

The Police Officer's Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals

Joshua Correll, Bernadette Park,
and Charles M. Judd
University of Colorado at Boulder

Bernd Wittenbrink
University of Chicago

Using a simple videogame, the effect of ethnicity on shoot/don't shoot decisions was examined. African American or White targets, holding guns or other objects, appeared in complex backgrounds. Participants were told to "shoot" armed targets and to "not shoot" unarmed targets. In Study 1, White participants made the correct decision to shoot an armed target more quickly if the target was African American than if he was White, but decided to "not shoot" an unarmed target more quickly if he was White. Study 2 used a shorter time window, forcing this effect into error rates. Study 3 replicated Study 1's effects and showed that the magnitude of bias varied with perceptions of the cultural stereotype and with levels of contact, but not with personal racial prejudice. Study 4 revealed equivalent levels of bias among both African American and White participants in a community sample. Implications and potential underlying mechanisms are discussed.

In February 1999, around midnight, four plain-clothes police officers were searching a Bronx, New York, neighborhood for a rape suspect. They saw Amadou Diallo, a 22-year-old West African immigrant, standing in the doorway of his apartment building. According to the police, Diallo resembled the suspect they were tracking. When they ordered him not to move, Diallo reached into his pants pocket. Believing he was reaching for a gun, the police fired a total of 41 shots, 19 of which hit and killed Diallo. Diallo was in fact unarmed. All four officers were later acquitted of any wrongdoing in the case.

The police could not have known for certain that Diallo was harmless. In the dark, they had ordered a potentially dangerous man to freeze, and that man reached for something. If Diallo had been armed, their decision to open fire would never have been questioned. But the decision to shoot a man who later proved to be unarmed did raise questions, one fundamental question in particular: Would the police have responded differently if Diallo had been White? Perhaps Diallo would have been given the benefit of the doubt, perhaps the order to freeze would have been repeated, perhaps a slight delay in the decision to fire would have given the officers time to recognize that this suspect was not reaching for a

gun. Though it is impossible to reach a definitive answer with respect to Diallo's case, the dilemma faced by these officers has important consequences for cities nationwide and warrants a systematic investigation. It seems crucial to understand whether or not the decision to shoot is influenced by the target's ethnicity, and if so, what this bias represents.

Social psychology has long held an interest in the way that schemata, including expectancies about social categories like ethnicity, guide the interpretation of ambiguous information (Duncan, 1976; Hilton & von Hippel, 1990; Jacobs & Eccles, 1992; Rothbart & Birrell, 1977; Sagar & Schofield, 1980). The quick and almost effortless classification of a unique individual into a broad social category (Brewer, 1988; Fiske, Lin, & Neuberg, 1999; Fiske & Neuberg, 1990) may lead people to assume that traits generally associated with the category also apply to this particular member. Either in the absence of individuating information (Darley & Gross, 1983; Locksley, Borgida, Brekke, & Hepburn, 1980; see Hamilton & Sherman, 1994, for a review) or in spite of it (Beckett & Park, 1995; Krueger & Rothbart, 1988), stereotypic associations can influence an observer's perceptions in a top-down fashion. A stereotype, in essence, can function as a schema to help clarify or disambiguate an otherwise confusing situation.

Of particular interest to the question of Diallo's death is the possibility that the officers' decision to fire was influenced by the stereotypic association between African Americans and violence. The ambiguity of Diallo's behavior (what was he reaching for?), which ironically provides a justification for the officers' decision, may have set the stage for bias, prompting the officers to draw on other sources of information, including stereotypes, in an effort to understand what was happening. Duncan (1976) showed that the same mildly aggressive behavior is perceived as more threatening when it is performed by an African American than when it is performed by a White person. A White person's light push seems like a violent shove when performed by an African American. Sagar and Schofield (1980), following Duncan, presented 6th-

Joshua Correll, Bernadette Park, and Charles M. Judd, Department of Psychology, University of Colorado at Boulder; Bernd Wittenbrink, Graduate School of Business, University of Chicago.

This material is based on work supported by a National Science Foundation graduate research fellowship awarded to Joshua Correll. Support for this work also came from National Institute of Mental Health Grant R01-45049 awarded to Bernadette Park and Charles M. Judd and from a sabbatical award from the James McKeen Cattell Fund awarded to Bernadette Park. We thank Paul G. Davies for his invaluable contributions to the conceptual development of this project, and Joshua Ingram and David M. Deffenbacher for their assistance in its implementation.

Correspondence concerning this article should be addressed to Joshua Correll, Department of Psychology, University of Colorado, Boulder, Colorado 80309-0345. E-mail: jcorrell@psych.colorado.edu

grade boys with line drawings and verbal accounts of ambiguous dyadic interactions, for example, two boys bumping into one another in the hallway, or one boy borrowing a pencil from a classmate without asking. To manipulate the ethnicity of the people interacting, the researchers simply shaded in the drawings. Like Duncan, they found that when an actor was depicted as African American, rather than White, his behavior seemed more mean and threatening to the participants. Sagar and Schofield further found that this bias in perception was similar for both White and African American participants. That is, the tendency to see an African American's behavior as more mean and threatening than a White person's did not depend on the observer's ethnicity. On the basis of this result, Sagar and Schofield argued that the bias reflects not the internalization of anti-African American attitudes, but rather the application of a widely known and cognitively derived stereotype about the group to the particular target individual.

Devine (1989) went on to demonstrate that the impact of ethnicity on interpretation could occur even without participants' awareness. She asked participants to rate a target's ambiguously hostile behavior after subliminally priming them with words related to both the social category and the stereotype of African Americans (but excluding words directly related to violence). Participants who were primed with a greater number of these words were more likely to interpret the behavior as hostile, even though the target's ethnicity was never mentioned. Lepore and Brown (1997) primed only the social category of African Americans (not the stereotype) and found that the effect of the primes on interpretation of behavior was only evident among the more prejudiced participants. In all of these studies, the association between the social category, African American, and the concept of violence seems to lead participants to interpret an ambiguous target as more dangerous.

Most recently, Payne (2001) demonstrated that participants were faster and more accurate in distinguishing guns from hand tools when they were primed with an African American face, as opposed to a White face. Using Jacoby's (1991; Jacoby, Toth, & Yonelinas, 1993) Process Dissociation Procedure, Payne then separated participants' errors into automatic and controlled components. The magnitude of the automatic estimate represents the degree to which the ethnicity of the prime influences participants' decisions when their ability to control that decision fails. Among participants who were low in motivation to control prejudiced responding, Payne found that greater prejudice was associated with a greater automatic effect.

The primary goal of the current research was to carry this line of inquiry one step further, investigating the effect of a target's ethnicity on participants' decision to "shoot" that target. We present data from a simplified videogame, which roughly simulates the situation of a police officer who is confronted with an ambiguous, but potentially hostile, target, and who must decide whether or not to shoot. In the game, images of people who are either armed or unarmed, and either African American or White, appear unexpectedly in a variety of contexts. Unlike previous research, this game requires participants to make a behavioral shoot/don't shoot decision similar to that of a police officer. And unlike a sequential priming study (such as Payne, 2001), this game simultaneously presents a target person's ethnicity and the object he is holding. A participant need not process ethnicity to determine

whether the target is armed. In spite of these differences, the research reviewed above strongly suggests that interpretation of the target as dangerous, and the associated decision to shoot, will vary as a function of the target's ethnicity. In Studies 1 and 2, we test this basic prediction. In Studies 3 and 4, we make an initial effort to understand the processes underlying this bias in the decision to shoot.

Study 1

Method

Participants and Design

Forty undergraduates (24 female, 16 male) at the University of Colorado at Boulder participated in this experiment in return for either \$8 or partial credit toward a class requirement.¹ One of the male participants was Latino. All other participants were White. The study used a 2 × 2 within-subject design, with Target Ethnicity (African American vs. White) and Object Type (gun vs. no gun) as repeated factors.

Materials

Using the PsyScope software package (Cohen, MacWhinney, Flatt, & Provost, 1993), we developed a simplistic videogame that presented a series of background and target images. The videogame used a total of 20 backgrounds and 80 target images. Twenty young men, 10 African American and 10 White, were recruited on college campuses to pose as models for the targets. Each of these models appeared in the game four times, twice as a target in the gun condition and twice as a target in the no-gun condition, with a different object and in a different pose each time (five basic poses were used in the game). There were four non-gun objects (a silver-colored aluminum can, a silver camera, a black cell phone, and a black wallet) and two guns (a silver snub-nosed revolver and a black 9-mm pistol). Each of the objects, within condition, appeared equally often in each of the five poses. The four target images for each model were superimposed on randomly determined backgrounds, constrained so that each background was used once in each of the four conditions and no target appeared on the same background more than once. Background images included an intentionally diverse assortment of photographs, such as train station terminals, parks, hotel entrances, restaurant facades, and city sidewalks. No people appeared in any of the original background scenes. Examples of the stimuli appear in Figure 1.

In total, there were 80 trials in the videogame, with 20 trials in each cell of the 2 × 2 design created by crossing the ethnicity of the target with whether the target held a gun or a non-gun. Each of the 80 trials began with the presentation of a fixation point, followed by a series of empty backgrounds, presented in slide-show fashion. The number of backgrounds on a given trial was randomly determined, ranging from 1 to 4. The duration of each was also random, ranging from 500 to 1,000 ms. The final background in the series was replaced by the target image, created by superimposing the target on the final background. From the perspective of the participant, a man seemed to simply appear on the background. The design of the game was intended to ensure that the participant never knew when or where the target would appear in the background or when a response would be required.

To play the game, the participant needed to decide as quickly as possible whether the object the man was holding was a gun or not. If it was a gun,

¹ Gender did not moderate any of the effects we report in this or subsequent studies. In Study 3, there was a main effect of gender, such that men had faster reaction times for all targets than did women, $t(43) = 2.31$, $p = .03$, but this effect did not replicate in the other studies.



Figure 1. Target and background example scenes from videogame. Color originals are available at psych.colorado.edu/~jcorrell/tpod.html

the man posed an imminent danger, and the participant needed to shoot him as quickly as possible by pushing the right button, labeled *shoot*, on a button box. If he was holding some object other than a gun, he posed no

danger, and the participant needed to press the left button, labeled *don't shoot*, as quickly as possible. Participants were instructed to use separate hands for each button and to rest their fingers on the buttons between trials.

Table 1
Means (and Standard Deviations) for Reaction Times and Error Rates as a Function of Target Ethnicity and Object Type (Studies 1, 2, and 3)

Study	Reaction times		Errors per 20 trials	
	White targets	Afr. Am. targets	White targets	Afr. Am. targets
Study 1				
Armed targets	554 (46)	544 (39)	0.70 (1.07)	0.40 (0.78)
Unarmed targets	623 (38)	634 (39)	1.23 (1.29)	1.45 (1.04)
Study 2				
Armed targets	449 (23)	451 (28)	2.46 (1.83)	1.48 (1.38)
Unarmed targets	513 (32)	523 (38)	2.40 (2.76)	3.29 (2.87)
Study 3				
Armed targets	550 (40)	539 (45)	0.76 (0.86)	0.49 (0.80)
Unarmed targets	607 (38)	620 (38)	0.33 (0.90)	0.65 (1.24)

Note. Afr. Am. = African American.

The game awarded and deducted points on the basis of performance. A hit (correctly shooting a target holding a gun) earned 10 points, and a correct rejection (not shooting a target holding some non-gun object) earned 5 points. A false alarm (shooting a target holding a non-gun) was punished by taking away 20 points, and a miss (not shooting a target holding a gun) resulted in our harshest penalty: a loss of 40 points.² This payoff matrix represented an effort to partially, if weakly, recreate the payoff matrix experienced by police officers on the street, where shooting an innocent suspect is a terrible mistake (as in the case of Amadou Diallo), but where the stronger motivation is presumably to avoid misidentifying an armed and hostile target, which could result in an officer's death. To minimize nonresponse, the game assessed a timeout penalty of 10 points if participants failed to respond to a target within 850 ms. This time window was selected to force participants to respond relatively quickly, while still allowing enough time such that errors in the game would be minimized. Participants' decisions ("shoot" or "don't shoot") and their reaction times were recorded for each trial. Each trial ended by giving participants feedback on whether they had made the correct decision on that trial and by showing them their cumulative point total.

Procedure

Participants, in groups of 1 to 4, were met by a male experimenter who outlined the study as an investigation of perceptual vigilance, or the ability to monitor and quickly respond to a variety of stimuli. A detailed set of instructions for the videogame task followed, including the point values for each of the outcomes. Participants were also informed that the people with the first, second, and third highest scores in the study would receive a prize (\$30, \$15, and \$10, respectively) and that 5 others, randomly selected from participants with scores in the top 30%, would each receive \$10. These prizes were intended to make the payoff matrix personally meaningful. Finally, participants were asked to pay attention to the faces of the targets, because they would be tested on their ability to recognize the targets at the end of the game. Participants then moved to individual rooms to play the game.

At the conclusion of the game, participants were presented with a series of 16 recognition trials in a paper-and-pencil task to determine whether facial characteristics of the targets had been attended to. For each of the 16 faces, participants had to indicate whether they believed it was the face of one of the targets that had been seen during the game or not. Half of the presented targets had in fact been seen previously; half had not. Additionally, half of the targets were African American and half were White.

Following the recognition task, participants were given a short questionnaire, which asked whether they valued the monetary incentives, whether

they remembered the point values for hits, misses, false alarms, and correct rejections. Participants were then fully debriefed, with the experimenter paying particular attention to alleviate any negative feelings aroused by the game.

Results and Discussion

To analyze the resulting reaction times, we excluded all trials on which the participant had either timed-out (i.e., failed to make a decision in the allotted 850-ms window) or made an incorrect response (e.g., shooting a target holding a non-gun). This resulted in the exclusion of data from 7% of the trials across participants, with a maximum of 20% of the trials for any one participant. Response latencies on the remaining trials were log-transformed and then averaged within subject across trials occurring in the same cell of the 2 × 2 within-subject research design. An analysis of variance (ANOVA) of the resulting mean latencies was then conducted, treating Target Ethnicity (White vs. African American) and Object Type (gun versus no gun) as within-subject factors.

This analysis revealed a significant main effect for Object, $F(1, 39) = 244.16, p < .0001$, and a significant Object × Ethnicity interaction, $F(1, 39) = 21.86, p < .0001$. The resulting cell means (converted back to the millisecond metric) appear in Table 1. As these means reveal, participants were significantly faster at making the correct decision to shoot, when the target held a gun, than the correct decision to not shoot, when the target did not hold a gun. More central to our predictions, the interaction suggests that the speed of responding on gun versus no-gun trials depended on target ethnicity. We decomposed this interaction by examining the simple effects of ethnicity separately for the gun and no-gun trials. Both were significant: Participants fired at an armed target more quickly if he was African American than if he was White, $F(1, 39) = 10.89, p < .005$, and they decided not to shoot an unarmed White target more quickly than an unarmed African American target, $F(1, 39) = 9.77, p < .005$.

² These point values should, objectively, create a bias to shoot: The two "don't shoot" options yield an average reward of -17.5 points, whereas the "shoot" options yield a less aversive average of -5 points.

We intentionally gave participants a long enough response window (850 ms) in this study to maximize correct responses to examine effects on response latencies. And, as we suspected, the proportions of errors were quite low, averaging 4% of the trials across participants. Nonetheless, it is possible to examine the error rates to see if they depended on Target Ethnicity, Object Type, or their interaction (see mean error rates in Table 1). This analysis revealed a main effect for Object, $F(1, 39) = 32.31, p < .0001$, such that errors in the no-gun condition (i.e., false alarms) were more frequent than errors in the gun condition (i.e., misses). The interaction between Ethnicity and Object was also significant, suggesting that the tendency to make more false alarms than misses was more pronounced for African American targets than for White targets, $F(1, 39) = 7.68, p < .01$. That is, whereas participants tended to shoot unarmed targets more frequently than they decided not to shoot armed targets, in general, this tendency was stronger when the target was African American than when the target was White. The simple effects were in the correct direction, but not statistically significant. Participants were marginally more likely to miss an armed target when he was White than when he was African American, $F(1, 39) = 3.66, p = .06$, but errors in response to unarmed targets did not seem to depend on ethnicity, $F(1, 39) = 1.68, p = .20$.

Both the latency and error results attest to the role of target ethnicity in disambiguating potentially threatening stimuli. Clearly, the responses of participants to these stimuli depended at some level on the ethnic category of the target, with potentially hostile targets identified as such more quickly if they were African American rather than White and benign targets identified as such more quickly if they were White rather than African American. Although these results are certainly consistent with our expectations, they are also somewhat surprising given the fact that the target ethnicity appeared at exactly the same time as the object that had to be identified as a gun or not. Certainly participants could have performed perfectly on the task by attending only to the object held in the target's hand and by completely ignoring the target's ethnicity or any other individuating information.

To examine whether a target's features, other than the object he held, were attended to by participants, we examined their ability to recognize the faces of the targets they had seen during the game. A signal detection analysis revealed that sensitivity to old versus new faces was not above chance level in these recognition data (mean $d' = 0.15$), $t(39) = 1.15, p = .26$. Separate analyses within target ethnicity revealed that participants were unable to recognize African American targets at a better than chance level (mean $d' = -0.08$), $t(39) = -0.48, p = .63$, although recognition sensitivity for the White targets did exceed chance levels (mean $d' = 0.33$), $t(39) = 2.26, p < .05$. Our data suggest, then, that target ethnicity affected participants' judgments even while participants remained largely incapable of recognizing the faces of the targets they had seen.

Study 2

Our first study allowed participants a sufficient response window so that they made correct decisions in the case of nearly all targets. That is, error rates were very low. As a result, the strongest results from the first study were found with decision latencies on correct responses, with faster decisions to armed African American

targets than to armed Whites, and faster decisions to unarmed White targets than to unarmed African Americans. Although significant, the interactive effects of Target Ethnicity and Object Type on response errors were substantially weaker (and the relevant simple comparisons were not significant).

In the second study, we sought to replicate the basic pattern of results from the first study, but this time to make the task substantially harder by shortening the amount of time during which participants had to respond. Clearly, if the effects that we are exploring are to be relevant to more real-life scenarios, such as those encountered by police officers, then we would like to show our effects on actual responses (and errors in responses) rather than simply on the speed with which correct responses are made. Additionally, to increase the importance of performance in the task, we recruited participants exclusively for pay in this study and we offered them incentives directly tied to the quality of their performance, paying up to \$20 for a study taking well less than an hour.

Method

Participants and Design

Forty-four undergraduates (33 female, 11 male) participated in this experiment in return for a minimum payment of \$10, with the opportunity to earn additional money (up to a total of \$20) by scoring points in the game. This incentive was intended to increase the personal significance of the rewards and penalties. One male participant was Latino, and 1 female was Asian. All other participants were White. We used the same 2×2 design, with Target Ethnicity (African American vs. White) and Object Type (gun vs. no gun) as within-subject factors.

Materials and Procedure

The materials and procedure were identical to those of Study 1, with the exception of the following modifications. First, we made clear to participants that they would be paid as a function of their performance. They were told that they started with an initial sum of \$14 to their credit. Each point earned or lost (according to the same payoff matrix used in Study 1) was worth 1 cent. It was made clear that if they performed perfectly across all 80 trials, they would earn \$20. If they lost points, they could lose up to \$4, but they were guaranteed a base pay of \$10. Second, we adjusted the game's response window from 850 ms to 630 ms to force participants to make decisions more quickly, with the goal of increasing error rates. Although a 630-ms response window may provide ample time to process simple stimuli such as faces or isolated objects, our images were fairly complex, and the shortened window proved a challenge for our participants. A pretest indicated that the shortened response window had the desired effect, increasing errors, but also dramatically increasing the proportion of trials on which participants failed to respond in time. Because the meaning of a timeout is ambiguous, a third change we made was to discourage timeouts by increasing their associated penalty from 10 to 50 points (i.e., 50 cents) and stressing the importance of responding quickly in the instructions. As participants' point totals directly affected the amount they were paid, this provided a considerable incentive. We also set an a priori limit, such that any participant with more than 10 timeouts would be excluded from the analysis. A final change we made was to the program used to record participants' data. In Study 1, for each trial, the program only recorded the response, response latency, target ethnicity, and target object (gun vs. no-gun), but the exact target and background for the trial were not recorded. We modified this in Study 2 so that we could identify particular stimuli that were associated with a greater number of errors.

Results and Discussion

Before conducting the primary analysis of error rates, we eliminated 5 participants (all female) who exceeded our a priori threshold of 10 timeouts (one eighth of all trials). Additionally, we examined error rates for particular targets to determine if correct responses were particularly difficult for some. In fact, there were a number of targets that were outliers in the overall distribution, inducing many more errors than the other targets. For instance, one unarmed African American target was shot by more than 90% of our participants. Additionally, one armed African American target and four unarmed White targets resulted in errors for more than one third of the participants. In each of these target images, some detail seemed potentially misleading. For example, one target had a stripe in his shorts that could be mistaken for a gun given the position of his arm. We suspect the substantially shorter time-out window was responsible for producing the unusually high error rates for these six targets. To deal with these outliers, we conducted all analyses twice, once with the full dataset and once deleting the six outlying targets. The analyses that we report are based on the partial dataset. However, with only one exception, as noted below, the results were unaffected by their inclusion/exclusion.

Participants' error rates (number of errors divided by the total number of valid trials) were subjected to a 2×2 ANOVA, with Target Ethnicity (White vs. African American) and Object Type (gun vs. no gun) as the independent variables. The relevant cell means are given in Table 1. The analysis revealed a significant effect for Object, such that the proportion of errors when a gun was present (i.e., misses) was lower than the proportion of errors when a gun was absent (i.e., false alarms), $F(1, 38) = 6.42, p < .02$. We also found the predicted interaction between Ethnicity and Object, $F(1, 38) = 17.83, p < .0001$. A test of the simple effects revealed that, when the target was unarmed, participants mistakenly shot him more often if he was African American than if he was White, $F(1, 38) = 6.53, p < .02$, though this effect was not significant when all targets were analyzed. When the target was armed, however, participants mistakenly decided not to shoot more often if he was White than if he was African American, $F(1, 38) = 13.31, p < .001$.

In addition to the analyses of the error rates, we also analyzed the decision latencies for correct responses, as in Study 1. Not surprisingly, given the considerably shorter response window in this study, there were no effects in the latencies. It seems that Study 1's interaction in response speed was, in this study, pushed over into error rates, due to the tightened response window.

As in Study 1, participants were unable to recognize presented targets above chance level. An analysis of the mean sensitivity to old versus new faces revealed a nonsignificant overall $d' = -0.02$, $t(38) = -0.15, p = .88$. Sensitivity was not above chance for either the White targets (mean $d' = 0.12$), $t(38) = 0.74, p = .46$, or for the African American targets (mean $d' = -0.16$), $t(38) = -0.97, p = .34$.

To understand the error rate results in greater detail, further analyses were conducted using the signal detection model (Green & Swets, 1966/1974; MacMillan & Creelman, 1991). Applied to the present context, the signal detection analysis assumes that targets encountered, both those with a gun and those without a gun, vary on a judgment-relevant dimension. For example, in the

present studies, the extent to which the targets appeared to be threatening might have served as a critical dimension. On average, targets with guns are more threatening than targets who possess other objects (to the extent that they are discriminated at all), but nevertheless, there is a distribution of targets within each set, and these vary in how threatening they subjectively appear to be. Thus, we have two distributions of targets, one of targets with guns and one of targets without guns, and the signal detection model assumes that these are both normal distributions with equal variances. To some extent, of course, these two distributions overlap and the question of sensitivity is the question of the extent to which this is true. That is, if participants are relatively sensitive or accurate, shooting those targets who have guns and not shooting those targets who don't have guns, then the two distributions are largely separated from each other.

Additionally, because participants make a choice between shooting and not shooting a target on the basis of the subjective sense of how threatening the target appears to be, they set a decision threshold somewhere along the continuum that underlies the two distributions. Above that threshold, they shoot the target; below threshold they do not. Where that threshold is set is commonly referred to as the *decision criterion*.

From the two kinds of errors (false alarms: shooting an unarmed target; misses: not shooting an armed target), one can derive estimates of both sensitivity, commonly defined as d' , and decision criterion, in this case defined as c . We estimated both of these parameters for our participants, once for the White targets and once for the African American targets. Unsurprisingly, given the relatively low percentages of errors, participants showed considerable accuracy (i.e., high levels of d') for both the White and African American targets (White $M = 2.47$ [$SD = 0.87$]; African American $M = 2.48$ [$SD = 0.85$]). A test of differential sensitivity between the two kinds of targets failed to reject the null hypothesis, $F(1, 38) = 0.01, p = .93$. There were differences, however, between the two kinds of targets in the response criterion (White $M = 0.03$ [$SD = 0.30$]; African American $M = -0.24$ [$SD = 0.31$]), such that a significantly lower decision criterion to shoot the target was found for African American targets, $F(1, 38) = 22.21, p < .0001$. These results are depicted graphically in Figure 2. In sum, from the perspective of the signal detection model, the differences between responses to the African American and White targets arose not from differences in the underlying accuracy with which the two kinds of targets, those with a gun and those without a gun, can be discriminated. Rather, in the case of the African American targets, participants simply set a lower threshold for the decision to shoot, being willing to shoot targets who seemed less threatening.³

³ The same pattern of signal detection results emerges both for Study 1 and for Study 3. For Study 1, sensitivity did not differ: African American $d' = 3.30$, White $d' = 3.28$, $F(1, 39) = 0.10, p = .75$; but the decision criterion did: African American $c = -0.17$, White $c = -0.09$, $F(1, 39) = 10.07, p = .003$. In Study 3, sensitivity did not differ: African American $d' = 3.54$, White $d' = 3.56$, $F(1, 44) = 0.12, p = .73$; but the decision criterion did: African American $c = -0.02$, White $c = 0.07$, $F(1, 44) = 6.96, p = .02$.




Figure 2. Hypothetical normal distributions representing unarmed and armed targets for signal detection analyses: White (top panel) and African Americans (bottom panel) targets.

Study 3

Studies 1 and 2 provide evidence that the decision to shoot an armed target is made more quickly and more accurately if that target is African American than if he is White, whereas the decision not to shoot is made more quickly and more accurately if the target is White. This pattern of results is fundamentally consistent with research suggesting that participants may use ethnicity to interpret an ambiguously threatening target. When ambiguous behavior is performed by an African American, it seems more hostile, more mean, and more threatening than when it is performed by a White person (Duncan, 1976; Sagar & Schofield, 1980). Participants also recognize a weapon more quickly and more accurately after seeing an African American face, rather than a White face (Payne, 2001). Here, we have shown that ethnicity can also influence a behavioral judgment with serious consequences for both target and shooter.

Simply documenting the existence of this bias does not clarify the mechanism by which ethnicity influences the decision to shoot. We suggested earlier that participants may use the stereotypic association between the social category, African American, and concepts like violence or danger as a schema to help interpret ambiguous behavior on the part of any given African American target. Through deductive inference, traits associated with the category may be applied to the individual category member. It is important to recognize that the proposed process does not require a participant to dislike African Americans, or to hold any explicit prejudice against them, nor does it require that the participant endorse the stereotype; it simply requires that, at some level, the

participant associates the two concepts “African American” and “violent.” Previous research is equivocal in its support of this possibility, suggesting that bias in the interpretation of an ambiguous stimulus may depend on both stereotypic associations and on prejudice. Sagar and Schofield (1980), for example, provide evidence for a stereotype-driven effect. Recall that these researchers found that both White and African American participants interpreted behavior as more threatening if it had been performed by an African American target. Reasoning that bias among the African American participants is not likely to reflect prejudice against African Americans, they concluded that it reflects instead a common belief, or stereotype, that African Americans are more violent than Whites. A culturally communicated stereotypic association may influence interpretations even if the observer does not personally endorse the stereotype or hold a prejudiced attitude (Devine, 1989). Data presented by both Lepore and Brown (1997) and Payne (2001), however, have shown that more prejudiced participants show greater bias in their interpretations of ambiguous stimuli (for Payne, 2001, this relationship was moderated by motivation to control prejudice). Of course, the effect of prejudice on perceptions may be indirect, operating chiefly through the stronger negative stereotypic associations that accompany prejudiced attitudes. The question is whether stereotypic associations predict bias over and above prejudice. To be clear, we hypothesize that although the magnitude of the bias evident in our videogame may covary with participants’ prejudice against African Americans, it is not a function of that prejudice, per se, but rather reflects the deductive application of stereotypic associations (often asso-

ciated with prejudice) between African Americans and violence. Because participants can use traits associated with the group to disambiguate a particular African American target, they may inappropriately perceive that target as threatening or hostile.

Study 3 represents a first attempt to test these predictions. After playing the videogame, participants completed a questionnaire designed to measure prejudice and two forms of the association between African Americans and violence. The first measure of this association assessed stereotypes that the participant personally endorses or believes. We refer to this as the *personal stereotype*. The second measure, called the *cultural stereotype*, is designed to assess the participant's awareness that a stereotype of African Americans as violent is present in U.S. culture, generally. Though we use the terms *personal stereotype* and *cultural stereotype*, this distinction maps cleanly on to the endorsement/knowledge distinction suggested by Devine (Devine, 1989; Devine & Elliot, 1995), who has shown that, although people often personally disavow negative stereotypes about African Americans, they are well aware that those stereotypes exist. Because this knowledge represents a psychological link between the social category and the trait, activating the concept of the group may predispose a participant to make use of the stereotypic trait in interpretations of an ambiguous target—even if he or she does not personally endorse the stereotype. Both personal and cultural forms of the stereotypic association, then, may influence interpretation of an ambiguous target.

Method

Participants and Design

Forty-eight undergraduates (26 female, 22 male) participated in this experiment in return for either \$10 or partial credit toward a class requirement. Two male participants were Latino, and 1 female was Asian. Another female was African American and was excluded from our analyses. All other participants were White. Two White females were also removed from the dataset, one because the game's *shoot* and *don't shoot* labels were reversed, and one because she was working as a research assistant on a different study of African American stereotypes. The final sample included 45 students. This study used the same 2×2 within-subject design, with Target Ethnicity (African American vs. White) and Object Type (gun vs. no gun) as repeated factors.

Materials

Videogame. In this study, we used the videogame parameters we had used in Study 1. The response window was set at 850 ms, and we expected effects primarily in the latency of correct responses, rather than in error rates.

Questionnaire. Study 3 added a battery of individual difference measures. First, participants completed the Modern Racism Scale (MRS; McConahay, Hardee, & Batts, 1981), the Discrimination (DIS) and Diversity Scales (DIV) (both from Wittenbrink, Judd, & Park, 1997), all of which are designed to measure prejudice against African Americans, as well as the Motivation to Control Prejudiced Responding Scale (MCP; Dunton & Fazio, 1997; Fazio, Jackson, Dunton, & Williams, 1995), which assesses participants willingness to express any prejudice they may feel. Items from these scales were intermixed (presented in a single, randomly determined order) and responses were given on 5-point scales, ranging from *strongly disagree* to *strongly agree*. The items were intermingled with filler items from the Right-Wing Authoritarianism Scale (RWA, Altemeyer, 1988) and the Personal Need for Structure Scale (PNS; Thompson, Naccarato, Parker, & Moskowitz, 2001), which are addressed below.

Second, to examine the degree to which participants endorsed a negative stereotype of African Americans as aggressive and dangerous, we asked them to estimate, on the basis of their personal beliefs, the percentage of both African Americans and Whites who are dangerous, violent, and aggressive (separate estimates were made for each trait by filling in a value from 0% to 100%). Third, we included a measure of participants' perceptions of the cultural stereotype that African Americans are aggressive and dangerous. Participants were asked to again consider the three attributes (dangerous, violent, and aggressive), giving prevalence estimates, not on the basis of their own personal beliefs, but rather on the basis of their perceptions of what most White Americans would estimate. These estimates were made by marking a 130-mm line anchored with the adjective (e.g., *dangerous*) on the right, and its negation (e.g., *not dangerous*) on the left.

In addition to these primary measures, the questionnaire included several exploratory components. We included the RWA scale (Altemeyer, 1988), which measures an individual's predisposition to think of social relations in terms of dominance and submission; the PNS scale (Thompson et al., 2001), which measures differences in the desire for a simple structure; and a five-item measure of contact with African Americans. These measures were included partly as filler items designed to mask the questionnaire's focus on prejudice and stereotyping, but also because these constructs have been shown to be related to prejudice or stereotyping in previous research (Neuberg & Newson, 1993; Pettigrew & Tropp, 2000; Pratto, Sidanius, Stallworth, & Malle, 1994). Responses to the contact items were made on 7-point scales. The first contact question asked participants to rate how many African Americans they know, using a scale anchored with *don't know any African Americans* and *know a lot of African Americans*. The second item asked for a rating of how well they know their African American acquaintances on a scale from *don't know well* to *know very well*. The third item asked about the degree of contact with African Americans in their neighborhood, when growing up. The fourth item asked about the number of African American friends they had while growing up. And the fifth item asked about the number of African Americans who had attended their high school. The last three items used a scale ranging from *none* to *many*.

Procedure

As before, a male experimenter greeted participants, in groups of 1 to 4, and introduced the study as an investigation of perceptual vigilance. He went on to note that, because the vigilance task did not require the entire time period, participants would work on a separate questionnaire study afterward. After learning about the rules of the game, participants moved to computer terminals in private rooms and played the videogame. As each participant completed the game, the experimenter moved him or her to a table (still in the private room) and administered the short questionnaire, from Studies 1 and 2, assessing basic reactions to the game. The experimenter subsequently announced that the videogame study was over and provided another consent form, ostensibly for the separate questionnaire study. After collecting the consent form, he handed the participant an envelope containing the questionnaire. We made every effort to stress the confidentiality of the responses on the questionnaire. The experimenter told participants not to put any identifying information on the forms, not even a code number, and to seal the packet in the envelope when they had finished. He then left them alone to complete the questions. As in Studies 1 and 2, participants were fully debriefed. During this process, the experimenter probed for suspicion about the relationship between the game and the subsequent questionnaire.

Results and Discussion

In the debriefing, 6 participants reported that they had noticed that both the game and the questionnaire involved ethnicity, and

that this awareness had prompted them to wonder if the two were related. Two of the 6 reported strong suspicion. The following results are based on the complete dataset, but exclusion of the 6 participants does not affect the analyses in either direction or significance. To analyze the videogame data, we submitted the log-transformed reaction times from correct trials to a 2 × 2 ANOVA, with Target Ethnicity and Object Type as the independent variables (see Table 1 for means converted back to milliseconds). The targets that had proved problematic in Study 2 were excluded from this analysis, though their inclusion does not substantially affect the results. Replicating the results from the first study, we found both a pronounced effect for Object, such that armed targets were responded to more quickly than unarmed targets, $F(1, 44) = 171.33, p < .0001$, and an Ethnicity × Object interaction, $F(1, 44) = 22.44, p < .0001$. Simple effects tests revealed that, when the target was armed, participants, on average, fired more quickly if he was African American than if he was White, $F(1, 44) = 4.15, p < .05$. When presented with an unarmed target, participants chose the “don’t shoot” alternative more quickly if he was White than if he was African American, $F(1, 44) = 22.72, p < .0001$.

Mean scores on the error rates were largely consistent with those from Study 1. The Ethnicity × Object interaction was significant, $F(1, 44) = 7.20, p = .01$. Simple effects tests showed an ethnicity effect only among targets without guns, $F(1, 44) = 5.76, p = .02$, such that these were incorrectly shot more often if they were African American. The simple effect for armed targets was not significant, $F(1, 44) = 2.31, p = .14$. A test of the mean recognition sensitivity for the presented targets was significant in this study (mean $d' = 0.25$), $t(44) = 2.51, p = .016$. As in Study 1, however, sensitivity was above chance only for the White targets (mean $d' = 0.62$), $t(44) = 4.71, p < .0001$, and not for the African American targets (mean $d' = -0.15$), $t(44) = -1.14, p = .26$.

Having replicated the Ethnicity × Object interaction in the response latency scores, we wanted to examine its correlates. Accordingly, for each participant we computed a within-subject contrast score, assessing the magnitude of the Ethnicity × Object interaction for that particular participant. Higher scores on this variable, which we refer to as *Shooter Bias*, indicate faster re-

sponses to unarmed White than to unarmed African American targets, and to armed African American than armed White targets.

Table 2 reports the correlations between this Shooter Bias measure and the various questionnaire measures. Table 2 also reports the means, standard deviations, and internal consistency (coefficient alpha) statistics for the various attitude scales in our data. With the exception of contact and the personal and cultural stereotype measures, all measures were collected on 5-point scales with higher numbers indicating greater endorsement of the construct. None of the explicit prejudice scales—MRS, DIS, and DIV—show significant correlations with the Shooter Bias from the videogame. That is, those who reported higher levels of prejudice on these scales did not show a stronger ethnicity bias in the videogame. Because these three measures are highly intercorrelated, we also combined them, averaging all items together. This composite scale was similarly uncorrelated with Shooter Bias.

To compute the personal stereotype measure of African Americans as aggressive, we calculated the degree to which participants rated African Americans as more violent than Whites, more dangerous than Whites, and more aggressive than Whites. These three difference scores were averaged together to form the personal stereotype index. The measure reflects perceptions of the prevalence of the negative stereotypic attributes among African Americans relative to Whites. Because this measure is based on percentage estimates, it can potentially range from -100 to 100. One participant chose not to complete the relevant items, so all tests of the personal stereotype are based on a sample of 44, rather than 45. The same process was followed in computing the extent to which participants believed there is a negative cultural stereotype of African Americans as dangerous and aggressive. Because the raw scores on the cultural stereotype items were made on 130-mm lines, the index potentially ranges from -130 to 130. As is clear in Table 2, the measure of personal endorsement of the negative stereotype of African Americans as aggressive and violent did not correlate with the Shooter Bias. However, the perception of a parallel negative cultural stereotype did correlate with the magnitude of the Shooter Bias in the videogame.

Of the exploratory measures (RWA, PNS, and contact), only contact was related to Shooter Bias. Contact scores were calculated

Table 2
Correlations of Shooter Bias in Videogame With Questionnaire Measures (Study 3)

Variable	<i>M</i>	<i>SD</i>	α	Shooter Bias	MRS	DIS	DIV	Prejudice comp.	Personal stereo.	Cultural stereo.	MCP	Contact	RWA
MRS	1.63	.66	.86	.15									
DIS	2.09	.73	.87	.16	.80**								
DIV	2.43	.64	.64	.05	.46**	.59**							
Prejudice comp.	2.09	.60	.91	.14	.85**	.95**	.78**						
Personal stereo.	1.43	7.56	.54	.05	.38**	.38**	.38**	.43**					
Cultural stereo.	41.37	24.15	.88	.37**	-.06	-.07	-.21	-.12	-.06				
MCP	3.23	.48	.72	.03	-.35*	-.29*	-.27†	-.34*	-.31*	.06			
Contact	2.56	1.00	.72	.38**	-.18	-.02	.11	-.03	-.07	.09	.15		
RWA	2.16	.61	.72	-.04	.26†	.24	.37**	.33*	.02	-.25†	.25†	.11	
PNS	2.78	.55	.77	-.15	.16	.00	-.16	-.01	-.04	.15	.17	-.06	.29*

Note. For all measures except personal stereotype, $n = 45$. All comparisons involving personal stereotype are based on $n = 44$. MRS = Modern Racism Scale; DIS = Discrimination Scale; DIV = Diversity Scale; Prejudice comp. = prejudice composite; stereo. = stereotype; MCP = Motivation to Control Prejudiced Responding Scale; RWA = Right-Wing Authoritarianism Scale; PNS = Personal Need for Structure Scale.
† $p \leq .10$. * $p \leq .05$. ** $p \leq .01$.

This document is copyrighted by the American Psychological Association or one of its allied publishers. This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.

by averaging participants' responses to the five 7-point contact items. This measure showed a significant and somewhat surprising correlation with the bias: Participants who reported more contact with African Americans exhibited a more pronounced Shooter Bias in the videogame. We discuss this intriguing effect in the General Discussion when we consider potential mechanisms that may give rise to Shooter Bias.

We suggested that the Shooter Bias evident in this videogame might be a consequence of participants using stereotypic associations about African Americans to help interpret ambiguous African American targets. The data from Study 3 suggest that the magnitude of the bias was related to participants' perceptions of the cultural stereotype about African Americans. The bias was not, however, related to either personally endorsed stereotypes or to prejudice. This is somewhat surprising, because, to the extent that people personally endorse the violent stereotype or hold prejudices against African Americans, we might suppose the negative associations to be stronger and more likely to influence their interpretations of, and behavior toward, an ambiguous target.

There are well-documented social desirability concerns associated with expressing prejudice or negative stereotypic beliefs about African Americans (Dunton & Fazio, 1997; McConahay et al., 1981; Plant & Devine, 1998), so it may be that participants simply refused to express their personal views. In his research, Payne (2001) found no zero-order correlation between prejudice (as measured by the MRS) and the automatic component in his weapon identification task. He did find a moderated relationship between the two variables, though, such that a positive correlation emerged only among participants who were low in MCP. A similar test in our data yielded no significant interaction between MRS and MCP, $F(1, 41) = 0.00, p < .95$, or between personal stereotype and MCP, $F(1, 40) = 0.95, p < .34$, when predicting Shooter Bias.

Unlike prejudice and personal stereotypes, our measure of cultural stereotype should be generally free from social desirability concerns. It involves participants' estimates of the stereotype held by American society. The fact that cultural stereotype correlates with Shooter Bias suggests that awareness of the stereotype, itself, even though a person may not believe that stereotype, can be sufficient to produce bias. One might argue, however, that our cultural stereotype measure was just another way of measuring personal prejudice, in a manner that allowed participants to express their own prejudices relatively free from normative constraints. That is, by attributing prejudicial beliefs to others, participants were now able to express more freely the prejudice that they themselves felt.

The bivariate correlation between the cultural stereotype measure and our composite personal prejudice scale was $-.12$ ($p = .41$), suggesting that this cultural stereotype measure is not a simple proxy for personal prejudice levels. However, it might be the case that the relationship between the cultural stereotype measure and personal prejudice depends on the participant's level of motivation to control prejudice, again following the theoretical arguments of Fazio et al. (1995). To examine this possibility, we regressed the cultural stereotype measure on our composite personal prejudice measure, MCP, and their interaction. The interaction proved to be a significant predictor, $F(1, 41) = 4.67, p < .05$. The direction of this interaction was as predicted: there was a more positive relationship between personal prejudice levels and the

cultural stereotype measure among those who were lower in motivation to control prejudice.

We were interested in whether cultural stereotype would continue to predict Shooter Bias once we removed the extent to which the cultural stereotype variable is a measure of personal prejudice, particularly among those low in motivation to control prejudice. Accordingly, we estimated a model with Shooter Bias as the criterion, regressing it on the cultural stereotype measure while controlling for our personal prejudice composite, MCP, and the interaction between personal prejudice and MCP. In this model, again, only the cultural stereotype measure related significantly to bias in the videogame, $F(1, 40) = 5.24, p < .03$. Thus, even removing personal prejudice levels from the cultural stereotype, and controlling for the fact that personal prejudice levels were more strongly related to the cultural stereotype among those low in MCP, the cultural stereotype measure continued to predict bias in our videogame.⁴ This suggests that it is truly knowledge of the cultural stereotype that is at work here, rather than simply an indirect measure of personal prejudice. We consider this a sobering prospect because it suggests that the bias may be endemic in American society.

A number of studies have shown that cultural stereotypes can be automatically activated even when a perceiver does not endorse them (Banaji & Greenwald, 1995; Devine, 1989; Gilbert & Hixon, 1991; Macrae, Milne, & Bodenhausen, 1994). Cultural influences, including television, movies, music, and newspapers provide a constant barrage of information that often depicts African Americans as violent (Cosby, 1994; Gray, 1989), and those depictions may shape our understanding of the world (Gerbner, Gross, Morgan, & Signorielli, 1986). Popular culture, including Gangsta Rap songs like the Notorious B.I.G.'s "Somebody's Gotta Die," Snoop Dogg's "Serial Killa," or Dr. Dre's "Murder Ink," and movies like *Colors* or *Training Day* may foster bias by enhancing detrimental stereotypic associations, in spite of the fact that the audience knows the characters and events are fictitious.

If cultural stereotypes associating African Americans with violence do, in fact, lead to Shooter Bias, any person exposed to American culture should be liable to demonstrate the bias, regardless of his or her personal views about African Americans. Research suggests that the very people who are targeted by cultural stereotypes are influenced by the media representations they see (Berry & Mitchell-Kernan, 1982; Stroman, 1986; SuberviVelez & Necochea, 1990), know full well that the stereotypes exist (Steele & Aronson, 1995), and even activate those stereotypes automatically (Banaji & Greenwald, 1995). Sagar and Schofield (1980), as noted above, found similar levels of bias among their African American and White participants using their interpretation task. To examine further the possibility that knowledge of the cultural stereotype may, in and of itself, lead to Shooter Bias, we sought to

⁴ The attempt to control for the prejudice composite measure, MCP, and their interaction only removes variance based on personal prejudice to the extent that these scales reliably measure that variance. There is reason to assume that these measures only partially assess prejudice, particularly for participants high in MCP. Thus, although the analysis represents our best attempt to examine the effects of cultural stereotypes over and above prejudice in the current dataset, it is nonetheless imperfect. Our thanks to Keith Payne for this insight.

Table 3
Means (and Standard Deviations) for Reaction Times and Error Rates as a Function of Target Ethnicity, Object Type, and Participant Ethnicity (Study 4)

Participants	Reaction times		Errors per 20 trials	
	White targets	Afr. Am. targets	White targets	Afr. Am. targets
White participants				
Armed targets	590 (43)	578 (36)	1.38 (1.36)	0.76 (0.77)
Unarmed targets	652 (40)	665 (41)	1.19 (0.93)	1.29 (1.49)
Afr. Am. participants				
Armed targets	578 (42)	567 (47)	2.00 (1.53)	1.52 (1.58)
Unarmed targets	645 (47)	659 (41)	1.64 (1.80)	1.44 (1.47)

Note. Afr. Am. = African American.

test for bias in a more diverse sample that included African American participants.

Study 4

Method

Participants and Design

Fifty-two adults from bus stations, malls, and food courts in Denver, Colorado, were recruited to participate in this study in return for \$5. The study followed the same 2×2 within-subject design used in Studies 1–3, with Target Ethnicity (African American vs. White) and Object Type (gun vs. no gun) as repeated factors, but in Study 4 we added a between-subject factor, namely Participant Ethnicity (African American vs. White). The final sample included 25 African Americans (6 females, 19 males) and 21 Whites (8 females, 13 males). One Asian and 4 Hispanic or Latino participants, and 1 participant who did not indicate his ethnicity, were excluded from the analyses, though the results do not differ if they are included in the White sample.

Materials

In this study, we used the videogame parameters from Studies 1 and 3. The response window was set at 850 ms and, again, we expected effects in the latency of correct responses, rather than in error rates. Before beginning this study, the targets identified as problematic in Study 2 were edited in Photoshop to clarify the object in the picture.

Procedure

At each location, two male experimenters set up 2–3 laptop computers equipped with the videogame program and earphones, to minimize distractions inherent in the nonlaboratory environment. Without a button box, participants pressed the *k* key on the laptop keyboard to indicate *shoot*, and the *d* key to indicate *don't shoot*. While one experimenter circulated and recruited participants, the other oversaw the experiment, giving instructions to each participant individually. After completing the videogame, participants were paid and debriefed. In this study, we did not include instructions to attend to target faces, nor did we test for recognition after the game.

Results and Discussion

Before analyzing the videogame data, we reexamined the targets that were problematic in Study 2. The targets no longer induced

unusually high numbers of errors, and they were therefore included in the analyses reported below. The results reported do not change in direction or magnitude if the targets are excluded. We submitted the log-transformed reaction times from correct trials to a $2 \times 2 \times 2$ mixed-model ANOVA, with Participant Ethnicity as a between-subject factor, and Target Ethnicity and Object Type as within-subject factors (see Table 3 for means converted back to milliseconds). Across all participants, we again found a pronounced effect for Object, such that armed targets evoked responses more quickly than unarmed targets, $F(1, 45) = 347.82$, $p < .0001$. The Target Ethnicity \times Object interaction, or Shooter Bias, was also significant, $F(1, 45) = 14.75$, $p < .001$. Crucially, though, the magnitude of the bias did not depend on Participant Ethnicity, $F(1, 44) = 0.10$, $p = .75$. Examining the African American and White samples separately, we found that the Target Ethnicity \times Object interaction was significant for both, $F(1, 24) = 6.55$, $p = .017$ and $F(1, 20) = 8.01$, $p = .01$, respectively.

Simple effects tests again showed that, when the target was armed, participants decided to shoot more quickly if he was African American than if he was White, $F(1, 45) = 7.62$, $p = .008$. When the target was unarmed, participants pressed the *don't shoot* button more quickly if he was White than if he was African American, resulting in an identical test statistic, $F(1, 45) = 7.62$, $p = .008$. Neither simple effect depended on Participant Ethnicity, $F(1, 44) = 0.07$, $p = .79$, for the unarmed targets, and $F(1, 44) = 0.42$, $p = .52$, for the unarmed targets.

An analysis of the error rates revealed that the Target Ethnicity \times Object interaction was only marginal, $F(1, 45) = 3.24$, $p = .08$, and its magnitude did not depend on Participant Ethnicity, $F(1, 44) = 0.66$, $p = .42$.

General Discussion

In four studies, we attempted to recreate the experience of a police officer who, confronted with a potentially dangerous suspect, must decide whether or not to shoot. Our goal was to examine the influence of the suspect's ethnicity on that decision. We used a simplified videogame to present African American and White male targets, each holding either a gun or a nonthreatening object. Participants were instructed to shoot only armed targets. We reasoned that participants might use the stereotype, or schema, that African Americans are violent to help disambiguate the target

stimuli, and would therefore respond with greater speed and accuracy to stereotype-consistent targets (armed African Americans and unarmed Whites) than to stereotype-inconsistent targets (armed Whites and unarmed African Americans).

In Study 1, participants fired on an armed target more quickly when he was African American than when he was White, and decided not to shoot an unarmed target more quickly when he was White than when he was African American. In Study 2, we attempted to increase error rates by forcing participants to make decisions very quickly. Participants in this study failed to shoot an armed target more often when that target was White than when he was African American. If the target was unarmed, participants mistakenly shot him more often when he was African American than when he was White. A signal detection analysis of these data revealed that, although participants' ability to distinguish between armed and unarmed targets did not depend on target ethnicity, participants set a lower decision criterion to shoot for African American targets than for Whites. That is, if a target was African American, participants generally required less certainty that he was, in fact, holding a gun before they decided to shoot him. In Study 3, we returned to an analysis of reaction times, replicating the Ethnicity \times Object Type interaction (Shooter Bias) obtained in Study 1, and examining individual difference measures associated with the magnitude of that effect. Shooter Bias was more pronounced among participants who believed that there is a strong stereotype in American culture characterizing African Americans as aggressive, violent and dangerous; and among participants who reported more contact with African Americans. Prejudice and personal endorsement of the stereotype that African Americans are violent failed to predict Shooter Bias in the simple correlations, and their predictive power was no stronger among participants low in motivation to control prejudice. The fact that Shooter Bias in Study 3 was related to perceptions of the cultural stereotype, rather than prejudice or personally endorsed stereotypes, suggests that mere knowledge of the stereotype is enough to induce this bias. In Study 4, we obtained additional support for this prediction. Testing both White and African American participants, we found that the two groups display equivalent levels of bias.

The results of these studies consistently support the hypothesized effect of ethnicity on shoot/don't shoot decisions. Both in speed and accuracy, the decision to fire on an armed target was facilitated when that target was African American, whereas the decision not to shoot an unarmed target was facilitated when that target was White. This Shooter Bias effect is consistent with the results reported by Payne (2001). Payne primed participants with African American and White faces, and asked them to identify subsequent target objects as either hand tools or weapons. His results suggest that responses to hand tools were faster (and, in a second study, more accurate) when preceded by White, relative to African American, primes, whereas responses to weapons were faster (but no more accurate) when preceded by African American primes. This priming effect maps nicely onto our results. The consistency between our results and those obtained by Payne is particularly striking given methodological differences between the two paradigms. Four primary differences stand out. Payne used small, decontextualized and relatively simple images of faces (the center portion of the face) and objects, whereas our stimuli were very complex, with target individuals appearing against realistic backgrounds. Payne used a sequential priming task, whereas we

used simultaneous presentation of ethnicity and object. A consequence of Payne's priming task, which used a constant 200-ms stimulus onset asynchrony, is that the appearance of a prime in his task should have clearly indicated to participants that a target was imminent. Our task, however, presented targets at random intervals, with no prime, so that participants were never certain about when they would appear. Finally, whereas Payne asked his participants to identify a target object as a tool or a weapon, we asked our participants to decide whether or not to shoot a target person. Although both decisions depend on the presence of a weapon, the psychological implications of the two tasks are quite different. Payne's task was framed as a categorization judgment, whereas our task was characterized as a behavioral response. In spite of these distinctions, both paradigms reveal a pronounced effect of target ethnicity on reactions to weapons.

In line with Sagar and Schofield (1980), we have argued that ethnicity influences the shoot/don't shoot decision primarily because traits associated with African Americans, namely "violent" or "dangerous," can act as a schema to influence perceptions of an ambiguously threatening target. The relationship between cultural stereotype and Shooter Bias obtained in Study 3 provides support for this hypothesis. The subsequent finding that African Americans and Whites, alike, display this bias further buttresses the argument. It is unlikely that participants in our African American sample held strong prejudice against their own ethnic group (Judd, Park, Ryan, Brauer, & Kraus, 1995), but as members of U.S. society, they are, presumably, aware of the cultural stereotype that African Americans are violent (Devine & Elliot, 1995; Steele & Aronson, 1995). These associations, we suggest, may influence reactions to the targets in our videogame. Though ambient cultural associations may impact most members of U.S. society, it is certainly plausible that personal endorsement of stereotypes, and perhaps prejudice, will lead to even stronger negative associations with African Americans, potentially magnifying bias. (Though the data in Study 3, specifically the lack of a relationship between Shooter Bias and personal stereotype, offer little support for this argument, at present.)

It seems appropriate at this juncture to speculate on mechanisms that may underlie Shooter Bias. Our basic findings indicate that a target's ethnicity, though technically irrelevant to the decision task at hand, somehow interferes with participants' ability to react appropriately to the object in the target's hand. This interference seems roughly analogous to a Stroop effect, and research on this extensively studied phenomenon may provide a useful perspective from which to consider our results. The common Stroop experiment presents participants with a word, and requires them to identify the color of the ink in which that word is written (e.g., green ink). Performance on this simple task can be disrupted when the word, itself, refers to a different color than the ink (e.g., *RED* printed in green ink), relative to performance when the color of the ink and the referent of the word are the same (e.g., *GREEN* printed in green ink) or when the word does not refer to a color at all (e.g., *EGGS* printed in green ink). The Stroop paradigm, like our videogame, simultaneously presents participants with information that is relevant to the judgment at hand (ink color and object, respectively) as well as information that is irrelevant (word name and ethnicity, respectively). Participants need not process the irrelevant information to perform the task, but in both cases, the presence of incongruent information on the irrelevant dimension interferes

with participants' ability to process the relevant information. Researchers have suggested that, because we so frequently read the words that we see, reading occurs quickly. Ink naming, though, is an unusual and relatively cumbersome task. If these two processes occur in parallel, the quicker word reading may produce interference by winning a kind of horse race, getting to the finish line and influencing responses ahead of the slower ink-naming process, which eventually provides the definitive answer (Cohen, Dunbar, & McClelland, 1990; Posner & Snyder, 1975). Similarly, the speedy categorization of people into ethnic categories, described by Brewer (1988) and Fiske and Neuberg (1990), should quickly activate stereotypes and interfere with the unfamiliar and less automatic gun/no-gun judgment (see Figure 3). This analogy is not perfect, of course. Although it may be natural to read the word *RED* when it appears, the typical day-to-day response to an African American does not involve gunfire. However, to the extent that a person spontaneously associates an African American target with violence, the ethnicity of the target should conflict with the judgment that he is unarmed, and it may therefore inhibit the "don't shoot" response.

Cohen et al. (1990) characterized Stroop interference as an interaction between two variables: attention to the irrelevant dimension and the strength of the association between the incongruent information and the incorrect response. Both of these variables can moderate Stroop effects independently (see Walley, McLeod, & Khan, 1997; Walley, McLeod, & Weiden, 1994, for research on attention; see Lu & Proctor, 2001, for research on the strength of association). Though it is only speculation at present, we suggest that the two significant predictors of Shooter Bias in Study 3, cultural stereotype and contact, are important because they capture these two components of Stroop interference. We have already presented the argument that a cultural stereotype represents an associative link between African Americans and traits related to violence and danger. We further suggest that the role of contact in predicting Shooter Bias may reflect, at least in part, the other component of Stroop interference: attention to irrelevant ethnic cues. People who have had extensive contact with African Americans may have, over the course of that experience, learned to naturally parse the world in terms of ethnic categories. They may

be essentially schematic for ethnicity. Greater attention to ethnicity combined with an association between African Americans and violence should, from the Stroop perspective, magnify Shooter Bias. In line with this prediction, Payne (Payne, Lambert, & Jacoby, in press) has shown that asking participants to use ethnic cues in their judgments (like a person engaged in racial profiling) increases the magnitude of the automatic component in error responses in his task, relative to control participants who receive no special instructions. Of greater interest, asking participants to avoid using ethnicity in their judgments also increases the magnitude of the automatic component. This suggests that attention to the irrelevant ethnic cue may produce interference.

The Stroop conceptualization offers another, perhaps more hopeful, prediction. If Shooter Bias is, in part, a function of the automaticity with which ethnic cues are processed relative to the automaticity of the object cues (i.e., ethnicity's ability to win the horse race against the relevant dimension), the bias should be minimized by interventions that speed up the gun/no-gun decision. As the relevant decision becomes more automatic, the effect of the irrelevant dimension should weaken. Experimental research (MacLeod, 1998; MacLeod & Dunbar, 1988) as well as computer simulations (Cohen et al., 1990) have demonstrated that repeated training on ink-naming tasks, which should render that process more automatic, reduces Stroop interference. Similarly, training participants to quickly and effortlessly distinguish guns from cell phones may reduce Shooter Bias.

Though we have characterized Shooter Bias as a result of distorted interpretations of an ambiguous target, there are several stages at which this bias may actually be functioning. Before shooting, a participant must (a) *perceive* the object, (b) *interpret* the object as a gun with some degree of certainty, and (c) *decide* to press the "shoot" button once a criterion of certainty has been reached. Stereotypic schemata may theoretically affect any or all of these processes, and it is difficult to disentangle them theoretically, let alone empirically. Figure 3 depicts the three processing stages and how faster, more automatic processing along the irrelevant dimension (as suggested by the Stroop research) might bias each stage of relevant processing (the solid arrows). Throughout this article, we have argued that bias impacts the second stage of

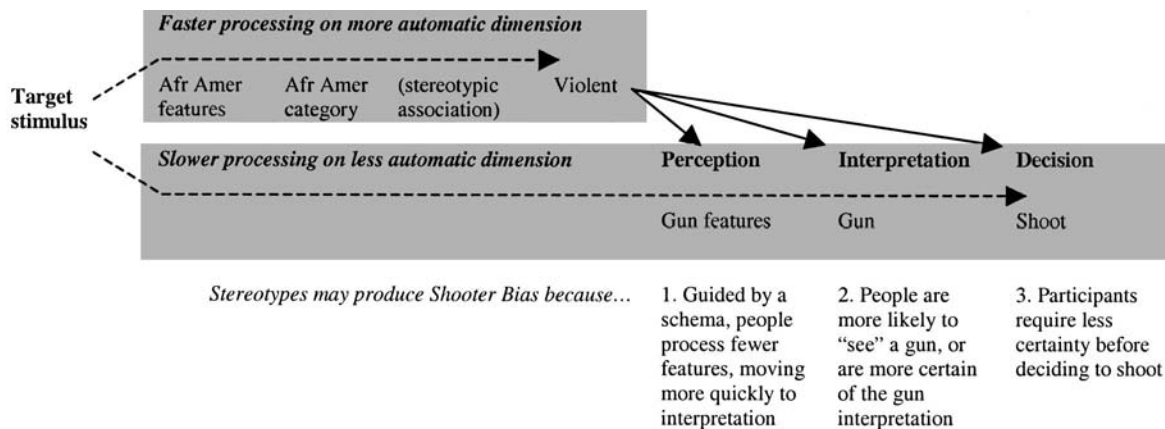


Figure 3. Faster, more automatic processing on the irrelevant ethnic dimension may bias participants' (a) perception of targets, (b) interpretation of targets, or (c) the criterion of certainty required for the "shoot" response. Afr Amer = African American.

this process, changing the interpretation of an ambiguous object. A participant who catches a glimpse of some elongated shape in the target's hand may draw on stereotypic associations, interpreting the shape as a gun if the target is African American but as a cell phone if he is White. Participants may almost "see" different objects.

One problem with this perspective is that we took pains to ensure that the objects presented in our target images were clearly identifiable. Even under time pressure, is it fair to characterize these objects as ambiguous? Certainly, very few of our participants actually misperceived the objects: our primary effects were in reaction times, not errors, and errors were consistently quite low. It is possible that the bias in reaction time represents the effects of stereotypes on actual perception of the object, not on its interpretation. von Hippel and his associates (von Hippel, Jonides, Hilton & Narayan, 1993) showed that when a participant has a relevant schema, he or she can infer the gist of a stimulus with very few perceptual details. Without a schema, though, more detailed perceptual encoding may be necessary. In the context of the current research, the stereotype of African Americans may influence the number features a participant must process to make the correct object identification. An African American target provides a schema relevant to guns, so participants who see just a few features of a gun quickly identify it and decide to shoot. A White target, perhaps, provides no useful schema, and participants must attend to more features of the gun in his hand before they recognize it, causing them to respond more slowly. von Hippel's research provides an elegant rationale for the differential speed required to shoot an armed target, but we are less confident that perceptual differences underlie reactions to the unarmed targets. Perceptual processes can only account for the simple effect of ethnicity among unarmed targets if we assume that White people stereotypically carry cell phones, wallets, coke cans, and cameras, and that this stereotype reduces the number of perceptual cues necessary to identify these objects relative to African Americans, where no cell phone stereotype exists. Empirically, it should be possible to test the viability of a perceptual encoding account of Shooter Bias. If perceptual differences drive ethnic bias, then memory for trivial details, such as the kind of gun or the color of the cell phone, which should reflect the extent of perceptual encoding, should differ as a function of target ethnicity.

Another, more macroscopic, alternative to our interpretation-based account is that Shooter Bias may reflect changes in the decision criterion that participants use. Bias would clearly emerge if participants require one level of certainty that the object is a gun when deciding to shoot African American targets, but have another, more stringent, criterion for Whites. Even if the perception and interpretation of an object do not differ as a function of target ethnicity (e.g., the participant is 75% certain that the object is a gun for both African American and White targets), a participant who requires 60% certainty for African American targets, but 80% certainty for Whites, will show Shooter Bias. Unfortunately, the current studies do not allow us to discern between the interpretation and decision criterion explanations. Though the signal detection terms, *sensitivity* and *criterion*, might foster an expectation that Study 2 should be able to resolve this question, that is not the case. Study 2 suggested that sensitivity was equal for African American and White targets, and that only the criterion differed. The criterion may differ, though, either because the certainty

needed to make the shoot/don't shoot decision differs with target ethnicity (bias in the decision stage), or because a given object in the hand of an African American target is *simultaneously* more likely to be perceived as a gun and less likely to be perceived as a non-gun, than the same object in a White target's hand (bias in the interpretation stage). The signal detection theory figure from Study 2 assumes that the average armed White target and the average armed African American target seem equally threatening, that the two gun distributions fall at the same point on the *x*-axis. As we have graphed it, the figure suggests that the criterion to shoot shifts down for African American targets. However, it is also possible that participants use the same criterion for White and African American targets, but generally perceive African Americans as more threatening. If this were the case, the criterion line in the chart for African American targets would have the same *x*-coordinate as the White criterion line, but the mean of the two African American distributions (both armed and unarmed) would seem to shift up on the dimension of perceived threat. Even using signal detection theory, we have no way to statistically disentangle these two possibilities in the current data.

Bias in the decision-making stage may be seen as consistent with ideomotor effects. Bargh, Chen, and Burrows (1996, Study 3), for example, found that participants primed with African American faces exhibited more aggressive behavior in response to a rude request from an experimenter. It is possible that the participants' behavior was still driven by bias in their interpretation, that those primed with African American actually perceived the experimenter as more hostile (along the lines of Devine, 1989). But Bargh and others (e.g., Chartrand & Bargh, 1996; Dijksterhuis, Aarts, Bargh, & van Knippenberg, 2000; Stapel & Koomen, 2001) have demonstrated direct behavioral priming effects in a number of situations designed to preclude interpretation-based bias. It is not unreasonable to suppose that participants in our studies were cued by a target's ethnicity to behave aggressively toward African American targets, shooting them more often and more quickly than Whites. Payne (2001), though, did not require a behavioral response. Because his task required that participants classify objects as guns or hand tools, rather than react violently, ideomotor effects cannot account for his findings. In the absence of more definitive evidence, and given the consistency between Payne's results and ours, parsimony argues for an interpretation-based explanation of Shooter Bias, rather than a criterion-based or ideomotor explanation.

These studies have demonstrated that the decision to shoot may be influenced by a target person's ethnicity. In four studies, participants showed a bias to shoot African American targets more rapidly and/or more frequently than White targets. The implications of this bias are clear and disturbing. Even more worrisome is the suggestion that mere knowledge of the cultural stereotype, which depicts African Americans as violent, may produce Shooter Bias, and that even African Americans demonstrate the bias. We understand that the demonstration of bias in an African American sample is politically controversial given the nature of this task, and we offer two considerations. First, the results of a single study are not definitive. Our findings should be replicated by researchers in other labs with different materials before generalizations are made. Second, our goals as psychologists include understanding, predicting, and controlling behavior. Ultimately, efforts to control (i.e., reduce or eliminate) any ethnic bias in the decision to shoot must

be based on an accurate understanding of how target ethnicity influences that decision, even if that understanding is politically or personally distasteful.

Though these studies suggest that bias in the decision to shoot may be widespread, it is not yet clear that Shooter Bias actually exists among police officers. The studies we report use exclusively lay samples, and there is no reason to assume that this effect will generalize beyond this population. There is even a possibility, suggested by literature on the Stroop effect, that police training may actually reduce Shooter Bias by rendering the gun/no-gun decision more automatic for officers. If this is the case, police might show less bias than the average college sophomore. Examining these sorts of effects in a sample of police officers is of the utmost importance.

The studies reported here suggest that Shooter Bias is present among White college students (Studies 1–3) and among a community sample that consists of both Whites and African Americans (Study 4). The effect is robust and clearly a cause for concern, no matter the underlying cause. On the basis of our data, though, bias does not seem to simply reflect prejudice toward African Americans, and there is reason to believe the effect is present simply as a function of stereotypic associations that exist in our culture. That these associations can have such potentially profound consequences for members of stigmatized groups is a finding worthy of great concern. Since the death of Amadou Diallo, New York has witnessed a number of similar, though less publicized, cases, and Cincinnati, Ohio, has added Timothy Thomas's name to the list of unarmed African American men killed by police officers. Social psychological theory and research may prove invaluable in the effort to identify, understand and eventually control processes that bias decisions to shoot (and possibly kill) a person, as a function of his or her ethnicity.

References

- Altemeyer, B. (1988). *Enemies of freedom: Understanding right-wing authoritarianism*. San Francisco: Jossey-Bass.
- Banaji, M. R., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology*, *68*, 181–198.
- Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology*, *71*, 230–244.
- Beckett, N. E., & Park, B. (1995). Use of category versus individuating information: Making base-rates salient. *Personality and Social Psychology Bulletin*, *21*, 21–31.
- Berry, G. L., & Mitchell-Kernan, C. (1982) Introduction. In G. L. Berry & C. Mitchell-Kernan (Eds.), *Television and the socialization of the minority child* (pp. 1–11). New York: Academic Press.
- Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.
- Chartrand, T. L., & Bargh, J. A. (1996). Automatic activation of impression formation and memorization goals: Nonconscious goal priming reproduces effects of explicit task instructions. *Journal of Personality and Social Psychology*, *71*, 464–478.
- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review*, *97*, 332–361.
- Cohen, J. D., MacWhinney, B., Flatt, M., & Provost, J. (1993). PsyScope: An interactive graphic system for designing and controlling experiments in the psychology laboratory using Macintosh computers. *Behavior Research Methods, Instruments & Computers*, *25*, 257–271.
- Cosby, C. O. (1994). *Television's imageable influences: The self-perceptions of young African-Americans*. Lanham, MD: University Press of America.
- Darley, J. M., & Gross, P. G. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology*, *44*, 20–33.
- 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.
- Dijksterhuis, A., Aarts, H., Bargh, J. A., & van Knippenberg, A. (2000). On the relation between associative strength and automatic behavior. *Journal of Experimental Social Psychology*, *36*, 531–544.
- Duncan, B. L. (1976). Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping of blacks. *Journal of Personality and Social Psychology*, *34*, 590–598.
- 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., 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., Lin, M., & Neuberg, S. L. (1999). The continuum model: Ten years later. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 231–254). New York: Guilford Press.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 1–74). New York: Academic Press.
- Gerbner, G., Gross, L., Morgan, M., & Signorielli, N. (1986). Living with television: The dynamics of the cultivation process. In J. Bryant & D. Zillman (Eds.), *Perspectives of media effects* (pp. 14–40). Hillsdale, NJ: Erlbaum.
- 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.
- Gray, H. (1989) Television, Black Americans, and the American dream. *Critical Studies in Mass Communication*, *6*, 376–386.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley. (Reprinted 1974, Huntington, NY: Krieger).
- Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (pp. 1–68). Hillsdale, NJ: Erlbaum.
- Hilton, J. L., & von Hippel, W. (1990). The role of consistency in the judgment of stereotype-relevant behaviors. *Personality and Social Psychology Bulletin*, *16*, 430–448.
- Jacobs, J. E., & Eccles, J. S. (1992). The impact of mothers' gender-role stereotypic beliefs on mothers' and children's ability perceptions. *Journal of Personality and Social Psychology*, *63*, 932–944.
- 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., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General*, *122*, 139–154.
- Judd, C. M., Park, B., Ryan, C. S., Brauer, M., & Kraus, S. (1995). Stereotypes and ethnocentrism: Diverging interethnic perceptions of

- African American and White American youth. *Journal of Personality and Social Psychology*, 69, 460–481.
- Krueger, J., & Rothbart, M. (1988). The use of categorical and individuating information in making inferences about personality. *Journal of Personality and Social Psychology*, 55, 187–195.
- Lepore, L., & Brown, R. (1997). Category and stereotype activation: Is prejudice inevitable? *Journal of Personality and Social Psychology*, 72, 275–287.
- Locksley, A., Borgida, E., Brekke, N., & Hepburn, C. (1980). Sex stereotypes and social judgment. *Journal of Personality and Social Psychology*, 39, 821–831.
- Lu, C.-H., & Proctor, R. W. (2001). Influence of irrelevant information on human performance: Effects of S–R association strength and relative timing. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 54(A), 95–136.
- MacLeod, C. M. (1998). Training on integrated versus separated Stroop tasks: The progression of interference and facilitation. *Memory & Cognition*, 26, 201–211.
- MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14, 126–135.
- MacMillan, N. A., & Creelman, C. D. (1991). *Detection theory: A users guide*. Cambridge, England: Cambridge University Press.
- Macrae, C. N., Milne, A. B., & Bodenhausen, G. V. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology*, 66, 37–47.
- 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.
- Neuberg, S. L., & Newson, J. T. (1993). Personal need for structure: Individual differences in the desire for simpler structure. *Journal of Personality and Social Psychology*, 65, 113–131.
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, 81, 181–192.
- Payne, B. K., Lambert, A. J., & Jacoby, L. L. (in press). Best laid plans: Effects of goals on accessibility bias and cognitive control in race-based misperceptions of weapons. *Journal of Experimental Social Psychology*.
- Pettigrew, T. F., & Tropp, L. R. (2000). Does intergroup contact reduce prejudice: Recent meta-analytic findings. In S. Oskamp (Ed.), *Reducing prejudice and discrimination* (pp. 93–114). Mahwah, NJ: Erlbaum.
- Plant, E. A., & Devine, P. G. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology*, 75, 811–832.
- 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.
- Pratto, F., Sidanius, J., Stallworth, L. M., & Malle, B. F. (1994). Social dominance orientation: A personality variable predicting social and political attitudes. *Journal of Personality and Social Psychology*, 67, 741–763.
- Rothbart, M., & Birrell, P. (1977). Attitude and perception of faces. *Journal of Research in Personality* 11, 209–215.
- Sagar, H. A., & Schofield, J. W. (1980). Racial and behavioral cues in black and white children's perceptions of ambiguously aggressive acts. *Journal of Personality and Social Psychology* 39, 590–598.
- Stapel, D. A., & Koomen, W. (2001). The impact of interpretation versus comparison mindsets on knowledge accessibility effects. *Journal of Experimental Social Psychology*, 37, 134–149.
- Steele, C. M., & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology*, 69, 797–811.
- Stroman, C. A. (1986). Television viewing and self-concept among Black children. *Journal of Broadcasting and Electronic Media*, 30, 87–93.
- SuberviVelez, F. A., & Necochea, J. (1990). Television viewing and self-concept among Hispanic American children—A pilot study. *Howard Journal of Communications*, 2, 315–329.
- Thompson, M. M., Naccarato, M. E., Parker, K. C. H., & Moskowitz, G. B. (2001). The personal need for structure and personal fear of invalidity measures: Historical perspectives, current applications, and future directions. In G. B. Moskowitz (Ed.), *Cognitive social psychology: The Princeton symposium on the legacy and future of social cognition* (pp. 19–39). Mahwah, NJ: Erlbaum.
- von Hippel, W., Jonides, J., Hilton, J. L., & Narayan, S. (1993). Inhibitory effect of schematic processing on perceptual encoding. *Journal of Personality and Social Psychology*, 64, 921–935.
- Walley, R. E., McLeod, B. E., & Khan, M. (1997). Tests of a translational model of Stroop interference: Translation or attention? *Canadian Journal of Experimental Psychology*, 51, 10–19.
- Walley, R. E., McLeod, B. E., & Weiden, T. D. (1994). Increased attention to the irrelevant dimension increases interference in a spatial Stroop task. *Canadian Journal of Experimental Psychology*, 48, 467–492.
- Wittenbrink, B., Judd, C. M., & 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 February 11, 2002

Revision received July 5, 2002

Accepted July 5, 2002 ■