CURRENT RESEARCH IN SOCIAL PSYCHOLOGY

http://www.uiowa.edu/~grpproc/crisp/crisp.html

Volume 11, Number 14
Submitted: July 11, 2005
First Revision: June 20, 2006
Second Revision: July 28, 2006
Accepted: July 28, 2006
Published: July 28, 2006

SELF-SERVING OUTCOME-BIASES IN TRAIT JUDGMENTS ABOUT THE SELF

Scott T. Allison and A. Neville Uhles
University of Richmond

Arlene G. Asuncion
San Jose State University

James K. Beggan
University of Louisville

Diane M. Mackie
University of California, Santa Barbara

ABSTRACT

Sociologists and psychologists have long been interested in how we acquire knowledge about the self. The present study investigated whether people form inferences about the self that are biased by performance outcomes they receive. We hypothesized that trait judgments about the self are outcome-biased when (a) outcomes imply the presence of a positive trait, and (b) outcomes imply traits that are peripheral to one's self-concept. The results of our experiment support these hypotheses. Participants made outcome-biased judgments about the self for positive outcomes but not for negative ones, whereas they were equally outcome-biased in their trait judgments about another target for both positive and negative outcomes. Moreover, the more peripheral the trait dimension implied by the outcome, the stronger the outcome-biased self-judgment. Implications of these results for theory and research on the cognitive and motivational bases of outcome-biased judgments about the self are discussed.

INTRODUCTION

Although a large body of work has shown that people are quick to draw outcome-biased dispositional inferences about others (see Allison, Mackie, & Messick, 1996, for a review), little research has explored how resistant one's self-image is to the outcomes of one's own behavior. The purpose of the present research was to determine whether perceivers display an oversensitivity to their own behavioral outcomes (particularly the positive ones) in the process of reaching conclusions about their own traits.

Historically, social psychologists have been more preoccupied with the processes underlying perceivers' use of behavior in making social judgments than with the processes underlying perceivers' use of the outcomes of behavior (Gilbert, 1989; Gilbert & Malone, 1995; Heider, 1958; Jones, 1979; Jones & Harris, 1967; Ross, 1977). In short, research has primarily focused on the tendency of perceivers to display a behavior bias, whereby a behavior is judged to be a reliable reflection of dispositions even when the behavior has been shaped by circumstance. This tendency of people to focus on behavior to infer the dispositions of other has been called the fundamental attribution error (Ross, 1977) or correspondence bias (Gilbert & Malone, 1995).

In a similar vein, a multitude of social psychological studies have revealed the numerous ways in which human judgment is more influenced by the outcomes of behavior than by the behavior itself. Outcomes bias our judgments about who is responsible for them (Walster, 1966), influence our estimates of how foreseeable they should have been (Fischhoff, 1975), bias our perceptions of how much they were deserved (Lerner & Miller, 1978), affect our beliefs about how likely they are to occur in the future (Kahneman & Tversky, 1973), and influence our evaluations of the behaviors that produce them (Baron & Hershey, 1988). Allison et al. (1996) have suggested that whereas behavior may enjoy greater perceptual significance than situations, the outcomes of behavior enjoy greater psychological significance than behavior itself. Behavior may engulf the field, as Heider (1958) noted, but outcomes appear to engulf behavior. The power of outcomes in molding human judgment is derived from the fact that they represent the final consequences of our actions; they are the irrevocable bottom line.

The idea that outcomes reign supreme in the dispositional inference process was demonstrated by Mackie, Ahn, Asuncion, and Allison (2001; Experiment 1). All participants in the study were informed that a teacher answered 70% of the questions correctly on a teacher competency test. Half the participants learned that the criterion for successful performance on the test was 65% (thus the teacher passed), whereas the other half learned that the criterion was 75% (thus she failed). Despite the fact that the teacher's behavior (in the form of her performance on the test) was identical in both conditions, participants judged her as more competent when she passed the test than when she failed.

Mackie et al. (2001; Experiment 2) replicated this finding using a within-subjects design in which participants were informed of two tests taken by the same teacher. Although she answered 70% of the questions correctly on both tests, her outcomes differed due to a change in the criterion for success (a change from either 65% to 75% or the reverse). Participants made outcome-biased judgments about each performance, inferring that she was competent when she passed one of the exams but incompetent when she failed the other. The within-subjects design

represented the most conservative test of the outcome-bias inasmuch as participants made two different trait inferences about the same target despite being aware of the target's invariant behavior and the shift in the outcome criterion. Moreover, participants made these outcome-biased judgments even though they knew that the change in the criterion was made arbitrarily. The results of numerous studies attest to the robustness of this outcome-bias in perceivers' trait inferences about both group and individual targets (Allison & Messick, 1985; Allison, Worth, & King, 1990; Beggan & Allison, 1993; Mackie & Ahn, 1998; Mackie & Allison, 1987; Mackie, Allison, Worth, & Asuncion, 1992a, 1992b; Mackie, Worth, & Allison, 1990; McHoskey & Miller, 1994; Schroth & Messick, 1994).

Although we may be quick to draw outcome-biased judgments about others, there are several theoretical reasons why we should be resistant to making such judgments about ourselves. First, people have access to a wider range of information about themselves than they do about others, such as their own thoughts, feelings, past behavior, and physiological reactions (Andersen & Ross, 1984). Second, mental representations of the self are far more elaborate, rich, and complex than are representations of others (Linville & Carlston, 1994). Third, people are motivated to process information about the self in a more systematic manner than information that is not relevant to the self (Chaiken et al., 1989).

At the same time, there are compelling theoretical reasons for proposing that at least two conditions are ripe for producing outcome-biased judgments about the self. First, people may form outcome-biased trait judgments about the self when outcomes suggest the possession of a desirable trait. This hypothesis is consistent with a model of outcome-biased inferences formulated by Mackie and her colleagues (Allison et al., 1996; Mackie & Ahn, 1998; Mackie et al., in press). Drawing heavily from theoretical work on motivated processing (e.g., Klein & Kunda, 1993; Kruglanski, 1989; Kunda, 1987), the model proposes that perceivers' motivational goals often determine the output of the inference process, and that these motivations (such as self-enhancement) direct processing towards achieving a desired outcome. For example, a teacher who passes a competency test after earlier failing it may be motivated to view her success as now revealing self-competency, ignoring the fact that her raw score performance on both tests was nearly the same.

The idea that positive outcomes should play a large role in self-concept formation than negative ones is consistent with the large literature attesting to people's tendency to form self-serving attributions about themselves from their own behavior (Zuckerman, 1979). The classic finding is that people tend to make internal attributions for their successful behaviors and external attributions for their failed behaviors (Bradley, 1978; Miller & Ross, 1975; Knight & Vallacher, 1981). Although these studies have examined attributions resulting from behavior rather than from the outcome of behavior (a very crucial distinction we make) the studies do suggest that people's self-judgments may be biased by the valence of any self-generated event, whether that event is a behavior or an outcome. From these theoretical considerations, and from our more recent work on outcome-biased motivational processing (Allison et al., 1996), we hypothesized that people may readily internalize outcomes implying the presence of a positive trait but may discount outcomes implying their possession of a negative trait.

We also propose that perceivers may be outcome-biased in their self-judgments to the extent that the trait (or traits) implied by outcomes are not central to their self-concepts. A trait may be peripheral to one's self-image because one is aschematic on the trait dimension (Markus, 1977), because one is uncertain whether one possesses the trait, because one views the trait as unimportant, or because one lacks knowledge about the meaning and implications of possessing the trait. In short, we propose that self-judgments may be biased by outcomes for the same reason that Bem (1972) suggests they are biased by behavior, namely, when people are uncertain where they stand on the trait dimension in question. Outcome information that implies the presence of these peripheral traits may elicit outcome-biased judgments about the self because the information fails to trigger the motivation to process systematically, or because perceivers are unable to assess adequately whether the outcome truly reflects internal dispositions.

The goal of the present experiment was to test these two theoretical propositions. Specifically, we hypothesize that people are more likely to form outcome-biased self-judgments when outcomes are positive rather than negative, and when outcomes imply a peripheral rather than a central trait.

METHOD

Participants

The participants were 86 undergraduates at the University of Richmond who received five dollars each for their participation in the experiment. Participants responded by telephone to an advertisement placed in the campus electronic newsletter. Participants were invited to the laboratory for an evening appointment and were tested individually. Participants ranged in age from 18 to 24; they included 41 males and 45 females. Of the 86 participants, 68 were Caucasian, 10 were African American, 5 were Asian, and 3 chose not to disclose their race. The strong majority of participants (77) identified themselves as middle to upper middle class Americans, while 9 participants chose not to disclose their economic status.

Valence of Trait Dimension

Participants were presented with an outcome implying the presence of the trait broad-minded or narrow-minded. Prior to the study, 33 undergraduates at the University of Richmond were asked to judge the social desirability of broad and narrow-mindedness. These participants were asked, How favorably would you view someone who was broad-minded (or narrow-minded)? Responses were recorded on a 1 (extremely unfavorably) to 9 (extremely favorably) rating scale. The results showed that participants rated a broad-minded individual more favorably ($M = 6.54$) than they did a narrow-minded individual ($M = 3.87$), $F(1, 30) = 11.26$, $p < .01$.

Manipulation of Target Person

Participants were randomly assigned to either the self or the other condition. Those in the self condition read that they would take a personality test, whereas those in the other condition were informed that they would read about a fellow Richmond student, Chris, who had earlier taken a personality test.

Participants in the self condition first received a questionnaire that requested demographic information such as their age, place of birth, major, and level in school. Participants then received a questionnaire asking them to rate themselves on fifteen personality trait dimensions. One of these fifteen was the target trait dimension of broad-minded versus narrow-minded. This trait rating represented a baseline measure of participants' self-rating on the target trait dimension prior to receiving information about the outcome of their performance on the personality test. The rating made on 9 point rating scale where 1 indicated extremely narrow-minded and 9 indicated extremely broad-minded. Participants rated how important each trait dimension was to them and how certain they were that they knew where they stood on each trait dimension. These ratings were also completed using rating scales ranging from 1 (not at all) to 9 (extremely). In the other target condition, participants received a questionnaire that included information about Chris including Chris' age, place of birth, major, and level in school. These items paralleled those given to participants in the self condition. Participants then made the same fifteen baseline ratings about Chris that participants in the self condition made about themselves.

Manipulation of Outcomes and Performance

After completing the baseline measures, all participants were given a personality test crafted by the experimenters entitled, The Multi Phasic Personality Test, which participants were led to believe assessed people's level of broad and narrow-mindedness. The test required participants to provide open-ended interpretations of numerous geometric shapes and to rate the extent to which those shapes could reasonably represent familiar objects such as houses, butterflies, rakes, leaves, and pillows. Participants were informed that all respondents of the personality test received percentile scores reflecting their performance in relation to other college students in America. Participants were given an example of a respondent who scored in the 58th percentile; this respondent's percentile score indicated that she was more broad-minded than 58 percent of her college student peers nationwide.

Half of the participants read that they (or Chris) had to score in the 65th percentile or higher to be labeled broad-minded, and that scores below this cutoff indicated the trait of narrow-minded. The other half read that they (or Chris) must score in the 75th percentile or higher to be considered broad-minded and that scores below this value indicated the trait of narrow-minded. Participants in the self condition were then given 15 minutes to complete the personality test, whereas those in the other condition read that Chris was given 15 minutes to complete the test.

All participants then learned that they or Chris had scored in the 70th percentile on the test. Thus, half of the participants learned that they (or Chris) had achieved the label of broad-minded (i.e., those who were given the 65th percentile cutoff), while the other half learned that they (or Chris) had achieved the label of narrow-minded (i.e., those who were given the 75th percentile cutoff). Therefore, the experimentally manipulated cutoff score varied the outcome of their personality test, but the performance score itself remained the same in all conditions.

After receiving this feedback about themselves (or Chris), participants were again asked to rate themselves (or Chris) on the trait dimension of narrow-minded versus broad-minded. Participants were also asked to judge how valid a measure of broad-mindedness versus narrow-mindedness they believed the test was on a rating scale ranging from 1 (not at all) to 9 (extremely). Moreover, they were asked to indicate which of the two traits (broad versus

narrow-mindedness) was more social desirable and which of t he two traits they would rather have. Participants were then debriefed and excused from the experiment.

RESULTS

Manipulation Checks

Participants were asked to recall the criterion that must be met to classified as broad-minded (65th or 75th percentile), the actual score received (70th percentile), and whether the results reflected the trait of broad-minded or narrow-minded. Moreover, participants were asked to define what their 70% percentile raw score indicated in terms of their relative position to their peers on the trait dimension of broad and narrow-mindedness. Six of the 86 participants responded incorrectly to one or more of these items, and these participants' data were discarded.

All 80 of the remaining participants indicated that broad-mindedness was a more socially desirable trait than narrow-mindedness, and all 80 checked a box indicating that they would prefer to be broad-minded rather than narrow-minded.

None of our demographic variables (gender, race, or income), when included as between-subjects factors in the analyses below, were associated with any statistically significant effects. Hence, our reporting of our results below do not include these variables.

Trait Inferences

The main question of interest was whether or not participants' trait inferences about themselves or about the other target person were sensitive to the outcome of the personality test. Participants' inferences were submitted to a 2 (target: self, other) x 2 (outcome: broad-minded, narrow-minded) x 2 (time of inference: pre vs. post outcome) ANOVA, with repeated measures on the last factor. The analysis revealed a significant three way interaction, $F(1, 76) = 4.43$, $p < .04$. The means associated with this effect, displayed in Table 1, suggest that our participants were outcome-biased in their self-inferences for positive outcomes but not for negative ones. In contrast, participants were outcome-biased in their inferences about Chris for both the positive and the negative outcome. To further understand the nature of this interaction, we analyzed participants' trait inferences separately for the self and for the other target person.

Table 1. Mean Trait Inferences as a Function of Target, Outcome, and Time.

	Pre Outcome	Post Outcome	Difference
Self, Broad-minded	6.36 (1.16)	7.75 (0.91)	+ 1.39
Self, Narrow-minded	6.50 (1.32)	6.32 (1.37)	- 0.18
Other, Broad-minded	5.88 (1.91)	7.38 (1.13)	+ 1.50
Other, Narrow-minded	5.73 (0.99)	4.17 (1.35)	- 1.56

Note: High numbers indicate stronger inferences of broad-mindedness. Standard deviations are reported in parentheses.

Inferences about the Other Target

The analysis of trait judgments about the other target (Chris) revealed a significant two way interaction between outcome and time of trait rating, $F(1, 38) = 48.46$, $p < .001$, indicating that participants' trait inferences about Chris were influenced by the outcome (both positive and negative) of the personality test. When the outcome revealed Chris to be broad-minded, they rated him as more broad-minded after learning this outcome ($M = 7.38$) than they did initially ($M = 5.88$), $p < .01$. In addition, when they learned that the outcome revealed Chris to be narrow-minded, they rated him as more narrow-minded after learning the outcome ($M = 4.17$) than they did before ($M = 5.73$), $p < .01$. Participants' inferences about Chris were biased by both positive and negative outcomes despite knowing in both outcome conditions that his raw score placed him in the 70th percentile among his peers on the trait dimension of broad-mindedness.

Inferences About the Self

A different pattern of results was found when participants were asked to make trait judgments about themselves. Although participants' trait inferences about themselves were also affected by outcomes, these inferences were outcome-biased only when the outcome implied the presence of a positive trait. The two way interaction between outcome valence and time of trait rating was significant, $F(1, 38) = 5.45$, $p < .03$. This interaction revealed that when participants were informed that their 70th percentile score signified a broad-minded outcome, they rated themselves as more broad-minded after receiving the outcome ($M = 7.75$) than prior to receiving it ($M = 6.38$), $p < .01$. However, when this same 70th percentile score signified that participants were narrow-minded, they were just as likely to rate themselves as moderately broad-minded after receiving this outcome ($M = 6.32$) as they were prior to receiving it ($M = 6.50$), $F < 1$. In short, our participants make trait judgments about themselves that were outcome-biased when the outcome suggested that they possessed a positive trait but not when it suggested their possession of a negative trait.

Judgments of Trait Certainty, Importance, and Validity

Were participants especially likely to draw outcome-biased self-inferences when they were uncertain about where they stood on the trait dimension? The answer appears to be yes. We computed the difference between participants' self-inferences before and after receiving outcome feedback, and these difference scores were found to be negatively correlated with their ratings of certainty, $r = -.26$, $p < .05$. Thus, the more confident participants were about their position on the trait dimension, the less their inferences were influenced by the outcome. In addition, participants' ratings of the importance of the trait dimension were marginally related to their sensitivity to outcomes, $r = -.21$, $p < .07$. The greater the importance that participants assigned to the trait, the smaller their outcome-biased self-judgments.

Did participants rate the personality test as more valid when it produced a positive outcome for them than when it yielded a negative outcome? The answer is yes. Participants judged the test as more valid when the outcome revealed that they were broad-minded ($M = 6.65$) than when it revealed narrow-mindedness ($M = 5.50$), $F(1, 38) = 4.76$, $p < .04$. Moreover, participants' validity ratings were negatively correlated with their inference difference scores, $r = -.45$, $p <$

.01, indicating that the more valid they judged the test to be, the more their inferences were likely to change in response to outcomes. To test whether participants' validity ratings may have mediated the effects of the outcome on their self-inferences, we next conducted an analysis of covariance with participants' validity ratings as the covariate. The results revealed that validity judgments significantly affected inferences, $F(1, 37) = 28.46$, $p < .001$, and that the previously significant effect of the outcome on inferences diminished considerably when controlling for validity ratings, $F(1, 37) = 1.40$, $p = .24$. This finding suggests that participants' perceptions of the test's validity may have assumed a mediating role in producing outcome-biased judgments about the self.

DISCUSSION

The results of our experiment suggest that at least two conditions exist that foster the formation of outcome-biased self-inferences. The first condition is the self-enhancing qualities of the trait dimension; the more the outcome reflects a trait dimension that enhances or protects the self, the greater the outcome-biased self-judgment. The second condition involves the degree to which the trait dimension implicated by the outcome is central to one's self-concept; the less central the dimension, the greater the outcome-biased self-judgment.

When presented with a socially desirable personality outcome, participants were quick to internalize the outcome and claim the positive trait to an even greater degree than they did prior to the outcome. In this instance, participants' processing goals were clearly centered on affirming and reinforcing a favorable self-image. Participants may have made outcome-biased self-inferences for the positive trait because they judged the personality test to have greater validity when it produced a positive outcome than when it produced a negative one. Our data are therefore consistent with the results of Mackie et al. (2001), who found that outcome-biased trait inferences were stronger in magnitude to the extent that participants judged the mechanism for generating the outcome to be diagnostic of the target trait. In making judgments about the self, people appear to use selectively information that will produce the most favorable judgment (even if it is a biased judgment) and they do so to the extent that they perceive the source of the information to be valid and reputable.

The results of our experiment are consistent with recent theoretical work on the processes by which people judge whether a trait is self-descriptive. Klein and Loftus (1993) proposed that the psychological mechanisms implicated in forming self-descriptive trait judgments differ as a function of people's familiarity and experience with the trait dimension in question (see also Klein, Sherman, & Loftus, 1996). When people have little familiarity with a trait dimension and are asked to judge themselves on the dimension, they access any trait relevant experiences that are available in memory and base their judgments on these specific experiences. Klein and Loftus labeled this an exemplar model of arriving at self-judgments. In contrast, when people are highly familiar with a trait dimension and are given the task of judging themselves, they access an abstract summary representation in memory that serves as the basis of their self-judgments. Klein and Loftus call this an abstraction model of reaching judgments about the self.

From these theoretical considerations, we propose that an exemplar based process of arriving at self-judgments is more likely than an abstraction based process to engender outcome-biased

inferences. Klein and Loftus (1993) argue that exemplar based processes consist of the act of accessing individual behaviors exemplifying the trait under consideration. We propose that the outcomes of a behavior serve as a psychologically more meaningful and potent exemplar than the behavior itself. If so, outcomes should play a more pivotal role in influencing self-judgments about traits with which people have little familiarity and experience. In contrast to exemplar based processes, abstraction based processes do not involve the retrieval of specific experiences associated with the trait dimension under consideration. As a result, outcome information should have little impact on self-judgment processes involving highly familiar trait dimensions insofar as these trait judgments implicate previously established summary representations only.

Given that outcome-biased inferences about social groups have been shown to be a useful mechanism for undermining damaging stereotypic conceptions of groups (Mackie et al., 1990, 1992a, 1992b), it seems reasonable that outcome-biased judgments about the self could prove useful as a therapeutic tool for promoting self-concept change. The tendency to focus on outcome information only when it is self-flattering may represent a fruitful strategy for preserving self-esteem, and it may represent one of many cognitive strategies that go awry among depressed individuals (Burns, 1980). Finding ways to alter dysfunctional self-concepts by manipulating or highlighting information bearing on outcomes would seem to be a promising direction for future research.

REFERENCES

Allison, S.T., Mackie, D.M., & Messick, D.M. (1996). Outcome-biases in social perception: Implications for dispositional inference, stereotyping, attitude change, and social behavior. In M. Zanna (Ed.), *Advances in Experimental Social Psychology*. (Vol. 28; pp. 53 94). San Diego: Academic Press.

Allison, S.T., & Messick, D.M. (1985). The group attribution error. *Journal of Experimental Social Psychology*, 21, 563 579.

Allison, S.T., Worth, L.T., & King, M.W.C. (1990). Group decisions as social inference heuristics. *Journal of Personality and Social Psychology*, 58, 801 811.

Andersen, S. M., & Ross, L. (1984). Self-knowledge and social inference: I. The impact of cognitive/affective and behavioral data. *Journal of Personality and Social Psychology*, 46, 280 293.

Baron, J., & Hershey, J. C. (1988). Outcome-bias in decision evaluation. *Journal of Personality and Social Psychology*, 54, 569 579.

Beggan, J. K., & Allison, S. T. (1993). The landslide victory that wasn't: The bias toward consistency in recall of election support. *Journal of Applied Social Psychology*, 23, 669 677.

Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.) *Advances in Experimental Social Psychology* (Vol. 6). New York: Academic Press.

Bradley, G. W. (1978). Self-serving biases in the attribution process: A reexamination of the fact or fiction question. *Journal of Personality and Social Psychology*, 36, 56 71.

Burns, D. D. (1980). *Feeling good: The new mood therapy*. New York: William Morrow.

Chaiken, S., Liberman, A., & Eagly, A.H. (1989). Heuristic and systematic information processing within and beyond the persuasion context. In J. S. Uleman & J.A. Bargh (Eds.), *Unintended thought* (pp. 212 252). New York: Guilford Press.

Fischhoff, B. (1975). Hindsight does not equal foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 288 299.

Gilbert, D. T. (1989). Thinking lightly about others: Automatic components of the social inference process. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought: Limits of awareness, intention, and control* (pp. 189 211). New York: Guilford.

Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin*, 117, 21 38.

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

Jones, E. E. (1979). The rocky road from acts to dispositions. *American Psychologist*, 34, 107 117.

Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology*, 3, 1 24.

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

Klein, S. B., & Loftus, J. (1993). Behavioral experience and trait judgments about the self. *Personality and Social Psychology Bulletin*, 19, 740 745.

Klein, S. B., Sherman, J. W., & Loftus, J. (1996). The role of episodic and semantic memory in the development of trait self knowledge. *Social Cognition*, 14, 277 291.

Klein, W. M. & Kunda, Z. (1993). Maintaining self-serving social comparisons: Biased reconstruction of one's past behaviors. *Personality and Social Psychology Bulletin*, 19, 732 739.

Knight, J. A., & Vallacher, R. R. (1981). Interpersonal engagement in social perception: The consquence of getting into the action. *Journal of Personality and Social Psychology*, 40, 990 999.

Kruglanski, A. W. (1989). *Lay epistemics and human knowledge: Cognitive and motivational biases*. New York: Plenum Press.

Kunda, Z. (1987). Motivated inference: Self-serving generation and evaluation of causal theories. *Journal of Personality & Social Psychology*. 53, 636 647.

Lerner, M. J., & Miller, D. T. (1978). Just world research and the attribution process: Looking back and ahead. *Psychological Bulletin*, 85, 1030 1051.

Linville, P. W., & Carlston, D. E. (1994). Social cognition of the self. In P. Devine, D. Hamilton, & T. Ostrom (Eds.), *Social cognition: Impact on social psychology* (pp. 143 193). San Diego: Academic Press.

Mackie, D. M., & Ahn, M. N. (1998). Ingroup and outcome inferences: When ingroup bias overwhelms outcome-bias. *European Journal of Social Psychology*, 28, 343 360.

Mackie, D. M., Ahn, M. N., Asuncion, A., & Allison, S. T. (2001). The impact of perceiver attitude on outcome-biased trait inferences. *Social Cognition*, 19, 71 93.

Mackie, D. M., & Allison, S. T. (1987). Group attribution errors and the illusion of group attitude change. *Journal of Experimental Social Psychology*, 23, 460 480.

Mackie, D.M., Allison, S., Worth, L., & Asuncion, A. (1992a). The generalization of outcome-biased counter stereotypic inferences. *Journal of Experimental Social Psychology*, 28, 43 64.

Mackie, D.M., Allison, S., Worth, L., & Asuncion, A. (1992b). The impact of outcome-biases on counterstereotypic inferences about groups. *Personality and Social Psychology Bulletin*, 18, 44 51.

Mackie, D.M., Worth, L.T., & Allison, S.T. (1990). Outcome-biased inferences and the perception of change in groups. *Social Cognition*, 8, 325 342.

Markus, H. (1977). Self-schemata and processing information about the self. *Journal of Personality and Social Psychology*, 35, 63 78.

McHoskey, J.W., & Miller, A.G. (1994). Effects of constraint identification, processing mode, expectancies, and intragroup variability on attributions toward group members. *Personality and Social Psychology Bulletin*, 20, 266 276.

Miller, D. T., & Ross, M. (1975). Self-serving biases in attribution of causality: Fact or fiction? *Psychological Bulletin*, 82, 213 225.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10). New York: Academic Press.

Schroth, H., & Messick, D. M. (1994). *Attribution of group attitudes from active versus passive decisions*. Unpublished manuscript, Northwestern University.

Walster, E. (1966). The assignment of responsibility for an accident. *Journal of Personality and Social Psychology*, 3, 73 79.

Zuckerman, M. (1979). Attribution of success and failure revisited, or: The motivational bias is alive and well in attribution theory. *Journal of Personality*, 47, 245 287.

APPENDIX A: Survey Instrument Used to Measure Broad & Narrow-Mindedness Self-Ratings

How friendly are you?
EXTREMELY UNFRIENDLY 1-2-3-4-5-6-7-8-9 EXTREMELY FRIENDLY

How smart are you?
EXTREMELY STUPID 1-2-3-4-5-6-7-8-9 EXTREMELY SMART

How talkative are you?
EXTREMELY QUIET 1-2-3-4-5-6-7-8-9 EXTREMELY TALKATIVE

How mature are you?
EXTREMELY IMMATURE 1-2-3-4-5-6-7-8-9 EXTREMELY MATURE

How flexible are you?
EXTREMELY STUBBORN 1-2-3-4-5-6-7-8-9 EXTREMELY FLEXIBLE

How ambitious are you?
EXTREMELY NOT 1-2-3-4-5-6-7-8-9 EXTREMELY AMBITIOUS

How energetic are you?
EXTREMELY LAZY 1-2-3-4-5-6-7-8-9 EXTREMELY ENERGETIC

How broad-minded are you?
EXTREMELY NARROW-MINDED 1-2-3-4-5-6-7-8-9 EXTREMELY-BROADMINDED

How intuitive are you?
EXTREMELY UNINTUITIVE 1-2-3-4-5-6-7-8-9 EXTREMELY INTUITIVE

How artistic are you?
EXTREMELY UNARTISTIC 1-2-3-4-5-6-7-8-9 EXTREMELY ARTISTIC

How empathetic are you?
EXTREMELY UNEMPATHETIC 1-2-3-4-5-6-7-8-9 EXTREMELY FRIENDLY

How moody are you?
EXTREMELY UNFRIENDLY 1-2-3-4-5-6-7-8-9 EXTREMELY EMPATHETIC

How introspective are you?
EXTREMELY NOT 1-2-3-4-5-6-7-8-9 EXTREMELY INTROSPECTIVE

How creative are you?
EXTREMELY UNCREATIVE 1-2-3-4-5-6-7-8-9 EXTREMELY CREATIVE

How grateful are you?
EXTREMELY UNGRATEFUL 1-2-3-4-5-6-7-8-9 EXTREMELY GRATEFUL

APPENDIX B: Descriptive Statistics and Correlations Among the Dependent Variables

		Mean	SD	1	2	3	4
1	Inferences	6.26	1.22	(.89)			
2	Certainty	4.11	1.31	.26**	(.75)		
3	Importance	4.63	1.08	.21	.10	(.82)	
4	Validity	4.01	1.44	.16	.04	.02	(.88)

**p < .05; Coefficient alphas appear on the diagonal

AUTHOR NOTE

This research was supported by the John Neasmith Dickinson Award to A. Neville Uhles from the University of Richmond, a Faculty Summer Research Fellowship awarded to Scott T. Allison from the University of Richmond, a CSU Affirmative Action Faculty Development Grant and a SJSU Psychology Department Research Committee Grant awarded to Arlene Asuncion, and NSF Grant SBR 9209995 awarded to Diane Mackie.
Correspondence concerning this article should be sent to Scott Allison, Department of Psychology, University of Richmond, VA 23173. Email can be sent to sallison@richmond.edu.

AUTHOR BIOGRAPHIES

Scott T. Allison received his Ph.D. in Social Psychology at the University of California, Santa Barbara, in 1987. He is currently MacEldin Trawick Professor and Chair of Psychology at the University of Richmond. His email address is sallison@richmond.edu.

A. Neville Uhles received her B.A. degree in Psychology from the University of Richmond. Her email address is nuhles@erol.com.

Arlene G. Asuncion received her Ph.D. in Social Psychology at the University of California, Santa Barbara, in 1993. She is an Associate Professor of Psychology at San Jose State University. Her email address is asuncia@email.sjsu.edu.

James K. Beggan received his Ph.D. in Social Psychology at the University of California, Santa Barbara, in 1989. He is an Associate Professor of Sociology at the University of Louisville. His email address is james.beggan@louisville.edu.

Diane M. Mackie received her Ph.D. in Psychology at Princeton University in 1984. She is Professor of Psychology at the University of California, Santa Barbara. Her email address is mackie@psych.ucsb.edu.