

# Understanding Implicit and Explicit Attitude Change: A Systems of Reasoning Analysis

Robert J. Rydell  
University of California, Santa Barbara

Allen R. McConnell  
Miami University

There is considerable controversy about how to conceptualize implicit and explicit attitudes, reflecting substantial speculation about the mechanisms involved in implicit and explicit attitude formation and change. To investigate this issue, the current work examines the processes by which new attitudes are formed and changed and how these attitudes predict behavior. Five experiments support a systems of reasoning approach to implicit and explicit attitude change. Specifically, explicit attitudes were shaped in a manner consistent with fast-changing processes, were affected by explicit processing goals, and uniquely predicted more deliberate behavioral intentions. Conversely, implicit attitudes reflected an associative system characterized by a slower process of repeated pairings between an attitude object and related evaluations, were unaffected by explicit processing goals, uniquely predicted spontaneous behaviors, and were exclusively affected by associative information about the attitude object that was not available for higher order cognition.

*Keywords:* implicit attitudes, explicit attitudes, attitude change

The study of attitudes—evaluations of the self, individuals, groups, and other objects—has a long and rich history in social psychology (Eagly & Chaiken, 1993). In recent years, the focus of attitude research has shifted from understanding explicit attitudes (i.e., attitudes that people can report and for which activation can be consciously controlled) to examining implicit attitudes (i.e., attitudes for which people do not initially have conscious access and for which activation cannot be controlled).<sup>1</sup> Past research has shown that relying on implicit rather than explicit measures of attitudes can circumvent self-presentational motives (e.g., Dunton & Fazio, 1997) and can often uniquely predict spontaneous behaviors (e.g., McConnell & Leibold, 2001); however, less is known about the processes underlying how implicit and explicit attitudes form and operate. The current work posited that there are important differences between them, especially in how they change. Specifically, we propose that explicit attitudes form and

change through the use of fast-learning, rule-based reasoning, whereas implicit attitudes form and change through the use of slow-learning, associative reasoning (Sloman, 1996).

Heretofore, implicit attitude change and explicit attitude change have been studied in relative isolation. Indeed, research on explicit attitude change has been one of the most productive areas of study in social psychology (Eagly & Chaiken, 1993; Petty & Wegener, 1998). Although some researchers have found that implicit attitudes are relatively difficult to change with conventional attitude change manipulations (e.g., Gawronski & Strack, 2004; Gregg, Seibt, & Banaji, 2006; Petty, Tormala, Briñol, & Jarvis, 2006), other research has demonstrated that implicit attitudes can change relatively quickly in response to contextual stimuli or social roles (e.g., Barden, Maddux, Petty, & Brewer, 2004; Dasgupta & Greenwald, 2001; Wittenbrink, Judd, & Park, 2001). But despite these demonstrations, the theory underlying implicit attitude change is relatively underdeveloped (see Devine, 2001; Fazio & Olson, 2003; Wilson, Lindsey, & Schooler, 2000), and experimental paradigms that can systematically examine the concurrent formation and change of implicit and explicit attitudes

---

Robert J. Rydell, Department of Psychology, University of California, Santa Barbara; Allen R. McConnell, Department of Psychology, Miami University.

This research was supported by National Institute of Mental Health Grant MH068279 and National Science Foundation Grant BCS 0516931. Portions of this work were submitted by Robert J. Rydell in partial fulfillment of the requirements for a doctoral degree from Miami University. We thank Doris Bergen, Heather Claypool, Dave Hamilton, Kurt Hugenberg, Diane Mackie, Jeff Sherman, and Laura Strain for their extremely helpful comments and guidance in this research.

Correspondence concerning this article should be addressed to Robert J. Rydell, Department of Psychology, University of California, Santa Barbara, Santa Barbara, CA 93106-9660. E-mail: rydell@psych.ucsb.edu

---

<sup>1</sup> Although there is disagreement about the use of the terms *implicit attitudes* and *explicit attitudes* in the literature (e.g., Fazio & Olson, 2003), we agree with Strack and Deutsch (2004) who note that “explicit and implicit measures are defined by the cognitive operations that they capture. In this sense, explicit measures tap into people’s knowledge or beliefs, implicit measures tap into their associative structures” (p. 239; see also, Wilson et al., 2000). Because we contrast and compare implicit and explicit measures, we use the terms *implicit attitudes* and *explicit attitudes* throughout this article.

is lacking. The current work seeks to address these issues. Such an advance would not only speak to important theoretical issues (e.g., are the processes underlying implicit and explicit attitude change fundamentally different?), but it has implications for topics ranging from persuasion (e.g., Petty et al., 2006) to intergroup relations (e.g., Jellison, McConnell, & Gabriel, 2004). We now turn to developing a framework that can account for how implicit and explicit attitudes change.

### Systems of Reasoning

Sloman (1996) proposed a systems of reasoning approach to account for how different cognitive systems affect thinking, language, and behavior (see also, Gawronski & Bodenhausen, 2006; Smith & DeCoster, 2000; Strack & Deutsch, 2004). He argued that two independent systems of reasoning use very different processes to operate, learn, and change. The first system of reasoning, the *slow-learning system*, operates by using paired associations based on similarity and contiguity. In this case, learning is based on the slow accrual of information over time to form and strengthen associations in memory. The second system of reasoning proposed by Sloman, the *fast-learning system*, relies on logical, verbal, or symbolic representations at a relatively higher order level of cognitive processing. Judgments and behaviors rendered by this system are based on processes requiring at least some degree of conscious control (Strack & Deutsch, 2004). Unlike the slow-learning system, which relies on the accretion of paired associations in memory, the fast-learning system can operate relatively quickly and flexibly to take into account new information that is not associative in nature, but rather, reflects abstractions, language, and logic. In summary, the slow-learning system is characterized by more automatic processes based on the slow accumulation of paired associations in memory, whereas the fast-learning system responds relatively more flexibly and deliberately to abstract information rather than accumulating associations in memory (Smith & DeCoster, 2000).

When imported into the attitudes literature, this systems of reasoning approach maps nicely onto implicit attitudes and explicit attitudes. That is, the slow-learning system can shed light on how implicit attitudes form and function because implicit attitudes are posited to follow the basic principles of similarity and paired associations across time (Olson & Fazio, 2001; Smith & DeCoster, 2000; Wilson et al., 2000). On the other hand, the fast-learning system is compatible with explicit attitudes, which can change quickly and often require some degree of cognitive resources in their production and revision (Fazio, 1995; Petty & Wegener, 1998). Indeed, it has been proposed that implicit attitudes and explicit attitudes are the products of different and distinct underlying cognitive processes (Wilson et al., 2000), and accordingly, empirical studies have demonstrated that implicit and explicit attitudes predict different kinds of behavior (spontaneous and nonverbal vs. deliberate and self-presentational, respectively; Dovidio, Kawakami, & Gaertner, 2002; Jellison et al., 2004; McConnell & Leibold, 2001).

On the basis of a systems of reasoning account, one would anticipate that implicit and explicit attitudes might be differentially responsive to particular types (nonconscious and associative vs. conscious and verbal, respectively) of attitude-object information.

Indeed, Rydell, McConnell, Mackie, and Strain (in press) recently demonstrated that explicit attitudes were formed in response to consciously available information, whereas implicit attitudes

formed in response to the valence of subliminally presented primes when both types of information were available. Specifically, participants were presented with a series of trials in which a target person was preceded by a subliminal prime (either positive or negative in valence), who was described in a sentence as having performed a particular behavior (the valence of which was always opposite of the subliminal prime). After a number of such trials, Rydell et al. found that implicit attitudes toward the person reflected the valence of the subliminal primes, whereas explicit attitudes responded to the valence of the verbally presented behaviors. For example, when presented with negative subliminal primes and positive behavioral sentences, participants reported negative implicit attitudes and positive explicit attitudes toward the same target. Consistent with a systems of reasoning account, the formation of implicit and explicit attitudes were independent of each other, with each reflecting the type of information (associative and nonconscious vs. verbal and conscious) assumed to influence the slow-learning and fast-learning systems, respectively. These findings are difficult to explain by attitude theories that do not assume that people can simultaneously hold different implicit and explicit attitudes about the same object (e.g., Fazio & Olson, 2003; Petty & Wegener, 1998).<sup>2</sup>

Despite this evidence of independent implicit and explicit attitude formation, the question remains as to what processes underlie the formation and change of implicit and explicit attitudes. In other words, although Rydell et al. (in press) established the independence of implicit and explicit attitudes, they did not evaluate whether slow-learning and fast-learning systems (respectively) account for these outcomes. The current research focuses directly on this issue.

It was anticipated that, in general, implicit attitudes would change more slowly than explicit attitudes in response to target-relevant information because implicit attitudes reflect the slow accrual of paired associations in memory. Although this should be the case when information is presented so that it can be acted on by higher order cognition, there should also be situations (e.g., information is presented outside of conscious awareness) in which attitude-relevant information will impact implicit attitudes but not explicit attitudes, reaffirming their dissociation (Rydell et al., in press). Another consequence of the dissociation between implicit and explicit attitudes should be revealed in the types of behaviors they predict, with implicit attitudes uniquely predicting subtle, less deliberate behavior (e.g., social distance) and explicit attitudes uniquely predicting more thoughtful actions (e.g., Dovidio et al., 2002; McConnell & Leibold, 2001).

Earlier, we noted that several studies have shown that reports of implicit attitudes can change relatively quickly. This raises the question of how a systems of reasoning approach would explain abrupt shifts in implicit attitudes in response to positive exemplars (Dasgupta & Greenwald, 2001), contextual features (Wittenbrink et al., 2001), or

<sup>2</sup> Instead, these theories assume that people hold an attitude about an object in memory whose expression can be adjusted to accommodate self-presentational concerns, differences in motivation and cognitive resources, societal norms, or persuasive communications. In summary, these models assume that implicit measures reflect an association between an attitude object and its evaluation in memory, whereas explicit measures elucidate more "downstream" consequences of accessing the attitude (Fazio & Olson, 2003).

social roles (Barden et al., 2004). In our view (see also, Gawronski & Bodenhausen, 2006; Mitchell, Nosek, & Banaji, 2003; Richeson & Trawalter, 2005), implicit attitudes, just like any other memorial structure, can be affected by priming manipulations that increase the accessibility of a subset of information associated with an attitude object (e.g., increasing the accessibility of positive members of a stigmatized group) or even affect how attitude objects are classified (e.g., encountering nonprototypic group members may temporally impact the categorization of subsequent group members). Thus, situational factors may alter the accessibility of associations related to an attitude object without changing the attitude in an enduring fashion. Instead, these factors may affect which information about an attitude object is activated in memory and may also change the standards used for category membership.

Although these effects are important to explore and can speak to the underlying mechanisms of attitudes and attitude activation (especially for well-established attitude objects where many of these outcomes, such as temporarily increasing the accessibility of a subset of group members, are possible), the current work was concerned with understanding how slow-learning and fast-learning processes can account for how implicit and explicit attitudes change. Thus, we had participants learn about a novel attitude object under conditions in which we could manipulate the learning history of the attitude object, allowing us to examine the basic mechanisms through which implicit and explicit attitudes form and change.

### Overview of the Current Work

Five experiments were conducted to understand whether slow-learning and fast-learning systems could account for implicit and explicit attitudes. To examine a systems of reasoning approach to attitudes, we gave participants information about a novel target person (*Bob*) in a learning paradigm that initially presented considerable behavioral information about *Bob* before revealing counterattitudinal behavioral information about him (i.e., behavioral information inconsistent with the valence of the initial information). Afterward, participants reported their implicit and explicit attitudes toward *Bob*. We sought to understand when and how counterattitudinal information affected implicit and explicit attitudes differently.

Experiment 1 examines the conditions under which quick changes in explicit attitudes, but not implicit attitudes, are found. Experiment 2 examines conditions under which implicit attitudes do change in response to counterattitudinal information and how these changes differ from those observed for explicit attitudes in response to the same information. In Experiment 3, we focus on how providing explicit processing goals for forming impressions affects explicit attitudes but not implicit attitudes. Experiment 4 examines how implicit and explicit attitudes toward *Bob* in the current paradigm uniquely predict different types of behaviors directed toward him. Finally, Experiment 5 uses a modified learning paradigm (incorporating subliminal priming and measuring attitudes at two different times; similar to Rydell et al., in press) to demonstrate conditions under which implicit attitudes, but not explicit attitudes, changed in the face of counterattitudinal information about *Bob*.

### Experiment 1

Experiment 1 was designed to demonstrate that implicit and explicit attitudes reflect different systems of reasoning by present-

ing counterattitudinal information to change explicit attitudes (assumed to be governed by the fast-learning system) but not implicit attitudes (presumably governed by the slow-learning system). This experiment was modeled after an elegant study by Kerpelman and Himmelfarb (1971) in which participants were randomly assigned to receive positive reinforcement (i.e., positive behaviors are characteristic of the target and negative behaviors are uncharacteristic of the target) about the behaviors performed by a target person 100%, 80%, or 70% of the time over a block of 100 trials (with the remaining trials providing negative counterattitudinal feedback in the latter two conditions). After this initial learning, participants reported their explicit attitudes toward the target person or they learned in 50 subsequent trials that the target person performed additional behaviors that were of the opposite valence from the majority of the first 100 and then reported their explicit attitudes. Participants in the 100% reinforcement condition displayed a drastic and almost immediate change in their evaluations of the target person in the direction opposite to the originally learned attitude. Thus, to the extent that original explicit attitudes were more extreme because of initially greater consistency in levels of reinforcement, participants showed greater shifts in their explicit attitudes in line with the counterattitudinal information presented.

However, because Kerpelman and Himmelfarb (1971) only examined how positive attitudes were changed by negative counterattitudinal information, it is also possible that the processes involved in explicit attitude change in this paradigm are more complex than they acknowledged. Indeed, there is reason to believe that negative counterattitudinal information (i.e., learning negative information following mostly positive initial information) about a target will change attitudes more strongly than positive counterattitudinal information (i.e., learning positive information following mostly negative initial information). Notably, for social judgments involving liking (like those used in the current work), negative information receives greater emphasis and is more crucial in forming impressions (Fiske, 1980; Skowronski & Carlston, 1987).

Although these negative asymmetries have been shown for explicit attitudes, it is an open question as to whether they also occur for implicit attitudes. For example, one could argue that implicit attitudes would also be more impacted by negative counterattitudinal information because of its greater attention-grabbing value (e.g., Pratto & John, 1991). However, although negative behaviors are more diagnostic for liking judgments, positive behaviors are more diagnostic for ability judgments (Skowronski & Carlston, 1987). Thus, a simple "negative information is given more weight" explanation seems insufficient. Also, because extracting a trait from behavior may rely on some amount of effortful processing (Bassili & Smith, 1986) and may require verbal processes (Carlston, 1994), it is possible that valence asymmetries are more likely for the fast-learning system than for the slow-learning system. With these latter points in mind, our prediction was that explicit attitudes were more likely to reveal a valence asymmetry (i.e., stronger attitude change following negative counterattitudinal exposure than that following positive counterattitudinal exposure) than implicit attitudes.

Thus, we used a learning paradigm similar to that of Kerpelman and Himmelfarb (1971) because it provides a useful way to study how implicit attitudes and explicit attitudes about the same attitude object are formed and changed differently on the basis of the same information. First, participants received a considerable amount of information about *Bob*, allowing them to form implicit attitudes

toward him. Second, the introduction of the counterattitudinal information provided a window in which explicit (which relies on the fast-learning system) but not implicit (which relies on the slow-learning system) attitudes should change in the face of new target-relevant information. Thus, we have an opportunity to assess and understand how this new information affects implicit attitudes and explicit attitudes differently, shedding light on the processes involved in their change.

We expected to observe that people would quickly change their explicit attitudes in the face of counterattitudinal information, especially when the initial learning was very consistent (Kerelman & Himmelfarb, 1971) and when the counterattitudinal information was negative (Fiske, 1980). However, we did not expect implicit attitudes to change as quickly in response to a modest amount of counterattitudinal information nor did we expect to observe a valence asymmetry for implicit attitudes.

## Method

**Participants.** A sample of 170 undergraduates at Miami University participated in return for research credit in their introductory psychology courses. They were randomly assigned to a 2 (valence of learned attitude: positive vs. negative)  $\times$  2 (level of reinforcement: 100%, 75%)  $\times$  2 (counterattitudinal condition: control vs. counterattitudinal conditioning) between-subjects factorial.

**Learning task.** The current work used a modified version of the attitude learning paradigm developed by Kerelman and Himmelfarb (1971). Specifically, participants were presented with a target person's behaviors that were either relatively positive or negative in valence, and participants judged whether each behavior was characteristic or uncharacteristic of him. As part of a between-subjects manipulation, participants were given different levels of reinforcement in their responses, leading them to form different attitudes toward him.

First, participants completed the learning task on a computer, in which they were told that they would be receiving information about a person named *Bob*. In the initial learning trials, participants read 100 behaviors performed by *Bob* while a picture of *Bob* was presented on the computer monitor directly above each behavior.<sup>3</sup> After reading each behavior, participants indicated whether they believed that the behavior was characteristic or uncharacteristic of *Bob* by pressing the *C* key (characteristic) or the *U* key (uncharacteristic). After they responded, participants were given feedback about whether the behavior was characteristic of *Bob* for 5 s. Specifically, feedback consisted of the word *correct* (in blue text) or *incorrect* (in red text) positioned in the center of the computer monitor and, at the same time, the behavior was restated "correctly," on the basis of the assigned reinforcement condition, at the bottom of the computer monitor (e.g., "Helping the neighborhood children is characteristic of Bob." or "Helping the neighborhood children is uncharacteristic of Bob."). In the initial 100 learning trials, the feedback given portrayed *Bob* as positive or as negative in 100% or in 75% of the behaviors (with 25 of the trials in the 75% reinforcement condition being counterattitudinal). The ordering of the behaviors and feedback were randomly determined (in accordance with the experimental condition) for each participant.

Following these 100 trials, participants in the control condition received 20 neutral trials (i.e., the behavior performed by *Bob* was neither positive nor negative; e.g., "Bob waited at the street corner."). However, participants in the counterattitudinal condition (20 CA) received counterattitudinal feedback about *Bob* on 20 trials (i.e., the behaviors that were described as characteristic or uncharacteristic of *Bob* were opposite of the valence presented during the initial learning trials). Finally, participants completed implicit and explicit attitude measures.<sup>4</sup>

**Explicit attitude measure.** To assess explicit attitudes, participants judged how likable *Bob* was on a scale ranging from 1 (*very unlikely*) to

9 (*very likable*). In addition, they completed five semantic differential scales, each using a 9-point scale to describe *Bob*: good–bad, pleasant–mean, agreeable–disagreeable, caring–uncaring, and kind–cruel. Further, participants provided their evaluation of *Bob* on a feeling thermometer that ranged in temperature from 0° to 100°. The response for each explicit measure was standardized and an overall mean was computed (in all experiments to be reported,  $\alpha$ s > .90). Then the standardized scores in the negative valence condition were reverse scored so that greater scores on this measure indicated that explicit attitudes were more extreme in the direction of initial learning.

**Implicit attitude measure.** The Implicit Associations Test (IAT; Greenwald, McGhee, & Schwartz, 1998) was used to assess implicit attitudes toward *Bob*. The IAT had 26 stimuli: 1 picture of *Bob*, 5 different pictures of White men who were not *Bob*, 10 positive adjectives (e.g., *wonderful*), and 10 negative adjectives (e.g., *disgusting*). All stimuli were presented in the center of the monitor, and the adjectives were always presented in lowercase letters.

This IAT task was a modified version of the task used by McConnell and Leibold (2001), featuring seven blocks with 20 trials per block. Participants were informed that the task involved making category judgments for a variety of stimuli (photos or words) presented on a computer monitor by using one of two responses (the *D* or *K* keys on the keyboard). During each block, category label reminders were displayed on the left and right sides of the display (assignment of particular labels to the *D* and *K* keys was counterbalanced across participants and produced no effects). Participants were instructed to complete that task quickly while also minimizing errors, and they were told to keep their index fingers on the *D* and *K* keys throughout the experiment to minimize delays in responding. There was a 250-ms intertrial interval.

In Block 1, participants judged photos of *Bob* or *not Bob* and in Block 2 they judged whether the adjectives were "negative" or "positive." In Blocks 3 and 4 (Combination 1), participants judged whether the stimuli were "Bob or negative" or "not Bob or positive." In Block 5, participants performed the same judgment task as Block 2 except the assignment of response keys assigned to the two valence categories was reversed. Finally, in Blocks 6 and 7 (Combination 2), participants judged whether the stimuli were "Bob or positive" or "not Bob or negative." As in past IAT research, half of the participants performed Combination 1 in Blocks 3–4 and Combination 2 in Blocks 6–7, whereas the rest performed Combination 2 in Blocks 3–4 and Combination 1 in Blocks 6–7 (this counterbalancing manipulation produced no effects).

In order to assess implicit attitudes toward *Bob*, we subtracted the mean response latencies of Combination 2 from the mean response latencies of Combination 1 (regardless of the order they were completed).<sup>5</sup> Again, the

<sup>3</sup> Photographs of one of 5 different White males were randomly presented as *Bob*. These 5 White males were judged as equal in attractiveness and the target used did not affect the results in any of the experiments. The positive and the negative behaviors used in the current work were borrowed from those developed by McConnell, Sherman, and Hamilton (1994a).

<sup>4</sup> In all of the experiments, half of the participants completed the implicit measure first and the other half completed the explicit measure first. This order variable produced no effects in any of the studies and thus is not discussed further.

<sup>5</sup> Following Greenwald et al. (1998), all trials in the critical blocks were retained, responses faster than 300 ms were recoded as 300 ms, and trials slower than 3,000 ms were recoded as 3,000 ms. After any such adjustments were made, each latency was then log transformed to reduce positive skew inherent in response latency data (Fazio, 1990). Alternative scoring techniques for the IAT (e.g., Greenwald, Nosek, & Banaji, 2003) produced the same results in all studies reported. Analyses were performed on the log-transformed values, but means are reported as standardized scores.



standardized scores in the negative valence condition were reverse scored so that greater scores on this measure indicated that implicit attitudes were more extreme in the direction of initial learning.

## Results

**Explicit attitudes.** To examine whether explicit attitudes changed in response to small amounts of counterattitudinal information and were more likely to show attitude change with greater initial reinforcement, a 2 (valence of learned attitude)  $\times$  2 (level of reinforcement)  $\times$  2 (counterattitudinal condition) analysis of variance (ANOVA) was conducted on explicit attitude extremity (see Figure 1). First, there were significant main effects of level of reinforcement,  $F(1, 162) = 39.22, p < .001$ , and of counterattitudinal condition,  $F(1, 162) = 89.90, p < .001$ . As one would expect, the main effect of level of reinforcement showed that explicit attitudes were more extreme in the direction of initial learning in the 100% reinforcement condition ( $M = 0.94, SD = 0.62$ ) than in the 75% reinforcement condition ( $M = 0.54, SD = 0.45$ ). Similarly, the main effect of counterattitudinal condition revealed that explicit attitudes were more extreme in the direction of initial learning in the control condition ( $M = 1.04, SD = 0.52$ ) than in the 20 CA condition ( $M = 0.44, SD = 0.47$ ). More important, the anticipated two-way interaction between level of reinforcement and counterattitudinal condition was significant,  $F(1, 162) = 19.06, p < .001$ . To examine this interaction, the simple effect of counterattitudinal condition was examined for each level of reinforcement. In the 75% reinforcement condition, there was a simple effect of counterattitudinal condition,  $F(1, 162) = 13.54, p < .001$ , showing that participants in the control condition had more extreme explicit attitudes toward Bob ( $M = 0.71, SD = 0.42$ ); than participants in the 20 CA condition ( $M = 0.38, SD = 0.41$ ). In the 100% reinforcement condition, there was an even stronger effect of counterattitudinal condition,  $F(1, 162) = 97.03, p < .001$ , indicating that although participants in the control condition had especially extreme explicit attitudes (in the direction of initial conditioning) toward Bob ( $M = 1.37, SD = 0.37$ ), counterattitudinal information led to far less extreme attitudes toward Bob ( $M = 0.51, SD = 0.52$ ). Thus, the interaction reflects the much larger effect of counterattitudinal condition on explicit

attitude extremity in the 100% reinforcement condition than in the 75% reinforcement condition (replicating Kerpelman & Himmelfarb, 1971). Consistent with negative asymmetries, the two-way interaction between counterattitudinal condition and valence of learned attitude was also significant,  $F(1, 162) = 16.64, p < .001$ . In the positive learned attitudes condition, those in the control condition had far more extreme explicit attitudes ( $M = 1.17, SD = 0.59$ ) than those in the 20 CA condition ( $M = 0.28, SD = 0.47$ ),  $F(1, 162) = 76.70, p < .001$ . In the negative learned attitudes condition, this effect was significant but weaker, with those in the control condition having more extreme explicit attitudes ( $M = 0.91, SD = 0.39$ ) than those in the 20 CA condition ( $M = 0.61, SD = 0.42$ ),  $F(1, 162) = 8.12, p < .005$ . In other words, negative counterattitudinal information had a greater impact on attitude extremity than did positive counterattitudinal information (e.g., Fiske, 1980; Skowronski & Carlston, 1987). No other effects were significant.

**Implicit attitudes.** As with the explicit attitude data, a 2 (valence of learned attitude)  $\times$  2 (level of reinforcement)  $\times$  2 (counterattitudinal condition) ANOVA was conducted on implicit attitude extremity (see Figure 2). In stark contrast to the explicit attitudes, the interaction of reinforcement and counterattitudinal condition and the interaction of valence of learned attitude and counterattitudinal condition were not significant for implicit attitudes ( $F_s < 1$ ). In fact, the only effect to obtain for implicit attitudes was an effect showing the that grand mean was significantly different than zero,  $F(1, 166) = 55.12, p < .001$  ( $M = 0.50, SD = 0.87$ ). This shows that participants formed implicit attitudes about Bob in accordance with the valence of their initial learning but that subsequent counterattitudinal information had no impact on them. It is important that this effect was not statistically moderated by any of the experimental manipulations, showing no evidence of changes in attitude extremity or negative asymmetries for implicit attitudes.<sup>6</sup>

## Discussion

A systems of reasoning conceptualization of attitude change was supported in this experiment because explicit attitudes were changed dramatically by the introduction of counterattitudinal information, whereas implicit attitudes were unaltered by this same information. This suggests that explicit attitudes are the product of a fast-learning system, whereas implicit attitudes reflect a slow-learning system. In this study, participants did form implicit attitudes about Bob, but,




Figure 1. Explicit attitude extremity as a function of reinforcement and counterattitudinal condition (20 CA) in Experiment 1. Values for the negative initial learning condition have been reverse scored to reflect attitude extremity.

<sup>6</sup> When implicit and explicit attitude measures were simply standardized (i.e., the standardized attitudes in the negative valence of learned attitude condition were not reverse scored) and submitted to a 2 (valence of learned attitude)  $\times$  2 (level of reinforcement)  $\times$  2 (counterattitudinal condition)  $\times$  2 (standardized attitude measure: implicit vs. explicit, a repeated measure) mixed-model factorial ANOVA, the expected four-way interaction was significant,  $F(1, 156) = 3.98, p < .05$ , reflecting differential responses to counterattitudinal feedback for explicit attitudes and implicit attitudes. In all subsequent experiments, similar omnibus analyses were conducted by using the standardized attitude measure as a within-subjects factor, and the highest order interaction obtained in each experiment ( $F_s > 3.88, p_s < .03$ ). These analyses reveal that examining implicit and explicit attitudes separately throughout the article is justified inferentially. In the current work, we present the data as examining attitude extremity by reverse scoring the negative learning condition attitude measures in order to simplify the presentation of how implicit and explicit attitudes are differentially affected by our manipulations.




Figure 2. Implicit attitude extremity as a function of reinforcement and counterattitudinal condition (20 CA) in Experiment 1. Values for the negative initial learning condition have been reverse scored to reflect attitude extremity.

unlike explicit attitudes, they were unaffected by the introduction of counterattitudinal information. In addition, the current experiment replicated past work on learned attitudes and research on negative asymmetries in impression formation (e.g., Kerpelman & Himmel-farb, 1971; Skowronski & Carlston, 1987), but these effects were observed only for explicit attitudes (and not for implicit attitudes). That is, explicit attitudes showed the greatest change in the face of contradictory information when negative counterattitudinal information followed initially positive feedback and when it came on the heels of consistent feedback in general.

### Experiment 2

Although Experiment 1 demonstrates that explicit and implicit attitudes were differentially impacted by counterattitudinal information (with explicit attitudes changing quickly and implicit attitudes remaining unaffected by it), it did not test the systems of reasoning derived supposition that implicit attitudes will change slowly if sufficient counterattitudinal information is encountered. Showing that implicit attitudes do change following substantial counterattitudinal information would provide important support for our systems of reasoning approach. If implicit attitudes are the product of a slow-learning system, they should change when enough counterattitudinal information is encountered.

Because there was no significant effect of counterattitudinal information on implicit attitudes in Experiment 1 through the use of just 20 counterattitudinal behaviors (when compared with the control group), Experiment 2 presented some participants with considerably more counterattitudinal information. We expected that explicit attitudes would change quickly in the face of a small amount of counterattitudinal information, whereas implicit attitudes would remain relatively unaffected (replicating Experiment 1). However, we expected that providing participants with a large amount of counterattitudinal information would eventually lead to implicit attitude change as well.

### Method

**Participants.** A sample of 186 undergraduates at Miami University participated in return for research credit in their introductory psychology courses.

They were randomly assigned to a 2 (valence of learned attitude: positive, negative)  $\times$  2 (level of reinforcement: 100%, 75%)  $\times$  3 (counterattitudinal condition: control, 20 CA, 100 CA) between-subjects factorial.

**Procedure.** All materials, methods, and measures (and scoring of the measures) paralleled those used in Experiment 1, with the exception that, in the current experiment, there are three levels of counterattitudinal learning and all participants received 100 initial learning trials followed by 100 additional trials. The control condition and the 20 CA condition were the same as those used in Experiment 1 (except that the final 100 or 80 descriptions of *Bob*, respectively, were neutral so that all participants received the same number of trials). In addition, there was another counterattitudinal condition (100 CA) in which participants received 100 trials of counterattitudinal feedback. Thus, in the 100 CA condition, participants had much more information that was inconsistent with the initially learned attitude than in the other two learning conditions, which should lead to implicit attitude change consistent with the valence of the counterattitudinal information.

### Results

**Explicit attitudes.** To examine explicit attitude change in response to counterattitudinal information, a 2 (valence of learned attitude)  $\times$  2 (level of reinforcement)  $\times$  3 (counterattitudinal condition) ANOVA was conducted on the explicit attitude extremity score (see Figure 3). First, main effects of level of reinforcement,  $F(2, 174) = 36.09, p < .001$ , and of counterattitudinal condition,  $F(2, 174) = 70.99, p < .001$ , were observed. Not surprisingly, explicit attitudes were more extreme in the direction of initial learning in the 100% reinforcement condition ( $M = 0.71, SD = 0.81$ ) than in the 75% reinforcement condition ( $M = 0.20, SD = 0.73$ ). Also, the main effect of counterattitudinal condition showed that explicit attitudes were more extreme in the direction of initial learning in the control condition ( $M = 1.09, SD = 0.64$ ) than in the 20 CA condition ( $M = 0.37, SD = 0.59$ ) and in the 100 CA condition ( $M = -0.09, SD = 0.70$ ), with all means significantly different.<sup>7</sup> It is important that the expected interaction of these two effects obtained,  $F(2, 174) = 9.24, p < .001$ . In the 75% reinforcement condition, there was a simple effect of counterattitudinal condition,  $F(2, 174) = 31.90, p < .001$ , showing that participants in the control condition had more extreme explicit attitudes toward *Bob* ( $M = 0.69, SD = 0.54$ ) than those in the 20 CA condition ( $M = 0.38, SD = 0.57$ ) and those in the 100 CA condition, who had a significantly less extreme view of *Bob* ( $M = -0.44, SD = 0.55$ ) than those in the control condition or in the 20 CA condition. In the 100% reinforcement condition, there was also a simple effect of counterattitudinal condition,  $F(2, 174) = 45.55, p < .001$ . This effect found that participants in the control condition had relatively extreme explicit attitudes toward *Bob* ( $M = 1.50, SD = 0.45$ ); however, the presentation of counterattitudinal information led participants to have less extreme attitudes toward *Bob*, which did not vary between the 20 CA ( $M = 0.37, SD = 0.62$ ) and 100 CA ( $M = 0.24, SD = 0.68$ ) conditions.

Revealing the expected negative asymmetries, the two-way interaction between counterattitudinal condition and valence of learned attitude was also significant,  $F(2, 174) = 3.94, p < .03$ . In the positive learned attitudes condition, those in the control condition had significantly more extreme explicit attitudes ( $M = 1.21, SD = 0.79$ ) than did those in the 20 CA condition ( $M = 0.36,$

<sup>7</sup> All post hoc tests described as significant differed at the .05 level with Tukey's honestly significant difference.




Figure 3. Explicit attitude extremity as a function of reinforcement and counterattitudinal conditions (20 CA, 100 CA) in Experiment 2. Values for the negative initial learning condition have been reverse scored.

$SD = 0.45$ ), who had significantly more extreme attitudes than those in the 100 CA condition ( $M = -0.25$ ,  $SD = 0.64$ ),  $F(2, 174) = 42.34$ ,  $p < .001$ . In the negative learned attitudes condition, those in the control condition had significantly more extreme explicit attitudes ( $M = 0.96$ ,  $SD = 0.40$ ) than did those in the 20 CA ( $M = 0.40$ ,  $SD = 0.71$ ) and 100 CA ( $M = 0.06$ ,  $SD = 0.74$ ) conditions. Thus, as in Experiment 1, explicit attitudes were changed more strongly in response to negative counterattitudinal information than to positive counterattitudinal information.

**Implicit attitudes.** Implicit attitude extremity was calculated in the same fashion as in Experiment 1, and it was analyzed in a 2 (valence of learned attitude)  $\times$  2 (level of reinforcement)  $\times$  2 (counterattitudinal condition) ANOVA (see Figure 4). The only effect to obtain was the expected main effect of counterattitudinal condition,  $F(2, 174) = 5.02$ ,  $p < .01$ . That is, implicit attitudes in the 100 CA condition ( $M = -0.23$ ,  $SD = 0.97$ ) were significantly less consistent with the direction of initial learning than were those in the control condition ( $M = 0.27$ ,  $SD = 0.88$ ) and 20 CA condition ( $M = 0.23$ ,  $SD = 1.09$ ), which did not differ. As expected, these results show that implicit attitudes did change when sufficient counterattitudinal information (100 CA) was presented. Yet replicating Experiment 1, there was no difference in implicit attitude extremity between the control condition and the 20 CA condition, and once again there was no evidence of negative asymmetry effects (i.e., stronger attitude change when negative information follows initial positive information).

## Discussion

The results of Experiment 2 show that implicit attitudes change if sufficient counterattitudinal information is encountered. Because we assume that implicit attitudes reflect the totality of the evaluative information associated with an attitude object, a small amount of counterattitudinal information should have little impact in modifying one's implicit attitudes (i.e., the 20 CA condition). However, once the totality of the counterattitudinal information increased sufficiently (i.e., the 100 CA condition), implicit attitudes did show substantial change. These findings provide evidence that different systems of reasoning are responsible for changing implicit attitudes and explicit attitudes. In response to counterattitudinal information, implicit atti-

tudes changed in line with the slow-learning system, whereas explicit attitudes changed more quickly, consistent with the fast-learning system. Further, we replicated the findings of Experiment 1 for valence asymmetries (e.g., Skowronski & Carlston, 1987) and for greater attitude change following relatively more consistent initial reinforcement (e.g., Kerpelman & Himmelfarb, 1971), but once again, only for explicit attitudes.

Although Experiment 2 showed that implicit attitudes were changed by the sufficient presentation of counterattitudinal information, the results of Experiment 2 also show that participants who received less consistent reinforcement (i.e., 75% condition) continued to show explicit attitude change in response to 100 pieces of counterattitudinal information, and those who received consistent reinforcement (i.e., the 100% condition) did not. Why might this occur? We propose that those in the 75% condition may have forestalled judgments of *Bob* and continued to effortfully process more counterattitudinal information about him. Although perceivers typically form fast on line impressions of individuals (McConnell, Sherman, & Hamilton, 1994b, 1997), they do so because they expect considerable consistency in their behaviors (McConnell, 2001; McConnell et al., 1997). Thus, it is possible that the current 75% level of reinforcement condition provides sufficient inconsistency as to lead perceivers to delay forming their impressions of *Bob*. In order to test this explanation, Experiment 3 experimentally manipulates the presumed impression formation theory involved to test whether "rushing to judgment" versus "forestalling judgment" could account for the pattern of explicit attitude data observed in the 100% and 75% reinforcement conditions, respectively. If participants are instructed to forestall judgments, then more linear (rather than asymptotic) explicit attitude change should be observed across the conditions (i.e., control, 20 CA, and 100 CA) regardless of level of reinforcement. And similarly, participants who receive rush to judgment instructions should show more asymptotic (than linear) explicit attitude change regardless of level of reinforcement.

But more important, Experiment 3 allows us to examine the extent to which explicit processing goals affect implicit and explicit attitudes. If explicit attitudes are the product of a fast-learning system, deliberate processing instructions should affect explicit attitudes but not implicit attitudes. Thus, even though




Figure 4. Implicit attitude extremity as a function of reinforcement and counterattitudinal conditions (20 CA, 100 CA) in Experiment 2. Values for the negative initial learning condition have been reverse scored.

participants received the exact same information about the attitude objects, holding different processing goals will likely alter higher order cognitive processes and influence explicit attitudes. Implicit attitudes, because they are based on associations and not on higher order logic, should be relatively unaffected by such deliberate processing goals.

### Experiment 3

A systems of reasoning account predicts that explicit attitudes will be affected by conscious processing goals, but implicit attitudes should not be affected by such goals. Because people are able to selectively use and effortfully give more emphasis to practical information about an attitude object and to devalue other types of information on the basis of processing goals, explicit attitudes should be especially amenable to change by altering processing goals. Alternatively, because implicit attitudes are based on the accrual of information about an attitude object and are not based on the selective use of information, they should be relatively immune to the effects of conscious goals because they are devoid of the higher order logic necessary to follow the goal.

This experiment also allowed us to examine the results for explicit attitudes of Experiment 2 in more detail. We hypothesized that those in the 75% condition were less inclined to rush to judgment to form an early, on line impression of *Bob* (leading them to process later information and to modify their attitudes accordingly), whereas those in the 100% condition relied on initial counterattitudinal information, modified their impression quickly and then were less impacted by later information. To evaluate this explanation, participants' processing goals (i.e., to rely on early information vs. to rely on later information) were experimentally manipulated in Experiment 3. If those in the 75% reinforcement condition of Experiment 2 did adopt the goal of forestalling impression formation of *Bob*, participants explicitly instructed to do so should be more impacted by later information and report relative greater overall attitude change regardless of the actual consistency of initial learning provided (i.e., 75% or 100% reinforcement). Conversely, those instructed to form an early impression should pay less attention to later information and show less overall attitude change, regardless of the consistency of initial learning.

### Method

**Participants.** A sample of 113 Miami University undergraduates participated in return for research credit and were randomly assigned to a 2 (first impressions: correct, incorrect)  $\times$  2 (level of reinforcement: 100%, 75%)  $\times$  3 (counterattitudinal condition: control, 20 CA, 100 CA) between-subjects factorial.

**Procedure.** All materials, methods, and measures paralleled those of Experiment 2 with three exceptions. First, only the positive valence condition was used (thus, the initial attitudes were positive and counterattitudinal information, when presented, was negative). Second, we manipulated instructions for the learning task such that participants were told to rely on either initial or later information in forming their impressions. Third, because there was no negative initial learning condition requiring reverse scoring, we discuss our data in terms of more positive attitudes, as opposed to greater learning-consistent attitude extremity, toward *Bob*.

**Manipulating the value of first impressions.** To manipulate the importance of early versus late information presented about the target, instructions provided before participants learned any information about *Bob* noted that first impressions are usually correct and rarely lead to errors (non-

bracketed version appearing below) or that first impressions are often incorrect and biased (bracketed version appearing below). Specifically, participants were told the following:

When forming your opinion about what type of person *Bob* is, you should [not] focus on your first impression of what *Bob* is like. First impressions are almost always [in]correct when forming an impression about a new person, and using your first impressions allows you to avoid [causes you to make] several biases (or errors in thinking) that are caused by thinking too much [not thinking enough] about what a person is like.

### Results

**Explicit attitudes.** The predicted three-way interaction between level of reinforcement by counterattitudinal condition by first impressions was found,  $F(2, 101) = 5.37, p < .005$ . Thus, level of reinforcement by counterattitudinal condition ANOVAs were conducted for explicit attitudes in the correct first impressions condition and in the incorrect first impressions condition separately. In the correct first impressions condition, the main effects of level of reinforcement and counterattitudinal condition were both significant,  $F(2, 101) = 7.96, p < .005$ , and  $F(2, 101) = 24.20, p < .001$ , respectively. The main effect of level of reinforcement showed that explicit attitudes were more positive in the 100% reinforcement condition ( $M = 0.49, SD = 0.67$ ) than in the 75% reinforcement condition ( $M = 0.12, SD = 0.39$ ). The main effect for counterattitudinal condition showed that explicit attitudes were more positive in the control condition ( $M = 0.84, SD = 0.47$ ) than in either the 20 CA ( $M = 0.01, SD = 0.40$ ) or 100 CA conditions ( $M = -0.02, SD = 0.34$ ). It is important that the interaction was also significant,  $F(2, 101) = 6.55, p < .005$ . As seen in Figure 5, participants in the 75% reinforcement condition showed more positive explicit attitudes toward *Bob* in the control condition ( $M = 0.44, SD = 0.31$ ) than in the 20 CA condition ( $M = -0.12, SD = 0.33$ ) and 100 CA condition ( $M = 0.02, SD = 0.27$ ),  $F(2, 101) = 8.44, p < .005$ . Participants in the 20 CA and 100 CA conditions did not differ. Participants in the 100% reinforcement condition showed more positive explicit attitudes toward *Bob* in the control condition ( $M = 1.20, SD = 0.23$ ) than in the 20 CA condition ( $M = 0.15, SD = 0.45$ ) and 100 CA condition




Figure 5. Explicit attitudes as a function of reinforcement and counterattitudinal conditions (20 CA, 100 CA) for the first impressions are the correct condition in Experiment 3.



( $M = -0.04$ ,  $SD = 0.39$ ),  $F(2, 101) = 34.56$ ,  $p < .001$ . In other words, replicating Experiments 1–2 and Kerpelman and Himmel-farb (1971), the effect of counterattitudinal information was stronger in the 100% reinforcement condition than in the 75% condition. However, when participants were explicitly told to rely on their first impression, any on-going drop in explicit attitudes in the 100 CA condition for those in the 75% reinforcement condition was not evidenced. As expected, adopting a “rely on first impressions” goal eliminated on-going attitude adjustments previously observed in the 100 CA condition for those receiving 75% reinforcement.

In the correct first impressions condition, the Level of Reinforcement  $\times$  Counterattitudinal Condition interaction was not significant,  $F(2, 101) = 2.50$ ,  $p > .10$ . Instead, the main effects of level of reinforcement and of the effect of counterattitudinal condition were both significant,  $F(2, 101) = 97.87$ ,  $p < .001$ , and,  $F(2, 101) = 22.54$ ,  $p < .001$ , respectively. As seen in Figure 6, participants in the 100% reinforcement condition had more positive explicit attitudes toward Bob ( $M = -0.08$ ,  $SD = 1.01$ ) than those in the 75% reinforcement condition ( $M = -0.56$ ,  $SD = 1.06$ ). Also, overall participants showed more positive explicit attitudes toward Bob in the control condition ( $M = 0.93$ ,  $SD = 0.62$ ) than in the 20 CA condition ( $M = -0.13$ ,  $SD = 0.66$ ), which were more positive than their attitudes in the 100 CA condition ( $M = -1.20$ ,  $SD = 0.62$ ). This stair-step pattern across counterattitudinal conditions indicates that, unlike the correct first impressions, explicit attitude change continued across the entire presentation of counterattitudinal information (regardless of level of initial reinforcement) and did not stop at the end of 20 counterattitudinal pieces of information.

**Implicit attitudes.** The three-way interaction for implicit attitudes was not significant ( $F < 1$ ). As Figure 7 reveals, the only significant effect to obtain was the predicted main effect of counterattitudinal condition,  $F(2, 101) = 19.89$ ,  $p < .001$ . Replicating Experiment 2, implicit attitudes were more positive (i.e., more strongly in the direction of initial learning) in the control condition ( $M = 0.39$ ,  $SD = 0.89$ ) and in the 20 CA condition ( $M = 0.28$ ,  $SD = 0.81$ ) than in the 100 CA condition ( $M = -0.75$ ,  $SD = 0.88$ ).

These results for explicit and implicit attitudes again showed that explicit attitudes were more quickly changed than were im-




Figure 6. Explicit attitudes as a function of reinforcement and counterattitudinal conditions (20 CA, 100 CA) for the first impressions are the incorrect condition in Experiment 3.




Figure 7. Implicit attitudes as a function of focus and counterattitudinal conditions (20 CA, 100 CA) in Experiment 3.

PLICIT attitudes when people encountered counterattitudinal information. The amount of explicit attitude change was not different between 20 CA and 100 CA when participants focused on forming first impressions. However, the increase in counterattitudinal information from 20 to 100 behaviors did lead to greater attitude change when participants were instructed not to rely on forming first impressions. Also, these results again showed that implicit attitudes changed more slowly and only when a sufficient amount of counterattitudinal information was encountered. And as expected, explicit impression formation goals did not affect implicit attitudes, although these verbal instructions had considerable impact on explicit attitudes.

## Discussion

Experiment 3 again found that explicit attitudes were changed by a different system of reasoning than were implicit attitudes. More specifically, explicit attitudes were altered by the introduction of conscious processing goals but implicit attitudes were not. Consistent with Experiment 2, implicit attitudes did change following the presentation of a substantial amount of counterattitudinal information. However, these implicit attitudes were not affected by verbal processing goals. In addition, the experimental manipulation of processing goals explained why differences in explicit attitudes, as a function of level of reinforcement, were found when a large amount of counterattitudinal information was presented in Experiment 2. It appears that less consistent reinforcement led participants to suspend early judgments of the target individual and attend to later information when forming their impression.

The preceding experiments provide evidence that implicit and explicit attitudes were formed and changed by slow-learning and fast-learning processes, respectively. However, what implications do these different attitudes have for behavior? Because an important function of attitudes is to predict behavior (Fazio, 1986), in Experiment 4 we sought to understand the relation between learned implicit and explicit attitudes and target-relevant behavior.

## Experiment 4

Recent work has begun to tease apart when implicit and explicit attitudes guide behavior. In general, this work has found that implicit attitudes predict subtle, spontaneous behavior, whereas explicit atti-

tudes predict more deliberative, intentional behavior (e.g., Jellison et al., 2004; McConnell & Leibold, 2001). We were interested in whether the attitudes toward *Bob* created in the current experiments could predict behavior in the same manner as past research. Specifically, would explicit attitudes toward *Bob* only predict deliberate judgments about him but not predict more subtle forms of behavior (i.e., seating distance)? Similarly, would implicit attitudes toward *Bob* only predict subtle behaviors but not explicit judgments about him? Experiment 4 tests these predictions, anticipating unique predictive value for implicit and explicit attitudes.

These findings could be important for at least three additional reasons. First, past research has shown such double dissociations on the basis of measures of group prejudice (e.g., Dovidio et al., 2002; Jellison et al., 2004); however, this would be the first time that such effects have been shown for a different type of attitude object (i.e., a target person). Second, this previous work has shown these outcomes for groups with preexisting attitudes, whereas this would be the first study to demonstrate such dissociation effects on the basis of attitudes engineered in a controlled laboratory setting. For example, it is possible that cultural prescriptions might shape both implicit prejudice and subtle forms of social behavior toward social group members, providing the appearance of an attitude-behavior relation when, in fact, other factors may produce both. By engineering attitudes in the laboratory without any other target-relevant knowledge, it is far more likely that behavior reflects the influence of attitudes directly. Finally, if we show that implicit attitudes have unique predictive utility for subtle behavior in this study, then the findings would argue against concerns that our implicit measure has poor sensitivity. One might argue that slow changes on our implicit measure may reflect a relatively weak measure (i.e., it is simply less responsive to change than our explicit measures) rather than a slow-learning system. By establishing that our implicit (but not explicit) attitude measure can uniquely predict theoretically derived types of behavior, we could provide evidence inconsistent with a position that our implicit attitude measure is simply a poor measure.

## Method

**Participants.** A sample of 29 undergraduates at Miami University participated in return for research credit in their introductory psychology courses. Participants were randomly assigned to receive no counterattitudinal information about *Bob* (control) or to receive 20 counterattitudinal pieces of information about *Bob* (20 CA).

**Procedure.** All materials, methods, and measures paralleled Experiment 1, with these exceptions. First, only the positive valence condition was used, and only the 100% reinforcement condition was used. The two experimental conditions (control and 20 CA) were selected to maximize the discrepancy between implicit and explicit attitudes. In Experiment 1 there was a drastic change in explicit attitudes between the control and the 20 CA conditions, however there was no difference in implicit attitudes between them. Additionally, as in Experiment 3, because there was no negative initial learning condition to reverse score, greater standardized measures of attitudes reflected more positive attitudes toward *Bob*.

In addition to the attitude measures, participants completed explicit judgments of desire for social contact with *Bob*. Specifically, participants rated the extent to which they would want to have *Bob* as a neighbor, friend, classmate, roommate, and family member, each on 100-point scales ( $\alpha = .92$ ). Greater scores on this measure indicated that they wanted more social contact with *Bob*.

After completing the attitude measures and the explicit social contact judgments, participants were told that they would "have a 2-min get acquainted session with *Bob*." They were escorted to a different room in which two chairs were set 221 cm apart. One chair had a book bag and a book next to it (where *Bob* was supposedly sitting), the other chair (for the participant) was on wheels and set against the wall of the room. The experimenter told each participant, "It looks like *Bob* has stepped out for a moment. Take that seat against the wall and move it so that you can have a face-to-face conversation with *Bob*." Participants took the seat and moved it into a position to converse with *Bob*. Afterward, they were told that they were not going to meet *Bob* and were then debriefed. The seating distance between the participant's chair and the chair where *Bob* had supposedly been sitting served as our measure of subtle, spontaneous behavior.

## Results

The attitude measures were examined with one-way ANOVAs of counterattitudinal condition. The only effect to obtain was the predicted effect of counterattitudinal condition for explicit attitudes,  $F(1, 27) = 12.86, p < .005$ . Replicating the findings of Experiment 1, explicit attitudes were more positive in the control condition ( $M = 0.48, SD = 0.80$ ) than in the 20 CA condition ( $M = -0.51, SD = 0.77$ ),  $F(1, 27) = 11.57, p < .005$ . In contrast, implicit attitude data did not show an effect of counterattitudinal condition ( $F < 1$ ).

The effect of counterattitudinal condition for social contact judgments was also examined with a one-way ANOVA. This analysis showed, as expected, that people reported wanting more social contact when they were in the control condition ( $M = 74.53, SD = 15.83$ ) than when they were in the 20 CA condition ( $M = 61.21, SD = 17.71$ ),  $F(1, 27) = 4.57, p < .05$ . Also, there was no effect of counterattitudinal condition on seating distance ( $F < 1$ ). Thus, the counterattitudinal condition manipulation affected deliberate behavior (i.e., desire for social contact) but not the subtle behavior (i.e., seating distance).

To examine the main hypotheses, the correlation between explicit attitudes, implicit attitudes, deliberate behavior (i.e., desire for social contact), and subtle behavior (i.e., seating distance) were calculated. As expected, more positive explicit attitudes were related to greater desire for social contact ( $r = .71, p < .001$ ) but were unrelated to seating distance ( $r = .04, ns$ ). It is important that more positive implicit attitudes were unrelated to desire for social contact ( $r = -.03, ns$ ) but were significantly related to closer seating distance ( $r = -.41, p < .03$ ). Moreover, two multiple regressions analyses were conducted in which explicit and implicit attitudes served to predict desire for social contact (first analysis) and seating distance (second analysis). As predicted, explicit attitudes ( $\beta = 0.70, p < .001$ ) but not implicit attitudes ( $\beta = -0.01, ns$ ) predicted desire for social contact. On the other hand, implicit attitudes ( $\beta = -0.41, p < .04$ ) but not explicit attitudes ( $\beta = 0.02, ns$ ) predicted seating distance. Thus, explicit attitudes uniquely predicted deliberate judgments and implicit attitudes uniquely predicted subtle, spontaneous behaviors.

## Discussion

Experiment 4 showed that the differential formation and change of implicit and explicit attitudes demonstrated in Experiments 1–3 have important implications for predicting behavior toward an attitude object, which in turn, reflect different systems of reason-

ing. As in the previous experiments, explicit attitudes were affected by the introduction of a small amount of counterattitudinal information, and in the current study, these attitudes uniquely predicted deliberate judgments of the target, whereas implicit attitudes did not. Conversely, implicit attitudes were unaffected by the presentation of a small amount of counterattitudinal information, and these implicit attitudes uniquely predicted spontaneous behaviors (i.e., seating distance) that explicit attitudes did not predict. Moreover, the current study provided clear evidence that the implicit measure is sensitive (i.e., it uniquely predicted subtle behavior), and it showed this double dissociation pattern of predicting behavior for the first time for an individual target and for an attitude object for which there were no preexisting beliefs. Indeed, coupling these results with those of Rydell et al. (in press), in which the same measures of implicit and explicit attitudes were opposite in valence when the valence of the subliminal prime and the behavioral information were inconsistent, strongly argues that the current findings (i.e., slow implicit attitude change) are not due to lack of sensitivity in the implicit measure.

### Experiment 5

Although Experiments 1–4 establish that implicit and explicit attitudes change at different rates in response to the same information, we more seriously consider whether the observed differential rate of implicit and explicit attitude change reflects a less sensitive implicit attitude measure. Thus, Experiment 5 sought to provide additional evidence showing that implicit, but not explicit, attitudes can be changed to further discredit this alternative account. Specifically, Experiment 5 sought to demonstrate that implicit, but not explicit, attitudes would be affected by nonconscious associations with the attitude object. In this study, two types of information about *Bob* that should be differentially attended to by a slow-learning system of reasoning (based on associations in memory) and a fast-learning system (based on higher order cognition and logic) were presented. Verbal behavioral information was presented, but unlike the previous experiments, all of this behavioral information was relatively neutral (which should produce relatively neutral explicit attitudes toward *Bob*). In addition, a valenced prime was presented subliminally before the presentation of *Bob*'s face (which should shape implicit attitudes toward *Bob*). A systems of reasoning explanation predicts that implicit attitudes should reflect the valence of the subliminal primes, whereas the fast-learning system should not be affected by the primes when other information is available. This pattern of results would provide compelling evidence that implicit attitudes are changed by associative processes and that the implicit measures in the current work are indeed sensitive.

### Method

**Participants.** A sample of 50 undergraduates at Miami University participated in return for research credit in their introductory psychology courses. Participants were randomly assigned to one of two conditions in which they were provided with all neutral behavioral information and were exposed to either positive primes first and then to negative primes later or to negative primes first and then to positive primes later.

**Procedure.** The materials and methods in this experiment were based on those used in Experiments 1–4 but differed in some key respects. Most notably, a subliminal prime was presented prior to *Bob*'s picture and the

behavioral information about *Bob*. Specifically, following a fixation point appearing in the center of the computer monitor for 200 ms, a positive or negative word was presented in the center of the computer monitor for 25 ms (serving as a prime). Next, participants saw a screen with only a picture of *Bob* for 250 ms, and finally, with the picture remaining on the screen, they were given neutral behavioral information that they judged to be characteristic or uncharacteristic of *Bob*. The valence of the primes that participants saw was varied systematically to be either unambiguously negative (e.g., death) or positive (e.g., love; Fazio, Sanbanmatsu, Powell, & Kardes, 1986). During the first 100 trials, half of the participants received 10 negative primes 10 times each, and the other half received 10 positive primes 10 times each. During the second 100 trials, the valence of the prime presented was switched such that those who had initially seen the positive primes now saw the negative primes and those who had seen the negative primes now saw the positive primes. Thus, overall all participants saw the same 20 primes (10 positive and 10 negative) 10 times each.

In another change from the previous experiments, participants' implicit and explicit attitudes were assessed at two different times during the session: after the first 100 trials (Time 1) and after the second set of 100 trials (Time 2). However, the attitude measures were identical at both times of assessment and paralleled those of Experiments 1–4 (in addition, these measures were counterbalanced at both times and this manipulation produced no effects). Finally, participants were given a recognition task for the positive and negative primes after the second assessment of attitudes. They were told that words were presented before *Bob*'s picture and that we were interested in their ability to detect them. To assess whether participants recognized the words that were flashed on the monitor, we gave them a list of 40 words presented alphabetically (20 actual words, 10 positive and 10 negative, and 20 filler words, 10 positive and 10 negative) from which they chose 20 that they believed could have been presented during the session.

### Results

The attitude measures were examined separately with 2 (prime order: negative prime first, positive prime first)  $\times$  2 (Time 1, Time 2) mixed-model ANOVAs. As expected, the two-way interaction between condition and time was not found for explicit attitudes,  $F(1, 48) = 1.66$ , *ns* (see Figure 8). Indeed, no effects were statistically significant for explicit attitudes. As Figure 9 reveals, a very different picture emerged for implicit attitudes. The predicted two-way interaction between condition and time was the only significant effect found for implicit attitudes,  $F(1, 48) = 10.02$ ,  $p < .005$ . Thus, simple effects analyses of time were conducted for implicit attitudes in each of the two between-subjects conditions.




Figure 8. Explicit attitudes as a function of condition and time in Experiment 5.




Figure 9. Implicit attitudes as a function of condition and time in Experiment 5.

In the negative prime first condition, implicit attitudes were more negative at Time 1 than at Time 2,  $F(1, 48) = 6.94, p < .03$  (Time 1:  $M = -0.85, SD = 1.13$ ; Time 2:  $M = 0.44, SD = 0.86$ ). In the positive prime first condition, implicit attitudes were more positive at Time 1 than at Time 2,  $F(1, 48) = 7.16, p < .02$  (Time 1:  $M = 0.85, SD = 0.67$ ; Time 2:  $M = -0.44, SD = 0.63$ ).

To ensure that participants did not recognize the words presented before *Bob's* picture in the learning task (i.e., to establish that the primes were subliminal), we assessed the mean accuracy for their identifying which 20 of the 40 words they thought were presented. As intended, participants were no better than chance ( $M = 0.51, SD = 0.07$ ) at recognizing the primes,  $t(49) = 1.01, ns$ , even though each prime was presented 10 times, indicating that the presentation of the primes was indeed subliminal.

## Discussion

Experiment 5 demonstrated that implicit, but not explicit, attitudes were sensitive to the subliminal presentation of the priming words. This is consistent with a slow-learning system that is based on the slow accrual of associations encountered across time and a fast-learning system that is based on mentally manipulating conscious information. Moreover, the current study provides additional evidence, supplementing that of Experiment 4 and Rydell et al. (in press), that the results involving implicit attitudes in Experiments 1–3 were not due to an insensitive IAT, but instead, reflect associations with *Bob* in memory.

## General Discussion

This research indicates that implicit and explicit attitudes change because they are governed by different systems of reasoning. Explicit attitudes changed more quickly in response to new information and were responsive to deliberate processing goals, consistent with a quick-learning, rule-based system of reasoning. Alternatively, implicit attitudes changed much more slowly and were unaffected by processing goals, consistent with a slow-learning, associative system of reasoning.

An old paradigm was modified in the current work that allows perceivers to form an initial attitude about an attitude object and that also allows for the presentation of counterattitudinal information

about that same attitude object. Consistent with distinct systems of reasoning (Sloman, 1996; Smith & DeCoster, 2000) and dual-attitude approaches (Wilson et al., 2000), these experiments show that explicit attitudes quickly changed following only a small amount of counterattitudinal information, whereas implicit attitudes about the same attitude object did not change in response to the same counterattitudinal information. Thus, counterattitudinal information, like that used in many persuasion paradigms (Petty & Wegener, 1998), may not erase the initial (implicit) attitude (see Petty et al., 2006; Wilson et al., 2000). However, consistent with a slow-learning system of reasoning, Experiments 2 and 3 show that implicit attitudes were changed when sufficient counterattitudinal information was presented. In addition, this work shows that people can hold different implicit and explicit attitudes about the same attitude object at the exact same time based on how the information they encountered impacts different systems of reasoning.

The results from Experiments 2 and 3 were clear in elucidating that implicit attitudes change by a slow-learning, associative system of reasoning. Specifically, by presenting substantial amounts of counterattitudinal information, implicit attitudes changed to reflect the accrual of copious amounts of counterattitudinal information. On the other hand, explicit attitudes showed a different pattern. That is, explicit attitude change was best explained by a fast-learning, rule-based system of reasoning. Experiment 1 shows that explicit attitudes changed following a relatively small amount of counterattitudinal information. Consistent with past research, Experiments 1 and 2 also show the greatest amount of attitude change following greater consistency of initial reinforcement (e.g., Kerpelman & Himmelfarb, 1971) and when negative counterattitudinal information followed positive information (e.g., Fiske, 1980). Interestingly, implicit attitudes did not reveal these results, indicating that classic attitude asymmetry effects may be more likely for explicit than implicit attitudes.

Moreover, the manipulation of reinforcement level also produced an interesting pattern with explicit attitude change. In Experiment 2, people who had less consistent initial learning showed continued attitude change as more counterattitudinal information was presented. Experiment 3 experimentally established that this pattern resulted from participants not rushing to form strong on line impressions following less consistent feedback (cf., McConnell et al., 1994b). More important, this experiment showed that manipulating impression formation goals (a deliberate process) changed explicit attitudes yet did not change implicit attitudes, further supporting a systems of reasoning explanation of implicit and explicit attitude change.

Experiment 4 shows that implicit and explicit attitudes developed in the laboratory predicted different types of behaviors. Consistent with past research (e.g., McConnell & Leibold, 2001), implicit attitudes predicted subtle, spontaneous behaviors toward *Bob* (i.e., seating distance) but not deliberative judgments toward him (i.e., desire for social contact). Conversely, explicit attitudes predicted deliberate judgments about *Bob* but not more subtle, spontaneous behaviors. This experiment shows that attitudes created and changed by different systems of reasoning have important implications for when attitudes correspond to behavior. Implicit attitudes only predicted spontaneous behaviors, whereas explicit attitudes only predicted deliberate target-relevant judgments; this double dissociation further supports a systems of reasoning account.



Finally, Experiment 5 shows that, consistent with a systems of reasoning prediction, implicit (but not explicit) attitudes were changed by counterattitudinal information that was associated with the attitude object subliminally. Explicit attitudes, on the other hand, reflected the neutral information that was consciously available about *Bob*. Whereas Experiments 1–4 showed that explicit attitudes changed more quickly than implicit attitudes, Experiment 5 revealed that implicit attitudes would change even when explicit attitudes did not because of the type of information available to each system of reasoning.

However, the results of Experiment 5 raise other important questions. First, why did we observe only implicit attitude change when other research (e.g., Murphy & Zajonc, 1993) has shown that explicit attitudes change in response to valenced primes? We believe that providing participants with a series of neutral behaviors occupied the fast-changing, verbal-based system with valence irrelevant information, which in turn, led to the expression of relatively neutral attitudes. Indeed, past subliminal priming research has not presented supraliminal behaviors in tandem that might engage the fast-verbal system. Another alternative is that the presentation of many neutral behaviors may have led participants to conclude that they should ignore or discount any affect generated by the subliminal primes, leading them to not use these feelings in their explicit judgments (Yzerbyt, Schandron, Leyens, & Rocher, 1994). However, it is unclear whether people would make such an attribution for explicit attitudes when the supraliminal information is neutral. This is not to argue that people do not use meta-informational cues such as “social judgability” (Yzerbyt et al., 1994), but it is not established whether people provided with neutral behaviors feel “unentitled” to render evaluations of a target. If people did feel unentitled, one would expect low-variance judgments around the midpoint. But given the relatively large variability in the current data, it seems more likely that people viewed the behaviors with some idiosyncratic degree of positivity and negativity rather than circling the midpoint because they felt they could not render a judgment. Certainly, future research should address this interesting possibility more directly.

Also related to this issue are the data of Rydell et al. (in press), who used the same methods as the current Experiment 5 but presented supraliminal behaviors that were always of the opposite valence to the subliminal primes (e.g., positive subliminal primes with negative supraliminal behaviors). As predicted by a systems of reasoning perspective, explicit attitudes toward *Bob* mirrored the valence of the supraliminal information, and implicit attitudes toward *Bob* reflected the valence of the subliminal information. In this work, a social judgability alternative seems untenable because the verbally available information about the target should seem coherent. Yet at the same time, explicit attitudes toward the target person were radically different than the implicit attitudes toward the target person, indicating that feelings from the subliminal presentations did not “spill over” on their explicit attitudes in a substantial manner.

More broadly, the current work has important implications for existing models of attitudes and persuasion. Indeed, it provides some of the clearest support for two of the dual attitudes model’s most important suppositions: People can hold different implicit attitudes and explicit attitudes about an attitude object at the same time, and implicit attitudes are not changed at the same rate as explicit attitudes (Wilson et al., 2000). This last point is extremely

important for understanding how the current research relates to other research on attitude change, specifically research on the equation likelihood model (ELM) of persuasion (Petty & Wegener, 1998). The ELM predicts that once an attitude is changed, usually by the presentation of compelling arguments or by a peripheral cue to persuasion (e.g., attractiveness), the original attitude no longer exists. Although this model is extremely powerful in predicting explicit attitude change and deliberate behavior toward an attitude object, it may not account for how implicit attitudes change (e.g., Petty et al., 2006). In addition, it does not account for the possibility that implicit attitudes and explicit attitudes predict different types of behavior (e.g., Dovidio et al., 2002; Jellison et al., 2004; McConnell & Leibold, 2001). Thus, models such as the ELM should consider how implicit attitudes change and how they guide behavior in order to provide a fuller account of persuasion.

Also, the motivation and opportunity to deliberate (MODE) model argues that differences between implicit and explicit attitudes are evidence that people differ in the extent to which they have the motivation and ability to modify the initial automatic activation of an attitude in memory (e.g., Fazio, 1995). Although this account is undoubtedly true in many circumstances, it may not capture the relation between implicit and explicit attitudes in all situations, especially those involving novel attitude objects and in situations in which attitude accessibility is low. In addition, this account of attitudes has difficulty explaining how implicit and explicit attitudes can differ at the exact same point in time devoid of some motivation to modify the expression of the explicit attitude (e.g., evaluating *Bob* carries far less social desirability concerns than expressing racial attitudes). However, a systems of reasoning approach predicts this outcome and fits the data obtained in the current work nicely.

In summary, the current work shows that implicit attitudes and explicit attitudes form and change on the basis of different processes that support a systems of reasoning approach to attitude change. Understanding attitude change, and more specifically the different processes underlying implicit and explicit attitude change, is extremely important for advancing theoretical conceptualizations of attitude formation and attitude change. The interplay between implicit and explicit attitudes is extremely important for diverse areas of social psychology (e.g., attitude formation, persuasion, prejudice, attitude–behavior correspondence), and the current research begins to disentangle the differences in how implicit attitudes and explicit attitudes form and respond to social information.

## References

- Barden, J., Maddux, W. W., Petty, R. E., & Brewer, M. B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality and Social Psychology*, 87, 5–22.
- Bassili, J. N., & Smith, M. C. (1986). On the spontaneity of trait attribution: Converging evidence for the role of cognitive strategy. *Journal of Personality and Social Psychology*, 50, 239–245.
- Carlston, D. (1994). Associated systems theory: A systematic approach to cognitive representations of persons. In R. S. Wyer, Jr., (Ed.), *Advances in social cognition* (Vol. 7, pp. 1–78). Hillsdale, NJ: Erlbaum.
- Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology*, 81, 800–814.

- Devine, P. G. (2001). Implicit prejudice and stereotyping: How automatic are they? *Journal of Personality and Social Psychology*, 81, 757-759.
- Dovidio, J. F., Kawakami, K., & Gaertner, S. L. (2002). Implicit and explicit prejudice and interracial interaction. *Journal of Personality and Social Psychology*, 82, 62-68.
- Dunton, B. C., & Fazio, R. H. (1997). An individual difference measure of motivation to control prejudiced reactions. *Personality and Social Psychology Bulletin*, 23, 316-326.
- Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Orlando, FL: Harcourt, Brace & Jovanovich.
- Fazio, R. H. (1990). A practical guide to the use of response latency in social psychological research. In C. Hendrick & M. S. Clark (Eds.), *Research methods in personality and social psychology* (pp. 74-97). Thousand Oaks, CA: Sage.
- 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: Erlbaum.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and uses. *Annual Review of Psychology*, 54, 297-327.
- Fazio, R. H., Sonbanmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50, 229-238.
- Fiske, S. T. (1980). Attention and weight in person perception: The impact of negative and extreme behavior. *Journal of Personality and Social Psychology*, 38, 889-906.
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, 132, 692-731.
- Gawronski, B., & Strack, F. (2004). On the propositional nature of cognitive consistency: Dissonance changes explicit but not implicit attitudes. *Journal of Experimental Social Psychology*, 40, 535-542.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, 74, 1464-1480.
- Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85, 197-216.
- Gregg, A. P., Seibt, B., & Banaji, M. H. (2006). Easier done than undone: Asymmetries in the malleability of implicit preferences. *Journal of Personality and Social Psychology*, 90, 1-20.
- Jellison, W. A., McConnell, A. R., & Gabriel, S. (2004). Implicit and explicit measures of sexual orientation attitudes: In-group preferences and related behaviors and beliefs among gay and straight men. *Personality and Social Psychology Bulletin*, 30, 629-642.
- Kerpelman, J. A., & Himmelfarb, S. (1971). Partial reinforcement effects in attitude acquisition and counterconditioning. *Journal of Personality and Social Psychology*, 19, 301-305.
- McConnell, A. R. (2001). Implicit theories: Consequences for social judgments of individuals. *Journal of Experimental Social Psychology*, 37, 215-227.
- McConnell, A. R., & Leibold, J. M. (2001). Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology*, 37, 435-442.
- McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1994a). Illusory correlations in the perception of groups: An extension of the distinctiveness-based account. *Journal of Personality and Social Psychology*, 67, 414-429.
- McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1994b). On-line and memory-based aspects of individual and group target judgments. *Journal of Personality and Social Psychology*, 67, 173-185.
- McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1997). Target entitativity: Implications for information processing about individual and group targets. *Journal of Personality and Social Psychology*, 72, 750-762.
- Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit evaluations. *Journal of Experimental Psychology: General*, 123, 455-469.
- Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology*, 64, 723-739.
- Olson, M. A., & Fazio, R. H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science*, 12, 413-417.
- Petty, R. E., Tormala, Z. L., Briñol, P., & Jarvis, W. B. G. (2006). Implicit ambivalence from attitude change: An exploration of the PAST model. *Journal of Personality and Social Psychology*, 90, 21-41.
- Petty, R. E., & Wegener, D. T. (1998). Attitude change: Multiple roles for persuasion variables. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., Vol. 1, pp. 323-390). New York: McGraw-Hill.
- 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.
- Richeson, J. A., & Trawalter, S. (2005). On the categorization of admired and disliked exemplars of admired and disliked racial groups. *Journal of Personality and Social Psychology*, 89, 517-530.
- Rydell, R. J., McConnell, A. R., Mackie, D. M., & Strain, L. M. (in press). Of two minds: Forming and changing valence inconsistent implicit and explicit attitudes. *Psychological Science*.
- 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.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3-22.
- Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review*, 4, 108-131.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, 8, 220-247.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107, 101-126.
- Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, 81, 815-827.
- Yzerbyt, V. Y., Schandron, G., Leyens, J. P., & Rocher, S. (1994). Social judgability: The impact of meta-informational cues on the use of stereotypes. *Journal of Personality and Social Psychology*, 66, 48-55.

Received June 17, 2005

Revision received February 27, 2006

Accepted March 12, 2006 ■