Reports

Dangerous enough: Moderating racial bias with contextual threat cues☆

Joshua Correll [a,*], Bernd Wittenbrink [a], Bernadette Park [b], Charles M. Judd [b], Arina Goyle [a]

[a] University of Chicago, USA
[b] University of Colorado, USA

ABSTRACT

Research shows that participants shoot armed Blacks more frequently and quickly than armed Whites, but make don't-shoot responses more frequently and quickly for unarmed Whites than unarmed Blacks. We argue that this bias reflects the perception of threat — specifically, threat associated with Black males. Other danger cues (not just race) may create a similar predisposition to shoot, and if these cues promote shooting when the target is White, they should attenuate racial bias. We embedded targets in threatening and safe backgrounds. Racial bias was evident in safe contexts but disappeared when context signaled danger, and this reduction was largely due to an increased tendency to shoot White targets.

© 2010 Elsevier Inc. All rights reserved.

In April, 2001, Timothy Thomas had several outstanding warrants for non-violent traffic violations. Late one night, an officer approached him, and Thomas ran, ultimately turning down a dark alley. Another officer entered the alley and fired a single round, which struck the unarmed Thomas in the chest, killing him. The officer later reported thinking that Thomas had a gun (Larson, 2004, April 10). Thomas' death – tragic in its own right – was complicated by the fact that the victim was Black. Like other police shootings of unarmed Black men in recent years (e.g., Amadou Diallo, Sean Bell, Officer Omar Edwards), Thomas' death sparked community protests and fueled speculation about racism, raising a simple question: would the officer have fired if Thomas had been White?

The difficulty of attributing Thomas' death to racism is that other situational factors may have affected the officer's decision. For example, the dark alley in which the shooting occurred may have created a sense of confusion or danger that facilitated the decision to fire. It is not surprising that police attend to cues about their surroundings, using more extreme force in neighborhoods with higher crime and poverty, and greater proportions of minority residents (Geller, 1982). Werthman and Piliavin's (1967) ecological contamination hypothesis proposed that the environment taints officers' perceptions: people in a dangerous neighborhood actually seem more dangerous, and police treat them accordingly. In an effort to test this idea, Terrill and Reisig (2003; but see Smith, 1986) examined officers in 24 beats in Indiana and Florida. They found that, although officers used more force with Black (rather than White) suspects, the effect of race disappeared when statistical models controlled for neighborhood disadvantage. They argued that police treat Black suspects more harshly not simply because the suspects are Black but because they typically encounter Black suspects in more dangerous neighborhoods. Inferring support for the ecological contamination hypothesis, the authors suggested that officers may feel threatened in those environments.

For several years, we have been using a first-person-shooter task (FPST) to investigate decisions to "shoot" (Correll, Park, Judd & Wittenbrink, 2002, 2007; Correll, Urland & Ito, 2006). This task presents images of young men, some armed, some unarmed, set against backgrounds like city streets or parks. The participant's goal is to shoot armed targets by pressing a button labeled *shoot*, but not to shoot unarmed targets (a decision indicated by pressing a second button labeled *don't shoot*). Half of the targets (armed and unarmed) are Black, and half are White.

Our work has largely focused on the influence of a target's race, and it has provided robust evidence of bias, both in the decisions participants make and in the speed with which they make them. First, participants typically shoot Black targets, armed and unarmed, more frequently than White targets. Second, even when participants respond correctly, they typically shoot an armed target more quickly when he is Black, but in response to an unarmed target, they indicate *don't shoot* more quickly when he is White. Other researchers have obtained similar results using a variety of paradigms (e.g., Greenwald, Oakes & Hoffman, 2003; Payne, 2001; Plant, Peruche & Butz, 2005).

Considering the processes that might give rise to bias in the FPST, we propose a simple model of associations (see Fig. 1). We argue that participants quickly categorize targets according to race (Brewer, 1989; Ito & Urland, 2003) and activate related stereotypes, including the concept of danger (Devine & Elliot, 1995). In a task like ours, which involves detection of hostile targets, rapid racial categorization of Black targets may increase participants' tendency to perceive them as threats, resulting in a predisposition to shoot.
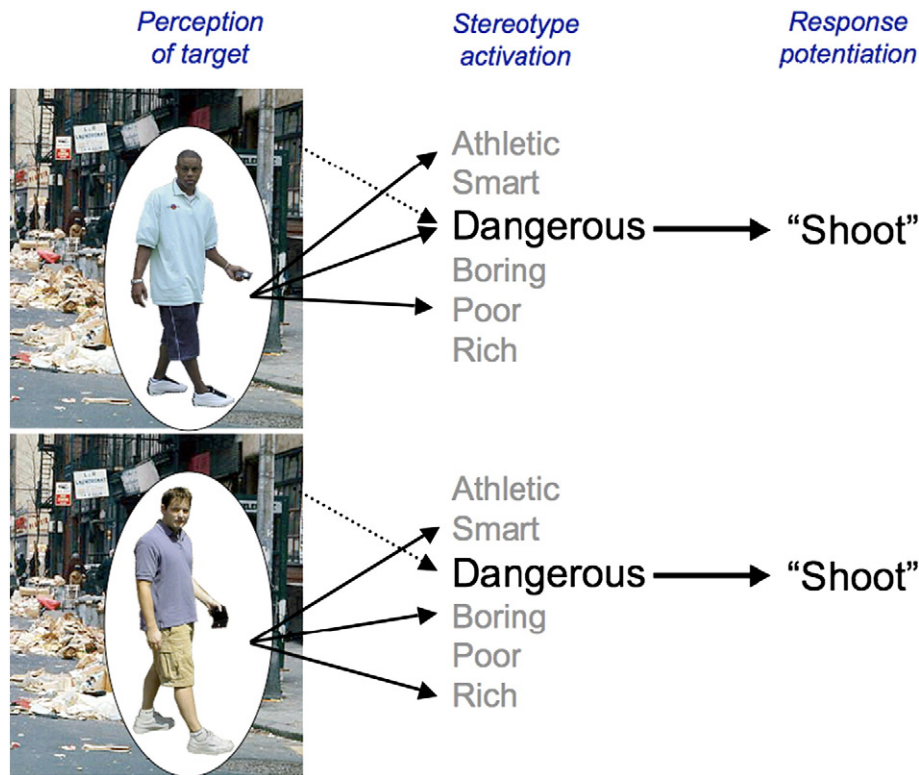
Fig. 1. A hypothetical model of processing operations, adapted from Correll et al., 2002. Targets are categorized by race, prompting differential stereotypic activation (solid lines). As a common component of racial stereotypes, activation of danger creates a predisposition to shoot that disproportionately affects Black targets. If danger is signaled by context (dashed lines), independently of race, it may create a predisposition to shoot White targets. The circles around each target have been added for the sake of exposition in this paper. In the task, targets are not isolated from the context.

From this model, we derived two key predictions. First, as the association between race and danger becomes more pronounced, bias should increase (see Correll et al., 2006, 2007, for support). A second prediction – the focus of the current research – suggests that any cue (i.e., not just target race) that sufficiently activates the concept of danger should create a predisposition to shoot. That is, if Black targets evoke a tendency to shoot *because* they seem threatening, then other (non-racial) cues that enhance perceptions of threat should similarly increase the tendency to shoot. For example, a dark alley in a dangerous part of town may create a context in which any unfamiliar person seems hostile.

Cacioppo and Berntson (1994; see also Kahneman & Tversky, 1984) contend that human beings (indeed, all organisms) must distinguish between opportunities and threats in their environment. To do so, they often integrate complex information, both positive and negative. But responses to multifaceted stimuli are often constrained to a unidimensional response. In the FPST, participants judge complex images — diverse people holding different objects situated in visually complicated environments. But participants have only two response options: *shoot* and *don't shoot*. Cacioppo's model involves two critical elements. First, it weights negative information more heavily than positive information. This *negativity bias* suggests that, given equal intensity, negative cues prompt withdrawal more than positive cues prompt approach. Second, the relationship between the intensity of stimulus information and behavior is nonlinear. Initial negative information has a powerful impact on behavior, motivating a defensive posture. But the incremental effect of each additional piece of negative information becomes less potent. In terms of our task, this model suggests that any highly salient threat cue may prompt a defensive orientation (i.e., a tendency to shoot), but additional threat cues may have little effect. In the absence of other cues, Black targets may signal greater danger than Whites, prompting the perceiver to shoot. But if contextual threats activate the concept of danger more broadly, all targets (Black and White) may be perceived in the light of this preexisting threat. Participants may therefore respond in a hostile fashion. Once participants adopt a defensive orientation due to an environmental threat cue, additional cues based on the target's race may have minimal incremental effects. Accordingly – if contextual threat cues cause White targets to receive the hostile treatment that Black targets *always* receive – racial bias should be attenuated in a hostile environment. This hypothesis (our primary prediction, H_1), derived from a psychological and neuroscientific perspective, perfectly matches the prediction of sociology's ecological contamination hypothesis: a dangerous context will increase the tendency to shoot White targets.

H_1 can be contrasted with an alternative hypothesis based on priming research. In a typical evaluative priming task, participants might view a prime stimulus (positive or negative) and then a target stimulus (positive or negative). In general, evaluatively consistent prime-target pairs facilitate responses: participants respond more quickly when the prime and the target are both positive (or both negative) and more slowly when one stimulus is positive and the other negative. In an interesting extension of this work, researchers included a third stimulus — a positive or negative stimulus that appeared before the prime. The results showed that the magnitude of the priming effect was enhanced when the pre-prime stimulus was evaluatively *inconsistent* with the prime. Essentially, participants showed a contrast effect as a function of the relationship between the pre-prime and the prime (Deutsch & Gawronski, 2009; Gawronski, Deutsch & Seidel, 2005). For example, a pleasant photograph of a kitten may generally facilitate classification of a positive target word like *gift*, and delay classification of negative words like *war*; but this pattern is exacerbated when an unpleasant image of a cockroach precedes the kitten. The idea is that the pre-prime image creates a psychological context against which the prime is contrasted. Translating this finding to the FPST, H_2 predicts that, in a threatening

context, non-threatening White targets will produce a contrast effect. In a dangerous environment, Whites may seem *even less dangerous* than they normally do. If so, participants should demonstrate a more pronounced reluctance to shoot White targets when they appear in a dangerous environment.

We tested these hypotheses by systematically manipulating the backgrounds in which FPST targets appeared. In our standard paradigm, targets are embedded in relatively innocuous backgrounds (e.g., parks, apartment buildings, train stations). For the present research, we assembled additional backgrounds featuring dumpsters and subway terminals covered with graffiti, dilapidated buildings, and inner-city streets (images evocative of poverty and crime). Targets from the original task were embedded in these threatening scenes. We then asked participants to perform the FPST, once with the original backgrounds and once with the dangerous backgrounds. In the neutral backgrounds, we predicted the standard pattern of bias (a tendency to shoot Black targets more frequently than Whites). However, the dangerous contexts raise two competing hypotheses. H_1 suggests that a dangerous environment will prime the concept of threat in general, evoking a tendency to shoot *both* White and Black targets. H_2 suggests that the contrast between a dangerous context and relatively non-threatening White targets may prompt participants to see Whites as even less dangerous, reducing (rather than increasing) the tendency to shoot them.

Pilot studies

We developed two versions of the FPST. The first, identical to the task used in previous research (Correll et al., 2002), presented targets in neutral backgrounds. The second differed only in that it employed threatening backgrounds. We conducted a preliminary study to equate the background images for visual complexity. We did not want performance to reflect systematic differences in participants' ability to discriminate objects in the neutral vs. dangerous contexts. We selected a set of 40 (20 neutral and 20 dangerous) backgrounds (640 pixels × 480 pixels). Using Adobe Photoshop, we created a set of 160 images by superimposing 4 non-threatening objects (two cell phones, a coke can, and a wallet), one at a time, in each background. Objects were randomly assigned to 1 of 9 possible positions defined by a 3 × 3 grid covering the central portion of the image (384 pixels × 288 pixels). The grid was used only to guide the creation of these stimuli, it was not visible to participants.

Thirty-four students participated. Each of the 40 backgrounds was presented 4 times with an object (once with each of the 4 objects) and 4 times with no object (320 trials). For each background, participants indicated whether or not an object was present as quickly as possible. We examined average reaction times (across participants and objects) for trials on which participants correctly detected the object. We ultimately selected 18 neutral backgrounds and 18 dangerous backgrounds, which did not differ in terms of latency (M's = 818.42 ms, 851.44 ms, respectively), $t(34) = 0.95$, $p = 0.35$. Latency to detect an object was thus roughly equivalent across contexts.

We also tested the 36 images to ensure that the ostensibly dangerous backgrounds were, in fact, perceived as more threatening than the ostensibly neutral scenes. 62 undergraduates rated each image on a scale from 1 (*extremely safe*) to 5 (*neither safe nor dangerous*) to 9 (*extremely dangerous*). A paired-samples test demonstrated that, as intended, the dangerous backgrounds ($M = 6.25$) were perceived as more threatening than the neutral backgrounds ($M = 3.30$), $t(61) = 29.44$, $p < 0.001$.

Method

Participants and design

Fifty-five non-Black undergraduates (mean age = 18.87; 33 female, 21 male, 1 missing; 46 White, 3 Latina/o, 2 Asian, 4 missing/

Table 1

n = 55	Neutral context				Dangerous context			
	White targets		Black targets		White targets		Black targets	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Unarmed error rate	0.134	0.118	0.160	0.149	0.149	0.116	0.147	0.117
Armed error rate	0.140	0.100	0.096	0.074	0.112	0.091	0.101	0.086
Difference	−0.005	0.127	0.063	0.137	0.037	0.131	0.046	0.131
Unarmed log latency	6.253	0.049	6.269	0.044	6.250	0.039	6.254	0.037
Armed log latency	6.134	0.065	6.131	0.073	6.139	0.068	6.146	0.060
Difference	0.118	0.060	0.138	0.067	0.112	0.060	0.107	0.063
Criterion (c)	0.037	0.278	−0.109	0.289	−0.063	0.307	−0.095	0.282
Sensitivity (d′)	2.419	0.789	2.529	0.831	2.486	0.743	2.533	0.739

other) participated in partial fulfillment of a course requirement.[1] Participants performed the task twice. The study involved a 2 (Task Order: dangerous context first vs. neutral context first) × 2 (Context: dangerous vs. neutral) × 2 (Target Race: Black vs. White) × 2 (Object Type: gun vs. no gun) mixed-model design, in which Order varied between participants and all other factors varied within participants.

Materials

Neutral-context task

The neutral task was essentially identical to the task used by Correll et al. (2002). Ten Black and ten White men were photographed holding guns and non-gun objects (e.g., a wallet or a cell phone). For each individual, we selected four images, two with guns and two with innocuous objects, resulting in 80 distinct images (20 of each type: armed White, armed Black, unarmed White, and unarmed Black). Using Photoshop, each target was embedded in a neutral background scene (in both the neutral and dangerous stimuli, two of the 18 backgrounds identified in the pilot test were randomly selected and used twice, yielding a set of 20). Targets were randomly assigned to backgrounds, with the restriction that each target type should be represented with equal frequency in each background.

The task was developed in PsyScope (Cohen, MacWhinney, Flatt & Provost, 1993), following a 2 × 2 within-subjects design, with Target Race (Black vs. White) and Object Type (gun vs. non-gun) as repeated factors. On any given trial, a series of scenes appeared, followed by a target image. Players were instructed to respond as quickly as possible whenever a target appeared, pressing a button labeled *shoot* if the target was armed, and pressing a button labeled *don't shoot* if he was unarmed. The game awarded points based on performance, and participants were required to respond to a target within 630 ms.[2] Point totals and feedback, both visual and auditory, were presented at the conclusion of every trial. The game consisted of a 16-trial practice block and an 80-trial test block (see Correll et al., 2002, for details).

Dangerous-context videogame task

Targets from the neutral-context game were embedded in the 18 backgrounds chosen in the pilot study (two of which were used twice). These images included dilapidated buildings, dumpsters, subway terminals with graffiti, etc. With the exception of these backgrounds, the task was identical to the neutral-context version.

---

[1] Two additional participants were excluded due to experimenter error, and two were excluded as outliers (Cook's D's for the three-way interaction reported below = .233 and .108; next highest value = .075; see Judd & McClelland, 1989).

[2] In the present research, our goal was to examine the effect of context on errors. Accordingly, we used a 630-ms window, which reduces variability in reaction times, but increases errors.
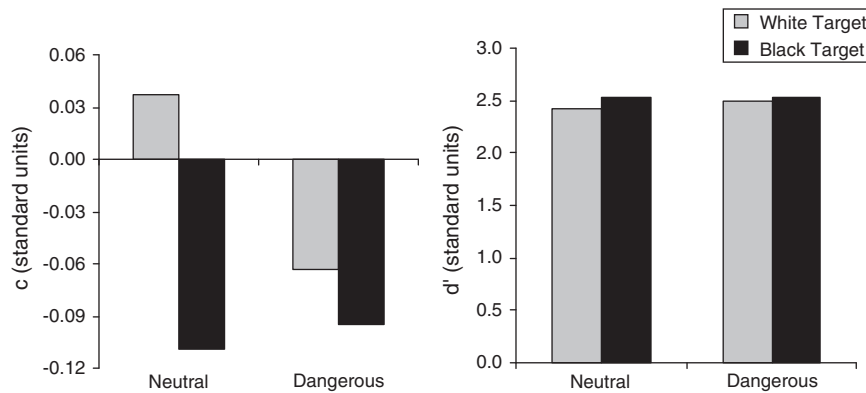
Fig. 2. Signal detection analysis. Participants set a lenient criterion (c, left panel) for Black targets in both neutral and dangerous contexts. For White targets, the criterion was low only in a dangerous context. Participants' sensitivity (d′, right panel) did not differ as a function of race or context.

Procedure

Participants, in groups of 1 to 6, were seated in cubicles equipped with iMacs. The experimenter introduced the study as a test of vigilance. Participants were told that they would perform the task twice. Between the two games, they were given a short break. Upon completion of the second game, participants were thanked and debriefed.

Results

On average, participants responded incorrectly on 11.52% of trials and timed out on 10.46% of trials. Error rates for each target type, in each game, were submitted to a 2 (Task Order: danger first vs. neutral first) $\times 2$ (Context: dangerous vs. neutral) $\times 2$ (Target Race: Black vs. White) $\times 2$ (Object Type: gun vs. no gun) mixed-model ANOVA. There was a main effect of object, $F(1,53) = 4.70$, $p < 0.035$, such that participants were more likely to incorrectly shoot an unarmed target (false alarm) than to incorrectly choose not to shoot in response to an armed target (miss). The Race $\times$ Object interaction (i.e., bias) was also significant, $F(1,53) = 8.37$, $p = 0.006$, but its magnitude depended on Context, $F(1,53) = 6.05$, $p < 0.018$, for the 3-way Race $\times$ Object $\times$ Context interaction.[3] In the neutral context, bias was significant, $F(1,53) = 14.08$, $p = 0.0004$, such that participants shot an unarmed Black target more frequently than an unarmed White, $F(1,53) = 3.56$, $p < 0.065$, and failed to shoot an armed White target more frequently than an armed Black, $F(1,53) = 10.64$, $p < 0.002$. In the dangerous context, bias was not significant, $F(1,53) = 0.27$, $p < 0.61$. No other effects approached significance, F's < 1.17, p's > 0.28, and, in particular, there were no effects of Order, F's < 0.95, p's > 0.3 (Table 1).

Although the focus of this research involved errors, we tested the effects of context on latency. Average log-transformed latencies from correct trials were submitted to an ANOVA identical to the error analysis. Parallel results emerged. The 3-way Race $\times$ Object $\times$ Context interaction suggested that context moderated bias, $F(1,53) = 4.04$, $p < 0.05$. Bias (the Race $\times$ Object interaction) was significant in neutral contexts, $F(1,53) = 5.47$, $p < 0.024$, but not dangerous contexts, $F(1,53) = 0.27$, $p = 0.61$. (This pattern is inconsistent with speed-accuracy tradeoff.)

Signal detection analysis

Signal Detection Theory (SDT; Green & Swets, 1966/1974; MacMillan & Creelman, 1991) offers a way to quantify and

conceptualize the effects of race and context in this task by disentangling two distinct factors that influence error rates. First, it provides a measure of participants' ability to differentiate armed targets from unarmed targets. Sensitivity was gauged with the statistic, d′. Second, SDT estimates the degree to which participants favored a shoot response over a don't-shoot response — did they set a very low bar for the decision to shoot (shooting often) or a very high bar (shooting rarely)? This criterion was assessed with the statistic, c. Our previous work suggests that, in neutral contexts, participants set a more lenient criterion when targets are Black, rather than White (i.e., c is lower for Blacks), but that sensitivity (d′) does not differ as a function of race.

We calculated c four times for each participant: White targets in neutral contexts, Blacks in neutral contexts, Whites in dangerous contexts, and Blacks in dangerous contexts (see Fig. 2). We then analyzed the criteria as a function of Order, Context and Target Race, submitting them to a 3-way mixed-model ANOVA. We obtained no effects involving Order, F's < 0.14, p's > 0.7. We did, however, obtain an effect of Target Race, $F(1,53) = 8.51$ $p < 0.006$, and an interaction between Context and Race, $F(1,53) = 4.14$ $p < 0.047$.[4] The criterion to shoot was lower for Black targets than for White targets in the neutral contexts, $F(1,53) = 11.51$, $p < 0.002$. In the dangerous context, however, the criteria did not differ for Black vs. White targets, $F(1,53) = 0.63$, $p < 0.44$. H₁ suggests, specifically, that the presence of contextual danger cues should induce participants to set a lower criterion for White targets when they appear in threatening environments. Indeed, estimates of c were lower when Whites appeared in dangerous contexts rather than neutral contexts, $F(1,53) = 4.98$, $p < 0.03$. But context had no effect on the criteria for Black targets, $F(1,53) = 0.11$, $p < 0.75$. In essence, whenever danger was salient – either due to race or due to context – participants adopted a more lenient, more trigger-happy criterion for the decision to shoot.

The criterion measure is valuable, not just because it allows us to test the effects of condition, but because the statistic, c, has meaning in an absolute sense. Values close to zero indicate that participants employed the two response options (*shoot* and *don't shoot*) with roughly equal frequency. Values below zero indicate that participants favored the *shoot* response over the *don't-shoot* response. Values above zero indicate the opposite. By comparing c to zero, we can assess the extent to which participants were predisposed to shoot. Mean values of c fell significantly below zero for Blacks in both contexts (neutral $t(54) = -2.80$, $p < 0.007$; dangerous $t(54) = -2.50$, $p < 0.016$). They also tended in a negative direction for Whites in the

---

[3] Including the outliers does not alter the effect ($p < .022$).

[4] Including the outliers slightly weakens the effect ($p < .073$).

dangerous context ($t(54) = -1.53$, $p < 0.133$), suggesting that participants favored the *shoot* response for these targets. Only when no danger cues were present (i.e., for White targets in a neutral context, $t(54) = 0.98$, $p < 0.34$) did the criterion rise, indicating no predisposition to shoot.

Estimates of sensitivity (participants' ability to distinguish between armed and unarmed targets, calculated as d') were submitted to an identical mixed-model ANOVA. There were no significant effects of either race, context or their interaction, F's = 0.25, 1.42, and 0.22, respectively, p's < 0.62, 0.24, and 0.65, respectively), suggesting that the independent variables did not affect participants' ability to distinguish between armed and unarmed targets. This result suggests that trials involving the dangerous contexts were not inherently more or less difficult than those involving the neutral contexts. Thus, participants set different criteria for the decision to shoot as a function of context even though they were equally able, in all conditions, to discriminate between armed and unarmed targets.

Discussion

The present research tested a single, very simple proposition: perception of threat fosters a predisposition to shoot. We have suggested that racial bias in decisions to shoot reflects the fact that most Americans associate Blacks (or, at least, young Black men) with danger (Correll et al., 2002; Devine, 1989; Devine & Elliot, 1995). If threat perception really mediates shooter bias, then *any* salient danger cue (not just race) should evoke a tendency to shoot. The present study manipulated racial and non-racial danger cues orthogonally in an effort to show that danger, not race, is the proximal cause of the predisposition to shoot. Consistent with our primary hypothesis, H$_1$, dangerous contexts reduced bias, and this reduction reflected an increased willingness to shoot (otherwise nonthreatening) White targets when they appeared in dangerous contexts. These findings clearly contradict H$_2$, which suggested that White targets will be viewed *in contrast to* a threatening environment, and that they will therefore seem less threatening in a ghetto than in an upscale neighborhood.

Though it is true that a dangerous context eliminated bias in the decision to shoot, the results by no means suggest that Blacks were treated with greater restraint in those environments. Rather, our data indicate that Whites were treated in a more hostile fashion. The reduction in bias was due to the fact that participants were predisposed to shoot *everybody* – Whites and Blacks, alike – when the context suggested danger. These data support a sufficiency model of threat detection, suggesting that participants adopt a defensive orientation in the presence of a sufficiently salient danger cue. Once danger is activated by that initial cue, additional cues have little, if any, incremental impact.[5]

Several recent studies have explored the impact of race in the context of threat-related cues. On the surface, this work seems inconsistent with the sufficiency model presented here. However, we believe a careful reading of prior work suggests almost no conflict with (and even offers a hint of support for) the present findings. For instance, Maner et al. (2005) suggested that when participants are in a self-protective frame of mind, they rate outgroup targets as more angry. As such, their studies examine how internal emotional or motivational states can bias perception (an effect with a well-established and interesting history Bruner, 1957). The focus of the current studies, however, is not this kind of incidental affect, but rather the integration of exogenous threat cues in a stimulus array. Thus the present results involve a fundamentally different question. In other work, Trawalter and her colleagues found that White partici-

pants preferentially attended to Black (rather than White) faces, and that this bias was enhanced for faces displaying direct (rather than averted) eye-gaze (Trawalter, Todd, Baird & Richeson, 2008). To the extent that eye-gaze represents a threat cue, these results suggest that a secondary cue exacerbates racial discrepancies, in clear conflict with the data reported here. However, Trawalter and her colleagues explicitly reject the idea that eye-gaze is a threat signal. Rather, they argue that eye-gaze is a signal of *approach motivation* (p. 1325), which can signify a host of intentions, ranging from very positive to very negative, depending on who is doing the gazing. An ingroup member's gaze may have very different connotations than direct gaze from a member of an outgroup. For White participants, the direct gaze of a White face may not constitute a threat cue at all. Schaller, Park and Mueller (2003), however, directly manipulate a secondary threat cue, so their data are the most directly relevant to the question at hand. In one study, they use the Implicit Association Test (Greenwald, McGhee & Schwartz, 1998) to examine the effect of ambient darkness on the association between Blacks and danger. The paper ultimately focuses on an interaction between responses to the Belief in a Dangerous World (BDW) scale and darkness, such that the correspondence between BDW and stereotyping is stronger in a dark room than in a light room.[6] But the authors also admirably report all of the lower-order effects in their design (p. 645), allowing us to examine the question that most directly bears on the present research: on average, does darkness exacerbate or attenuate the Black-danger stereotype? Though the effect is not statistically significant, the data suggest that, if anything, a darkened room tends to *reduce* racial bias in danger-based stereotypes rather than exacerbate it (p < 0.17). Though the IAT does not allow us to determine whether the reduction in bias is due to an increase in White-danger associations, it is important to note that this trend is perfectly consistent with the results reported here. Based on our reading, then, prior work generally seems to address different questions than the current study. Moreover, in the one paper with a similar focus, the results offer tentative support for our conclusions.

It may also be useful to speculate on the factors that differentiate the present findings, reflecting a reduction in bias in the more threatening context, from research showing contrast effects or exacerbated bias in negative contexts (e.g., Deutsch & Gawronski, 2009; Gawronski et al., 2005; Wittenbrink, Park & Judd, 2001). We would like to highlight two issues. One potentially important issue is the nature of the stimulus presentation. In many studies that report contrast effects, researchers have presented primes and targets in a sequential fashion, whereas in the current work, we presented the target and the contextual information simultaneously. In fact, previous work has shown that mode of presentation can moderate the nature of priming effects, with contrast emerging for sequential primes and assimilation for simultaneous primes (Wedell, Parducci & Geiselman, 1987).[7] This seems like a fruitful path for future research. A second difference between earlier studies and the current research is the type of judgment performed by participants. Contrast effects have all been observed in studies assessing evaluation ("is the stimulus good or bad?"). The current work focuses on threat detection ("is there a threat in the environment?"). It has been suggested that contrast effects in rapid response paradigms reflect differences in the salience of task-relevant stimulus features (Deutsch & Gawronski, 2009). Negatively valenced context information is thought to increase the likelihood that negative stereotypes or evaluations associated with the target impact participants' judgments. One reason the current results show a distinctly different pattern may be that the context information in our studies has quite different implications for the judgment task. The context, in and of itself, may be seen as diagnostic

---

[5] It may be important that the current study employed clear and visually salient danger cues (race and environment). If less salient cues are *insufficient* (individually) to prompt a defensive posture, a sufficiency model would predict more additive effects.

[6] BDW is a measure of generalized fear, but it has also been associated with racial prejudice (Duckitt, Wagner, du Plessis & Birum, 2002) so this interaction may, in part, reflect different attitudes toward Black and White people.

[7] We wish to thank the editor for this citation.

for the judgment at hand, namely to detect threat. This reasoning of course is entirely consistent with our proposal that threat detection relies on a sufficiency model; once a sufficient threat level has been reached, participants adopt a defensive posture (and additional cues have little incremental impact). As a result, the addition of threatening cues to the situation leads to an overall reduction in bias.

It is also important to recognize that the current study focuses explicitly on the effects of additional threat cues whereas the evaluative work examines both negative and positive contexts. This constraint raises a question about the effect of contextual cues that might signal safety. Positive situational cues may mitigate bias by impacting responses to Black targets in a positive fashion, rather than making responses to Whites more negative.

As discussed in the Introduction, sociologists and criminologists have a longstanding interest in race and its effects on police behavior. Based on a range of data sources (including the Federal Bureau of Investigation, citizen complaints, local police agency records and social observation), research suggests that police use of force, including the use of firearms, is applied more frequently and more severely to suspects who are young, male, and either Black or Latino (Geller, 1982; DOJ, 2001). Researchers like Terrill and Reisig (2003) have begun to explore how police shootings relate to the neighborhood in which an encounter occurs, and their initial results perfectly match the data reported here: a potentially dangerous and disadvantaged neighborhood may prompt more extreme use of force regardless of the suspect's race.

Clearly, racial cues can and do signal threat. Young Black male targets prompt a defensive orientation. But racial threat perception may be one manifestation of a more comprehensive threat-detection process — a process that monitors the environment for a variety of threats. In this more comprehensive framework, race represents one of many possible signals, each of which can prompt defensive action. Several scholars suggest that humans have evolved mechanisms for detecting threat, which – like the speeding bus, charging bear or snake in the grass – must be processed quickly and efficiently (Öhman, Lundqvist & Esteves, 2001; Öhman, Flykt & Esteves, 2001). It may be reasonable to suggest that this mechanism responds to context as well as race. The current research only begins to touch on this question. All the same, the decelerating impact of multiple danger cues, documented here, highlights the complexity of the threat-detection process and hints at a fast-acting system capable of processing a complicated, multifaceted environment and triggering defensive behavior whenever the situation is deemed "dangerous enough." That the presence of a Black man is, in some circumstances, sufficient to stimulate such a response remains a source of great concern, but the dynamic nature of this system holds promise for the study, understanding, and (perhaps eventually) control of these biases.

Acknowledgments

References

Brewer, M. B. (1989). A dual process model of impression formation. In T. K. Srull, & R. S. Wyer (Eds.), Advances in social cognition, Vol. 1, Hillsdale, NJ: Erlbaum.

Bruner, J. S. (1957). On perceptual readiness. Psychological Review, 64, 123−152.

Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. Psychological Bulletin, 115, 401−423.

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.

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 & Social Psychology, 83, 1314−1329.

Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2007a). The influence of stereotypes on decisions to shoot. European Journal of Social Psychology, 37, 1102−1117.

Correll, J., Park, B., Judd, C. M., Wittenbrink, B., Sadler, M. S., & Keesee, T. (2007b). Across the thin blue line: Police officers and racial bias in the decision to shoot. Journal of Personality & Social Psychology, 92, 1006−1023.

Correll, J., Urland, G. L., & Ito, T. A. (2006). Event-related potentials and the decision to shoot: The role of threat perception and cognitive control. Journal of Experimental Social Psychology, 42, 120−128.

Department of Justice (2001). Policing and homicide, 1976–98: Justifiable homicide by police, police officers murdered by felons. Bureau of Justice Statistics, NCJ 180987. .

Deutsch, R., & Gawronski, B. (2009). When the method makes a difference: Antagonistic effects on "automatic evaluations" as a function of task characteristics of the measure. Journal of Experimental Social Psychology, 45(1), 101−114.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. Journal of Personality & 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.

Duckitt, J., Wagner, C., du Plessis, I., & Birum, I. (2002). The psychological bases of ideology and prejudice: Testing a dual process model. Journal of Personality and Social Psychology, 83(1), 75−93.

Gawronski, B., Deutsch, R., & Seidel, O. (2005). Contextual influences on implicit evaluation: A test of additive versus contrastive effects of evaluative context stimuli in affective priming. Personality and Social Psychology Bulletin, 31(9), 1226−1236.

Geller, W. A. (1982). Deadly force: What we know. Journal of Police Science and Administration, 10, 151−177.

Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics.New York: Wiley Reprinted 1974 by Krieger, Huntington, NY.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. Journal of Personality and Social Psychology, 74, 1464−1480.

Greenwald, A. G., Oakes, M. A., & Hoffman, H. G. (2003). Targets of discrimination: Effects of race on responses to weapons holders. Journal of Experimental Psychology, 39, 399−405.

Ito, T. A., & Urland, G. R. (2003). Race and gender on the brain: Electrocortical measures of attention to race and gender of multiply categorizable individuals. Journal of Personality and Social Psychology, 85, 616−626.

Judd, C. M., & McClelland, G. H. (1989). DATA ANALYSIS: A model comparison approach. San Diego: Harcourt, Brace, & Jovanovich.

Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. American Psychologist, 39, 341−350.

Larson, J. (2004, April 10). Behind the death of Timothy Thomas: Shooting of 19-year-old brings to light pattern of ticketing that raises questions of racial profiling. Dateline NBC. Retrieved May 12, 2004, from http://msnbc.msn.com/id/4703574/.

MacMillan, N. A., & Creelman, C. D. (1991). Detection theory: A users guide. Cambridge: Cambridge University Press.

Maner, J. K., Kenrick, D. T., Becker, D. V., Robertson, T. E., Hofer, B., Neuberg, S. L., et al. (2005). Functional projection: How fundamental social motives can bias interpersonal perception. Journal of Personality and Social Psychology, 88(1), 63−78.

Ohman, A., Flykt, A., & Esteves, F. (2001). Emotion drives attention: Detecting the snake in the grass. Journal of Experimental Psychology: General, 130(3), 466−478.

Öhman, A., Lundqvist, D., & Esteves, F. (2001). The face in the crowd revisited: A threat advantage with schematic stimuli. Journal of Personality & Social Psychology, 80, 381−396.

Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. Journal of Personality & Social Psychology, 81, 181−192.

Plant, E. A., Peruche, B. M., & Butz, D. A. (2005). Eliminating automatic racial bias: Making race non-diagnostic for responses to criminal suspects. Journal of Experimental Social Psychology, 41, 141−156.

Schaller, M., Park, J. H., & Mueller, A. (2003). Fear of the dark: Interactive effects of beliefs about danger and ambient darkness on ethnic stereotypes. Personality and Social Psychology Bulletin, 29(5), 637−649.

Smith, D. A. (1986). The neighborhood context of police behavior. In A. J. Reiss, & M. Tonry (Eds.), Communities & Crime (pp. 313−341). Chicago: University of Chicago Press.

Terrill, W., & Reisig, M. D. (2003). Neighborhood context and police use of force. Journal of Research in Crime & Delinquency, 40, 291−321.

Trawalter, S., Todd, A. R., Baird, A. A., & Richeson, J. A. (2008). Attending to threat: Race-based patterns of selective attention. Journal of Experimental Social Psychology, 44 (5), 1322−1327.

Wedell, D. H., Parducci, A., & Geiselman, R. E. (1987). A formal analysis of ratings of physical attractiveness: Successive contrast and simultaneous assimilation. Journal of Experimental Social Psychology, 23, 230−249.

Werthman, C., & Piliavin, I. (1967). Gang members and the police. In D. Bordua (Ed.), The police: Six sociological essays. New York: Wiley.

Wittenbrink, B., Park, B., & Judd, C. M. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. Journal of Personality & Social Psychology, 81, 815−827.