

Weapon Bias

Split-Second Decisions and Unintended Stereotyping

B. Keith Payne

University of North Carolina at Chapel Hill

ABSTRACT—*Race stereotypes can lead people to claim to see a weapon where there is none. Split-second decisions magnify the bias by limiting people's ability to control responses. Such a bias could have important consequences for decision making by police officers and other authorities interacting with racial minorities. The bias requires no intentional racial animus, occurring even for those who are actively trying to avoid it. This research thus raises difficult questions about intent and responsibility for racially biased errors.*

KEYWORDS—*implicit; attitude; stereotyping; prejudice; weapon*

The trouble with split-second decisions is that they seem to make themselves. It is not simply that snap decisions are less accurate than “snail” decisions; it is easy to understand why people might make random errors when thinking fast. If you only have 30 seconds, it is probably a bad idea to do your taxes, pick a stock, or solve any problem beginning with “Two trains leave the station . . .” The real puzzle is when snap judgments show systematic biases that differ from our considered decisions. Should I consider those decisions *my* decisions if they differ from my intentions? Who is responsible?

These questions are asked most loudly when decisions have immense consequences, as when a split-second decision has to be made by a surgeon, a soldier, or a police officer. Four New York City police officers had to make that kind of decision while patrolling the Bronx on a February night in 1999. When the officers ordered Amadou Diallo to stop because he matched a suspect's description, Diallo reacted unexpectedly. Rather than raising his hands, he reached for his pocket. The Ghanaian immigrant may have misunderstood the order, or maybe he meant to show his identification. The misunderstanding was mutual: One officer shouted, “Gun!” and the rest opened fire.

Address correspondence to Keith Payne, Department of Psychology, University of North Carolina, Chapel Hill, Campus Box 3270, Chapel Hill, NC 27599; e-mail: payne@unc.edu.

Only after the shooting stopped was it clear that Diallo held only his wallet.

Many in the public were outraged. Some accused the NYPD of racial bias. Congress introduced legislation. Protests followed the officers' acquittal, in which the defense successfully argued that at the moment of decision, the officers believed their lives were in danger and that they therefore did not have the conscious intent, the *mens rae* (literally, “guilty mind”) to commit a crime. The court did not consider the mechanisms that might produce such a belief.

The death of Amadou Diallo dragged into the spotlight some of the disquieting questions that have run through implicit social cognition research for some time. Can stereotypes about race influence such split-second decisions? And can that kind of race bias take place without intent to discriminate? To answer these questions, it is necessary to move away from the particulars of the Diallo case and toward controlled studies in which causes and mechanisms can be identified. What are the psychological factors that would lead a person, in the crucial moment, to shout, “Gun”?

THE WEAPON BIAS

To study these questions, we developed a laboratory task in which participants made visual discriminations between guns and harmless objects (hand tools). A human face flashed just before each object appeared: a black face on some trials, a white face on others (see Fig. 1). The task for participants was to ignore the faces and respond only to the objects (Payne, 2001). There were two versions of the experiment. In one version, participants responded at their own pace. In the other version they had to respond within half a second on each trial. In the self-paced condition, accuracy was very high regardless of race. However, participants detected guns faster in the presence of a black face. This suggested that the black face readied people to detect a gun but did not distort their decisions.

In the snap-judgment condition, race shaped people's mistakes. They falsely claimed to see a gun more often when the face was black than when it was white (Fig. 2). Under the pressure of

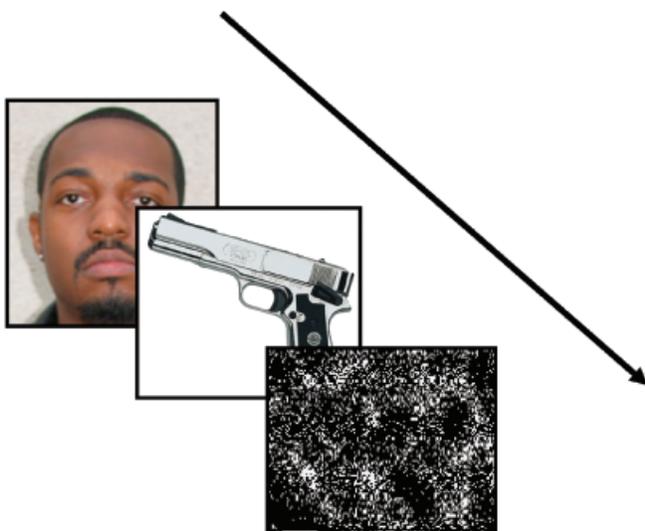


Fig. 1. Schematic illustration of weapons-priming procedure. On each trial, a white or black face appears first, followed by a gun or hand tool, followed by a visual mask. Participants' task is to indicate, as quickly as possible, whether they saw a gun or a tool.

a split-second decision, the readiness to see a weapon became an actual false claim of seeing a weapon.

These effects are not bound to the details of a particular experimental paradigm. Several independent lab groups have reported strikingly similar results using a variety of different procedures. For example, one procedure presented photos of black and white men who appeared on a computer screen holding a variety of objects such as guns, bottles, or cell phones (Correll, Park, Judd, & Wittenbrink, 2002). Participants were told to “shoot” any armed person by pressing one button, and to “not shoot” unarmed persons by pressing a different button. Another procedure presented pictures of white and black men

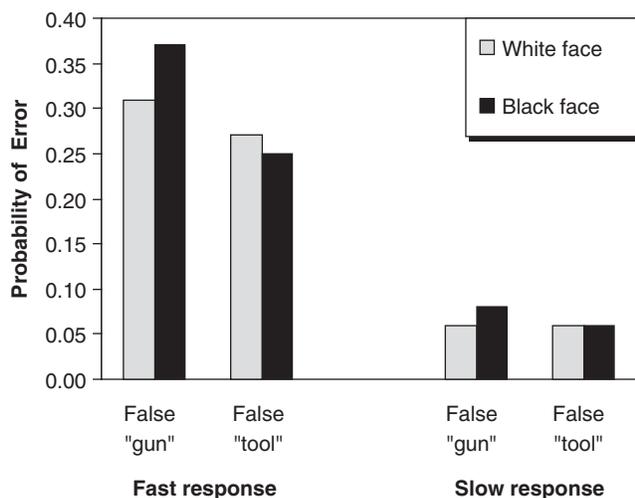


Fig. 2. Probability of falsely identifying a “gun” or “tool” depending on the race of the person shown prior to the object and whether participants were under time pressure to respond. Data adapted from Payne (2001).

popping out from behind obstacles, again holding either guns or harmless objects (Greenwald, Oakes, & Hoffman, 2002). During some phases of the study, participants were instructed to shoot if a white person, but not a black person, was armed. In other phases, the instructions were reversed. All of these procedures have provided evidence of race bias in both response times and errors. Although the samples in these studies have often been convenience samples, the data suggest that the bias is widespread. Responses made by African American participants in one study were indistinguishable from those of European American participants: both groups were biased toward claiming weapons in black hands more than in white hands (Correll et al., 2002).

Though participants did not need to use race to make their judgments, these studies provide no proof that the bias is unintentional in the strong sense of happening despite intentions to the contrary. Another study tested whether intentional use of race was necessary to produce bias (Payne, Lambert, & Jacoby, 2002). In a baseline condition, participants completed the weapon task under instructions to ignore the faces altogether. A second group was told that the faces might bias them and was instructed to try to avoid being influenced by race. Finally, a third group was also told about the biasing potential of the faces but was instructed to intentionally use the race of the faces as a cue to help them identify guns.

Results showed that although participants' goals affected their self-reported intentions, such goals did not improve their performance. Reliable race bias emerged in all three conditions and was in fact greater in both the “avoid race bias” and the “use race bias” conditions than in the baseline condition. Ironically, directing attention to race had exactly the same effect whether participants attended to race with the intent to discriminate or with the intent to avoid discrimination. In this and other studies, the weapon bias seems largely independent of intent. This is important because it means that the bias can coexist with conscious intentions to be fair and unbiased.

WHAT DRIVES THE WEAPON BIAS?

Why is it that people use stereotypes in their decisions both when they intend to and when they intend not to? And if we are not to turn intelligent people into caricatures or automatons, shouldn't intentions play a role somewhere? Integrating intentional and unintentional aspects of behavior is the job of dual-process theories, which attempt to explain when, how, and why behavior is driven by automatic versus intentionally controlled aspects of thought. My collaborators and I have proposed a particular dual-process theory to account for both intentional control over decisions and the patterns of unintended bias seen in snap judgments (Payne, 2001; Payne, 2005; Payne, Lambert, & Jacoby, 2002).

The first factor is a stereotypic association that, for some people, links African Americans to violence and weapons. These

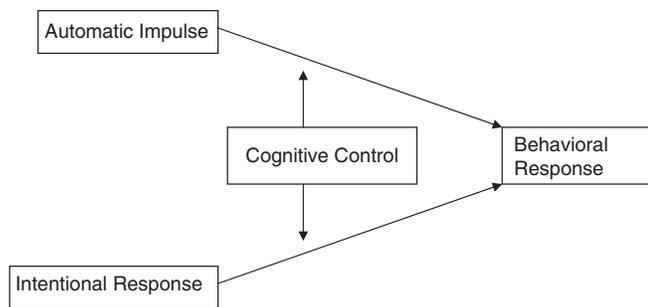


Fig. 3. A dual-process model of weapon bias. When people have full control of their behavior, they respond as intended. When control is impaired, automatic impulse drives responses.

stereotypic links can include both purely semantic associations and emotions such as fear or anger. These associations serve as an impulse that automatically drives responses whenever a person is unable to control a response. The second factor is the degree of intentional control participants have over how they respond (see Fig. 3). To predict whether someone will show the weapon bias, it is critical to know the answers to two questions. First, what is the person's automatic impulse that will drive responses when behavioral control fails? Second, how likely is it that control will fail? Research using a variety of behavioral and neuroscience methods has provided support for the key claims.

Behavioral Evidence

Evidence for the role of stereotypic associations comes from studies of individual differences. One study found that individuals with more negative self-reported attitudes toward blacks showed greater race bias in their weapon claims (Payne, 2001). In another study, weapon bias correlated with individual differences in perceptions of cultural stereotypes about African Americans (Correll et al., 2002). To avoid the limitations of self-reports, a recent study had participants complete two popular implicit-attitude measures in addition to the weapons task (Payne, 2005). Because implicit measures assess attitudes indirectly, without asking for a self-report, they avoid problems of introspection and social-desirability bias that affect explicit or self-report measures. Individuals with more negative implicit attitudes toward Blacks showed greater weapon bias. Finding consistent correlations using multiple measures provides converging evidence for the important role of stereotypic associations.

The finding that people with stronger stereotypes tend to show greater weapon bias is deceptively simple. It is deceptive because it tempts us to conclude that automatic stereotyping is all there is to the story. But that conclusion leaves out the important factor of how much intentional control people have over their responses. In the first studies described above, there was only one key difference between the snap-judgment and the slow-

judgment conditions: how much time participants had to respond. Snap judgments didn't change people's stereotypes. Snap judgments allowed those stereotypes to spill out into overt behavioral errors.

Time pressure is only one way to limit control over responses. Govorun and Payne (2006) showed similar effects as a result of self-regulation depletion. When people are required to self-regulate in one way, they are less likely to control themselves in other ways (Muraven & Baumeister, 2000). We depleted resources for one group of participants by requiring them to persist for several hundred trials on the tedious Stroop color-naming task. The Stroop task presents color words (e.g., *red*, *green*) in font colors that are either congruent or incongruent with the word meanings. When participants name the font color, incongruent word meanings interfere, requiring cognitive control. A nondepleted group saw a few trials of the Stroop task but did not exert sustained effort. The depleted group showed greater weapon bias, a result of reduced control over their responses.

Neuroscience Evidence

Several studies have examined the neural underpinnings of the weapon bias. Event related potentials (ERP) are more useful than other methods such as functional magnetic resonance imaging or positron emission tomography for this split-second effect because ERPs have greater temporal resolution. ERP studies examine fluctuations in electrical brain activity as a person processes information. Because different ERP components reflect specific cognitive functions, researchers can use those components to reveal processes underlying behavior.

One informative study examined an ERP component called error-related negativity (ERN), which is associated with detecting conflicts between goals and ongoing mental activity (Amodio et al., 2004). Conflict detection is a critical part of mental control because detecting a conflict between current and intended states is necessary for implementing self-control. Individuals showing the greatest ERN activity showed the fewest false weapon claims, and this effect was mediated by the ability to control responses.

A second study using ERP methods found several additional ERP components associated with weapon biases (Correll, Urland, & Ito, 2006). Of particular interest were two components, known as the P200 and the N200. The P200 is associated with emotional reactions to threatening stimuli, whereas the N200 is associated with conflict detection and cognitive control—similar to what was found with the ERN. Consistent with the two-factor theory, participants with greater P200 responses to black individuals, and those with lesser N200 responses, showed greater race bias.

Modeling the Weapon Bias

The evidence reviewed here converges to suggest that both automatic stereotype activation and failures of control are im-

portant in the weapon bias. Dual-process theories are commonly tested by comparing implicit and explicit tests, on the assumption that implicit tests measure only automatic responses and explicit tests measure only controlled responses. That assumption is not likely to be realistic, however, as virtually any task reflects a combination of automatic and controlled components (Jacoby, 1991). An alternative approach is to use a formal model to separate component processes within the same task. The value in this approach is that each component process can be studied individually without confounding underlying processes with different test formats.

My collaborators and I have used the process-dissociation procedure (Jacoby, 1991) as a tool to model automatic and controlled factors in the weapon bias. By that model, if a process is automatic, it influences responses regardless of whether it is consistent with intent or inconsistent with intent. In contrast, when a process is controlled, it influences responses only when intended, but not otherwise. When a black face precedes a gun, stereotypes and intent are in concert. Responding based on either will lead to the correct response. When a black face precedes a harmless object, stereotypes and intent are in opposition. The relationships among intentional control, automatic stereotyping, and behavioral responses can be formalized using algebraic equations (Jacoby, 1991; Payne, 2001). We can then decompose responses into numeric estimates of two processes: automatic stereotyping and cognitive control.

Applying the model to the studies just reviewed sheds light on the factors driving the weapon bias. For example, time pressure (Payne, 2001) and self-regulation depletion (Govorun & Payne, 2006) affected only the controlled component but not the automatic component. In other cases, differences in automatic stereotype activation were key. For example, implicit measures of race attitudes correlated with the automatic but not the controlled component (Payne, 2005). The evidence from these studies supports the two-factor account of the weapon bias and provides a means of measuring the underlying factors. The utility of modeling the underlying processes becomes apparent when considering strategies to reduce the race bias.

REDUCING WEAPON BIAS

Bias-reduction strategies might take either of two approaches. On one hand, they can try to change the automatic impulse. On the other hand, they can try to maximize behavioral control. One intriguing study compared police officers and civilians drawn from the same communities and found that both groups showed weapon bias, though officers showed somewhat less bias than civilians (Correll, Park, Judd, Wittenbrink, Sadler, & Keese, 2006). Even more important, the officers with the most firearms training showed the least race bias. This finding suggests that the routine training that officers receive may effectively reduce weapon bias. There is evidence that practice in identifying weapons may have beneficial effects on both controlled and

automatic components of responses and that these benefits extend to police officer volunteers (Plant & Peruche, 2005; Plant, Peruche, & Butz, 2005).

Finally, a recent study shows that although people cannot simply will the weapon bias away, certain specific strategies may be able to eliminate the automatic component of the bias. Stewart and Payne (2006) had participants form simple plans that linked racial categories to specific counterstereotypic thoughts (Gollwitzer, 1999). For example, participants made the plan, "when I see a black face I will think 'safe.'" Unlike participants who simply tried to avoid bias, those who formed specific plans showed no automatic race bias. Together, these studies offer clues to how and why specific strategies may succeed or fail.

IMPLICATIONS AND FUTURE DIRECTIONS

Research on the weapon bias has been consistent in answering several basic questions. Race can bias snap judgments of whether a gun is present, and that bias can coexist with fair-minded intentions. Although overt hostility toward African Americans is probably sufficient to produce this bias, it is not necessary. The bias happens not just because of racial animus but because of stereotypical associations that drive responses when people are unable to fully control them.

The answers to these questions suggest many more questions. One question is how well, and under what conditions, these findings generalize to the decisions police and other authorities make. Samples of police officers provide some evidence that the effect generalizes to a critical population. However, all of the existing studies have used computer tasks, even the most realistic of which do not capture the complexity facing an actual police officer. Future studies might incorporate manipulations of suspects' race into real-time, three-dimensional simulations of the sort that are used in police firearms training.

A second question concerns the mechanisms underlying the weapon bias. Evidence suggests that both emotional responses to and semantic associations with race play a role (Correll, Urland, & Ito, 2006; Judd, Blair, & Chapleau, 2004). But it is unknown under what conditions one or the other is likely to be influential. Do emotional and semantic responses act in identical ways, or do they have different consequences? And do the mechanisms of control differ for emotional versus semantic responses?

Another important question concerns how people attribute responsibility for biases that demonstrably contradict intent. I received two letters shortly after the first paper on the topic was published. A retired police officer rejected the conclusion that race may bias weapon decisions, concerned that the research might lead to unjustified allegations that police, who must make the best decisions they can under terrible conditions, are prejudiced. A second letter writer objected to the conclusion that the weapon bias may happen without intent, concerned that the research might be used to excuse race bias among police officers rather than holding them accountable for their decisions.

It is difficult to dismiss the worries of either writer, though they are polar opposites. Each expresses some of the thorny possibilities that may reasonably follow from a complex situation. Do ordinary people consider this a case of diminished capacity and therefore diminished responsibility? Or do they perceive the bias to reflect hidden malice? Are their judgments biased by their own racial attitudes or their attitudes toward police?

Empirical research will not settle the hard normative questions of ethics and responsibility. But it can shed light on how ordinary people actually reason about such unintended biases. Because juries and other decision-making bodies are made up of these same people, the answers are important for how social and political institutions will treat unintended race biases. Understanding the psychology of the weapon bias is a prelude to a better-informed conversation about the hard questions.

Recommended Reading

- 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.
- Payne, B.K., Jacoby, L.L., & Lambert, A.J. (2005). Attitudes as accessibility bias: Dissociating automatic and controlled components. In R. Hassin, J. Bargh, & J. Uleman (Eds.), *The new unconscious*. Oxford, England: Oxford University Press.
-

REFERENCES

- 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*, 225–232.
- Correll, J., Park, B., Judd, C.M., & Wittenbrink, B. (2002). The police officer's dilemma: Using race to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology, 83*, 1314–1329.
- Correll, J., Park, B., Judd, C.M., Wittenbrink, B., Sadler, M.S., & Keesee, T. (2006). *Across the thin blue line: Police officers and racial bias in the decision to shoot*. Unpublished manuscript, University of Chicago.
- 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.
- Gollwitzer, P.M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist, 54*, 493–503.
- Govorun, O., & Payne, B.K. (2006). Ego depletion and prejudice: Separating automatic and controlled components. *Social Cognition, 24*, 111–136.
- Greenwald, A.G., Oakes, M.A., & Hoffman, H.G. (2002). Targets of discrimination: Effects of race on responses to weapon holders. *Journal of Experimental Social Psychology, 39*, 399–405.
- Jacoby, L.L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513–541.
- Judd, C.M., Blair, I.V., & Chapleau, K.M. (2004). Automatic stereotypes vs. automatic prejudice: Sorting out the possibilities in the Payne, 2001 weapon paradigm. *Journal of Experimental Social Psychology, 40*, 75–81.
- Muraven, M., & Baumeister, R.F. (2000). Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychological Bulletin, 126*, 247–259.
- 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.
- Payne, B.K. (2005). Conceptualizing control in social cognition: How Executive Functioning Modulates the Expression of Automatic Stereotyping. *Journal of Personality and Social Psychology, 89*, 488–503.
- 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.
- Plant, E.A., & Peruche, B.M. (2005). The consequences of race for police officers' responses to criminal suspects. *Psychological Science, 16*, 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*, 141–156.
- Stewart, B.D., & Payne, B.K. (2006). *Counterstereotypical thought plans reduce automatic stereotyping*. Unpublished manuscript, Ohio State University.