

Targets of discrimination: Effects of race on responses to weapons holders

Anthony G. Greenwald,* Mark A. Oakes, and Hunter G. Hoffman

Department of Psychology, University of Washington, P.O. Box 351525, Seattle, WA 98195-1525, USA

Received 10 June 2002; revised 22 October 2002

Abstract

Rapid actions to persons holding weapons were simulated using desktop virtual reality. Subjects responded to simulated (a) criminals, by pointing the computer's mouse at them and left-clicking (simulated shooting), (b) fellow police officers, by pressing the spacebar (safety signal), and (c) citizens, by inaction. In one of two tasks Black males holding guns were police officers while White males holding guns were criminals. In the other, Whites with guns were police and Blacks with guns were criminals. In both tasks Blacks or Whites holding harmless objects were citizens. Signal detection analyses revealed two race effects that led to Blacks being incorrectly shot at more than Whites: a perceptual sensitivity effect (when held by Blacks guns were less distinguishable from harmless objects) and a response bias effect (objects held by Blacks were more likely to be treated as guns).

© 2003 Elsevier Science (USA). All rights reserved.

Reporting a non-existent fire is a false alarm. More generally, false alarms are errors of acting as if an important event has occurred when it has not. This paper is concerned with a specific type of false alarm—acting as if a weapon is present when it is not. There have been many such false alarms by police officers who believed themselves to be in the presence of a threatening, armed antagonist when they were confronted only by an unarmed citizen. The error in these cases can be the tragic error of mistakenly shooting at, perhaps killing or seriously injuring, a non-threatening citizen. A disturbingly frequent component of these erroneous shootings has been that the unarmed citizen/victim was Black (Fyfe, 1982).¹ In trying to understand why these false alarms occur, it is natural to think of the social psychological concept of stereotype, defined as the association of a group (such as elderly, Asian, or female) with a trait (such as conservatism, diligence, or nurturance). In the case of erroneously acting as if Blacks are armed, the relevant stereotype is an association that links the group, Blacks, to the trait of dangerousness or criminality.

The concept of stereotype captures the existence of group-trait associations, but stops short of explaining how these associations can lead to erroneous actions such as accidental shootings. Banaji and Greenwald (1995) proposed to fill this explanatory gap with a well-established psychophysical method for analyzing perceptual decisions, signal detection theory (SDT; Green & Swets, 1967). SDT permits analysis of repeated encounters of a critical situation to determine the extent to which false alarms are due to (a) failures of *perceptual sensitivity*—limited ability to distinguish the critical event's presence from its absence, or (b) *response bias*—increased reporting of the critical event whether or not it has occurred.

To illustrate: Suppose that a radiologist's record shows a history in which lung biopsies, recommended after examining X-rays, have revealed lung cancer less often for men than for women patients. Signal detection analysis can determine whether the radiologist's greater false alarms (orders of biopsies that reveal no disease) for men are due to greater difficulty in distinguishing healthy from malignant tissue on men's X-ray images (a perceptual sensitivity cause) or, alternately, to a generally greater tendency to advise biopsies for men (a response bias cause).²

* Corresponding author. Fax: 1-206-685-3157.

E-mail address: agg@u.washington.edu (A.G. Greenwald).

¹ See also www.washingtonpost.com/wp-dyn/metro/md/prince-georges/government/police/shootings/; www.prospect.org/webfeatures/2001/05/sobel-l-05-08.html.

² In order actually to apply SDT to the radiologist example it would be necessary also to know the malignant versus non-malignant status for patients for whom biopsies were not ordered.

This paper follows Banaji and Greenwald (1995) in merging the social psychological construct of stereotype with the analytic methods of SDT. Banaji and Greenwald used SDT to analyze false alarms of incorrectly judging familiar-sounding names as famous more often when the names were male than when they were female. Their experiment built upon the observation of a *false fame effect* by Jacoby, Kelley, Brown, and Jasechko (1989): A day after being exposed to a list of non-famous names, Jacoby et al.'s subjects judged those names (incorrectly) as belonging to famous people more than comparable names that had not been seen the day before. Banaji and Greenwald (1995) found that such false-fame alarms occurred more when the names were male than when they were female. With the aid of SDT, Banaji and Greenwald further showed that the increased male-name false alarms were due to a response bias of being more ready to judge familiar-sounding male than familiar-sounding female names as famous. SDT simultaneously showed that these false alarms did not involve impaired ability to perceptually discriminate presence versus absence of fame for female (relative to male) names.

The present research used SDT to determine whether increased weapons false alarms (WFAs, hereafter) for Blacks are due to reduced perceptual sensitivity (i.e., less ability to distinguish weapons from harmless objects when these objects are held by Blacks than by Whites) or increased response bias (i.e., increased tendency to respond to any object held by a Black as a weapon).

Some recent studies have suggested that increased WFAs to Blacks are due to response bias (Correll, Park, Judd, & Wittenbrink, 2002; Payne, 2001; Payne, Lambert, & Jacoby, 2002). Payne and colleagues asked subjects to rapidly distinguish computer-presented images of weapons from images of tools. Each image was preceded by a briefly flashed *prime*, which was either a Black face or a White face. These experiments found that subjects were more likely to press a computer key that indicated identification of the object as a weapon when the face prime was Black than when it was White, regardless of whether the object was a tool or a weapon. At the same time, subjects showed no difference in perceptual ability to distinguish tools from weapons as a function of the prime's race. The WFAs observed by Payne and colleagues were therefore shown to involve a response bias (increased tendency to give the weapon response after the Black prime), rather than reduced (perceptual) ability to distinguish tools from weapons after the Black prime. A similar conclusion was reported by Correll et al. (2002).

In the studies by Payne (2001), Payne et al. (2002), and Correll et al. (2002) the only cue to which subjects were obliged to attend on each trial was a weapon or non-weapon. Race was always an incidental cue. The task of the present research was deliberately made more

complex. Subjects were instructed not simply to detect the presence of weapons, but to discriminate threatening from non-threatening weapons—a discrimination that is also required in police activities. This greater complexity was implemented in a task in which both simulated police and simulated criminals had weapons. Additionally, in order to assure that subjects would attend to race, race was used as the cue that discriminated police from criminals, with criminals sometimes White and sometimes Black.

Experiments 1 and 2

The research used a *weapons task*, a desktop virtual-reality simulation of responding rapidly to armed and unarmed persons who varied in race. In two experiments, subjects were asked to react rapidly to guns and harmless objects held by Blacks and Whites (these human figures in the simulation are hereafter described as *targets*). Data in the form of hit rates (proportions of correct responses to the presence of a gun) and false alarm rates (proportions of errors of responding to a harmless object as if it were a gun) were used to compute SDT's perceptual sensitivity (d') and response bias ($\log \beta$ ['beta']) measures. These two measures were then used in further analyses to determine if subjects' time-pressured judgments showed effects of target race on perceptual sensitivity or on response bias. If reduced perceptual sensitivity underlies race effects on weapon judgments, subjects should be less able to discriminate a gun from a harmless object when held by a Black than when held by a White. This result would appear as reduced hit rates or increased false alarm rates (or both) when the target is Black. If response bias underlies stereotype effects, subjects should show generally elevated gun-present judgments for Blacks compared to Whites. This result would appear as an elevation in both hit and false alarm rates when the target is Black.³

Method

There was one difference between Experiments 1 and 2. Experiment 1's weapons task required subjects to

³ In addition, both experiments examined whether individual differences in stereotype effects on weapons judgments could be predicted by measures of implicit race preference and by self-report measures of race prejudice (cf. Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio, Jackson, Dunton, & Williams, 1995). Although, as will be shown, the weapons task proved to be reliably sensitive to race effects, its sensitivity and bias measures lacked sufficient internal consistency to be useful in correlational analyses. Unsurprisingly (for measures with low internal consistency), correlations of these measures with all other measures were generally weak and non-significant. Accordingly, they are not discussed further.

respond to gun-holding targets within 800 ms. Experiment 2 extended this deadline to 900 ms in order to increase accuracy.

Subjects

Sixty-nine subjects from the University of Washington's undergraduate Psychology subject pool participated in Experiment 1 and 63 participated in Experiment 2. Data for six of these 132 subjects were not obtained, four due to apparatus failure, and two who withdrew because of discomfort with the task on the basis of its initial description. Data for 20 others proved unusable, three because of insufficient fluency with English, eight due to an experimenter's procedural error, eight because their non-use of one of the three instructed response options suggested that they had misunderstood the instructions, and one who (it was later discovered through a routine check) had previously participated in the experiment. This left 106 subjects, 54 in Experiment 1 (28 male, 25 female, and one sex not reported), and 52 in Experiment 2 (27 male, 21 female, and four sex not reported).

Desktop virtual-reality weapons task

Subjects were asked to play the role of a plainclothes police officer. Their task was to take rapid action in response to three categories of simulated targets—criminals, fellow police officers, and citizens. Subjects were given less than a second (800 ms in Experiment 1 and 900 ms in Experiment 2) to: (a) point and shoot at criminals (pointing and left-clicking the computer mouse) or (b) respond with a designated safety signal to fellow police officers (pressing the spacebar). They were instructed to make no response to citizens, who were distinguishable from the other two targets because they were holding harmless objects (a camera, a beer bottle, or a flashlight) rather than guns.

Targets were indistinguishable by dress—all appeared in casual street clothes. Subjects could distinguish police officers and criminals from citizens by virtue of the former holding guns and the latter holding harmless objects. The only variable that distinguished police officers from criminals was race. Each subject performed two variations of the weapons task, one in which White targets were criminals and Blacks were police officers (White criminal condition, W), and the other in which these roles were switched (Black criminal condition, B). Subjects performed two blocks of trials for each of the W and B conditions. As a counterbalancing variation, the four blocks were administered in either WBBW or BWBW order. The 52 trials in each block were a random sequence of 18 criminals, 18 police officers, and 16 citizens (eight Black and eight White).

Subjects were informed that the targets would appear from behind one of two screen objects that had the ap-

pearance of garbage dumpsters, one in the middle-left and one in the middle-right of the screen. To assure that mouse movement would start equidistant from the two possible target locations on each trial, subjects were required to position the mouse pointer over a small green box at mid-screen before pressing the right mouse key to begin a trial. Headphone audio feedback for each response was designed to keep subjects both involved and mindful of the non-laboratory situation that was being simulated. A mouse left-click made within the deadline (800 or 900 ms) of a figure rising from one of the dumpsters produced the sound of a silencer-equipped gun being fired. If the subject failed to respond within the deadline to a gun-wielding target of either race, a loud gunshot was heard, indicating either that the criminal had fired at the subject or that the fellow police officer had fired because of not receiving the safety signal (spacebar response) in time. Incorrectly shooting at a police officer or a citizen within the deadline produced a loud scream, which indicated the unintended injury to the target. A timely spacebar response to a police officer produced an audible acknowledgment, "Okay." Pressing a spacebar in response to a citizen (for which the response should have been inaction) produced an audible, question-inflected "Huh?"

Subjects participated in individual booths at PC-type desktop computers. Experiment 1 administered auxiliary self-report and implicit attitude measures (see Footnote 3) before the weapons task, whereas Experiment 2 administered these after the weapons task.

Results

Weapons task

The main hypotheses concerned the effect of target race on signal detection measures of perceptual sensitivity and response bias. Computation of d' (sensitivity) and $\log \beta$ (bias) required first classifying subject responses as hits and false alarms. Hits were defined as giving the correct response to a target who was holding a gun (either left-clicking at a criminal or pressing the spacebar for a police officer, and doing so within the allowed deadline). False alarms were defined as making one of these rapid responses inappropriately to a citizen (a person holding a harmless object).

The design required identification of hits and false alarms separately for two targets (criminals and police) in each of the two conditions (White and Black criminal). In the White criminal condition, *criminal hits* were rapid shooting (mouse-clicking) at Whites holding guns and *criminal false alarms* were rapid shooting at Whites holding harmless objects. For the Black criminal condition, the corresponding hit and false alarm measures were computed for trials with Blacks holding guns and

harmless objects, respectively. *Police hits* in the White criminal condition were trials on which the subject rapidly pressed the space bar in response to a Black with a gun and *police false alarms* were rapid spacebar responses to Blacks holding harmless objects. For the Black criminal condition, the corresponding hit and false alarm measures were obtained from trials with Whites holding guns and harmless objects, respectively.^{4,5}

The analyses included Experiment 1 versus Experiment 2 as a design factor. There were some expected main effects of experiment due to the greater accuracy of subjects with Experiment 2's longer (900 ms) deadline. However, because the effects of experiment did not in any way qualify the findings to be reported, it is appropriate to present results combined over the two experiments.

Hit and false alarm rates. Fig. 1 shows mean hit and false alarm rates, classified by task and block. Hit rates rose through the first three blocks of the experiment, averaging 55, 63, 70, and 69% for Blocks 1, 2, 3, and 4, respectively. False alarm rates similarly improved (i.e., decreased), averaging 39, 37, 37, and 33% for Blocks 1–4.

Format for describing results. It was initially expected that the design would allow a within-subjects analysis of all design factors. However, the within-experiment improvements in performance shown in Fig. 1 precluded this strategy, because these improvements were confounded with race-of-target effects. Effects of this unexpected confounding were avoided by analyzing the data as two sub-experiments—one consisting of Blocks 1 and 4, and the other consisting of Blocks 2 and 3. (Recall that all subjects performed either the White or Black criminal condition in both Blocks 1 and 4, and the other condition in Blocks 2 and 3.) Each of the hypothesis tests to be reported was conducted in both sub-experiments. Consequently, all results are reported with two significance tests—the first for analysis of the sub-experiment in Blocks 1 and 4, and the second for analysis of the sub-experiment in Blocks 2 and 3. Results are presented as statistically significant only if they were separately significant ($\alpha = .05$, two-tailed) in both sub-experiments. For none of the reported findings was there a statistically significant interaction effect with the ex-




Fig. 1. Hit rates (A) and false alarm rates (B) for four categories of targets. For police targets, hits were pressing the spacebar prior to a deadline and false alarms were pressing the spacebar prior to the deadline in response to unarmed citizens of the same race. For criminal targets, hits were firing (left-clicking the mouse) prior to the deadline and false alarms were pre-deadline firing at unarmed citizens of the same race. Data are combined from Experiments 1 and 2, and are presented separately by blocks of trials. Each mean is based on an N of 53. Error bars are standard errors.

periment factor—that is, all reported results were similarly apparent in both Experiments 1 and 2.

Perceptual sensitivity (d') results. For the perceptual sensitivity measures (see Fig. 2A) higher scores indicated greater ability to discriminate weapons from harmless objects. The main result apparent in Fig. 2A was greater perceptual sensitivity for criminal targets than for police targets. This appears as greater height of the two right-side than the two left-side bars in each of the four blocks, $F_s(1, 102) = 70.6$, $p = 10^{-13}$, and 124.9 , $p = 10^{-19}$. Greatest interest was in the possibility of higher perceptual sensitivity for White than Black targets. This was indeed found, appearing as greater height of the right than the left bar in each adjacent pair of bars for both police and criminal targets, $F_s(1, 102) = 13.3$, $p = .0004$, and 18.9 , $p = 10^{-5}$. Averaged over both police and criminal targets and all blocks of trials, effect size of the effect of race on perceptual sensitivity was

⁴ Signal detection measures are not computable when hit or false alarm rates are either 0 or 1. These 'end-point' values occurred several times in the data set and were managed by replacing the value of 0 with $.25/N$ (where N is the number of trials on which the hit or false alarm rate is based) and by replacing a value of 1 with $1 - .25/N$.

⁵ Occasional spacebar responses to criminals and mouse-click responses to police were treated as the equivalent of inaction for the analyses to be reported. Additional analyses indicated that very similar conclusions were supported when these responses were instead treated as the equivalent of correct actions.

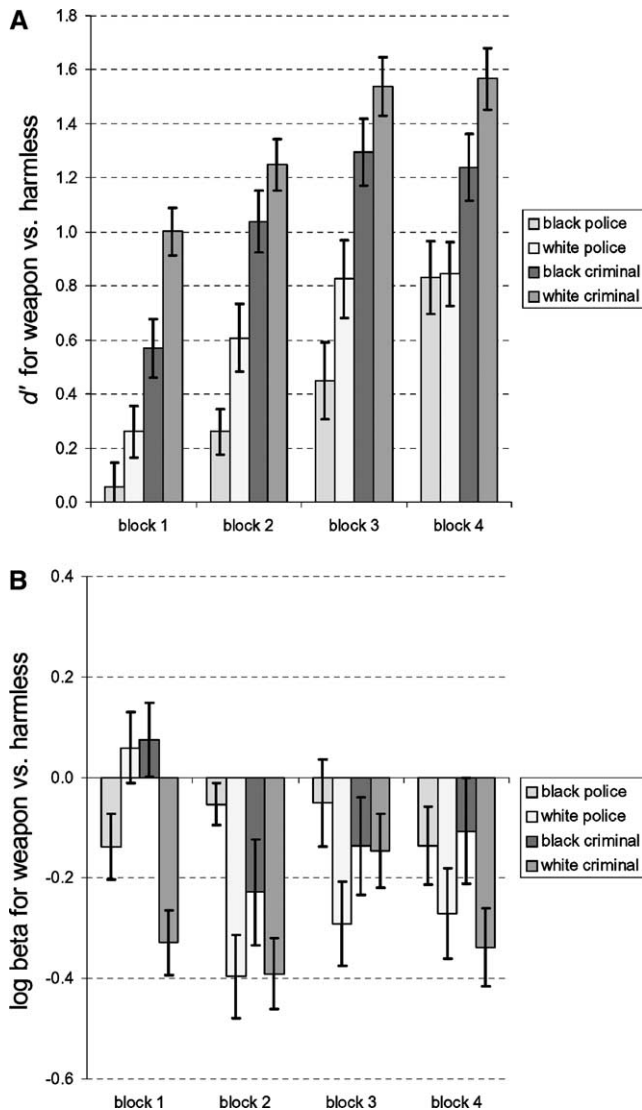


Fig. 2. Perceptual sensitivity (A; SDT's d' measure) and response bias (B; SDT's $\log \beta$ measure) for four categories of targets. Higher values of d' indicate greater accuracy in discriminating weapons from harmless objects. Higher values of $\log \beta$ indicate greater readiness to respond to the object in the target's hand as a weapon. Data are combined from Experiments 1 and 2, and are presented separately by blocks of trials. Each mean is based on an N of 53. Error bars are standard errors.

$d = .31$.⁶ The improvement of performance as the experiment progressed was reflected in greater perceptual sensitivity in the second half than the first half of the experiment, appearing as greater heights for Block 4

⁶ The test of significance for race of target (Black versus White) was provided by the interaction effect of target role and race of criminal. To explain: For subjects in the Black criminal condition, the police target was White and the criminal target was Black; for those in the White criminal condition, the police target was Black and the criminal target was White. With this design, race difference is synonymous with the interaction between race of criminal and target role (police versus criminal).

than Block 1, $F(1, 102) = 74.8, p = 10^{-13}$, and for Block 3 than Block 2, $F(1, 102) = 10.9, p = .001$. Finally, and not shown in Fig. 2A, Experiment 2's longer response deadline expectably yielded greater perceptual sensitivity in Experiment 2 than Experiment 1, $F_s(1, 102) = 10.2, p = .002$, and $7.30, p = .008$.

Response bias ($\log \beta$) results. The SDT bias measure (β) is a ratio measure that typically has non-normal distributions. Accordingly, it was analyzed (and is displayed in Fig. 2B) in the form of its log transformation, which has distributional properties better suited to statistical analyses. Higher values of $\log \beta$ indicate greater readiness to give the response appropriate for presence of a gun (i.e., left-click for criminal targets or spacebar for police targets). Fig. 2B shows that the response bias measure was, on average, numerically negative, which indicates overall greater readiness to give the harmless-object-appropriate response than the weapon-appropriate response. Greatest interest was in the possibility of finding higher response bias (greater tendency to give the weapon-appropriate response) for Black than White targets. With the exception of the two police targets in Block 1 (leftmost two bars in Fig. 2B) values of the $\log \beta$ measure were indeed higher when targets were Black than White (greater height of left than right bar in each adjacent pair of bars). This corresponded to statistically significant effects of race, $F_s(1, 102) = 5.43, p = .02$, and $8.35, p = .005$. Averaged over both types of targets and all blocks, effect size for the effect of race on response bias ($\log \beta$) was $d = .27$, slightly smaller than the effect size (0.31) for the effect of race on perceptual sensitivity (d'). Additionally, $\log \beta$ was lower in Experiment 1 than Experiment 2, $F_s(1, 102) = 4.19, p = .04$, and $6.98, p = .01$, which most likely indicates the greater difficulty of giving the weapon-appropriate response in timely fashion with Experiment 1's shorter response deadline.

Discussion

The results of two simulation experiments, analyzed using SDT (signal detection theory), demonstrated both that (a) subjects had greater difficulty distinguishing weapons from harmless objects when the weapons were in the hands of simulated Blacks than Whites and (b) subjects were response-biased in the sense of giving the weapon-appropriate response more readily to Black than to White targets. Both findings were consistent with the spate of disturbing episodes (see Footnote 1) in which law enforcement officers have mistakenly responded to Blacks holding harmless objects as if they were brandishing threatening weapons. At the same time, the perceptual sensitivity finding was surprising because it diverged from recent results that had identified only response bias as the source of weapons false alarms (WFAs) associated with race stereotypes.

Perceptual sensitivity and response bias

Three previous studies with aims very similar to the present research had shown that stereotype-related WFAs were due only to response bias—that is, greater readiness to treat an object as a gun when the situation included a Black (compared to a White) racial cue (Correll et al., 2002; Payne, 2001; Payne et al., 2002). The present research, in part, replicated those three findings by finding the same effect of race on response bias. However, the present experiments obtained a very clear additional effect of race on perceptual sensitivity. That is, subjects found it harder to distinguish weapons from harmless objects when the person holding one of these objects was Black. This perceptual sensitivity finding was observed approximately equally no matter whether Blacks with guns were understood to be police officers (effect size $d = 0.28$) or criminals ($d = 0.34$).

In attempting to understand why the present studies found a perceptual sensitivity effect that did not occur in the three related studies, the obvious strategy is to look for procedural differences. The main difference of the present research from the three other studies was the greater complexity of subjects' task in the present research. The weapons task of the present experiments required subjects to discriminate non-threatening from threatening weapons and to use race as a cue in making this choice—Blacks with guns had to be distinguished from Whites with guns.⁷ Another part of the added complexity of the present experiments was their requirement of a move-and-click mouse response rather than a key press. The possible role of this added processing load of the present task in producing the perceptual sensitivity effect will require additional research to evaluate.

The studies by Payne (2001) and Payne et al. (2002) were analyzed using process dissociation analysis (PD; see Jacoby, 1991) rather than SDT. PD yields two parameters that are analogous to SDT's measures of sensitivity and bias. However, because PD and SDT use different underlying decision process assumptions, their measures are computed differently. To assess the possible effect of the different computational procedures, additional analyses of the present data were conducted using the PD measures. These additional analyses agreed with the SDT analyses in showing statistically significant effects of race on both the PD discriminability measure (analogous to SDT's perceptual sensitivity measure), and the PD bias measure (analogous to SDT's response bias measure). Accordingly, it is clear that the difference in conclusions from the other recent studies is not due to these computational differences.

Why did the present experiments find perceptual sensitivity effects that did not appear in previous research? Although this question cannot be answered confidently, a plausible suggestion can be based on the assumption that, for many or most subjects, Black targets demanded more attention than White targets. With the less complex tasks of previous research, this extra attentional demand may not have sufficed to impair discrimination of objects held by Black targets. However, the added processing demand of the present research's more complex task may have done just that.

Ecological validity of research models

Research such as the present experiments and those of Payne (2001), Payne et al. (2002), and Correll et al. (2002) aim in part to provide a research model of the reactions of law enforcement officers to threatening situations involving weapons. It is therefore useful to consider the extent to which procedures of these various experiments succeed in capturing important properties of relevant natural situations. The tasks in all of these studies obliged subjects to respond rapidly. Therefore, all were intended to capture the time pressure experienced by police when facing a possibly armed suspect. The respects in which the tasks differed were in their complexity and in whether or not race of target was a cue to which attention was required.

To the extent that natural situations in which police encounter possibly armed suspects: (a) are themselves complex and (b) contain race as a cue that must be attended, the present task may provide an appropriate model. However, race is *not* an obligatory cue of the natural situation—police officers can function effectively without having to determine the race or ethnicity of a person holding a gun. Indeed, police may be most effective in identifying threatening weapons if they are not simultaneously distracted by attending to race. At the same time, law enforcement personnel are trained to observe and to remember information that is useful in identifying suspects. This information prominently includes race, along with sex, age, height, weight, clothing, and hair style. Even without considering police training to attend to race as an identifying cue, race may inevitably be attended by virtue of automatically activated stereotypes that pervade the general population (e.g., Devine, 1989). The present experiments' requirement that subjects attend to race may therefore have enhanced their ecological validity. However, this suggestion should be treated skeptically.

Conclusion

In combination with the findings of Payne (2001), Payne et al. (2002), and Correll et al. (2002), the present research suggests that race stereotypes have two paths

⁷ The authors thank Keith Payne for pointing out the possibility of explaining differences among results in terms of the greater complexity of the task in the present research.

for influencing actions in the face of a possible weapons threat. Race of target can affect both (a) perceptual ability to discriminate a weapon from a harmless object and (b) bias to respond as if a weapon is present. This two-effect conclusion does not simplify the task of those who seek to reduce false-alarm errors of firing at unarmed suspects. However, even if only one of the two processes were involved, the task of designing a corrective strategy would not be simple because the time pressure of natural situations in which WFAs occur likely brings into play cognitive processes that are automatic and therefore difficult to control.

In seeking practical benefit from this collection of laboratory experiments, perhaps their methods will prove more useful than their theoretical explanations. The several experiments combine to show the possibility of constructing simulations that can both (a) reveal effects of race on weapons-related errors and (b) indicate whether race operates through perceptual impairment or response bias. Use of these or similarly conceived simulations may provide law-enforcement officials the means to assess effectiveness of training in overcoming race-influenced errors.

Acknowledgments

This research was supported by grants from National Institute of Mental Health, MH-41328, MH-01533, and MH-57672. The authors thank Mahzarin R. Banaji and Donald P. Green for useful discussions before the research was initiated; and Ari Hollander for computer programming of the simulation task used in this experiment.

References

- Banaji, M. R., & Greenwald, A. G. (1995). Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology*, *68*, 181–198.
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology*, *83*, 1314–1329.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.
- Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology*, *33*, 510–540.
- 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 bonafide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027.
- Fyfe, J. J. (1982). Blind justice: Police shootings in Memphis. *Journal of Criminal Law and Criminology*, *73*, 707–722.
- Green, D. M., & Swets, J. A. (1967). *Signal detection theory and psychophysics*. New York: Wiley.
- 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., Kelley, C. M., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology*, *56*, 326–338.
- 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. (2002). Best laid plans: Effects of goals on accessibility bias and cognitive control in race-based misperceptions of weapons. *Journal of Experimental Social Psychology*, *38*, 384–396.