

Are Police Racially Biased in the Decision to Shoot?

Tom S. Clark, Emory University

Elisha Cohen, New York University

Adam N. Glynn, Emory University

Michael Leo Owens, Emory University

Anna Gunderson, Louisiana State University

Kaylyn Jackson Schiff, Yale University and Purdue University

We present a theoretical model predicting that racially biased policing produces (1) more use of potentially lethal force by firearms against Black civilians than against White civilians and (2) lower fatality rates for Black civilians than White civilians. We empirically evaluate this second prediction with original officer-involved shooting data from 2010 to 2017 for eight local police jurisdictions, finding that Black fatality rates are significantly lower than White fatality rates and that this significance would survive an omitted covariate three times as strong as any of our observed covariates. Furthermore, using outcome test methodology and a comparability assumption, we estimate that at least 30% of Black civilians shot by the police would not have been shot had they been White. An omitted covariate would need to be at least three times as strong as any of our observed covariates to eliminate this finding. Finally, any omitted covariate would have to affect Black fatality rates substantially more than Hispanic fatality rates in order to be consistent with the data.

Police are agents of the state, exercising a high degree of autonomy and discretion when implementing policy (Brown 1981; Wilson 1978). But, unlike other domestic agents of the state, “the police are . . . a mechanism for the distribution of situationally justified force in society” (Bittner 1970, 39). Consequently, the character of their interactions with the public differs greatly from those of other “street-level bureaucrats” (Lipsky 1980): police-civilian encounters are more unpredictable, with greater potential for violence and death, for civilians and police. Accordingly, policing is “profoundly involved with the most significant questions facing any political order, those pertaining to justice, order, and equity” (Brown 1981, 6–7). It is especially true when police use their discretion to shoot civilians.

While police use force against civilians more in some nations than others, police shootings of civilians are more common in the United States relative to other advanced, liberal

democracies (Zimring 2017). Furthermore, racial disparities in police use of force in the United States seem common and particularly wide between Blacks and Whites, and there are racial disparities in policing, generally, including in deployment, surveillance, involuntary contact by stop and frisk, arrest, and jailing (Bittner 1970; Brown 2019; Soss and Weaver 2017). Given the fraught history and contemporary realities of race in the United States, racial disparities in police shootings raise concerns about racial bias influencing officers’ discretion to shoot during police-civilian encounters. Whether racial bias causes racial disparities in policing, and how much, however, remains an academic and civic puzzle.

It is empirically difficult to discern how many police shootings of Black Americans result from their disproportionate contact with police versus disproportionate use of force by police against them versus racial bias by patrol officers and their departments (e.g., Fryer 2016; Knowles, Persico, and

Tom S. Clark (tclark7@emory.edu) is Charles Howard Candler Professor at Emory University, Atlanta, GA. Elisha Cohen (ec1302@nyu.edu) is Center for Data Science faculty fellow at New York University, New York, NY. Adam N. Glynn (adam.glynn@emory.edu) is professor at Emory University, Atlanta, GA. Michael Leo Owens (mowens4@emory.edu) is professor at Emory University, Atlanta, GA. Anna Gunderson (agunderson@lsu.edu) is assistant professor at Louisiana State University, Baton Rouge, LA. Kaylyn Jackson Schiff (schiffk@purdue.edu) is postdoctoral researcher at Yale University, New Haven, CT, and assistant professor at Purdue University, West Lafayette, IN.

Replication files are available in the *JOP* Dataverse (<https://dataverse.harvard.edu/dataverse/jop>). The empirical analysis has been successfully replicated by the *JOP* replication analyst. An appendix with supplementary material is available at <https://doi.org/10.1086/723973>.

Published online May 19, 2023.

The Journal of Politics, volume 85, number 3, July 2023. © 2023 Southern Political Science Association. All rights reserved. Published by The University of Chicago Press for the Southern Political Science Association. <https://doi.org/10.1086/723973>

Todd 2001; Knox, Lowe, and Mummolo 2020). Further, neither police departments nor agencies overseeing them track or report all lethal and nonlethal police shootings of civilians, especially by race (Zimring 2017). Consequently, depending on their data, measures, and methods, studies draw contradictory conclusions, ranging from significant differences in the likelihood and speed of shooting Black civilians compared to other civilians (Mekawi and Bresin 2015) to no racial differences in fatal shootings of civilians by police (Johnson et al. 2019). Therefore, even when relatively good data are available for social scientists to observe and describe racial patterns in policing, scholarly consensus on whether and how much police discriminate by the race of civilians when using lethal force, let alone nonlethal force, remains elusive.

To better assess whether there is evidence of racial bias in the use of force by police against civilians, measured by shootings, lethal and nonlethal, we develop a model of police-civilian encounters that yields empirical implications for evaluating racial bias in officer-involved shootings (OIS). In our model, informed by studies of the transactional nature and iterative process of police-civilian encounters (Binder and Scharf 1980; Kahn et al. 2017; Terrill 2005), civilians and police engage in behaviors, covering actions that may and can escalate their encounters toward harm, including police violence against civilians (and civilian violence against police). Ultimately, our model predicts that racially biased police officers will be more likely to use force against Black civilians than against White civilians. Moreover, police shootings of Black civilians should result in more nonfatalities than fatalities.

We test the implication of our model with OIS data from eight local police jurisdictions in the United States. Our data, covering 2010 through 2017, and obtained through public records requests, include all instances of police reporting they shot civilians—fatally and nonfatally—and the race of civilians, along with other attributes of the police-civilian encounters. Consistent with our theoretical expectation, we find that Black civilians are significantly more likely to survive an OIS, reflecting, we posit, a higher degree of racial bias in the decisions by officers to shoot Black civilians compared to non-Black civilians. Furthermore, we show that both the estimate and significance of this disparity would survive an omitted covariate three times as strong as any of our observed covariates.¹

Additionally, we estimate a lower bound on the magnitude of racial bias in the decision to shoot a civilian, guided by Cohen (2021) and Knox et al. (2020). Borrowing their techniques, we conceptually divide Black civilians who were shot by police into two groups—(1) Black civilians who would have

been shot had they been White and (2) Black civilians who would not have been shot had they been White. The proportional size of the second group is our parameter of racial bias. To estimate a lower bound for this quantity, we evaluate the difference in fatality rates of White and Black civilians shot by the police in the eight local jurisdictions relative to their White fatality rates, where we posit fatal shootings are more likely to be justified as “reasonable” shootings from the perspective of police departments and that nonfatal shootings are more prevalent among Black civilians compared to other groups. Using the techniques from Cohen (2021) and VanderWeele and Ding (2017), we estimate that at least 30% of Black civilians shot would not have been shot had they been White and that to eliminate this estimate, an omitted covariate would again need to be three times a strong as any of our observed covariates.² Finally, such an omitted covariate would have to affect Black fatality rates and not Hispanic fatality rates in order to be consistent with the data.³

Our theory and findings provide novel evidence of racial bias in police decision-making, buttressing other research (Knowles et al. 2001; Knox and Mummolo 2020; Knox et al. 2020; Persico and Todd 2006). That alone is important in light of the continuing need to understand discretion by the police as street-level bureaucrats and how much race affects policing, including use and severity of force. Plus, our theory and findings about the most extreme form of police use of force bear on classic concerns in political science, including but not limited to the exercise of power by the state, democratic accountability, and equality under the law (Brown 1981).

POLICE DISCRETION IN USE OF FORCE

Encounters with the police are among the most common encounters civilians have with government agents (Brown 1981; Jacob 1972; Soss and Weaver 2017). A key contrast with other civilian encounters with government agents is that police-civilian contact, whether initiated by police or initiated by civilians, has the potential for violence. How officers exercise their discretion to use force and violence during police-civilian encounters and why it may cause racial disparities are important considerations (e.g., Terrill 2011). “In the police shooting context,” in particular, “there is a concern that officers, despite their best intentions or conscious beliefs, will subconsciously let preconceived ideas about certain individuals influence their decision processes” (Worrall et al. 2018, 1176). This includes

2. Strength is defined in terms of percentage change in the bias factor of VanderWeele and Ding (2017).

3. We lack data on all instances of police drawing their weapons, but including moments where police drew guns without firing would likely increase the estimate of the lower bound (Worrall et al. 2018).

1. We use the techniques of Cinelli and Hazlett (2020) and a linear probability model.

their racial beliefs, which may bias their behaviors during police-civilian encounters. Inferring racial bias, however, is challenging.

Racial disparities in use of force

Generally, social scientists expect that police are more likely to use force and more of it against Black civilians than against White civilians (Goff et al. 2016; James, Vila, and Daratha 2013; Jetelina et al. 2017). Whether police do is well studied experimentally and observationally, often finding that officers are more willing to use force against Black civilians than against White civilians (Baumgartner, Epp, and Shoub 2018; Buehler 2017; Correll et al. 2007; Eberhardt et al. 2004; Engel and Calnon 2004; Johnson et al. 2019; Mekawi and Bresin 2015; Schuck 2004; Sikora and Mulvihill 2002; Terrill 2005; Worden 2015). Furthermore, the recent availability of “big data” on police-civilian encounters at the incident level (e.g., New York City’s stop, question, and frisk program) has enabled rigorous social science to deepen evidence of racial disparities in police use of force (e.g., Fryer 2016; Gelman, Fagan, and Kiss 2007; Goel, Rao, and Shroff 2016; Mummolo 2018; Pierson et al. 2017; Voigt et al. 2017).

However, some studies temper or contradict claims and the expectation of racial bias in police use of force, particularly shootings (e.g., Worrall et al. 2018). In other words, racial bias in policing may not necessarily increase the likelihood of use of force against Black civilians. Some evidence, drawn typically from observational studies, and limited by concerns about unmeasured confounding or misapplied methods (Alpert and Dunham 2004; Fryer 2016; Garner and Maxwell 1999; Garner, Maxwell, and Heraux 2002; Garner et al. 1995; Johnson et al. 2019), suggests we should expect and observe either smaller-scale or no racial disparities in police use of force (e.g., shootings). Plus, a “counter bias” may exist, inducing officers to be extra sensitive to the potential negative consequences of using force against racial minorities, especially Black civilians (James et al. 2013). The negative consequences of using force and more of it against Black civilians might be higher, not lower, than they are for using force against White civilians, even as the strength of evidence of that effect is debatable (Johnson et al. 2019; Knox and Mummolo 2020).

Challenges to inferring racial bias

Different conceptions of racial bias can exist. On the one hand, we could focus on the potential bias of the patrol officer who shoots a civilian. On the other hand, we could focus on the police department (and supervisors) of the officer. As Bittner (1970, 10) posited, “The ecological deployment of police work at the level of departmentally determined concentrations of deployment, as well as in terms of the orientations of indi-

vidual police officers, reflects a whole range of public prejudices.” For this study, we focus on bias by the patrol officer, acknowledging the potential of administrative control and bureaucratic bias to affect the context of police-civilian encounters (Brown 1981). However, we must acknowledge that the race of an individual is not randomly realized during police encounters with civilians.⁴ As a consequence, any inference about the causal effect of the race of a civilian on police use of force, or other police behaviors (e.g., driver or pedestrian stops), depends on the comparability of incidents.

Confounds in the use of force can be difficult to measure. Even if one can account for the lack of observed outcomes for officer-civilian encounters that never take place, empirical tests for racial bias still require accounting for confounds affecting contact and use of force (Knox et al. 2020). Race, for example, may be correlated with other characteristics (e.g., income, education, geography, employment, social networks) that might cause disparate rates of contact with police, thereby influencing civilian exposure to police use of force. Therefore, racially disparate patterns in the use of force and its severity may spuriously relate to characteristics of police-civilian encounters that explain use of force (e.g., Cesario, Johnson, and Terrill 2019; Jetelina et al. 2017; Knowles et al. 2001; Worrall et al. 2018). To best study the effect of race on the propensities of civilians to experience police use of force requires conditioning on a range of civilian characteristics that may confound the relationship. Furthermore, there is the matter of selection into contact with police and how it challenges drawing inferences about racial bias during citizen-police interactions (Johnson et al. 2019; Knox and Mummolo 2020; Knox et al. 2020).

Assuming racial bias in police shootings exists, there are at least two theoretical mechanisms, one circumstantial and the other psychological (for a brief discussion, see Ross [2015], 3). The first mechanism is that racial minorities, especially Black Americans, are circumstantially associated with conditions that give rise to police using greater force against them. They are more likely to come into contact with police because police officers racially profile them, or they are more proximate to high-crime or highly policed environments.⁵ The second mechanism is that police officers differentially perceive the stakes for using force against civilians depending on the race of the civilians. Officers might, for example, anticipate differential downstream consequences from using force against Black civilians than from using force against White civilians or interpret behaviors differently for Black and White civilians. In its most nefarious expression, regardless of the race of the

4. By “race” of civilians, we mean the officer’s perception of their race.

5. Racial profiling as a mechanism of racial disparities in use of force, however, is potentially circular.

officer, police may devalue the lives of Black civilians relative to the lives of White civilians.

A RACIAL BIAS MODEL OF POLICE SHOOTINGS

Our racial bias model of police shootings stems from the model Knowles et al. (2001) employ to examine police stops of drivers. It seeks to capture “the transactional, or step-by-step unfolding, of police-public encounters” and the “micro process of the police-suspect encounter,” in which civilian noncompliance, be it actual or perceived, can be pivotal to the decisions and discretion of police officers to use force (e.g., Terrill 2005, 100, quoting National Institute of Justice).

The first stage of our model is a selection stage. It allows for disparate rates of civilian encounters with police officers across civilian racial groups. Such an allowance is important. Encounters with police in which civilians are “suspect” are unequal. Differences in the deployment of and exposure to police in the United States are historic, with some races (and places) receiving greater surveillance, intervention, and state-sanctioned violence by the police, even when unmerited. In particular, studies from across the social and public health sciences of police contact with civilians, drawing on varied data from police records, public opinion surveys, face-to-face interviews, and focus groups, demonstrate that, generally, police devote greater—often needless—attention to Black civilians relative to White civilians for the same activities (e.g., traffic and pedestrian stops and outcomes of searches for contraband; Baumgartner et al. 2018; Epp, Maynard-Moody, and Haider-Markel 2014; Pierson et al. 2020; Prowse, Weaver, and Meares 2020).

Modeling the first stage allows us to make empirical predictions about behavior implied by racial bias that should manifest even in the presence of selection into encounters with the police. The selection stage captures, conceptually, every element of the police-civilian interaction that takes place up until the civilian and the officer reach the point of violence. It includes quotidian inequalities such as “attentional biases” to Black civilians in public and differential perceptions of “suspicious” and “threatening” civilian behavior by the race of civilians (Eberhardt et al. 2004), along with the social construction of the “Black symbolic assailant” (Bell 2017) and differences in civilian experiences with police discretion by skin color and phenotype (e.g., Kahn et al. 2016; Monk 2019).

In the second stage, we model a conflict subgame. It seeks to capture the kinds of split-second choices that police make at the point of using force. “During high-pressure situations, including some police-citizen encounters, . . . officers may not have the luxury of making slow, considered analytical decisions and, instead, rely on intuition and experience” (Hine et al. 2018, 1785). The same may be true for civilians. Nevertheless, the heightened pace of decision making, the urgency with

which individuals—both civilian and police—respond to real or perceived threats to their dignity and physical safety, and the uncertainty about each other (e.g., does the civilian have a gun or a wallet?), suggest that this process is accurately captured by a simultaneous structure.

In our model, conflict takes the form of escalating or accumulating aggression in the demeanor and deed of the civilian (actual or perceived by the officer) and the use of force by the officer, following initial interaction(s) between the civilian and officer (e.g., stopping the civilian, civilian disregard of verbal commands). We use “escalation” in a specific way: civilian demeanor or deed perceived by a police officer to be threatening, where real or misperceived aggression could “harm another person who is motivated to avoid that harm” (Allen and Anderson 2017, 1). It includes nonphysical noncompliance with police directives, inclusive of verbal hostility and antagonism (e.g., cursing or berating an officer), and physical noncompliance (e.g., turning from or striking an officer). Escalation by the civilian risks the dignity, respect, authority, or safety of an officer (or another civilian).

The choices of police during police-civilian encounters may partially result from the demeanors and deeds of civilians. Certainly, however, not all uses of force by police, especially shootings, or civilian deaths by police are entirely or at all affected by civilian behavior. A civilian may comply with a directive from an officer, displaying neither defiance nor belligerence, but an officer may mistake or misperceive the behavior of the civilian and use deadly force. Examples include the 1967 and 2014 nonfatal shootings of Huey Newton and Levar Smith and the 1999 and 2016 fatal shootings of Amadou Diallo and Philando Castille. Or, situational factors beyond the influence and control of civilians may influence shootings by police and civilian deaths by police. Informational priming by 911 dispatchers or other civilians, for instance, may exaggerate the degree of threat a “suspect” civilian poses for police, quickening lethal use of force by police when none was necessary (e.g., Tamir Rice, Breonna Taylor, and John Crawford). Also, civilians may be impaired by intoxicants or untreated mental illness, preventing them from making decisions or acting to reduce their appearance of threat to an officer (or other civilians), inclusive of nonresponse to police directives, resulting in civilian harm, inclusive of death (e.g., the fatal police shootings of Eleanor Bumpurs in 1984 and Daniel Prude in 2020). Additionally, training and socialization of police officers to expect immediate compliance with directives and to assume violence against them looms may influence the use of force in the absence of civilian escalation (Oberfield 2012; Sierra-Arévalo 2021). Finally, differences in the demeanor of police (e.g., tone, tenor, courtesy, and respect) when dealing with different civilians (Epp et al. 2014; Voigt et al.

2017) may test civilian patience, increase their aggravation, and possibly play a role in civilian escalation of conflict during encounters with police.

From the perspective of the “objectively reasonable” officer, civilian escalation of conflict may heighten the stakes of police-civilian encounters. At a minimum, conflict escalation can create “a type of strain that may also have situational effects, increasing officers’ anger and frustration toward specific civilians within individual encounters” (Nix, Pickett, and Mitchell 2019, 615). Plus, it may strengthen officer assumptions that conflict escalation signifies danger and “a greater likelihood of violence” (618).⁶

Together, perception, emotion(s), and assumptions likely account, in part, for the scholarly consensus that “non-compliant citizens face a greater likelihood of being treated disrespectfully by the police . . . [and] are more likely to experience other negative outcomes, such as arrest and the use of force” (Nix et al. 2017, 1155). We assume, therefore, that if civilian escalation of conflict may increase the severity of police use of force, it, in part, should increase the likelihood of death following police shooting a civilian. Studies that statistically associate the degree of civilian noncompliance (e.g., resistance) with police directives and the degree of police use of force against civilians buttress our assumption (e.g., Engel, Sobol, and Worden 2000; Garner et al. 2002; James, James, and Vila 2018; McCluskey and Terrill 2005; McElvain and Kposowa 2008; Sun, Payne, and Wu 2008; Wheeler et al. 2017).

We model the possibility of racial bias by allowing officer perceptions of the cost of fatally shooting a civilian to vary by the race of civilians. Our formal representation captures emotional reactions, anxiety, and threat perception associated with racial bias and the use of force (e.g., Correll et al. 2002; Kleider, Parrott, and King 2010; Nieuwenhuys, Savelbergh, and Oudejans 2012; Welch 2007), along with a more dispassionate cost-benefit analysis by the officer about the anticipated consequences of killing a civilian.

Primitives

Players, sequence of play, and strategies. The model is played between a civilian, C , and an officer, O . The civilian is characterized by a type, which is a pair, $\tau = \langle \kappa, \rho \rangle$. This pair includes a racial identity, $\rho \in \{B, W\}$, and observable civilian characteristics, denoted $\kappa \in \mathbb{R}$. The latter include dress, demeanor, location, time, or any other characteristic. We denote the probability density function of κ , conditional on ρ , as $g(\kappa|\rho)$.

6. The likelihood that police will use force may be greater, too, when officers have evidence that an offense or crime occurred (McCluskey and Terrill 2005; McCluskey, Terrill, and Paoline 2005; Sun and Payne 2004) or civilians possess weapons (Johnson 2011; McCluskey et al. 2005; Sun and Payne 2004). However, the seriousness of an offense or crime may not influence the likelihood that police use force (Friedrich 1980; Lawton 2007).

That is, the distribution of observable characteristics in the population can be different for any racial group. When we turn to the empirical implications of our model, we consider a population of civilians, \mathcal{P} , characterized by the density function, $g(\cdot)$, from whom the civilian in the interaction is drawn.

Figure 1 summarizes the play sequence. The game begins with the civilian, who engages in behavior the “objectively reasonable” officer could perceive as questionable or suspicious. Crucially, the behavior the civilian engages in need not actually be suspicious; it may be any kind of activity that an officer has the ability to further investigate (e.g., “loitering” or “furtive movement”). Let $s \in \{0, 1\}$ denote that choice, where $s = 1$ indicates the choice to engage in an activity, which could potentially be perceived as questionable or suspicious by an officer (or another civilian). If the “suspect” civilian chooses $s = 0$, the game ends. However, if the “suspect” civilian chooses $s = 1$, then the officer must use discretion to decide whether to engage the civilian for purposes of order maintenance or law enforcement (e.g., stop, question, and frisk). Let $l \in \{0, 1\}$ denote this choice, with $l = 1$ denoting engaging the civilian. If the officer chooses $l = 0$, the game ends; if he chooses $l = 1$, the game proceeds to the next stage, with simultaneous interactions by civilian and officer. Specifically, both players must decide how to engage the other, whereby each can choose behaviors that could escalate to violence. The civilian must choose to escalate or not, $t \in \{0, 1\}$, where $t = 1$ denotes escalating. (Reiterating an earlier point, escalation can be in the eye of the beholder, especially that of the police officer, influenced by different factors.) The officer must choose whether to use lethal force or not, $f \in \{0, 1\}$, where $f = 1$ denotes lethal force. If the officer chooses lethal force, the civilian dies with probability $\delta(t)$, where we assume $1 \geq \delta(1) > \delta(0) \geq 0$. That is, the probability the civilian dies when an officer uses lethal force is strictly greater when the civilian is escalating than when he is not, recognizing there can be exceptions, which we identified earlier. If neither player escalates conflict (i.e., $t = 0$ and $f = 0$), then less adverse, nonfatal outcomes follow. In either event, the game ends after these choices are made and payoffs are realized.

Let $\pi(\tau)$ denote a probability distribution over r , conditional on the civilian’s type, $\tau = \langle \kappa, \rho \rangle$, and let $\sigma(\tau)$ denote a probability distribution over f conditional on the civilian’s observable characteristics and race. A strategy profile for the civilian is, therefore, a tuple, $\mathcal{C} = \langle s, \pi(\tau) \rangle$, and a strategy profile for the officer is a tuple, $\mathcal{O} = \langle l, \sigma(\tau) \rangle$.

Preferences and utilities. Civilians have preferences over their behavior and the outcome of their interaction with the officer. Specifically, we assume that a civilian of type τ who chooses to engage in suspicious behavior, $s = 1$, receives a




Figure 1. Sequence of play in the model

payoff $c(\tau) > 0$ if the officer chooses not to engage in law enforcement activity (i.e., $l = 0$). This source of utility represents the value of engaging in whatever kind of behavior a citizen of type τ would like to engage in, without having to deal with the police. This payoff can depend on the individual's type (i.e., her race and observable characteristics). If the officer chooses to engage, though, $l = 1$, then we assume the civilian's payoff depends on whether the officer chooses to apply lethal force, as well as whether the civilian chooses a behavior that escalates conflict. If the officer chooses $l = 1$, then the civilian pays a cost, $-w(\tau)$, where we assume $w(\tau) > 0$, for all τ . This source of utility represents the cost of being subjected to policing and, as with the value of potentially suspicious behavior, can depend on the civilian's type. In addition to the cost of being subjected to policing, we assume the civilian pays a cost $-d(\tau)$ if he dies. That is, if the officer chooses to use lethal force (i.e., $f = 1$), then the civilian pays, in expectation, $-\delta(r) \cdot d(\tau)$, where we assume $d(\tau) > 0$. This source of utility represents the cost associated with the loss of life, which can depend on civilian type (i.e., some civilians may value living more than others such as the suicidal). To avoid considering unreasonable situations, we assume that the cost of dying is worse than the cost of being subjected to policing for all types of civilians.

Assumption 1 (Civilians prefer not to die). $d(\tau) > w(\tau), \forall \tau$.

If the civilian escalates, and the officer chooses less than lethal force, we assume the civilian receives positive utility

$b(\tau) > 0$. The source of utility represents the value of engaging in escalation against an officer and can vary by type. The civilian's expected utility function is given by

$$EU_C(s, t | \tau) = \begin{cases} 0 & \text{if } s = 0 \\ c(\tau) & \text{if } s = 1 \& l = 0 \\ -w(\tau) & \text{if } s = 1 \& l = 1 \& t = 0 \& f = 0 \\ b(\tau) - w(\tau) & \text{if } s = 1 \& l = 1 \& t = 1 \& f = 0 \\ -w(\tau) - \delta(r) \cdot d(\tau) & \text{if } s = 1 \& l = 1 \& f = 1 \end{cases}$$

The officer has preferences over conducting policing work, stopping suspects and criminals, fatally wounding civilians, and his own physical well-being. Specifically, we assume the officer pays a cost $-c_o(\tau)$, where $c_o(\tau) \in (0, 1)$, whenever the civilian chooses to engage in potentially suspicious activity (i.e., $s = 1$) and the officer does not engage in law enforcement (i.e., $l = 0$). This cost represents the cost of allowing potentially criminal activity to go overlooked or a forsaking of duty. Importantly, we allow this cost to vary by civilian type. Allowing an officer's disutility from permitting potentially criminal activity to occur is a function of everything the officer can observe about the civilian. In addition, the officer pays a cost $-d_o$, where we assume $k_p \in (0, 1)$ for all ρ , whenever he fatally wounds a civilian of race ρ . By contrast, the officer pays a cost, $-d_o$, where $d_o > 0$ whenever a civilian is escalating and he does not use lethal force (i.e., $f = 0$). Substantively, this cost can represent injury to the officer, disutility from not stopping a "suspect" civilian acting aggressively, or another adverse consequence. Finally, we assume the officer receives positive utility 1 from using force to stop a "suspect" civilian who

escalates conflict. This represents the utility of exercising authority, maintaining order, and stopping a potentially dangerous person. Therefore, the officer's expected utility function is given by

$$EU_O(\gamma, \lambda | \tau) = \begin{cases} -c_o(\tau) & \text{if } s = 1 \& l = 0 \\ -d_o & \text{if } s = 1 \& l = 1 \& t = 1 \& f = 0 \\ -w_o - \delta(0) \cdot k_\rho & \text{if } s = 1 \& l = 1 \& t = 0 \& f = 1 \\ 1 - \delta(1) \cdot k_\rho & \text{if } s = 1 \& l = 1 \& t = 1 \& f = 1 \\ 0 & \text{otherwise} \end{cases}$$

Analysis

We characterize a mixed-strategy subgame perfect Nash equilibrium. There can exist a pure strategy equilibrium if officers are never willing to use lethal force, which we rule implausible by assumption. For the officer to be willing to play a mixed strategy, the civilian must choose a probability distribution over her decision to escalate that makes the officer indifferent between using lethal force and not. There is a probability that satisfies this requirement:

$$\pi^*(\tau) = \frac{w_o + \delta(0)k_\rho}{1 + w_o + d_o - (\delta(1) - \delta(0))k_\rho}. \quad (1)$$

Notice that $\pi^*(\tau)$ is increasing in k_ρ . As an officer perceives it to be costlier to kill a civilian of race ρ , the civilian will be more likely to escalate. In addition, $\pi^*(\tau)$ is decreasing in $(\delta(1) - \delta(0))$. Hence, as the civilian's behavior has a larger impact on the probability of dying when the officer uses force, the equilibrium probability of a civilian escalating conflict will decrease. Intuitively, this makes sense: if the civilian's behavior does not matter, fatality becomes irrelevant for her calculation, and fatality is the major factor deterring her from being contentious. At the same time, the officer's equilibrium probability distribution over using lethal force, $\sigma^*(\tau)$, must make the civilian indifferent about choosing to escalate. That probability is given by

$$\sigma^*(\tau) = \frac{b(\tau)}{b(\tau) + d(\delta(1) + \delta(0))}. \quad (2)$$

Thus, in any equilibrium that reaches the conflict subgame, there exists a mixed-strategy subgame perfect Nash equilibrium where civilians probabilistically escalate and officers probabilistically use lethal force.⁷

Proposition 1. In any subgame perfect Nash equilibrium where players reach the conflict subgame, the

civilian and officer play mixed strategies whereby a civilian of type $\tau = \langle \kappa, \rho \rangle$ chooses to escalate with probability $\pi^*(\tau)$, and the officer chooses use lethal force with probability $\sigma^*(\tau)$.

Empirical implications

How does racial bias by police officers affect equilibrium behavior? We offer a simple definition of bias, guided by Knowles et al. (2001). Specifically, we say that *an officer is racially biased if he perceives the cost of shooting an individual to vary by racial groups*. If an officer thinks it is less costly to shoot a Black civilian than a White civilian, then we say the officer is biased against Black civilians.

Definition 1. An officer is racially biased if $k_B \neq k_W$.
An officer is racially unbiased if $k_B = k = k_W$.

With this definition in hand, proposition 1 is instructive about evidence of racial bias by police in OIS. Given definition 1, we can identify the probability that a civilian should die, conditional on being involved in an OIS, when the police are not racially biased and when they are racially biased.

Importantly, the model yields implications for how we can infer bias without having to make judgments about how to measure group traits, benefits to crime, or the distribution of traits in a group. That is, we are able to draw inferences from OIS outcomes among those who are actually involved in a shooting, without having data on the selection process that leads individuals into OIS events. Specifically, let $K(\rho)$ represent the set of characteristics for which an individual of race ρ would choose $s = 1$. Then, the fatality rate among people who are shot is given by

$$\mathcal{F}(\rho) = \int_{K(\rho)} (\delta(1) \cdot \pi^*(\tau) + \delta(0) \cdot (1 - \pi^*(\tau))) \frac{\sigma^*(\tau)g(\kappa|\rho)}{\int_{K(\rho)} \sigma^*(z|\rho)g(z|\rho)dz} d\kappa. \quad (3)$$

Notice that this fatality rate is not the fatality rate for all civilians of a given race but only for those who are shot by a police officer.

Notice that by definition 1, if an officer is not racially biased, then $k_B = k = k_W$. Given the civilian's equilibrium strategy, $\pi(\tau)^* = (w_o + \delta(0)k_\rho) / (1 + w_o + d_o - (\delta(1) - \delta(0))k_\rho)$, from above, then we can substitute $(w_o + \delta(0)k_\rho) / (1 + w_o + d_o - (\delta(1) - \delta(0))k_\rho)$ for $\pi^*(\tau)$. Because this quantity is independent of κ , equation (3) reduces to

$$\mathcal{F}(\rho) = \delta(0) + (\delta(1) - \delta(0)) \left(\frac{w_o + \delta(0)k_\rho}{1 + w_o + d_o - (\delta(1) - \delta(0))k_\rho} \right). \quad (4)$$

7. In the appendix, we show that the civilian and officer reach the conflict subgame under intuitive conditions.

Notice the only way this quantity varies with civilian race is if the officer's perceived cost of taking a civilian life varies by race. Therefore, differential fatality rates can only arise as a result of racially biased policing.

Proposition 2. In equilibrium, different fatality rates by racial groups arise only when the officer is racially biased.

The consequence is that if police are not racially biased then the probability a civilian is killed in an OIS, conditional on being involved in a shooting, should be independent of her race, even accounting for all other observable characteristics that might influence her incentive to engage in noncompliance or resistance, as well as the officer's incentive to use force in the first instance. That is, equation (3) provides the theoretical foundations for a sufficient test of racial bias in the use of lethal force in OIS. It is important to underscore that this implication of our model allows us to evaluate evidence of racial bias, even taking into account unobservable behavioral differences across racial groups that might take place during a police-civilian encounter. This result is parallel in logic to the way Knowles et al. (2001) study racial disparities in traffic stops and Alesina and La Ferrara (2014) study bias in capital sentencing. It allows us to assess evidence of racial bias without having to measure observable or behavioral characteristics of either civilians or officers. It is sufficient to evaluate variation in ultimate consequences—namely, patterns of fatality.

Implication 1. If police officers are racially biased in favor of shooting Black civilians, then, conditional on being involved in an OIS, Black civilians will be less likely to die than will non-Black civilians.

The core logic underlying this implication is that officers will be more likely to use force in less dangerous situations involving Black civilians than in similar situations involving White civilians. As a consequence, a greater proportion of OIS involving Black civilians will not lead to a fatal outcome.

A corollary implication of our model is that White civilians should be more likely than Black civilians to escalate conflict with officers. That implication helps clarify the underlying theoretical mechanism we posit: Black civilians are induced to be more cautious during an interaction with police than are White civilians.

Implication 2. If police officers are racially biased against shooting White civilians, then, conditional on being subjected to law enforcement activity, White civilians will be more likely to engage in threatening behavior, such as resisting arrest, disobeying officer

commands, or behaving belligerently than will non-White civilians.

It is beyond the limits of this article to fully investigate that implication because of insurmountable data limitations, particularly data on the perceptions or degree of civilian escalation, but its verisimilitude is important for establishing the mechanism that drives the analysis we present. To that end, we note that beyond anecdotal support for the mechanism, there is some evidence from extant literature to support the implication. Kavanagh (1997) studies more than 1,000 encounters between civilians and officers in New York City's Port Authority Bus Terminal between 1990 and 1991 and finds suggestive evidence that White civilians are more likely to resist arrest than are non-White civilians. Matrofski, Snipes, and Supina (1996) compare civilian-officer race combinations as predictors of civilian compliance with officer requests for orderly behavior. They find that, compared to White civilians interacting with White officers, White civilians interacting with minority officers are less likely to comply with officer instructions. At the same time, they find that minority civilians interacting with White officers are more likely to comply with officer instructions. They also find that minority civilians interacting with minority officers are more likely to comply, although this difference is not statistically significant. Finally, according to the FBI's Law Enforcement Officers Killed and Assaulted data, as of July 2017, 55% of officers killed by civilians were killed by White civilians and 58% of officers assaulted by civilians were assaulted by White civilians. While far from constituting a systematic evaluation, those descriptive findings provide initial evidence to corroborate the underlying mechanism we posit. However, for the remainder of the article, we evaluate the primary implication of the mechanism articulated above.

EMPIRICAL ASSESSMENT

Our empirical assessment of the implications for racial bias in police shootings proceeds in four steps. First, we describe our method—the outcome test. Second, we describe an original data set we built that includes all OIS (fatal and nonfatal) in eight local police jurisdictions but, because of data limitations imposed by police reporting, excludes sufficient data on civilian behavior (or police perceptions of it) during the police-civilian encounter. Third, we focus on an evaluation of implication 1 that predicts that racial bias among police officers will produce disparities in fatalities across racial groups. We underscore that this prediction is not intended to estimate the effect of civilian race on the decision to use force; it is designed to demonstrate evidence implied by any such bias. In the fourth step, therefore, we directly engage the size of the bias. Assuming no omitted covariates, we calculate a lower bound

for the magnitude of racial bias in the decision of an officer to shoot a civilian in our sample of localities. We then calculate how strong an omitted covariate would need to be in order to eliminate the findings.

Discerning racial bias: The outcome test method

To evaluate implication 1, we employ an outcome or “hit rate” test, which is capable of observing disparate impact and identifying bias in decision-making (e.g., Alesina and La Ferrara 2014; Knowles et al. 2001; Persico and Todd 2006). Mortgage lending illustrates the general logic of the approach. Mortgage lenders may care about timely repayment of loans. If we observe that non-White lendees repay mortgages on time at higher rates than Whites lendees, then that would suggest that qualified non-White applicants are being denied loans (Ayres 2002). If the same standard were applied for mortgage lending, independent of borrowers’ race, we should expect similar default rates across racial categories. However, because lenders were willing to lend to less qualified White borrowers than to Black borrowers, the default rate would be higher for White borrowers. For policing, we may see similar systematic differences by race, in the other direction. Stops may be considered successful, for instance, if they lead to arrest, perhaps because of the discovery of contraband or the harmful behavior of drivers. Gelman et al. (2007), for example, found that 1 in 7.9 White people police stopped were arrested, compared to 1 in 9.5 Black people. That suggests the discretion threshold police use to decide whom to stop is lower or more indiscriminate for Black drivers than for White drivers. Our logic similarly implies that if officers have a lower threshold for deciding to shoot Black civilians than White civilians, then there will be a greater proportion of Black civilians who will choose to not threaten and, therefore, survive an OIS.

Importantly, in many traditional settings, hit-rate tests are used to evaluate the presence of a latent trait to uncover evidence of bias. In our setting, as in Knowles et al. (2001), the latent trait is a choice by another player. In Knowles et al. (2001), drivers strategically choose whether to carry contraband; in our model, civilians strategically decide how to behave during police-civilian encounters. Anticipating bias by officers, Black civilians will be less likely in equilibrium to behave in ways that escalate a confrontation toward police-civilian violence than will White civilians. That feature is a result, not an assumption. The motivating assumption, as we noted above, is that the risk of death should be higher during a police-civilian encounter involving civilian escalation than one without it.

A note on causality

Before presenting our analysis, we underscore the causal pathway at the heart of our argument. Our claim is not that

racial bias directly causes differential fatality rates. Our argument is instead that racial bias causes officers to use force differently in different situations across racial groups. Anticipating that, civilians interact differently with officers in a way correlated with the civilian racial identities. The effects of those behaviors in conjunction is a distribution of force-civilian action combinations that vary by civilian race. Our analysis reveals that differential fatality rates are evidence consistent with that effect, not the effect itself. Just as we would not argue that differential default rates by race are a direct effect of racial bias in mortgage lending, we do not argue that differential fatality rates are a direct effect of racial bias in the decision to use force. Thus, as we proceed to our empirical analysis, we do not set out to demonstrate a causal effect of bias on fatality rates because the path from bias to fatality rates runs through myriad immeasurable intermediate mechanisms.

Data on officer-involved shootings

To evaluate racial disparities in fatality rates among different racial groups, we require data on every single OIS, not just fatal shootings. Data on OIS—even just fatal ones—are notoriously difficult to acquire (Zimring 2017). Recent efforts have begun to compile extensive data on fatal encounters between officers and civilians. They typically rely on media reports and crowd-sourced data, making it difficult to assess how comprehensive and systematic the data are. Moreover, existing data typically do not include instances of OIS that do not include a fatality. Thus, we collected original OIS data by filing public records requests with individual police departments.

We sent public records requests to police departments and sheriffs’ offices in the 50 largest local jurisdictions in the United States, measured by population. We requested records of every single instance of an officer discharging a weapon between 2010 and 2017. Although most policing agencies were positively responsive to our requests, most that responded with data did not provide racial information about civilians involved in OIS. Our data, therefore, comprise eight jurisdictions—Charlotte; Houston; King County, WA; Los Angeles; Orlando; San Jose, Seattle; and Tucson—that provided comprehensive racial information in response to our public records requests.⁸ The unit of analysis for each incident is the civilian/officer pair.⁹

We constructed all civilian/officer pairs, yielding 1,274 total pairs, representing 748 unique incidents. Overall, 48% of our

8. Unfortunately, the departments could not provide objective data on observed officer interactions with civilians or civilian behavior during all interactions with police officers, and they often could not provide even subjective data for interactions involving use of force by officers and civilian behavior leading up to it.

9. San Antonio also provided such information, but the sample size was too small to make Black/White comparisons. Results are available on request.




Figure 2. Logged number of officer-involved shootings per month in eight locations, 2010–17

OIS incidents represent fatal shootings, varying considerably by department. Charlotte had the lowest rate of fatalities from OIS, where 9 out of 45 observations were fatal (20%). Los Angeles had the highest number of reported OIS (663), where 58% of them were fatal. Our data demonstrate we have considerable variation in OIS incidents, not just by department and by time (see fig. 2) but by fatality, too.

Figure 2 shows the frequency of OIS in each of the jurisdictions. Because there is considerable variation in the size of the jurisdictions, there is considerable variation in the total number of OIS. The most come from Los Angeles, the second-largest police jurisdiction in the country. Therefore, we log the number of observations per month, to prevent scale differences from skewing the temporal patterns and cross-jurisdiction variation. Notably, with the exception of an increase in OIS in Houston at the end of the series, there is little within-city variation in the frequency of OIS.

Furthermore, the spatial distribution and concentration of OIS within jurisdictions show intuitive but instructive patterns. Figure 3 shows the distribution of fatal and nonfatal shootings in our eight jurisdictions. Los Angeles and Houston, by far the largest localities in our data set, experience the most OIS, whereas cities like Charlotte and Tucson experience relatively few. Additionally, it appears there is a higher fatality rate among OIS in localities like Los Angeles and Houston, which is less of an issue in jurisdictions like San Antonio and Charlotte. Overall, figure 3 highlights the geographical diversity in these fatal OIS, that they do not appear to systematically occur in only certain parts of certain localities, and that fatality rates vary across geographies.

Analysis and results

We begin our empirical analysis of implication 1 by simply comparing the distribution of fatalities across racial groups, conditional on being involved in an OIS. Table 1 summarizes the frequencies among the observations in our data. The columns break down OIS by the race of the civilian involved, and the rows distinguish between fatal and nonfatal OIS.

The evidence is startling, revealing considerable dependence between fatalities and the race of the civilian ($\chi^2 = 76.888$, $p \leq .001$). In particular, a majority of Black civilians survive OIS, whereas a majority of civilians of all other races do not. Of course, demographics and police behavior both vary across jurisdictions, and we might worry that the correlation detected in table 1 is spurious. To speak to this we estimate a series of logistic regression specifications on all observations of OIS for which the departments we sampled provided race information. The unit of analysis is the civilian involved in an OIS, and the outcome variable is an indicator for whether the civilian was fatally wounded. For 17 observations, the outcome was recorded as “undetermined” or “unknown.” We treat these observations as missing data. Our primary explanatory variable of interest is the race of the civilian involved.

We also consider specifications in which we include as explanatory variables the distance from each OIS to the nearest trauma center as well as year fixed effects (see the appendix for the specifications with year fixed effects).¹⁰ We also include

10. Some observations lacked adequate location information to calculate the distance to the nearest trauma center, which has been shown to be a particularly important factor for the chances of survival of a gunshot




Figure 3. Locations of fatal shootings (filled circles) and nonfatal shootings (open circles) in our sample of eight locations. Triangles mark level 1 trauma centers.

fixed effects for the cities from which we have data, which are likely correlated with the distance to a trauma center and the racial indicator. This is because trauma centers have fixed locations in cities, and demographic characteristics of populations vary across cities. Unfortunately, for 24 of our 1,274 observations, the address of the OIS was too imprecise to calculate a reliable distance measure. We consider specifications both with and without this control variable.

The main results of our analysis are reported in table 2. The primary result appears in the top row. In each of our specifications, among those civilians shot by an officer, Black civilians are less likely to die than are White civilians. This difference is statistically significant in each specification. In our main specification, reported in the first column of results, White civilians have a predicted probability of 0.52 of dying, whereas Black civilians have a predicted probability of dying of

0.32—a 20 percentage point decrease. This relationship supports the primary empirical implication of our theoretical model of racial bias. It is consistent with the claim that police officers have a lower threshold for deciding to use lethal force against Black civilians than against White civilians. Notably, the relationship between being a Hispanic civilian and a reduced probability of dying does not emerge even after we include jurisdiction and year fixed effects. This functions as a placebo test and implies that any problematic unmeasured covariates would have to have different relationships for Black

Table 1. Summary of Officer-Involved Shootings by Race and Fatality

	White	Black	Hispanic	Asian
Not fatal	118 (48%)	329 (67%)	208 (42%)	11 (27%)
Fatal	126 (52%)	162 (33%)	290 (58%)	30 (73%)

Note. $\chi^2 = 76.888$, $p \leq .001$.

wound (Crandall et al. 2013). Therefore, in the models including distance to the nearest trauma center as a control variable, we only have 1,250 observations, covering 729 unique incidents.

Table 2. Estimated Relationship between Civilian Race and Probability of Fatality Conditional on an Officer-Involved Shooting

	Model 1	Model 2	Model 3	Model 4
Black	-.77*	-.70*	-.74*	-.67*
	(.32)	(.33)	(.32)	(.33)
Hispanic	.27	.10	.29	.13
	(.31)	(.32)	(.31)	(.31)
Asian/AI/AN/PI	.94	.82	.97	.91
	(.62)	(.58)	(.63)	(.58)
Distance			.14	.41
			(.19)	(.21)
Houston	.00			-.21
		(.57)		(.63)
King County	.27			-.12
		(.76)		(.81)
Los Angeles	1.26*		1.15	
		(.57)		(.62)
Orlando	.59		.52	
		(.65)		(.70)
San Jose	.18		.16	
		(.65)		(.69)
Seattle	1.25		1.27	
		(.73)		(.77)
Tucson	1.61*		1.64*	
		(.69)		(.73)
Intercept	.07	-.80	-.04	-1.00
	(.25)	(.59)	(.28)	(.64)
N	1,274	1,274	1,250	1,250

Note. Logit coefficients with cluster-robust standard errors (in parentheses). Omitted category is White civilians and Charlotte. Distance is in tens of miles. AI = American Indian; AN = American Native; PI = Pacific Islander.

* $p < .05$.

** $p < .01$.

*** $p < .001$.

and Hispanic civilians (e.g., concerns about characteristics that affect the probability of death—such as police behavior, training, and medical attention—would be largely ruled out by this analysis).

As we do not observe a depression of the relationship between being a Black civilian and the probability of survival after we include jurisdiction and time fixed effects, a spurious correlation between race and jurisdiction does not drive the observed relationship. This pattern—while not necessarily causal—is precisely what we expect if police are racially biased in favor of shooting Black civilians, given the logic of our model. In order to explore the possibility that the relationship would be eliminated by an omitted covariate, we conduct a number of sensitivity analyses based on the methodology

presented in Cinelli and Hazlett (2020).¹¹ These analyses consider how strong an unmeasured confounding variable would have to be in order to wipe out the effects we are finding for Black civilians. One method to measure such strength is to benchmark any potential unmeasured confounder against measured covariates in the model. In analysis presented in the appendix, we show that in order to eliminate the significance of the apparent effect, there would need to be an unmeasured confound that is more than three times as strong as any of the variables currently included in the model (jurisdiction fixed effects, time fixed effects, and distance to trauma center). For example, we do not have race of the officer in the data set. One could propose a theory whereby there are more Black officers in Black neighborhoods and perhaps propose that Black officers were more likely to nonfatally shoot Black civilians (conditional on shooting them). However, in order to eliminate the estimated effect, one would have to simultaneously claim that the strength of this relationship was at least three times as strong as distance to trauma center or any cross-jurisdiction variation. This would also need to be true for any conceptually unmeasurable variable to eliminate these findings.

How big of an effect could racial bias have on officer-involved shootings?

Our analysis revealed evidence consistent with racial bias, per our definition, in the decision of police officers to use lethal force. However, we have not yet quantified the size of the bias, substantively. Accordingly, we estimate a lower bound on the magnitude of racial bias in OIS, relying on logic and assumptions paralleling Cohen (2021), Cohen and Glynn (2021), and Knox et al. (2020) for identifying racial bias in police contact with civilians. The approach we adopt has two steps. First, we define the fatality rate for Black civilians that police shot, comprising two components—those that would not have been shot had they been White and those who would have also been shot were they White. Second, we define the fatality rates of groups relative to each other.

The magnitude of racial bias in the decision to shoot a civilian is the proportion of Black civilians shot who would not have been shot had they been White. The intuition behind this is that the observed fatality rate of Black civilians is made up of two components—Black civilians who were shot but would not have been shot had they been White and Black civilians who would have been shot had they instead been White. Our quantity of interest, p , is the proportion that are in the former, that is, the proportion of Black civilians shot who would not

11. Although the Cinelli and Hazlett (2020) analysis is based on a linear probability model, we generally find small differences for these data between analyses based on the logit model and the linear model.

have been shot had they been White. By using the principal strata defining these groups we can derive a lower bound for p as the ratio of a difference of fatality rates: rate for Black people shot who would have been shot if White minus rate for Black people over rate for Black people shot who would have been shot if White minus rate for Black people who would not have been shot if White:

$$p = \frac{\mathcal{F}_{s(b)=s(w),b} - \mathcal{F}_b}{\mathcal{F}_{s(b)=s(w),b} - \mathcal{F}_{s(b)>s(w),b}}. \quad (5)$$

Equation (5) contains counterfactual quantities, so to derive an empirically estimable lower bound for p , we assume that the fatality rate for White civilians is no greater than the fatality rate for Black civilians who would have been shot if White (Cohen and Glynn 2021). Furthermore, we do not observe $\mathcal{F}_{s(b)>s(w)}$; however, substituting 0 for $\mathcal{F}_{s(b)>s(w)}$ yields a lower bound on the true value of p :

$$p \geq \frac{\mathcal{F}_w - \mathcal{F}_b}{\mathcal{F}_w - \mathcal{F}_{s(b)>s(w),b}} \geq \frac{\mathcal{F}_w - \mathcal{F}_b}{\mathcal{F}_w}. \quad (6)$$

Equation (6) expresses p as a function of \mathcal{F}_w and \mathcal{F}_b , the observed fatality rates among White and Black civilians shot. See the appendix for formal assumptions, definitions, and derivation of our quantity of interest, p . To estimate the lower bound on p , we first estimate a logistic model. We estimate it with a subset of OIS data containing only Black and White civilians, including our main covariate of interest, namely, race (White equal to 1, Black equal to 0), along with binary indicator variables for locality and a continuous variable of distance to closest trauma center in miles. Using this model we estimate the regression coefficient on White to be 0.70 (see the appendix for full regression specification results), the associated fatality difference between White civilians and Black civilians controlling for city fixed effects. The lower bound estimate, p , follows from the estimated risk ratios (Cohen 2021) as in equation (7) (see the appendix for the derivation) and uses a Poisson regression to estimate the risk ratio:

$$p \geq 1 - \frac{1}{RR}. \quad (7)$$

Thus, we estimate that 30% is the lower bound on the proportion of Black civilians that police would not have shot had they been White. Potentially there are unmeasured confounders that would affect both race and the likelihood of being fatally shot. As a sensitivity analysis we use the techniques of Cohen (2021) and VanderWeele and Ding (2017). These analyses indicate that an omitted covariate would have to produce a bias factor with a percentage change three times

stronger than any covariate in our data set in order to eliminate the proportion p .

Substantively, our estimate of 30% is considerable and, given it is a lower bound, may be higher. Our estimate implies that police would not have shot 156 Black civilians had they been White, from the 497 Black civilians in our eight localities over the years we study. Extrapolating this estimate to the larger population of the United States, however, is beyond the limits of our data. Moreover, significant intralocality variation suggests police behavior, measured by OIS, is not uniform across the country. Additionally, comparing Hispanic civilians and Asian civilians to White civilians yielded no statistically significant differences. That is consistent with what we would expect—police officers differentially exercise discretion against Black civilians as compared to all other groups. Given the extant debate about whether the use of force by police is tainted with racial bias, these findings suggest there is a substantively significant problem. Quantifying the magnitude of its effect, though, requires richer administrative data beyond what police departments, generally, in the United States currently provide. Specifically, the important matter of how much police violence is attributable to racial bias requires knowing how often police fire their weapons, as well as how often they draw their weapons (e.g., Worrall et al. 2018), which is not universally known across local police departments.

DISCUSSION

A significant challenge to credible inferences about the influence of racial bias in policing is that empirical observations typically need to condition on a wide range of difficult-to-measure confounds. For example, if civilian race is correlated with factors that directly affect contact with police—such as income, locality, employment rates or sectors, education level, or any possible factor—then it will be challenging to disentangle the causal effect of one's race from the effects of those other confounding forces. However, our approach helps overcome that challenge by identifying an empirical implication of racial bias in the use of force that is conditional on contact with the police, allowing social scientists to sidestep the challenges of selection bias due to racial rates of police contact with civilians (e.g., Knox et al. 2020).

What is more, our theory, analysis, and results help make better sense of seemingly contradictory findings in the contemporary use of force literature. For example, some studies show that the probability of being Black, conditional on being shot, is not statistically different from the probability of being White, conditional on being shot (Johnson et al. 2019). In our theoretical model, however, this pattern is completely consistent with racial bias by officers in favor of shooting Black civilians. Such a pattern could emerge because Black civilians

are aware of such bias and systematically avoid escalation during encounters with the police that could lead to fatal OIS. Therefore, the probability of being shot, conditional on being Black, might still be higher than it is conditional on being White, even while the observed rates of being fatally wounded are the same. Similarly, our analysis can reconcile the distinction Fryer (2016) documents between lethal and nonlethal force against civilians.¹² If Black civilians are aware (or believe) that police officers are biased in favor of using force against them, then they should be less likely to engage in threatening behavior that would escalate a situation from a nonlethal outcome to a lethal outcome. We would expect, then, that Black civilians should be disproportionately subject to nonlethal force but not necessarily disproportionately represented in lethal encounters with police.

At the same time, while our analysis helps explain racial differences across the observed patterns in police use of force, all we can demonstrate is evidence consistent with racial bias. The primary implication of our model, and the one we subject to empirical scrutiny, is a statement of an empirical regularity that is implied if civilians and officers behave as though the latter are racially biased. Lower fatality rates among Black civilians shot by the police than among White civilians shot by the police are a secondary form of evidence—a pattern implied by racial bias in the decision to shoot in the first instance. Those rates, however, do not in and of themselves tell us anything about the magnitude of the effect of bias.

However, given what we know about the existence of racial bias, we are able to calculate a lower bound on the effect size under an assumption of no omitted covariates. Still, the bounds we estimate cannot tell us about the upper limit on the effect or the lower limit if covariates are considerably stronger than our observed covariates. Refining these bounds, while not necessarily an impossible task, remains one of the most salient limitations that research on the subject faces.

As we noted above, we have not investigated implication 2. Doing so would require objective data on observed officer interactions with civilians. In particular, we would need data on civilian behavior during all interactions with police officers, not just those involving use of force by officers. Such data are difficult to come by. However, it bears noting that there is some evidence in the extant literature that is potentially consistent with the expectation. It predicts that, if officers are racially biased against Black civilians, White civilians will be more likely to engage in escalating behavior than will Black civilians. While doing so would require the collection of rich new data

that are not currently available, we believe it is a worthy endeavor as scholars continue to work out the mechanisms underlying disparate outcomes in civilian-officer interactions.

CONCLUSION

Police-civilian encounters have special implications for the study of democratic governance and equality of citizenship. Police are perhaps the most common government officials with whom civilians have contact (e.g., Jacob 1972), and, distinct from other bureaucrats, interactions with police officers always have the potential for violence. Consequently, the modal contact a civilian has with police relative to other government agents in the United States is one that might involve the use of physical force, including fatal and nonfatal shootings. Yet, whether justified or not, whether garnering mass and elite attention or not, whether we know enough about correlates and causes or not, police shootings (and other forms of police use of force such as use of compliance holds, pepper spray, and canines) are moments that “raise fundamental questions of governmental responsiveness and state power, and they are frequently at the heart of grievances that generate political demands and protests” (Soss and Weaver 2016, 83). Police shootings, along with predatory and extractive policing (Sances and You 2017), police “militarization” (Lawson 2019), and broader practices of policing, inclusive of surveillance, order maintenance, and arrests, coupled with choices by local prosecutors and judges (e.g., requiring bail and jailing arrestees for low-level offenses), invite political scientists to ask “questions about police authority, state projects of social control, and daily encounters with local governance” (Soss and Weaver 2017, 568). They also invite questions about the influence of bias, especially racial bias.

Racial bias on the part of government officials has the distinct potential to undermine the legitimacy of the state and civilian cooperation and engagement with government. To the extent, then, that police officers engage in racially biased use of force, that behavior has potentially profound consequences for the maintenance of a well-functioning democratic order. In light of these observations, recent analyses of racial disparities in the use of force by police officers have set out to address whether and how much racial bias influences policing in the United States. The implications of the findings are far-reaching.

Our results raise concern about racial bias in the use of force by police. They also highlight the need for more research and more comprehensive data about OIS, including, among other things, officer attributes and situational and contextual factors. For example, to understand the mechanisms by which racial bias affects civilian and police behavior, scholars need to study all civilian interactions with police, not just those encounters ending in fatalities or even just the encounters where the use of

12. Of course, Knox et al. (2020) also suggest that the analysis in Fryer (2016) is flawed because of selection bias.

force occurred. Of course, as others have pointed out (e.g., Knox et al. 2020) and as our model considers, there is potentially racial bias in the initial selection of civilians into contact with police. To the extent racial bias systematically affects not just how police interact with civilians but which civilians they interact with, our analysis underscores the extent to which training, recruiting, and monitoring of police officers have implications beyond public and officer safety.

Although our empirical study provides evidence consistent with racial bias in the use of force and a lower bound on the magnitude of racial bias in the decision to shoot, more research is necessary to assess the magnitude of the effect. We also need more research on racial bias in policing to assess the efficacy of policies designed to minimize racial disparities in policing, as well as to determine the underlying mechanisms that produce such racial bias. While normatively we might believe that, independent of its cause, racial disparities are problematic, what to do about them depends on identifying the causes. In particular, whether racial disparities are a result of circumstantial factors or systematic bias by police officers affects what kinds of remedies are desirable and the implications of the disparities for the legitimacy and integrity of the police as a key law enforcement institution.

But better research will require richer administrative data on police practices, ranging across both the use of force continuum (e.g., no guns, guns drawn, guns fired) and outcomes (i.e., lethal and nonlethal consequences), as well as civilian behavior (e.g., resistance). The current nature and contents of use of force and consequences record keeping by many police departments, however, present serious challenges to improving research and establishing consensus in weighting across the varied factors associated with OIS. Decentralization of law enforcement and varied discretion across localities in the United States further complicates research. Nonetheless, police departments, elected officials, and institutions of civilian oversight of police departments may become more interested in research about policing practices and outcomes, more anticipatory of scholarly needs, more transparent about and willing to share data with scholars and others through digitization and open access, and interested in replication and extension of academic studies. If so, causal research on police behavior, from the spectacular to the mundane, may flourish, perhaps improving policy making for public safety and improving policing (and police legitimacy) in the United States.

ACKNOWLEDGMENTS

We thank Chuck Cameron, Anna Harvey, Hakeem Jefferson, John Kastellec, Beatriz Magaloni-Kerpel, Pablo Montagnes, John Patty, Zac Peskowitz, Dan Thompson, Georg Vanberg, and Kirsten Widner for helpful comments. Earlier versions

of this article were presented at the 2019 meeting of the American Political Science Association, the Politics of Policing Conference at Princeton University, and seminars at Emory University.

REFERENCES

Alesina, Alberto, and Eliana La Ferrara. 2014. "A Test of Racial Bias in Capital Sentencing." *American Economic Review* 104 (11): 3397–33.

Allen, Johnie J., and Craig A. Anderson. 2017. "Aggression and Violence: Definitions and Distinctions." In Peter Sturme, ed., *The Wiley Handbook of Violence and Aggression*. Chichester: Wiley, 1–14.

Alpert, Geoffrey P., and Roger G. Dunham. 2004. *Understanding Police Use of Force: Officers, Suspects, and Reciprocity*. Cambridge: Cambridge University Press.

Ayres, Ian. 2002. "Outcome Tests of Racial Disparities in Police Practices." *Justice Research and Policy* 4 (1–2): 131–42.

Baumgartner, Frank R., Derek A. Epp, and Kelsey Shoub. 2018. *Suspect Citizens: What 20 Million Traffic Stops Tell Us about Policing and Race*. Cambridge: Cambridge University Press.

Bell, Jeannine. 2017. "Dead Canaries in the Coal Mines: The Symbolic Assailant Revisited." *Georgia State University Law Review* 34 (3): 513.

Binder, Arnold, and Peter Scharf. 1980. "The Violent Police-Citizen Encounter." *Annals of the American Academy of Political and Social Science* 452 (1): 111–21.

Bittner, Egon. 1970. *The Functions of the Police in Modern Society: A Review of Background Factors, Current Practices, and Possible Role Models*. Public Health Service publication no. 2059. Chase, MD: National Institute of Mental Health.

Brown, Michael K. 1981. *Working the Street: Police Discretion and the Dilemmas of Reform*. New York: Russell Sage.

Brown, Robert A. 2019. "Policing in American History." *Du Bois Review: Social Science Research on Race* 16 (1): 189–95.

Buehler, James W. 2017. "Racial/Ethnic Disparities in the Use of Lethal Force by US Police, 2010–2014." *American Journal of Public Health* 107 (2): 295–97.

Cesario, Joseph, David J. Johnson, and William Terrill. 2019. "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* 10 (5): 586–95.

Cinelli, Carlos, and Chad Hazlett. 2020. "Making Sense of Sensitivity: Extending Omitted Variable Bias." *Journal of the Royal Statistical Society B* 82 (1): 39–67.

Cohen, Elisha. 2021. "Sensitivity Analysis for Outcome Tests with Binary Data." Unpublished manuscript.

Cohen, Elisha, and Adam Glynn. 2021. "Estimating Bounds on Selection Bias with Outcome Tests." Unpublished manuscript.

Correll, Joshua, Bernadette Park, Charles M. Judd, and Bernd Wittenbrink. 2002. "The Police Officer's Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals." *Journal of Personality and Social Psychology* 83 (6): 1314–29.

Correll, Joshua, Bernadette Park, Charles M. Judd, Bernd Wittenbrink, Melody S. Sadler, and Tracie Keesee. 2007. "Across the Thin Blue Line: Police Officers and Racial Bias in the Decision to Shoot." *Journal of Personality and Social Psychology* 92 (6): 1006–23.

Crandall, Marie, Douglas Sharp, Erin Unger, David Straus, Karen Brasel, Renee Hsia, and Thomas Esposito. 2013. "Trauma Deserts: Distance from a Trauma Center, Transport Times, and Mortality from Gunshot Wounds in Chicago." *American Journal of Public Health* 103 (6): 1103–9.

Eberhardt, Jennifer L., Phillip Atiba Goff, Valerie J. Purdie, and Paul G. Davies. 2004. "Seeing Black: Race, Crime, and Visual Processing." *Journal of Personality and Social Psychology* 87 (6): 876–93.

Engel, Robin Shepard, and Jennifer M. Calnon. 2004. "Examining the Influence of Drivers' Characteristics during Traffic Stops with Police: Results from a National Survey." *Justice Quarterly* 21 (1): 49–90.

Engel, Robin Shepard, James J. Sobol, and Robert E. Worden. 2000. "Further Exploration of the Demeanor Hypothesis: The Interaction Effects of Suspects' Characteristics and Demeanor on Police Behavior." *Justice Quarterly* 17 (2): 235–58.

Epp, Charles R., Steven Maynard-Moody, and Donald P. Haider-Markel. 2014. *Pulled Over: How Police Stops Define Race and Citizenship*. Chicago: University of Chicago Press.

Friedrich, Robert J. 1980. "Police Use of Force: Individuals, Situations, and Organizations." *Annals of the American Academy of Political and Social Science* 452 (1): 82–97.

Fryer, Roland G., Jr. 2016. "An Empirical Analysis of Racial Differences in Police Use of Force. Technical report, National Bureau of Economic Research, Cambridge, MA.

Garner, Joel, and Christopher Maxwell. 1999. "Measuring the Amount of Force Used by and against the Police in Six Jurisdictions." Technical report, Bureau of Justice Statistics and National Institute of Justice.

Garner, Joel H., Christopher D. Maxwell, and Cedrick G. Heraux. 2002. "Characteristics Associated with the Prevalence and Severity of Force Used by the Police." *Justice Quarterly* 19 (4): 705–46.

Garner, Joel H., Thomas Schade, John Hepburn, and John Buchanan. 1995. "Measuring the Continuum of Force Used by and against the Police." *Criminal Justice Review* 20 (2): 146–68.

Gelman, Andrew, Jeffrey Fagan, and Alex Kiss. 2007. "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* 102 (479): 813–23.

Goel, Sharad, Justin M. Rao, and Ravi Shroff. 2016. "Precinct or Prejudice? Understanding Racial Disparities in New York City's Stop-and-Frisk Policy." *Annals of Applied Statistics* 10 (1): 365–94.

Goff, Philip Atiba, Tracey Lloyd, Amanda Geller, Steven Raphael Steven, and Jack Glaser. 2016. *The Science of Justice: Race, Arrests, and Police Use of Force*. Los Angeles: Center for Policing Equity.

Hine, Kelly A., Louise E. Porter, Nina J. Westera, Geoffrey P. Alpert, and Andrea Allen. 2018. "Exploring Police Use of Force Decision-Making Processes and Impairments Using a Naturalistic Decision-Making Approach." *Criminal Justice and Behavior* 45 (11): 1782–801.

Jacob, Herbert. 1972. "Contact with Government Agencies: A Preliminary Analysis of the Distribution of Government Services." *Midwest Journal of Political Science* 16 (1): 123–46.

James, Lois, Stephen James, and Bryan Vila. 2018. "Testing the Impact of Citizen Characteristics and Demeanor on Police Officer Behavior in Potentially Violent Encounters." *Policing: An International Journal* 41 (1): 24–40.

James, Lois, Bryan Vila, and Kenn Daratha. 2013. "Results from Experimental Trials Testing Participant Responses to White, Hispanic and Black Suspects in High-Fidelity Deadly Force Judgment and Decision-Making Simulations." *Journal of Experimental Criminology* 9 (2): 189–212.

Jetelina, Katelyn K., Wesley G. Jennings, Stephen A. Bishopp, Alex R. Piquero, and Jennifer M. Reingle Gonzalez. 2017. "Dissecting the Complexities of the Relationship between Police Officer–Civilian Race/Ethnicity Dyads and Less-than-Lethal Use of Force." *American Journal of Public Health* 107 (7): 1164–70.

Johnson, David J., Trevor Tress, Nicole Burkel, Carley Taylor, and Joseph Cesario. 2019. "Officer Characteristics and Racial Disparities in Fatal Officer-Involved Shootings." *Proceedings of the National Academy of Sciences* 116 (32): 15877–82.

Johnson, Richard R. 2011. "Suspect Mental Disorder and Police Use of Force." *Criminal Justice and Behavior* 38 (2): 127–45.

Kahn, Kimberly Barsamian, Phillip Atiba Goff, J. Katherine Lee, and Diane Motamed. 2016. "Protecting Whiteness: White Phenotypic Racial Stereotypicality Reduces Police Use of Force." *Social Psychological and Personality Science* 7 (5): 403–11.

Kahn, Kimberly Barsamian, Joel S. Steele, Jean M. McMahon, and Greg Stewart. 2017. "How Suspect Race Affects Police Use of Force in an Interaction over Time." *Law and Human Behavior* 41 (2): 117–26.

Kavanagh, John. 1997. "The Occurrence of Resisting Arrest in Arrest Encounters: A Study of Police–Citizen Violence." *Criminal Justice Review* 22 (1): 16–33.

Kleider, Heather M., Dominic J. Parrott, and Tricia Z. King. 2010. "Shooting Behaviour: How Working Memory and Negative Emotionality Influence Police Officer Shoot Decisions." *Applied Cognitive Psychology* 24 (5): 707–17.

Knowles, John, Nicola Persico, and Petra Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy* 109 (1): 203–29.

Knox, Dean, Will Lowe, and Jonathan Mummolo. 2020. "Administrative Records Mask Racially Biased Policing." *American Political Science Review* 114 (3): 619–37.

Knox, Dean, and Jonathan Mummolo. 2020. "Making Inferences about Racial Disparities in Police Violence." *Proceedings of the National Academy of Sciences* 117:1261–62.

Lawson, Edward, Jr. 2019. "TRENDS: Police Militarization and the Use of Lethal Force." *Political Research Quarterly* 72 (1): 177–89.

Lawton, Brian A. 2007. "Levels of Nonlethal Force: An Examination of Individual, Situational, and Contextual Factors." *Journal of Research in Crime and Delinquency* 44 (2): 163–84.

Lipsky, Michael. 1980. *Street-Level Bureaucracy: Dilemmas of the Individual in Public Service*. New York: Russell Sage.

Matrofski, Stephen D., Jeffrey B. Snipes, and Anne E. Supina. 1996. "Compliance on Demand: The Public's Response to Specific Police Requests." *Journal of Research in Crime and Delinquency* 33 (3): 269–305.

McCluskey, John D., and William Terrill. 2005. "Departmental and Citizen Complaints as Predictors of Police Coercion." *Policing: An International Journal of Police Strategies and Management* 28 (3): 513–29.

McCluskey, John D., William Terrill, and Eugene A. Paoline III. 2005. "Peer Group Aggressiveness and the Use of Coercion in Police–Suspect Encounters." *Police Practice and Research* 6 (1): 19–37.

McElvain, James P., and Augustine J. Kposowa. 2008. "Police Officer Characteristics and the Likelihood of Using Deadly Force." *Criminal Justice and Behavior* 35 (4): 505–21.

Mekawi, Yara, and Konrad Bresin. 2015. "Is the Evidence from Racial Bias Shooting Task Studies a Smoking Gun? Results from a Meta-analysis." *Journal of Experimental Social Psychology* 61:120–30.

Monk, Ellis P. 2019. "The Color of Punishment: African Americans, Skin Tone, and the Criminal Justice System." *Ethnic and Racial Studies* 42 (10): 1593–612.

Mummolo, Jonathan. 2018. "Modern Police Tactics, Police–Citizen Interactions, and the Prospects for Reform." *Journal of Politics* 80 (1): 1–15.

Nieuwenhuys, Arne, Geert J. P. Savelsbergh, and Raoul R. D. Oudejans. 2012. "Shoot or Don't Shoot? Why Police Officers Are More Inclined to Shoot When They Are Anxious." *Emotion* 12 (4): 827–33.

Nix, Justin, Justin T. Pickett, and Renée J. Mitchell. 2019. "Compliance, Noncompliance, and the In-Between: Causal Effects of Civilian Demeanor

on Police Officers' Cognitions and Emotions." *Journal of Experimental Criminology* 15 (4): 611–39.

Nix, Justin, Justin T. Pickett, Scott E. Wolfe, and Bradley A. Campbell. 2017. "Demeanor, Race, and Police Perceptions of Procedural Justice: Evidence from Two Randomized Experiments." *Justice Quarterly* 34 (7): 1154–83.

Oberfield, Zachary W. 2012. "Socialization and Self-Selection: How Police Officers Develop Their Views about Using Force." *Administration and Society* 44 (6): 702–30.

Persico, Nicola, and Petra Todd. 2006. "Generalising the Hit Rates Test for Racial Bias in Law Enforcement, with an Application to Vehicle Searches in Wichita." *Economic Journal* 116 (515): F351–F367.

Pierson, Emma, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Vignesh Ramachandran, Cheryl Phillips, and Sharad Goel. 2017. "A Large-Scale Analysis of Racial Disparities in Police Stops across the United States." Stanford University Working paper.

Pierson, Emma, et al. 2020. "A Large-Scale Analysis of Racial Disparities in Police Stops across the United States." *Nature Human Behaviour* 4 (7): 736–45.

Prowse, Gwen, Vesla M. Weaver, and Tracey L. Meares. 2020. "The State from Below: Distorted Responsiveness in Policed Communities." *Urban Affairs Review* 56 (5): 1423–71.

Ross, Cody T. 2015. "A Multi-Level Bayesian Analysis of Racial Bias in Police Shootings at the County-Level in the United States, 2011–2014." *PLoS One* 10 (11): e0141854.

Sances, Michael W., and Hye Young You. 2017. "Who Pays for Government? Descriptive Representation and Exploitative Revenue Sources." *Journal of Politics* 79 (3): 1090–94.

Schuck, Amie M. 2004. "The Masking of Racial and Ethnic Disparity in Police Use of Physical Force: The Effects of Gender and Custody Status." *Journal of Criminal Justice* 32 (6): 557–64.

Sierra-Arévalo, Michael. 2021. "American Policing and the Danger Imperative." *Law and Society Review* 55 (1): 70–103.

Sikora, Andrew G., and Michael Mulvihill. 2002. "Trends in Mortality due to Legal Intervention in the United States, 1979 through 1997." *American Journal of Public Health* 92 (5): 841–43.

Soss, Joe, and Vesla Weaver. 2016. "Learning from Ferguson." In Alvin B. Tillery and Juliet Hooker, eds., *The Double Bind: The Politics of Racial and Class Inequalities in the Americas*. Report of the APSA Task Force on Racial and Class Inequalities in the Americas. Washington, DC: American Political Science Association.

Soss, Joe, and Vesla Weaver. 2017. "Police Are Our Government: Politics, Political Science, and the Policing of Race-Class Subjugated Communities." *Annual Review of Political Science* 20:565–91.

Sun, Ivan Y., and Brian K. Payne. 2004. "Racial Differences in Resolving Conflicts: A Comparison between Black and White Police Officers." *Crime and Delinquency* 50 (4): 516–41.

Sun, Ivan Y., Brian K. Payne, and Yuning Wu. 2008. "The Impact of Situational Factors, Officer Characteristics, and Neighborhood Context on Police Behavior: A Multilevel Analysis." *Journal of Criminal Justice* 36 (1): 22–32.

Terrill, William. 2005. "Police Use of Force: A Transactional Approach." *Justice Quarterly* 22 (1): 107–38.

Terrill, William. 2011. "Police Coercion: Application of the Force Continuum." PhD diss.

VanderWeele, Tyler J., and Peng Ding. 2017. "Sensitivity Analysis in Observational Research: Introducing the E-Value." *Annals of Internal Medicine* 167 (4): 268–74.

Voigt, Rob, Nicholas P. Camp, Vinodkumar Prabhakaran, William L. Hamilton, Rebecca C. Hetey, Camilla M. Griffiths, David Jurgens, Dan Jurafsky, and Jennifer L. Eberhardt. 2017. "Language from Police Body Camera Footage Shows Racial Disparities in Officer Respect." *Proceedings of the National Academy of Sciences* 114 (25): 6521–26.

Welch, Kelly. 2007. "Black Criminal Stereotypes and Racial Profiling." *Journal of Contemporary Criminal Justice* 23 (3): 276–88.

Wheeler, Andrew P., Scott W. Phillips, John L. Worrall, and Stephen A. Bishopp. 2017. "What Factors Influence an Officer's Decision to Shoot? The Promise and Limitations of Using Public Data." *Justice Research and Policy* 18 (1): 48–76.

Wilson, James Q. 1978. *Varieties of Police Behavior*. Cambridge, MA: Harvard University Press.

Worden, Robert E. 2015. "The 'Causes' of Police Brutality: Theory and Evidence on Police Use of Force." *Criminal Justice Theory: Explaining The Nature and Behavior of Criminal Justice* 2:149–204.

Worrall, John L., Stephen A. Bishopp, Scott C. Zinser, Andrew P. Wheeler, and Scott W. Phillips. 2018. "Exploring Bias in Police Shooting Decisions with Real Shoot/Don't Shoot Cases." *Crime and Delinquency* 64 (9): 1171–92.

Zimring, Franklin E. 2017. *When Police Kill*. Cambridge, MA: Harvard University Press.