

JAMES P. MURPHY, ROLAND NEIL, JESSICA WELBURN PAIGE

Methodological Challenges for Research on Racial Bias in Police Shootings

Shootings of racial minority civilians by police officers have long been a major social issue in the United States, the subject of extensive media coverage, and a reality that has galvanized social movements, such as Black Lives Matter. Police in the United States shoot people at far higher rates than police in other economically developed countries (Zimring, 2017). The victims of these shootings are disproportionately Black and, to a lesser extent, Hispanic, compared with their shares of the population. These shootings occur in a society that continues to grapple with its oftentimes racist history, one that, to this day, remains beset with pervasive racial inequalities. This is especially true of the relationship between police in the United States and the

country's minority communities, which, although complex, has long been characterized by high degrees of both mistreatment and mistrust. It is perhaps no surprise, then, that this issue has become a national flashpoint and that high-profile police shootings weigh so heavily on the psyche of Americans, especially Black Americans.

In this context, researchers have increasingly sought to answer a fundamental question: To what extent are police shootings of civilians driven by racial bias? This is research that is high profile and high stakes, being published in prestigious journals, getting coverage in leading media

KEY FINDINGS

- The impact of police shootings is severe—particularly for communities of color in the United States—and more-nuanced research is needed to understand the dynamics of racial bias.
- Thinking of racial bias in police shootings as a process with a series of stages can improve research in this area. These stages include officers' encounters with civilians, use of force, and the type of force used.
- Fully unpacking how racial bias operates in police shootings is challenging due to data limitations. These limitations include unreliable federal and crowdsourced datasets, and the problem is particularly severe for non-fatal shootings.
- The most common forms of tests—benchmark and outcome tests—face significant challenges in inferring bias that research has only started to confront.
- We recommend several ways to improve understanding of the dynamics of racial bias in police shootings: better data, clearer definitions and estimands, better methods, and the use of bounding and sensitivity analyses.

outlets, and gaining exposure among a public in search of answers. However, research on this topic faces limitations: the quality of data available and the methods used to make inferences about racial bias. These limitations are significant enough to call into question the findings of this body of work.¹

In this report, we review the evidence to date (as of early 2024) on racial bias in police shootings, focusing on the data and methodological challenges that exist for research on this topic. In particular, we explore the importance of thinking of racial bias as a process with a series of stages, the data that are available to study bias in shootings, the main methods that have been used and their limitations, and—crucially—how researchers ought to proceed to arrive at stronger and more-informative conclusions. While we employ a critical eye, our goal is productive: to draw attention to the data and methodological improvements most needed to strengthen research on this pivotal topic.

What Is Racial Bias?

Research has consistently shown that racialized processes affect all areas of life for people of color in the United States. As a result, when it comes to police encounters—which can reflect the many ways in which race has affected one’s life to that point—it

Police shootings in racially segregated neighborhoods are embedded within a larger set of racial inequalities, many of which could reasonably be said to reflect racial bias.

is challenging to conceptualize and operationalize what aspects do or do not reflect racial bias (Sampson and Neil, 2023). For example, racial bias affects and constrains residential choices. As a result, people of color—particularly African American individuals—are more likely to reside in racially and economically segregated neighborhoods with limited resources and opportunities. This includes limited opportunities for educational and occupational attainment (Charles, 2003; Massey and Denton, 1993; Sharkey, 2013; Wilson, 1987; Wilson, 1996). In some cases, limited opportunities in segregated neighborhoods can lead to increased involvement in illegal activities, potentially increasing residents’ contact with law enforcement (Anderson, 1999; Bourgois, 2003; Wilson, 1987; Wilson, 1996). Research has also shown that residents of segregated neighborhoods may be targeted by law enforcement, leading to what Rios (2011) describes as over-policing of people of color. Thus, police shootings in racially segregated neighborhoods are embedded within a larger set of racial inequalities, many of which could reasonably be said to reflect racial bias.²

That said, in this report, we focus on a narrow form of racial bias that corresponds to the target of much research to date: whether officers treat similar people differently because of their perceived race.³ As we will discuss, more-formal definitions are often useful because they clarify the exact population of interest and under what conditions the methods used can estimate the quantity of interest. Also, those who prefer a formal definition should consider the objections of Kohler-Hausmann (2019), who questions whether it is sensible to talk about otherwise similar people of different races and reinforces that this definition of racial bias is ultimately a narrow one. Notably, our definition does not incorporate racially biased department- or city-level decisionmaking. For example, departmental policy or practices that lead to the over-policing of minority neighborhoods could, in a sense, reflect racial bias, which is an important possibility to consider, but one that lies beyond the scope of this report. Finally, racial bias in who is shot by police is distinct from differences among officers of different races because while the two may be related, White officers are not the only ones susceptible to having racial biases.

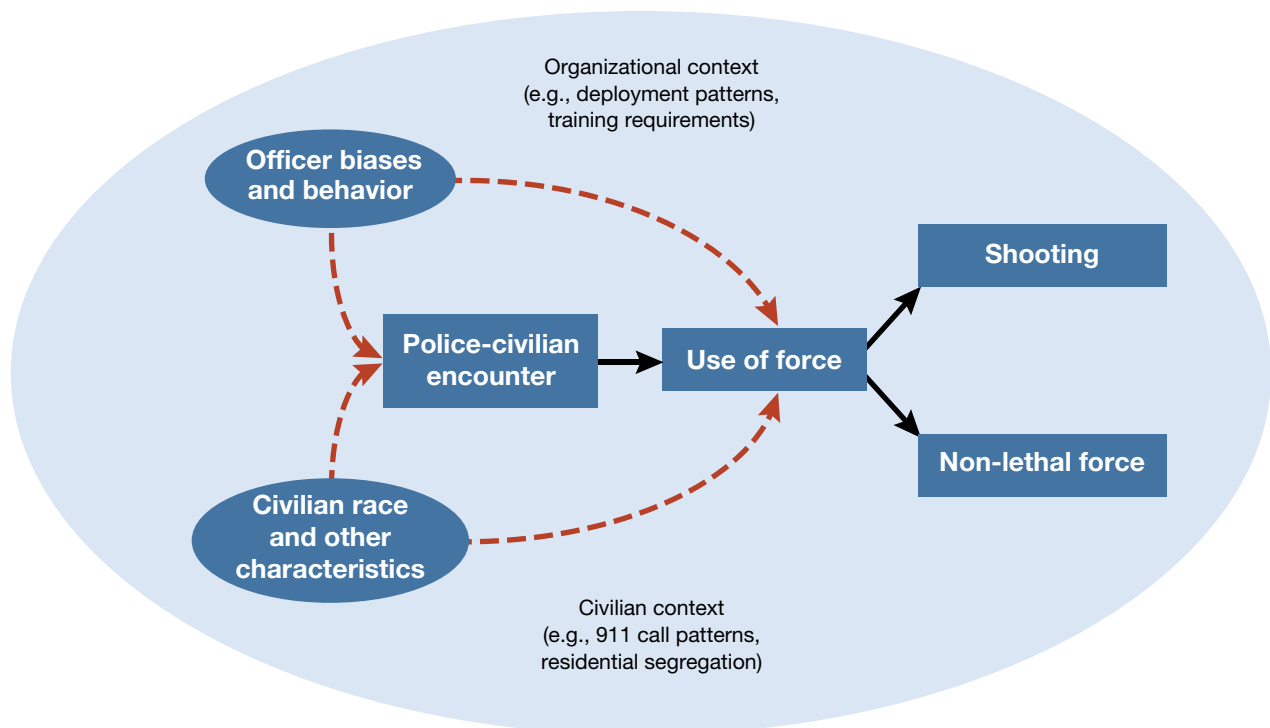
Racial Bias as a Process

How racial bias contributes to police shootings can be thought of as a multi-stage process, as illustrated in Figure 1. Racial bias can enter at each of these three stages. In the first stage, officers encounter a civilian, which is a prerequisite for officers to use force. If police have a lower threshold of suspicion for minority civilians when they make stops on patrol or decide toward whom to direct attention when responding to a call, there is racial bias in the encounter.⁴ Having encountered the civilian, an officer's racial bias might emerge at the next stage—the use of force—if the officer is more willing to use force against minority civilians than similar White civilians. This might occur, for example, if there is a greater expectation that a weapon will be drawn when a Black man's hand comes near his pocket than for a White man. Finally, the type of force used might vary by the civilian's race. If an officer is disproportionately likely to think that the White suspect is reaching for a knife but the Black suspect is reaching for a gun, they may be more likely to use a Taser or only point a gun at the former but quickly pull the gun's trigger when facing

the latter. Bias can exist to varying degrees at any of these three stages. The point at which one has data or decides to begin an analysis can shape a researcher's findings and the resulting policy implications.

While our review and assessment of research on racial disparities in police shootings focuses on these three stages, this dynamic unfolds in a broader organizational and societal context. Law enforcement agencies can influence racial disparities through biased policies and practices even if individual officers themselves are not biased. For example, bias in the deployment of (even unbiased) officers to neighborhoods likely results in racial disparities in encounters. Likewise, even without individual officer bias, racial bias in the broader civilian population can result, for example, in racial disparities in 911 calls for similar incidents. Such factors form the context in which the three-step process occurs. While we recognize their importance for policy and practice, these macro-level considerations lie beyond the scope of this review. We focus on the micro-level dynamic between individual police officers and civilians, which has dominated policy debates to date. Under-

FIGURE 1
Stages of Bias in Police Shootings



standing processes at this level of analysis can inform institutional reforms.

Recognizing racial bias as a process is important for at least two reasons. First, from a normative and policy perspective, the fact that an “unbiased” shooting occurs after a racially biased stop is of little comfort. If one is trying to reduce racial inequalities in police shootings, ignoring biased encounters obscures a potential avenue for effective policy solutions.

A second, more subtle methodological implication of ignoring racial discrimination in the encounter stage is known as *post-treatment bias*. This stems from the fact that biased selection into an encounter can lead to unobserved differences among the civilians with whom police interact, which may confound or suppress the association between civilian race and shootings. For example, suppose that at the encounter stage there is a lower threshold of suspicion—either by an officer or a civilian making a call for service—for encountering a Black civilian than a White civilian. This would lead to the set of White civilians that police interact with being “riskier” on average. If so, then an equal probability of Black and White civilians being shot after an encounter *does not* mean there is no racial bias in shootings—it means quite the opposite.

Biased selection into an encounter can lead to unobserved differences among the civilians with whom police interact, which may confound or suppress the association between civilian race and shootings.

The issue of post-treatment bias has spurred debate about a prominent paper by Fryer (2019).⁵ Fryer drew on a diverse set of administrative records: stops in New York City, weapons discharges against civilians in 16 jurisdictions, and a random sample of Houston arrests for violence against an officer or resisting arrest. Fryer concluded that there was racial bias in whether officers used force, but there was no evidence of bias in shootings. The latter conclusion was based primarily on the analysis of the Houston arrests. After adjusting for a large set of individual and contextual factors, such as officer and suspect demographics and reason for officer response, Fryer found that there was no statistically significant difference between White and minority suspects’ probabilities of being shot. While Fryer’s study was innovative in its use of rich contextual information on encounters, subsequent critiques (Durlauf and Heckman, 2020; Knox, Lowe, and Mummolo, 2020) highlighted that his reliance on arrest records and stops could introduce post-treatment bias if officers are racially biased in stopping or arresting civilians. In re-analyses accounting for post-treatment bias in the publicly available New York City Police Department (NYPD) stop data—which did not include shootings as an outcome—Knox, Lowe, and Mummolo (2020) argued that Fryer severely underestimated racial bias in officers’ decisions to draw weapons and point them at civilians. While the authors did not have access to the Houston data, these findings cast doubt on Fryer’s conclusions about the lack of racial bias in shootings.

An important question in studying racial bias in shootings is as follows: How much racial bias is there in police encounters? There is much research on this topic; we focus on a few high-profile, informative studies. Grogger and Ridgeway (2006) studied racial bias in vehicle stops using what they termed the *veil of darkness* hypothesis. Their test takes advantage of how driver race is less perceptible during darker hours to derive a benchmark of what unbiased policing would look like. They did not find evidence of racial bias in Oakland, Calif., vehicle stops. However, Pierson et al. (2020) more recently used this method on a dataset that they assembled on more than 100 million traffic stops from across the United States, and found evidence of racial bias against Black drivers. In another major study, Epp, Maynard-Moody, and

Haider-Markel (2014) surveyed more than 2,000 drivers from in and around the Kansas City, Mo., metropolitan area about their traffic interactions with police, driving behaviors, and other potentially relevant information. Using a variety of methods, they concluded that anti-Black racial bias in vehicle stops exists and is driven by pretextual, investigatory stops as opposed to situations in which a clear traffic-safety violation occurred.

Most research on racial bias in pedestrian encounters has focused on the NYPD's practice of Stop, Question, and Frisk. Gelman, Fagan, and Kiss (2007) used a complex form of benchmark test, finding patterns consistent with racial bias in pedestrian stops of Black civilians. However, Ridgeway (2007) built on that analysis using a wider range of benchmarks and data from a different period, but he did not find much evidence of bias in NYPD pedestrian stops. Goel, Rao, and Shroff (2016) found that Black and Hispanic pedestrians stopped by NYPD officers because there was suspicion that they were carrying a weapon were less likely to be carrying a weapon than White pedestrians, suggestive of a higher threshold in stopping White pedestrians.

While these studies employ relatively strong designs, testing for racial bias in police encounters is difficult (see Neil and Winship, 2019; Ridgeway and MacDonald, 2010). Also, the extent of racial bias in police encounters may be heterogeneous across place and time, such that there is no one answer to the question of how much racial bias drives police encounters. Large-scale datasets, such as that built by Pierson et al. (2020), are promising in this respect because they can be used to examine variability across contexts.

Importantly, it is unclear how relevant research on pedestrian or vehicle stops is in understanding racial bias in police shootings. Police officers spend much time responding to 911 calls and typically have little to no discretion in deciding whether they will do so, although they may have discretion in who they target once they respond to a call.⁶ Generally, different types of police encounters might be driven by racial bias to different degrees. Two empirical questions follow from these observations: How does racial bias in police encounters vary by encounter type? And what types of encounters tend to precede

How does racial bias in police encounters vary by encounter type? And what types of encounters tend to precede police shootings?

police shootings? Answering these questions is crucial to understand the role that racial bias in police encounters plays in racial bias in police shootings. Bias during and following 911 calls is understudied. Recently, Hoekstra and Sloan (2022) obtained 911 call data from two anonymous cities, but they used these data to study differences by the race of officers in the use of force rather than by the race of the victim. Thus, while this study provides evidence that, under certain conditions, Black, Hispanic, and White officers' shooting behaviors differ, the authors do not directly speak to whether minority and White civilians are treated differently.

While racial bias is a process and bias in encounters may exist, policy solutions do not necessarily need to intervene at the encounter stage to reduce police shootings. For example, Nix (2020) suggests that studying whether officers are biased in the decision to shoot once an encounter has occurred may be particularly useful to find solutions. Relatedly, while it is possible that there is racial bias in which individuals have 911 called on them, this might not be relevant in estimating racial bias in shootings because the police do not control who calls 911, reinforcing the importance of clearly specifying what is being estimated—that is, of defining racial bias.

Data on Police Shootings

Foundationally, estimating racial bias requires having high-quality data on police shootings. Ideally, such data should (1) be comprehensive, capturing a large and representative portion of law enforcement agencies (LEAs) and officers, (2) possess consistent measurement and reporting both over time and across jurisdictions, and (3) include contextual information on the circumstances under which shootings happened. Table 1 lists prominent national-level data sources that receive ongoing updates. It omits several sources used in past research that have not received public updates during the past three years.⁷ Many state and local LEAs have begun releasing data on fatal and non-fatal shootings at varying levels of granularity and detail, which may be useful, but because there is no central repository of this infor-

mation, we exclude these data from this overview of national datasets.⁸

National datasets can be divided into two types: *federal data* based on reports from state and local authorities and *crowdsourced data* gathered by media organizations and citizens. Both types of data suffer from potential biases that—for the moment—limit rigorous research estimating national rates of police violence and racial bias:

- Federal datasets intended to measure shootings and other uses of force by police rely on voluntary self-reports by LEAs. Recordkeeping protocols vary across agencies, and, without reporting requirements, there is no incentive for jurisdictions with high levels of police violence to participate. LEA reports also may be unreliable because officers have incentives

TABLE 1
Active or Recently Updated National-Level Datasets That Include Police Shootings

Dataset	Source	Scope	Time Span Available
Crowdsourced			
Fatal Encounters	News reports and public federal, state, local government data	Fatal	2000–2021
Gun Violence Archive	News reports and public federal, state, local government data	Fatal and non-fatal	2014–present
Mapping Police Violence	News reports, government sources, other crowdsourced databases	Fatal only	2013–present
<i>Washington Post</i> fatal police shootings database	News reports, social media, government sources, other crowdsourced databases	Fatal only	2015–present
Federal			
CDC National Violent Death Reporting System	Death certificates, coroner and medical examiner reports, police reports	Fatal only	2002–2020 ^a
CDC National Vital Statistics System	Death certificates	Fatal only	1980–2022
CPSC National Electronic Injury Surveillance System (NEISS)	Sample of hospitals (~100 per year)	Non-fatal only	2001–2020
FBI National Incident-Based Reporting System	Voluntary reports by LEAs	Fatal only	1991–2022 (FBI standard since 2021)
FBI National Use of Force Database	Voluntary reports by LEAs	Fatal and non-fatal	2019–present
FBI Supplemental Homicide Reports	Voluntary reports by LEAs	Fatal only	1976–2020

NOTE: CDC = Centers for Disease Control and Prevention; CPSC = Consumer Product Safety Commission; FBI = Federal Bureau of Investigation.
^a Since 2018, all 50 states, Puerto Rico, and the District of Columbia participate in the National Violent Death Reporting System, but coverage is still partial in some states.

to describe their use of force as more justified than it might have been. LEA-based data may underreport fatal shootings by up to 50 percent at the national level (Barber et al., 2016). In the near term, violent death data collected by health authorities may hold more promise for studies of fatal shootings.

- Crowdsourced data may suffer from unmeasurable biases because of geographic differences in media coverage and temporal variation in attention paid to police violence. There are no clear ways to assess the extent of such bias. Additionally, there are concerns over these databases' level of detail and coding of contextual information (e.g., whether a civilian was carrying a weapon).
- The most-comprehensive sources—both federal and crowdsourced—track only fatal shootings. Research based on nonrepresentative samples of jurisdictions that make non-fatal shooting data available suggests that about half of police shootings result in non-fatal injuries and that racial disparities are larger for non-fatal than fatal shootings (Clark et al., 2023; Nix and Shjarback, 2021). Thus, analysis based only on fatal outcomes may underestimate racial bias.

In this section, we provide a high-level overview of the strengths and weaknesses of each type of dataset.

Federal Data Sources

Various federal datasets rely on self-reports from state and local LEAs and public health authorities. Both types of datasets have insufficiently documented the prevalence of fatal police shootings. On the public health side, the CDC's National Vital Statistics System depends on death certificates—which frequently have limited information—and has undercounted law-enforcement-involved deaths by as much as 50 percent (Barber et al., 2016; Sharara et al., 2021). The CDC's newer National Violent Death Reporting System (NVDRS) supplements death certificate data with data from other sources. Comparisons with crowdsourced data suggest that NVDRS has

historically been much more reliable for participating states than other federal sources (Barber et al., 2016; Conner et al., 2019), but these analyses were based on periods during which NVDRS covered no more than 27 states. We are unaware of any similar assessments since NVDRS expanded in 2016 and 2018; these expansions added states that accounted for almost half of fatal police shootings between 2018 and 2022 in the crowdsourced Mapping Police Violence website (Mapping Police Violence, 2024).⁹ It is an open question as to whether NVDRS can be considered a relatively comprehensive source for fatal shootings.

LEA-based federal sources rely on voluntary participation and self-reports, which pose clear concerns about incentives for selective participation and reporting. The FBI's Supplementary Homicide Reports (SHR) database has compiled information on all homicides in participating jurisdictions since 1976, but Barber et al. (2016) found that SHR undercounted the number of police shooting fatalities by about 50 percent between 2005 and 2012, compared with NVDRS. Created in 2019, the FBI's National Use-of-Force Data Collection (NUFDC) was intended to provide a comprehensive database of the circumstances, subjects, and officers involved in all use-of-force incidents that resulted in the death or serious bodily injury of a civilian or involved firearm discharge by an officer (FBI, 2020e). However, NUFDC also has suffered from low participation rates (Jackman, 2021), and it has not provided useful public information on either the prevalence or context of shootings or use of force more broadly. The FBI's 2021 transition to the National Incident-Based

About half of police shootings result in non-fatal injuries, and racial disparities are larger for non-fatal than fatal shootings.

Reporting System (NIBRS) as the standard mechanism of crime data collection promises to make more-detailed incident-level contextual information available. As of 2022, participating LEAs covered almost 80 percent of the U.S. population, although non-participation is unevenly distributed throughout the country (FBI, 2022d).¹⁰

Crowdsourced Data

Faced with limited and suspect government data, over the past decade, journalists and citizens have mounted significant efforts to compile alternatives, such as the website Fatal Encounters, the *Washington Post's* Fatal Police Shootings Database, and the Mapping Police Violence website. Such efforts follow similar methodologies. Although their procedures differ slightly, generally, these aggregators conduct automated or (to a lesser extent) manual searches of local and national news databases and state and local police reports before manual inspection and coding by human reviewers. These crowdsourced efforts are more comprehensive than federal sources for fatal police shootings, but we may still expect systematic biases across jurisdiction and time because changes to the media landscape have reduced coverage of suburban rural areas (Abernathy, 2018). Systematic biases could also arise if certain police shootings were considered more newsworthy than others. Schroedel and Chin (2020) argued that the police's use of lethal force against Native Americans is underreported in

The very need to create these databases suggests that police shootings may have been under-documented in the past, even by the media.

major media outlets, whereas Zuckerman et al. (2019) found that coverage of police killings of unarmed Black Americans increased in 2014. Neither of these studies necessarily means that crowdsourced data are biased, although they do suggest that victim characteristics influence coverage. The very need to create these databases suggests that police shootings may have been under-documented in the past, even by the media. Sustained media attention and consistent, standardized data collection practices raise hope that crowdsourced databases will be valuable sources to track fatal shooting trends more reliably and evaluate the effects of future policy interventions to reduce fatalities (Cahill, Labriola, and Taylor, 2021).¹¹

However, crowdsourced data are limited in the contextual information they provide about an incident (Fryer, 2018). For example, the Fatal Encounters, Mapping Police Violence, and *Washington Post* databases have either considerable missing data or lack entirely such information as the officer's race, age, gender, or past use of force, or the incident time of day. These are all important factors that may influence inferences about whether a shooting is racially biased. This type of information is generally more extensively documented in police administrative records, which are only selectively made public. The most-rigorous claims about racially biased use of force have been based on administrative data from a small number of local urban and state police agencies, which either make these data public or agree to share the data with researchers.

Undocumented Non-Fatal Police Shootings

Barring increased participation and greater transparency in the NUFDC, the United States continues to lack a reliable national picture of non-fatal police shootings. Research based on nonrepresentative samples of states and local jurisdictions that make non-fatal shooting data available suggests that about half of shootings result in non-fatal injuries and that racial disparities are larger for non-fatal than fatal shootings (Clark et al., 2023; Nix and Shjarback, 2021).¹² Ward et al. (forthcoming) find similar results nationally using the crowdsourced Gun Violence

Archive.¹³ Poor documentation of non-fatal shootings is a problem for gun violence and policy research broadly, not only police shootings. In other work within the Gun Policy in America initiative, RAND researchers have generated more-comprehensive and improved longitudinal estimates of non-fatal firearm injuries by pooling information from multiple data sources, such as state inpatient data (Smart et al., 2021a). If the involvement of police is adequately recorded, similar efforts—such as using hospitalization data and combining data sources—could develop a more accurate picture of non-fatal police shootings in America.

Estimating Racial Bias: Benchmark and Outcome Tests

Research on racial bias in police shootings tends to use some form of benchmark or outcome test, both of which face important inferential challenges.¹⁴

Benchmark Tests

With *benchmark tests*, the number of people shot by police is divided by the number of people who would be at risk of being shot by police were the police unbiased separately for each race.¹⁵ Racial differences in the resulting numbers are interpreted as evidence of bias.

The greatest challenge with this method is finding the correct benchmark (Neil and Winship, 2019; Ridgeway and MacDonald, 2010; Tregle, Nix, and Alpert, 2019): that is, a metric that adequately proxies the number of people at risk of being shot. Reflecting that this is the denominator in benchmarking tests, this is known as the *denominator problem*. A common choice of denominator is the population of each race, typically drawn from the Census (Mentch, 2020; Ross, 2015; Schwartz and Jahn, 2020). This design invariably yields results suggesting bias in fatal police shootings of Black civilians, with more-mixed findings for Hispanic-White comparisons. However, benchmarking with population assumes that shootings would be randomly distributed among the population if the police were unbiased. Put differently, it does not incorporate any risk factors, such as civilian behavior, that influence the decision to use

Fatal shootings rarely actually follow arrests; they are more accurately thought of as different ways that encounters can end.

lethal force. If such factors vary by race, population benchmarks would yield inaccurate inferences about racial bias.

Consequently, researchers often use crime (Cesario, Johnson, and Terrill, 2019; Ross et al., 2021) or arrests (Mentch, 2020; Mullainathan, 2015) as their benchmarks instead to proxy risky encounters with the police. Most often, findings of racial bias are reduced or eliminated when using these instead of population as a benchmark. Unfortunately, this is not as clear-cut or promising a methodological solution as it might initially appear. As we have seen, insofar as racial bias in police encounters is part of the process that generates racial bias in police shootings, conditioning on police encounters could mask bias. What's more, fatal shootings rarely actually follow arrests; they are more accurately thought of as different ways that encounters can end. Thus, arrest benchmarks may not measure the right encounters for fatal shootings. This possible bias in fatal shooting data underscores the need for more-reliable data on non-fatal shootings, where an arrest is a more likely outcome following a shooting. Similarly, it is not clear why we should think that crimes reflect the proper risk set for either fatal or non-fatal incidents, because neither police witnessing a crime nor police encountering a criminal are sufficient reasons for them to shoot someone—even suspects of violent crimes.

Moreover, benchmark tests assume that people shot (the numerator) were part of the risk set used (the denominator). This assumption can be far off the mark, compromising inferences about racial bias

(Neil and Winship, 2019). Imagine, for example, that a researcher uses violent crime as a benchmark, but a nontrivial portion of those shot to death by police were not committing violent crimes. The larger that portion of people, the more similar this test is to arbitrarily dividing the number of shootings for different racial groups by different numbers. Given higher recorded rates of violent crime among African American individuals generally, the result would likely be that racial bias in police shootings against African American individuals is masked.¹⁶ Ross et al. (2021) devised a statistical bias correction to address this problem, showing that the finding of no racial bias in fatal police shootings from an earlier study using crime as a benchmark (Cesario, Johnson, and Terrill, 2019) is reversed once the bias correction is applied, leading Ross et al. to conclude that there is racial bias against Black civilians in fatal police shootings.

Ultimately, it is not clear whether any commonly used benchmark can measure the number of people whom unbiased police would have encountered and who were posing enough of a threat to warrant the use of deadly force against them. Certainly, estimates of racial bias are highly sensitive to the choice of benchmark (Tregle, Nix, and Alpert, 2019), and it is not clear which of these choices gives the correct answer, if any. It is likely that population benchmarks will continue to find racial bias, whereas crime and arrest benchmarks will find little to no bias.

While there is a relatively clear justificatory outcome for police stops—finding contraband, weapons, or both—there is no equivalent for shootings.

While such benchmarks as the racial composition of the local population or arrestees are relatively naïve, more-plausible, more-sophisticated benchmarking is possible, but it requires more-detailed information than what is typically contained in publicly available police shooting data. Although susceptible to post-treatment bias, Fryer (2018) was admirable in this regard. In related areas of policing research, quasi-experimental studies have used contextual features to estimate officer differences in use of force (e.g., Ba et al., 2021). Natural experiments (e.g., Hoekstra and Sloan, 2022; Legewie, 2016) are similarly implicit benchmark tests. Despite benchmarks' challenges, they range in plausibility. Richer data on the context of encounters between police and citizens and inventive natural experiments can improve benchmark test research.

Outcome Tests

Outcome tests involve looking for racial differences in the rate at which shootings are justified as a means of inferring racial bias. If shootings of White civilians are more frequently justifiable than shootings of Black civilians, this would suggest that the police have a lower threshold in the decision to shoot Black civilians. Most non-benchmark test research employs outcome tests. Unlike benchmark tests, outcome tests do not require measurement of the number of people at risk of being shot by police, thereby circumventing the denominator problem.

However, outcome tests face three primary methodological challenges. The first is what constitutes the *correct* outcome. While there is a relatively clear justificatory outcome for police stops—finding contraband, weapons, or both—there is no equivalent for shootings. The ideal outcome would be whether the civilian shot posed a threat that justified (or possibly, necessitated) the use of this level of force (or, better yet, a threat level rather than its binary counterpart). No such measures exist, so proxies are necessary.

A few outcomes have been used to date. The most common is whether the suspect was armed or unarmed. Fryer (2019) found no racial differences in the probability of finding a weapon on people shot by the police in 16 local jurisdictions. Conversely, Nix et al. (2017) and DeAngelis (2021) found racial differ-

ences in the probability of recovering weapons from victims of fatal police shootings nationally, consistent with antiminority bias. Importantly, carrying a weapon is not in and of itself a justification for lethal force, nor is it the only legal justification for the use of lethal force. More generally, no one proxy outcome always justifies lethal force, and there are multiple factors that could justify its use (Cahill, Labriola, and Taylor, 2021). Results need to be interpreted with this caveat in mind.

Nix et al. (2017) also reported racial differences in the probability of those shot dead having been attacking the police, finding that White individuals were more likely to have been attacking the police, whereas DeAngelis (2021) found that Black victims who were fatally shot were less likely than White victims to be carrying a firearm, less likely to exhibit signs of mental illness, and more likely to be fleeing police. DeAngelis interpreted this as officers having a lower threshold for shooting Black civilians, which suggests racial bias, but did not find Hispanic-White differences in these respects. Clark et al. (2023) examined racial differences in the probability of being fatally shot conditional on being shot as a way to test for racial bias in the threshold that police apply in deciding to shoot civilians. Using administrative data from eight local LEAs, Clark et al. found that fatality rates are higher for Black shooting victims than for White or Hispanic victims, consistent with evidence of racial bias against Black civilians. They estimate that at least 30 percent of Black civilians shot would not have been shot if they were White.

Researchers must be cognizant of outcome test interpretations of studies, even if they do not set out to use outcome tests. For example, Wheeler et al. (2017) and Worrall et al. (2018) studied racial differences in the probability of being shot by police conditional on the officers involved having drawn their guns. They found that Black civilians were less likely to be shot and interpreted this as there being no evidence of bias against them. However, these studies might be more suitably conceived of as outcome tests, revealing racial bias against Black civilians in the threshold for police to draw their guns. Worrall, Bishopp, and Terrill (2021) did not find a higher probability of weapons being drawn against Black suspects. However, given those authors' reliance on

use-of-force data alone, this finding could simply reflect racial bias in the use of force generally and is consistent with the possibility of racial bias against Black suspects in police drawing weapons on them.

The second challenge, which the aforementioned studies confront, is omitted variable bias (Neil and Winship, 2019). Factors that differ by race and that are correlated with the outcome could reflect reasons—other than racial bias—that outcome rates differ by race. For instance, if police have a lower threshold to shoot in the areas with the most gun violence, and White civilians are rarely encountered in such spots, then racial differences in outcome rates might reflect place-based differences in perceptions of danger rather than racial bias. Which factors might be confounders requires careful consideration because some control variables might be mechanisms through which bias operates as opposed to confounders. Omitted variable bias may be difficult to address because police shootings are, in a statistical sense, rare events, thereby limiting the availability of highly comparable incidents.

The third inferential challenge is called *infra-marginality* (Ayres, 2002; Neil and Winship, 2019; Simoiu, Corbett-Davies, and Goel, 2017). Outcome rates are compared across race to estimate whether there are racial differences in police officers' thresholds to shoot. However, outcome rates are simply an average, and they do not necessarily equal the threshold trying to be inferred. To see the difference, imagine that for one racial group there was an exception-

Factors that differ by race and that are correlated with the outcome could reflect reasons—other than racial bias—that outcome rates differ by race.

ally large proportion of people (compared with other racial groups) who posed a high risk to police and thus the use of lethal force against them was justified. This group of people would pull the outcome rate for this group upward. Giving a threshold interpretation to outcome rates in this situation would identify that racial group as having a high threshold (suggesting the police are unbiased against that group), although it could be consistent with the evidence for that group's threshold to, in fact, be the same or even lower than that of other racial groups, meaning that the police might be biased against that group. Generally, differences in the shape of the risk distribution by race mean that outcome rates do not necessarily reveal against which groups (if any) the police are biased. Conditioning on variables that drive these racial differentials in risk can mitigate this problem, but it is not a guaranteed fix. Importantly, if a group has a lower outcome rate than other groups, that group having more high-risk people will not change conclusions about the presence of bias. For instance, if an outcome rate is especially low for Black civilians, pointing to bias against them, this conclusion will stand even if there is an especially risky subset of Black offenders.

Where Do We Go from Here?

Research on racial bias in police shootings faces challenges in both data quality and modeling. Despite these obstacles, policy researchers cannot resign themselves to pleas of ignorance. How, then, should researchers proceed? In this section, we review four ways in which research on racial bias in police shootings can be advanced, drawing on recent examples when available.

Police shootings are as much organizational failures as individual ones.

Better Data

The lack of comprehensive, detailed data on fatal and non-fatal shootings is perhaps the biggest impediment. Crowdsourced databases may now provide a reasonably complete picture of the contemporary prevalence of fatal police shootings, but they contain limited information on these incidents' circumstances, and the need to create them in the first place raises questions about their coverage further back in time. There is a critical lack of reliable national data on non-fatal shootings. Ultimately, these shortcomings are likely to be solved only with federally mandated reporting at the level of detail that the NUFDC was intended to collect.

More-comprehensive data on police shootings will be particularly important for understanding heterogeneity across LEAs. Rates of fatal police violence and racial disparities relative to the local population show significant heterogeneity along regional, socioeconomic, and urban-rural dimensions (Edwards, Esposito, and Lee, 2018; Medina et al., 2022; Ross, 2015). Other racial disparities in policing likewise vary significantly across jurisdictions (Anwar et al., 2021). There is some evidence that organizational characteristics, such as department size and training requirements, are correlated with misconduct (Eitle, D'Alessio, and Stolzenberg, 2014), but the reliability of misconduct data on a national scale is even more suspect than that of the data on police shootings. A police shooting is the outcome of complex chains of decisions made under limited information in limited time. Police shootings are as much organizational failures as individual ones (Sherman, 2018). Going forward, research on race and police shootings will be improved by asking *when and where* there is racial bias rather than *whether* there is such bias.

Clearer Definitions and Estimands

Defining what constitutes racial bias is important. Researchers should consider taking this a step further and specifying what exactly they are trying to estimate—their *estimand*—with their models. Doing so offers several benefits, such as clearly delineating the population of interest (e.g., the whole population versus those who are encountered) and

making it clear what assumptions are necessary to identify racial bias with the model used, two issues that underlie many of the points above. Lundberg, Johnson, and Stewart (2021) made this point generally, while Knox and Mummolo (2020) did so for the study of racial bias in policing. In an applied example, Clark et al. (2023) developed a theoretical model of selection into encounters and that model's empirical implications for shootings. Specifically, the authors posited a behavioral model of police-civilian interactions both leading up to and during encounters. The justification for fatality as an outcome test rests on several assumptions about how civilians and police anticipate and react to one another's behavior if racial bias exists. For instance, the argument presupposes that escalating aggression by a civilian will be met by escalating severity of force imposed by an officer and, consequently, more-aggressive civilians who are shot are more likely to die. If this is true, then racial differences in police shooting fatality rates would imply that officers have different thresholds for shooting civilians of each race. While it is possible to take issue with this behavioral model, there is much merit in the clear specification of what is trying to be estimated and why the method used is expected to estimate that and under what assumptions.

Better Methods

Given the severity of the aforementioned methodological challenges, more-extensive data alone would not yield strong answers about the degree of racial bias in police shootings; innovations to methods are needed as well. Recent work has seen innovations to outcome tests, such as in Clark et al. (2023), and in benchmark tests, as in the bias correction developed by Ross et al. (2021). Further refinements of these tests may be possible, although researchers should consider whether new methods are possible. Analogously, research on other forms of racial bias in policing was advanced when researchers confronted types of methodological challenges and developed novel research designs, such as veil of darkness tests (Grogger and Ridgeway, 2006) and computational methods (Goel, Rao, and Shroff, 2016; Pierson et al., 2020; Simoiu, Corbett-Davies, and Goel, 2017).

More-extensive data alone would not yield strong answers about the degree of racial bias in police shootings; innovations to methods are needed as well.

Bounding and Sensitivity Analysis

Given concerns about data and the assumptions required for any outcome or benchmark test, research could be improved with greater transparency about the sensitivity of the authors' conclusions and more recognition of uncertainty. Clark et al. (2023) and Knox, Lowe, and Mummolo (2020) employed two closely related practices that may be valuable for researchers because of the challenges of studying racial bias in police shootings. The first practice is to estimate a range of estimates for racial bias (*bounds*) that one would find under different violations of a model's assumptions, a practice known as *partial identification* (Imbens and Rubin, 2015; Tamer, 2010). Both Clark et al. (2023) and Knox, Lowe, and Mummolo (2020) constructed bounds on how much racial bias one would find under different levels of post-treatment bias. Concretely, they examined the sensitivity of estimates of racial bias in the use of force to varying degrees of racial bias in the encounter process. Clark et al. (2023) supplemented their bounding strategy with a *sensitivity analysis* to explore the possibility that unobserved factors related to race may also affect the study's outcome (fatality). Their sensitivity analysis examined how strongly an unobserved confounding factor (e.g., victim's age or pre-shooting physical health) would have to be associated with both race and fatality to eliminate evidence of race's effects on fatality. The analysis indicates that evidence of racial bias would exist even if the potential

Strong data on shootings are likely inadequate on their own because of methodological challenges in inferring racial bias, and because of the fact that racial bias is a process that requires a consideration of more than just the moment a shooting occurs.

confounder influenced fatality three times as much as any variable in their data set. Similar to bounds on racial bias, sensitivity analysis is a transparent way of conveying uncertainty about one's conclusions and a study's susceptibility to some of the research challenges that we have discussed.

Summary and Conclusions

The extent to which police shootings are driven by racial bias among police officers in the United States is a difficult question to answer. There is no comprehensive and detailed data source with which

to answer it, especially when it comes to non-fatal shootings. Furthermore, strong data on shootings are likely inadequate on their own because of methodological challenges in inferring racial bias, and because of the fact that racial bias is a process that requires a consideration of more than just the moment a shooting occurs. However, it is too important a topic to abandon. Not only is racial bias in police violence a major contemporary issue, it is also one that exists within—and stems from—a wider context of racial inequality, the legacy of racism, and problematic police-minority community relations in the United States.

We have surveyed four broad ways through which research on racial bias in police shootings can be improved: better data, better methods, clearer definitions and estimands, and the use of bounding and sensitivity analyses. To reiterate, such improvements should be used not only to study racial bias in shootings overall, but also to understand heterogeneity—the times, places, and circumstances in which racial bias in shootings is especially common or uncommon. Finally, as important as these questions are, they are not the only policy-relevant questions in addressing the shooting and killing of minority civilians by police in the United States (Cahill, Labriola, and Taylor, 2021). Rather, these questions are part of a broader agenda that ought to be pursued, one that studies the impacts of policy interventions, examines how to reduce unnecessary police shootings, and confronts implementation challenges in translating insights on the nature of bias and the causes of police shootings into actionable and scalable interventions that can make a real impact.

Notes

¹ In other work, RAND researchers have detailed more-general challenges in the study of police violence beyond racial bias (Cahill, Labriola, and Taylor, 2021).

² While our definition of racial bias in this essay is limited in scope, there is an expansive literature in this space. For a more complete review of racial bias and how it operates in daily life, see Richeson and Sommers (2016) and Small and Pager (2020).

³ This report focuses on studies that use observational data, which speak more to the existence of real-world racial bias than to its underlying mechanisms in individual decisionmaking (Spencer, Charbonneau, and Glaser, 2016). A small body of experimental research in laboratory settings has attempted to assess officers' racial bias in the decision to shoot. Such studies generally use variations on implicit association tests (and have recently used virtual reality simulators) to evaluate officers' decisions to shoot and reaction times when faced with White or minority suspects (Correll et al., 2002; Correll et al., 2014; Cox et al., 2014; Eberhardt et al., 2004; James, 2018; James, James, and Vila, 2016; Plant and Peruche, 2005). Social psychologists consider racial categorization and the activation of racial prejudice to be context-dependent (Richeson and Sommers, 2016), which may affect the external validity of findings from laboratory studies. For example, James (2018) found that responses on implicit association tests are sensitive to officers' sleep deprivation, and Correll et al. (2014) reported more evidence of racial bias in virtual reality simulations with officers who work in jurisdictions that are more urban, have higher violent crime rates, and have more-concentrated minority populations. Consequently, the external validity of laboratory findings on police bias and their sensitivity to design choices and sample selection remain open to debate (Roussell et al., 2019; Terrill, 2016). Going forward, researchers should consider the consistency of experimental and observational studies and their implications for understanding racial disparities.

⁴ A lower threshold for suspicion of racial minorities could be shaped by preexisting stereotypes about race and criminality. Research has shown that White people are more likely to believe that minorities—particularly African American individuals—exhibit characteristics associated with criminality (Alexander, 2012). In addition, perceptions may vary by skin tone: Individuals with darker skin within minority populations are more likely to be stereotyped negatively than lighter-skinned individuals (Monk, 2019).

⁵ Fryer's work attracted public attention prior to peer review, including coverage in *The New York Times* (Bui and Cox, 2016; Carey and Goode, 2016).

⁶ Data on calls for service can include calls that the shooting officer initiated or that another officer initiated. It is important to be attentive to what exactly is being measured in such datasets because it is not necessarily 911 calls from civilians exclusively.

⁷ In 2003, the federal Bureau of Justice Statistics (BJS) began the Arrest-Related Deaths (ARD) program to obtain information on deaths that occur in the custody of state and local law enforcement agencies. BJS suspended the ARD program in 2014 because of concerns over its coverage and reliability. BJS has conducted pilot assessments for a revised ARD program (Banks et al., 2019), but, as of August 2023, the ARD program remains dormant.

⁸ The Police Data Initiative (undated) maintains a repository of open data on police use of force that are volunteered by LEAs, including those in many small and mid-sized cities, which are understudied sites. These data are not comprehensive, and LEAs differ in the information they report.

⁹ This is derived from the authors' calculations. The new states include California, Florida, and Texas, which were the three states with the most fatal police shootings.

¹⁰ Historically, LEAs not reporting to NIBRS were from larger cities and jurisdictions with higher crime rates than participating LEAs (McCormack, Pattavina, and Tracy, 2017).

¹¹ Cahill, Labriola, and Taylor (2021) reviewed findings on specific interventions to reduce fatal police violence.

¹² The most easily accessible national-level source is NEISS, which collects injury data annually from a nationally representative sample of approximately 100 hospitals with emergency departments and includes injuries' causes. However, NEISS estimates of firearm injuries in general are unreliable and imprecise (Smart et al., 2021a). One should expect these issues to be even more severe when using NEISS to estimate a relatively rare event like police shootings. See Smart et al. (2021a) for a review of issues with other national data sources on non-fatal shootings.

¹³ Attentional biases in crowdsourced data are likely even more significant for non-fatal shootings than fatal shootings. Nonetheless, Ward et al. (forthcoming) have painted the most comprehensive national picture of non-fatal police shootings to date, in terms of both demographics and incident characteristics.

¹⁴ While our focus is on police shootings generally, most research to date has focused specifically on bias in fatal shootings.

¹⁵ There are more-complex versions of benchmark tests, such as using regression models (e.g., Ross, 2015), but the basic intuition and challenges hold.

¹⁶ Data on racial differences in violent crime rates reflect historical and contemporary inequalities in the United States. One such historical inequality is racial bias in how ideas about criminality and violent behavior are constructed and structural inequalities that contribute to disproportionate rates of violent crime among some populations and in some geographic areas (Sampson and Neil, 2023). For example, research has shown that African American men are particularly likely to be stereotyped as violent and criminal (Alexander, 2012; Dow, 2016; Smiley and Fakunle, 2016). This can contribute to racial bias in violent crime reporting, for instance, if it affects civilian reporting or police practices.

References

- Abernathy, Penelope Muse, *The Expanding News Desert*, University of North Carolina Press, 2018.
- Alexander, Michelle, *The New Jim Crow: Mass Incarceration in the Age of Colorblindness*, New Press, 2012.
- Anderson, Elijah, *Code of the Street: Decency, Violence, and the Moral Life of the Inner City*, W. W. Norton & Company, 1999.
- Anwar, Shamena, Patrick Bayer, Randi Hjalmarsson, and Matthew L. Mizel, *Racial Disparities in Misdemeanor Speeding Convictions*, RAND Corporation, RR-A1317-1, 2021. As of January 29, 2024:
https://www.rand.org/pubs/research_reports/RRA1317-1.html
- Ayres, Ian, "Outcome Tests of Racial Disparities in Police Practices," *Justice Research and Policy*, Vol. 4, Nos. 1–2, Fall 2002.
- Ba, Bocar A., Dean Knox, Jonathan Mummolo, and Roman Rivera, "The Role of Officer Race and Gender in Police-Civilian Interactions in Chicago," *Science*, Vol. 371, No. 6530, 2021.
- Banks, Duren, Michael G. Planty, Lance Couzens, Philip Lee, Connor Brooks, Kevin M. Scott, and Anthony Whyde, *Arrest-Related Deaths Program: Pilot Study of Redesigned Survey Methodology*, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics, July 2019.
- Barber, Catherine, Deborah Azrael, Amy Cohen, Matthew Miller, Deonza Thymes, David Enze Wang, and David Hemenway, "Homicides by Police: Comparing Counts from the National Violent Death Reporting System, Vital Statistics, and Supplementary Homicide Reports," *American Journal of Public Health*, Vol. 106, No. 5, May 2016.
- Bourgois, Philippe, *In Search of Respect: Selling Crack in El Barrio*, 2nd ed., Cambridge University Press, 2003.
- Bui, Quoc Trung, and Amanda Cox, "Surprising New Evidence Shows Bias in Police Use of Force but Not in Shootings," *New York Times*, July 11, 2016.
- Cahill, Meagan, Melissa M. Labriola, and Jirka Taylor, *Police Killings: Road Map of Research Priorities for Change*, RAND Corporation, RR-A1525-1, 2021. As of January 29, 2024:
https://www.rand.org/pubs/research_reports/RRA1525-1.html
- Carey, Benedict, and Erica Goode, "Police Try to Lower Racial Bias, but Under Pressure, It Isn't So Easy," *New York Times*, July 11, 2016.
- Cesario, Joseph, David J. Johnson, and William Terrill, "Is There Evidence of Racial Disparity in Police Use of Deadly Force? Analyses of Officer-Involved Fatal Shootings in 2015–2016," *Social Psychological and Personality Science*, Vol. 10, No. 5, 2019.
- Charles, Camille Zubrinsky, "The Dynamics of Racial Residential Segregation," *Annual Review of Sociology*, Vol. 29, No. 1, 2003.
- Clark, Tom S., Elisha Cohen, Adam N. Glynn, Michael Leo Owens, Anna Gunderson, and Kaylyn Jackson Schiff, "Are Police Racially Biased in the Decision to Shoot?" *Journal of Politics*, Vol. 85, No. 3, July 2023.
- Conner, Andrew, Deborah Azrael, Vivian H. Lyons, Catherine Barber, and Matthew Miller, "Validating the National Violent Death Reporting System as a Source of Data on Fatal Shootings of Civilians by Law Enforcement Officers," *American Journal of Public Health*, Vol. 109, No. 4, 2019.
- Correll, Joshua, Sean M. Hudson, Steffanie Guillermo, and Debbie S. Ma, "The Police Officer's Dilemma: A Decade of Research on Racial Bias in the Decision to Shoot," *Social and Personality Psychology Compass*, Vol. 8, No. 5, 2014.
- Correll, Joshua, Bernadette Park, Charles M. Judd, and Bernd Wittenbrink, "The Police Officer's Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals," *Journal of Personality and Social Psychology*, Vol. 83, No. 6, 2002.
- Cox, William T. L., Patricia G. Devine, E. Ashby Plant, and Lauri L. Schwartz, "Toward a Comprehensive Understanding of Officers' Shooting Decisions: No Simple Answers to This Complex Problem," *Basic and Applied Social Psychology*, Vol. 36, No. 4, 2014.
- DeAngelis, Reed T., "Systemic Racism in Police Killings: New Evidence from the Mapping Police Violence Database, 2013–2021," *Race and Justice*, 2021.
- Dow, Dawn Marie, "The Deadly Challenges of Raising African American Boys: Navigating the Controlling Image of the 'Thug,'" *Gender and Society*, Vol. 20, No. 2, 2016.
- Durlauf, Steven N., and James J. Heckman, "An Empirical Analysis of Racial Differences in Police Use of Force: A Comment," *Journal of Political Economy*, Vol. 128, No. 10, October 2020.
- Eberhardt, Jennifer L., Valerie J. Purdie, Phillip Atiba Goff, and Paul G. Davies, "Seeing Black: Race, Crime, and Visual Processing," *Journal of Personality and Social Psychology*, Vol. 87, No. 6, 2004.
- Edwards, Frank, Michael H. Esposito, and Hedwig Lee, "Risk of Police-Involved Death by Race/Ethnicity and Place, United States, 2012–2018," *American Journal of Public Health*, Vol. 108, No. 9, September 2018.
- Eitle, David, Stewart J. D'Alessio, and Lisa Stolzenberg, "The Effect of Organizational and Environmental Factors on Police Misconduct," *Police Quarterly*, Vol. 17, No. 2, 2014.
- Epp, Charles R., Steven Maynard-Moody, and Donald Haider-Markel, *Pulled Over: How Police Stops Define Race and Citizenship*, University of Chicago Press, 2014.
- FBI—See Federal Bureau of Investigation.
- Federal Bureau of Investigation, "FBI Releases 2019 Participation Data for the National Use-of-Force Data Collection," press release, July 27, 2020e.
- Federal Bureau of Investigation, "Crime Data Explorer," homepage, 2022d. As of November 17, 2023:
<https://cde.ucr.cjis.gov/>
- Fryer, Roland G., Jr., "Reconciling Results on Racial Differences in Police Shootings," *AEA Papers and Proceedings*, Vol. 108, May 2018.
- Fryer, Roland G., Jr., "An Empirical Analysis of Racial Differences in Police Use of Force," *Journal of Political Economy*, Vol. 127, No. 3, June 2019.
- Gelman, Andrew, Jeffrey Fagan, and Alex Kiss, "An Analysis of the New York City Police Department's 'Stop-and-Frisk' Policy in the Context of Claims of Racial Bias," *Journal of the American Statistical Association*, Vol. 102, No. 479, 2007.
- Goel, Sharad, Justin M. Rao, and Ravi Shroff, "Precinct or Prejudice? Understanding Racial Disparities in New York City's Stop-and-Frisk Policy," *Annals of Applied Statistics*, Vol. 10, No. 1, March 2016.

- Grogger, Jeffrey, and Greg Ridgeway, "Testing for Racial Profiling in Traffic Stops from Behind a Veil of Darkness," *Journal of the American Statistical Association*, Vol. 101, No. 475, September 2006.
- Hoekstra, Mark, and Carly Will Sloan, "Does Race Matter for Police Use of Force? Evidence from 911 Calls," *American Economic Review*, Vol. 112, No. 3, March 2022.
- Imbens, Guido W., and Donald B. Rubin, *Causal Inference in Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press, 2015.
- Jackman, Tom, "FBI May Shut Down Police Use-of-Force Database Due to Lack of Police Participation," *Washington Post*, December 9, 2021.
- James, Lois, "The Stability of Implicit Racial Bias in Police Officers," *Police Quarterly*, Vol. 21, No. 1, March 2018.
- James, Lois, Stephen M. James, and Bryan J. Vila, "The Reverse Racism Effect: Are Cops More Hesitant to Shoot Black Than White Suspects?" *Criminology and Public Policy*, Vol. 15, No. 2, May 2016.
- Knox, Dean, Will Lowe, and Jonathan Mummolo, "Administrative Records Mask Racially Biased Policing," *American Political Science Review*, Vol. 114, No. 3, 2020.
- Knox, Dean, and Jonathan Mummolo, "Toward a General Causal Framework for the Study of Racial Bias in Policing," *Journal of Political Institutions and Political Economy*, Vol. 1, No. 3, 2020.
- Kohler-Hausmann, Issa, "Eddie Murphy and the Dangers of Counterfactual Causal Thinking About Detecting Racial Discrimination," *Northwestern University Law Review*, Vol. 113, No. 5, 2019.
- Legewie, Joscha, "Racial Profiling and Use of Force in Police Stops: How Local Events Trigger Periods of Increased Discrimination," *American Journal of Sociology*, Vol. 122, No. 2, September 2016.
- Lundberg, Ian, Rebecca Johnson, and Brandon M. Stewart, "What Is Your Estimand? Defining the Target Quantity Connects Statistical Evidence to Theory," *American Sociological Review*, Vol. 86, No. 3, 2021.
- Mapping Police Violence, homepage, last updated January 14, 2024. As of February 1, 2024: <https://mappingpoliceviolence.org/>
- Massey, Douglas S., and Nancy A. Denton, *American Apartheid: Segregation and the Making of the Underclass*, Harvard University Press, 1993.
- McCormack, Phillip D., April Pattavina, and Paul E. Tracy, "Assessing the Coverage and Representativeness of the National Incident-Based Reporting System," *Crime and Delinquency*, Vol. 63, No. 4, 2017.
- Medina, Richard M., Simon Brewer, Andrew M. Linke, Emily A. Nicolosi, Marco Allain, and Douglas Tharp, "The Sociospatial Ecology of Deaths During Police Interactions in the United States, 2016–2020," *Political Geography*, Vol. 98, 2022.
- Mentch, Lucas, "On Racial Disparities in Recent Fatal Police Shootings," *Statistics and Public Policy*, Vol. 7, No. 1, 2020.
- Monk, Ellis P., "The Color of Punishment: African Americans, Skin Tone, and the Criminal Justice System," *Ethnic and Racial Studies*, Vol. 42, No. 10, 2019.
- Mullainathan, Sendhil, "Police Killings of Blacks: Here Is What the Data Say," *New York Times*, October 16, 2015.
- Neil, Roland, and Christopher Winship, "Methodological Challenges and Opportunities in Testing for Racial Discrimination in Policing," *Annual Review of Criminology*, Vol. 2, 2019.
- Nix, Justin, "On the Challenges Associated with the Study of Police Use of Deadly Force in the United States: A Response to Schwartz and Jahn," *PLoS One*, Vol. 15, No. 7, 2020.
- Nix, Justin, Bradley A. Campbell, Edward H. Byers, and Geoffrey P. Alpert, "A Bird's Eye View of Civilians Killed by Police in 2015: Further Evidence of Implicit Bias," *Criminology and Public Policy*, Vol. 16, No. 1, February 2017.
- Nix, Justin, and John A. Shjarback, "Factors Associated with Police Shooting Mortality: A Focus on Race and a Plea for More Comprehensive Data," *PLoS One*, Vol. 16, No. 11, 2021.
- Pierson, Emma, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Daniel Jenson, Amy Shoemaker, Vignesh Ramachandran, Phoebe Barghouty, Cheryl Phillips, Ravi Shroff, and Sharad Goel, "A Large-Scale Analysis of Racial Disparities in Police Stops Across the United States," *Nature Human Behaviour*, Vol. 4, No. 7, July 2020.
- Plant, E. Ashby, and B. Michelle Peruche, "The Consequences of Race for Police Officers," *Psychological Science*, Vol. 16, No. 3, March 2005.
- Police Data Initiative, homepage, undated. As of March 27, 2024: <https://www.policedatainitiative.org/>
- Richeson, Jennifer A., and Samuel R. Sommers, "Toward a Social Psychology of Race and Race Relations for the Twenty-First Century," *Annual Review of Psychology*, Vol. 67, 2016.
- Ridgeway, Greg, *Analysis of Racial Disparities in the New York Police Department's Stop, Question, and Frisk Practices*, RAND Corporation, TR-534-NYCPF, 2007. As of September 5, 2023: https://www.rand.org/pubs/technical_reports/TR534.html
- Ridgeway, Greg, and John MacDonald, "Methods for Assessing Racially Biased Policing," in Stephen K. Rice and Michael D. White, eds., *Race, Ethnicity, and Policing: New and Essential Readings*, New York University Press, 2010.
- Rios, Victor M., *Punished: Policing the Lives of Black and Latino Boys*, New York University Press, 2011.
- Ross, Cody T., "A Multi-Level Bayesian Analysis of Racial Bias in Police Shootings at the County-Level in the United States, 2011–2014," *PLoS One*, Vol. 10, No. 11, 2015.
- Ross, Cody T., Bruce Winterhalder, and Richard McElreath, "Racial Disparities in Police Use of Deadly Force Against Unarmed Individuals Persist After Appropriately Benchmarking Shooting Data on Violent Crime Rates," *Social Psychological and Personality Science*, Vol. 12, No. 3, 2021.
- Roussell, Aaron, Kathryn Henne, Karen S. Glover, and Dale Willits, "Impossibility of a 'Reverse Racism' Effect: A Rejoinder to James, James, and Vila," *Criminology and Public Policy*, Vol. 18, No. 1, February 2019.
- Sampson, Robert J., and Roland Neil, "The Social Foundations of Racial Inequalities in Arrest over the Life Course and in Changing Times," draft ahead of publication, 2023.
- Schroedel, Jean Reith, and Roger J. Chin, "Whose Lives Matter: The Media's Failure to Cover Police Use of Lethal Force Against Native Americans," *Race and Justice*, Vol. 10, No. 2, April 2020.

- Schwartz, Gabriel L., and Jaquelyn L. Jahn, "Mapping Fatal Police Violence Across U.S. Metropolitan Areas: Overall Rates and Racial/Ethnic Inequities, 2013–2017," *PLoS One*, Vol. 15, No. 6, 2020.
- Sharara, Fablina, Eve E. Wool, Gregory J. Bertolacci, Nicole Davis Weaver, Shelly Balassyano, Alexandra Warson, Ilse N. Dippenaar, Matthew Cunningham, John E. Fuller, Laure B. Marczak, et al., "Fatal Police Violence by Race and State in the USA, 1980–2019: A Network Meta-Regression," *The Lancet*, Vol. 398, No. 10307, October 2, 2021.
- Sharkey, Patrick, *Stuck in Place: Urban Neighborhoods and the End of Progress Toward Racial Equality*, University of Chicago Press, 2013.
- Sherman, Lawrence W., "Reducing Fatal Police Shootings as System Crashes: Research, Theory, and Practice," *Annual Review of Criminology*, Vol. 1, January 2018.
- Simoiu, Camelia, Sam Corbett-Davies, and Sharad Goel, "The Problem of Infra-Marginality in Outcome Tests for Discrimination," *Annals of Applied Statistics*, Vol. 11, No. 3, 2017.
- Small, Mario L., and Devah Pager, "Sociological Perspectives on Racial Discrimination," *Journal of Economic Perspectives*, Vol. 34, No. 2, Spring 2020.
- Smart, Rosanna, Samuel Peterson, Terry L. Schell, Rose Kerber, and Andrew R. Morral, *Inpatient Hospitalizations for Firearm Injury: Estimating State-Level Rates from 2000 to 2016*, RAND Corporation, TL-A243-3, 2021a. As of January 31, 2024: <https://www.rand.org/pubs/tools/TLA243-3.html>
- Smiley, Calvin John, and David Fakunle, "From 'Brute' to 'Thug': The Demonization and Criminalization of Unarmed Black Male Victims in America," *Journal of Human Behavior in the Social Environment*, Vol. 26, Nos. 3–4, 2016.
- Spencer, Katherine B., Amanda K. Charbonneau, and Jack Glaser, "Implicit Bias and Policing," *Social and Personality Psychology Compass*, Vol. 10, No. 1, 2016.
- Tamer, Elie, "Partial Identification in Econometrics," *Annual Review of Economics*, Vol. 2, No. 1, 2010.
- Terrill, William, "Deadly Force: To Shoot or Not to Shoot," *Criminology and Public Policy*, Vol. 15, No. 2, May 2016.
- Tregle, Brandon, Justin Nix, and Geoffrey P. Alpert, "Disparity Does Not Mean Bias: Making Sense of Observed Racial Disparities in Fatal Officer-Involved Shootings with Multiple Benchmarks," *Journal of Crime and Justice*, Vol. 42, No. 1, 2019.
- Ward, Julia A., Javier Cepeda, Dylan B. Jackson, Odis Johnson, Jr., Daniel W. Webster, and Cassandra K. Crifasi, "National Burden of Injury and Deaths from Shootings by Police in the United States, 2015–2020," *American Journal of Public Health*, forthcoming.
- Wheeler, Andrew P., Scott W. Phillips, John L. Worrall, and Stephen A. Bishopp, "What Factors Influence an Officer's Decision to Shoot? The Promise and Limitations of Using Public Data," *Justice Research and Policy*, Vol. 18, No. 1, 2017.
- Wilson, William Julius, *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy*, 1st ed., University of Chicago Press, 1987.
- Wilson, William Julius, *When Work Disappears: The World of the New Urban Poor*, Alfred A. Knopf, 1996.
- Worrall, John L., Stephen A. Bishop, and William Terrill, "The Effect of Suspect Race on Police Officers' Decisions to Draw Their Weapons," *Justice Quarterly*, Vol. 38, No. 7, 2021.
- Worrall, John L., Stephen A. Bishopp, Scott C. Zinser, Andrew P. Wheeler, and Scott W. Phillips, "Exploring Bias in Police Shooting Decisions with Real Shoot/Don't Shoot Cases," *Crime and Delinquency*, Vol. 64, No. 9, 2018.
- Zimring, Franklin E., *When Police Kill*, Harvard University Press, 2017.
- Zuckerman, Ethan, J. Nathan Matias, Rahul Bhargava, Fernando Bermejo, and Allan Ko, "Whose Death Matters? A Quantitative Analysis of Media Attention to Deaths of Black Americans in Police Confrontations, 2013–2016," *International Journal of Communication*, Vol. 13, 2019.

Abbreviations

ARD	Arrest-Related Deaths
BJS	Bureau of Justice Statistics
CDC	Centers for Disease Control and Prevention
FBI	Federal Bureau of Investigation
LEA	Law Enforcement Agency
NEISS	National Electronic Injury Surveillance System
NIBRS	National Incident-Based Reporting System
NUFDC	National Use-of-Force Data Collection
NVDRS	National Violent Death Reporting System
NYPD	New York City Police Department
SHR	Supplementary Homicide Reports



RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors. **RAND**[®] is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

For more information on this publication, visit www.rand.org/t/RR-A243-8.

© 2024 RAND Corporation

www.rand.org

About This Report

In this short report, we explore challenges in measuring racial bias in police shootings of civilians. We discuss the importance of thinking about bias as a process to understand *how* racial bias operates in police shootings. We also discuss existing research in this area and data limitations, and provide recommendations for future research.

Justice Policy Program

RAND Social and Economic Well-Being is a division of the RAND Corporation that seeks to actively improve the health and social and economic well-being of populations and communities throughout the world. This research was conducted in the Justice Policy Program within RAND Social and Economic Well-Being. The program focuses on such topics as access to justice, policing, corrections, drug policy, and court system reform, as well as other policy concerns pertaining to public safety and criminal and civil justice. For more information, email justicepolicy@rand.org.

Funding

Funding for the Gun Policy in America initiative was originally provided through unrestricted gifts from RAND supporters and income from operations. Since June 2018, this initiative has been supported by a grant from Arnold Ventures.

Acknowledgments

We thank James Anderson, Amanda Charbonneau, Andrew Morral, and Justin Nix for their feedback on earlier versions of this report.