

Racial Bias in Police Investigations

Jeremy West*

UC Santa Cruz

February 2018

Abstract

Endogenous selection into police encounters typically complicates evaluations of law enforcement discrimination. This study overcomes selection concerns by examining automobile crash investigations, for which officer dispatch is demonstrably exogenous to drivers' race. I find State Police officers issue significantly more traffic citations to drivers whose race differs from their own. This bias is evident for both moving and nonmoving violations, the latter indicating a preference for discriminatory leniency towards same-race individuals. I show this treatment is unmitigated by socioeconomic factors: officers cite other-race drivers more frequently regardless of their age, gender, vehicle value, or characteristics of the local community.

JEL: J15, K42

Keywords: Law enforcement, racial discrimination, automobile crashes

*West: Economics Department, University of California at Santa Cruz, and The E2e Project (email: westj@ucsc.edu). I am grateful for valuable comments from Amanda Agan, Joshua Angrist, David Autor, Peter Blair, Howard Bodenhorn, George Bulman, Jamein Cunningham, Gregory DeAngelo, Carlos Dobkin, Jennifer Doleac, Jonathan Gruber, Benjamin Hansen, Rick Harbaugh, Mark Hoekstra, Sally Hudson, Christopher Knittel, Steven Levitt, Jason Lindo, Justin Marion, Jonathan Meer, James Poterba, Steven Puller, Daniel Sacks, Allison Shertzer, Glen Waddell, Randall Walsh, Casey Wichman, and numerous seminar participants. Any errors, 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. I have no material financial or legal interests related to the research described in this paper.

1 Introduction

There is strong public sentiment that law enforcement is characterized by racial bias, with the majority of Americans believing that the country’s criminal justice system “favors whites over blacks” (CNN/ORC Poll, 2014).¹ This sentiment is embodied in activist movements such as the “black lives matter” campaign and in phrases such as “driving while black” (Harris, 1999). Descriptively, the scholarly literature supports this sentiment. For example, Pierson et al. (2017), using 60 million traffic stops across 20 states, find that black drivers are stopped more often than white drivers relative to their share of the driving age population.

Empirical studies are much less conclusive in demonstrating that racial bias is a causal factor in generating racial disparities in law enforcement. The fundamental identification challenge is that analysts only observe data on encounters that police officers actually initiate – if an officer decides not to stop a driver or pedestrian, then this (non)encounter is never recorded. Empirical tests for racial bias attempt to overcome these selection concerns by making statistical assumptions about the distribution of unobserved drivers, either overall (e.g. Knowles et al., 2001; Sanga, 2009) or in a relative sense across officers of different race (e.g. Anwar and Fang, 2006; Antonovics and Knight, 2009; Goncalves and Mello, 2017). These modeling assumptions are ultimately untestable, and the endogeneity of police encounters could substantially affect researchers’ ability to accurately quantify racial bias.

This concern is widely acknowledged in the literature. Anwar and Fang (2006) note that “the trooper 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,” 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 (2016) use the “veil of darkness” surrounding dusk as a clever identification strategy that accounts explicitly for the endogeneity in traffic stops; however, even the veil of darkness test appears to suffer from selection bias (Kalinowski et al., 2017).

In contrast to previous studies, I assess racial bias in automobile crash investigations, a setting in which police encounters are *demonstrably* exogenous with respect to drivers’

¹Throughout this manuscript, I use the term *race* in the cultural or identity sense, not as a biological or genealogical descriptor. I use the term *racial bias* to refer to disparities in treatment that are causally attributable to an individual’s race and discretionary on the part of the decision-maker. In my empirical study, race is a categorical variable defined already for individuals included in the data sets I analyze, and I consistently, if imprecisely, refer to race rather than ethnicity.

race. The support for my empirical identification is straightforward: unlike in a traffic stop, police officers are dispatched to crash investigations. The specific driver(s) to be investigated are determined by the occurrence of the automobile crash, in advance and independently of any officer involvement. I validate the exogeneity of these officer-driver interactions using data on 440,000 crash investigations conducted by a single State Police Department (SPD) during 2006-2012. When conditioning on the geographic area 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 indeed uncorrelated with that of the driver(s) involved in the crash.

This exogenous 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 that accounts for racial variation 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](#); [Parsons et al., 2011](#); [Anbarci and Lee, 2014](#)). Although 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 same-race drivers, on an average citation rate of 45 percent. This finding is robust to including detailed controls specific to the crash or driver, to using within-officer variation, and even to using within-crash variation identified from drivers of differing race crashing with each other. In exploring heterogeneity in this racial bias, I find significant evidence of discrimination for both moving and nonmoving violations, but not for felony violations. Moreover, the extent of racial bias appears to be invariant to drivers’ other sociodemographic and vehicle characteristics, and present systematically across communities in which the SPD operates.

I draw several types of inference from this set of findings. First, these results help to illuminate *why* police officers exhibit racial bias, supporting a mechanism of [Becker \(1957\)](#) preference-based discrimination rather than [Arrow \(1973\)](#) statistical discrimination. For one, the official objective in a crash investigation is to determine the factors relevant to the collision, not to maximize a “hit rate” for finding contraband, minimizing the utility of statistical discrimination. More directly, the evidence of racial bias in citations for nonmoving

violations cannot feasibly be consistent with only statistical discrimination: these offenses, such as expired vehicle registration, are extremely salient, objective, and the associated documentation is required to be recorded as a routine part of all crash investigations. Although my findings cannot fully rule out statistical discrimination, they do provide ample support for attributing at least some of the officers' racial bias to a preference-based explanation.

Second, the results illustrate *when* police officers practice racial discrimination. One facet is that the estimated racial bias is essentially invariant to driver characteristics. Officers cite other-race drivers more frequently regardless of their age, gender, vehicle value, or characteristics of the local community surrounding the crash. Strikingly, the estimated racial bias is virtually identical towards drivers of expensive new vehicles as towards those driving inexpensive older vehicles, in contrast to theories that higher wealth insulates individuals from racial discrimination (Cole, 1999). A less direct facet is that officers' racial bias is more prevalent for lower-stakes offenses in which there is minimal explicit and implicit monitoring of their actions. Parsons et al. (2011) show that monitoring plays a key role in determining when agents practice racial bias, and my findings are also consistent with Fryer's (2016) evidence on when police officers decide to use physical force.

Finally, the results highlight one aspect of *how* police officers discriminate: by acting comparatively more lenient in their interactions with same-race individuals. The evidence for nonmoving violations again provides the most compelling support. It is extremely unlikely that an officer would issue a citation for a nonmoving violation to a driver who presents valid paperwork; thus, officers can discriminate in these citations only via leniency towards guilty individuals. These findings are particularly policy-relevant as, with 86 percent of police officers stating that recent events "have made their jobs harder," the incentive becomes stronger for a biased officer to discriminate using lenience rather than harshness (Simon et al., 2017). More broadly, these results exemplify how discriminatory leniency serves as a channel through which economic agents can practice preference-based discrimination.

To summarize, I document significant and systematic racial bias by police officers that is causal, that is more consistent with preference-based than statistical discrimination, that is more prevalent in actions involving less monitoring, and that operates in part via discriminatory leniency towards same-race individuals. The extent to which the behavior identified here applies in other police settings is an open question, but the pattern of evidence supports a more general argument regarding racial discrimination in law enforcement. Clearly demonstrating the existence of this type of behavior in a large police agency serves as a step towards policies that encourage more transparent, effective, and equitable policing.

2 Data

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

2.1 State Police crash investigations and personnel records

The automobile crash investigation data initially consist of the population of police-reported vehicle crashes handled by SPD. Consistent with the role of 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 are more concentrated on Interstates and highways.² The crash data include detailed information about the crash (location, time, 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 main 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 SPD employee's full name, job title, hire date, race, and gender. For several supporting analyses I additionally include data on Census Block Group characteristics from the 2008-2012 American Community Survey 5-Year Estimates, data on local presidential election voting from the Harvard Election Data Archive, and data on motor vehicle attributes purchased from DataOne Software and merged to records using the 11-digit prefix of the Vehicle Identification Number.³

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

²The 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.

³Respective URLs for these data: <http://www.census.gov/geo/maps-data/data/tiger-data.html>, <http://dataverse.harvard.edu/dataverse/eda>, <http://www.dataonesoftware.com/vehicle-data-vin-decoding>.

record, and I only incorporate matches for which the personnel data include a unique record of that full name, which excludes several names that combine a popular first name with 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 sample loss 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 19 percent are due to significant data entry errors.⁴ My identification strategy also requires dropping records with missing data about 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/Latina is coded as its own race, and throughout I refer imprecisely to race rather than ethnicity). In the crashes handled by SPD, this is a very minor restriction.

2.3 Summary statistics of crashes and citations

My final sample for analysis consists of roughly 440,000 encounters between a driver and a SPD police officer. Summary statistics are in Table 1. The first panel presents statistics of the drivers and crashes: Column (1) includes all officers in the sample, while Columns (2) and (3) disaggregate to officers who are respectively white or non-white. 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, providing little statistical power for 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 State Police handle primarily highway collisions. Roughly half of driver-crashes involve two drivers; the remaining 20 percent involve three or more vehicles.⁵

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 officer should vary with the race of the driver involved in the crash.

⁵Of total crashes, rather than driver-crashes: 57.6 percent involve only a single vehicle, 32.6 percent involve two vehicles, and the remaining 9.8 percent involve three or more vehicles.

The second panel shows variation in citations by driver and vehicle characteristics. On average, drivers are cited about 45 percent of the time. Many of the uncited drivers are involved in multivehicle crashes with other drivers who are cited, but officers issue no citations for more than 38 percent of *crashes* in the data. Additionally, multiple drivers are simultaneously cited in about 8 percent of multivehicle crashes. Together, these statistics underscore that officers *de facto* regularly have discretion over whether or not they cite a driver.

Disaggregated by driver characteristics, perhaps unsurprisingly male drivers and younger drivers are cited comparatively more often than female and older drivers. Likewise, officers cite drivers of older and lower-priced vehicles more often than drivers of newer and higher-priced vehicles. By race, 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. This pattern continues to hold for statistics disaggregated by officer race in Columns (2) and (3). Thus, a naive comparison of the raw data for black and white drivers would yield a conclusion of no racial bias. However, as shown in the top panel, white officers investigate crashes by white drivers nearly 74 percent of the time, while non-white officers handle crashes by white drivers only 55 percent of the time. Moreover, black and white drivers might not actually commit infractions at equal rates, and officers may differ in their leniency by driver race. These factors motivate the need for a more convincing 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 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 geographic area, the race of the dispatched investigating officer is uncorrelated with that of any drivers involved in a crash. Second, I use a generalized difference in differences

estimator to identify the effect on citations of being exogenously assigned an other-race officer relative to an officer of the driver’s own race.

3.1 Demonstrate exogenous officer assignment to crashes

The causal inference of this study is facilitated by the process of dispatching officers to crash investigations. Unlike in a traffic stop, police officers in most jurisdictions are *assigned* to investigate particular automobile crashes, with dispatch typically based on their proximity to the scene and on other factors unrelated to driver race. Thus, as an institutional practice officer dispatch should yield exogenous encounters of different officer-driver racial combinations.⁶ Importantly, because my test for racial bias employs a difference in differences framework, the identification requires only that driver race is uncorrelated with investigating officer race – not that the officer is dispatched exogenously with respect to any other factors of the crash. More formally, the (testable) identification assumption is that $P(\text{officer race is } r) = P(\text{officer race is } r \mid \text{driver race is } r)$ for any given race r .

Testing the exogeneity of these encounters in data is straightforward, but the identification will hold 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.⁷ Ignoring the location of the crash, a black driver is unconditionally more likely than a white driver to be dispatched a black officer, and a Hispanic driver is unconditionally more likely than a white driver to see a Hispanic crash 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 exogenous assignment of officer-driver racial combinations. Note that the use of Census Block Groups as the geospatial unit is an arbitrary one, as all that the identification test requires is a sufficiently small geographic area in which driver race is uncorrelated with officer race. I explicitly test the exogeneity of this assignment using linear probability regressions such as:

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

⁶As the data include all police-reported crashes, the question of exogeneity regards *which* SPD officer the driver encounters, in contrast to settings such as traffic stops for which the first-order exogeneity question is *whether* or not a police encounter occurs.

⁷As [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 local police officers.

In Equation (1), $I\{\text{white officer}\}_{aij}$ is an indicator variable for whether officer j responding 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 .⁸

Results from these tests of exogenous officer-driver combinations are shown 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 above, 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 unequivocal 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(s) 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). This serves as very compelling evidence of exogenous assignment of officer-driver interactions and provides a clean framework for causal inference in the analyses of racial bias that follow.⁹

⁸An 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 SPD crash records name the officer who is the official decision-making agent, and any additional officers at the scene are present only to direct traffic and handle other supporting activities.

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 identifying variation is provided by Appendix Figure B.1.

3.2 Difference in differences to identify racial bias in citations

This demonstrably exogenous 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 average citation rates for drivers of a given race across different officer races. For example, one could compare how often black drivers are cited by white officers to how often black drivers are cited by black officers. But, this approach ignores that officers of differing races may vary in their propensities to cite drivers regardless of driver race. In this example, perhaps white officers are comparatively more likely to issue citations to *all* drivers, not just to black 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 black drivers. This second approach circumvents the concern 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 estimator 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 (shown here for black and white drivers and officers):

$$\begin{aligned} \text{I}\{\text{citation}\}_{aij} &= \beta_B \cdot \text{I}\{\text{black driver}\}_i + \gamma_B \cdot \text{I}\{\text{black officer}\}_j \\ &+ \delta \cdot \text{I}\{\text{black driver \& white officer}\}_{ij} + \mu_a + \text{controls}_{ij}\Psi + \epsilon_{aij} \end{aligned} \quad (2)$$

The β_B coefficient captures differences in citation rates of black drivers by black officers relative to white drivers by black officers, allowing for drivers to differ in their overall propensities to deserve citations. The γ_B coefficient captures differences in citation rates of white drivers by black officers relative to white drivers by white officers, allowing for officers to differ in their overall propensities to issue citations. δ is the difference in differences coefficient of interest and captures the degree of racial bias. If there is no racial bias, then $\delta = 0$. A positive sign for δ indicates same-race leniency (or other-race harshness) in citations. Importantly, the difference in differences estimator provides only a symmetric and *relative* measure of

racial bias, a point of discussion I expand on later in this section.¹⁰

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 study, 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 outcome 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 including all three racial categories. The coefficient of interest δ continues to quantify the degree of racial bias, just as it does in Equation (2), and in practice this pooled δ coefficient is akin to a weighted average of the two δ estimates from separately estimating Equation (2) for black/white and Hispanic/white.¹¹ I cluster standard errors for all estimations at the officer level, as this seems *ex ante* to be the appropriate cluster for most autocorrelation concerns. In practice this is also a more conservative approach, yielding 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 assumptions. 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 same-race drivers, adjusted for racial propensities of drivers to receive citations and officers to issue citations, but this estimate cannot distinguish whether

¹⁰That is, an exactly identical estimate of racial bias as in Equation (2) is obtained from estimating: $\text{I}\{\text{citation}\}_{aij} = \beta_W \cdot \text{I}\{\text{white driver}\}_i + \gamma_W \cdot \text{I}\{\text{white officer}\}_j + \delta \cdot \text{I}\{\text{white driver \& black officer}\}_{ij} + \mu_a + \text{controls}_{ij}\Psi + \epsilon_{aij}$. In the tables I denote the label for this coefficient of interest as “other-race officer.”

¹¹Strictly speaking, the pooled δ is not *exactly* a weighted average of the δ from black/white and δ from Hispanic/white, because black drivers are sometimes assigned Hispanic officers and vice versa. However, as white officers handle the large majority of investigations and interactions between black/Hispanic or Hispanic/black are relatively sparse, the pooled δ is effectively a weighted average in practice. The causal inference of any treatment effect from being dispatched an *other*-race officer holds regardless.

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 credible causal evidence regarding the existence of racial bias in law enforcement.

4 Results

As discussed and shown in Section 3, the institutional setting provides for exogenous assignment of racial encounters between police officers and drivers: that is, $P(\text{officer race is } r) = P(\text{officer race is } r \mid \text{driver race is } r)$ for any given race r . My identification strategy is to leverage this exogenous dispatch of police officers to crash investigations and use a difference in differences estimator across these officer-driver interactions to estimate the degree of racial bias in citations. Specifically, the “other-race officer” difference in differences coefficients reported in the results tables show how much more or less likely a driver is to be cited by an officer of a different race relative to one of the same race as the driver, adjusted for overall racial propensities of drivers to receive citations and officers to issue citations.

4.1 Officers issue more citations to other-race drivers

The first hypothesis I test is whether being exogenously assigned an other-race officer increases (or decreases) the likelihood that a driver is issued *any* citation in the crash investigation. 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 per the difference in differences identification, and all columns include at least Block Group spatial fixed effects to maintain the exogenous assignment demonstrated earlier in Table 2. Figure 1 uses these results from Table 3 to graph the likelihoods 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

(and vice versa). 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.¹² In addition, the estimates are economically significant. Given a mean of 44.88 percent, the pooled result indicates that racial bias causes a driver assigned an other-race officer to be nearly six percent more likely to be given any citation in a crash investigation.

The additional columns of Table 3 address potential confounders to this racial bias estimate. For example, 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 from these officers 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). Indeed, 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 same-race versus other-race drivers across different encounters. 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 $p < 0.05$.

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 multivehicle 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 race crash with each other, the race of the exogenously assigned officer plays a decisive 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 persuasive to see such similar results from an identification strategy that uses a rather different source of variation and captures

¹²As discussed in Section 3, I cluster standard errors by officer, a comparatively conservative approach.

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 exogenously assigned to interact with civilians. The remainder of this manuscript examines heterogeneity in these effects to better characterize the nature of this racial bias and shed light on the likely mechanisms.

4.2 Estimated racial bias indicates discriminatory leniency

My study so far has focused on documenting the overall degree of racial bias in automobile crash investigations. An instinctive next step is to evaluate the nature of this bias. One important question is whether the disparity in treatment operates primarily via officers exhibiting additional “harshness” when dealing with other-race drivers – such as more aggressively issuing them citations based on subjective evaluation – or via officers acting relatively more lenient in their interactions with same-race individuals.

The literature provides robust support for a hypothesis that the racial bias in crash investigations occurs primarily through a channel of same-race leniency. For one, studies of bias in other contexts often find that “endophilia dominates exophobia” (Feld, Salamanca, and Hamermesh, 2016). In addition, there is substantial evidence that police officers in general exhibit leniency in enforcing traffic laws, partly because officers “wish to avoid the unpleasantness of dealing with citizens who express their anger and hostility at being on the receiving end of the officer’s punitive choices” (Schafer and Mastrofski, 2005). More directly, Anbarci and Lee’s (2014) “findings suggest that officers are racially biased in speed discounting” and issue lower-valued speeding citations to same-race motorists, a finding echoed by Goncalves and Mello (2017).¹³

One rationale for practicing discrimination via leniency rather than harshness is that it is more difficult to detect and punish discriminatory leniency. Previous literature has shown that monitoring, whether explicit or implicit, dramatically decreases the incidence of evaluators’ discriminatory bias (e.g. Parsons et al., 2011; Borcan, Lindahl, and Mitrut, 2017). In the context of my study there is little explicit monitoring, as “most traffic enforcement occurs at a low level of visibility in the organization,” but implicit monitoring is a significant factor, as “every act of punishment [can] increase the risk that the offenders will either

¹³Evidence is more mixed elsewhere in the justice system. Anwar, Bayer, and Hjalmarsson (2012) document same-race leniency in court decisions by juries, but Depew, Eren, and Mocan (2016) find judges sentence juvenile defendants of their *own* race more severely, perhaps due to “concern about creating the impression of being prejudiced towards the defendants who are of the opposite race.”

complain to the police department about something the officer did or demand their right to a hearing before a magistrate” (Schafer and Mastrofski, 2005).

The crash investigation data include no objective measure of drivers’ guilt for offenses, so it is infeasible to directly assess officer leniency. However, citations for nonmoving violations offer a convincing indirect method of testing the role of leniency in biasing officers’ decisions. Many violations in a crash investigation require subjective evaluation that a driver might challenge in court, such as speeding violations (the officer cannot measure vehicle speed with a radar device in a crash investigation). In contrast, nonmoving violations such as expired vehicle registration are perfectly objective to evaluate and extremely salient.¹⁴ Indeed, the officer is required to record driver and vehicle information as a routine part of every crash investigation. A guilty driver who is cited for an objectively-verifiable infraction has little ground to challenge the citation in court, even if other guilty drivers are not charged for the same offense. Moreover, whereas an investigating officer might subjectively issue a moving violation citation to a driver whose actual speed was uncertain, an officer is extremely unlikely to issue a nonmoving violation citation to a driver who presents valid paperwork. Thus, there is little explicit or implicit monitoring of (non)citations for nonmoving violations, and officers may discriminate in these citations *only by leniency*.

Table 4 presents results, with specifications using Equation 3 and directly corresponding to those in Panel [C] of Table 3, separately for nonmoving violations, moving violations, and felony violations.¹⁵ There is significant racial bias 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 around 14-15 percent of the mean; that for moving violations is around 5-6 percent of the mean. As with the results for racial bias in overall citations, estimates by citation type remain very stable and statistically significant across the specifications.

I find very small and statistically insignificant effects for felony violations, but it is unsurprising to see estimated racial bias only for less severe violations. Donohue and Levitt (2001) also find that racial bias by police is most pronounced for minor offenses, and the

¹⁴For instance, one can determine the validity of vehicle registration by looking at the exterior of a vehicle, and all of these documents have a plainly indicated expiration date.

¹⁵As discussed in Section 2, 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 miscellaneous other 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 shown in Appendix Table B.1.

findings in Table 4 are consistent with those of Fryer (2016), who finds “blacks and Hispanics are more likely to experience some form of force in interactions with police,” but “finds no racial differences [in] the most extreme use of force.”¹⁶

As discussed above, racial bias in citations for highly salient and objective nonmoving violations is convincing evidence of same-race leniency. More so, it provides compelling evidence of *preference-based* discrimination via same-race leniency. For moving violations, it is possible that police officers are statistically discriminating in their citations for (unobserved) speeding or other offenses, especially given that “police often choose who to investigate based on previous offenses or convictions” (Hansen, 2015). For nonmoving violations this cannot be a factor as the current details of the driver’s information must be recorded anyway as part of the investigation. 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 incorrect information or false beliefs (Georgakopoulos, 2004; Burke, 2007).¹⁷ The simplicity and salience of guilt for nonmoving violations leaves virtually no scope for incomplete information or false belief, further underscoring a mechanism of officers’ preference for same-race leniency in generating this racial disparity.

4.3 The racial bias appears to be systematic

The remainder of the empirical results examine heterogeneity in the estimated racial bias by characteristics of the drivers, vehicles, and crash locations. Because the difference in differences estimator can only identify relative racial bias across officers, I cannot examine heterogeneity at the officer level or by fixed officer characteristics such as officer race.¹⁸ Broadly, these analyses show that socioeconomic factors do not appear to play a role in

¹⁶There 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 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.

¹⁷To be clear, the *theory* of statistical discrimination is based on an assumption that decision-makers use accurate statistical information in making their decisions, but in practice agents may have poor quality or inaccurate information, or may be extrapolating statistical relationships to make out-of-sample decisions. As an example, Gelman et al. (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 local propensities of different racial groups to commit crimes, this incorrect belief could reconcile their behavior as being (misguided) statistical discrimination rather than prejudice.

¹⁸I find no economically or statistically significant relationship between officers’ time-varying accumulated experience and estimated magnitude of racial bias.

either augmenting or mitigating officers’ racial bias in this context. In addition, they bolster the evidence for a mechanism of preference-based leniency towards same-race drivers. The literature has noted that “those who are disrespectful to the police are more likely to have their behavior reciprocated” (Reisig et al., 2004), and that “officers expect citizens to be contrite and acknowledge responsibility for their infractions” (Schafer and Mastrofski, 2005). Results in this section make it challenging to maintain that officers are only reciprocating driver behavior in deciding not to afford other-race drivers the same leniency as they afford drivers of the officers’ own race.¹⁹

4.3.1 By driver and vehicle characteristics

Table 5 includes estimates for racial bias in different types of citations using subsets of the data by driver or vehicle characteristics. On balance, the evidence is that there is very little heterogeneity across subgroups: police officers issue significantly more citations to other-race drivers regardless of their age, gender, vehicle age, or vehicle value. Column (1), which includes all types of citations, shows point estimates ranging from 2.51 to 4.1 percentage points, a fairly tight range from 5.5-7.5 percent relative to the respective subgroup means. In addition, every estimate in Column (1) is statistically significant at $p < 0.01$ despite the substantially smaller subgroup samples, which drop to including as few as 20 percent of the total observations (due to data limitations, vehicle MSRP is only available for observations from 2010 and later). Results by citation type in Columns (2) - (4) are similarly largely invariant to driver or vehicle characteristics and broadly consistent with the full-sample results in Table 4, showing racial bias in moving and nonmoving, but not felony, violations.²⁰

The results by vehicle age and value are particularly interesting. The estimate for any racial bias towards drivers of vehicles with below-median original Manufacturer Suggested Retail Price (\$23,350 in-sample) is 3.04 percentage points – virtually identical to the estimate of 3.01pp for drivers with above-median-MSRP vehicles. Estimates for drivers of vehicles below/above median in-sample age (of 6 years) are likewise very close at respectively 2.67 and 2.79 percentage points. I further explore heterogeneity by these vehicle characteristics

¹⁹Price and Wolfers (2010) consider (and rule out) that the racial bias they document instead could be attributed to National Basketball Association 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. The consideration here is that the racial combination might affect how respectful a driver acts towards the police officer, and that driver disrespect in turn could influence whether the officer writes a citation for an observed violation.

²⁰One apparent exception is the smaller estimate and statistical imprecision for nonmoving violations by drivers of above-median MSRP vehicles, but this is largely because drivers of higher-value vehicles tend to be current in their vehicle registrations and other paperwork, so there is little scope for (differential) leniency.

in Figures 2 and 3. For these figures, I estimate officer racial bias separately for subgroups of much more granular composition of vehicle age or original MSRP, still using Equation 3 throughout. The figures plot these local average treatment effects (LATE) relative to the average citation rate for same-race encounters by vehicle age or MSRP bin. The baselines show a nearly monotonic increase (decrease) in average citation rates as vehicle age (value) increases, consistent with drivers of older and less expensive vehicles having higher citation rates in general. Despite these baseline trends, the LATE are strikingly uniform in magnitude across vehicle age and value, and every single coefficient in these figures indicates an increased likelihood of being cited by an other-race officer.

Substantial commentary in the mainstream press and research literature suggests that, while not immune to racial bias in the justice system, individuals of higher economic means are often buffered from its effects (Cole, 1999). Generally speaking, there is almost certainly some truth to the sentiment held by many higher-income minorities that “their socioeconomic status and education insulates them from the problems of poorer and uneducated blacks” (Simpson, 1998). The findings shown here suggest, however, at least in a context in which officers are exogenously assigned to interact with civilians, racial bias is neither augmented nor mitigated by drivers’ economic means. Additionally, that the racial bias seems invariant to drivers’ demographic and (proxies for) economic characteristics undermines support for the more plausible driver-based explanations for the racial disparity documented in this study.

4.3.2 By geographic area and community characteristics

This final results section evaluates heterogeneity in racial bias by the location of the crash. To do so, I interact the “other-race officer” indicator with each Block Group indicator to estimate Block Group specific local average treatment effect coefficients for racial bias. The results of this exercise are shown in Figure 4. The mean of the distribution is around three percentage points, as reflected in the statewide average treatment effect estimates shown earlier in Panel [C] of Table 3. The interesting aspect of this distribution is that it is essentially a normal distribution centered to the right of zero, suggesting very systematic racial bias geographically within the state.

There is visible weight in the tails, but this is due to the idiosyncratic draws of which (and how many) crashes occur in each Block Group – statistical noise, in other words. In particular, a number of Block Groups have relatively few minority drivers or officers. A formal evaluation of these more extreme LATE is shown in Figure 4(b), which plots the

distribution of p-values from tests of the hypotheses that each Block Group specific LATE is equal to the average treatment effect statewide. If the more extreme LATE values are due primarily to sampling variation, then these p-values should be distributed uniformly, which as shown is largely the case. Although there is some additional mass at low-valued p-scores, as Figure 4(c) shows this is not attributable to systematically positive or negative large degrees of estimated racial bias.²¹

This exercise yields additional compelling evidence that the racial bias I estimate is a systematic effect, not attributable to a just a handful of officers or types of situations. Given the distribution in Figure 4, it is thus unsurprising to see in Figure 5 that the Block Group specific estimates of racial bias are uncorrelated with local community characteristics. Figure 5(a) and 5(b) plot the estimated LATE against the in-sample fraction of white drivers and officers for each Block Group, an exercise motivated in part by heterogeneity documented in [Donohue and Levitt \(2001\)](#) and [Shayo and Zussman \(2011\)](#). Figure 5(c) presents a similar exercise based on the racial representativeness of the set of local officers, defined here as one minus the absolute value difference of the fraction of white drivers and white officers. [Keating et al. \(2014\)](#) and others voice concerns regarding the “common discrepancy between the racial makeup of a police force and the community it serves.” In my study context, there is no apparent relationship between the local degree of racial bias and the “racial imbalance” of local officers. In Figure 5(d) I explore how the LATE relate to the Obama ’08 and Kerry ’04 voting shares, a proxy for community-specific racial bias as shown in [Stephens-Davidowitz \(2014\)](#). As with subfigures (a)-(c), there is a flat correlation. Finally, Figure 5(e)-(f) plot the estimated LATE against Block Group statistics on income and education from the American Community Survey, again showing no pattern. To summarize: all six of the linear fit lines in Figure 5 remain fully above zero, but none have a slope that is visibly nor statistically different from zero, further illustrating that the racial bias is systematic in this context.

5 Conclusions

A question of significant social consequence is the extent to which law enforcement is characterized by discrimination. Although a large and active literature seeks to identify and

²¹I thank Dan Sacks for this suggestion. My inspection of the frequencies of encounters for these more “extreme” Block Groups shows that these larger coefficients occur almost exclusively in areas that have a fairly large number of total encounters but very few cross-racial encounters. For instance, one LATE coefficient of 0.51 with a standard error of 0.04 is for a Block Group with more than 200 total observations, of which only 5 are for a minority driver with an other-race officer – a statistically significant, but not economically meaningful, difference from the statewide ATE.

quantify the degree of racial bias by police officers, these studies share in common a substantial limitation in that encounters between civilians and police are endogenous.

My study overcomes selection concerns by examining police officer behavior in automobile crash investigations. Because officers are dispatched to investigate crashes, with this dispatch based on factors unrelated to drivers' race, these interactions are demonstrably exogenous. 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(s). This allows for a causal interpretation of how the assigned officer-driver racial combination affects police officers' discretionary behavior. Using a difference in differences estimator, I leverage this exogenous assignment to quantify racial bias in citations for traffic violations, adjusted for overall racial propensities of drivers to receive citations and officers to issue citations.

I find that police officers issue significantly more citations to other-race drivers for both moving and nonmoving violations. In addition, I find that the racial bias appears to be systematic throughout communities in which the police agency operates, and that it is qualitatively invariant to driver demographics or vehicle characteristics. The bias in citations for nonmoving violations – essentially, expired paperwork infractions that are salient and indisputable – is particularly compelling evidence of an officer preference for discriminatory leniency in situations with little explicit or implicit monitoring. More broadly, the findings serve as decisive evidence of systematic racial bias by police officers that is apparently unmitigated by civilians' sociodemographic attributes or economic means.

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 increased monitoring of officer behavior will help to curb racial disparities in law enforcement ([Harvard Law Review, 2015](#)). The findings of my study support that increased monitoring may serve as an effective step towards more transparent, effective, and equitable policing.

References


- N. Anbarci and J. Lee. Detecting racial bias in speed discounting: Evidence from speeding tickets in Boston. *International Review of Law and Economics*, 38:11–24, June 2014.
- K. Antonovics and B. 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.

- S. Anwar and H. Fang. An alternative test of racial prejudice in motor vehicle searches: Theory and evidence. *The American Economic Review*, 96(1):127–151, March 2006.
- S. Anwar, P. Bayer, and R. Hjalmarsson. The impact of jury race in criminal trials. *The Quarterly Journal of Economics*, 127(2):1017–1055, May 2012.
- K. J. Arrow. The theory of discrimination. In O. Ashenfelter and A. Rees, editors, *Discrimination in Labor Markets*, pages 3–33. Princeton University Press, 1973.
- G. S. Becker. *The Economics of Discrimination*. University of Chicago Press, Chicago, 1957.
- O. Borcan, M. Lindahl, and A. Mitrut. Fighting corruption in education: What works and who benefits? *American Economic Journal: Economic Policy*, 9(1):180–209, Feb. 2017.
- A. Burke. Neutralizing cognitive bias: An invitation to prosecutors. *NYU Journal of Law and Liberty*, 2:512–530, 2007.
- CNN/ORC Poll. CNN and Opinion Research Corporation poll, December 18-21 2014.
- D. Cole. *No equal justice: Race and class in the American criminal justice system*. The New Press, New York, NY, 1999.
- B. Depew, O. Eren, and N. Mocan. Judges, juveniles and in-group bias. NBER Working Paper No. w22003, February 2016.
- D. Dharmapala and S. L. Ross. Racial bias in motor vehicle searches: Additional theory and evidence. *Contributions in Economic Analysis & Policy*, 3(1):1538–0645, September 2004.
- J. J. Donohue and S. D. Levitt. The impact of race on policing and arrests. *The Journal of Law and Economics*, 44(2):367–394, October 2001.
- J. Feld, N. Salamanca, and D. S. Hamermesh. Endophilia or exophobia: Beyond discrimination. *The Economic Journal*, 126(594):1503–1527, August 2016.
- R. G. Fryer. An empirical analysis of racial differences in police use of force. NBER Working Paper No. w22399, July 2016.
- A. Gelman, J. Fagan, and A. 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, September 2007.
- N. L. Georgakopoulos. Self-fulfilling impressions of criminality: Unintentional race profiling. *International Review of Law and Economics*, 24(2):169–190, June 2004.
- F. Goncalves and S. Mello. A few bad apples? Racial bias in policing. Princeton University working paper 608, March 2017.


- J. Grogger and G. Ridgeway. Testing for racial profiling in traffic stops from behind a veil of darkness. *Journal of the American Statistical Association*, 101(475):878–887, Sept. 2006.
- B. Hansen. Punishment and deterrence: Evidence from drunk driving. *The American Economic Review*, 105(4):1581–1617, April 2015.
- D. A. Harris. *Driving while Black: Racial profiling on our nation’s highways*. American Civil Liberties Union, Washington, D.C., 1999.
- Harvard Law Review. Considering police body cameras. *Harvard Law Review*, 128(6):1794–1817, April 2015.
- W. C. Horrace and S. M. Rohlin. How dark is dark? Bright lights, big city, racial profiling. *The Review of Economics and Statistics*, 98(2):226–232, May 2016.
- J. Kalinowski, S. L. Ross, and M. B. Ross. Endogenous driving behavior in veil of darkness tests for racial profiling. University of Connecticut working paper 2017-03, February 2017.
- D. Keating, E. M. Badger, and K. Elliott. Where police forces don’t resemble the community. *The Washington Post*, August 14 2014.
- J. Knowles, N. Persico, and P. Todd. Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy*, 109(1):203–229, February 2001.
- C. A. Parsons, J. Sulaeman, M. C. Yates, and D. S. Hamermesh. Strike three: Discrimination, incentives, and evaluation. *The American Economic Review*, 101(4):1410–1435., June 2011.
- E. Pierson, C. Simoiu, J. Overgoor, S. Corbett-Davies, V. Ramachandran, C. Phillips, and S. Goel. A large-scale analysis of racial disparities in police stops across the United States. Stanford University working paper, August 2017.
- J. Price and J. Wolfers. Racial discrimination among NBA referees. *The Quarterly Journal of Economics*, 125(4):1859–1887, November 2010.
- M. D. Reisig, J. D. McCluskey, S. D. Mastrofski, and W. Terrill. Suspect disrespect toward the police. *Justice Quarterly*, 21(2):241–268, June 2004.
- S. Sanga. Reconsidering racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy*, 117(6):1155–1159, December 2009.
- J. A. Schafer and S. D. Mastrofski. Police leniency in traffic enforcement encounters: Exploratory findings from observations and interviews. *Journal of Criminal Justice*, 33(3): 225–238, May 2005.
- M. Shayo and A. Zussman. Judicial ingroup bias in the shadow of terrorism. *The Quarterly Journal of Economics*, 126(3):1447–1484, August 2011.

- D. Simon, S. O'Key, and M. Schuman. Poll finds stark disconnect between police and the public. *CNN*, January 11 2017.
- A. Y. Simpson. *The tie that binds: Identity and political attitudes in the post-civil rights generation*. New York University Press, New York, NY, 1998.
- S. Stephens-Davidowitz. The cost of racial animus on a black candidate: Evidence using Google search data. *Journal of Public Economics*, 118:26–40, October 2014.

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: Estimated likelihood of citation by racial combination and vehicle age




Notes: Figure 2 shows the predicted likelihood that a driver is cited in a crash investigation depending on the age of the driver’s vehicle and the racial combination of the driver and the dispatched investigating officer. The solid line depicts the average citation rate by vehicle age for drivers whose race is the same as that of the officer. The circular markers add to this baseline the difference in differences coefficients from separately estimating by vehicle age the “other-race officer” indicator in the specification of Column (1) of Table 3 Panel [C], with the size of the markers representing the distributional mass at that vehicle age. This effectively yields a measure of racial bias in citations that is specific to vehicles of each particular age. For clarity and statistical power reasons, I topcode age at 16 years, corresponding to roughly the 95th percentile of vehicle age in the data. The dashed lines show the upper and lower 95 percent confidence interval around the estimated local average treatment effects.

Figure 3: Estimated likelihood of citation by racial combination and vehicle price




Notes: Figure 3 shows the predicted likelihood that a driver is cited in a crash investigation depending on the original Manufacturer Suggested Retail Price of the driver’s vehicle and the racial combination of the driver and the dispatched investigating officer. The solid line depicts the average citation rate by vehicle MSRP for drivers whose race is the same as that of the officer. The circular markers add to this baseline the difference in differences coefficients from separately estimating by vehicle MSRP bins (of \$2500) the “other-race officer” indicator in the specification of Column (1) of Table 3 Panel [C], with the size of the markers representing the distributional mass at that vehicle MSRP. This effectively yields a measure of racial bias in citations that is specific to vehicles of each particular original MSRP. For clarity and statistical power reasons, I bottomcode MSRP at \$15000 and topcode at \$40000, corresponding respectively to roughly the 10th and 95th percentiles of vehicle MSRP in the data. The dashed lines show the upper and lower 95 percent confidence interval around the estimated local average treatment effects.


Figure 4: Local average treatment effect estimates of racial bias by Census Block Group



(a) Distribution of estimated coefficients (weighted by crash counts)



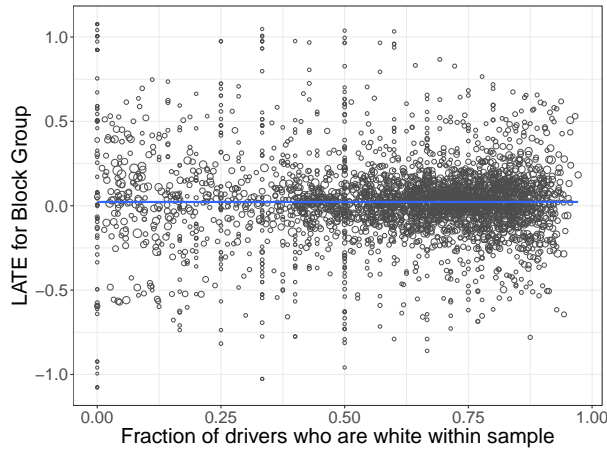
(b) Tests for H_0 : LATE = ATE




(c) LATE for BG with this $p < 0.05$

Notes: Figure 4(a) shows estimated coefficients from interacting the “other-race officer” indicator with each Block Group indicator in the specification of Column (1) of Table 3 Panel [C]. This effectively yields a difference in differences estimate of racial bias in citations that is specific to each Block Group. Figure 4(b) presents the distribution of p-scores from tests of hypotheses that each Block Group LATE is equal to the statewide ATE. Figure 4(c) zooms in on the Block Groups with this p-score < 0.05 , showing these LATE.

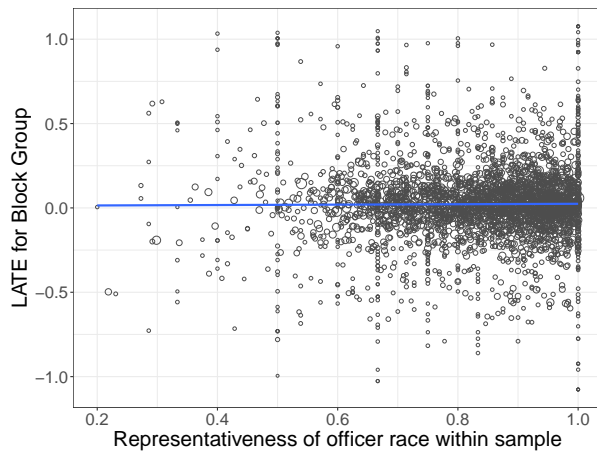
Figure 5: Block Group LATE estimates of racial bias by Block Group characteristics




(a) Driver racial composition




(b) Officer racial composition




(c) Officer racial representativeness



(d) Voting-based measure of racial bias



(e) Median household income (ACS 2008-2012)



(f) Adults with Bachelor's (ACS 2008-2012)

Notes: Figure 5 plots the Block Group LATE estimates of racial bias – i.e. the values underlying the distribution shown in Figure 4(a) – against various local community characteristics. The fit lines show the first-order linear least-squares fit, weighted by total crash counts per Block Group. None of the fit lines have a slope that is statistically different from zero at any conventional significance level.

Table 1: Summary statistics

	Fraction of driver-crashes		
	All officers (1)	White officers (2)	Non-white off. (3)
Drivers and crashes:			
Number of observations	439,605	301,723	137,882
White officer	0.686	1	0
Male officer	0.959	0.962	0.952
White driver	0.677	0.736	0.548
Male driver	0.65	0.649	0.651
Driver younger than 25	0.271	0.269	0.274
One vehicle in crash	0.329	0.345	0.295
Two vehicles in crash	0.473	0.455	0.514
Three or more vehicles	0.197	0.2	0.191
Multiple driver races in crash	0.222	0.217	0.233
Citations:			
Driver cited	0.449	0.444	0.459
White driver cited	0.432	0.432	0.432
Black driver cited	0.434	0.435	0.429
Hispanic driver cited	0.504	0.502	0.506
Male driver cited	0.462	0.457	0.472
Female driver cited	0.429	0.424	0.44
Driver under 25 cited	0.573	0.565	0.591
Driver over 25 cited	0.403	0.4	0.41
Vehicle age below median cited	0.397	0.394	0.403
Vehicle age above median cited	0.504	0.497	0.518
Vehicle MSRP below median cited	0.469	0.464	0.48
Vehicle MSRP above median cited	0.399	0.394	0.409
Cited: nonmoving violations	0.11	0.101	0.128
Cited: moving violations	0.346	0.345	0.348
Cited: felony violations	0.0165	0.017	0.0152

Notes: Table 1 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 values reported in this table are unconditional average rates of the variables indicated in rows. The median vehicle age in-sample is 6 years and median original vehicle MSRP is \$23,350.

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
R ² for Panel [B]	0.0041	0.2805	0.2820	0.2822	0.3942
R ² for Panel [C]	0.0754	0.3978	0.3989	0.3991	0.4887

*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						
Other-race 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)
Observations	283,844	283,844	283,844	283,844	283,844	283,844
R ²	0.0332	0.0543	0.0829	0.1427	0.1527	0.6297
[B] Hispanic & white						
Other-race 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)
Observations	365,092	365,092	365,092	365,092	365,092	365,092
R ²	0.0367	0.0587	0.0901	0.1411	0.1509	0.6185
[C] Pooled 3 races						
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)
Observations	439,605	439,605	439,605	439,605	439,605	439,605
R ²	0.0332	0.0541	0.0855	0.1302	0.1396	0.5833
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

*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)						
Other-race officer	0.0164*** (0.0038)	0.0168*** (0.0038)	0.0165*** (0.0037)	0.0153*** (0.0036)	0.0163*** (0.0037)	0.0173* (0.0097)
<i>Estimate/mean</i>	0.150	0.153	0.151	0.140	0.149	0.158
[B] Moving violations (mean = 0.3460)						
Other-race officer	0.0164*** (0.0042)	0.0166*** (0.0040)	0.0196*** (0.0039)	0.0219*** (0.0041)	0.0223*** (0.0041)	0.0420** (0.0164)
<i>Estimate/mean</i>	0.047	0.048	0.057	0.063	0.064	0.121
[C] Felony violations (mean = 0.0164)						
Other-race officer	0.0009 (0.0010)	0.0011 (0.0010)	0.0006 (0.0010)	0.0002 (0.0010)	0.00003 (0.0010)	0.0005 (0.0025)
<i>Estimate/mean</i>	0.055	0.067	0.037	0.012	0.002	0.030
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	0.5593
R ² for Panel [C]	0.0216	0.0385	0.0413	0.1127	0.1200	0.8145

*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: Estimates of racial bias in types of citations by driver and vehicle characteristics

Data subset	Dependent variable: driver cited for type of violation				Obs.
	All (1)	Nonmoving (2)	Moving (3)	Felony (4)	
All	0.0282*** (0.0043)	0.0165*** (0.0037)	0.0196*** (0.0039)	0.0006 (0.0010)	439,605
<i>Estimate/mean</i>	0.063	0.151	0.057	0.037	
Male driver	0.0290*** (0.0052)	0.0190*** (0.0041)	0.0202*** (0.0048)	0.0007 (0.0013)	284,354
<i>Estimate/mean</i>	0.063	0.162	0.058	0.034	
Female driver	0.0256*** (0.0070)	0.0106** (0.0052)	0.0172*** (0.0067)	0.0008 (0.0012)	153,481
<i>Estimate/mean</i>	0.060	0.111	0.049	0.087	
Driver age under 25	0.0410*** (0.0080)	0.0357*** (0.0074)	0.0252*** (0.0078)	0.0014 (0.0023)	118,914
<i>Estimate/mean</i>	0.072	0.228	0.056	0.064	
Driver age over 25	0.0251*** (0.0047)	0.0109*** (0.0035)	0.0184*** (0.0043)	0.0004 (0.0011)	320,691
<i>Estimate/mean</i>	0.062	0.118	0.060	0.028	
Vehicle age < median	0.0267*** (0.0056)	0.0109*** (0.0036)	0.0181*** (0.0053)	0.0014 (0.0012)	218,892
<i>Estimate/mean</i>	0.067	0.167	0.056	0.114	
Vehicle age > median	0.0279*** (0.0058)	0.0174*** (0.0051)	0.0204*** (0.0054)	-0.0005 (0.0015)	218,620
<i>Estimate/mean</i>	0.055	0.113	0.055	-0.024	
Vehicle MSRP < median	0.0304*** (0.0092)	0.0184*** (0.0068)	0.0097 (0.0087)	0.0015 (0.0022)	88,530
<i>Estimate/mean</i>	0.065	0.164	0.029	0.099	
Vehicle MSRP > median	0.0301*** (0.0094)	0.0040 (0.0063)	0.0171* (0.0088)	-0.0014 (0.0021)	88,586
<i>Estimate/mean</i>	0.075	0.055	0.058	-0.105	
Spatial fixed effect	Block Group	Block Group	Block Group	Block Group	
Crash-level controls	Yes	Yes	Yes	Yes	
Driver-level controls	Yes	Yes	Yes	Yes	

*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 Equation 3. Reported coefficients are for “other-race officer” using data subsets indicated in rows. Fixed effects for driver race and officer race are included, so these specifications correspond to those in Column (3) of Tables 3 - 4. 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)	0.0051*** (0.0018)	0.0051*** (0.0018)	0.0051*** (0.0018)	0.0054*** (0.0018)	0.0052*** (0.0019)	0.0117* (0.0062)
Passing (mean = 0.0090)	0.0008 (0.0008)	0.0009 (0.0008)	0.0009 (0.0008)	0.0008 (0.0008)	0.0008 (0.0008)	-0.0023 (0.0031)
Right of way (mean = 0.0241)	0.0011 (0.0012)	0.0014 (0.0012)	0.0019 (0.0012)	0.0015 (0.0012)	0.0022* (0.0013)	0.0046 (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)	0.0001 (0.0001)	0.0001 (0.0001)	0.0001 (0.0001)	0.0001 (0.0002)	0.0002 (0.0002)	0.0005 (0.0006)
Speeding (mean = 0.1911)	0.0079** (0.0034)	0.0073** (0.0033)	0.0092*** (0.0033)	0.0107*** (0.0033)	0.0092*** (0.0033)	0.0125 (0.0098)
Traffic signs (mean = 0.0345)	0.0002 (0.0015)	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)	0.0014 (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)	0.0009* (0.0005)	0.0013 (0.0015)
Wrong way (mean = 0.0069)	0.0008 (0.0006)	0.0008 (0.0006)	0.0008 (0.0006)	0.0005 (0.0006)	0.0008 (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)	-0.0005 (0.0013)
Reporting (Hit-Run) (mean = 0.0140)	0.0012 (0.0009)	0.0013 (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. All reported coefficients are for “other-race officer.” All columns and table notes correspond to those in Table 4.