

The Police Officer's Dilemma: A Decade of Research on Racial Bias in the Decision to Shoot

Joshua Correll^{1*}, Sean M. Hudson¹, Steffanie Guillermo¹
and Debbie S. Ma²

¹University of Colorado, Boulder

²California State University, Northridge

We review sociological, correlational, and experimental research that examines the effect of a target's race on the decision to shoot. Much of this work involves computer-based simulations of a police encounter, in which a participant must decide whether or not to shoot a potentially hostile target who is either Black or White. Experimental work with undergraduate participants reveals a clear pattern of bias (a tendency to shoot Black targets but not Whites), which is associated with stereotypes linking Blacks with the concept of danger. Subsequent work with police officers presents a more complex pattern. Although police are affected by target race in some respects, they generally do *not* show a biased pattern of shooting. We suggest that police performance depends on the exercise of cognitive control, which allows officers to overcome the influence of stereotypes, and we conclude with potential implications of this research for law enforcement.

Does Suspect Race Impact Police Shooting?

In the early hours of 4 February 1999, in a notorious neighborhood in the South Bronx, four White plainclothes police officers spotted Amadou Diallo, a 22-year-old Black man, in front of his apartment. The officers approached, Diallo turned to enter his building, and the officers pursued. In the vestibule, they ordered Diallo to show his hands and freeze (Flynn, 1999; Fritsch, 2000). Diallo reportedly reached into his pocket. One of the officers identified the object Diallo pulled out as a gun and yelled, "Gun! He's got a gun!" (Fritsch, 2000). The officers opened fire, ultimately discharging 41 rounds and fatally wounding Diallo. A wallet – not a gun – was found at the scene.

In 1994, 13-year-old Nicholas Heyward, Jr., was shot and killed by an officer while playing with a toy gun in the stairwell of his apartment building. In 2003, Orlando Barlow was shot while surrendering on his knees. The officer stated that he feared Barlow was feigning surrender and was going to reach for a gun. In 2010, Aaron Campbell was walking *backward toward police with his hands behind his head* (presumably to surrender) and was shot by an officer who believed he was going to reach for a gun. Although each of these cases involves different circumstances, in each case, the suspect was a Black man, and in each case, officers stated that they believed their lives were in immediate danger because the suspect was reaching for, or had, a gun.

Are police more likely to shoot Black suspects? Police officers face situations in which (a) behaviors and objects are unclear (e.g., it may be dark, the object in question may be hidden from plain sight), and (b) they must make sense of that ambiguity quickly in order to protect themselves and those around them. Because of this lack of clarity and time pressure, they may rely on salient secondary cues (e.g., the suspect's race or gender, the context provided by the neighborhood) to interpret the situation.

Abundant evidence supports the argument that stereotypes – including racial stereotypes – guide expectations and interpretations of ambiguous stimuli. For example, people are typically

faster when responding to stereotype-congruent information rather than stereotype-incongruent information, and they interpret ambiguous behavior as more violent when the actor is Black rather than White (Duncan, 1976; Dovidio, Evans, & Tyler, 1986; Fazio, et al., 1995; Gaertner & McLaughlin, 1983; Greenwald, McGhee, & Schwartz, 1998; Kunda & Sherman-Williams, 1993; Sagar & Schofield, 1980; Wittenbrink, Judd, & Park, 1997). It is important to note that one of the most prevalent stereotypes about Blacks, particularly young, Black men, is that they are dangerous (Devine, 1989; Wittenbrink, Judd, & Park, 1997).

Sociological Evidence Concerning Officer-involved Shootings

Data from the Department of Justice (DOJ, 2001) indicate that, per capita, police are roughly five times more likely to shoot a Black person than a White person. That is, knowing nothing but a person's race, the likelihood of being killed by police quintuples if the person is Black rather than White. The specific cases described above are clearly part of a broader trend. Indeed, sociological research consistently shows a discrepancy such that officers use greater force (both lethal and non-lethal) when the suspect is Black rather than White (Inn, Wheeler, & Sparling, 1977; Jacobs & O'Brien, 1998; Smith, 2004; Sorenson, Marquart, & Brock, 1993; Terrill & Mastrofski, 2002; Terrill & Reisig, 2003). For example, Jacobs and O'Brien (1998) analyzed archival data from 170 US cities between 1980 and 1986 and found evidence that police killed Black suspects at a significantly higher rate in cities with a greater proportion of Black residents, even after statistically controlling for violent crime rates (see also Smith, 2004; Sorensen et al., 1993). This evidence seems to suggest that law enforcement may be biased when employing lethal force against Black suspects.

At the same time, others have argued that the discrepancy in use of force is not driven by race but rather by the types of activities in which Blacks and Whites engage. Specifically, if Blacks conduct more criminal activity than Whites, it stands to reason that Black suspects will be overrepresented in police encounters. Inn, Wheeler, and Sparling (1977) showed that statistically controlling for racial differences in arrest rates (taken as a proxy for criminal behavior) eliminated the effect of race on shootings. Terrill and Reisig (2003) found that, although officers were more likely to use force against Blacks, this was largely due to the fact that officers were more likely to encounter Blacks in dangerous neighborhoods. Findings such as these seem to suggest that the discrepancy in use of force is warranted and can be explained in terms of differences between Whites and Blacks in criminal activity.

Observational data provide a rich, naturalistic glimpse at the association between race and policing. However, as these divergent interpretations suggest, it is difficult to demonstrate a causal relationship between suspect race and police use of force based solely on correlational data. Though laboratory studies (including those reviewed below) involve impoverished simulations of complex behavior, in order to establish causality, it can be useful to experimentally manipulate the variables of interest.

Examining Racial Bias in the Laboratory

Experimental paradigms and initial demonstrations of bias

Payne's (2001) research provided the first experimental test of whether race influences weapon identification. His computer task briefly presented either a Black or White male face as a prime, followed by a target object, which was either a gun or a tool (Figure 1a). Participants were asked to classify that object. In the most common version of this task, participants have very little time to respond (500 ms). Payne found that participants identified guns more quickly and more accurately after they had been primed with a Black face (rather

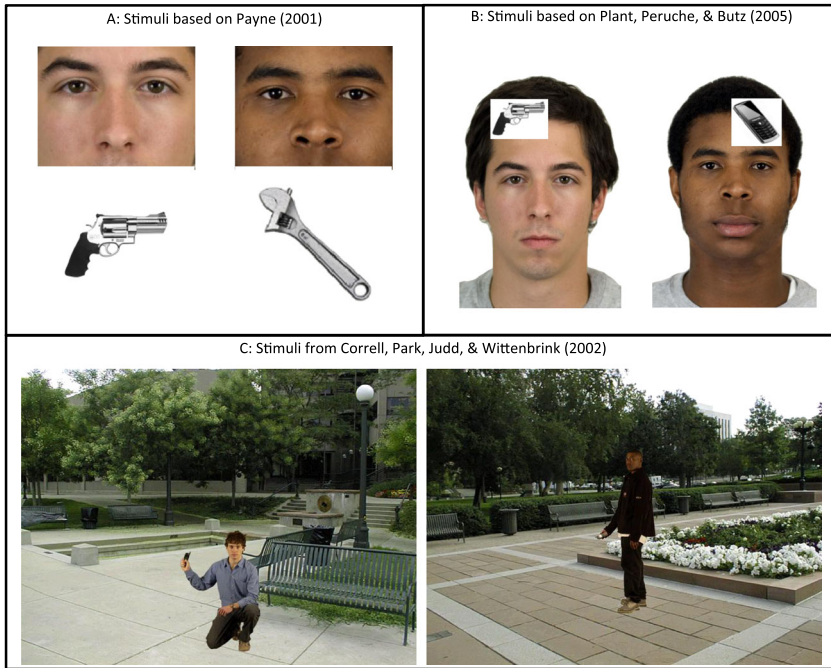


Figure 1 Example stimuli for three different tasks used to measure racial bias in weapon identification and decisions to shoot.

than a White face). At the same time, participants identified *tools* more quickly and more accurately after they had been primed with a White face (rather than a Black face). These findings suggest that the Black prime enhanced the tendency to respond as if a threat were present, implying an association between Blacks and threat.

We investigated a similar question, focusing on the impact of race on participants' decisions to shoot (Correll, Park, Judd, & Wittenbrink, 2002). Participants performed a first-person-shooter task (FPST) designed to simulate the experience of a police officer. In this task, on any given trial, a male target (either Black or White) appears against a realistic background (Figure 1c). Participants must decide whether or not to shoot, either pressing a key labeled "shoot" or a separate key labeled "don't shoot". Participants receive feedback after each trial and are rewarded with points for correct responses or penalized for errors or responses that are too slow. In one version of the task, participants must respond within 850 ms (allowing us to examine variation in response time). In another version, they have only 630 ms (which induces high error rates and allows us to examine variation in the nature of those mistakes).

Participants in the original studies showed bias in both response latencies and error rates. They shot an armed target more quickly if he was Black (rather than White) but decided not to shoot an unarmed target (pressing the "don't shoot" key) more quickly if he was White rather than Black. Regarding error rates, participants erroneously shot unarmed Black targets more frequently than unarmed Whites. Conversely, they mistakenly chose not to shoot armed White targets more often than armed Blacks (Figure 2).

Signal detection analysis provides a useful lens for interpreting these errors. Signal detection theory (SDT) assumes that armed and unarmed targets vary along some judgment-relevant dimension (e.g., perceived threat). While both types of targets may vary, armed targets are presumably more threatening, in general, than unarmed targets (Figure 2).

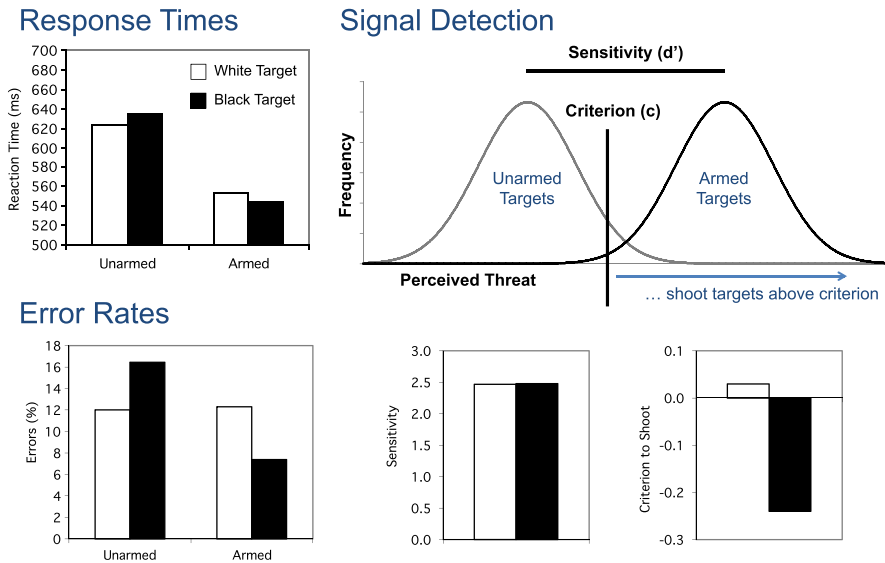


Figure 2 Typical patterns of bias in the FPST: response times, error rates, and signal detection sensitivities and criteria (derived from error rates).

SDT distinguishes between two factors that influence the participant's response: sensitivity and criterion. Sensitivity involves a participant's ability to differentiate armed targets from unarmed targets. It estimates the psychological distance between the armed and unarmed distributions on the perceived threat dimension. In our work, we generally find that race has little effect on sensitivity: regardless of target race, participants generally distinguish between armed and unarmed targets quite well. The SDT estimate of the criterion assesses a participant's predisposition to shoot in general. The criterion represents the point on the perceived threat dimension at which the participant decides to shoot. A lower, more lenient criterion suggests that participants will open fire frequently, whereas a higher, more stringent criterion suggests that participants will only fire at extremely threatening targets. Our work suggests that race affects the criterion, such that participants employ a more lenient criterion for Black targets (relative to White targets), indicating a greater willingness to open fire.

Exploring psychological process: threat-relevant stereotypes

Correll and colleagues (2002) argued that racial bias in the decision to shoot stems from stereotypic associations connecting Blacks to the concepts of danger and threat in our culture – a proposition supported by correlational, neural, and experimental evidence. For example, Correll, Urland, and Ito (2006) measured cultural stereotypes and recorded event-related brain potentials (ERPs) during the FPST. ERPs measure fluctuations in the brain's electrical activity. In the FPST, an ERP component associated with threat (P200) differentiated between White and Black targets. P200 responses were larger to Black targets than White targets, suggesting that Blacks were perceived as more threatening. Moreover, the magnitude of this differentiation depended on cultural stereotypes. Participants who expressed greater awareness of stereotypes linking Blacks with danger showed more pronounced threat-like responses to Black targets in the P200, which in turn led them to show more bias when making the decision to shoot. Correll, Park, Judd, and Wittenbrink (2007) experimentally

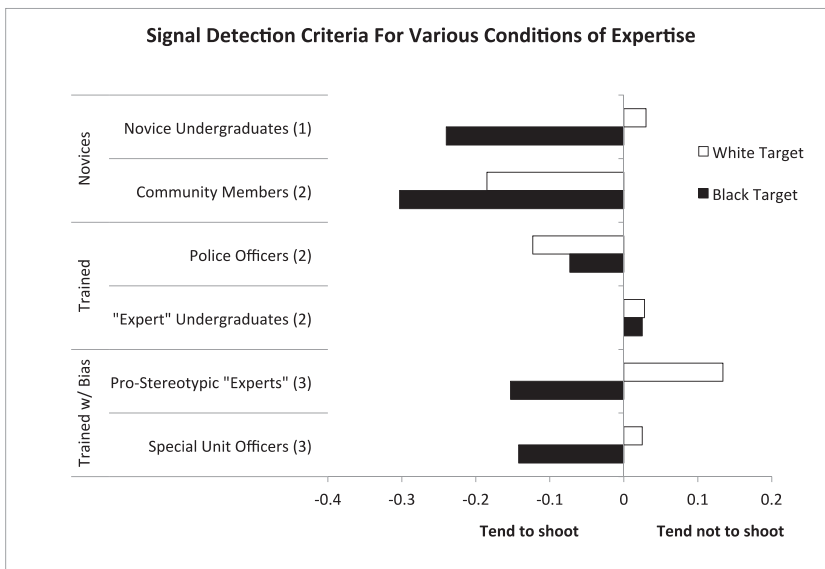


Figure 3 Signal detection criteria for several different conditions, including novices and experts. Note that expertise does not always reduce bias. 1, Correll et al., (2002). 2, Correll et al., 2007a. 3, Sim et al., 2013.

manipulated the association between race and threat. Manipulations that enhanced the association between Blacks and danger increased the magnitude of bias. Newspaper articles that highlighted Black (rather than White) criminals and stimulus sets in which Blacks (rather than Whites) were disproportionately likely to be armed increased the tendency to shoot Black targets. Bias favoring the decision to shoot has also been observed toward Latino targets (who are stereotyped as threatening), but not Asians (who are not typically associated with danger) (Sadler, Correll, Park, & Judd, 2012).

Research using the FPST strongly suggests that race affects the decision to shoot (in terms of both RTs and error rates) and that stereotypic associations between Blacks and danger drive this bias. However, the initial research employed college students as participants. A second generation of work on this topic examined whether police officers, trained individuals whose role requires them to make these decisions, exhibit similar patterns of bias.

Practice, Expertise, and Police Officer Performance

It was difficult to predict how police would perform on the FPST. Police are experts with firearms and with the decision to shoot, they receive regular training, and their day-to-day experience affords constant practice in assessing potential threats. To the extent that practice generally improves decision-making processes (MacLeod, 1998; MacLeod & Dunbar, 1988; Plant & Peruche, 2005), we reasoned that police might show reduced bias. However, there was also reason to suspect that officers would show *more pronounced* racial bias due to on-the-job socialization and frequent exposure to minorities and crime, alleged to promote negative views of Blacks (e.g., Geller, 1982; Smith, 2004; Teahan, 1975).

The research of Plant and colleagues

In 2005, Ashby Plant and her colleagues examined the effect of training on a task that was, in many respects, similar to our own paradigm (Plant & Peruche, 2005; Plant, Peruche, & Butz,

2005). However, rather than presenting participants with a complex image of a person situated in a realistic background and holding an object, Plant presented participants with a face (either Black or White) paired with an object (a gun or a neutral object) (Figure 1b). Because our research suggests that stimulus complexity affects performance on the FPST (Correll, Wittenbrink, Crawford, & Sadler, 2014), we differentiate this task from our own. When undergraduates performed Plant's task, they initially demonstrated bias similar to the studies described above. For example, they chose "shoot" more often when Black faces (rather than Whites) were paired with neutral objects. However, over the course of the task (which was fairly long: 160 trials), stereotypic errors diminished until – in the second half of the task – there was no evidence of racial bias. This suggests that brief experience with Plant's task can reduce bias. In an experimental test, participants practiced either the race-weapon task or (in a control condition) a task unrelated to race. Twenty-four hours later, all participants completed the critical race-weapon task. Participants in the untrained control condition showed significant bias, but participants who practiced the race-weapon task did not. Plant and Peruche (2005) conducted studies with police officers as participants and found similar results: initial bias that could quickly be trained away.

The research of Correll and colleagues

Our own work explored the effects of training and expertise using the more complex stimuli in Figure 1c. In one study, we compared 127 Denver police officers, a diverse sample of 113 officers from across the nation, and 124 Denver community members (Correll et al., 2007b, Study 1). The community members' performance was similar to that of untrained undergraduates: lay people made many errors, and they showed pronounced racial bias in both response times and decisions. Police outperformed the community in several ways. First, they were simply more likely to make correct decisions. In signal detection terms, they demonstrated enhanced sensitivity, or an ability to more effectively differentiate between armed and unarmed targets. Second, officers were faster to make those correct decisions. Third, and most critically, whereas the community made a biased pattern of errors, officers did not. Police occasionally shot unarmed targets, but they were no more likely to shoot an unarmed Black target than they were to shoot an unarmed White. This effect was driven by the fact that the community sample set a more lenient criterion for Black targets, favoring the *shoot* response, whereas officers employed strict criteria for both Blacks and Whites.

However, police were not immune to race. Although they showed no bias in their *decision criteria*, we observed clear, robust evidence of bias in officers' *response times*. Just like the community, officers were faster to shoot armed targets when they were Black (rather than White), and they were faster to choose a *don't-shoot* response if an unarmed target was White (rather than Black). This bias certainly suggests that, when confronted with a Black target, officers may activate racial stereotypes related to threat. In line with this possibility, response time bias was greatest among officers whose districts were characterized by a large (urban) population, a high rate of violent crime, and a greater concentration of Blacks and other minorities – environments likely to reinforce racial stereotypes. We wish to highlight the puzzle presented by the police officers' data. On one hand, response times indicate that police are aware of and responsive to target race. On the other hand, police show no bias in the decisions they ultimately make: when deciding to shoot, they set statistically equivalent criteria for Blacks and Whites.

We were concerned that our initial results might reflect the fact that the task was too easy for police (who, as noted above, were fast and highly accurate). We therefore conducted

another study, reducing the response window. We expected that this shorter window would prove more challenging, resulting in more errors and (potentially) the emergence of bias in SDT criteria among officers. Thirty-one Denver police officers and 45 community members participated. As expected, reducing the time window dramatically increased the number of errors. In fact, the officers in Study 2 performed somewhat worse than the community in Study 1. Even for the highly trained officers, then, this task presented a challenge. As in Study 1, community members showed bias, setting a lower decision criterion for Black targets. Despite the difficult task, however, police still showed no evidence of bias in their criteria.

Hypothesizing that practice accounted for the differences between officers and civilians, we attempted to assess its effect in the laboratory. Fifty-eight undergraduate participants performed two rounds (200 trials) of the FPST on 1 day. They then returned 48 h later and (again) performed two rounds. On the first round, participants showed clear evidence of bias in both response time and SDT criteria. However, the magnitude of the SDT bias decreased from Round 1 to Round 2, suggesting that training reduces bias in the decision to shoot (Plant & Peruche, 2005). It is interesting to note that training also improved overall performance, increasing estimates of sensitivity and reducing response time. However, when participants returned to the lab for their second session, they performed as though they were novices. The results demonstrate that practice can reduce bias and improve performance (making undergraduates look like police officers), but the effects were short lived.

Like undergraduates in our original research, these studies showed that laypeople demonstrate racial bias in both response times and SDT criteria. Police officers, however, differed from laypeople in several ways. They were faster, more accurate, and less biased in terms of the errors they commit. But racial bias was not *eliminated* among the officers. Rather, it emerged exclusively in response times. Police, like undergraduates and laypeople, were faster to shoot armed Blacks and faster to choose *don't shoot* for unarmed Whites. This pattern suggests that police attend to race and to racial stereotypes, even though the decisions they ultimately make are unbiased (Figure 3). How does this dissociation occur?

The Role of Cognitive Control in Expert Performance

We suggest that, when officers perform the FPST, practice and expertise enable them to minimize the behavioral consequences of stereotypes. We propose that, in the context of this threat-detection task, stereotypes engender a kind of prepotent response. To the extent that participants (including officers) associate Blacks with threat, the appearance of a Black target on the computer screen may promote a tendency to shoot. But, as with other prepotent responses, participants can learn to override that tendency.

In many ways, this process parallels the operations thought to be involved in responding accurately in the Stroop (1935) task, in which participants attempt to name the color of the typeface in which a word is printed. Because word reading is highly routinized, when participants see the word *BLUE*, they tend to respond by indicating, "blue". When the text is printed in a red typeface, however, participants must override this tendency to comply with task instructions. This is an effortful process that requires cognitive control (Friedman & Miyake, 2004).

We similarly propose that, when a Black target appears in the FPST, all participants (police officers, experts, and novices, alike) experience a prepotent tendency to shoot. But unlike novices, experts overcome that tendency through the exercise of cognitive control. There are at least two important implications of the prospect that police rely on control to minimize bias. First, if the nature of training fails to promote control, it should not reduce bias. Second, when the situation makes control difficult to exercise (e.g., by inducing stress or fatigue), it

should undermine the benefits of expertise. In stressful situations, then, even highly trained officers may demonstrate bias. We explore these implications in turn.

When training and experience fail to promote control

Typically, when “expert” undergraduates practice the FPST, they realize that race is not a useful cue. They quickly determine that they should focus on the object (“is it a gun?”) instead of the target person’s race, and they learn to minimize bias (Correll et al., 2007b), presumably by exercising control. But it is possible to practice the FPST in a manner that does *not* promote controlled processing. In a recent study, Sim, Correll, and Sadler (2013) asked participants to practice one of several versions of the FPST. Some participants performed a pro-stereotypic version in which (stereotypic) armed Blacks and unarmed Whites were overrepresented. In essence, these participants encountered a world where race was a valid, diagnostic cue – Black targets *were actually* more likely to pose a threat. Another group encountered a counterstereotypic environment during training in which *White targets* were more likely to be armed. All participants had the same opportunity to practice the judgment in question (shoot/don’t-shoot decisions), but some practiced in an environment that allowed them to take advantage of stereotypic shortcuts and still arrive at the correct answer (requiring little control), while others practiced in an environment that required them to avoid stereotypes (using control to resist the prepotent impulse). We reasoned that practice should only promote controlled processing when participants were forced to overcome prepotent response tendencies. As predicted, the results showed that training with the counterstereotypic (rather than the pro-stereotypic) task reduced bias.

A real-world demonstration of this idea comes from work comparing different groups of police officers. Many police might be called “beat cops”. They patrol a specified area, responding to calls from dispatch and interacting with a wide range of people – some criminals, but others who may be victims or people who are just going about their business. These officers have to carefully assess *every* person in *every* situation. This should promote control. These patrol officers can be contrasted with officers serving on units dedicated to gangs and street crime. Street crime units often take a more proactive (some might say aggressive) approach, stopping and frisking young men, conducting surveillance and raids of narcotics operations, and often focusing heavily on minorities. The New York Times wrote the following about the NYPD street crimes unit:

They make up less than 2 percent of the police force, but they seize 40 percent of all illegal guns confiscated in the city. They proudly proclaim, ‘We own the night,’ and... express their devotion to hunting down armed criminals. With an exceptional blend of accomplishment and bravado, the Police Department’s street crimes unit stands as the most striking example of both the success and the risk of the city’s aggressive approach to law enforcement. (Kocieniewski, 1999)

Ten years later, the Chicago Tribune echoed,

Officers assigned to such units... tend to make a lot of arrests and take violent drug dealers off the streets. But the relative freedom on the street given to the units has also fostered rogue tendencies in some instances and created a string of misconduct scandals. (Heinzmann, 2008)

Due to the relatively unconstrained nature of their interactions with the community, special-unit officers might not develop the same kind of controlled responses that patrol

officers do. We tested a small group of officers drawn from gang and street crime units. As described above, beat cops generally show no evidence of SDT bias in our task. The special-unit officers, however, showed robust bias. Similar to untrained members of the community, these officers showed clear SDT bias (Sim et al., 2013) (Figure 3).

When circumstances compromise control

If experts activate racial stereotypes but avoid a biased pattern of errors (or bias in the SDT criteria) by exerting cognitive control, then depriving them of the ability to exercise control should exacerbate SDT bias. To test this hypothesis, we randomly assigned participants to either an expert condition (in which they received extensive practice on the FPST) or a novice condition (in which they received no practice) (Correll, Wittenbrink, Axt, Miyake, & Goyle, 2014). We then asked both groups to complete a test phase of the FPST. In the test phase, participants performed the task while listening to a series of auditory stimuli (numbers ranging from 0 to 9). In a no-load condition, participants were told to simply ignore the numbers. In a medium-load condition, the task was made more demanding by asking participants to indicate whether each number was greater than or less than the number 5. Finally, in a high-load condition, we increased the difficulty even further. Participants were asked to compare each number they heard to the number that preceded it. For example, in the series, ...4, 7, 6, 3..., they would need to compare 7 to 4, then 6 to 7, and then 3 to 6. This is known as a 1-back task, and it requires participants to store one stimulus in memory (e.g., remember 4), listen for the next number and compare it to the stored value ($7 > 4$), and then update the stored number (forget 4, remember 7) in preparation for the subsequent stimulus. Because a 1-back task requires multiple cognitive operations, it is thought to induce a greater cognitive load, depriving participants of the capacity to process other information (Jonides & Smith, 1997).

In the no-load condition, performance mirrored previous training studies. Novices showed evidence of SDT bias; experts did not. When the demands of the situation were low, experts were no more likely to shoot a Black target than they were to shoot a White target. The core question, however, involves the impact of cognitive load. If the experts require cognitive resources to override racial stereotypes, then when they are deprived of those resources by the number-judgment task, they should be unable to exert control. In a demanding task, bias should re-emerge – even among experts. As predicted, cognitive load eliminated the beneficial effect of training. As resources were taxed, experts showed more and more bias. In the high-load condition, experts not only showed statistically significant SDT bias (i.e., a more lenient criterion to shoot Blacks), but they also showed as much bias as the novices in the no-load condition.

Conclusions: Translations from Lab to Street

The death of Amadou Diallo stimulated more than a decade of research examining the psychological processes involved in weapon identification and decisions to shoot. Our review has focused primarily on experimental laboratory research, and we want to highlight the overlap between a concrete real-world scenario and experimental social psychology. For police and students of social cognition, both, it is crucial to consider (a) the need to respond quickly; (b) the ambiguity of relevant information; (c) the presence of peripheral cues like race, gender, and environment; and (d) the potential influence of difficult-to-control associations between those cues and threat. This overlap offers hope that lab-based research may ultimately help us understand something crucial about the processes that unfold in the

streets. But applications have been hard to pin down. We must acknowledge that the experience of an officer can never be ethically simulated in a lab and that there are many factors that complicate police decision-making. It is only now – after more than 10 years of research – that we have reached a point where we feel somewhat confident commenting on police work. As we conclude this review, we consider factors that dramatically affect police decision making (fatigue and fear), which may compromise controlled processes critical to the reduction of racial bias.

Fatigue

When considering the conditions officers face, we must highlight fatigue. According to a study conducted by the DOJ (Vila, Kenney, Morrison, & Reuland, 2000), 41% of police officers showed signs of severe sleep deprivation. Research investigating police fatigue and its association with racial bias is scant. However, police shootings tend to occur at night when police officers report being most fatigued (Vila & Kenney, 2002) and experimentally induced fatigue increases bias in the weapon identification task and FPST (Govorun & Payne, 2006; Ma, Correll, Wittenbrink, Bar-Anan, Sriram, & Nosek, 2013). Fatigue presumably compromises cognitive control, increasing an individual's reliance on heuristics, including stereotypes (Macrae, Milne, & Bodenhausen, 1994).

Fear and Arousal

Perhaps the starkest contrast between laboratory settings and real-world police encounters involves fear and arousal. Accounts of officer-involved shootings suggest that these events involve a stunning departure from normal psychological functioning (Grossman & Christensen, 2004). Shootings may trigger a state known as “hypervigilance”, in which participants frantically seek escape and/or engage in a variety of seemingly nonsensical behaviors. According to Olson (1998), “officers experiencing hypervigilance might repeatedly pull the trigger of an empty weapon, misidentify innocuous items as weapons, or not see or hear innocent bystanders in the line of fire” (p. 5). Laboratory work presumably never comes close to inducing this kind of psychological state (for work that comes closer, see James, Vila, & Daratha, 2012).

Over a century ago, researchers identified an inverted u-shaped function relating arousal to performance (Yerkes & Dodson, 1908). Moderate levels of arousal improve performance, in part by enhancing attention to relevant cues in the environment (Cornsweet, 1969; Lorist & Tops, 2003; Reeves & Bergum, 1972). However, low levels of arousal (e.g., boredom) and very high levels of arousal (e.g., panic) can impair performance. High arousal impairs performance on tasks that require inhibition of a prepotent response (Foreman, Barraclough, Moore, Mehta, & Madon, 1989; Lorist & Tops, 2003) and facilitates the participant's dominant response (Hull, 1943; Zajonc, 1965). Lambert, Payne, Jacoby, Shaffer, Chasteen, and Khan (2003) argued that stereotypes constitute a dominant response and showed that aroused participants made more stereotypic inferences about a Black individual and expressed more racist attitudes on explicit measures of prejudice. It is plausible that the intense arousal experienced during an officer-involved shooting may impair control, leading even highly trained officers to show bias.

If training eliminates racial bias only for officers who are well rested, calm, and focused, then the benefits of training may evaporate when they are needed most. In real-world encounters, optimal conditions rarely apply: police are often fatigued and scared. Is this part

of the reason for persistent race-based discrepancies in police shooting? One important hypothesis derived from this line of research is that low-intensity training may not be sufficient to eliminate bias outside the lab. To reduce bias in high-stakes situations, it may be important for police to *train* in situations that are similar to actual officer-involved shootings (Beilock & Carr, 2001). For example, intense video or live-action training simulations that induce higher levels of arousal may help officers develop the capacity to focus on relevant information (e.g., the nature of the object in a suspect's hand) in a real encounter, when stress is high. It may also be useful to consider forms of control that do not rely on cognitively intensive operations. For example, Mendoza, Gollwitzer, and Amodio (2010) showed that implementation intentions (thought to induce reflexive control) improved accuracy, even for novices. Further research is clearly required on this issue. As experimental work progresses, improving our understanding of both low-level psychological mechanisms (e.g., Amodio et al., 2004; Correll et al., 2006) and performance in more externally valid scenarios (James et al., 2012), we hope to realize the potential of this “applied” work to make more meaningful contact with public policy.

Note

* Correspondence: Department of Psychology & Neuroscience, University of Colorado, Boulder, Campus Box 345, Boulder, Colorado, 80309-0345, USA. E-mail: correll@colorado.edu

Work Cited

- Amodio, D. M., Harmon-Jones, E., Devine, P. G., Curtin, J. J., Hartley, S. L., & Covert, A. E. (2004). Neural signals for the detection of unintentional race bias. *Psychological Science*, *15*(2), 88–93.
- Beilock, S. L., & Carr, T. H. (2001). On the fragility of skilled performance: What governs choking under pressure?. *Journal of experimental psychology: General*, *130*(4), 701.
- Cornsweet, D. M. (1969). Use of cues in the visual periphery under conditions of arousal. *Journal of Experimental Psychology: Journal of Experimental Psychology*, *80*(1), 14.
- Correll, J., Wittenbrink, B., Crawford, M. T., & Sadler, M. S. (2014). Connecting the dots: Stereotypes disambiguate complex visual stimuli. Manuscript in preparation.
- 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*(6), 1314.
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2007a). The influence of stereotypes on decisions to shoot. *European Journal of Social Psychology*, *37*(6), 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 and Social Psychology*, *92*(6), 1006.
- Correll, J., Urland, G. R., & 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*(1), 120–128.
- Correll, J., Wittenbrink, B., Axt, J., Miyake, A., & Goyle, A. (2014). When practice makes prudent: Training, task complexity, and control over racial bias in the decision to shoot. Manuscript in preparation.
- Department of Justice. (2001). *Policing and homicide, 1976–98: Justifiable homicide by police, police officers murdered by felons (NCJ 180987)*. Washington, DC: Bureau of Justice Statistics.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.
- Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology*, *22*(1), 22–37.
- 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.
- 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.
- Flynn, K. (1999). The focus of Diallo murder trial: Testimony of four police officers. *New York Times*, **CXLIX 51732**, A1.
- Foreman, N., Barraclough, S., Moore, C., Mehta, A., & Madon, M. (1989). High doses of caffeine impair performance of a numerical version of the Stroop task in men. *Pharmacology Biochemistry and Behavior*, *32*(2), 399–403.

- Friedman, N. P., & Miyake, A. (2004). The relations among inhibition and interference cognitive functions: A latent variable analysis. *Journal of Experimental Psychology: General*, **133**, 101–135.
- Fritsch, J. (2000). The Diallo verdict: The overview; 4 officers in Diallo shooting are acquitted of all charges. *The New York Times*, **CXLIX 51310**, A1–B6.
- Gaertner, S. L., & McLaughlin, J. P. (1983). Racial stereotypes: Associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly*, **46**, 23–30.
- Geller, W. A. (1982). Deadly force: What we know. *Journal of Police Science and Administration*, **10**, 151–177.
- Govorun, O., & Payne, B. K. (2006). Ego – Depletion and prejudice: Separating automatic and controlled components. *Social Cognition*, **24**(2), 111–136.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*, **74**(6), 1464.
- Grossman, D., & Christensen, L. W. (2004). *On Combat: The Psychology and Physiology of Deadly Conflict in War and Peace*. Belleville, IL: PPCT Research Publications.
- Heinzmann, D. (2008) After disbanding scandal-plagued Special Operations Section, Chicago police to start special gang unit. *Chicago Tribune*. Retrieved from http://articles.chicagotribune.com/2008-10-07/news/0810060694_1_fight-gangs-scandal-unit.
- Hull, C. (1943) Principles of behavior.
- Inn, A., Wheeler, A. C., & Sparling, C. L. (1977). The effects of suspect race and situation hazard on police officer shooting behavior. *Journal of Applied Social Psychology*, **7**(1), 27–37.
- Jacobs, D., & O'Brien, R. M. (1998). The determinants of deadly force: A structural analysis of police violence. *The American Journal of Sociology*, **103**(4), 837–862.
- James, L., Vila, B., & Daratha, K. (2012). Results from experimental trials testing participant responses to White, Hispanic and Black suspects in high-fidelity deadly force judgment and decision-making simulations. *Journal of Experimental Criminology*, **9**, 1–24.
- Jonides, J., & Smith, E. E. (1997). The architecture of working memory. In M. D. Rugg (Ed.), *Cognitive Neuroscience* (pp. 243–276). Cambridge, MA: MIT Press.
- Kocieniewski, D. (1999). Success of Elite Police Unit Exacts a Toll on the Streets. *New York Times*, **15**, 1.
- Kunda, Z., & Sherman-Williams, B. (1993). Stereotypes and the construal of individuating information. *Personality and Social Psychology Bulletin*, **19**(1), 90–99.
- Lambert, A. J., Payne, B. K., Jacoby, L. L., Shaffer, L. M., Chasteen, A. L., & Khan, S. R. (2003). Stereotypes as dominant responses: On the “social facilitation” of prejudice in anticipated public contexts. *Journal of Personality and Social Psychology*, **84**(2), 277.
- Lorist, M. M., & Tops, M. (2003). Caffeine, fatigue, and cognition. *Brain and cognition*, **53**(1), 82–94.
- Ma, D. S., Correll, J., Wittenbrink, B., Bar-Anan, Y., Sriram, N., & Nosek, B. (under review). When fatigue turns deadly: The association between fatigue and racial bias in the decision to shoot. Manuscript submitted for publication.
- 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.
- 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.
- Mendoza, S. A., Gollwitzer, P. M., & Amodio, D. M. (2010). Reducing the expression of implicit stereotypes: Reflexive control through implementation intentions. *Personality and Social Psychology Bulletin*, **36**(4), 512–523.
- Olson, D. T. (1998). Improving deadly force decision making. *FBI Law Enforcement Bulletin*, **67**(2), 1–9.
- 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**(2), 181.
- Plant, E. A., & Peruche, B. M. (2005). The consequences of race for police officers' responses to criminal suspects. *Psychological Science*, **16**(3), 180–183.
- 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**(2), 141–156.
- Reeves, F. B., & Bergum, B. O. (1972). Perceptual narrowing as a function of peripheral cue relevance. *Perceptual and Motor Skills*, **35**(7), 19–724.
- Sadler, M. S., Correll, J., Park, B., & Judd, C. M. (2012). The world is not black and white: Racial bias in the decision to shoot in a multiethnic context. *Journal of Social Issues*, **68**, 286–313.
- 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.
- Sim, J. J., Correll, J., & Sadler, M. S. (2013). Understanding police and expert performance when training attenuates (vs. exacerbates) stereotypic bias in the decision to shoot. *Personality and Social Psychology Bulletin*, **39**(3), 291–304.

- Smith, B. W. (2004). Structural and organizational predictors of homicide by police. *Policing: An International Journal of Police Strategies & Management*, **27**(4), 539–557.
- Sorensen, J. R., Marquart, J. W., & Brock, D. E. (1993). Factors related to killings of felons by police officers: A test of the community violence and conflict hypotheses. *Justice Quarterly*, **10**(3), 417–440.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of experimental psychology*, **18**(6), 643.
- Teahan, J. E. (1975). A longitudinal study of attitude shifts among black and white police officers. *Journal of Social Issues*, **31**(1), 47–56.
- Terrill, W., & Mastrofski, S. D. (2002). Situational and officer-based determinants of police coercion. *Justice Quarterly*, **19**(2), 215–248.
- Terrill, W., & Reisig, M. D. (2003). Neighborhood context and police use of force. *Journal of Research in Crime and Delinquency*, **40**(3), 291–321.
- Vila, B., & Kenney, D. J. (2002). Tired Cops: The Prevalence and Potential Consequences of Police Fatigue. National Institute of Justice.
- Vila, B. J., Kenney, D. J., Morrison, G. B., & Reuland, M. (2000). Evaluating the effects of fatigue on police patrol officers: Final report. US Department of Justice (DOJ), National Criminal Justice Reference Service (NCJRS).
- 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.
- Yerkes, R. M., & Dodson, J. D. (1908). The relation of strength of stimulus to rapidity of habit-formation. *Journal of comparative neurology and psychology*, **18**(5), 459–482.
- Zajonc, R. B. (1965). Social facilitation. Research Center for Group Dynamics, Institute for Social Research, University of Michigan.