

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.

Eiser, J. R., & Mower White, C. J. (1975). Categorization and congruity in attitudinal judgment. *Journal of Personality and Social Psychology*, *31*, 769–775.

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. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty, & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 247–282). Mahwah, NJ: Lawrence Erlbaum.

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp

Fazio, R. H., Chen, J., McDonel, E. C., & Sherman, S. J. (1982). Attitude accessibility, attitude-behavior consistency, and the strength of the object-evaluation association. *Journal of Experimental Social Psychology*, *18*, 339–357.

Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, *87*, 293–311.

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

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

Shook, N. J., Fazio, R. H., & Eiser, J. R. (2007). Attitude generalization: Similarity, valence and extremity. *Journal of Experimental Social Psychology.* In press

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

Copyright © 2007 John Wiley & Sons, Ltd.

DOI: 10.1002/ejsp