

Racial Bias in Police Investigations

Jeremy West*

MIT

November 2015

JOB MARKET PAPER

Abstract

Law enforcement discrimination studies are usually qualified by selection concerns regarding which encounters transpire between individuals and police. This paper overcomes these problems by examining automobile crash investigations by a State Police Department. Because officers are dispatched to investigate crashes on the basis of factors unrelated to driver race, these interactions are effectively exogenous. I show that, conditional on the Census Block Group area of a crash, the race of the investigating officer is uncorrelated with that of the driver. For these investigations, I find that police officers exhibit significant bias in issuing citations to other-race drivers for both moving and nonmoving violations, though no bias is evident for felony violations such as hit-and-run. Because racial bias is present even for easily observable offenses such as expired vehicle registration, these findings are consistent with preference-based rather than statistical discrimination.

*I am a postdoctoral associate in the Department of Economics at the Massachusetts Institute of Technology. Email: westj@mit.edu. I am grateful for valuable comments from Joshua Angrist, David Autor, Jonathan Gruber, Mark Hoekstra, Sally Hudson, Christopher Knittel, Jonathan Meer, James Poterba, Steven Puller, and numerous seminar participants. Any opinions and conclusions in this study are my own and do not necessarily represent the position of the police agency I study, the State in which it operates, or any other State agencies. Any errors in this paper are also my own. I have no relevant or material financial or legal interests related to the research described in this paper.

1 Introduction

Commentators have long argued that law enforcement is characterized by racial prejudice, or what Gary Becker (1957) termed “preference-based” discrimination (e.g. Horsey, 2014). This sentiment is embodied in catchphrases such as “driving while black” (Harris, 1999) and the “black lives matter” movement (Ross, 2015), and it is shared by much of the public (CNN/ORC Poll, 2014). Policy focus on this question has escalated as recent events “attracted international attention and added momentum to a national debate over the treatment of black people by white police officers” (Montgomery, 2015).

In contrast to public opinion, the empirical literature is much less conclusive that racial prejudice is a systemic factor in law enforcement. The scholarly literature is clear regarding the *existence* of disparities in how police officers treat individuals of differing races. For instance, Kochel, Wilson, and Mastroski’s (2011) meta study finds that “with strong consistency, minority suspects are more likely to be arrested than white suspects.” But, the literature is very mixed regarding whether preference-based discrimination explains any cross-racial differences in law enforcement outcomes. One possibility is that officers are discriminating based on their personal preferences; an alternate explanation is that officers use racial profiling, or statistical discrimination (Arrow, 1973), to improve the effectiveness of law enforcement.

Researchers have tested for racial bias in a variety of police settings, including vehicle drug searches (Knowles, Persico, and Todd, 2001; Hernández-Murillo and Knowles, 2004; Anwar and Fang, 2006; Close and Mason, 2007; Antonovics and Knight, 2009; Sanga, 2009), stop-and-frisk (Gelman, Fagan, and Kiss, 2007), speeding citations (Anbarci and Lee, 2014), and driving-under-the-influence enforcement (Horn, McCluskey, and Mittelhammer, 2014).¹ The results range broadly, from finding preference-based discrimination, statistical discrimination only, to no evidence of any racial bias.

Whereas methodologies and conclusions vary extensively in these studies, a nearly-universal factor in the literature is that analysts only observe data on encounters that police officers decide to initiate: if an officer does not stop a driver or pedestrian, then this (non)encounter is never recorded. Such endogeneity of police encounters could substantially affect researchers’ ability to quantify racial bias. Anwar and Fang (2006) note that “the trooper

¹This is only a sample from this literature. Some examples of non-police studies of racial bias include automobile sales (Ayres and Siegelman, 1995; Langer, 2012), mortgage lending (Munnell, Tootell, Browne, and McEneaney, 1996), employment (Altonji and Blank, 1999), airport security (Persico and Todd, 2005), sports officiating (Price and Wolfers, 2010; Parsons, Sulaeman, Yates, and Hamermesh, 2011), jury criminal trials (Anwar, Bayer, and Hjalmarsson, 2012), and rental apartments (Ewens, Tomlin, and Wang, 2014).

must first stop the motorist prior to a search,” allowing the possibility “that the racial prejudice of police officers is reflected in their stop decisions as well as (or instead of) their search decisions.” [Dharmapala and Ross \(2004\)](#) show that “because potential offenders are frequently not observed by the police,” the same data are “consistent with prejudice against African-American males, with no prejudice, and with reverse discrimination, depending on the assumption[s].” [Grogger and Ridgeway \(2006\)](#) and [Horrace and Rohlin \(2015\)](#) use the “veil of darkness” surrounding dusk as a clever attempt to account explicitly for this endogeneity in traffic stops.

In this paper, I overcome these endogeneity concerns by studying racial bias in automobile crash investigations, a setting in which police officers do not initiate encounters with civilians. Because officers are dispatched to investigate crashes, and this dispatch is based on factors unrelated to driver race, these officer-driver racial combinations are effectively random. I test the exogeneity of these interactions using data on 440,000 crash investigations by a single State Police Department (SPD) during years 2006-2012. When conditioning on the Census Block Group location of an automobile crash – which is necessary in order to account for cross-community variation in the racial composition of residents and police officers – I show that the race of the dispatched SPD officer is uncorrelated with that of the driver(s) involved in the crash.

This effectively random assignment affords a causal interpretation of how officer-driver racial combinations affect police behavior such as issuing citations for traffic violations. I quantify racial bias in these encounters using a generalized difference in differences framework, which adjusts for racial variations in drivers’ propensities to commit, and officers’ propensities to issue citations for, traffic infractions. This approach is commonly used in the broader literature as a test for racial bias (e.g. [Price and Wolfers, 2010](#)). Though it identifies a causal relationship under minimal assumptions, the limitation of this approach is that it can only detect *relative* bias and cannot distinguish whether this bias is attributable to the actions of white or minority police officers.

The results show significant racial bias in traffic citations. The difference in differences estimates indicate that police officers issue citations to other-race drivers about three percentage points more often than they do to own-race drivers (on a mean citation rate of 45 percent). This finding is very robust and holds when including detailed controls specific to the crash or driver, to using within-officer variation, and even to using within-crash variation – when drivers of differing race crash with each other, the race of the exogenously dispatched officer plays an important role in determining which (if any) driver is cited.

Additional results indicate that this pattern of behavior is best explained by preference-based discrimination, rather than statistical discrimination. Estimating the difference in differences by type of citation shows that officers are more likely to issue citations to other-race drivers for both moving violations, e.g. speeding or following too closely, and nonmoving violations, such as expired vehicle registration.² The existence of racial bias in citations for nonmoving violations – essentially, expired documents – is particularly compelling evidence of preference-based discrimination. These infractions are highly salient from the perspective of the officer, who must record details of driver and vehicle documents as a routine part of any crash investigation. Additionally, these charges offer no scope for subjectivity in evaluation as a document is either expired or it is current. It is challenging to reconcile own-race bias in citations for these violations as any form of statistical discrimination, even if more generally allowing for false beliefs or incomplete information.

I also assess heterogeneity in racial bias by driver and crash characteristics, and across areas in which the SPD operates. These analyses support that this racial bias is systemic, rather than being concentrated among a small subset of officers or drivers. The literature notes concerns over possible reciprocity in the aggressiveness of individuals and the officials who evaluate them (e.g. [Reisig, McCluskey, Mastrofski, and Terrill, 2004](#); [Price and Wolfers, 2010](#)). That the estimates are qualitatively invariant to driver and crash characteristics provides further support for a mechanism of police officer preference-based discrimination. It is hard to defend a driver-based explanation for these own-race effects that would hold across a large range of driver demographics and crash settings. Examining spatial heterogeneity in local average treatment effects also underscores that the racial bias is not attributable to only a small subset of SPD officers.

To summarize, I document significant and systemic racial bias by police officers that is consistent with preference-based rather than statistical discrimination. The extent to which the results documented here apply in other police settings is an open question, but the fact that the behavior appears to be both systemic and preference-based supports a more general argument regarding racial discrimination in law enforcement. Clearly demonstrating the existence of this type of behavior in a police agency serves as a step towards policies that encourage more transparent, effective, and equitable policing.

²I find no evidence of racial bias in citations for felony violations such as hit-and-run or vehicular assault. It is unclear whether this indicates a lack of racial bias in high-stakes situations or increased oversight of officers' responses for these types of offenses. Due to this ambiguity, as well as a low baseline rate, little inference should be drawn regarding the apparent lack of racial bias in citations for felony offenses.

2 Data

The data in my study include seven years (2006-2012) of State Police Department (SPD) motor vehicle crash investigations within a single state of the United States, with the investigating police officers linked to administrative employee personnel records. Data arrangements preclude me from disclosing which specific state is in this study.

2.1 State Police crash investigations and personnel records

The automobile crash data initially consist of the population of police-reported vehicle crashes handled by SPD. Consistent with their role as State Police, this subset of vehicle crashes is largely outside of the jurisdictions of municipal police agencies, meaning that they are somewhat less urban and more concentrated on Interstates and highways.³ The crash data are a very rich dataset, including detailed information about the crash (location, weather and road conditions, etc.), the involved vehicles, the drivers and passengers (demographics, injury status), and any citations issued by the investigating officer.

The other source of data for this study is administrative employee records maintained by the state's personnel office. These records, which the state provided under a Freedom of Information Act, include each employee's full name, job title, hire date, race, and gender.

2.2 Data merging and sample formation

I merge the crash records to the personnel records by string-matching the investigating officers' names. I prioritize avoiding measurement error from false positives, and thus take a very conservative approach in the matching process. In particular, I never allow for there to be simultaneous entry errors in both the officer's name and badge number for a crash record, and I only incorporate matches for which the personnel data include a unique record of that full name, which excludes several names that have a popular first name and a very common surname.

Ultimately, I am able to match 81 percent of crashes to a unique officer in the personnel records; some of this is due to excluding ambiguous full name matches, but spot-checks of random samples of unmatched records indicate that the vast majority of the unmatched

³The official jurisdiction of the State Police extends statewide, but State Police primarily operate in areas outside the jurisdictions of local agencies such as municipal police and sheriff departments.

19 percent are due to significant data entry errors.⁴ My identification strategy also requires dropping records with missing data on the crash location or driver race. Appendix A describes the details of the data cleaning and string-matching process and rules out the possibility of data restrictions causing the final analysis sample to be subset in a meaningfully nonrandom way.

Finally, because of statistical power considerations – as I am conditioning on small geographic areas – I restrict my analysis to interactions in which the driver and the officer are either white, black, or Hispanic (in both data sets, Hispanic/Latino is coded as its own race, and throughout this article I refer simply to race rather than ethnicity). In the crashes handled by SPD, this is a very minor restriction.

2.3 Summary statistics

My final sample for analysis consists of roughly 440,000 encounters between a driver and a police investigator. Summary statistics are in Table 1. The first panel presents statistics of the drivers and crashes. Overall, white police officers are dispatched to 68.6 percent of crashes statewide, though there is considerable heterogeneity across the state (as shown in Appendix Figure B.1). Similarly, about 68 percent of drivers in crashes are white. More than 95 percent of crashes are investigated by a male officer, precluding a similar exercise by gender as I do in this study by race. About a third of driver-crashes are single-vehicle, which may seem a high portion, but recall that these are State Police who handle primarily highway collisions. Roughly half of driver-crashes involve two drivers; the remaining 20 percent involve three or more vehicles.

The second panel shows variation in citations by driver characteristics. Unsurprisingly, male drivers and younger drivers are cited comparatively more often than female and older drivers. Hispanic drivers are cited about 50 percent of the time on average, but white drivers and black drivers are both cited about 43 percent of the time. Thus, a naive comparison of the raw data for black and white drivers would yield a conclusion of no racial bias. This motivates the need for a stronger identification strategy, as discussed in the next section.

My primary analysis includes all citations, but I additionally estimate racial bias by type of citation. For these analyses, I use three mutually exclusive citation categories: nonmoving violations, moving violations, and felony violations. Nonmoving violations consist of

⁴The degree of entry errors in the crash data is somewhat staggering. For instance, one investigator’s name was misspelled in 72 *uniquely* different ways. Fortunately for this study, even if certain names are more likely to be entered incorrectly, there is no reason to think that entry errors for the same investigator should vary with the race of the driver involved in the crash.

expired/nonexistent driver license and expired vehicle registration, inspection, or insurance. Moving violations consist of charges related to passing, right-of-way, signal intention, speeding, traffic signs, seatbelt, equipment defects, and other miscellaneous moving violations. Felony violations consist of vehicular assault, manslaughter, and hit-and-run, which are all potentially felony offenses.

3 Methods

My identification strategy consists of two steps. First I show that, conditional on a sufficiently small spatial area, the race of the dispatched investigating officer is uncorrelated with that of a driver involved in a crash. Second, I use a generalized difference in differences strategy to identify the effect on citations of being exogenously assigned a same-race officer relative to an officer of differing race.

3.1 Demonstrate exogenous officer assignment

The causal inference of this study is facilitated by exogenous assignment of officer-driver interactions. Unlike in a traffic stop, police officers are dispatched to investigate crashes on the basis of their proximity to the scene and on other factors unrelated to driver race.

This provides effectively random assignment of officer-driver racial combinations, but only when conditioning on relatively small geographies as, generally speaking, areas which have higher prevalence of minority drivers are also likely to have higher prevalence of minority police officers.⁵ Ignoring the location of the crash, a black driver is more likely than a white driver to be dispatched a black officer, and a Hispanic driver is more likely than a white driver to see a Hispanic investigator. By conditioning on a sufficiently fine area such as a Census Block Group – roughly equivalent in size to a typical neighborhood – I control for these cross-community correlations to recover the effectively random assignment of officer-driver racial combinations. I explicitly test the exogeneity of this assignment using linear probability regressions such as:

$$I\{\text{white officer}\}_{aij} = \alpha_B \cdot I\{\text{black driver}\}_i + \alpha_H \cdot I\{\text{Hispanic driver}\}_i + \mu_a + \epsilon_{aij} \quad (1)$$

In Equation (1), $I\{\text{white officer}\}_{aij}$ is an indicator variable for whether officer j responding

⁵As Keating, Badger, and Elliott (2014) demonstrate, this is far from a universal truth, but there is certainly a positive cross-community correlation in the races of residents and police officers.

to a crash of driver i in area a is white. As I include only white, black, and Hispanic drivers in my analysis, α_B and α_H respectively identify how much more or less likely black and Hispanic drivers are than white drivers to be assigned a white crash investigator. If the dispatch of officers within a given small area is invariant to driver race (as should be expected), then these coefficients will equal zero when controlling for Block Group fixed effects μ_a .⁶

3.2 Difference in differences to identify own-race bias

This effectively random assignment of officer-driver racial combinations affords a causal interpretation of how each *racial combination* affects crash outcomes such as driver citations. One naive approach to estimating racial bias is to simply compare citation rates for drivers of a given race across different officer races. For example, one could compare how often minority drivers are cited by white officers to how often minority drivers are cited by minority officers. But, this approach ignores that officers of differing races may vary in their propensities to cite drivers regardless of driver race. For instance, perhaps white officers are more likely to issue citations to *all* drivers, not just to minority drivers.

Instead, a second naive approach is to compare how often officers of a given race issue citations to drivers of different races. For example, one could compare how often white officers cite white drivers to how often white officers cite minority drivers. This approach circumvents the limitation of the first one. But, it ignores that drivers of differing races may not have the same true propensities to be guilty of infractions.

The difference in differences method combines both of these naive approaches to leverage their strengths while overcoming both shortcomings. This approach is often used in the broader literature as a test for racial bias (Antonovics and Knight, 2009; Price and Wolfers, 2010; Parsons, Sulaeman, Yates, and Hamermesh, 2011; Anbarci and Lee, 2014). I first consider the two-by-two case which includes only officers and drivers who are either white or a *particular* racial minority, an approach which uses as a specification:

$$\begin{aligned} \text{I}\{\text{citation}\}_{aij} &= \beta_M \cdot \text{I}\{\text{minority driver}\}_i + \gamma_M \cdot \text{I}\{\text{minority officer}\}_j \\ &+ \delta_M \cdot \text{I}\{\text{minority driver} \& \text{white officer}\}_{ij} + \mu_a + \text{controls}_{ij}\Psi + \epsilon_{aij} \end{aligned} \quad (2)$$

⁶An implicit assumption in this approach is that only one officer is dispatched to investigate each crash. Unfortunately, the data do not consistently indicate all of the officers who are present. Multiple officers are listed in fewer than one percent of crash investigations, typically very serious multivehicle collisions; I exclude these crashes from my analysis (they are included in the “unmatched” 19 percent of records). My understanding is that records which list a single officer include the officer who is the decision-making agent, and any additional officers at the crash scene are present to direct traffic and handle other supporting activities.

The β_M coefficient captures differences in citation rates of minority drivers by minority officers relative to white drivers by minority officers, allowing for drivers to differ in their overall propensities to deserve citations. The γ_M coefficient captures differences in citation rates of white drivers by minority officers relative to white drivers by white officers, allowing for officers to differ in their overall propensities to issue citations. δ_M is the difference in differences coefficient of interest and captures the degree of racial bias. If there is no racial bias, then $\delta_M = 0$. A positive sign for δ_M indicates own-race leniency (or other-race harshness) in citations.⁷

Equation (2) facilitates a two-by-two comparison of either black/white or Hispanic/white racial bias in citations. For most of the empirical results in this article, I pool the data to include all three racial groups, rather than conducting separate two race comparisons. I do this primarily for statistical power, a significant consideration especially for results which examine heterogeneity in types of citations or by characteristics of drivers or crashes (many of which have relatively small numbers of observations). Additionally, pooling the data allows for a single coefficient of interest for each relationship and thereby heightens clarity in inference and exposition. The pooled estimation replaces Equation (2) with the following:

$$\begin{aligned}
\text{I}\{\text{citation}\}_{aij} &= \beta_B \cdot \text{I}\{\text{black driver}\}_i + \gamma_B \cdot \text{I}\{\text{black officer}\}_j \\
&+ \beta_H \cdot \text{I}\{\text{Hispanic driver}\}_i + \gamma_H \cdot \text{I}\{\text{Hispanic officer}\}_j \\
&+ \delta \cdot \text{I}\{\text{black or Hispanic driver \& other-race officer}\}_{ij} \\
&+ \mu_a + \text{controls}_{ij}\Psi + \epsilon_{aij}
\end{aligned} \tag{3}$$

Equation (3) retains the difference in differences framework, but collapses the estimation into a single specification. The coefficient of interest δ continues to quantify the degree of racial bias, just as δ_B does for black/white or δ_H does for Hispanic/white comparisons using Equation (2), and in practice this δ coefficient is effectively a weighted average of the δ_B and δ_H estimates.⁸ I cluster standard errors for all estimations at the officer level, as this seems ex ante to be the appropriate cluster for any autocorrelation concerns. In practice, this is also a more conservative approach and yields larger standard errors than those from clustering by other groupings such as Census Block Group.

This difference in differences estimator identifies racial bias under a minimal set of as-

⁷The difference in differences framework provides only a symmetric and relative measure of racial bias. That is, the same racial bias estimate as in Equation (2) is obtained from estimating: $\text{I}\{\text{citation}\}_{aij} = \beta_W \cdot \text{I}\{\text{white driver}\}_i + \gamma_W \cdot \text{I}\{\text{white officer}\}_j + \delta_W \cdot \text{I}\{\text{white driver \& minority officer}\}_{ij} + \mu_a + \text{controls}_{ij}\Psi + \epsilon_{aij}$.

⁸Strictly speaking, δ is not purely a weighted average of δ_B and δ_H , because black drivers are sometimes assigned Hispanic officers, and vice versa, but white officers handle the large majority of investigations.

sumptions. The limitation of this method is that it can only identify *relative* racial bias, not absolute racial bias. A positive value for δ indicates that officers cite other-race drivers relatively more than they cite own-race drivers, adjusted for racial propensities of drivers to receive citations and officers to issue citations, but this estimate cannot distinguish whether this is due to white officers citing minority drivers more harshly or minority officers citing minority drivers more leniently. I join others such as [Price and Wolfers \(2010\)](#) in acknowledging this limitation, but view it as a worthwhile tradeoff in order to obtain clear causal evidence regarding the existence of racial bias in law enforcement.

4 Results

As discussed in Section 3, the institutional setting provides for exogenous assignment of encounters between police officers and drivers. My identification strategy is to demonstrate this effectively random assignment of police officers to crash investigations and use difference in differences for these exogenously-formed matches to estimate the racial bias in citations.

4.1 Officers are exogenously assigned to crashes

I demonstrate the effectively random assignment of officer-driver combinations in Table 2. The table is divided into three panels, respectively testing how assignment of white, black, or Hispanic officers varies with driver race. Each panel-column presents estimates from a separate linear probability regression, with all specifications corresponding generally to Equation (1).

Column (1) of Table 2 includes no fixed effects or controls. As the large and highly statistically significant coefficients indicate, there are strong correlations between driver race and officer race in the state overall. Specifically, drivers of all races are relatively more likely to be assigned a same-race officer. As discussed in Section 3, this expected result is due to cross-community differences in the racial composition of local residents and police officers.

In Column (2), which adds fixed effects only for the Census Block Group of each crash, and no other controls, the coefficients are all very small in magnitude and statistically indistinguishable from zero. This is clear evidence of passing the test for exogenous assignment of officer-driver racial combinations: the race of the dispatched officer is completely uncorrelated with that of the driver in the crash. Unsurprisingly, this uncorrelated relationship continues to hold when adding in control terms specific to the crash and driver in Columns

(3) and (4), and to switching to Block Group by year spatial fixed effects in Column (5).⁹ This serves as very compelling evidence of effectively random assignment of officer-driver interactions and provides a clean framework for causal inference in the analyses of racial bias that follow.

4.2 Officers issue fewer own-race citations

Given this exogenous assignment of officer-driver racial combinations, I estimate racial bias using difference in differences. Table 3 presents three comparisons. Panel [A] includes only black and white drivers and officers, and Panel [B] includes only Hispanic and white drivers and officers, so the specifications for both of these panels directly correspond to Equation (2). Panel [C] pools all three racial groups using Equation (3). All three panels include fixed effects for driver race and officer race to facilitate the difference in differences identification, and all columns include at least Block Group fixed effects to maintain the exogenous assignment demonstrated in Table 2. Figure 1 uses results from the first two panels to graphically show the likelihood of citation by racial combination.

Column (1) of Table 3 includes Block Group fixed effects and no other controls. The estimated degree of racial bias for citations in the black/white comparison is 1.75 percentage points. That is, after adjusting for race-specific differences in propensities for drivers to commit infractions and adjusting for race-specific differences in officers' propensities to issue citations, a black driver who is exogenously assigned a white officer is 1.75 percentage points more likely to be cited relative to a white driver who is exogenously assigned a white officer. For the Hispanic/white comparison in Panel [B], the estimate is qualitatively similar, though somewhat larger in magnitude, at 2.96 percentage points. And, as expected of what is essentially a weighted average, the pooled racial bias estimate of 2.61 percentage points in Panel [C] falls in between those from Panels [A] and [B] in magnitude. All three results are statistically significant at the five percent level. In addition, the estimates are economically significant. Given a mean of 44.88 percent, the pooled result indicates that racial bias causes a black or Hispanic driver with an other-race officer to be nearly six percent more likely to be given any citation in a crash investigation.

The additional columns address potential confounders to this racial bias estimate. For ex-

⁹One potential concern is that these coefficients show a lack of correlation only because there is *just enough* identifying variation. At the extreme, if there was only one officer per Census Block Group, then of course a driver's race would not predict that of the dispatched officer. An indirect rebuttal is that such a lack of identifying variation would be reflected in the standard errors, both here and in the other results tables. A more direct demonstration of sufficient variation is provided by Appendix Figure B.1.

ample, despite an overall exogenous assignment of racial interactions, if officers of a particular race are relatively more experienced, then they might be assigned to investigate disproportionately more severe crashes. And, if drivers of a different race are disproportionately likely to be involved in more severe crashes, then failure to control for crash severity could result in a spurious racial bias estimate. Columns (2) - (3) add crash-specific and driver-specific controls and Column (4) changes the spatial fixed effect to Block Group by year; results in these columns remain very similar to those in Column (1). In fact, the estimates are largely unchanged even when including officer-specific fixed effects in Column (5), which identifies racial bias in how the *same officer* cites own-race versus other-race drivers. Across these five columns, the black/white estimates range from 1.71 to 2.27 percentage points, and the Hispanic/white estimates from 2.96 to 3.31 percentage points. The pooled estimates range from 2.61 to 2.95 percentage points. All fifteen of these estimates are statistically significant at $p < 0.05$.¹⁰

Finally, I include crash-specific fixed effects in Column (6). This alternative test for racial bias uses variation in citations by the same officer *at the same crash*. As such, the identifying variation for these regressions requires a multi-vehicle crash in which the drivers differ in race. Difference in differences estimates for racial bias using these specifications are if anything larger. This means that if two drivers of different races crash with each other, the race of the exogenously assigned officer plays an important role in determining which (if any) driver receives a citation. The data provide much less identifying variation in these specifications and the standard errors become appreciably larger. The black/white estimate is no longer statistically significant at conventional levels, but the estimates for Hispanic/white and the pooled comparison remain significant. Broadly, it is reassuring to see such qualitatively similar results from an identification strategy that leverages a rather different source of variation and captures so much unobserved heterogeneity.

To summarize, the primary finding of this study is that there is meaningful racial bias by police officers in a context in which officers are effectively randomly assigned to interact with civilians. The remainder of this article examines heterogeneity in these effects to better characterize the nature of this racial bias and shed light on the likely mechanisms.

¹⁰I cluster standard errors for all estimations at the officer level, as this seems *ex ante* to be the appropriate cluster for any autocorrelation concerns. In practice, this is also a more conservative approach and yields larger standard errors than those from clustering by other groupings such as Census Block Group.

4.3 Racial bias exists for both moving and nonmoving violations

A common concern in studies of racial bias is that it is generally difficult to distinguish preference-based discrimination from statistical discrimination that is founded on false beliefs (Georgakopoulos, 2004; Burke, 2007). For example, Gelman, Fagan, and Kiss (2007) estimate that minorities are more likely to be targeted for stop-and-frisk in New York City, even controlling for race-specific rates of crime participation. If NYC police officers are misinformed of the true propensities of different racial groups to commit crimes, this incorrect belief could reconcile their behavior as being statistical discrimination rather than prejudice.

Fortunately, crash investigations as a setting provide opportunity to examine racial bias in behavior that is very difficult to reconcile as being founded on incorrect information or false beliefs: citations for nonmoving violations, essentially expired documents. These types of violations are highly salient and objective from the perspective of the investigating officer. For instance, one can determine the validity of vehicle registration simply by looking at the exterior of a vehicle, and all of these documents have a clearly indicated expiration date. Moreover, details about driver license, vehicle registration, and insurance are required to be recorded by the officer as a routine part of any crash investigation.

Table 4 presents results, with specifications directly corresponding to those in Panel [C] of Table 3, separately for nonmoving violations, moving violations, and felony violations.¹¹ There are significant own-race effects in citations for both moving and nonmoving violations. Relative to the means for each category, the effect is much larger for nonmoving violations than for moving violations. Specifically, the racial bias in citations for nonmoving violations is nearly 15 percent (1.64 percentage points on a mean of 10.95 percent); that for moving violations is around 6 percent. As with the racial bias in overall citations, results remain very stable and statistically significant across the specifications. This evidence of racial bias in citations for highly salient and clearly objective violations is strong evidence of preference-based discrimination.

I find very small and statistically insignificant effects for felony violations. There are several candidate explanations for the lack of an effect on felony citations. One possibility is that police are not racially biased in cases that have the most at stake. An alternate explanation is that police responses to felony-level offenses are given more official scrutiny

¹¹Nonmoving violations consist of expired/nonexistent driver license and expired vehicle registration, inspection, or insurance. Moving violations consist of charges related to passing, right-of-way, signal intention, speeding, traffic signs, seatbelt, equipment defects, and other miscellaneous moving violations. Felony violations consist of vehicular assault, manslaughter, and hit-and-run, which are all potentially felony offenses. Results for each of the fourteen underlying charge categories are in Appendix Table B.1.

and that these cases have higher chance of resulting in lawsuits and other judicial activity during which officers' actions might be questioned. This ambiguity in possible explanations – along with the low baseline citation rate for felony violations – precludes drawing much inference from the apparent lack of racial bias in citations for felony violations.

4.4 The racial bias appears to be systemic

Next, I explore heterogeneity in racial bias by types of drivers or crashes. The literature has suggested that there might be reciprocity in the aggressiveness of suspects and police officers (e.g. [Reisig, McCluskey, Mastroski, and Terrill, 2004](#)). In their study of racial bias by National Basketball Association referees, [Price and Wolfers \(2010\)](#) consider – and rule out – that the racial bias they document could be instead attributed to NBA players adapting their behavior in response to the assigned refereeing crew. Unlike in a basketball game, the nature of an automobile crash means that any infractions actually committed must occur prior to the individual's interaction with the official. However, it is still possible that the racial combination affects how aggressively a driver acts towards the police officer, and that such aggressiveness in turn influences whether the officer writes a citation for a known violation. If true, this would mean that the results might be attributable to *driver* behavior rather than racial bias by police officers.

Table 5 presents estimates for different types of citations using subsets of drivers or crashes. Male drivers appear to receive somewhat more racial bias than female drivers, and younger (than 25 years) drivers appear to receive somewhat more racial bias than older drivers, but these groups also have comparatively larger baseline citation rates. The extent of racial bias appears to be somewhat larger during the “daytime” hours (6:00 a.m. - 9:00 p.m.) relative to nighttime hours. And, the bias is greater in the last week of each month, potentially consistent with officers issuing additional citations in order to meet quotas. However, the broad takeaway from Table 5 is that the racial bias in citations seems systemic across different types of drivers and crashes. Qualitatively, the estimates are invariant across driver and crash characteristics, and this systemic nature of the results strengthens inference about the mechanism of the racial bias. It is hard to reconcile driver-based explanations with results that hold so broadly across demographic groups and crash settings.

The other source of heterogeneity I explore is by the location of the crash. To do so, I interact the “black or Hispanic driver and other-race officer” indicator with each Block Group indicator to estimate Block Group specific coefficients (local average treatment effects) for racial bias. The resulting distribution is presented in Figure 2. The mean of the distribution

is just under three percentage points, as reflected in the average treatment effect estimates in Panel [C] of Table 3. The weight in the tails is due to the idiosyncratic draws of which crashes occur in each Block Group – in particular, a number of Block Groups have relatively few minority drivers or officers. The interesting aspect of this distribution is that it is basically a normal distribution centered to the right of zero. This is additional (strong) evidence that the racial bias I estimate is a systemic effect, not attributable to a just a handful of officers.

5 Conclusions

A fundamental question of significant social consequence is the extent to which law enforcement is characterized by [Becker \(1957\)](#)-type preference-based discrimination. A large and active literature seeks to identify and quantify the degree and nature of racial bias by police officers. These studies share in common a substantial limitation in that encounters between individuals and police are endogenous.

I overcome these endogeneity concerns by examining police officer behavior in automobile crash investigations. Because officers are dispatched to investigate crashes, and this dispatch is based on factors unrelated to driver race, these interactions are effectively random. For automobile crash investigations handled by a State Police Department during 2006-2012, I show that conditional on a sufficiently fine geographic area, the race of the dispatched officer is uncorrelated with that of the driver. This allows for a causal interpretation of how the assigned racial combination affects police behavior. Via a difference in differences identification strategy, I use this exogenous assignment to quantify racial bias in citations for traffic violations.

For these investigations, I find that police officers exhibit significant bias in issuing citations to other-race drivers for both moving and nonmoving violations. The bias in citations for nonmoving violations – essentially, expired paperwork, an infraction that is salient and indisputable – is particularly compelling evidence of preference-based discrimination. In addition, I find that the racial bias appears to be largely systemic throughout areas in which the police agency operates, and that it is qualitatively invariant to driver or crash characteristics. Overall, the findings serve as clear evidence of systemic racial bias by police officers that is consistent with preference-based rather than statistical discrimination.

At the surface level, these findings may appear to have fairly minor economic consequence: police officers' bias in crash investigations causes some individuals to be about six percent more likely to receive traffic citations of several hundred dollars. There are of course

indirect consequences – for example, even small citations add “points” on a driver’s license and exacerbate existing disparities in minorities’ access to reliable transportation (Gautier and Zenou, 2010). However, the fact that the behavior appears to be both systemic and preference-based supports a more general argument regarding racial discrimination in law enforcement, some aspects of which have substantial, even lethal, consequence.

These results of systemic preference-based racial bias in law enforcement have significant policy implications. Societal attention is focused on the behavior of police officers, particularly in respect to their interactions with racial minorities. Many policies currently being explored, such as mandatory officer body cameras, are motivated in part by an assumption that statistical discrimination alone does not fully explain the cross-racial disparities in undesirable outcomes from police encounters (Harvard Law Review, 2015). The results of this study suggest that these disparities are unlikely to be merely an artifact of attempts to improve police efficiency, providing support for policies to change the practice of law enforcement.

References

- Joseph G. Altonji and Rebecca M. Blank. Race and gender in the labor market (chapter 48). *Handbook of Labor Economics*, 3(C):3143–3259, 1999.
- Nejat Anbarci and Jungmin Lee. Detecting racial bias in speed discounting: Evidence from speeding tickets in Boston. *International Review of Law and Economics*, 38:11–24, 2014.
- Kate Antonovics and Brian G. Knight. A new look at racial profiling: Evidence from the Boston police department. *The Review of Economics and Statistics*, 91(1):163–177, February 2009.
- Shamena Anwar and Hanming Fang. An alternative test of racial prejudice in motor vehicle searches: Theory and evidence. *The American Economic Review*, 96(1):127–151, March 2006.
- Shamena Anwar, Patrick Bayer, and Randi Hjalmarsson. The impact of jury race in criminal trials. *The Quarterly Journal of Economics*, 127(2):1017–1055, May 2012.
- Kenneth J. Arrow. The theory of discrimination. In Orley Ashenfelter and Albert Rees, editors, *Discrimination in Labor Markets*, pages 3–33. Princeton University Press, 1973.
- Ian Ayres and Peter Siegelman. Race and gender discrimination in bargaining for a new car. *The American Economic Review*, 85(3):304–321, 1995.
- Gary S. Becker. *The Economics of Discrimination*. University of Chicago Press, Chicago, 1957.

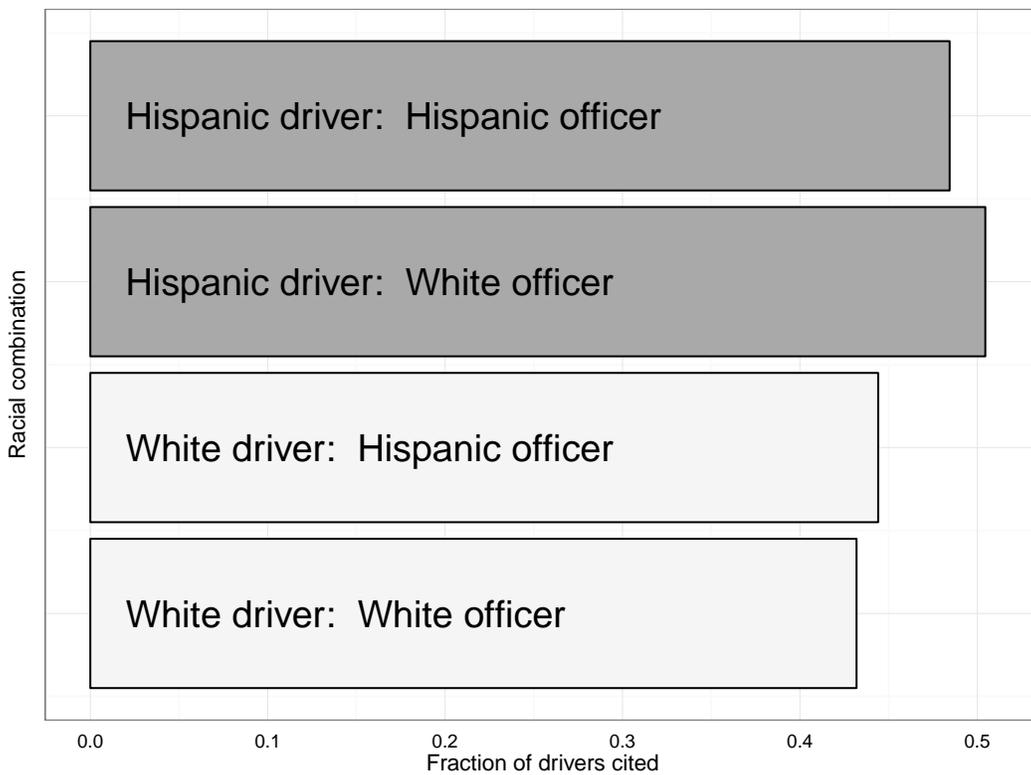
- Alafair Burke. Neutralizing cognitive bias: An invitation to prosecutors. *NYU Journal of Law and Liberty*, 2:512–530, 2007.
- Billy R. Close and Patrick Leon Mason. Searching for efficient enforcement: Officer characteristics and racially biased policing. *Review of Law & Economics*, 3(2):263–321, 2007.
- CNN/ORC Poll. CNN and Opinion Research Corporation poll: December 18-21, 2014, December 2014.
- Dharmika Dharmapala and Stephen L Ross. Racial bias in motor vehicle searches: Additional theory and evidence. *Contributions in Economic Analysis & Policy*, 3(1):1538–0645, September 2004.
- Michael Ewens, Bryan Tomlin, and Liang Choon Wang. Statistical discrimination or prejudice? A large sample field experiment. *The Review of Economics and Statistics*, 96(1):119–134, March 2014.
- Pieter A. Gautier and Yves Zenou. Car ownership and the labor market of ethnic minorities. *Journal of Urban Economics*, 67:392–403, 2010.
- Andrew Gelman, Jeffrey Fagan, and Alex Kiss. An analysis of the New York City Police Department’s “stop-and-frisk” policy in the context of claims of racial bias. *Journal of the American Statistical Association*, 102(479):813–823, 2007.
- Nicholas L. Georgakopoulos. Self-fulfilling impressions of criminality: unintentional race profiling. *International Review of Law and Economics*, 24(2):169–190, 2004.
- Jeffrey Grogger and Greg Ridgeway. Testing for racial profiling in traffic stops from behind a veil of darkness. *Journal of the American Statistical Association*, 101(475):878–887, 2006.
- David A. Harris. The stories, the statistics, and the law: Why “driving while black” matters. *Minnesota Law Review*, 84(265):265–326, 1999.
- Harvard Law Review. Considering police body cameras. *Harvard Law Review*, April 2015.
- Rubén Hernández-Murillo and John Knowles. Racial profiling or racist policing? Bounds ttest in aggregate data. *International Economic Review*, 45(3):959–989, 2004.
- Brady P. Horn, Jill J. McCluskey, and Ron C. Mittelhammer. Quantifying bias in driving-under-the-influence enforcement. *Economic Inquiry*, 52(1):269–284, January 2014.
- William C. Horrace and Shawn M. Rohlin. How dark is dark? Bright lights, big city, racial profiling. *The Review of Economics and Statistics*, Forthcoming, 2015.
- David Horsey. Racial prejudice and militarized cops are a bad mix in America. *Los Angeles Times*, August 2014.

- Dan Keating, Emily M. Badger, and Kennedy Elliott. Where police forces don't resemble the community. *The Washington Post*, August 14 2014.
- John Knowles, Nicola Persico, and Petra Todd. Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy*, 109(1):203–229, 2001.
- Tammy Rinehart Kochel, David B. Wilson, and Stephen D. Mastrofski. Effect of suspect race on officers' arrest decisions. *Criminology*, 49(2):473–512, 2011.
- Ashley Langer. Demographic preferences and price discrimination in new vehicle sales. Mimeo, 2012.
- David Montgomery. Texas trooper's behavior called 'catalyst' in Sandra Bland's death. *The New York Times*, July 2015.
- Alicia H. Munnell, Geoffrey M. B. Tootell, Lynn E. Browne, and James McEneaney. Mortgage lending in Boston: Interpreting HMDA data. *The American Economic Review*, 86(1):25–53, 1996.
- Christopher A. Parsons, Johan Sulaeman, Michael C. Yates, and Daniel S. Hamermesh. Strike three: Discrimination, incentives, and evaluation. *The American Economic Review*, 101(4):1410–1435., 2011.
- Nicola Persico and Petra E. Todd. Passenger profiling, imperfect screening, and airport security. *The American Economic Review*, 95(2):127–131, May 2005.
- Joseph Price and Justin Wolfers. Racial discrimination among NBA referees. *The Quarterly Journal of Economics*, 125(4):1859–1887, November 2010.
- Michael D. Reisig, John D. McCluskey, Stephen D. Mastrofski, and William Terrill. Suspect disrespect toward the police. *Justice Quarterly*, 21(2):241–268, June 2004.
- Janell Ross. How Black Lives Matter moved from a hashtag to a real political force. *The Washington Post*, August 2015.
- Sarith Sanga. Reconsidering racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy*, 117(6), 2009.

Figure 1: Estimated likelihood of citation by racial combination – from Table 3 Column (4)

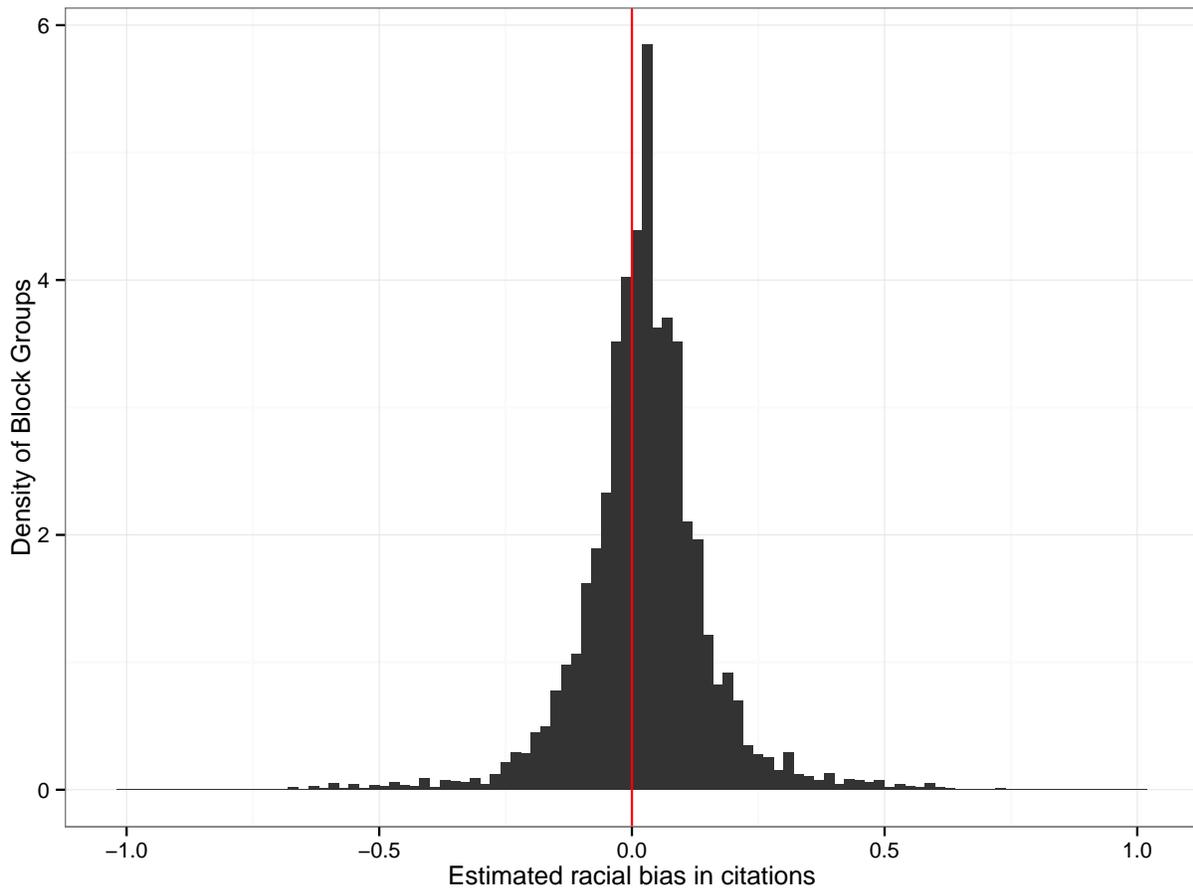


(a) Black and white officers and drivers



(b) Hispanic and white officers and drivers

Figure 2: Block Group variation in estimated racial bias (weighted by crash counts)



Notes: Figure 2 shows estimated coefficients from interacting the “black or Hispanic driver & other-race officer” indicator with each Block Group indicator in the specification of Column (4) of Table 3 Panel [C]. This yields a difference in differences estimate of racial bias in citations that is specific to each Block Group.

Table 1: Summary statistics

	Fraction
Drivers and crashes:	
White officer	0.686
Male officer	0.959
White driver	0.677
Male driver	0.650
Driver younger than 25	0.271
One vehicle in crash	0.329
Two vehicles in crash	0.473
Three or more vehicles	0.197
Citations:	
Driver cited	0.449
White driver cited	0.432
Black driver cited	0.434
Hispanic driver cited	0.504
Male driver cited	0.462
Female driver cited	0.429
Driver under 25 cited	0.573
Driver over 25 cited	0.403
Cited: nonmoving violations	0.110
Cited: moving violations	0.346
Cited: felony violations	0.016
Number of observations	439,605

Notes: Table includes all geocodable driver-crashes during 2006-2012 for a single State Police department for which the drivers and officers are white, black, or Hispanic. All variables as included in this table are specified as binary indicator variables.

Table 2: Tests of exogeneity of officer-driver racial combinations

	Dependent variable: race of investigating officer				
	(1)	(2)	(3)	(4)	(5)
[A] White officer					
Black driver	-0.0256** (0.0116)	0.0005 (0.0023)	-0.0001 (0.0023)	0.00002 (0.0023)	-0.0002 (0.0021)
Hispanic driver	-0.2477*** (0.0145)	-0.0043 (0.0035)	-0.0048 (0.0035)	-0.0046 (0.0035)	-0.0027 (0.0031)
[B] Black officer					
Black driver	0.0528*** (0.0104)	0.0005 (0.0018)	0.0008 (0.0017)	0.0007 (0.0017)	0.0010 (0.0016)
Hispanic driver	-0.0189*** (0.0062)	0.0029 (0.0018)	0.0029 (0.0018)	0.0029 (0.0018)	0.0016 (0.0015)
[C] Hispanic officer					
Black driver	-0.0272*** (0.0082)	-0.0010 (0.0018)	-0.0007 (0.0018)	-0.0007 (0.0018)	-0.0008 (0.0018)
Hispanic driver	0.2666*** (0.0152)	0.0015 (0.0036)	0.0019 (0.0035)	0.0017 (0.0036)	0.0011 (0.0031)
Spatial fixed effect	None	Block Gp.	Block Gp.	Block Gp.	BG-year
Crash-level controls	No	No	Yes	Yes	Yes
Driver-level controls	No	No	No	Yes	Yes
Observations	439,605	439,605	439,605	439,605	439,605
R ² for Panel [A]	0.0497	0.3542	0.3558	0.3559	0.4520

*p<0.1; **p<0.05; ***p<0.01 Standard errors are clustered by officer. Each panel-column consists of a separate regression using Equation 1. Regressions include only drivers and officers who are white, black, or Hispanic. The omitted category is a white driver. The coefficients for Black driver and Hispanic driver indicate how much more or less likely a black or Hispanic driver is than a white driver to be assigned an officer of the race indicated by the panel titles. Crash-level controls consist of fixed effects for the crash's calendar month, day of week, hour of day, light condition, road class, road surface condition, traffic control situation, weather condition, and whether the crash was in an intersection or construction zone. Driver-level controls are fixed effects for the driver's age and gender.

Table 3: Estimates of racial bias in all citations

	Dependent variable: driver cited (mean = 0.4488)					
	(1)	(2)	(3)	(4)	(5)	(6)
[A] Black/white						
Black driver & white officer	0.0175** (0.0084)	0.0171** (0.0082)	0.0208*** (0.0080)	0.0182** (0.0084)	0.0227*** (0.0086)	0.0520 (0.0359)
[B] Hispanic/white						
Hispanic driver & white officer	0.0296*** (0.0057)	0.0298*** (0.0056)	0.0301*** (0.0054)	0.0323*** (0.0055)	0.0331*** (0.0056)	0.0516** (0.0236)
[C] Pooled						
Black or Hisp. driver & other-race officer	0.0261*** (0.0045)	0.0268*** (0.0044)	0.0282*** (0.0043)	0.0294*** (0.0044)	0.0295*** (0.0045)	0.0448*** (0.0171)
Spatial fixed effect	Block Gp.	Block Gp.	Block Gp.	BG-year	BG-year	Crash
Driver race FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer race FE	Yes	Yes	Yes	Yes	—	—
Officer fixed effect	No	No	No	No	Yes	—
Crash-level controls	No	Yes	Yes	Yes	Yes	—
Driver-level controls	No	No	Yes	Yes	Yes	Yes
Total observations	439,605	439,605	439,605	439,605	439,605	439,605
R ² for Panel [A]	0.0332	0.0543	0.0829	0.1427	0.1527	0.6297
R ² for Panel [B]	0.0367	0.0587	0.0901	0.1411	0.1509	0.6185
R ² for Panel [C]	0.0332	0.0541	0.0855	0.1302	0.1396	0.5833

*p<0.1; **p<0.05; ***p<0.01 Standard errors are clustered by officer. Each cell presents a difference in differences estimate of racial bias, using separate regressions of Equation 2 in Panels [A] and [B] and Equation 3 in Panel [C]. Regressions in Panel [A] include only drivers and officers who are white or black. Regressions in Panel [B] include only drivers and officers who are white or Hispanic. Regressions in Panel [C] include only drivers and officers who are white, black, or Hispanic. Crash-level controls consist of fixed effects for the crash's calendar month, day of week, hour of day, light condition, road class, road surface condition, traffic control situation, weather condition, and whether the crash was in an intersection or construction zone. Driver-level controls are fixed effects for the driver's age and gender.

Table 4: Estimates of racial bias in types of citations

	Dependent variable: driver cited for type of violation					
	(1)	(2)	(3)	(4)	(5)	(6)
[A] Nonmoving violations (mean = 0.1095)						
Black or Hisp. driver & other-race officer	0.0164*** (0.0038)	0.0169*** (0.0038)	0.0165*** (0.0037)	0.0153*** (0.0036)	0.0164*** (0.0037)	0.0174* (0.0097)
[B] Moving violations (mean = 0.3460)						
Black or Hisp. driver & other-race officer	0.0163*** (0.0042)	0.0166*** (0.0040)	0.0196*** (0.0040)	0.0218*** (0.0041)	0.0223*** (0.0041)	0.0419** (0.0164)
[C] Felony violations (mean = 0.0164)						
Black or Hisp. driver & other-race officer	0.0009 (0.0010)	0.0011 (0.0010)	0.0006 (0.0010)	0.0002 (0.0010)	0.00002 (0.0010)	0.0004 (0.0025)
Spatial fixed effect	Block Gp.	Block Gp.	Block Gp.	BG-year	BG-year	Crash
Driver race FE	Yes	Yes	Yes	Yes	Yes	Yes
Officer race FE	Yes	Yes	Yes	Yes	—	—
Officer fixed effect	No	No	No	No	Yes	—
Crash-level controls	No	Yes	Yes	Yes	Yes	—
Driver-level controls	No	No	Yes	Yes	Yes	Yes
Observations	439,605	439,605	439,605	439,605	439,605	439,605
R ² for Panel [A]	0.0570	0.0659	0.0872	0.1473	0.1579	0.7489
R ² for Panel [B]	0.0279	0.0560	0.0831	0.1275	0.1370	0.5592
R ² for Panel [C]	0.0216	0.0386	0.0414	0.1127	0.1201	0.8143

*p<0.1; **p<0.05; ***p<0.01 Standard errors are clustered by officer. Each cell presents a difference in differences estimate of racial bias using separate regressions of Equation 3. Crash-level controls consist of fixed effects for the crash’s calendar month, day of week, hour of day, light condition, road class, road surface condition, traffic control situation, weather condition, and whether the crash was in an intersection or construction zone. Driver-level controls are fixed effects for the driver’s age and gender. Nonmoving violations consist of expired/nonexistent driver license and expired vehicle registration, inspection, or insurance. Moving violations consist of charges related to passing, right-of-way, signal intention, speeding, traffic signs, seatbelt, equipment defects, and other miscellaneous moving violations. Felony violations consist of vehicular assault, manslaughter, and hit-and-run, which are all potentially felony offenses.

Table 5: Heterogeneity in estimates of racial bias in types of citations

Data subset	Driver cited for type of violation			
	All	Nonmoving	Moving	Felony
	(1)	(2)	(3)	(4)
All observations (mean = 0.4488)	0.0294*** (0.0044)	0.0153*** (0.0036)	0.0218*** (0.0041)	0.0002 (0.0010)
Driver:				
Male (mean = 0.4619)	0.0297*** (0.0055)	0.0179*** (0.0042)	0.0219*** (0.0051)	0.0004 (0.0013)
Female (mean = 0.4285)	0.0266*** (0.0078)	0.0088 (0.0055)	0.0205*** (0.0074)	0.0020 (0.0013)
Younger than 25 (mean = 0.5731)	0.0425*** (0.0088)	0.0339*** (0.0080)	0.0283*** (0.0091)	0.0002 (0.0025)
Older than 25 (mean = 0.4027)	0.0261*** (0.0050)	0.0095*** (0.0036)	0.0203*** (0.0045)	0.0002 (0.0011)
Crash:				
9:00 p.m. - 6:00 a.m. (mean = 0.4640)	0.0216** (0.0094)	0.0124* (0.0066)	0.0166** (0.0084)	-0.0017 (0.0031)
6:00 a.m. - 9:00 p.m. (mean = 0.4442)	0.0311*** (0.0050)	0.0156*** (0.0041)	0.0209*** (0.0047)	0.0004 (0.0009)
Last week in month (mean = 0.4478)	0.0357*** (0.0082)	0.0206*** (0.0058)	0.0330*** (0.0078)	0.0001 (0.0019)
Spatial fixed effect	BG-year	BG-year	BG-year	BG-year
Crash-level controls	Yes	Yes	Yes	Yes
Driver-level controls	Yes	Yes	Yes	Yes
Observations	439,605	439,605	439,605	439,605

*p<0.1; **p<0.05; ***p<0.01 Standard errors are clustered by officer. Each cell results from a separate regression. Coefficients are for “black or Hispanic driver with other-race officer,” using the data subsets as indicated in rows. All regressions include fixed effects for driver race and officer race, so these specifications correspond to those in Column (4) of Tables 3 - 4. Regressions include only drivers and officers who are white, black, or Hispanic. Reported means are for all citations. For other definitions see Table 4 notes.

A Data appendix

This appendix describes the details of the data cleaning and matching process used to form the analysis sample. The objective is to merge the investigator names from the crash data to employee personnel records which include each officer’s race. In forming these matches, I heavily prioritize avoiding false positives, seeking to minimize measurement error at the cost of yielding a lower, but more accurate, match-rate. As such, I require that there not be simultaneous data entry errors in both the investigator’s badge number and name.

The personnel data as I obtained them are already very “clean.” In particular, these records have separate fields for first and last name without missing data, and any middle initials are in a separate field. The only data cleaning I do to the personnel data is to strip them of the occasional name suffix such as “JR” or “III,” because such designations are inconsistently recorded in the crash data. Additionally, I restrict the available set of personnel to only those with a unique first-last name combination, which excludes a few records that have a very common first name and surname.

Investigator names in the crash data, by contrast, require extensive cleaning. The raw crash data include only a single “Investigator Name” field, which has some combination of officers’ first name, first initial, middle name, middle initial, last name (or multiple, in the case of compound last names), suffix, and title (e.g. “SGT:”). By far the most common record format is a first initial followed by the last name. In addition, there is a field for the officer’s badge number. There are extensive data entry errors (typos) in both the name and the badge number fields. The badge number field I standardize by typically just stripping it of spaces, special characters, and inconsistently reported non-numerical content (e.g. “BADGE12345” or “BDG12345” rather than “12345”). For the name field I remove special characters, perform similar “regular expression” cleaning (e.g. removing surplus spaces), and remove any identifiable suffix, as with the personnel data.

Next, I parse the name field into string components based on spaces. Any (cleaned) name fields which parse into more than two components – less than one percent of records – are manually merged to the personnel data. Those with one or two components are string merged to the personnel data, allowing for several permutations of the components (e.g. field 1 is first name and field 2 is last name, versus field 1 is last name and field 2 is first name). Again, to minimize measurement error I require a strict match, meaning a correct and full first and last name. Using the names which match, I then “roll” these matches through other observations of the same badge number and same “fuzzy” (dynamic Damerau-Levenshtein string distance) last name. The utility of this rolling of matches is extensive, because there

are many cases in which the crash data consist only of a first initial with a last name, but there might be one or a few records for that badge number in which the full first and last name are spelled out. The resulting crosswalk matches 80.76 percent of crash records. Spot checking random subsets of the unmatched crash records indicates that the vast majority are due to significant typos being made for both the badge number and investigator’s name.

Observation counts at the various steps in the sample formation process are as follows. During 2006-2012, this State Police Department investigated almost 400,000 unique crashes involving motor vehicles for which the crash location is geocodable and matches to a 2010 Census Block Group. For these crash records, I was able to match 321,753 to a unique officer in the personnel records, a match rate of 80.76 percent. Of these, 313,912 crashes have a person record for at least one involved driver without missing information on the driver’s race. In total, these usable crashes include 454,184 unique usable driver records, as there are often multiple drivers involved in the same crash. I exclude 14,579 records for which either the driver or the officer was not white, black, or Hispanic/Latino, yielding a final analysis sample of 439,605 records of driver-officer encounters.

Table A.1 underscores that – for the crashes that have available geocoordinates and a driver who is white, black, or Hispanic – the analysis sample is not a meaningfully nonrandom subset due to the imperfect match rate to the personnel data or the exclusion of officers who are classified other than as one of these three races.

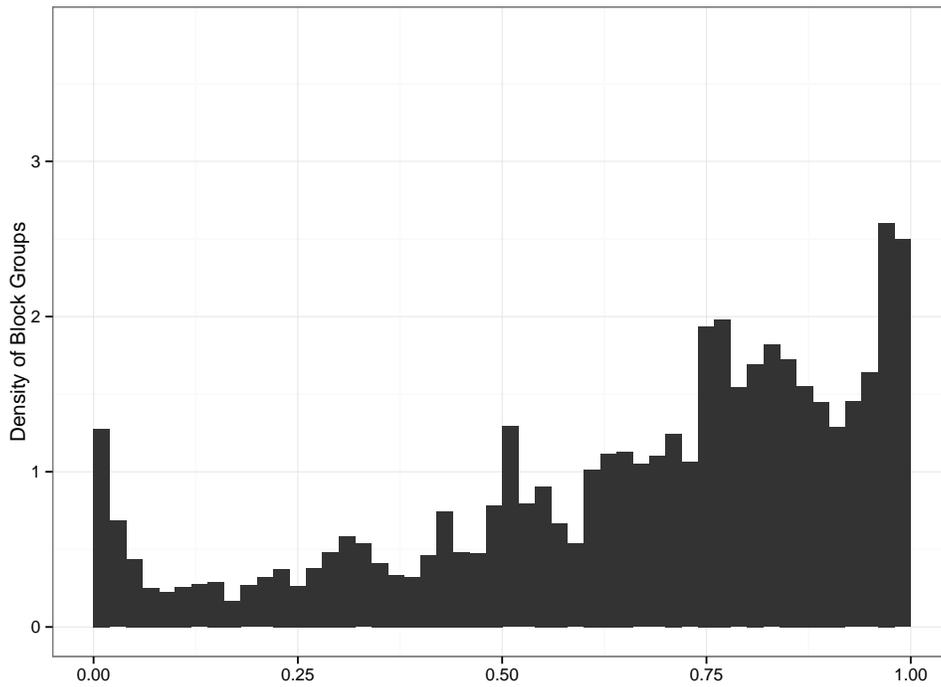
Table A.1: Test of importance of missing data

	Observation is included in sample (mean = 0.7873)
Black driver	0.0023 (0.0019)
Hispanic driver	0.0015 (0.0015)
Observations: total / in sample	558,322 / 439,605
R ²	0.1282

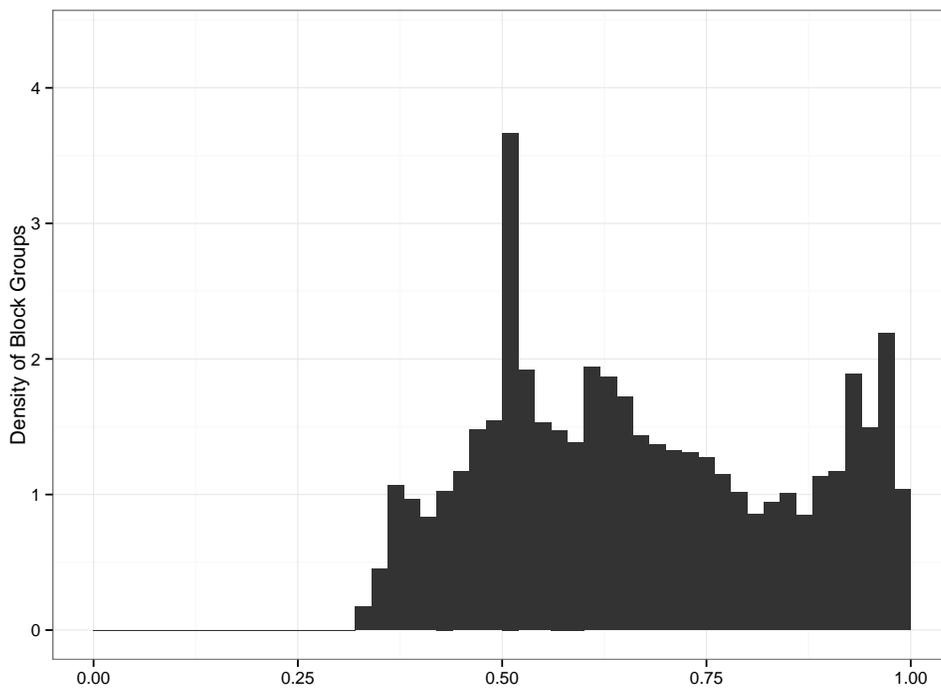
*p<0.1; **p<0.05; ***p<0.01 Standard errors are not clustered. Regression includes all observations with geocoordinates and a driver who is white, black, or Hispanic. The omitted category is a white driver. Regression includes fixed effects for Block Groups and includes no other controls.

B Additional figures and tables

Figure B.1: Block group variation in officer diversity (weighted by crash counts)



(a) Fraction of crashes within Block Group with white officer



(b) Officer race Herfindahl index for Block Group

Table B.1: Estimates of racial bias in specific citations

	Dependent variable: Driver cited for type of violation					
	(1)	(2)	(3)	(4)	(5)	(6)
[A] Nonmoving violations (mean = 0.1095)						
Insurance/inspection (mean = 0.0582)	0.0025 (0.0026)	0.0026 (0.0025)	0.0022 (0.0025)	0.0018 (0.0025)	0.0018 (0.0024)	0.0058 (0.0068)
License/registration (mean = 0.0720)	0.0179*** (0.0034)	0.0183*** (0.0034)	0.0181*** (0.0033)	0.0176*** (0.0033)	0.0188*** (0.0034)	0.0136 (0.0085)
[B] Moving violations (mean = 0.3460)						
Miscellaneous (mean = 0.0511)	0.0050*** (0.0018)	0.0050*** (0.0018)	0.0050*** (0.0018)	0.0053*** (0.0018)	0.0051*** (0.0019)	0.0115* (0.0063)
Passing (mean = 0.0090)	0.0008 (0.0008)	0.0009 (0.0008)	0.0009 (0.0008)	0.0008 (0.0008)	0.0008 (0.0008)	-0.0024 (0.0031)
Right of way (mean = 0.0241)	0.0011 (0.0012)	0.0015 (0.0012)	0.0019* (0.0012)	0.0016 (0.0012)	0.0023* (0.0013)	0.0047 (0.0053)
Seatbelt (mean = 0.0083)	-0.0006 (0.0007)	-0.0005 (0.0007)	-0.0004 (0.0007)	-0.0005 (0.0007)	-0.0004 (0.0007)	-0.0007 (0.0016)
Signaling (mean = 0.0003)	0.00004 (0.0001)	0.00004 (0.0001)	0.00004 (0.0001)	0.0001 (0.0001)	0.0001 (0.0002)	0.0005 (0.0006)
Speeding (mean = 0.1911)	0.0080** (0.0034)	0.0073** (0.0033)	0.0092*** (0.0033)	0.0108*** (0.0033)	0.0093*** (0.0033)	0.0125 (0.0098)
Traffic signs (mean = 0.0345)	0.0002 (0.0014)	0.0006 (0.0013)	0.0011 (0.0014)	0.0010 (0.0014)	0.0012 (0.0015)	0.0040 (0.0064)
Turning (mean = 0.0253)	0.0013 (0.0013)	0.0013 (0.0013)	0.0015 (0.0013)	0.0024* (0.0013)	0.0031** (0.0014)	0.0094* (0.0055)
Vehicle defects (mean = 0.0029)	0.0007* (0.0004)	0.0007* (0.0004)	0.0006 (0.0004)	0.0008* (0.0005)	0.0010** (0.0005)	0.0013 (0.0015)
Wrong way (mean = 0.0069)	0.0008 (0.0006)	0.0008 (0.0006)	0.0007 (0.0006)	0.0005 (0.0006)	0.0007 (0.0006)	0.0034 (0.0024)
[C] Felony violations (mean = 0.0164)						
Assault/manslaughter (mean = 0.0025)	-0.0004 (0.0004)	-0.0004 (0.0004)	-0.0004 (0.0004)	-0.0004 (0.0004)	-0.0004 (0.0004)	-0.0006 (0.0013)
Reporting (Hit-Run) (mean = 0.0140)	0.0012 (0.0009)	0.0014 (0.0009)	0.0009 (0.0009)	0.0004 (0.0009)	0.0003 (0.0010)	0.0010 (0.0022)
Spatial fixed effect	Block Gp.	Block Gp.	Block Gp.	BG-year	BG-year	Crash
Officer fixed effect	No	No	No	No	Yes	—
Crash-level controls	No	Yes	Yes	Yes	Yes	—
Driver-level controls	No	No	Yes	Yes	Yes	Yes
Observations	439,605	439,605	439,605	439,605	439,605	439,605

*p<0.1; **p<0.05; ***p<0.01 Standard errors are clustered by officer. Coefficients are for “black or Hispanic driver with other-race officer.” All columns and table notes correspond to Table 4.