

Prejudice and Perception: The Role of Automatic and Controlled Processes in Misperceiving a Weapon

B. Keith Payne
Washington University

Two experiments used a priming paradigm to investigate the influence of racial cues on the perceptual identification of weapons. In Experiment 1, participants identified guns faster when primed with Black faces compared with White faces. In Experiment 2, participants were required to respond quickly, causing the racial bias to shift from reaction time to accuracy. Participants misidentified tools as guns more often when primed with a Black face than with a White face. L. L. Jacoby's (1991) process dissociation procedure was applied to demonstrate that racial primes influenced automatic (A) processing, but not controlled (C) processing. The response deadline reduced the C estimate but not the A estimate. The motivation to control prejudice moderated the relationship between explicit prejudice and automatic bias. Implications are discussed on applied and theoretical levels.

In February 1999, four White New York Police officers shot and killed Amidou Diallo, an unarmed Black immigrant from West Africa, in a hail of 41 bullets (McFadden & Roane, 1999). The controversy surrounding that incident sparked protests from civil rights groups across the nation, fueling charges of *racial profiling*, the practice of considering race as a factor when police officers stop citizens. Critics alleged that racial bias played a role in the confrontation. The officers were acquitted of all charges on the grounds that although they made a mistake, their actions were justified at the time. The shooting was judged to be justified because at the moment that police officers ordered Diallo to stop, the victim moved, producing an object that later turned out to be a wallet. The police defendants contended that in this ambiguous situation, they acted on the information available, sincerely believing that they were in danger (Fritsch, 2000).

This case is interesting to psychologists, not because of its legal or ethical implications, but because of the psychological processes that it dramatically highlights. Research on the relationship between automatic and controlled cognition has recently made a strong impact on social psychology (e.g., Devine, 1989; Fazio, 1990a; Greenwald & Banaji, 1995; Wilson, Lindsey, & Schooler, 2000). One intriguing implication of this work is that both the critics and the defendants in the Diallo case could be right: It is possible that racial bias plays a role in such situations, but that the

individuals involved sincerely believe that their judgments are accurate. Specifically, several lines of research have shown that group stereotypes may be activated outside of awareness and may influence behavior without the knowledge or intent of the perceiver (e.g., Devine, 1989; Wittenbrink, Judd, & Park, 1997). Therefore, one can ask at least two separate questions about cases such as this one. One is the legally crucial question: What was the conscious intent of the officers? This is a question about the subjective psychological state of the perceiver. The second question is the more psychologically important one: What are the causes of that psychological state, the cognitive processes, and the behaviors that accompany it?

The purpose of the present research was to integrate insights and techniques from social and cognitive psychology to help understand these important issues. First, I adopt Jacoby's (1991) process dissociation procedure (PDP) to investigate the influence of racial cues on the visual identification of weapons within a priming paradigm. The PDP estimates the distinct contributions from automatic bias and controlled perception to task performance. This procedure has been applied successfully in the domain of memory (e.g., Hay & Jacoby, 1999; Jacoby, 1991, 1999; Jacoby, Toth, & Yonelinas, 1993), and it affords a new and theoretically interesting approach to distinguishing the automatic and controlled processes at work when group stereotypes are activated. The second goal of this research was to examine how automatic bias and controlled perceptual processing relate to explicit racial attitudes and motivations.

Priming research has shown that primes related to stereotyped group members tend to facilitate responses to negatively valenced (Gaertner & McLaughlin, 1983; Fazio, Jackson, Dunton, & Williams, 1995) and stereotypical target words (Banaji & Hardin, 1996; Wittenbrink et al., 1997). The typical finding with respect to racial groups is that White participants are faster to respond to

I thank Mahzarin Banaji for making available the pictures used as prime stimuli. Thanks also are due to Larry Jacoby, Alan Lambert, and Michael Strube for their helpful comments on drafts of this article and to Tyra Nelson for help in collecting data for this project.

Correspondence concerning this article should be addressed to B. Keith Payne, Department of Psychology, Washington University, 1 Brookings Drive, Box 1125, St. Louis, Missouri 63130. Electronic mail may be sent to bkpayne@artsci.wustl.edu.

positive target words paired with White primes and negative target words paired with Black primes, compared with Black–positive or White–negative pairs.

Racial primes may affect perceivers' responses in other ways besides producing differences in response latencies. Draine and Greenwald (1998) have argued that respondents may use at least two strategies when completing a priming task. They may adopt a certainty criterion, by which they wait to respond to a target until they are confident that their response is correct. Alternatively, the respondent may adopt a speed criterion, by which he or she attempts to respond very quickly, certain or not. For respondents using this latter strategy, primes influence the number of errors the respondent makes. Draine and Greenwald have demonstrated that when participants are required to respond within a limited time window, priming causes an increase in semantically consistent errors.

Taken together, the evidence from racial-priming paradigms and prime-induced errors suggests that when forced to respond rapidly, racial cues may cause perceivers to make stereotype-consistent errors. Consider the situation alluded to earlier, in which such stereotypic errors are important: when law enforcement officers identify weapons. Violent traits such as hostility, aggression, and criminality are consistently included in White Americans' stereotypes of Black Americans (Devine & Elliot, 1995; Dovidio, Evans, & Tyler, 1986). To the extent that guns are semantically associated with violence and aggression, the race of a suspect may influence White Americans' judgments of what is, and what is not, a weapon.

The present research tests whether priming with Black versus White faces biases the reaction time (RT; Experiment 1) or accuracy (Experiment 2) of identifying a weapon. However, Study 2 is not only an extension of Study 1 from RT to errors, it is aimed at linking social–cognitive theory and measurement procedures to an important domain of performance in which accuracy is critically important. I use Jacoby's (1991) PDP to estimate the separate contributions of automatic and controlled perceptual processing to people's performance. Finally, I investigate the correlations among explicit racial attitudes, motivation to control prejudice, and racial bias in identifying weapons. The strength of the present paradigm is that it estimates both automatic and controlled processes within a single task. These estimates are then used to clarify the mechanisms that produce a bias with important theoretical and practical consequences.

Automatic and Controlled Influences in Social Judgments

Researchers have taken several approaches in studying the automatic and controlled processes that jointly contribute to social judgments. Researchers during the 1970s first specified several criteria that were useful in distinguishing automatic from controlled processes (Posner & Snyder, 1975; Shiffrin & Schneider, 1977). Automatic processes operate outside of an individual's awareness and begin without conscious intent. Once begun, a participant cannot interrupt an automatic process. Finally, automatic processes are efficient, in the sense that they operate very quickly and do not compete with other operations for limited attentional resources. Controlled processes were defined by the opposites of these standards: They are conscious, intentional,

controllable, and are executed by a limited-capacity attentional system.

Contemporary and historical models of stereotyping have often held that stereotypes are activated automatically on exposure to a member of the stereotyped group (Allport, 1954; Brewer, 1988; Devine, 1989; Fiske & Neuberg, 1990; but see Gilbert & Hixon, 1991; Macrae, Bodenhausen, Milne, Thorne, & Castelli, 1997). Psychologists researching automatic and controlled processes in social judgments and attitudes have relied heavily on the *stimulus onset asynchrony* (SOA) to distinguish automatic from controlled processing in priming paradigms. The SOA is the amount of time that elapses between the onset of a prime and the beginning of the target stimulus. The length of the SOA is thought to determine whether the perceiver has the opportunity to override the automatic processing of the prime with intentionally controlled responses. Neely (1977) demonstrated that when participants had a conscious strategy that opposed the semantic relationship between a prime and target, they were able to respond in accordance with their conscious strategy only when the SOA was longer than 500 ms. At shorter SOAs, the semantic content of the prime influenced responses despite participants' intentional strategies. Priming effects that take place at short SOAs have often been considered relatively automatic for two reasons. First, because they proceed quickly, they are considered relatively efficient. Second, Neely's (1977) finding that priming effects occur despite conscious intentions to the contrary at short SOAs suggests that they are unintentional and uncontrollable.

However, it is not clear that any particular SOA (e.g., 500 ms.) is a sufficient criterion for discriminating between automatic and controlled processes. Several recent experiments have yielded results suggesting that processes seemingly under intentional control may influence responses at SOAs shorter than 500 ms. Blair and Banaji (1996) investigated the activation of gender stereotypes. Participants with a counter-stereotypic strategy were able to use it completely at a long (2,000 ms) SOA, responding faster to stereotype-incongruent pairs than to congruent pairs. However, at a short (250 ms) SOA, participants with a counter-stereotypic strategy showed no difference in RTs to congruent versus incongruent pairs, thus eliminating any priming effect. One interpretation of these results is that people's expectancies can counteract stereotypes at the automatic level. Another interpretation is that respondents were able to apply controlled strategies in this task even with a short SOA.

A series of experiments by Glaser and Banaji (1999) raised a similar issue. They investigated evaluative priming, using sequential primes, with SOAs from 150 ms to 300 ms. Across six experiments, they found that moderately valenced primes facilitated responses to targets of the same valence. However, extremely valenced primes produced a *reverse* priming effect. That is, responses were slower when an extremely positive prime was paired with a positive target than when the same prime was paired with a negative target. Glaser and Banaji interpreted these results as evidence that participants were "correcting" their judgments. Such corrective processes have typically been considered controlled and effortful (Martin, Seta, & Crelia, 1990; Wilson & Brekke, 1994). Glaser and Banaji suggested that judgmental correction may take place automatically. An alternative interpretation is that judgmental correction is a controlled process, but that the time required for participants to execute such control over their responses may

depend on specific features of the task, as well as the goals and strategies of the perceiver. Although Glaser and Banaji's procedure used a short SOA, the corrective processes they described appear to have important characteristics of control, including intention and mutability. Because a neat dichotomy between automatic and controlled processes is difficult to establish by using a particular SOA, it is important to specify the properties of automaticity or control on which one is focusing (see Bargh, 1989, 1994). The approach taken here has been to integrate evidence that converges from different perspectives. The present experiments combined short SOAs with Jacoby's (1991) PDP to examine automatic and controlled influences of racial primes on performance.

The Process Dissociation Approach

Memory researchers have grappled with parallel issues, trying to understand the contribution of implicit and explicit memory processes to performance on memory tests. Performance on implicit tasks has been dissociated from explicit memory, as measured by traditional direct memory tests, such as recall or recognition (for a review, see Roediger & McDermott, 1993). Jacoby and colleagues (Jacoby, 1991; Jacoby et al., 1993) have argued that the practice of identifying separate processes with different kinds of tests is problematic because on one hand, performance on indirect tests may be contaminated by consciously controlled recollection. On the other hand, performance on direct memory tests may also be contaminated by more automatic uses of memory.

An alternative approach has been to arrange experiments in which automatic and controlled processes are placed in opposition to one another (Jacoby & Dallas, 1981; Jacoby et al., 1993). To estimate the contributions of automatic and controlled processes to task performance, experiments must include both *congruent* conditions, in which they act in concert, and *incongruent* conditions, in which they oppose one another. Intentional control is measured as the difference between performance when a person intends to respond a certain way, and performance when the person intends not to respond in that way. To the extent that people can produce a particular response when they intend to, but not produce that response when they intend not to, they are exercising control. Control can be estimated from performance in congruent and incongruent conditions by using a set of simple algebraic equations (Jacoby, 1991). Consider the example of the police officer mentioned before to illustrate this procedure.

When a Black suspect possesses a gun, the officer is faced with a congruent condition. Both controlled perceptual processing (accurate identification of the gun) and automatic processing (stereotypic associations between Blacks and guns) should lead the officer to identify the object as a gun. The probability of responding "gun" on a congruent trial is the probability of control, C , plus the probability of an automatic association between the Black suspect and guns, when control fails, $A(1 - C)$:

$$\text{Congruent} = C + A(1 - C). \quad (1)$$

In situations where a Black suspect has an object that is not a gun, the officer is faced with an incongruent condition. The probability that the officer will respond "gun" is the probability that an automatic association favors the "gun" response, A , when there is a failure to properly control the response, $(1 - C)$:

$$\text{Incongruent} = A(1 - C). \quad (2)$$

Given these two equations, automatic and controlled components can be estimated separately. The estimate of controlled responding is the difference between responding "gun" on congruent and incongruent trials:

$$C = \text{Congruent} - \text{Incongruent}. \quad (3)$$

Having developed an estimate of control, the automatic estimate can be solved:

$$A = \text{Incongruent}/(1 - C). \quad (4)$$

The PDP defines automaticity by the relationships between performance and intentions. Automatic processes are those that operate regardless of whether they facilitate intentional performance or interfere with it. In contrast, controlled processes are those in which responses are successfully modulated by intentions. Within this procedure then, the A estimate reflects an automatic bias in which responses are systematically influenced by the race of the prime. The C estimate reflects correct perceptual processing of the target, that is, visual discrimination between guns and lures. An assumption of the PDP is that automatic and controlled processes are two independent bases for responding (for discussions of this assumption, see Curran & Hintzman, 1997; Hintzman & Curran, 1997; Jacoby, Begg, & Toth, 1997; Jacoby & Shrout, 1997). I discuss evidence for the validity of this assumption within the present paradigm and its implications in the General Discussion section.

Experiment 1

Overview

Experiment 1 used a priming paradigm to test whether pairing target stimuli with Black versus White faces biased participants' identification of those targets as weapons. Participants first completed a set of questionnaire materials, two of which are relevant to this study: the Modern Racism Scale (McConahay, Hardee, & Batts, 1981) and the Motivation to Control Prejudiced Reactions Scale (Dunton & Fazio, 1997). After completing these explicit attitude measures, participants performed the computerized priming task. In the priming task, digital photographs of White and Black male faces were used as primes, followed by targets that were either handguns or hand tools. Tools were selected as filler items because they are evaluatively neutral, and they were similar in size to handguns. The design of the study was a 2 (prime race: Black vs. White) \times 2 (target type: gun vs. tool) factorial design, with both factors manipulated within participants.

The first hypothesis was that participants would respond faster to guns when they were primed by a Black face compared with a White face. Second, correct responses and errors in the congruent and incongruent conditions served as a basis for deriving estimates of automatic and controlled influences, using the equations outlined above. In line with previous research, showing that stereotype activation is often automatic, it was expected that the racial prime would exert its influence in the automatic estimate, while leaving the controlled estimate unaffected. By experimentally dissociating the A and C components within the same procedure, this paradigm provides the important advantage of studying each com-

ponent separately, without confounding cognitive processes with task requirements. Third, I expected to conceptually replicate the results of Fazio et al. (1995). Using their "bona fide pipeline" priming procedure, these researchers found that modern racism scores and racial bias on the priming task were positively correlated only for individuals low in the motivation to control prejudice. Individuals high in the motivation to control prejudice displayed a negative correlation between explicitly reported racial attitudes and prejudice scores derived from the priming measure.

Fazio and colleagues (Fazio et al., 1995) interpreted this pattern of results as evidence that individuals can present themselves as unprejudiced on self-report measures when they are motivated to do so. Fazio et al. (1995) argued that the negative correlation between self-reported prejudice and the unobtrusive measure among the highly motivated reflected an overcompensation. Those participants who held relatively prejudiced attitudes, but were motivated to control them, completed the self-report questionnaire so as to appear very unprejudiced. Consistent with these results, I expected the motivation to control prejudice to moderate the relationship between explicit racial attitudes and racial bias in RT, with a positive correlation between explicit racial attitudes and RT bias only for those who were low in the motivation to control their prejudiced reactions.

Method

Participants

Thirty-one undergraduates (24 women, 7 men) participated in return for course credit. None of the participants was Black.

Stimuli

Photographs of two Black male and two White male faces were used as primes. The photos were selected from those used by Greenwald and Banaji (1995) for their Internet-based implicit association test.¹ They were black and white images presented at 5.3 cm × 4 cm in size. Each face wore a neutral expression, and the picture was cropped so that peripheral features (e.g., hair, clothes) were not visible. The photographs were chosen so that the only feature that varied systematically was race. Target stimuli included four photographs of handguns and four photographs of hand tools, each the same size as the primes. The tools included two kinds of pliers, one socket wrench, and an electric drill. Figure 1 displays examples of the stimuli used. The visual mask consisted of a rectangular pattern in the same size as the primes and targets. The pattern was irregularly covered with white and black color.

Procedure

Explicit measures. Participants were told that they would participate in two unrelated experiments. First, participants completed two explicit questionnaire measures, along with filler measures assessing general social attitudes and cognitive style. The Modern Racism Scale (MRS; McConahay et al., 1981) measured explicit racial prejudice. The MRS consists of seven items intended to measure subtle racism. Participants also completed the Motivation to Control Prejudiced Reactions Scale (MCP; Dunton & Fazio, 1997). This 17-item scale was developed to measure the extent to which individuals feel it is important not to experience or express prejudiced responses. Sample items include "I get angry with myself when I have a thought or feeling that might be considered prejudiced," and "It's never acceptable to express one's prejudices." Both measures used a




Figure 1. Examples of prime and target stimuli.

9-point Likert style response scale ranging from -4 (*strongly disagree*) to $+4$ (*strongly agree*).

Priming task. After the experimenter informed participants that they were finished with the first study, she asked them to perform a computerized task. The experimenter explained that the task tested speed and accuracy. The experimenter told participants that they would see pairs of pictures flashed briefly on the monitor. They were instructed to do nothing with the first picture, which would always be a face. It was explained that the face would signal that the target picture was about to be presented. They were instructed to respond to the second picture, which would always be either a gun or a tool. The participants' task was to classify each target object as either a gun or a tool by pressing one of two keys. The experimenter instructed participants that "You have to respond as quickly and accurately as you can. If you make a mistake, don't worry. Just keep going to the next trial. The first round of pictures is a practice trial." Before the active trials began, participants received 48 practice trials to become acquainted with the targets and practice classifying them quickly. During these practice trials, no primes appeared. Participants simply learned to classify the target objects by using a keypress.

Once the critical trials began, the priming task exposed participants to pairs of pictures. The first picture (the prime) was always a White or Black face. The second picture (the target) was always either a handgun or a hand tool. The prime remained on the screen for 200 ms and then was replaced immediately by the target. Thus, the SOA was 200 ms. After the target was presented for 200 ms, it was replaced by the visual mask. The mask remained on the screen until the participant responded. Response latencies were then recorded to the nearest millisecond, from the onset of the target stimulus. For each trial, the next prime appeared 500 ms after the previous response. Following the practice trials, participants completed 192 critical trials. Prime-target pairs were presented in a random order determined by the computer program. After participants completed the priming task, they were thoroughly debriefed and then dismissed.

Results

To determine whether the racial primes affected the perception of weapons, I first examine the response latencies with which participants identified weapons and tools within each priming condition. The top panel of Table 1 reports the mean RTs for guns and tools as a function of prime race. Second, I examine the effect of the primes on error rates. Next, I address the role of automatic

¹ See <http://buster.cs.yale.edu/implicit/> for more information.

Table 1
Mean Reaction Times (in Milliseconds) in Identifying Guns and Tools in Experiments 1 and 2

Target	Prime			
	Black		White	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Experiment 1				
Gun	423	64	441	73
Tool	454	57	446	60
Experiment 2				
Gun	299	28	295	31
Tool	307	29	304	29

and controlled processes by investigating the impact of the primes on automatic and controlled estimates separately. Finally, I explore the relationships between explicit attitudes, motivations, and perceptual task performance.

RTs

Because outliers can distort RT measures (Fazio, 1990b), a priori cut-off standards were adopted at 100 ms and 1,000 ms. Reaction times lying outside these limits were excluded from analyses. Because these criteria are beyond 3 *SDs* from the overall mean, less than 2% of the data were trimmed. Also dropped were RTs for incorrect responses. After trimming outliers, a log transformation was performed to reduce the positive skew of the RT distribution. Supplemental analyses performed on untransformed data produced results that were highly similar. For ease of interpretation, the raw RTs are reported in milliseconds.

To test the hypotheses, I computed the mean RT for each prime–target combination and performed a repeated measures analysis of variance (ANOVA). The analysis revealed a main effect of target, $F(1, 30) = 7.88, p < .009$, indicating that participants identified guns more quickly than tools. That effect was qualified by the predicted Prime \times Target interaction, $F(1, 30) = 16.45, p < .0003$. Simple effects tests revealed that participants identified guns faster when they were primed by a Black face than by a White face, $F(1, 30) = 13.46, p < .001$. In addition, participants identified tools more quickly when primed with a White face, compared to a Black face, $F(1, 30) = 6.13, p < .02$. Thus, the race of priming stimuli did affect the identification of weapons: The presence of Black faces facilitated the identification of guns relative to the presence of White faces.

Error Rates

Error rates in this experiment were very low overall ($M = 6\%$). A repeated measures ANOVA revealed no significant effects of prime or target conditions (all $F_s < 1.33, ns$). The top panel of Table 2 presents mean error rates for each condition. Note that participants in this study were allowed unlimited time to respond. As such, they may have used a certainty criterion, in which they waited to respond until they were relatively confident that their response was correct. With unlimited time to respond, the racial

prime exerted its effect by reducing the time required to reach that threshold.

Automatic and Controlled Estimates

The RT results indicated that priming with faces of different races did indeed influence participants' visual identification of weapons. Because this effect occurred at a relatively short SOA, it could be argued that the bias reflects an automatic process. However, the difficulties in relying on SOA as a criterion for automaticity make converging evidence desirable. If the bias introduced by the racial primes was automatic, then PDP estimates should reflect the influence of primes solely in the A estimate.

Using Equations 1 through 4 above, I computed automatic and controlled estimates for White and Black prime conditions. For the Black prime condition, the controlled estimate was calculated by subtracting the probability of false alarms when a tool was primed with a Black face (incongruent condition) from correct responses when a gun was primed with a Black face (congruent condition). Given that estimate of control, the automatic estimate was computed as the probability of false alarms when a tool was primed with a Black face (incongruent condition) divided by $(1 - C)$.

For the White prime condition, the controlled estimate was calculated by subtracting the probability of false alarms when a tool was primed with a White face from the probability of correct responses when a gun was primed with a White face. Having this estimate of control, I calculated the automatic estimate in this condition as the probability of false alarms when a tool was primed with a White face, divided by $(1 - C)$. Three participants received a score of $C = 1$ for the controlled estimate. Because $(1 - C)$ serves as the denominator when calculating the A estimate, these individuals would receive an undefined value for A. As a result, a correction that has been used to handle similar problems with signal detection and high-threshold memory models was applied for these three participants' data. Methodologically, the best approach is to design experiments in which control is less than perfect, which is taken up in Experiment 2. However, this correction is effective when a small portion of the data require adjustment. For a description of the adjustment procedure, see Snodgrass and Corwin (1988).

The estimates are displayed in the top panel of Table 3. To test whether the automatic estimates differed as a function of the prime, I performed a repeated measures ANOVA. The automatic

Table 2
Mean Proportion of Errors by Prime and Target Conditions in Experiments 1 and 2

Target	Prime			
	Black		White	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Experiment 1				
Gun	.06	.09	.06	.07
Tool	.08	.10	.06	.09
Experiment 2				
Gun	.25	.09	.27	.11
Tool	.37	.18	.31	.22

Table 3
Automatic and Controlled Estimates by Prime and Target
Conditions in Experiments 1 and 2

Process estimate	Prime			
	Black		White	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Experiment 1				
Automatic	.61	.29	.48	.34
Controlled	.86	.18	.88	.14
Experiment 2				
Automatic	.57	.13	.49	.17
Controlled	.40	.19	.44	.26

estimate was higher for the Black prime condition, compared with the White prime condition, $F(1, 30) = 4.13, p < .05$. However, when the control estimate was analyzed by using the same model, there was no difference between prime conditions, $F(1, 30) = 1.17, ns$. In fact, the controlled estimates are nearly identical (.86 vs .88). The priming manipulation affected the automatic component, leaving the controlled component unchanged.

Explicit Attitudes and Motivations

Having demonstrated that racial primes affect the visual identification of weapons, and that they do so on the automatic level, it is important to ask how individual differences in this bias relate to individual differences in consciously expressed racial attitudes. In this section, I explore the roles of the MRS and the MCP in predicting RT bias, A, and C estimates. Both the MRS ($M = -2.64, SD = 1.29, \alpha = .89$) and the MCP ($M = 0.35, SD = 1.20, \alpha = .86$) showed good reliability. Four individuals failed to complete the MCP scale. Because MRS and MCP were negatively correlated ($r = -.49, p < .005$), I used multiple regression to test the independent relationships between the MRS, MCP, and each of the dependent measures derived from the priming procedure. Three separate regression equations predicted individual variability in (a) RT differences between Black and White prime conditions, (b) automatic estimates, and (c) controlled estimates as dependent variables.

To examine the relationship between the explicit questionnaire measures and RTs on the priming task, I computed a difference score by subtracting RTs when identifying a weapon in the Black prime condition from RT when identifying weapons in the White prime condition. Higher scores on this index reflect greater racial bias. Raw scores are displayed in Figure 2. This contrast of the identification of guns in the Black versus White prime condition is the most theoretically interesting one for the present purposes, because the tool items were included as race-neutral lures. The main effects were entered on the first step. Also included on the first step was a dummy coded variable representing participants' sex, because initial analyses revealed a main effect of sex, $\beta = 0.60, p < .001$.² Men exhibited a larger racial bias in RT than women did. The two-way interaction terms were entered on the second step.

The main effect for the MCP was significant, $\beta = -1.00, p < .005$. However, this effect was qualified by a significant interaction

between the MRS and MCP, $\beta = -0.53, p < .05$. To display this interaction, I split individuals into three equal groups according to their scores on MCP. Next the regression lines were plotted describing the relationship between the MRS and RT bias separately for each group (see Figure 2). As Figure 2 shows, racial attitude as measured by the self-report MRS is weakly positively related to racial bias in RT for the low MCP group ($\beta = +0.12$). Among low MCP participants, greater prejudice as measured by self-report is associated with somewhat greater RT bias in the priming task. However, the medium ($\beta = -0.25$) and high ($\beta = -0.61$) MCP groups both showed a negative relationship between MRS and RT bias. This pattern conceptually replicates the results of Fazio et al. (1995). For individuals who were not motivated to control their prejudiced reactions, explicit and implicit measures were positively related, though the relationship was small. For individuals who were at least somewhat motivated to control prejudiced reactions, the relationship was negative: Individuals who displayed greater bias in the RT task scored lower in modern racism. Consistent with the interpretation of Fazio et al. (1995) it is likely that this interaction was driven by overcompensation on the self-report measure of prejudice among participants with a strong motivation to control their prejudices.³

Next, I tested the relationships between the MRS, MCP, and automatic and controlled estimates. Regression equations similar to the model described above were used. First, the automatic estimate was regressed on the MRS, MCP, and their interaction. Neither MRS, MCP, nor the interaction term was significantly related to the automatic estimate (full model $F < 1$). Next, I included the controlled estimate as the dependent variable. Note that the estimate of control in this paradigm does not refer to the controlled processing of the racial prime. Rather, it estimates participants' ability to successfully discriminate guns from tools. That is, it indexes the ability to identify an actual gun as a gun, and an actual tool as a tool. None of the variables were reliably associated with the controlled estimate (full model $F = 1.6, ns$). Automatic and controlled estimates were based on accuracy, which showed low variability in this study. Therefore, it would be difficult to detect reliable correlations with these estimates within this study (for a discussion of floor and ceiling effects in the PDP, see Jacoby et al., 1997). Experiment 2 addresses this problem by increasing the variability in accuracy.

² Preliminary analyses conducted throughout both experiments revealed no main effects or interactions of sex in any other comparisons. As a result, sex was not included and is not reported as a factor in subsequent analyses. Men were outnumbered by women in both studies, reflecting the composition of the participant pool from which they were drawn. However, the fact that sex did not qualify any of the results suggests that this may not pose a serious threat to the generality of the conclusions reached.

³ As suggested by an anonymous reviewer, a term representing RTs for the full Prime Race \times Target interaction was also computed as a dependent variable in this analysis. The interaction term was represented as Black weapon + White tool - Black tool - White weapon. The main effects of the MRS and MCP were significant, but smaller than in the reported analysis, but the MRS \times MCP interaction was not significant. Apparently, the variance most reliably related to the individual difference measures was primarily captured in the White weapon-Black weapon contrast.




Figure 2. Regression lines predicting reaction time bias (in milliseconds) from modern racism scores, as a function of motivation to control prejudice (MCP) in Experiment 1.

Discussion

Experiment 1 extended previous research on the activation of stereotypes by showing that the presence of racial information systematically biases an important perceptual judgment: the identification of weapons. Specifically, non-Black participants were faster to identify guns when they were primed by Black versus White faces. The fact that this effect took place at a relatively short SOA (200 ms) suggests that the impact of the racial prime had properties of automaticity. In particular, the effect occurred very rapidly and affected performance at a task that was ostensibly unrelated to race.

This is one of the first experiments to apply Jacoby's (1991) PDP in disentangling the automatic and controlled influences of stereotypes.⁴ Results, using PDP estimates, provided further evidence that the racial primes exerted an automatic influence, independent of controlled processing, which remained unaffected. Specifically, PDP results suggest that the effect was automatic because the presence of racial primes influenced responses regardless of whether that processing aided or interfered with intentional task performance. The RT results and PDP estimates from this study provide converging evidence that the presence of racial information biased perceivers' ability to detect and classify target stimuli as weapons versus nonweapons. Moreover, both lines of evidence support the claim that this bias exerts its effect in a largely automatic way, independent of the perceivers' conscious intentions.

Correlational analyses showed that racial bias in performance on the perceptual identification task was not directly related to explicit racial attitudes. However, the motivation to control prejudice appeared to moderate the relationship between self-reported attitudes and task performance. This finding is consistent with recent work by several researchers showing that the relationship between implicit and explicit measures of attitudes depends on self-presentational strategies (Fazio et al., 1995; Dunton & Fazio, 1997; Greenwald, McGhee, & Schwartz, 1998). Such strategies may

influence the ways that individuals respond to explicit questionnaires, but they are less likely to affect implicit measures. This is because implicit measures, by their nature, limit the opportunity to strategically manipulate one's performance. As a result, implicit and explicit attitudes may be related in conditions where individuals are willing to be candid. However, implicit and explicit measures in more socially sensitive areas, such as racial attitudes, may be unrelated or even inversely related (see Wilson et al., 2000).

In this procedure, the explicit measures were completed before the implicit priming measure. It is possible that filling out the explicit questionnaires may have caused participants to think about race, and thus may have affected performance on the RT task. However, the converse is also possible: Had participants completed the race-related priming task first, they may have responded differently to the questionnaires. Self-report measures are often criticized for their vulnerability to self-presentational and other strategic concerns. In contrast, one advantage of implicit measures is that they may be more difficult to control, by their very nature. The order chosen here was selected because the priming measure was deemed more resistant to outside influences than the self-report measure. A small pilot study ($N = 11$) in which no questionnaires were completed beforehand showed that the Prime Race \times Target interaction was still present ($p = .07$). Thus, explicit measurement does not appear to be a necessary condition to obtain the priming effects. However, the possible impact of the order of the measures remains an issue that warrants further investigation.

A second important question remaining is whether the bias introduced by racial primes can cause racially biased errors in the identification of weapons. Experiment 1 demonstrated that racial information affected the speed with which participants identified weapons, but not their accuracy. However, participants had unlimited time to respond, and error rates were very low. In Experiment 2, a response deadline was imposed, requiring participants to respond within 500 ms. This addition to the experimental design served three purposes. First, the response deadline was expected to increase the overall error rate, producing more variability that would allow a powerful test of the effects of racial primes on accuracy. By imposing the deadline, this procedure forced participants to use a speed criterion rather than a certainty criterion. As a result, the effects of priming should emerge as a systematic pattern of stereotype-consistent errors. Second, the deadline introduced pressure to respond quickly, creating a more realistic analog to situations in which law enforcement officers must make decisions rapidly. Third, and most important, the deadline should be expected to inflate error rates by reducing the opportunity to exert conscious control over responses, but not affecting the automatic influences of the racial primes.

⁴ An interesting study by Hense, Penner, and Nelson (1995) used the PDP to investigate implicit memory for adjectives stereotypical of older adults. However, the substantive issues addressed by that article are quite different from those addressed in the present research.

Experiment 2

Overview

The design of Experiment 2 was identical to Experiment 1, with the addition of a response deadline. The first prediction was that participants would make more stereotypic errors when primed with a Black face than with a White face. Specifically, I expected participants to falsely identify a tool as a gun more often when it was primed by a Black face, compared with a White face. Second, I predicted that the primes would exert their effect at the automatic level, replicating Experiment 1. Third, it was predicted that the effect of the response deadline would reduce the controlled estimate, while leaving the automatic estimate unchanged. That is, having to respond rapidly should make discrimination between weapons and lures more difficult, but should not change the automatic association between Black primes and guns. Finally, the response deadline should constrain variability in RT, while increasing variability in accuracy rates. As a result, I expected the moderated relationship observed in Experiment 1 among MRS, MCP, and racial bias to be replicated in Experiment 2. However, this meaningful variability should be exhibited in the automatic estimate, which is based on differences in accuracy, rather than on RT.

Method

Participants

Thirty-two new non-Black undergraduates (25 women, 7 men) participated in return for course credit.

Procedure

Experiment 2 was a direct replication of Experiment 1, with the addition of the deadline. All participants were instructed that the task required both speed and accuracy, but that they were required to respond quickly. If they did not, then they would receive feedback indicating that they were not responding fast enough. For each trial, participants were allowed 500 ms after the onset of the target stimulus to respond. If they did not respond within that limit, then a series of large red Xs appeared on the screen for 1 s before they were allowed to go on to the next trial. Participants again performed 48 practice trials in which they were encouraged to “be fast and accurate” at responding. As in Experiment 1, participants completed the MRS and MCP scales before performing the priming task.

Results

RTs

Data were prepared by using the same criteria as in Study 1. First, outliers and incorrect responses were dropped, then RTs were log transformed. The bottom panel of Table 1 presents the untransformed mean RTs for each prime and target condition. A repeated measures ANOVA revealed no significant differences between conditions. This result is not surprising because the response deadline forced participants to respond within a narrow window of time, restricting the variance in RTs. Consistent with this account, analysis of mean RTs across all conditions showed that responses were faster in Experiment 2 than in Experiment 1 ($M_s = 300.98$ vs. 440.99), $t(1, 61) = 12.03$, $p < .0001$. In addition, the standard deviation was smaller in the second study ($SD = 27.31$) than in the first ($SD = 59.70$). If the response

deadline forced participants to use a speed criterion in their responses, then priming effects should be evident in the pattern of errors.

Error Rates

Error rates for each Prime \times Target condition were analyzed by using a repeated measures ANOVA. As shown in the bottom panel of Table 2, the overall rate of errors was considerably higher in Study 2 than in Study 1 (29% vs. 6%). Because each trial required a binary response, correct identification performance is simply the complement of the error rate reported. Analyses revealed a main effect for target, $F(1, 31) = 4.31$, $p < .05$, indicating that participants misidentified tools more often than they misidentified guns. That main effect was qualified by a significant Prime \times Target interaction, $F(1, 31) = 12.02$, $p < .002$. Simple effects tests revealed that, as predicted, participants were more likely to falsely identify a tool as a gun when the target was primed with a Black face, compared with a White face, $F(1, 32) = 10.12$, $p < .003$. The race of the prime did not affect the likelihood of misperceiving a gun as a tool, $F(1, 32) = 2.19$, *ns*. Thus, whereas Experiment 1 showed that racial primes bias the speed with which participants identify weapons, Experiment 2 showed that racial primes can cause systematic errors when there is pressure to respond quickly. The critical finding is that simply priming participants with a Black rather than a White face was sufficient to make them call a harmless item a gun.

Automatic and Controlled Estimates

The racial bias in error rates observed above may be described as automatic in the important sense that it was clearly against the will of the perceivers. By using estimates derived from the PDP, one can explore the magnitude of the automatic and controlled influences that combined to produce this pattern of errors. Most important, the effect of the primes and the response deadline on automatic and controlled estimates can be examined separately.

To examine the effect of the racial primes on automatic processing, I computed estimates for each prime condition, as in Experiment 1. The estimates are shown in the bottom panel of Table 3. A repeated measures ANOVA revealed that automatic estimates were significantly different in the two prime conditions, $F(1, 31) = 13.65$, $p < .001$. As expected, the automatic estimate was larger in the Black prime condition than in the White prime condition. The difference in the controlled estimate did not approach significance ($p > .15$). Again, the effect of the Black prime was to increase the automatic estimate, leaving the controlled estimate unaffected.

Because the designs of the two studies were identical except for the deadline, the data are directly comparable when the deadline is used as an independent variable.⁵ I next compared the automatic and controlled estimates across Studies 1 and 2, to test the effect of the primes and response deadline. This analysis allows a simultaneous comparison not only of A and C estimates across prime

⁵ Strictly speaking, the participants compared from Study 1 to Study 2 were not randomly assigned to each study. However, the analyses were performed across studies because the results are informative and important to understanding the processes underlying the main pattern of results.

conditions, but also of the effect of the deadline on the estimates and the interaction of the prime and deadline conditions. If the automatic and controlled effects observed are independent, then each manipulation should exert an effect on one estimate but not the other, and there should be no interaction between the two manipulations.

Two 2 (prime race) \times 2 (deadline condition) ANOVAs were performed, first with the A estimate and then with the C estimate as the dependent variable. It was expected that the racial prime manipulation would affect the A estimate, whereas the deadline manipulation would affect the C estimate. As predicted in the first analysis, the A estimate was greater in the Black prime condition ($M = .59$) than in the White prime condition ($M = .49$), $F(1, 61) = 10.40, p < .002$. There was no main effect of deadline, $F(1, 61) = .11, ns$. Thus, the automatic estimate was not affected by the response deadline. Finally, the Prime Race \times Deadline interaction was not significant, $F(1, 61) = .41, ns$, indicating that the effect of the prime on the automatic estimate was similar whether participants were forced to respond quickly or not. Next, the same ANOVA model was used to analyze the C estimate across both studies. A main effect of deadline emerged, $F(1, 61) = 89.93, p < .0001$. Participants exhibited greater control in the no deadline condition than in the deadline condition. Neither the prime race, $F(1, 61) = 3.24$, nor the Prime Race \times Deadline interaction, $F(1, 61) = .44$, was significant.

Taken together, these results demonstrate a double dissociation between automatic and controlled processes. The racial-priming manipulation increased automatic activation of the "gun" response, leaving controlled processing unaffected. In contrast, the response deadline manipulation reduced controlled discrimination dramatically, leaving automatic bias estimates in place. The lack of interactions between prime and deadline manipulations suggests that the effects of each manipulation generalize across levels of the other, an important fact for the claim of independence between the two processes.

Explicit Attitudes and Motivations

As in Experiment 1, the MRS ($M = -2.65, SD = 1.09, \alpha = .84$) and MCP scales ($M = 0.67, SD = 1.06, \alpha = .80$) proved reliable. In Experiment 1, the racial bias evident in RT was (slightly) positively related to MRS scores only for participants low in the motivation to control prejudice. For those motivated to control prejudice, the correlation was negative. In Experiment 2, the variability in RT was constrained by the response deadline, but the variability in accuracy increased, compared to Experiment 1. As a result, RT bias scores may reflect less meaningful individual variation in racial bias than the A estimate, which is derived from accuracy.

To test the relationships among explicit prejudice scores, motivation to control prejudice, and RT performance, I repeated the set of regression analyses performed for Experiment 1. On the first step, the MRS and MCP scores were entered. On the second step, the MRS \times MCP interaction term was entered. None of the terms in the model reached significance, overall $F(3, 28) = 1.67, p > .20$. Next, the estimate of control was regressed on MRS, MCP, and their interaction using the same equation. No significant relationships emerged (overall $F < 1$).

Next, a regression analysis was performed predicting the A estimate as the dependent variable. Main effects were evident for MRS ($\beta = 0.51, p < .04$) and MCP ($\beta = -0.94, p < .05$). Overall, participants who scored higher in explicit prejudice showed higher automatic bias estimates. Participants who scored higher in the motivation to control prejudice tended to show less automatic bias. However, as expected, these effects were qualified by an MRS \times MCP interaction, ($\beta = -1.16, p < .03$). As in Experiment 1, for purposes of display, participants were split into three equal groups according to their scores on the MCP. Figure 3 displays the regression lines that predict the automatic estimate from the MRS scores separately for each group. The form of this interaction is generally similar to the one reported in Experiment 1 for RT bias. Among the low MCP group, MRS and the automatic estimate were positively correlated ($\beta = 0.51$). For the middle MCP group, the relationship was intermediate ($\beta = 0.28$). Finally, for participants high in MCP, MRS was inversely related to the automatic estimate ($\beta = -0.23$).

Similar to Experiment 1, a term representing the full Prime Race \times Target interaction was computed, this time for error scores rather than RT. Results of this analysis were similar to, though weaker than the results, using the A estimate as the dependent variable. A main effect for MRS ($\beta = 0.48, p < .06$) and the MRS \times MCP interaction ($\beta = -0.88, p < .10$) were marginally reliable. The PDP estimate of automatic processing may be a more sensitive measure because it estimates the pattern of bias in a particular direction, while correcting for overall level of accuracy.

Discussion

The results of Experiment 2 supported the hypothesis that racial primes would cause stereotype-consistent errors in the identification of weapons when participants were required to respond quickly. The presence of Black faces made participants more likely to misidentify a hand tool as a handgun, compared to the presence of a White face. Whereas Experiment 1 demonstrated a racial bias




Figure 3. Regression lines predicting automatic bias from modern racism scores, as a function of motivation to control prejudice (MCP) in Experiment 2.

in RT, Experiment 2 showed that the bias shifted from RT to accuracy when participants were forced to respond rapidly.

Process dissociation estimates replicated the pattern obtained in Experiment 1. The effect of the racial-priming manipulation was isolated in the automatic component, revealing automatic associations between Black male faces and guns. When the two experiments were compared, the effect of the response deadline was isolated in the controlled component. Having to respond quickly impaired participants' ability discriminate between tools and weapons, but did not change the magnitude of the automatic bias. Recall that automatic estimates were very similar for Experiments 1 and 2. Yet the only significant differences in error rates between race conditions occurred in Experiment 2. It appears that the automatic influence of the racial prime exerted its effect only when controlled identification failed.

It is interesting that the correlates of the RT bias observed in Experiment 1 also shifted in Experiment 2 from RT to the estimate of automatic influences. Explicit racial attitudes were positively correlated with automatic estimates only for individuals who were unmotivated to control their prejudiced responses. For those who were highly motivated to behave in an unprejudiced way, automatic estimates were negatively related to the racial attitudes that individuals reported explicitly. This pattern of results is consistent with an emerging body of literature that delineates the conditions under which implicit and explicit attitudes are likely to be related, and those under which they are not (Devine, 1989; Devine & Monteith, 1999; Fazio, 1990a; Fazio et al., 1995; Greenwald et al., 1998; Macrae et al., 1997; Wilson et al., 2000).

General Discussion

Results of this research strongly support the hypothesis that the race of faces paired with objects does influence the perceptual identification of weapons. Experiment 1 showed that when time was unlimited, Black primes facilitated the identification of guns, relative to White primes. Experiment 2 showed that when response time was constrained, Black primes caused race-specific errors. Harmless distracters were more likely to be classified as guns when primed by a Black face than when primed by a White face.

Beyond demonstrating the existence of a racial bias in the perception of weapons, these studies experimentally dissociated automatic and controlled influences that contribute to that bias. Dissociations such as this have theoretical importance, because they allow researchers to test fine-grained predictions about specific processes. In the present studies, it was predicted and found that two requirements must be met to bias participants' error rates. First, stereotypic cues must be present. Second, the opportunity to consider and control one's response must be limited. Unfortunately, these may often be precisely the conditions present when police officers enter into a confrontation with a stereotyped suspect.

Results showed that racial primes biased the perception of weapons through relatively automatic processes, without changing controlled processing. Further, requiring participants to respond rapidly reduced people's ability to control their responses, without changing the automatic bias caused by the primes. Rather, when control was reduced in this way, an automatic bias of the same magnitude as that observed in Experiment 1 was sufficient to produce a reliable bias in error rates. Finally, correlational analysis

specified the relationships among automatic biases, explicit attitudes, and the motivation to control prejudice. The motivation to control prejudice moderated the correlations between explicit attitudes and implicit bias.

The dissociations observed between automatic and controlled estimates provide evidence that validates the assumption of independence. In particular, the fact that a manipulation affects one parameter, but not another, suggests functional independence between the two parameters. Similarly, the fact that the correlates of automatic and controlled estimates were quite different provides further evidence that the two processes may operate independently. In addition, the correlation between automatic and controlled estimates across both studies was near zero, $r = .07$, *ns*. This low correlation provides yet another source of evidence that the two estimates are stochastically independent. Therefore, the independent dual-process model assumed by the PDP appears to be justified within the present paradigm.

Using the PDP as applied here represents an important new alternative for social cognition research. The PDP operationalizes automaticity as an influence that impacts people's performance regardless of whether it facilitates or interferes with their conscious goals. In the present paradigm, participants had the goal of responding "gun" when the target was, in fact, a gun; they had the goal of responding "tool" only when the target was actually a tool. Control was operationalized as the ability to flexibly monitor and control one's responses, therefore to successfully discriminate between guns and lures. This approach avoids the methodological problem of mapping the distinction between automaticity and control onto specific time intervals or separate measures. In addition, results obtained by using PDP estimates converged nicely with results that used response latency, bolstering the validity of these estimates as applied to the present paradigm.

Memory research that uses signal-detection or threshold models has often emphasized the discriminability parameter, while treating the bias parameter as a factor to be controlled or corrected for. For example, Bellezza and Bower (1981) used a one high-threshold signal-detection model to show that the effects of stereotypes on person memory can be isolated in the bias parameter rather than discriminability. Unfortunately, such findings have often been interpreted as "only bias," to the exclusion of "real memory." The A estimate in the present paradigm is analogous to a bias parameter in these other models. However, bias effects such as these may be at least as interesting as effects found in more controlled processes.

The reported pattern of results has applied implications for law enforcement. These data suggest that, because the bias caused by race is largely automatic, it may be difficult to control directly, especially when cognitive resources are limited. Returning to the example of the police officer in a confrontation with a possibly armed suspect, we can draw several conclusions about the automatic and controlled processes that may serve as independent bases for responding. If the officer is like the average participant in our experiments, he or she will experience some degree of automatic bias when interacting with a Black suspect. That is, the officer will be more prone to respond as if a Black suspect is armed, compared to a White suspect. In situations where a Black suspect is actually armed, this bias will facilitate performance: The officer will be faster to respond, and less likely to make an error, compared to the case in which a White suspect is armed. However,

in situations where a Black suspect is unarmed, the automatic bias may tragically interfere with performance.

The automatic bias experienced by our prototypic officer may only be a danger if that officer fails to exert intentional control. In Experiment 1, with unlimited time to respond, control was very high. Though the automatic bias estimates were very similar across the two experiments, a reliable bias in error rates was evident only in Experiment 2, where control was dramatically reduced. Thus, the automatic bias may serve as a basis for behavioral outcomes only when visual discrimination is difficult. Unfortunately, it is difficult to think of a situation in which time pressure is more intense, or the task is more demanding than in high-risk confrontations with possibly armed suspects, particularly under nonoptimal lighting conditions.

Efforts to reduce racial bias in such cases might proceed with one of two approaches. One approach would attempt to minimize the automatic psychological association between Black people and guns. The other approach would aim at maximizing cognitive control (Jacoby, Kelley, & McElree, 1999). The procedure that has been used here provides one way to test the effectiveness of intervention strategies on each component separately, as well as on the overall pattern of errors in performance. A future direction for research in our laboratory will be to investigate factors that exacerbate or remediate these racial biases. At a theoretical level, a crucial goal will be to specify the ways that automatic and controlled processing interact to produce changes in behavioral performance. The interplay between automatic biases caused by racial cues and the subjective conscious states in which perceivers "sincerely believe" their judgments and intend to behave consistently with them remains fertile ground for social cognition research.

References

- Allport, G. W. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.
- Banaji, M., & Hardin, C. (1996). Automatic stereotyping. *Psychological Science, 7*, 136-141.
- Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influences in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3-51). New York: Guilford.
- Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., pp. 1-40). Hillsdale, NJ: Erlbaum.
- Bellezza, F. S., & Bower, G. H. (1981). Person stereotypes and memory for people. *Journal of Personality and Social Psychology, 41*, 856-865.
- Blair, I., & Banaji, M. (1996). Automatic and controlled processes in stereotype priming. *Journal of Personality and Social Psychology, 70*, 1142-1163.
- Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 1, pp. 1-36). Hillsdale, NJ: Erlbaum.
- Curran, T., & Hintzman, D. L. (1997). Violations of the independence assumption in process dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 531-547.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5-18.
- Devine, P. G., & Elliot, A. J. (1995). Are racial stereotypes really fading? The Princeton trilogy revisited. *Personality and Social Psychology Bulletin, 21*, 1139-1150.
- Devine, P. G., & Monteith, M. J. (1999). Automaticity and control in stereotyping. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology*. New York: Guilford.
- Dovidio, J., Evans, N., & Tyler, R. (1986). Racial stereotypes: The contents of their cognitive representation. *Journal of Experimental Social Psychology, 22*, 22-37.
- Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General, 127*, 286-303.
- 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.
- Fazio, R. H. (1990a). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. *Advances in Experimental Social Psychology, 23*, 75-109.
- Fazio, R. H. (1990b). 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*. Newbury Park: Sage.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology, 69*, 1013-1027.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum model of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. *Advances in Experimental Social Psychology, 23*, 1-74.
- Fritsch, J. (2000, February 26). The Diallo verdict: The overview; 4 officers in Diallo shooting are acquitted of all charges. *The New York Times*, p. A1.
- Gaertner, S., & McLaughlin, J. (1983). Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly, 46*, 3-30.
- Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology, 60*, 509-517.
- Glaser, J., & Banaji, M. R. (1999). *Journal of Personality and Social Psychology, 77*, 669-687.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4-27.
- 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.
- Hay, J. F., & Jacoby, L. L. (1999). Separating habit and recollection in young and elderly adults: Effects of elaborative processing and distinctiveness. *Psychology and Aging, 14*, 122-134.
- Hense, R. L., Penner, L. A., & Nelson, D. L. (1995). Implicit memory for age stereotypes. *Social Cognition, 13*, 399-415.
- Hintzman, D. L., & Curran, T. (1997). More than one way to violate independence: Reply to Jacoby and ShROUT (1997). *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 511-513.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513-541.
- Jacoby, L. L. (1999). Ironic effects of repetition: Measuring age-related differences in memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*, 3-22.
- Jacoby, L. L., Begg, I. M., & Toth, J. P. (1997). In defense of functional independence: Violations of assumptions underlying the process-dissociation procedure? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 484-495.
- Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobio-

- graphical memory and perceptual learning. *Journal of Experimental Psychology: General*, 110, 306-340.
- Jacoby, L. L., Kelley, C. M., & McElree, B. D. (1999). The role of cognitive control: Early selection vs. late correction. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 383-400). New York: Guilford.
- Jacoby, L. L., & ShROUT, P. E. (1997). Toward a psychometric analysis of violations of the independence assumption in process dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 505-510.
- Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General*, 122, 139-154.
- Macrae, C. N., Bodenhausen, G. V., Milne, A. B., Thorn, T. M., & Castelli, L. (1997). On the activation of social stereotypes: The moderating role of processing objectives. *Journal of Experimental Social Psychology*, 33, 471-489.
- Martin, L. L., Seta, J. J., & Crelia, R. A. (1990). Assimilation and contrast as a function of people's willingness and ability to expend effort in forming an impression. *Journal of Personality and Social Psychology*, 59, 27-37.
- McConahay, J. B., Hardee, B. B., & Batts, V. (1981). Has racism declined in America? It depends on who is asking and what is asked. *Journal of Conflict Resolution*, 25, 563-579.
- McFadden, R. D., & Roane, K. R. (1999, February 6). U.S. examining killing of man in police volley. *The New York Times*, p. A1.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited capacity attention. *Journal of Experimental Psychology: General*, 106, 226-254.
- Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solo (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55-85). Hillsdale, NJ: Erlbaum.
- Roediger, H. L., & McDermott, K. B. (1993). Implicit memory in normal human subjects. In H. Spinnler & F. Boller (Eds.), *Handbook of neuropsychology* (Vol. 8, pp. 63-131). Amsterdam: Elsevier.
- Shiffrin, R., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127-190.
- Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, 117, 34-50.
- Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin*, 116, 117-142.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107, 101-126.
- Wittenbrink, B., Judd, C., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, 72, 262-274.

Received June 30, 2000

Revision received October 2, 2000

Accepted October 9, 2000 ■

ORDER FORM

Start my 2001 subscription to *Journal of Personality and Social Psychology!* ISSN: 0022-3514

_____ \$170.00, APA Member/Affiliate _____
 _____ \$340.00, Individual Non-Member _____
 _____ \$863.00, Institution _____
In DC add 5.75% sales tax / In MD add 5% sales tax _____
TOTAL AMOUNT ENCLOSED \$ _____

Subscription orders must be prepaid. (Subscriptions are on a calendar basis only.) Allow 4-6 weeks for delivery of the first issue. Call for international subscription rates.

SEND THIS ORDER FORM TO:
 American Psychological Association
 Subscriptions
 750 First Street, NE
 Washington, DC 20002-4242



Or call (800) 374-2721, fax (202) 336-5568.
 TDD/TTY (202)336-6123. Email: subscriptions@apa.org

Send me a Free Sample Issue Check Enclosed (make payable to APA)Charge my: VISA MasterCard American Express

Cardholder Name _____

Card No. _____ Exp. date _____

Signature (Required for Charge) _____

Credit Card _____

Billing Address _____

City _____ State _____ Zip _____

Daytime Phone _____

SHIP TO:

Name _____

Address _____

City _____ State _____ Zip _____

APA Customer # _____

GAD01

PLEASE DO NOT REMOVE - A PHOTOCOPY MAY BE USED