

# Systemic Racism and Policing

Asya Magazinnik\*

October 27, 2022

## Abstract

Concerns about racial bias in American policing have generated enormous scholarly literatures in political science, sociology, economics, statistics, criminology, social psychology, and law. The purpose of this review is to summarize the methodological and substantive knowledge these literatures have accumulated over the past two decades. Topics covered include how race shapes micro-level encounters between civilians and the police; the intersection between law enforcement and structural racism, including residential segregation and income inequality; the role of policing in a democratic society; and promising approaches to reduce racial harms. A key takeaway from this exercise is that the empirical study of racial bias in policing is not value-neutral: any approach requires the researcher to furnish a benchmark of no discrimination, inscribing (often unstated) standards of justice and fairness into the methodology. I propose avenues for future research for computational social scientists in light of my findings.

---

Prepared for the IDSS Systemic Racism and Computation Workshop at MIT. This is a work in progress; please do not cite without permission. Comments welcome. Natalie Li provided invaluable research assistance.

\*Assistant Professor, Department of Political Science, MIT. [asyam@mit.edu](mailto:asyam@mit.edu)

# 1 Overview

Concerns over racial bias in American policing have generated enormous scholarly literatures in political science, sociology, economics, statistics, criminology, social psychology, and law. The purpose of this review is to distill the knowledge these literatures have accumulated over the past two decades for scholars embarking on original empirical research on this subject — with the statistical tools, computational resources, and data-collection capabilities they now have at their disposal.

I begin at the micro level: with an encounter between a civilian and a police officer. Detecting racial bias in this basic state-citizen interaction has commanded significant scholarly attention. But I begin here not because of the staggering volume of work produced on the subject, but because the debates in this literature will help clarify how many social scientists have come to understand precisely what racial bias means. The focus of Section 2 is not on findings but on methodology, as the study of racial bias in policing to date has essentially been the study of *how* to study racial bias in policing. While the discussion can get somewhat technical, it never escapes being deeply normative, as any notion of bias is defined in reference to some ideal standard of justice or equity. Methodologically, this section covers economic notions of statistical discrimination and taste-based prejudice, benchmarking, outcome tests, and the potential outcomes framework. Substantively, it covers police stops of motorists and pedestrians as well as arrests, incidents of police violence, and shootings of civilians.

While the police-civilian encounter is a good place to start thinking about racial bias, it is not where the inquiry ends. Officers' behavior is conditioned not just by the characteristics of the civilians they interact with, but by the neighborhoods and social contexts in which these interactions arise. By the same token, civilians' experiences of law enforcement are more than the sum of the interactions they have had with officers; they are also shaped by the total police presence in their communities, as well as the relative presence compared to other communities; by the role the police play in civilians' lives and the lives of their neighbors;

and by narratives about the police in the public discourse. Bias arises not just from micro-level contact but from macro-level allocations of resources: between some communities and others, between punitive functions and social services. Thus, Section 3 reviews the literature on structural racism in policing, situating the police-civilian interaction in a set of larger overlapping systems including residential segregation and income inequality.

Bringing together insights from Sections 2 and 3, Section 4 discusses how racial bias in policing from both micro- and macro-level perspectives constructs the relationship between the citizen and the state. I review recent research on how experiences of racial bias in law enforcement decrease political participation, depress civic engagement, and erode perceptions of legitimacy and trust in government. Finally, Section 5 reviews what we currently know about avenues for reform: how proposed approaches such as increasing racial diversity on police forces, proactive and community policing models, and predictive policing moderate racial bias.

Three caveats are in order before proceeding. The first concerns the racial categories relevant to this study. While some perspectives emphasize discrimination against all people of color and the need for cross-racial solidarity in conversations about policing, others point to a history of state-sanctioned violence against Black people in particular, motivating the need for focused discussions about anti-Black bias (Kim 2020). Fortunately, the subject of this review is sufficiently general that I can respect both views. Many of the studies and issues covered here relate specifically to anti-Black bias, and when this is the case I use the term “Black” accordingly. Other times, I use “minority” and “majority” or “white,” including in more theoretical discussions that are in principle generalizable to any dominant/subordinate racial group dichotomy. That said, I do not reach for generalizability beyond the American context. While some of the methods and theoretical constructs covered here are surely transportable to other places, my focus is on the specific institution of American policing and the specific racial dynamics of the United States in the twenty-first century.

A second caveat relates to the scope of this review. A literature on the “carceral state”

situates policing in a larger set of overlapping and mutually reinforcing institutions that collude to disproportionately punish Black citizens, including the prison and parole systems, the courts, and even punitive social programs that rely on various forms of surveillance and supervision (Hinton 2016). I wholeheartedly agree with this view, and recognize that analyzing policing in isolation from this larger ecosystem is bound to overlook important insights. That said, a narrower focus on policing — by which I mean interactions between front-line officers and civilians as well as between police forces and the neighborhoods they serve — allows for a deeper discussion of the relevant theoretical and methodological issues, some of which are unique to this institution and would not get sufficient coverage if I broadened my scope. This review thus consciously excludes research about incarceration (beyond initial arrest), about the courts, and about crime as an outcome. Other recent work has risen to the challenge of reviewing research on the carceral state as a whole (Gottschalk 2016; Lerman and Weaver 2016).

Finally, this is a survey of peer-reviewed academic books and journal articles across the social sciences, predominantly published by American journals and academic presses. As such, it reflects any racial bias of the context that produced this work. While the search for rigorous, scientific evidence undoubtedly motivates many of the scholars represented here, my review will demonstrate repeatedly that, in this literature, methodologies are normative commitments. People of color, people from low-income communities, and people who have first-hand experience with the coercive arm of the state are not well-represented in academia, so their views have not informed academic debates as much as they should. I have made an effort to include a broad range of work by Black scholars in this review, while faithfully conveying ideas that have emerged from white-dominated fields. With these cards on the table, let us begin.

## 2 Racial Bias in Police-Civilian Interactions

### 2.1 What Constitutes Racial Bias and How Do We Detect It?

What does it mean for the police to be biased in their treatment of racial minorities, and what are the observable implications of racial bias? Establishing a theoretical definition of racial discrimination and agreeing on a set of measurable quantities that can directly speak to this construct has been a major challenge for social scientists studying policing, but a necessary first step for progress. I therefore begin by reviewing the major perspectives on police bias in the social science literature from the past two decades, alongside the empirical methodologies implied by each view. To fix ideas, and to delay the introduction of further methodological complications until the basic theoretical framework is in place, I focus in this discussion on *police stops*: the first point of contact between civilians and the police, and the necessary condition for many other important outcomes of police-civilian interaction such as arrests, incidents of use of force, and shootings.

An obvious starting point for detecting bias is to note racial disparities in the probability of being stopped by the police. Of course, such disparities alone do not constitute sufficient evidence of police bias. As a consequence of pervasive racial bias across a range of other social institutions (for instance, racism in education, employment, housing, health care, and wealth accrurement), crime is more prevalent in communities where racial minorities live. What is more, front-line police officers are often tasked with not only law enforcement responsibilities but basic social service provision (Beck and Goldstein 2018), which could give rise to significant racial disparities in police contact even in the absence of racial bias *within the police force*. In short, extrapolating racial discrimination from disparities in police contact is prone to a form of *omitted variable bias*: the failure to observe important factors correlated with both race and the probability of contact with the police.

What analysts seek, then, is a hypothetical *benchmark* against which police bias can be detected: given the socioeconomic and demographic facts on the ground, what are the

patterns of interaction we would expect to observe between an *unbiased* police force and Black citizens? White citizens? Plausible benchmarks — considered the “holy grail” in the economics, statistics, and criminology literatures — are extremely difficult to come by. Attempts (reviewed more thoroughly below, and in Engel and Calnon (2004), Ridgeway and MacDonald (2010), and Goff and Kahn (2012)) have generally fallen into one of two categories. The first approach seeks to approximate the racial composition of the population available for police stops, using either Census data for pedestrian stops or more sophisticated methods of capturing the local driving population for traffic stops. However, this first approach does not fully address the omitted variable problem highlighted above. To approximate the racial distribution of the “at-risk” population — those who would justifiably be stopped by the police in an unbiased system — a second approach utilizes local arrest data or crime rates (see, e.g., Gelman, Fagan, and Kiss (2007)). However, the validity of this approach crucially depends on crime or arrest data providing a neutral benchmark of criminal behavior. Since a stop is usually a precursor to an arrest, arrest records will reflect any bias in first contact with the police, rendering such benchmarks logically and statistically suspect (Goff and Kahn 2012; Knox and Mummolo 2020b).

An alternative way forward that avoids the benchmarking problem entirely is the *outcome test* (Becker 1957), which has been applied broadly in the literature on discrimination in economics, from lending and hiring decisions to law enforcement. Rather than focusing on the rate at which decisions are made, the outcome test assesses the success rate of those decisions. For instance, to find evidence of discrimination in lending decisions, Becker (1957) argued, one only needs to observe subsequent repayment rates. Even if there are underlying racial differences in borrowers’ ability to repay their loans — unobservable to the analyst, but observed by the bank — “non-discriminatory” lending should imply that the marginal minority borrower — the one just over the threshold for being approved for a loan — successfully repays their loan at the same rate as the marginal white borrower. By contrast, if the marginal minority loan has a *higher* repayment rate, then lenders are discriminating

against minority borrowers by making them clear a higher bar of creditworthiness.

If the logic of the outcome test is sound, then the seemingly intractable problem of detecting racial bias in police stops is reduced to the simple task of comparing white and minority “hit rates,” broadly defined as the proportion of police stops that turns up contraband or leads to an arrest. By similar logic, if searches of minority drivers’ vehicles turn up contraband less often than for white drivers, then it must take less for a minority driver to arouse suspicion, indicating racial prejudice on the part of officers.<sup>1</sup> Thus, a second major strand of the literature has departed from the search for benchmarks and devoted itself instead to outcome tests: developing the behavioral theories and assumptions under which they are valid, and applying them to data.

In Section 2.2.2 below, I review the methodological issues around the rigorous application of outcome tests to policing data. Severe inferential challenges arise from the *infra-marginality problem*: the fact that analysts can usually only observe *average* hit rates by race where what they need to know are the *marginal* rates, and the two quantities are not guaranteed to be equal (Ayres 2002). In an influential study, Knowles, Persico, and Todd (2001) (henceforth KPT) proposed a game-theoretic model of police-citizen interaction in which average and marginal hit rates are the same in equilibrium, meaning the infra-marginality problem does not arise. A sizable literature following KPT clarified the behavioral assumptions about police-citizen interactions needed to generate this equilibrium, rendering racial differences in average hit rates interpretable as evidence of racial bias.

However, even in the best case when these rather strenuous assumptions are satisfied, outcome tests comparing hit rates across racial groups can only speak to a narrow definition of discrimination widely accepted in economics but contested by criminologists and legal scholars. The economic theory of discrimination from which the outcome test arises distinguishes between *statistical discrimination* and *taste-based prejudice* (Becker 1957; Arrow

---

<sup>1</sup>A critical and heroic assumption here is that, even if there is racial bias in initial stops, finding contraband or making arrests is a mechanical consequence of the guilt or innocence of the driver and not itself a process prone to racial bias.

1972, 1973, 1998; Phelps 1972). The theory takes as a starting point that police seek to maximize searches that uncover crime, i.e. turn up contraband. If the police know that Black drivers are more likely than white drivers to carry contraband, then *statistical discrimination* results from officers stopping more Black drivers in order to maximize their probability of finding contraband, whereas *taste-based prejudice* is any additional, direct utility they may derive from policing Black people. This view regards statistical discrimination as tolerable in the service of the assumed goals of law enforcement, and indeed frames the detection of racial bias through the use of outcome tests as the search for taste-based prejudice unconfounded by statistical discrimination (Borooah 2001; Knowles, Persico, and Todd 2001; Anwar and Fang 2006).

Critical responses take issue with a theory of discrimination derived from maximizing “successful” searches rather than minimizing the social costs of crime — and of policing on the profiled community. Some scholars are primarily concerned with ensuring procedural equity and protecting civil liberties during police-citizen encounters (Calderon 2000; Gross and Barnes 2002); others emphasize the burden that coercive law enforcement actions impose on citizens and define an “unbiased” system as one in which those costs are equally distributed across racial groups (Harcourt 2004; Engel 2008). Such costs include but are not limited to “stigmatiz[ing] the group, leading others to view members in the group with suspicion or fear, and... brand[ing] the group as a bunch of criminals, thus reinforcing offensive stereotypes and perpetuating racist ideology” (Shelby 2018). Neither view accepts statistical discrimination as a neutral consequence of the pursuit of economic efficiency, and both reject methodologies that can only detect bias born of personal animus. Rather, Tomaskovic-Devey, Mason, and Zingraff (2004) argue for a more expansive definition of racial bias in policing that includes both taste-based and statistical discrimination, but also unconscious cognitive biases and cultural stereotypes as well as police deployment patterns.

Recognizing the limitations of both benchmarking approaches and outcome tests — and moving beyond distinctions between sources of bias — a burgeoning literature in political



science and statistics has built on the potential outcomes framework (Rubin 1974) to frame racial bias in policing as a causal inference problem. Taking a police-civilian interaction as the basic unit of analysis, Knox, Lowe, and Mummolo (2020) define the theoretical quantity of interest as the hypothetical difference in the encounter that would arise if one were to change the race of the civilian from Black to white (or vice-versa), all else equal. The thought exercise of manipulating a civilian’s racial identity poses some serious philosophical challenges, and remains deeply contested in the causal inference literature (Holland 1986; Hernán 2016; Sen and Wasow 2016; Pearl 2018). To some extent, focusing on the interaction rather than the civilian circumvents this debate. Here, “the manipulation of race is conceptualized as the counterfactual substitution of an individual with a different racial identity into the encounter, while holding the encounter’s objective context — location, time of day, criminal activity, etc. — fixed” (Knox, Lowe, and Mummolo 2020). Building on Sen and Wasow’s (2016) formulation of race as a “bundle of sticks,” Knox et al. conceive of the racial manipulation as potentially encapsulating multiple characteristics — for example, skin tone, dialect, and clothing — but leave the specific operationalization in a given context to the analyst. The payoff is the ability to rigorously define a *causal racial effect* (absorbing statistical discrimination, conscious prejudice, and cognitive bias) on any number of outcomes, from a civilian’s propensity to be stopped by the police to their risk of experiencing police violence. I defer a full discussion of the causal inference approach to Section 2.3, in which I review the research on outcomes beyond police stops such as arrests and use of force, as that is where it provides the greatest analytical leverage.

## 2.2 Police Stops: Measurement Strategies and Key Findings

### 2.2.1 Benchmarking Approaches

The first stage of many police-civilian interactions is an officer’s decision to pull over a car on the road or to stop and search a pedestrian. A large empirical literature has been interested in whether police are biased against minority drivers and pedestrians in making this initial

decision, where bias is broadly defined as the deviation in observed stops by race from some fair or neutral benchmarks. Properly specifying the appropriate benchmark is the central challenge of this literature (Goff and Kahn 2012), and consequently it has not yet, as a whole, produced conclusive findings on the presence or extent of racial bias in police stops.

A survey by Engel and Calnon (2004) highlights six benchmarking strategies for traffic stops: census data, observations of roadway usage, official accident data, assessments of traffic-violating behavior, citizen surveys, and internal departmental comparisons. Because the residential population of an area does not necessarily reflect the population that drives through it, approaches relying exclusively on census data have fallen out of favor; however, creative approaches to estimate the racial distribution of drivers or traffic violators include using non-responsible drivers involved in two-vehicle crashes (Alpert, Smith, and Dunham 2004) and data from police early warning systems (Walker 2001), as well as surveying drivers about their traffic violation patterns (Lange, Johnson, and Voas 2005). Internal benchmarking, which compares officers to other officers who patrol the same place at the same time of day (see, for example, Ridgeway and MacDonald (2009)), is useful for detecting outliers but cannot speak to overall levels of racial bias on a police force.

An additional benchmarking exercise involves the “veil of darkness,” or the hypothesis that police are less likely to observe a motorist’s race before stopping them at night than during the day. Assuming that racial differences in traffic patterns, driving behavior, and exposure to law enforcement do not vary by time of day, Grogger and Ridgeway (2006) compare the racial distribution of traffic stops during the day to that at night, with the latter serving as a race-blind benchmark. Applying this method to traffic stop data from Oakland, California, Grogger and Ridgeway (2006) do not find evidence of racial profiling. However, Horrace and Rohlin (2016) argue that urban areas like Oakland are well-lit at night, eroding the power of the “veil of darkness” test. When they refine the test using streetlight location data from Syracuse, New York, they find evidence of racial profiling of Black drivers, a finding supported in a large-scale nationwide study by Pierson et al. (2020).

Still, even carefully designed benchmark tests can fail if deployment patterns lead police to encounter racial minorities more often than white civilians. As Knox and Mummolo (2020*b*) point out, if Black and white civilians exhibit the same behavior, and police stop Black and white civilians at the same rate, but they encounter more Black than white civilians, then even a test that characterizes the correct benchmark distributions will detect bias where there is none. To account for deployment patterns — and as much additional local information as possible — Gelman, Fagan, and Kiss (2007), analyzing New York City’s stop-and-frisk program, use a hierarchical Bayesian model that controls for precinct-level variability in stops, past arrests, and demographic characteristics. They find that Black and Hispanic people were more likely to be stopped by police than whites under the program even after adjusting for local variation in crime and police deployment.

Although the benchmarking literature has become more sophisticated over the past two decades, it still finds itself between a rock and a hard place. On the one hand, the closer one gets to approximating the racial distribution of the “at-risk” population, the more believable the inference from this research design will be. But outside of clever but imperfect designs such as the “veil of darkness” test, it is difficult to construct good benchmarks without crime rates, and crime statistics will necessarily reflect any racial bias in police stops or in the process by which stops turn into arrests. One fruitful way forward, proposed by Knox and Mummolo (2020*b*), is a sensitivity test to help analysts evaluate just how erroneous their benchmark distribution can be and still yield an estimate of racial bias in the same direction. This sensitivity test is embedded in a causal inference framework that unites benchmark tests and outcome tests under a common interpretation, a contribution I will revisit in Section 2.3.

### **2.2.2 Outcome Tests and the Infra-marginality Problem**

I turn now to a second major strand of the empirical social science literature on racial bias in police stops: studies based on an “outcome test” comparing hit rates across racial groups.

While the proceeding discussion is somewhat technical, I devote some close attention to the methodological issues because the stakes of the debate are high. If the logic behind the outcome test in the context of policing is sound, then simple comparisons between hit rates by race are at the very least informative about one form of racial bias, “taste-based discrimination,” and a conceptually complicated construct can practically be measured with a simple statistical test and readily available data.<sup>2</sup> On the other hand, if the assumptions required are too unrealistic, then a wholesale reevaluation of what we know about bias in policing is in order, given the large number of studies whose findings are based on some form of outcome test. In this case, scholars should invest in developing rigorous theoretical and methodological frameworks for studying bias in policing — which would ideally be common across studies (Knox and Mummolo 2020*b*) — before conducting any further empirical research. On the whole, criminology experts — and even some economists — have come down on the side of the latter view: outcome tests require assumptions about police and citizen decisionmaking that are unrealistic for the vast majority of law enforcement interactions (see Engel and Tillyer (2008) for a summary of these arguments).

To give some intuition for the outcome test and the associated infra-marginality problem, I present a concrete example from Simoiu, Corbett-Davies, and Goel (2017). Suppose officers seek to maximize the contraband they uncover from traffic stops, and that Black and white drivers have different probabilities of carrying contraband. Further suppose that there are two readily observable (to officers) “types” in each racial category: for white drivers, those who have a 1% chance of carrying contraband and those who have a 75% chance; for Black drivers, those who have a 1% chance of carrying contraband and those who have a 50% chance. If officers have a race-neutral policy of stopping any individual with more than a 10% probability of carrying contraband, then the “success rate” of searches (the hit rate)

---

<sup>2</sup>However, I emphasize again that even valid measures of “taste-based discrimination” do not give a complete picture of racial bias and that “statistical discrimination” — the quantity that the economics literature on discrimination seeks to purge from its estimates — is equally relevant for conversations about civil liberties, state-citizen relations and overpolicing, and the well-documented adverse physical, psychological, and political consequences of police contact for racial minorities.

will be 75% for white drivers and 50% for Black drivers. This infra-marginality problem severely limits the rigorous application of the outcome test for racial bias, particularly when the decision or the outcome is dichotomous (Simoiu, Corbett-Davies, and Goel 2017).

An influential paper by Knowles, Persico, and Todd (2001) (KPT) provides a theoretically grounded solution by formulating police stops as a strategic interaction between officers seeking to maximize the success rate of searches and motorists deciding whether to carry contraband. In the Nash equilibrium of this game-theoretic model — where citizens and officers, having full information about each others’ strategies and utility functions, choose their own actions as a best response to the other players — all motorists end up carrying contraband with equal probability, thus getting around the infra-marginality problem. KPT’s application of the outcome test to vehicle searches along a strip of I-95 in Maryland yields no evidence of “taste-based” discrimination against Black motorists, although a subsequent replication on all Maryland searches using data that became available after KPT’s original analysis does detect evidence of discrimination against Black and Hispanic motorists (Sanga 2009). (See also Gross and Barnes (2002) for a discussion of why the Maryland data is idiosyncratic and unrepresentative of police stops in general.)

A sizable literature has interrogated the robustness of the KPT model — specifically its key prediction that motorists carry contraband with equal probabilities in equilibrium — to alternative assumptions, with the goal of understanding the conditions under which the outcome test remains valid. The necessary equilibrium holds up to a few modifications, including allowing for heterogeneity among police officers in their degree of racial bias and their search costs, as well as heterogeneity among motorists in their costs and benefits of carrying contraband (Persico and Todd 2006). Hernández-Murillo and Knowles (2004) develop an adapted outcome test based on a model that incorporates both discretionary and non-discretionary searches, given that in practice the majority of searches are not fully discretionary as assumed under the KPT model (Engel 2008).

However, relaxing KPT’s assumptions in various reasonable ways wholly or partially in-

validates the outcome test, leading some scholars to conclude that it is simply inappropriate for the policing context (Engel 2008). Such adaptations include relaxing the requirement that police observe all motorists (i.e., that every motorist has some nonzero probability of detection), introducing variation in the severity of the offense (rather than a carry/no carry dichotomy) (Dharmapala and Ross 2004), and allowing the police to observe a signal of guilt that depends on whether the individual is carrying contraband (Bjerk 2007). Importantly, the validity of the test also relies on the officer’s assumed objective function; in the KPT model, the police seek to maximize searches that turn up contraband, an assumption that has raised both practical and normative objections (Durlauf 2006; Harcourt 2004; Manski 2006). If the police are motivated by minimizing total crime (the number of drivers carrying contraband) or unpunished crime (carriers of contraband going undetected), then more stringent — and difficult to verify — assumptions are required (Dominitz and Knowles 2006).

Because the validity of the outcome test is so contested in the context of policing, empirical studies based on this methodology should be taken with some degree of caution. Moreover, findings from this literature are mixed. While some detect evidence of statistical but not taste-based discrimination against racial minorities (Knowles, Persico, and Todd 2001; Persico and Todd 2006; Horn, McCluskey, and Mittelhammer 2014), others find evidence of both kinds (Hernández-Murillo and Knowles 2004; Close and Mason 2007). Several empirical applications of the KPT framework test for differences in hit rates by officer race in addition to civilian race, arguing that such differences should not arise from purely statistical discrimination (Antonovics and Knight 2009; Anwar and Fang 2006; Anbarci and Lee 2014). Again, their empirical findings are contradictory. Analyzing traffic stops by the Boston Police, Antonovics and Knight (2009) find that searches are more likely when the race of the motorist differs from the race of the officer. Also looking at Boston Police, Anbarci and Lee (2014) find that minority officers are harsher than white officers in issuing speeding tickets, but especially toward minority drivers. Anwar and Fang (2006), looking at Florida State Highway Patrol, find that white officers search more motorists than Black officers, but they

fail to detect evidence of relative racial prejudice between groups.<sup>3</sup>

A promising recent innovation that unites some of these disparate findings, offered by Knox, Lowe, and Mummolo (2020), is to show that under some far less strenuous assumptions, the outcome test yields a *lower bound* on a well-defined measure of racial discrimination, albeit a broader one than the “taste-based” form that concerns the economics literature: “the decrease in detainment that would counterfactually result if white civilians were substituted into the circumstances where minority stops occurred” (Knox and Mummolo (2020*b*), p. 29). I now turn to a discussion of the theoretical framework that makes this interpretation possible.

## 2.3 Arrests, Use of Force, and Police Shootings

Given evidence of raced-based selection into the first encounter with the police, detecting racial bias in outcomes that *ensue* from stops — arrests, use of force, or lethal encounters — is fraught with even greater methodological difficulty. As Knox and Mummolo (2020*b*) point out, before the usual inferential challenges discussed in Section 2.2 can even arise, the foremost task of the researcher is to precisely state the question. For instance, one might be interested in a Black civilian’s likelihood of being shot by the police *conditional on being stopped* compared to that of a white civilian. This is an important and policy-relevant question, but to properly understand police *bias* one has to contend with the fact that if there was racial discrimination in the initial stop decision, then the populations of Black and white stopped citizens are fundamentally incomparable. Indeed, for any related research question one might ask, the first-stage selection process that leads to an officer pointing a gun at a white or Black civilian is of utmost relevance.

To greatly elucidate the exposition of extant approaches to dealing with (or ignoring) this multi-stage selection problem, I adopt the potential outcomes framework introduced by Rubin (1974) and applied to the study of policing by Knox and Mummolo (2020*b*). To start,

---

<sup>3</sup>However, such approaches are incapable of detecting bias that is common to both groups.

take as the unit of analysis a police-civilian encounter  $i$ , in which a police officer simply observes — and has not yet engaged — a citizen. Then let  $D_i \in \{0, 1\}$  represent whether a civilian is white ( $D_i = 0$ ) or a racial minority ( $D_i = 1$ ). Further, let an individual’s minority status determine whether they are stopped and engaged by the police, denoted by the potential outcome notation  $M_i(D_i) \in \{0, 1\}$ . Since we are holding everything about the encounter other than race fixed,  $M_i(0) = M_i(1)$  implies that civilian race has no causal effect on police stopping decisions. Finally, define an observed outcome such as an arrest or shooting  $Y_i(D_i, M_i(D_i))$ , which is both a direct function of race and a function of whether there was a stop, which is also determined by race.

This simple framework lends itself to the formulation of clear causal questions, and provides a starting point for deriving the conditions under which the answers are identifiable from data. For instance, a causal formulation of the question posed in Section 2.2 is  $ATE^M \equiv \mathbb{E}[M_i(1) - M_i(0)]$ , the average treatment effect of civilian minority status (as compared to being white) on the probability of being stopped by the police, holding everything else about the encounter constant. In this section, we are interested in statements about  $Y$ , many of which are in principle possible. A quantity that is both interesting and tractable is the *average treatment effect among stops*,  $ATE_{M=1}^Y \equiv \mathbb{E}[Y_i(1, M_i(1)) - Y_i(0, M_i(0)) | M_i = 1]$ , which answers the question: “among the subset of encounters that led to detainment ( $M_i = 1$ ), and thus have some record of occurring, what is the hypothetical average difference in police behavior if minority civilians had been present, as opposed to white civilians?” (Knox and Mummolo 2020b). An interesting but practically challenging alternative is the overall *average treatment effect*,  $ATE^Y \equiv \mathbb{E}[Y_i(1, M_i(1)) - Y_i(0, M_i(0))]$ , which gives the expected difference in outcomes if minorities were inserted into every encounter versus white citizens. The specification of the causal quantity of interest determines the *type of data* that is needed to answer the question (all encounters? encounters that culminate in a stop? incidents of the use of force?) along with the *assumptions for identification* of the desired causal effect, i.e., what potential confounders the researcher should adjust for to get unbiased estimates



(Elwert and Winship 2014).

Another important contribution of this framework is to highlight the enormous gap between the data one would ideally observe to answer most causal questions of interest and the administrative records that are generally available to researchers. Without imposing further assumptions, identifying  $ATE^Y$  from data requires that the researcher observe *every time a police officer encounters a civilian* and whether or not they engage in any contact. Of course, barring innovations such as data from body-worn cameras, most such encounters produce no administrative trace. Thus, the proceeding discussion lays out the different approaches scholars have explicitly — or, more often, implicitly — used to bridge the gap between their data and their causal questions. Following Knox and Mummolo (2020b), I cover three common approaches in the empirical literature: (1) predicting civilian race from a dataset with only positive instances of the use of force; (2) assuming away racial bias in stops; and (3) benchmarking. The first approach — making inferences about racial bias in the use of force using data on only positive instances of the use of force ( $Y_i = 1$ ) — is a case of selection on the dependent variable and is invalid for the desired causal quantities. However, with data on positive instances of stops ( $M_i = 1$ ) and variation in the use of force, researchers can make progress either by assuming away bias in stops (2) — a highly implausible strategy — or through the use of benchmarking strategies (3).

In the service of generalizability and national scope, researchers have examined datasets on police shootings and killings across the U.S. for insight into how characteristics of the encounter — including the race of the civilian — affect the deadly use of force (Ross 2015; Nix et al. 2017; Menifield, Shin, and Strother 2019; Streeter 2019). A recent prominent study using this approach (Johnson et al. 2019) that was subsequently retracted (Johnson et al. 2020) reported finding no racial disparity in “the degree to which Black civilians are more likely to be fatally shot than white civilians” (Johnson et al. (2019), p. 15877) and, “if anything, found anti-white disparities” (Ibid., p. 15880). However, this statement about  $ATE^Y$  requires knowing the total number of encounters for each racial group (including

non-shootings), as shown by an application of Bayes’ Rule by Knox and Mummolo (2020a). Intuitively, suppose a use of force dataset has minority and white civilians in equal number. If police encounter more white civilians due to their majority status, then this implies anti-white bias, but if they encounter more minority civilians due to deployment patterns then this implies anti-minority bias. Importantly, this missing data problem cannot be solved with a covariate adjustment, however rich the dataset. A first takeaway, then, is that despite the convenient availability of nationwide datasets on the deadly use of force, to make statements about racial bias researchers must supplement these sources with additional data and assumptions.

A second approach is to take datasets of police stops and to compare by race the *rates* at which such encounters lead to police violence. A prominent example is Fryer (2018), which claims to find racial disparities in sublethal but not lethal use of force. The issue is that if officers stop Black and white citizens according to different criteria, then these rates are not computed on comparable populations; to justify the conclusion from the design, there must be no racial bias in the decision to detain civilians. As Knox and Mummolo (2020b) put it, “*analysts must assume away racial bias in a study of racial bias*” (p. 25).

Benchmarking approaches offer a more promising way forward, though all of the issues covered in Section 2.2.1 still apply. With the help of the causal framework elaborated here, it is easy to see that benchmarking simply provides a proxy for the missing information on total encounters by race, which can be done through the use of local demographics (Edwards, Esposito, and Lee 2018; Edwards, Lee, and Esposito 2019; Mentch 2020); observations on police-citizen interactions (Tregle, Nix, and Alpert 2019), though this breaks down if police exhibit racial bias in their choice of which citizens to engage; and crime rates (Cesario, Johnson, and Terrill 2019) and arrests (Tregle, Nix, and Alpert 2019), though as discussed in Section 2.2.1 this has the flaw of reproducing past patterns of bias rather than serving as ground truth; see Section 5.6 (Predictive Policing and Big Data) for a related discussion of these issues.

### 3 Structural Racism in Policing

I have begun this inquiry into racial bias in policing by situating the discussion in the context of a micro-level *police-civilian interaction*. The intense scholarly focus on this context requires no normative justification: individual contact with the police shapes democratic citizenship, personal safety and freedom, and in the extreme (but not at all uncommon) case, the very course of people’s lives. Beyond its substantive importance, the micro-level interaction is also — with a great deal of careful thought and methodological ingenuity — ultimately *tractable* for social scientists who want to ask precise causal questions and answer them using readily available administrative records and statistics.

I now turn to an entirely different perspective on racial bias: the interaction between law enforcement *systems* and the communities they serve and surveil, as embedded in a larger set of overlapping social structures including residential segregation, income inequality, health, education, and the carceral state. Scholarship in this vein brings to the table an understanding of how racial bias permeates society and reproduces inequalities even if most police officers personally exhibit neither “taste-based discrimination” nor even the broadest definition of a causal race-based disparity in their treatment of civilians.

Siegel (2020) begins by noting a massive “disparity in racial disparities” across spatial contexts that seems unlikely to be driven by officer-level factors alone. For instance, whereas a Black person in Chicago is thirty-eight times more likely than a white person to be killed by the police, the ratio in Las Vegas is closer to two-to-one (Ibid.). Siegel (2020) explores structural racism as an explanation, including residential segregation and racial gaps in incarceration rates, educational attainment, employment status, and indicators of economic status. Economic disparity, employment disparity, and residential segregation all emerge as statistically and substantively significant drivers of racial disparities in police shootings of unarmed people across both cities and states; residential segregation has by far the most dramatic effect.

The interaction of segregation and policing has been explored at length in a recent lit-

erature on spatial context, which shifts the unit of analysis up to a neighborhood or other geography with a common set of social characteristics. Beyond mere racial separation, *segregation* refers to residential patterns that concentrate racial minorities in enclaves marked by stigmatization, exclusion from important social and legal institutions, and diminished access to economic resources and opportunity (Bell 2020). In the presence of segregation thus defined, the state-citizen relationship is fundamentally different in predominantly white and Black areas, and accordingly policing follows different scripts.

Specifically, Bell (2020) highlights some of the key mechanisms that link policing to segregation. First, mass criminalization traps already disadvantaged people in cycles of arrest, misdemeanor prosecution, and supervision (Hinton 2016; Goffman 2014). Second, law enforcement is implicated in the patrolling of borders between white and Black spaces (Anderson 2015; Bell 2015). Policing is also intertwined with other bureaucracies such as authorities who control housing and eviction (Desmond and Valdez 2013). Finally, police forces construct jurisdictions based on symbolic notions of the violence of Black neighborhoods and the economic value of white neighborhoods — ideas that become cemented in daily practices of service provision and social control (Gordon 2020).

Empirical studies of local context and policing have found that predominantly Black areas experience more vehicle searches (Wright, Gaozhao, and Snow 2021), more traffic citations (Ingram 2007), and more police-caused homicide (Holmes, Painter II, and Smith 2019) than other parts of the city, even after adjusting for individual-level features of the encounter. (However, there is also evidence that Black motorists are policed more heavily than white motorists when driving through white-majority areas (Meehan and Ponder 2002).) Local economic inequality is also predictor of police killings (Helms and Costanza 2020).

Another strand of the structural racism literature, from social psychology, explores how deeply embedded stereotypes about Black criminality and entanglement with the criminal justice system serve to reinforce the very policies that created them (Hetey and Eberhardt 2014; Carbado and Rock 2016). Such stereotypes operate at a subconscious level (Trawalter

et al. 2008), creating differences in how both civilians and officers react to Black and white faces in life-and-death scenarios (Payne 2001; Bradley and Kennison 2012). While the implications of these psychological mechanisms are certainly reflected in analyses of micro-level police-civilian interactions, this body of research suggests that their *mitigation* is unlikely to lie with interventions that target the individual; long-term reform must address the mechanisms that hard-wire racial bias into individuals’ psychology at their source.

## 4 The Democratic Dimension of Policing

An important growing literature in political science and sociology examines the central role of policing in constructing the relationship between the citizen and the state. In stark contrast to the liberal-democratic “first face” of the state, which includes citizenship-*reinforcing* mechanisms such as elections, interest groups, and the provision of public goods and social services, policing belongs to the citizenship-*eroding* “second face” of the state: the “governing institutions and officials that exercise social control and encompass various modes of coercion, containment, repression, surveillance, regulation, predation, discipline, and violence” (Soss and Weaver (2017), p. 567). Because it is by far the most frequent (and sometimes only) point of contact between American citizens and their government — especially for low-income and minority communities — this “second face” arguably plays the bigger role in shaping people’s political identities (Lerman and Weaver 2014*a,b*; Justice and Meares 2014). Beyond being important in their own right, citizens’ experiences with the coercive arm of the state in turn reconstruct their relationship with the liberal-democratic first arm, in a process that Lerman and Weaver (2014*a*) term “denaturalization.” This section covers three well-documented denaturalization processes resulting from negative encounters with the police: decrease in voter turnout and political participation; civic disengagement and alienation from other social institutions; and erosion of trust in the legitimacy of law enforcement and, more broadly, government.

## 4.1 Political Participation

In general, frequent contact with the criminal justice system, including police stops and searches, leads citizens to withdraw from political life: they become less likely to organize around common issues with other members of their communities, to contact their political representatives, or to turn out to vote, even if they are legally allowed to do so (Weaver and Lerman 2010; Lerman and Weaver 2014a). However, some conditional *mobilization* effects appear also possible. Looking at both people who have had personal contact with the police — including being arrested, charged, or questioned — and those with proximal contact via someone they know, Walker (2014) finds that self-reported non-voting political behaviors such as signing a petition, donating to a political cause, or taking part in a demonstration *increase* with contact for both groups, though personal contact depresses validated voter turnout. Using geocoded data from police stops in New York City together with local election returns, Laniyonu (2019) also finds both mobilizing and demobilizing effects of intense police contact. Whereas concentrated policing was associated with lower turnout in the 2006 and 2010 midterm elections, it was associated with higher turnout in the 2008 presidential election and the 2013 Democratic primary and general for mayor. In the 2013 mayoral election, heavily policed neighborhoods were more likely to support the candidate who advocated eliminating the stop-and-frisk program. Thus a fruitful avenue for future work is exploring when policing mobilizes people politically, including the potential role of candidate platforms and characteristics.

## 4.2 Civic Engagement

Several studies have examined the effects of different forms of police contact on 911 and 311 calls, which reflect citizens' willingness to report emergencies and request non-emergency assistance from government, respectively. In Milwaukee, the beating of Frank Jude, an unarmed man, by the police, had a chilling effect on 911 calls that lasted over a year and resulted in an estimated net loss of over 22,000 calls for service (Desmond, Papachristos, and

Kirk 2016); however, officer-involved shootings did not have a similar effect on 911 or 311 calls in Los Angeles (Cohen et al. 2019). Lerman and Weaver (2014*b*) find that police stops reduce 311 calls specifically when they include searches or the use of force, and Chenane, Wright, and Gibson (2020) find that stops for traffic violations depress emergency and non-emergency contact more than other face-to-face police encounters, especially when the interaction is perceived to be unfair. Thus, punitive interactions have important spillover effects, imposing costs not only on targeted individuals but on their neighbors and communities (Burch 2013).

In New York City, Legewie and Fagan (2019) found that exposure to police surges affected educational outcomes, significantly reducing test scores for African American boys (but not for African American girls or Hispanic students). However, importantly for social scientists, such effects are not always detectable from administrative data, even when they are present. In a phenomenon Brayne (2014) terms “system avoidance,” civic disengagement often manifests as a missing data problem, as individuals who have come to distrust the state after punitive interactions seek to render themselves invisible to its instruments of control (Rosenthal 2019). Specifically, Brayne (2014) finds that individuals who have been stopped by police, arrested, convicted, or incarcerated withdraw from institutions that keep formal records, such as hospitals, banks, schools, and employers, creating a cycle of social stratification and marginalization.

### **4.3 Perceptions of Legitimacy and Trust in Government**

Across all racial groups, trust in the police is driven by judgments about the fairness of the procedures they follow in exercising their authority, more so than by effectiveness in crime control or distribution of police resources (Tyler 2005). However, Black citizens are likely to experience different treatment than white citizens in otherwise similar situations: for instance, in traffic stops, they are more likely to report impolite officer demeanor, to be threatened with arrest and search, and to be handcuffed and arrested (Epp, Maynard-Moody, and Haider Markel 2014; Baumgartner, Epp, and Shoub 2018). As a consequence,

Black drivers who have been subjected to traffic stops are more likely than their white counterparts to feel as if they are “second-class citizens” and to question the legitimacy of police authority (Ibid.). This feeling of being surveilled without being protected or represented to the same degree as white citizens contributes to a condition of “distorted responsiveness”: the contradictory sense that police are “everywhere when surveilling people’s everyday activity and nowhere if called upon to respond to serious harm” (Prowse, Weaver, and Meares 2020). Importantly, such attitudes are not confined to perceptions of the police, but inform broader assessments of people’s confidence in government (Silva et al. 2020).

## **5 Avenues for Reform**

### **5.1 Racial Composition of the Force**

Since the 1960s, police forces have gradually diversified along lines of race, gender, and sexual orientation (Sklansky 2005), and there are reasons to believe that simply becoming more representative on these key dimensions may lead police to exhibit less discrimination and greater empathy toward marginalized groups. Evidence that racial diversity affects racial bias in micro-level encounters is, however, limited.

Donohue III and Levitt (2001) find that an increase in the proportion of white officers in metropolitan police departments leads to more minority arrests, and McCrary (2007) finds that the effect also holds in the other direction: increasing the proportion of officers who are Black decreases the proportion Black among serious arrestees. Consistent with these findings, Bulman (2019) shows that the ratio of Black-to-white arrests is significantly higher under white sheriffs. It is important to note that in all these studies, Black officers provide a benchmark for white officers, so the results cannot speak to any bias exhibited by both groups.

Several more recent studies using both quantitative and qualitative approaches have found null effects of the composition of the force on racial disparities in stops (Shjarback



et al. 2017) and arrests (Eitle, Stolzenberg, and D'Alessio 2005); police-involved homicides of Black people (Nicholson-Crotty, Nicholson-Crotty, and Fernandez 2017); citizen complaints (Hickman and Piquero 2009); perceptions of police professionalism and legitimacy among Black people (Benton 2020); attitudes about police responsiveness and misconduct (Brunson and Gau 2015); and perceptions of police fairness (Socia et al. 2021). Taken together, these findings indicate that macro-level socioeconomic and institutional factors greatly outweigh any influence of officer identity.

## 5.2 Political Representation

Another hypothesis is that minority descriptive representation in local government makes law enforcement more accountable to minority citizens, either via a direct electoral mechanism in the case of county sheriffs, or indirectly through city council control of police budgets and hiring/firing decisions. Studies that look at this question do, for the most part, find reductions both in punitive policing overall and in racial disparities in outcomes. Looking at 30 of the 60 largest U.S. cities, Ochs (2011) finds that the political incorporation of Black elected officials reduces the incidence of police use of force, and Christiani et al. (2021) show that in cities and towns with majority-Black city councils, traffic stops are less likely to result in a search, particularly for Black drivers. Focusing on elected sheriffs, Facchini, Knight, and Testa (2020) find that Black arrest rates fell after the passage of the Voting Rights Act in counties that both were covered by the legislation and had a large number of newly enfranchised Black voters, with no such change observed for white arrest rates.

Unpacking the mechanisms by which political representation changes policing outcomes is a promising avenue for future work. While Hopkins and McCabe (2012) find that Black mayors increase the Black share of the force, there is not strong evidence that changes in policing travel through this channel (see Section 5.1.1 above). Indeed, analyzing Black representation on the force and in government within the same study, Sharp (2014) shows that the former has no effect on Black order maintenance arrests without the latter. There

is, however, evidence of an interaction between the two: Choi and Hong (2020) find that increasing minority representation on city councils decreases racial disparities in outcomes only when minorities are also well-represented on the force. These findings raise important questions about how exactly local governments exercise control over front-line bureaucracies and the conditions under which they can achieve meaningful reform.

### **5.3 Training, Monitoring, and Internal Reform**

Responding to public outcry over shootings of Black citizens by the police, police departments have instituted various internal reforms, in particular officer trainings in deescalation tactics, the use of body-worn cameras, and the adoption of “community policing” models. Voluntary selection into these programs creates serious challenges for rigorous evaluation — and there have not yet been many randomized control trials — but taken as a whole such strategies do not seem to reduce the incidence of police violence or create long-run changes in police behavior.

Reviewing 64 de-escalation training evaluations conducted over a 40-year period, Engel, McManus, and Herold (2020) find mild to moderate positive changes in self-reported outcomes but inconclusive results on the number of incidents of police violence. However, they urge caution in interpreting most of these findings, as they do not have research designs that can support causal claims, and call for investments in rigorous evaluation of police training interventions before they are implemented broadly. Similarly, Fagan and Campbell (2020) find no evidence that enhanced police training focused on mental health reduces fatal police shootings of people in crises or racial disparities in shootings overall. They suggest that such trainings may be more effective in these goals if they incorporate explicit race-specific components in their curricula.

Another intervention, body-worn cameras (BWCs), shows more promise, and has been subjected to a few more rigorous evaluations in recent years. Randomized control trials conducted in Tempe, Arizona (Gaub, Todak, and White 2020), Hallandale Beach, Florida

(Headley, Guerette, and Shariati 2017), and Rialto, California (Ariel, Farrar, and Sutherland 2015) all showed reductions in police use of force. Hedberg, Katz, and Choate (2017) found a reduction in citizen complaints, though not arrests, even though in practice officers only had the cameras activated 32% of the time. However, taking an interrupted time series approach using post-hoc administrative data to avoid potential Hawthorne effects, Koslicki, Makin, and Willits (2020) observe an initial drop in use-of-force incidents immediately after adoption, but effects weakening over time as BWCs become normalized with daily use. This is one of the few studies to look over a multi-year time horizon, and its findings are less optimistic.

Finally, some departments have experimented with “community policing,” a conceptual approach that emphasizes community partnerships, increased police-community interaction, greater autonomy for front-line officers, and procedural justice for citizens (President’s Task Force on 21st Century Policing 2015). Practically, community policing is a heterogeneous and poorly defined treatment, ranging from trainings emphasizing a “guardian mindset” over a “warrior ethos” (Koslicki et al. 2021) to door-to-door non-enforcement contact with citizens (Peyton, Sierra-Arévalo, and Rand 2019). Taken as a whole, such interventions can improve citizens’ perceptions of and trust in the police (Gill et al. 2014; Peyton, Sierra-Arévalo, and Rand 2019), but do not affect the incidence of deadly use of force (Koslicki et al. 2021; Gill et al. 2014).

## **5.4 Reliance on Fines and Fees**

In the wake of the killing of Michael Brown, an 18-year-old Black man, by the police in Ferguson, Missouri, a Department of Justice report concluded that Ferguson had been imposing a “predatory system of government” on its poor Black citizens (Department of Justice Civil Rights Division 2015). The jurisdiction was using police surveillance and coercion to target and arrest Black people for civil-ordinance violations, charging prohibitive fines and fees and entrapping them in a perpetual cycle of financial debt and entanglement with the criminal

justice system. While the Michael Brown killing put Ferguson in the national spotlight, its policies are not unique. Across the U.S., fines, fees, and forfeitures constitute a larger share of local revenues in municipalities with more Black residents, and where there is no Black representation in local elected office (Sances and You 2017).

Extractive law enforcement institutions exacerbate racial disparities in policing outcomes. Recent work by Harvey (2020) shows that how much revenue a municipality captures from fines and fees impacts its policing intensity and priorities. Harvey uses a regression-discontinuity design based on a fiscal rule that reduced the share of traffic fine revenue captured by the local government in towns of 500 people or more in Saskatchewan, Canada. Towns just above the threshold saw more traffic enforcement than comparable towns just below the threshold; towns just below the threshold also gave drivers less time to pay their fines, leading to higher risks of late fees and license suspensions. Consistent with other literatures reviewed here, more aggressive traffic enforcement does not have race-neutral consequences: greater reliance on fines and fees also leads to wider racial disparities in vehicle searches Shoub et al. (2020) and more arrests of Blacks and Hispanics — but not whites — for drugs and driving under the influence (Makowsky, Stratmann, and Tabarrok 2019).

## 5.5 Proactive Policing

Policies that give officers broad discretion to stop and search citizens are causally and inextricably linked to racial profiling and racial disparities in outcomes (Epp, Maynard-Moody, and Haider Markel 2014; Fagan et al. 2016; Rushin and Edwards 2021), a finding most clearly documented under New York City’s Stop, Question, and Frisk (SQF) program (Goel, Rao, and Shroff 2016). However, some legal and administrative controls can mitigate disparities arising from “proactive policing” regimes. Mummolo (2018) shows that an internal memo to NYC police officers that raised their perceived level of supervisory oversight, issued contemporaneously with a class-action lawsuit against the city alleging unconstitutional policing under SQF (*Floyd, et al. v. City of New York, et al.*), led to an immediate, discontinuous,

and sustained increase in hit rates, consistent with more conservative policing tactics. MacDonald and Braga (2019) also find that racial disparities in hit rates became smaller after the *Floyd, et al.* settlement.

## 5.6 Predictive Policing and Big Data

Police departments in some of the largest U.S. cities have experimented with using big data and analytics to predict where crimes will occur and thereby inform police deployment patterns and practices. Proponents of data-driven decisionmaking argue that if officers' discretion is what generates racial disparities in outcomes, algorithmic approaches have the potential to reduce any inequalities driven by prejudice and to improve procedural fairness. For example, algorithms that dynamically use data to make stop and search decisions can be designed to maintain equal hit rates across racial groups (Goel et al. 2017), essentially building an outcome test into the police officer's work in real time. Such approaches have been adopted — and subsequently discontinued, following lawsuits and citizen complaints — in Los Angeles (Brantingham, Valasik, and Mohler 2018; Brayne 2017) and Chicago; in New York City proprietary algorithms continue to contribute to deployment decisions (Lau 2020).

But scholars and citizens have expressed deep concerns that predictive policing reinforces rather than mitigates preexisting racial disparities. Far from being an accurate or scientific ground truth of criminality, data generated by past patterns of bias create a self-fulfilling prophecy: individuals and communities profiled as criminal are subjected to more intense policing, leading to greater contact with the criminal justice system that reinforces their criminality and feeds back into the data-generating process (Harcourt 2007; Jefferson 2018). Thus, insofar as predictive algorithms reconstruct the very reality they are meant to respond to, evaluating their impact is a challenge beyond even the usual problems of causal inference.

## 6 Where Do We Go from Here?

I conclude by suggesting two broad avenues by which social scientists — in particular empirical researchers with statistical training, computational resources, and data-collection capabilities at their disposal — can make progress in the study of race and policing: (1) fostering intellectual exchange between causal inference approaches grounded in micro-level interactions and the scholarship on structural racism in policing; and (2) remedying the most critical missing data problems that currently stand in the way of further progress.

### 6.1 Causal Frameworks beyond the Micro-Level Interaction

The literature on police-civilian interactions has contributed not only to our understanding of an inherently important phenomenon, but to a rigorous conceptualization of the very idea of racial bias. However, an *exclusive* focus on the micro level risks missing some of the key mechanisms by which racism manifests and reproduces itself across society. One basic limitation that is readily acknowledged by the micro-level literature is that police deployment patterns generate inequalities in outcomes absent the broadest definition of officer bias or even any race-based differences in behavior (Knox and Mummolo 2020*b*; Streeter 2019), but they nonetheless constitute a form of biased policing. But a deeper critique takes seriously insights from the structural racism literature to reveal the limitations of a potential-outcomes analysis of police-civilian interactions. If, as Bell (2020) and others have convincingly argued, Black citizens tend to live in fundamentally different *communities* from white citizens, each with a different relationship with the second face of the state, and if the police in turn perceive civilians based not only on their personal characteristics but on their characteristics *interacted* with their spatial and social context, then the thought exercise of substituting a minority civilian for a white civilian “holding everything else about the encounter constant” starts to ring substantively hollow. In other words, segregation and multiple overlapping systems of inequality assure that the racial disparities observed in the world are not constituted

of all-else-equal comparisons. Understanding the ways in which *all else covaries* should be placed at the center of the research enterprise.

The conclusion is not that we should abandon our causal frameworks. Rather, there is great potential in applying potential outcomes logics to community- and system-level questions, while acknowledging the difficulties of making all-else-equal statements in a context where the “treatment” of racism is inherently structural. Some obvious ways forward suggest themselves, though others will require deeper consideration. For instance, one might take a policy-oriented view: holding facts on the ground (population, segregation, inequality and crime) constant, what is the causal effect of one particular deployment pattern or policing approach over another? Although it will require challenging assumptions, our models can also explicitly build in relationships between race and other factors, allowing us to think about the effect of race as channeled through entire overlapping systems — not unlike how Knox and Mummolo (2020*b*) encourage us to think about effects of race as channeled through multi-stage causal processes *within* a police-citizen interaction.

## 6.2 Data Collection: Beyond Administrative Records

One of the most important contributions of Knox, Lowe, and Mummolo (2020) is to highlight just how acute a missing data problem researchers interested in causal claims about policing face. Fortunately, government initiatives to enhance police accountability often go hand-in-hand with publicly available microdata. In California, for instance, The Racial and Identity Profiling Act of 2015 (RIPA) now requires every state law enforcement agency to report all their stop data to the Attorney General, and the data is gradually becoming available.<sup>4</sup> Thus, at the very least, the worst kind of selection on the dependent variable problem in studying arrests and police use of force can be avoided.

In principle, technological advances in collecting, processing, and analyzing audiovisual and geospatial data have the potential to go one step further: to chip away at the problem

---

<sup>4</sup><https://openjustice.doj.ca.gov/exploration/stop-data>

of detecting bias in officers' initial decision to engage civilians. Of course, access to such data will likely require cooperation between police forces and academics. But this research enterprise does not need to proceed on a large or nationwide scale; even an incidental window into this opaque first stage of the data-generating process can be quite helpful in informing researchers' assumptions when they use conventional administrative records.

Another particularly important scholarly contribution reviewed here has illuminated how policing constructs people's identities as democratic citizens. Continued attention to this question, with a particular emphasis on the voices of marginalized groups, should be a part of the social science research agenda and of conversations about policy reform. One of the most normatively and epistemologically troubling findings reviewed here is that people who have had negative contact with the police literally disappear from our data. Thus a valuable contribution for social scientists to work on is reconstituting the data from the ground up, through traditional surveys, text data mining and natural language processing, and other innovative strategies (e.g., Prowse, Weaver, and Meares (2020)) that shed light on people's experiences with the police.

## References

- Alpert, Geoffrey P., Michael R. Smith, and Roger G. Dunham. 2004. "Toward a Better Benchmark: Assessing the Utility of Not-at-Fault Traffic Crash Data in Racial Profiling Research." *Justice Research and Policy* 6(1): 43–69.
- Anbarci, Nejat, and Jungmin Lee. 2014. "Detecting Racial Bias in Speed Discounting: Evidence from Speeding Tickets in Boston." *International Review of Law & Economics* 38: 11–24.
- Anderson, Elijah. 2015. "'The White Space'." *Sociology of Race and Ethnicity* 1(1): 10–21.
- Antonovics, Kate, and Brian G. Knight. 2009. "A New Look at Racial Profiling: Evidence from the Boston Police Department." *Review of Economics & Statistics* 91(1): 163–177.



- Anwar, Shamena, and Hanming Fang. 2006. "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence." *The American Economic Review* 96(1): 127–151.
- Ariel, Barak, William Farrar, and Alex Sutherland. 2015. "The Effect of Police Body-Worn Cameras on Use of Force and Citizens' Complaints Against the Police: A Randomized Controlled Trial." *Journal of Quantitative Criminology* 31(3): 509–535.
- Arrow, Kenneth J. 1972. "Models of Job Discrimination." In *Discrimination in Economic Life*, ed. Anthony Pascal. Lexington, MA: D.C. Heath pp. 83–102.
- Arrow, Kenneth J. 1973. "The Theory of Discrimination." In *Discrimination in Labor Markets*, ed. Orley Ashenfelter, and Albert Rees. Princeton: Princeton University Press p. 3–33.
- Arrow, Kenneth J. 1998. "What Has Economics to Say about Racial Discrimination?" *The Journal of Economic Perspectives* 12(2): 91–100.
- Ayres, Ian. 2002. "Outcome Tests of Racial Disparities in Police Practices." *Justice Research and Policy* 4(Special Issue): 132–142.
- Baumgartner, Frank R., Derek A. Epp, and Kelsey Shoub. 2018. *Suspect Citizens: What 20 Million Traffic Stops Tell us about Policing and Race*. Chicago: University of Chicago Press.
- Beck, Brenden, and Adam Goldstein. 2018. "Governing Through Police? Housing Market Reliance, Welfare Retrenchment, and Police Budgeting in an Era of Declining Crime." *Social Forces* 96(3): 1183–1210.
- Becker, Gary S. 1957. *The Economics of Discrimination*. Chicago: University of Chicago Press.
- Bell, Jeannine. 2015. "Can't We Be Your Neighbor: Trayvon Martin, George Zimmerman, and the Resistance to Blacks as Neighbors Symposium: The Civil Rights Act of 1964 at 50: Past, Present and Future: Panel III: Reshaping Public and Private Space: Public Accommodations, Neighborhoods, and Housing." *Boston University Law Review* 95(3): 851–872.
- Bell, Monica C. 2020. "Anti-Segregation Policing." *New York University Law Review* 95.

- Benton, Mark. 2020. "Representation Is Not Enough: Symbolic Representation and Perceptions of the Police." *Administration & Society* 52(5): 794–822.
- Bjerk, David. 2007. "Racial Profiling, Statistical Discrimination, and the Effect of a Colorblind Policy on the Crime Rate." *Journal of Public Economic Theory* 9(3): 521–545.
- Borooah, Vani K. 2001. "Racial Bias in Police Stops and Searches: An Economic Analysis." *European Journal of Political Economy* 17(1): 17–37.
- Bradley, Kristopher I., and Shelia M. Kennison. 2012. "The Effect of Mortality Salience on Weapon Bias." *International Journal of Intercultural Relations* 36(3): 403–408.
- Brantingham, P. Jeffrey, Matthew Valasik, and George O. Mohler. 2018. "Does Predictive Policing Lead to Biased Arrests? Results From a Randomized Controlled Trial." *Statistics and Public Policy* 5(1): 1–6.
- Brayne, Sarah. 2014. "Surveillance and System Avoidance: Criminal Justice Contact and Institutional Attachment." *American Sociological Review* 79(3).
- Brayne, Sarah. 2017. "Big Data Surveillance: The Case of Policing." *American Sociological Review* 82(October): 977–1008.
- Brunson, Rod K., and Jacinta M. Gau. 2015. "Officer Race Versus Macro-Level Context: A Test of Competing Hypotheses About Black Citizens' Experiences With and Perceptions of Black Police Officers." *Crime and Delinquency* 61(2): 213–242.
- Bulman, George. 2019. "Law Enforcement Leaders and the Racial Composition of Arrests." *Economic Inquiry* 57(4): 1842–1858.
- Burch, Traci. 2013. *Trading Democracy for Justice: Criminal Convictions and the Decline of Neighborhood Political Participation*. Chicago: University of Chicago Press.
- Calderon, Tovah Renee. 2000. "Race-Based Policing from Terry to Wardlow: Steps down the Totalitarian Path." *Howard Law Journal* 44(1): 73–106.

- Carbado, Devon W., and Patrick Rock. 2016. "What Exposes African Americans to Police Violence?" *Harvard Civil Rights-Civil Liberties Law Review* 51(1): 159–187.
- 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): 586–595.
- Chenane, Joselyne L., Emily M. Wright, and Chris L. Gibson. 2020. "Traffic Stops, Race, and Perceptions of Fairness." *Policing and Society* 30(6): 720–737.
- Choi, Heasun, and Sounman Hong. 2020. "Linking Political and Bureaucratic Representation: Does Minority Representation on City Councils Affect the Outcomes of Street-Level Service?" *Public Administration* .
- Christiani, Leah, Kelsey Shoub, Frank R. Baumgartner, Derek A. Epp, and Kevin Roach. 2021. "Better for Everyone: Black Descriptive Representation and Police Traffic Stops." *Politics, Groups, and Identities* 0(0): 1–10.
- Close, Billy R., and Patrick L. Mason. 2007. "Searching for Efficient Enforcement: Officer Characteristics and Racially Biased Policing." *Review of Law and Economics* 3(2): 263–322.
- Cohen, Elisha, Anna Gunderson, Kaylyn Jackson, Paul Zachary, Tom S. Clark, Adam N. Glynn, and Michael Leo Owens. 2019. "Do Officer-Involved Shootings Reduce Citizen Contact with Government?" *Journal of Politics* 81(3): 1111–1123.
- Department of Justice Civil Rights Division. 2015. *The Ferguson Report: Department of Justice Investigation of the Ferguson Police Department*. New York: New Press.
- Desmond, Matthew, and Nicol Valdez. 2013. "Unpolicing the Urban Poor: Consequences of Third-Party Policing for Inner-City Women." *American Sociological Review* 78(1).
- Desmond, Matthew, Andrew V. Papachristos, and David S. Kirk. 2016. "Police Violence and Citizen Crime Reporting in the Black Community." *American Sociological Review* 81(5): 857–876.

- Dharmapala, Dhammika, and Stephen L. Ross. 2004. "Racial Bias in Motor Vehicle Searches: Additional Theory and Evidence." *Contributions in Economic Analysis & Policy* 3(1): 1–21.
- Dominitz, Jeff, and John Knowles. 2006. "Crime Minimisation and Racial Bias: What Can We Learn from Police Search Data?" *Economic Journal* 116(515): F368–F384.
- Donohue III, John J., and Steven D. Levitt. 2001. "The Impact of Race on Policing and Arrests." *The Journal of Law and Economics* 44(2): 367–394.
- Durlauf, Steven N. 2006. "Assessing Racial Profiling." *Economic Journal* 116(515): F402–F426.
- Edwards, Frank, Hedwig Lee, and Michael Esposito. 2019. "Risk of Being Killed by Police Use of Force in the United States by Age, Race–Ethnicity, and Sex." *Proceedings of the National Academy of Sciences* 116(34): 16793–16798.
- Edwards, Frank, Michael H. Esposito, and Hedwig Lee. 2018. "Risk of Police-Involved Death by Race/Ethnicity and Place, United States, 2012-2018." *American Journal of Public Health* 108(9): 1241–1248.
- Eitle, David, Lisa Stolzenberg, and Stewart J. D'Alessio. 2005. "Police Organizational Factors, the Racial Composition of the Police, and the Probability of Arrest." *Justice Quarterly* 22(1): 30–57.
- Elwert, Felix, and Christopher Winship. 2014. "Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable." *The Annual Review of Sociology* 40: 31–53.
- Engel, Robin S. 2008. "A Critique of the Outcome Test in Racial Profiling Research An Exchange on Racial Profiling Research." *Justice Quarterly* 25(1): 1–36.
- Engel, Robin S., and Jennifer M. Calnon. 2004. "Comparing Benchmark Methodologies for Police-Citizen Contacts: Traffic Stop Data Collection for the Pennsylvania State Police." *Police Quarterly* 7(1): 97–125.
- Engel, Robin S., and Rob Tillyer. 2008. "Searching for Equilibrium: The Tenuous Nature of the Outcome Test An Exchange on Racial Profiling Research." *Justice Quarterly* 25(1): 54–71.

- Engel, Robin S., Hannah D. McManus, and Tamara D. Herold. 2020. "Does De-escalation Training Work?" *Criminology & Public Policy* 19(3): 721–759.
- Epp, Charles, Steven Maynard-Moody, and Donald Haider Markel. 2014. *Pulled Over: How Police Stops Define Race and Citizenship*. New York: Cambridge University Press.
- Facchini, Giovanni, Brian G. Knight, and Cecilia Testa. 2020. The Franchise, Policing, and Race: Evidence from Arrests Data and the Voting Rights Act. Technical Report w27463 National Bureau of Economic Research.
- Fagan, Jeffrey, and Alexis Campbell. 2020. "Race and Reasonableness in Police Killings." *Boston University Law Review* 100(Columbia Public Law Research Paper No. 14-655).
- Fagan, Jeffrey, Anthony A. Braga, Rod K. Brunson, and April Pattavma. 2016. "Stops and Stares: Street Stops, Surveillance, and Race in the New Policing." *Fordham Urban Law Journal* 43(3): 539–614.
- Fryer, Roland G. 2018. "An Empirical Analysis of Racial Differences in Police Use of Force." *Journal of Political Economy* 127(3): 1210–1261.
- Gaub, Janne E., Natalie Todak, and Michael D. White. 2020. "The Distribution of Police Use of Force across Patrol and Specialty Units: A Case Study in BWC Impact." *Journal of Experimental Criminology*.
- 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): 813–823.
- Gill, Charlotte, David Weisburd, Cody W. Telep, Zoe Vitter, and Trevor Bennett. 2014. "Community-Oriented Policing to Reduce Crime, Disorder and Fear and Increase Satisfaction and Legitimacy among Citizens: A Systematic Review." *Journal of Experimental Criminology* 10(4): 399–428.

- Goel, Sharad, Justin M. Rao, and Ravi Shroff. 2016. "Precinct or Prejudice? Understanding Racial Disparities in New York City's Stop-and-Frisk Policy." *The Annals of Applied Statistics* 10(1): 365–394.
- Goel, Sharad, Maya Perelman, Ravi Shroff, and David Alan Sklansky. 2017. "Combatting Police Discrimination in the Age of Big Data." *New Criminal Law Review* 20(2): 181–232.
- Goff, Phillip Atiba, and Kimberly Barsamian Kahn. 2012. "Racial Bias in Policing: Why We Know Less Than We Should." *Social Issues and Policy Review* 6(1): 177–210.
- Goffman, Alice. 2014. *On the Run: Fugitive Life in an American City*. Chicago: University of Chicago Press.
- Gordon, Daanika. 2020. "The Police as Place-Consolidators: The Organizational Amplification of Urban Inequality." *Law & Social Inquiry* 45(1): 1–27.
- Gottschalk, Marie. 2016. "Inequality and the Carceral State." *The Oxford Handbook of Historical Institutionalism* .
- Grogger, Jeffrey, and Greg Ridgeway. 2006. "Testing for Racial Profiling in Traffic Stops from behind a Veil of Darkness." *Journal of the American Statistical Association* 101(475): 878–887.
- Gross, Samuel R., and Katherine Y. Barnes. 2002. "Road Work: Racial Profiling and Drug Interdiction on the Highway." *Michigan Law Review* 101(3): 651–754.
- Harcourt, Bernard E. 2004. "Rethinking Racial Profiling: A Critique of the Economics, Civil Liberties, and Constitutional Literature, and of Criminal Profiling More Generally." *The University of Chicago Law Review* 71(4): 1275–1381.
- Harcourt, Bernard E. 2007. *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age*. Chicago: University of Chicago Press.
- Harvey, Anna. 2020. "Fiscal Incentives in Law Enforcement." *American Law and Economics Review* 22(1): 173–210.

- Headley, Andrea M., Rob T. Guerette, and Auzeen Shariati. 2017. "A Field Experiment of the Impact of Body-Worn Cameras (BWCs) on Police Officer Behavior and Perceptions." *Journal of Criminal Justice* 53: 102–109.
- Hedberg, E. C., Charles M. Katz, and David E. Choate. 2017. "Body-Worn Cameras and Citizen Interactions with Police Officers: Estimating Plausible Effects Given Varying Compliance Levels." *Justice Quarterly* 34(4): 627–651.
- Helms, Ronald, and S. E. Costanza. 2020. "Contextualizing race: a conceptual and empirical study of fatal interactions with police across US counties." *Journal of Ethnicity in Criminal Justice* 18(1): 43–71.
- Hernán, Miguel A. 2016. "Does Water Kill? A Call for Less Casual Causal Inferences." *Annals of Epidemiology* 26(10): 674–680.
- Hernández-Murillo, Rubén, and John Knowles. 2004. "Racial Profiling or Racist Policing? Bounds Tests in Aggregate Data." *International Economic Review* 45(3): 959–989.
- Hetey, Rebecca C., and Jennifer L. Eberhardt. 2014. "Racial Disparities in Incarceration Increase Acceptance of Punitive Policies." *Psychological Science* 25(10): 1949–1954.
- Hickman, Matthew, and Alex Piquero. 2009. "Organizational, Administrative, and Environmental Correlates of Complaints About Police Use of Force: Does Minority Representation Matter?" *Crime & Delinquency* 55: 3–27.
- Hinton, Elizabeth K. 2016. *From the War on Poverty to the War on Crime: The Making of Mass Incarceration in America*. Cambridge: Harvard University Press.
- Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81(396): 945–60.
- Holmes, Malcolm D., Matthew A. Painter II, and Brad W. Smith. 2019. "Race, Place, and Police-Caused Homicide in U.S. Municipalities." *JQ: Justice Quarterly* 36(5): 751–786.

- Hopkins, Daniel J., and Katherine T. McCabe. 2012. "After It's Too Late: Estimating the Policy Impacts of Black Mayoralities in U.S. Cities." *American Politics Research* 40(4): 665–700.
- Horn, Brady P., Jill J. McCluskey, and Ron C. Mittelhammer. 2014. "Quantifying Bias in Driving-Under-the-Influence Enforcement." *Economic Inquiry* 52(1): 269–284.
- Horrace, William C., and Shawn M. Rohlin. 2016. "How Dark Is Dark? Bright Lights, Big City, Racial Profiling." *Review of Economics & Statistics* 98(2): 226–232.
- Ingram, Jason R. 2007. "The Effect of Neighborhood Characteristics on Traffic Citation Practices of the Police." *Police Quarterly* 10(4): 371–393.
- Jefferson, Brian Jordan. 2018. "Predictable Policing: Predictive Crime Mapping and Geographies of Policing and Race." *Annals of the American Association of Geographers* 108(1): 1–16.
- 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).
- Johnson, David J., Trevor Tress, Nicole Burkel, Carley Taylor, and Joseph Cesario. 2020. "Retraction for Johnson et al., Officer Characteristics and Racial Disparities in Fatal Officer-Involved Shootings." *Proceedings of the National Academy of Sciences* 117(30): 18130–18130.
- Justice, Benjamin, and Tracey L. Meares. 2014. "How the Criminal Justice System Educates Citizens." *The Annals of the American Academy of Political and Social Science* 651(1): 159–177.
- Kim, E. Tammy. 2020. "The Perils of 'People of Color'."
- Knowles, John, Nicola Persico, and Petra Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy* 109(1): 203–229.
- Knox, Dean, and Jonathan Mummolo. 2020a. "Making inferences about racial disparities in police violence." *Proceedings of the National Academy of Sciences* 117(3): 1261–1262.
- Knox, Dean, and Jonathan Mummolo. 2020b. "Toward a General Causal Framework for the Study of Racial Bias in Policing." *Journal of Political Institutions and Political Economy* 1: 1–38.



- Knox, Dean, Will Lowe, and Jonathan Mummolo. 2020. "Administrative Records Mask Racially Biased Policing." *American Political Science Review* 114(3): 619–637.
- Koslicki, Wendy M., Daniel J. Lytle, Dale W. Willits, and Rachael Brooks. 2021. "'Rhetoric without Reality' or Effective Policing Strategy? An Analysis of the Relationship between Community Policing and Police Fatal Force." *Journal of Criminal Justice* 72: 101730.
- Koslicki, Wendy M., David A. Makin, and Dale Willits. 2020. "When No One is Watching: Evaluating the Impact of Body-Worn Cameras on Use of Force Incidents." *Policing and Society* 30(5): 569–582.
- Lange, James E., Mark B. Johnson, and Robert B. Voas. 2005. "Testing the Racial Profiling Hypothesis for Seemingly Disparate Traffic Stops on the New Jersey Turnpike." *Justice Quarterly* 22(2): 193–223.
- Laniyonu, Ayobami. 2019. "The Political Consequences of Policing: Evidence from New York City." *Political Behavior* 41(2): 527–558.
- Lau, Tim. 2020. *Predictive Policing Explained*. Brennan Center for Justice.
- Legewie, Joscha, and Jeffrey Fagan. 2019. "Aggressive Policing and the Educational Performance of Minority Youth." *American Sociological Review* 84(April): 220–247.
- Lerman, Amy E., and Vesla M. Weaver. 2014a. *Arresting Citizenship: The Democratic Consequences of American Crime Control*. Chicago: University of Chicago Press.
- Lerman, Amy E., and Vesla M. Weaver. 2016. "The Carceral State and American Political Development." *The Oxford Handbook of American Political Development* .
- Lerman, Amy E., and Vesla Weaver. 2014b. "Staying out of Sight? Concentrated Policing and Local Political Action." *The Annals of the American Academy of Political and Social Science* 651(1): 202–219.

- MacDonald, John, and Anthony A. Braga. 2019. "Did Post-Floyd et al. Reforms Reduce Racial Disparities in NYPD Stop, Question, and Frisk Practices? An Exploratory Analysis Using External and Internal Benchmarks." *Justice Quarterly* 36(5): 954–983.
- Makowsky, Michael D., Thomas Stratmann, and Alex Tabarrok. 2019. "To Serve and Collect: The Fiscal and Racial Determinants of Law Enforcement." *The Journal of Legal Studies* 48(1): 189–216.
- Manski, Charles F. 2006. "Search Profiling With Partial Knowledge of Deterrence." *Economic Journal* 116(515): F385–F401.
- McCrary, Justin. 2007. "The Effect of Court-Ordered Hiring Quotas on the Composition and Quality of Police." *The American Economic Review* 97(1): 37.
- Meehan, Albert J., and Michael C. Ponder. 2002. "Race and Place: The Ecology of Racial Profiling African American Motorists." *Justice Quarterly* 19(3): 399–430.
- Menifield, Charles E., Geiguen Shin, and Logan Strother. 2019. "Do White Law Enforcement Officers Target Minority Suspects?" *Public Administration Review* 79(1): 56–68.
- Mentch, Lucas. 2020. "On Racial Disparities in Recent Fatal Police Shootings." *Statistics and Public Policy* 7(1): 9–18.
- Mummolo, Jonathan. 2018. "Modern Police Tactics, Police-Citizen Interactions, and the Prospects for Reform." *The Journal of Politics* 80(1): 1–15.
- Nicholson-Crotty, Sean, Jill Nicholson-Crotty, and Sergio Fernandez. 2017. "Will More Black Cops Matter? Officer Race and Police-Involved Homicides of Black Citizens." *Public Administration Review* 77(2): 206–216.
- Nix, Justin, Bradley A. Campbell, Edward H. Byers, and Geoffrey P. Alpert. 2017. "A Bird's Eye View of Civilians Killed by Police in 2015 Civilians Killed by Police: Research Article." *Criminology and Public Policy* 16(1): 309–340.

- Ochs, Holona Leanne. 2011. “The Politics of Inclusion: Black Political Incorporation and the Use of Lethal Force.” *Journal of Ethnicity in Criminal Justice* 9(3): 238–265.
- Payne, B. Keith. 2001. “Prejudice and Perception: The Role of Automatic and Controlled Processes in Misperceiving a Weapon.” *Journal of personality and social psychology* 81(2): 181–192.
- Pearl, Judea. 2018. “Does Obesity Shorten Life? or Is It the Soda? on Non-Manipulable Causes.” *Journal of Causal Inference* 6(2): 1–7.
- 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.” *The Economic Journal* 116(515): F351–F367.
- Peyton, Kyle, Michael Sierra-Arévalo, and David G. Rand. 2019. “A Field Experiment on Community Policing and Police Legitimacy.” *Proceedings of the National Academy of Sciences* 116(40): 19894–19898.
- Phelps, Edmund S. 1972. “The Statistical Theory of Racism and Sexism.” *The American Economic Review* 62(1): 659–61.
- Pierson, Emma, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Daniel Jenson, Amy Shoemaker, Vignesh Ramachandran, Phoebe Barghouty, Cheryl Phillips, Ravi Shroff, and Sharad Goel. 2020. “A Large-Scale Analysis of Racial Disparities in Police Stops across the United States.” *Nature Human Behaviour* 4(7): 736–745.
- President’s Task Force on 21st Century Policing. 2015. *Final Report of the President’s Task Force on 21st Century Policing*. Office of Community Oriented Policing Services.
- Prowse, Gwen, Vesla M. Weaver, and Tracey L. Meares. 2020. “The State from Below: Distorted Responsiveness in Policed Communities.” *Urban Affairs Review* 56(5): 1423–1471.
- Ridgeway, Greg, and John M. MacDonald. 2009. “Doubly Robust Internal Benchmarking and False Discovery Rates for Detecting Racial Bias in Police Stops.” *Journal of the American Statistical Association* 104(486): 661–668.

- Ridgeway, Greg, and John M. MacDonald. 2010. "Methods for Assessing Racially Biased Policing." In *Race, Ethnicity, and Policing*, ed. S. K. Rice, and M. D. White. New York: New York University Press pp. 180–205.
- Rosenthal, Aaron. 2019. "Investment and Invisibility: The Racially Divergent Consequences of Political Trust." 16(2): 511–533.
- 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).
- Rubin, Donald B. 1974. "Estimating Causal Effects of Treatments in Randomized and Non-Randomized Studies." *Journal of Educational Psychology* 66(5): 688–701.
- Rushin, Stephen, and Griffin Edwards. 2021. "An Empirical Assessment of Pretextual Stops and Racial Profiling." *Stanford Law Review* 73(3): 637–726.
- Sances, Michael W., and Hye Young You. 2017. "Who Pays for Government? Descriptive Representation and Exploitative Revenue Sources." *The Journal of Politics* 79(3): 1090–1094.
- Sanga, Sarath. 2009. "Reconsidering Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy* 117(6): 1155–1159.
- Sen, Maya, and Omar Wasow. 2016. "Race as a Bundle of Sticks: Designs that Estimate Effects of Seemingly Immutable Characteristics." *Annual Review of Political Science* 19: 499–522.
- Sharp, Elaine B. 2014. "Minority Representation and Order Maintenance Policing: Toward a Contingent View." *Social Science Quarterly* 95(4): 1155–1171.
- Shelby, Tommie. 2018. *Dark Ghettos: Injustice, Dissent, and Reform*. Cambridge: Harvard University Press.
- Shjarback, John, Scott Decker, Jeff J. Rojek, and Rod K. Brunson. 2017. "Minority Representation in Policing and Racial Profiling: A Test of Representative Bureaucracy vs Community Context." *Policing* 40(4): 748–767.

- Shoub, Kelsey, Leah Christiani, Frank R. Baumgartner, Derek A. Epp, and Kevin Roach. 2020. "Fines, Fees, Forfeitures, and Disparities: A Link Between Municipal Reliance on Fines and Racial Disparities in Policing." *Policy Studies Journal* .
- Siegel, Michael. 2020. "Racial Disparities in Fatal Police Shootings: An Empirical Analysis Informed by Critical Race Theory Symposium: Beyond Bad Apples: Exploring the Legal Determinants of Police Violence." *Boston University Law Review* 100(3): 1069–1092.
- Silva, Andrea, Diego Esparza, Valerie Martinez-Ebers, and Regina Branton. 2020. "Perceived Police Performance, Racial Experiences, and Trust in Local Government." *Politics, Groups, and Identities* 0(0): 1–24.
- Simoiu, Camelia, Sam Corbett-Davies, and Sharad Goel. 2017. "The Problem of Infra-Marginality in Outcome Tests for Discrimination." *The Annals of Applied Statistics* 11(3): 1193–1216.
- Sklansky, David Alan. 2005. "Not Your Father's Police Department: Making Sense of the New Demographics of Law Enforcement Criminology." *Journal of Criminal Law and Criminology* 96(3): 1209–1244.
- Socia, Kelly M., Melissa S. Morabito, Brenda J. Bond, and Elias S. Nader. 2021. "Public Perceptions of Police Agency Fairness and the Willingness to Call Police." *The American Review of Public Administration* .
- 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(1): 565–591.
- Streeter, Shea. 2019. "Lethal Force in Black and White: Assessing Racial Disparities in the Circumstances of Police Killings." *The Journal of Politics* 81(3): 1124–1132.
- Tomaskovic-Devey, Donald, Marcinda Mason, and Matthew Zingraff. 2004. "Looking for the Driving While Black Phenomena: Conceptualizing Racial Bias Processes and Their Associated Distributions." *Police Quarterly* 7(1): 3–29.

- Trawalter, Sophie, Andrew R. Todd, Abigail A. Baird, and Jennifer A. Richeson. 2008. "Attending to Threat: Race-Based Patterns of Selective Attention." *Journal of Experimental Social Psychology* 44(5): 1322–1327.
- Tregle, Brandon, Justin Nix, and Geoffrey P. Alpert. 2019. "Disparity Does Not Mean Bias: Making Sense of Observed Racial Disparities in Fatal Officer-Involved Shootings with Multiple Benchmarks." *Journal of Crime and Justice* 42(1): 18–31.
- Tyler, Tom R. 2005. "Policing in Black and White: Ethnic Group Differences in Trust and Confidence in the Police." *Police Quarterly* 8(3): 322–342.
- Walker, Hannah L. 2014. "Extending the Effects of the Carceral State: Proximal Contact, Political Participation, and Race." *Political Research Quarterly* 67(4): 809–822.
- Walker, Samuel. 2001. "Searching for the Denominator: Problems with Police Traffic Stop Data and an Early Warning System Solution." *Justice Research and Policy* 3(1): 63–95.
- Weaver, Vesla M., and Amy E. Lerman. 2010. "Political Consequences of the Carceral State." *American Political Science Review* 104(4): 817–833.
- Wright, James E., Dongfang Gaozhao, and Meagan A. Snow. 2021. "Place plus Race Effects in Bureaucratic Discretionary Power: An Analysis of Residential Segregation and Police Stop Decisions." *Public Performance & Management Review* 44(2): 352–377.