WORKING PAPER #608
PRINCETON UNIVERSITY
INDUSTRIAL RELATIONS SECTION
MARCH 2017
http://arks.princeton.edu/ark:/88435/dsp01z890rw746

A Few Bad Apples? Racial Bias in Policing*

Felipe Goncalves Steven Mello †

March 6, 2017

Abstract

We provide new evidence on the presence and distribution of racial bias in the criminal justice system. In many states, the punishment for speeding increases discontinuously with the speed of the driver, exhibiting large jumps in fine amounts. It is a common practice for officers to reduce the charged speed to just below this jump, avoiding an onerous punishment for the driver. Using data from the Florida Highway Patrol, we find evidence of significant bunching in ticketed speeds below a jump in punishment for all drivers but significantly more for whites than for blacks and Hispanics. We estimate the bias of each officer by comparing his lenience towards whites and non-whites, allowing us to recover the full distribution of bias. The total disparity in lenience across races can be explained by a small percentage (~20%) of officers. Officers tend to favor drivers of their own racial group, and younger, female, and college-educated officers are less biased. We then estimate a model that allows for both heterogeneity in officer preferences and driver speeds across races. Because minorities tend to live in areas where officers are harsher to all drivers, policies targeting bias have little effect on the aggregate speed gap. We find that racial bias in lenience explains 16% of the minority-white speed gap, and spatial differences in race-blind lenience explain 30% of the gap.

JEL Classification: J71, K42
Keywords: Discrimination, Racial Bias, Police, Traffic Enforcement.

*We are grateful to Ilyana Kuziemko and Alex Mas for guidance and support throughout this project. We benefitted from helpful comments by Peter Bergman, Jessica Brown, Janet Currie, Rebecca Diamond, Nik Engbom, Kirill Evdokimov, Jeremy Fox, Nathaniel Hendren, Daniel Herbst, Bo Honoré, Neale Mahoney, Michael Makowsky, Christopher Neilson, Jakob Schlokermann, and participants of the Princeton Labor Lunch, Public Finance Working Group, Center for Health and Wellbeing Working Group, and the Hamilton-Colgate Economics Seminar. We thank Beth Allman, Jeffrey Bissainthe, Kiara Guzzo, Wilton Johnson, Stacy Lehmann, Brenda Paige, and Aundrea Wilder for assistance with data from various agencies. The Industrial Relations Section, Princeton University, provided generous financial support. Any errors are our own.

†Industrial Relations Section, Princeton University, Princeton, NJ 08544. Goncalves: fmg@princeton.edu. Mello: smello@princeton.edu

1 Introduction

There are large racial disparities in criminal justice outcomes in the United States. Blacks and Hispanics are significantly more likely to be stopped, arrested, and imprisoned than whites. A central question in research on the criminal justice system, therefore, is whether the disparate outcomes of minorities are due to discrimination on the part of law enforcement agents. This view has gained traction in recent years following several highly publicized police killings of minority victims (West, 2015). For example, a 2013 Gallup poll revealed that half of black adults agreed that racial differences in incarceration rates are "mostly due to discrimination."[1] The existence of racial bias in the criminal justice system is, however, difficult to establish empirically. Disparate outcomes may reflect underlying differences in criminality across racial groups. Demonstrating conclusively that individuals of different races receive different treatment requires controlling for all relevant contextual factors across incidents, which is generally not possible. Further, much of the interaction between citizens and police goes unrecorded in data, and one of the central candidate dimensions of bias is whom police choose to stop.

In this paper, we examine whether police officers discriminate when setting punishments for speeding. Traffic stops are the most common form of civilian-police interaction, with about 41 million speeding tickets given and over $6B in fines paid annually.[2] Although officers typically observe a driver's speed via radar before stopping them, they are free to choose what speed to *charge*. In many states, the punishment for speeding increases discontinuously with the speed of the driver, exhibiting "jumps" in harshness. A jump may involve not only a higher fine, but also a mandated court appearance or permanent mark on the driver's record. As shown in Figure 1, the distribution of speeds ticketed by the Florida Highway Patrol between 20005 and 2015 shows substantial excess mass at speeds just below the first fine increase. Meanwhile, a remarkably small portion of tickets are issued for speeds just above. We take this bunching as compelling evidence that officers systematically manipulate the charged speed, commonly charging speeds just below fine increases after observing a higher speed, perhaps to avoid an onerous punishment for the driver.[3] While the practice of lenience appears to be widespread, a substantial portion of officers exhibit no such bunching, ruling out the possibility that the aggregate bunching reflects drivers choosing their speeds just below the fine increase.

We then test whether officers are systematically less lenient towards minority drivers when

[1] See `http://www.gallup.com/poll/175088/gallup-review-black-white-attitudes-toward-police.aspx`.

[2] National Highway Traffic Safety Administration (2014). For reference, there were about 13.1 million arrests made in the U.S. in 2010, according to the FBI.

[3] Past researchers have termed this phenomenon *speed discounting* (Anbarci and Lee, 2014).

choosing to reduce the charged speed. Figure 3 presents the speed histograms, broken down by race, and clearly shows a smaller mass of minorities charged at the bunch speed. These racial disparities remain after controlling for an array of stop and driver-level characteristics, including speed limit and stop location, age, gender, vehicle type, ZIP code income level, and prior tickets, which we treat as evidence that, on average, officers behave less favorably towards minority drivers.

Next, we leverage the unique features of our setting to estimate the racial bias of individual officers. Specifically, we compute each officer's lenience towards minorities relative to his own treatment of white drivers, adjusting for other features of the stop, and treat the difference as the officer's bias. While other papers have examined the presence of racial bias in policing, few have been able to say anything about the distribution of bias across agents.[4] Using a deconvolution technique to account for estimation error, we convert the officer-level bias measures into a non-parametric estimate of the full distribution of racial bias. This exercise reveals that the majority of officers exhibit no bias, with the aggregate disparity in treatment explained by the behavior of a small minority of officers composing about 20% of the patrol force. We also explore how bias varies with officer-level characteristics, documenting that officers exhibit own-race preferences and that younger, female, and college-educated officers are less likely to be biased.

One concern with our strategy to identify officer bias is the possibility that racial groups vary systematically in their driving speeds, with minorities less likely to be stopped at speeds eligible for lenience. We address this issue by explicitly modeling the officer's charging decision. Given a stop with an observed speed above the fine increase point, we suppose the officer charges either the true speed or the bunching speed as a function of both the true speed and his taste for the driver. The model makes clear predictions about the ticketed speed distribution and we rely on these predictions to estimate the model parameters via maximum likelihood. Importantly, the model allows for the identification of officers' racial preferences after accounting for unobserved differences in driving speeds across racial groups. Although we impose functional forms on officers' taste shocks and the underlying speed distributions to simplify the estimation, we again allow for a fully nonparametric distribution of racial bias.

The model estimates imply that when stopped by the average officer, minority drivers receive a fine reduction at the same rate as white drivers traveling four MPH faster. Surprisingly, forcing all officers to treat minority drivers as they treat white drivers removes only 33% of the gap in the probability of receiving a fine reduction, and 16% of the gap in charged speed. The majority

[4] Anwar and Fang (2006) and Antonovics and Knight (2009) see the race of the officer involved in a traffic vehicle search, and look at whether the behavior of officers is homogenous across their race. With their tests, it is possible to know that at least one group is biased but not which one.

of the disparity is due to the fact that the most lenient officers patrol in counties with the fewest minorities – 47% of the white-nonwhite speed gap disappears without bias or sorting of officers across counties; the remainder is due to differences in driven speed. By construction, all of the disparity in discount treatment vanishes in the absence of racial bias and officer sorting.

Finally, we use the model estimates to consider the effectiveness of various policies aimed at mediating the racial gap in treatment while holding the speed differences fixed. For the reasons discussed above, policies targeting bias directly are only mildly effective. Firing the most biased officers (both for and against minorities) reduces the gap, as does increasing the presence of minority or female officers, but the gains are limited. Imposing either minimum or maximum levels of *lenience* can substantially reduce bias with small requirements. Perhaps most effective and easily implemented, reassigning officers across counties within their troops so that minorities are exposed to more lenient officers can remove essentially the entire white-minority lenience gap.

While discrimination in speeding enforcement may seem innocuous, several features of our institutional setting are ideal for studying bias. In many settings, a central concern is accounting for unobserved differences in criminality across individuals. In the context of speeding tickets, *guilt* is summarized by the driving speed, which is both one-dimensional and typically observed by the ticketing officer. Further, in many criminal justice contexts, lenience is relative and possibly non-monotonic, while in our setting we can exploit an absolute lenience benchmark and assume a monotonic structure on lenience based on the observed distribution of charges. Perhaps most importantly, we observe agents making many decisions in very similar contexts, which allows us to construct a measure of bias for each officer by comparing his treatment of nonwhites and whites. Existing studies are unable to achieve individual-level measures of bias either because of data constraints or because agents' behavior must be benchmarked against agents of a different race (i.e. *external benchmarking*), as in Price and Wolfers (2010). Such studies rely on the assumption that a well-defined set of agents behave in an unbiased way and cannot distinguish whether all officers exhibit a mean bias level or whether aggregate disparities are generated by the behavior of a subset of especially biased agents, as we find.

This paper falls in a broad category of recent research using "bunching" estimators to recover behavioral parameters. Predominantly used in the literature on taxation, these studies attempt to estimate the hypothetical distribution of interest in the absence of bunching by looking at the distribution outside of a region around the manipulated area and inferring out-of-sample how the distribution should look at the discontinuity.[5] They then estimate bunching to be the difference between the true and hypothetical distribution around the bunch point. We improve on the

[5] See Kleven (2016) for a review of the bunching literature.

existing bunching methodology by exploiting the heterogeneity in behavior across officers around the bunch point. Because a significant share of officers practice no lenience, we can use these officers to estimate the true distribution of speeds across racial groups.

The rest of the paper is organized as follows. We summarize the existing literature in Section 2. Section 3 provides some institutional background on the Florida Highway Patrol and describes the data. In Section 4, we present a basic conceptual framework. Section 5 describes our empirical strategy and we discuss the results in Section 6. In Section 7, we present and estimate a model of officer behavior and discuss the results. Section 8 concludes.

2 Summary of Relevant Literature

Our paper follows a long line of research exploring the role of law enforcement discretion in generating racial disparities. Police officers are given enormous freedom in who they choose to arrest [Reiss (1973), Muir (1979)]. Once an arrest is made, prosecutors have discretion in choosing who to charge with a crime (Kessler and Piehl, 1998). A central challenge confronting this research is that individuals treated leniently are often not observed in the data. Police do not record arrests that are not made, and court records often do not include information on cases prosecutors declined to try.[6] An additional challenge facing research on racial discrimination is controlling for contextual factors that vary by incident. Some criminal justice disparities across racial groups may be due to differences in criminality, which are difficult to observe.

The current literature has offered various methods for detecting racial bias in the presence of these challenges. The most prominent is the "hit rate test." This test, pioneered by Becker (1957) stipulates that an agent who is trying to maximize some outcome regardless of race should have the same marginal "success" in the outcome for every racial group. Knowles *et al.* (2001), henceforth KPT, use this methodology to study police officers in the choice to search a motor vehicle. When no bias is present, officers should be equally successful in finding contraband across racial groups. Using their method, KPT find no evidence of bias against blacks and significant bias against Hispanics in the searching of motor vehicles.

One contention with the results from KPT is that it can only explain average disparities in treatment and implies homogenous behavior by the police force. If there are any differences across officers in how disparate their treatment is, it cannot be explained by statistical discrimination. Anwar and Fang (2006) follow this logic to develop a test of bias which we call *external bench-*

[6] Rehavi and Starr (2014) examine the racial disparity in federal criminal sentencing and provide a major innovation by linking federal criminal arrests to the prosecutor's choice of charges and to the final sentencing. This linkage resolves the above challenge and allows them to show that prosecutors choose harsher charges against blacks relative to comparable white defendants.

marking. If there is no bias, and disparities in treatment are due only to statistical discrimination or unobserved heterogeneity, then there should be no difference in behavior between white and minority officers. Anwar and Fang find that white officers are significantly more likely to search minority motorists than minority officers, violating the KPT model and suggesting bias by some group of officers.

A central issue with the external benchmarking methodology is that it does not pinpoint which officers are biased. While Anwar and Fang (2006), Antonovics and Knight (2009) and Price and Wolfers (2010) all find evidence of bias by agents in their settings, they are unable to say which individuals or even which race of individuals is biased. Our methodology, which we refer to as *internal benchmarking*, uses differences within officer to estimate the bias at an individual-level.

The setting we examine, where officers choose to reduce the ticketed speed of a driver, in large part avoids the unobservability of lenience and the difficulty in controlling for criminality. Firstly, we are still able to see drivers who are given a reduced speed, allowing us to quantify the degree of leniency by each officer. Secondly, traffic tickets are a setting where the "underlying criminality" is simply the true speed rather than a more complex set of contextual factors.

Our study is not the first to examine the practice of speed discounting by traffic officers. Anbarci and Lee (2014) first document the phenomenon using citations from the Boston Police Department, showing that a significant proportion of tickets list 10 mph over the limit, right below a jump in charged fine. Utilizing a "rank order test" similar to Anwar and Fang (2006), Antonovics and Knight (2009), and Price and Wolfers (2010), they show that black and Hispanic officers are relatively more harsh than white officers when ticketing minority drivers.

This study departs from Anbarci and Lee (2014) in several respects. While their empirical strategy relies on comparisons across officer and driver race, our method constructs an officer-by-officer estimate of discrimination. This approach allows us to see the entire distribution of police preferences and determine how many officers account for the aggregate disparity, which to our knowledge has not been done before in a study of discrimination. Our data include a more expansive set of driver characteristics that may be correlated with driver race, including vehicle description and previous ticketing history, allowing us to rule out several alternative hypotheses for the observed disparities. We further use our estimates to determine what percentage of the overall racial speed gap can be explained by differential leniency and examine various policies aimed at reducing the treatment gap.

3 Institutional Background and Data

3.1 Institutions of the Florida Highway Patrol

State-level patrols are the primary enforcers of traffic laws on interstates and many highways. When on patrol, officers are given an assigned zone, within which they combine roving patrol and parked observation patrol. During the course of a traffic stop for speeding, officers have two primary ways to exercise discretion. They can give a written or verbal warning, which leads to no fine or points on the driver's license, or they can reduce the speed charged on the ticket. The Florida Highway Patrol officers are told explicitly in their training manuals that no enforcement actions during a traffic stop can be based upon any demographic characteristics, including race and gender.

In Florida, driving 10 MPH over the limit leads to about a $75 higher fine than at 9 MPH over.[7] While drivers receive points on their license for speeding, tickets received for 9 and 10 MPH over the limit carry the same number of license points. While it is common to find a jump in fine between 19 and 20 MPH over as well, the data strongly suggest that officers prefer to reduce the ticket to 9 MPH over.

3.2 Data

From the Florida Court Clerks and Comptrollers, we obtained data on traffic citations issued by the Florida Highway Patrol (FHP) for the years 2005-2015. These data include all information provided on the stopped motorist's driver's license – name, address, race, gender, height, date of birth, as well as driver's license state and number. The make, model, and year of the stopped automobile is provided, but this information is recorded inconsistently. In the final sample of citations, 69% of tickets list the vehicle make and year. The citing officer is identified by name, rank, troop number, and badge number.

To supplement the citations data, we obtained officer demographic information from the Florida Department of Law Enforcement (FDLE). These data include officer race, sex, age, education level, and Florida law enforcement employment history of all law enforcement officers employed in the state of Florida. It further includes every misconduct investigation made by the state against an officer, recording the type of alleged violation and the ultimate verdict of the state.

We restrict the sample to citations where the main offense is speeding, no accident is reported, and the cited speed is between zero and 40 above the posted speed. To link the citations and officer information, we first narrowed the list of FDLE personnel to include only officers with an

[7] The actual fine schedule depends on the county in Florida, though the jump point is the same across all counties and always includes at least a $50 jump in fine.

employment spell as a sworn officer with the FHP covering some portion of the 2005-2015 period. We then match the list of candidate officers with the citations data using the officer name. We dismiss stops that cannot be matched to an officer. Lastly we restrict the sample to officers issuing at least 100 citations, with at least 20 given to minorities and 20 to whites. The final sample includes 988,096 citations issued by 1,334 officers. In all of our analysis, we consider speed relative to the speed limit (or posted speed) rather than absolute speed. We often refer to this quantity as *MPH Over* or simply as the speed. We do this so that each driver's speed is rescaled relative to a fine increase point, and after this rescaling, we can pool together citations issued in different speed limit zones.

While the citations record the driver race, there appear to be inconsistencies in the recording of Hispanic. For example, Miami-Dade County has less than 1% of their tickets issued to Hispanic drivers. To deal with this issue, we match the drivers' names to Census records, which record all names that appear more than 1,000 times and the share white, black, Hispanic, and other that carry that name. If an individual in our data has a name that is more than 80% Hispanic, we record them as such.

3.3 Summary Statistics

Table 1 presents summary statistics for the sample, broken out by driver race. 70% of drivers are white, 20% are black, and about 10% are Hispanic. Drivers are 35% female and about 36 years old on average, with Hispanics less likely to be female and minority drivers typically younger. In-state drivers account for 84% of tickets, and the average driver has been cited about 0.04 times in the past year. On average, minority drivers are charged with higher speeds than whites, just over one MPH higher for blacks and almost three MPH higher for Hispanics.

In Table 2, we compare the racial distributions of speeding tickets with the racial distribution of residents and drivers in Florida using the 2006-2010 American Community Survey 1% samples.[8] These data demonstrate that whites account for about 62.5% of Florida's population and 60% of its drivers (an individual is considered a *driver* if they indicate that they drive to work in the ACS), while representing about 58% of tickets. Blacks represent around 14% of the population and driving population, but 18% of tickets. Similarly, Hispanics are 20% of the population, almost 22% of the driving population, and 24% of tickets. In columns 4 and 5, we present the racial distribution of black, white, and Hispanic drivers involved in crashes and crashes with injuries over the 2006-2010 period. These shares are computed from records provided by the Florida Division of Motorist

[8]We obtained these data from IPUMS. So that the samples are parallel, we use only citations from 2006-2010 and keep only white, black, or Hispanic individuals aged 16 or over in the ACS. We use sampling weights when computing the shares from the ACS data.

Services that contain information on all auto accidents known to police. Relative to the citations data, Hispanics are underrepresented in crashes, which may suggest that Hispanic drivers are targeted for citations relatively more often. Blacks are slightly overrepresented in crashes relative to citations, while the white shares in citations and serious crashes are nearly identical.

3.4 Evidence that Officers Use Discretion

As highlighted in the introduction, our study begins with the observation presented in Figure 1. 31% of tickets are written for exactly 9 MPH over the limit, just below a large fine increase. Less than one percent of tickets are for 8 MPH over, while just over 1% are written for exactly 10. We posit that this bunching is due to systematic lenience, with officers choosing to reduce the charged speed, and therefore the fine faced by the driver, after observing a higher speed. In Figure 2, we present evidence for this theory. Panel A plots the officer-level distribution of lenience, defined as the share of tickets written for 9 MPH or above that are for exactly 9 MPH. A large share of officers appear to exhibit very little lenience. About 16% of officers write no tickets for exactly 9 over, while 30% write less than one percent of tickets for this bunching speed.

Of course, this apparent dispersion in officer-level lenience could be due to differences in speeds and driver characteristics across patrol areas and shifts. To account for this possibility, we compute a residualized measure of officer-level lenience by regressing an indicator for a 9 MPH charge on year, month, day-of-week, posted speed, and county fixed effects, computing residuals, and averaging by officer. Panel B plots the officer-level distribution of these measures and demonstrates that substantial variation in lenience remains after adjusting for the time and location of stops.

We provide further evidence that lenience is an officer-level phenomenon by showing that an officer's residualized share of stops with a 9 MPH charge is highly correlated across time and space. Specifically, we residualize lenience using the same procedure as above and average at the officer × year level (Panel C) or officer × county level (Panel D). In Panel C, we plot each officer's residualized lenience in his year with the second most stops (y-axis) against his residualized lenience in year with the most stops (x-axis). A strong correlation is evident – an officer who charges 9 MPH relatively more often in one year also does so in other years. In Panel D, we plot lenience in the county with the second most stops against lenience in the county with the most stops. The story here is the same, with officers who charge 9 MPH relatively often in one county also likely to do so in other counties. We take this as compelling evidence that bunching in the ticketed speed distribution is generated by the behavior of the officers.

3.5 Racial Disparities

Given the argument that bunching in the ticketed speed distribution results from systematic lenience on the part of officers, we examine whether the extent of bunching differs across racial groups as a first pass test for the presence of bias. Figure 3 plots the speed histograms by driver race (white versus nonwhite) and demonstrates an apparent racial disparity in the likelihood of benefitting from lenience. 35% of white drivers are charged 9 MPH over, while just 25% of nonwhites are charged at the bunching speed.

To assess the magnitude, robustness, and statistical significance of racial differences in charging outcomes, we first estimate regressions of the charged speed (relative to posted) on indicators for black and Hispanic driver. Table 3 presents these results. Black and Hispanic drivers are charged faster 1.2-2.86 MPH faster on average. Adding controls and fixed effects shrinks these magnitudes, particularly for Hispanics, but disparities persist and are highly statistically significant, with estimates implying that black (Hispanic) drivers are charged speeds 0.72 (0.69) MPH faster.

Next, we examine the statistical significance and robustness of racial differences in the probability of being charged at the bunching speed. In particular, we estimate linear probability models of the form

$$y_i = \alpha + \theta_B B_i + \theta_H H_i + \beta X_i + \epsilon_i$$

where y is an indicator for whether the charged speed is the bunching speed. B_i and H_i are indicators for driver race, and Z_j is a vector of stop and driver characteristics. Our goal in these regressions is for θ_R to capture the racial differences in the probability of receiving a charge at the bunching speed conditional on being observed at a higher speed. Therefore, we consider only citations for speeds at or above the bunching speed (9 MPH) in these regressions.

Table 4 presents the estimates. In Column 1, we estimate that black and Hispanic drivers are 3.8 and 14.9 percentage points less likely than white drivers to be cited at 9 MPH above the limit. In Column 2, we add controls for driver gender, an indicator for in-State driver, and a linear and quadratic term in driver age. Columns 3 through 5 progressively saturate the model with County-Speed Zone fixed effects, month and day of week fixed effects, and hour of day fixed effects, respectively.[9] In all regressions, we find at least a 2 percentage point difference in treatment between whites and blacks and a 1.37 percentage point difference between whites and Hispanics.

Table 5 presents further controls for car type, ZIP code income, and previous ticketing history. Because these records are each available for only a subset of drivers, the odd columns present

[9] Hour is missing for about 5% of stops. Further, we are skeptical of its reliability because hour is recorded using both 12 and 24 hour methods, and although the data include an AM/PM field, it does not appear to be reliably used. After adjusting to a 24 hour method using this field, we still find a sharp drop in the number of citations at between 12 noon and 1 pm.

regression (4) from Table 4 on the restricted samples for which each variable is available. The broad message is that the racial disparities persist when controlling for these characteristics. In column (2), we control for a quadratic in vehicle age and vehicle make fixed effects (e.g. Ford, BMW). In column (4), we control for the log of per-capita income in the driver's zip code. To do this, we matched the driver's zip code of residence indicated on the driver license to publicly available IRS data on total earnings and number of tax returns filed in each zip code. If anything, racial disparities increase when accounting for vehicle characteristics or income, suggesting that race is not picking up unobserved differences in income that are dictating the officer lenience. Perhaps most importantly, the gaps remain when accounting for previous tickets. In column (6), we control for a quadratic in the number of tickets a driver has received in the past three years. We compute this number directly by linking drivers across tickets using the driver's license state and number. The coefficients are essentially unchanged from Column 5. While we see only tickets written by the Florida Highway Patrol, these regressions suggest that differences in driving record are not generating the treatment disparity.

4 Conceptual Framework

Here we introduce a simple framework of officer decision-making that can explain several features of the data and motivates our strategy for estimating officer-level bias. Officer j stops motorist i for speeding. His observed speed x' is drawn from some discrete distribution $F(\cdot)$. We assume a simple discontinuous fine structure, where the fine for speeding depends on the charged speed x according to

$$Fine(x) = \begin{cases} p_L & \text{if } x \leq x_d \\ p_H & \text{if } x > x_d \end{cases}$$

with $p_H > p_L$. If the driver's speed is above x_d, the officer has the choice to reduce the charged speed to x_d to reduce the fine the driver will face. Otherwise the speed is set to x'. When deciding whether to reduce the ticket, we assume the officer weighs a mix of personal and policing objectives, such as the blameworthiness of the individual and the potential deterrence effect of ticketing the individual. We represent the benefit to reducing the ticket as $V_j(r_i, \xi_i, \epsilon_i)$, where r_i is the driver's race, ξ_i are all other characteristics observed by the officer, and ϵ_i is a random taste shock drawn from a mean-zero symmetric distribution $G(\cdot)$. We assume the cost is an increasing function of the observed speed, $c(x')$.

The officer discounts the drivers ticketing if the value is greater than the cost:

$$x = \begin{cases} x_d & \text{if } V(r_i, \xi, \epsilon) \geq c(x_i) \\ x' & \text{otherwise} \end{cases}$$

To add structure to the framework, we suppose officers' value of discounting is separable between observables and the taste shock:

$$V(r_i, \xi, \epsilon) = t_{rj} + \xi_i \cdot \beta_j + \epsilon_{ij}$$

For a given speed x', the probability an officer discounts the driver is given by

$$\begin{aligned} \Pr(\text{Discount} \mid \text{x}) &= \Pr(t_{rj} + \xi_i \cdot \beta_j + \epsilon_{ij} \geq c(x)) \\ &= G(t_{rj} + \xi_i \cdot \beta_j - c(x)) \end{aligned} \tag{1}$$

In this framework, it is natural to define bias in the following way: An officer is biased against group r relative to group r' if $t_{rj} < t_{r'j}$. This definition implies that a biased officer, when faced with i and i' with the same speed and observables but different races, has a higher probability of discounting i' if $t_{rj} < t_{r'j}$.

The first empirical step we take is to model the likelihood of an individual appearing at the discount point, given his observables. In our model, the probability of being charged the discount speed is the summed likelihood of appearing at or above that speed times the likelihood of being discounted:

$$\Pr(X_i = x_d) = F_i(x_d) + \sum_{k=x_d+1}^{\bar{x}} F_i(k) \cdot G(t_{rj} + \xi_i \cdot \beta_j - c(k)) \tag{2}$$

5 Empirical Strategy

If the driver's true speed x′ could be observed, the framework above suggests estimating officer bias by fitting an equation in the form of equation (1):

$$y_{ij} = \gamma \xi_j + f(x_j) + \sum_i \left(\theta_i + \theta_i^B B_j + \theta_i^H H_j \right) + \epsilon_{ij} \tag{3}$$

where y_{ij} is an indicator that motorist j stopped by officer i is charged x_d (receives the low punishment), using a sample of stops where the true speed $x > x_d$. The vector ξ_j is a set of driver characteristics analogous to those appearing in the officer value function. θ_i is a an officer fixed effect, which captures variation in the cost function C_i across officers. $\theta_i^B B_j$ and $\theta_i^H H_j$ are officer fixed effects interacted with driver race indicators, which capture variation across officers in the

λ_i's.

We cannot estimate (3) directly because we cannot observe the true speed x. The alternative is to estimate equation (2) accounting for the non-linear relationship between officer preferences and the drivers' distribution of speeds F(·). While we perform exactly this procedure in Section 7, we first approximate this equation with a linear probability model, restricting to stops with $x \geq x_d$ and estimating

$$y_{ij} = \gamma \xi_j + \sum_i \left(\theta_i + \theta_i^B B_j + \theta_i^H H_j\right) + \epsilon_{ij} \tag{4}$$

The coefficients θ_i^B and θ_i^H are our measure of officer i's bias against black and Hispanic drivers, respectively. Note that under the assumption that the speed distributions are equal across races, as well as some linearity assumptions about the functional forms of V and F, it can be shown that $\theta_i^B \propto t_i(w) - t_i(b)$. The driver covariates include gender, whether the driver is in-state, and linear and quadratic terms for driver's age and previous number of tickets. In the estimation, we also include year, month, day-of-week, posted speed, and county fixed effects.

As noted in the teacher value-added literature, the distributions of the $\hat{\theta}_i$'s will, in general, be too dispersed relative to the true distribution due to estimation error (Koedel *et al.*, 2015). To account for estimation error, we use standard Empirical Bayes techniques (Morris, 1983) to shrink the estimates towards the mean. However, the distribution of shrunken bias estimates is not equivalent to a measurement-error adjusted estimate of the distribution of bias. With this in mind, we follow a standard deconvolution procedure from Delaigle *et al.* (2008), described in detail in the Appendix to estimate a non-parametric distribution of bias accounting for estimation error in the individual bias estimates. Such procedures are often used to estimate the distribution of persistent wage differences across individuals, as in Postel-Vinay and Robin (2002). To account for the estimation error in regressions where the θ_i's are on the left hand side, i.e. when we show which officer characteristics predict bias, we implement a weighting procedure similar to Aaronson *et al.* (2007), also described in the appendix.

6 Results

6.1 Officer Bias

The distributions of officer-level bias estimates are plotted in Figure 4. The solid lines correspond to our measure of bias against black and Hispanic drivers. The dashed lines plot the distribution of bias where the estimates are shrunken using the Bayes Shrinkage procedure from Morris (1983).[10]

[10] The Bayes Shrinkage approach takes each estimate and adjusts it towards a mean, in the following way: $\tilde{\theta} = \frac{\sigma_\epsilon^2}{\sigma_\theta^2+\sigma_\epsilon^2} \cdot \mu + \frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\epsilon^2} \cdot \hat{\theta}$. How we calculate the mean μ and variance σ_ϵ^2 is in the appendix section *Calculating*

The further right an officer is in the distribution of bias against blacks, the greater his level of bias. The unit of our bias measure is probability difference in percentage points. An officer whose bias against blacks is 0.1, for example, is ten percentage points more likely to offer a fine reduction to a white than black driver. The bias coefficient distributions both peak close to zero but are clearly skewed to the right. The distribution of Hispanic bias, in particular, has a wide right tail.

Figure 5 shows the deconvoluted estimates of bias against blacks and Hispanics. Without measurement error, the skewness of the distribution is more clear. While the mode of the distribution is centered at zero, there is a small percentage of officers with positive bias. This argument is formalized in Figure 6, which shows the CDF's of bias. Slightly more than 20% of the distribution of officers have bias greater than zero against both races.

It is worth noting here the novelty of these distributions relative to the existing literature on racial bias. Table 7 shows the results of using the external benchmarking method of Anwar and Fang (2006) and Antonovics and Knight (2009) on our data. We do the same linear probability regressions as before on the likelihood of appearing at the discount point, now interacting driver race with officer race. The coefficients that identify bias in their framework are the interaction of Black Driver with Black Officer and Hispanic Driver with Hispanic Officer. These estimates reflect how much more often black officers discount black drivers than white officers and how much more often Hispanic officers discount Hispanic drivers than white officers, respectively. The values are .03 and .016, respectively, which are closely in line with our mean estimate of bias using our method. However, they miss the distribution of bias and the degree to which the aggregate disparity is due to a small percentage of officers.

6.2 Do Officer Characteristics Predict Bias?

Given an officer-level measure of racial bias, a natural question is whether officer characteristics are predictive of bias. We can tackle this question using the personnel records collected from the Florida Department of Law Enforcement and the Florida Highway Patrol.

In line with the analysis of Anwar and Fang (2006) and other existing studies, Figure 7 shows how our measure of bias varies by officer race. Perhaps consistent with intuition, white officers are much more likely to be biased against minority drivers. Among black officers, a very small percentage are biased in favor of black driver. This insight is another advancement beyond the current literature. Using the external benchmarking framework, we can know that some race of officers is acting in a biased way, but not which group. Here we can see the magnitude of bias separately for each officer race.

Weights for Bias Prediction Regressions.

In Table 8, we present regressions of officer-level bias on officer characteristics. All observations are weighted by the variance of the noise in our estimate of the officer's bias, as explained in the Appendix. In columns (1)-(3) the dependent variable is an officer's bias against black drivers, while in columns (4)-(6), officer bias against Hispanic drivers is on the left hand side. Columns (1) and (4) use as the outcome variable our measure of bias, columns (2) and (5) use an indicator for whether the measure is above 0, and columns (3) and (6) use an indicator for wether the the bias is above zero and statistically significant (t-value greater than 1.96).

As with the density plots, the clear takeaway from the regressions is that minority officers are more lenient towards drivers of their own race, as we might expect. Female officers appear less biased against both black and Hispanic drivers. Older officers exhibit more bias against both groups, though the standard errors are large. Officers with any higher education are less likely to be biased towards blacks but not Hispanics. Neither complaints nor seeking promotion have consistent predictive power on bias.

Table 9 presents the same results but only for officers who seek promotion to sergeant and in which we include the promotional exam score (averaged if the officer has multiple attempts). The results look similar though noisier due to the smaller sample of officers. It does not appear that an officer's test score has consistent predictive power for the level of bias and, surprisingly, the point estimates go in opposite directions for blacks versus Hispanics.[11]

6.3 Potential Concerns

In the following section, we analyze a model that accounts for differences in speeds across racial groups. Here we consider two other potential concerns that must be addressed in our analysis, specifically whether there is selection into who is ticketed and the possibility that our results are detecting statistical discrimination.

6.3.1 Selection

All of our analysis is conditional on an individual being stopped and ticketed. There is certainly the potential for racial bias in the officer's choice of who to stop and who to give a warning, a possibility that may bias our estimates of discrimination in the ticketing discount choice.[12] However, we think

[11]It is unclear whether these regressions should also include the average lenience of the officers, so that comparisons are for a given probability of discounting. Though not shown, we have conducted these regressions with average lenience included, and the results look nearly identical.

[12]Several studies have examined whether officers practice racial profiling when deciding which drivers to pull over. Grogger and Ridgeway (2006) find little evidence of racial profiling in Oakland, CA by comparing the racial distributions of stopped motorists during day and night, when the race of the driver is less likely to be known to the officer ex ante. However, Horrace and Rohlin (2016) augment the test used in Grogger and Ridgeway (2006) by considering the location of streetlights in Syracuse, NY. They find that black drivers are

these concerns do not invalidate our results. As discussed in Fryer (2016), the overrepresentation of blacks and Hispanics among stopped drivers will likely bias downwards any measure of later discrimination, as this increases the pool of *good* minority drivers who end up ticketed. Further, the coefficient on prior tickets goes in the direction we would expect and suffers from the same selection problem as race. We therefore conclude that selection on being stopped and ticketed is not driving the racial disparity we observe.

If there is selection in the first step of deciding to ticket a driver, we should perhaps expect that officers who are more biased in terms of discounting will also have a greater share of drivers who are minority. Figure 8 presents how the share black of an officer's ticketed drivers changes with the bias of the officer, where share black is first residualized for each driver's day of week, month, speed zone, and county. While there is a significant difference across officer the distribution of officer bias, suggesting potential selection, the magnitude is exceedingly small. We find that a two standard deviation change in officer's bias against blacks leads to a .075 standard deviation change in the share drivers who are black, and a -.075 standard deviation effect for Hispanics.[13]. These calculations suggest that, while there may be selection, the effect on our estimates of bias should be small.

6.3.2 Statistical Discrimination

We have thus far defined bias as the differential treatment of drivers by race who are stopped for the same speed. This definition is not innocuous, as there may be some reasons for differential treatment unrelated to observed driving speed that, while contentious in their use, are not specifically racial bias. For example, officers may be choosing treatment on the basis of how an individual's future driving responds to punishment [Gehrsitz (2015), Hansen (2015)] or the likelihood of paying a ticket [Rowe (2010), Makowsky and Stratmann (2009)]. If individuals systematically differ by race in these characteristics, the racial disparities we observe may reflect the fact that are officers are statistically discriminating by using race as a proxy for deterrability.

However, as noted in Anwar and Fang (2006), statistical discrimination can only explain behaviors that are uniform across officers, as they are due to the relationship between race and unobserved heterogeneity rather than anything specific to the officer. The median amount of bias in our setting is small, but we find a significant and skewed right tail. Such a distribution of disparate treatment cannot be explained by statistical discrimination.

about 15% more likely to be stopped than nonblack drivers in lighted areas/hours. Antonovics and Knight (2009) find evidence that officers are more likely to stop drivers of difference races in Boston.

[13]These effects are found by performing a local linear regression of residualized share black on officer bias and comparing the value for predicted share black for officers bias one standard deviation above and below the mean.

7 Model of Officer Discount Decision

In the above analysis, we proceeded under the assumption that individuals' driving speed do not vary systematically across racial groups. This section presents a model that allows us to simultaneously estimate officers' taste parameters for each racial group and speed parameters for each race-by-location. By doing so, we can also perform counterfactuals using this model to quantify the effect of various policies that change the distribution of bias across officers. The model setup is as follows.

Officer j encounters individual i driving at a speed drawn from a poisson distribution $x \sim P_{\lambda_i}(s)$, where the poisson parameter depends on the driver's county and race, $\lambda_i = \lambda_{rc}$. The officer faces the choice to either charge the driver his measured speed x or, if the speed is above the jump in fine, discount the speed to x_d. He makes this decision by weighing a cost to discounting, which we impose to have the form $c(x) = b \cdot x$, against the value of discounting, $t_{ij} = t_{rj} + \epsilon_{ij}$. So the driver has her speed reduced to x_d if

$$t_{rj} + \epsilon_{ij} > a + bx_i$$

The officer's preference is allowed to vary by race r. For simplicity, we pool black and Hispanic drivers into a single nonwhite, or minority, group when estimating the model parameters. The noise term ϵ_{ij} is assumed to be a standard normal variable. Thus, for an individual driving at speed x, her probability of discount is

$$Pr(\text{Discount}|x) = \Phi(t_{rj} - b \cdot x_i)$$

Conditional on officer, county, and driver race, the likelihood for each speed is the following:

$$Pr(X = x) = \begin{cases} P_{\lambda_{rc}}(x) & \text{if } x < x_d \\ P_{\lambda_{rc}}(x_d) + \sum_{k=x_d+1}^{X} P_{\lambda_{rc}}(k) \cdot \Phi(t_{rj} - b \cdot k) & \text{if } x = x_d \\ P_{\lambda_{rc}}(x) \cdot \Phi(t_{rj} - b \cdot x) & \text{if } x > x_d \end{cases}$$

7.1 Estimation

While the setup of the model is simple, non-parametrically estimating the distribution of bias is computationally complex. The model parameters to be identified are the 67×2 county-race speeds λ_{rc}, 1327×2 officer average racial preferences, t_{jr}, and the slope of the cost function b, totaling 2,789 parameters. This complexity makes estimation directly through maximum likelihood challenging.

We estimate the model by iterating across the groups of parameters until the solution converges.

We solve first for a set of initial guesses by assuming that speeds and officer preferences are uniform and performing maximum likelihood on b, λ, t. We then calculate the slope, county-specific speeds, and officer-specific preferences in the following way:

1. Guess speed values $\lambda_{rc}^{(0)} = \hat{\lambda}$ and slope $b^{(0)} = \hat{b}$ from the aggregate MLE estimation.
2. Solve for officer preferences $t_{ij}^{(0)}$ by maximizing likelihood $\mathcal{L}\left(t_{jr}|x_{jr}, \{\lambda_{rc}^{(0)}\}, b^{(0)}\right)$ conditional on speed parameters and slope estimate.
3. Solve for speed parameters $\lambda_{rc}^{(1)}$ by maximizing likelihood $\mathcal{L}\left(\lambda_{rc}|x_{jr}, \{t_{jr}^{(0)}\}, b^{(0)}\right)$ conditional on officer preferences and slope estimate.
4. Solve for slope parameter $b^{(1)}$ by maximizing likelihood $\mathcal{L}(b|x_{jr}, \{t_{jr}^{(0)}\}, \{\lambda_{rc}^{(1)}\}\big)$
5. return to step 2 and repeat until parameter guesses converge.

Because in every iteration the total likelihood increases[14], the process will converge at least to a local optimum. We check different starting values to confirm that our results achieve a global maximum.

Conditional on speed and slope, the officer parameters are separable and thus can be easily solved, and similarly for the speed parameters when conditioning on officer parameters. Further, the conditional likelihood functions are unimodal in the parameters; this means that the score functions only cross zero once, simplifying the search for an optimum. Standard errors are calculated by estimating the information matrix via the variance of the parameters' score functions.

7.2 Results

Table 10 presents the estimates of the model parameters. The columns present the mean and variance of each class of parameters, broken down by race, and the final column compares differences across racial groups in the mean parameter estimates. The slope parameter is positive and significant at .023. Consistent with our intuition, officers face an upwards sloping cost with respect to speed, meaning that tickets are less likely to be discounted the higher the observed speed. The parameter t represents an officer's mean valuation of a racial group. We find both significant heterogeneity and a significant disparity across whites and minorities in how officer's value discounting drivers, officers' mean valuation for whites being .1 higher than for minorities. While the values of t are by themselves hard to interpret, dividing them by the slope parameter b gives the interpretation of the valuation in terms of miles per hour driven. The difference of .1 in

[14]this approach is very similar to the EM algorithm, most commonly used to deal with missing data.

valuation between whites and minorities, scaled by .0228, tells us that the average officer treats a minority driver like a white driver stopped for driving 4 MPH faster.

These differences in treatment are more easily understood in terms of the probability of discount (i.e. fine reduction). *P(Discount)* represents the likelihood of receiving a reduced ticket if the driver is at the speed right above the bunching speed. Consistent with the reduced form evidence, the average officer is substantially lenient, with a large variance across officers. Officers are 3.3 percentage points less likely to discount minorities than whites, off a baseline of 35.7% likelihood of discount. Figure 9 further shows this disparity, highlighting how racial bias results in a decreased mass of officers with very high lenience and an increase in mass of officers with very low lenience.

The λ estimates tell us how races-by-counties differ in their underlying speeds prior to officers choice of lenience. We find that minorities on average drive significantly faster than whites, on the order of .5 to .7 MPH. Figure 11 presents this gap by county, showing that minority speeds stochastically dominate white speeds. These results are in line with previous studies of highway patrol ticketing, which argue that much of the gap in ticketing between whites and minorities can be explained by higher speeds by minorities [Smith *et al.* (2004), Lange *et al.* (2005)]. However, these previous studies and the news coverage that followed implicitly argued that the racial difference in speeds rules out the presence of bias by officers. Our study highlights how this thinking is incorrect by showing that disparities in driving and racial bias coexist in our setting.

Note that our estimates of driver speeds are conditional on the officer choosing to ticket the driver in question. As discussed above, there is a possibility that officers are biased in the original choice of whether to ticket a driver. If so, the threshold for minority drivers to be in the sample would be lower, suggesting that our observed gap is smaller than it would be without the presence of bias in the choice to ticket.

A central question in most studies of racial bias is the extent to which an aggregate racial disparity can be explained by the measured amount of bias. Table 11 seeks to answer this question by decomposing the measured discount probability and speed disparities across races. The first column presents the racial gap in likelihood that each individual's officer would ticket him were he at the speed right above the bunch point:

$$\text{Probability Gap} = \frac{1}{N_w}\sum_{r(i)=w} \text{Prob}(\text{Discount}_{ij}|x_i = x^* + 1) - \frac{1}{N_h}\sum_{r(i)=m} \text{Prob}(\text{Discount}_{ij}|x_i = x^* + 1) \quad (5)$$

While the average Florida officer only has a .03 racial difference in probability of discounting, the effective difference in the data is .09 because minority drivers are in counties with less lenient officers. Figure 12 shows how counties vary in their share minority, average bias, and white lenience. The most striking fact of these figures is that the regions with the greatest lenience towards whites are also the regions with the least minorities. This disconnect implies that even removing the bias of every officer will leave substantial disparities in treatment based on the geographic distribution of minorities and lenient officers. We formally explore the role of sorting by simulating the model in the case where individuals draw officers from a state-wide distribution in column 3, and the probability gap reduces to .03. The speed gaps is reduced by 47% with no bias or sorting, reflecting the total share of the disparity that can be explained by officer bias.

7.3 Counterfactual Analysis

The model we presented is useful for ruling out that the reduced form estimates of bias are not due to unobserved differences in speeds across races. While we find that minorities drive faster on average, we continue to find bias on the part of some officers. The model is also useful for decomposing the role of bias, lenience sorting, and speed differences in explaining the aggregate racial speed gap.

We now use the estimates to conduct a series of policy counterfactuals exploring how best to curb bias in speeding tickets. Because we provide a non-parametric estimate of the distribution of bias across officers and locations, we can explore a rich set of counterfactuals whose outcomes depend on the full distribution of bias. Because no previous studies have been able to provide distributional information on bias, we are the first to be able to explore the potential effects of such policies.

7.3.1 Firing Biased Officers

The first counterfactual we consider is firing the worst officers in the sample. While it is commonly believed that officers are difficult to fire, we want to explore the degree to which aggregate disparities can be reduced by small adjustments at the extremes of bias. To make the exercise symmetric, we fire officers of similar bias both against and for minorities (i.e. reverse discrimination). We carry out the exercise as follows: we calculate the p-th percentile of most biased officers, which we denote by C, and remove all officers with bias greater than C and lower than $-C$. We re-draw each individual's ticketing officer from the distribution of remaining officers within each driver's county, weighted by their number of tickets. If there are too few remaining unbiased officers in that driver's county, we draw from the remaining officers in the troop (which encompasses 8 counties). We then calculate the aggregate disparity in probability of being discounted if the driver were hypothetically driving at 10 MPH over the limit. Figure 13 explains graphically how the truncation is conducted.

The results are presented in Figure 14. The x-axis plots the percentile of bias being truncated, and the y-axis plots the racial gap in probability of discount. We see that truncating bias leads to a reduction in the probability gap. Note that this reduction is not mechanical. Officers are being removed from both directions of bias, so the decline in the racial gap is due to the greater mass of officers with pro-white bias. Removing officers who are at the 15th percentile or worse of bias reduces the gap from .093 to .08, a 14% reduction. In the limit, only officers who treat each race equally are left, and bias is cut in half to .05. Note that, because the distribution of bias is skewed and centered at a positive level of bias, the limit of no bias is at the 57th percentile. As before, the gap is not reduced to zero because minorities are in areas where, regardless of race, officers are less lenient to drivers. This fact suggests that the target of interventions should not be bias per se but absolute levels of lenience. Note that the results are non-monotonic both because of simulation error and because the cutoffs at different levels of the x-axis may differentially affect the presence of pro-white or anti-white bias. Specifically, a new cutoff may remove only one additional officer, and it may be an officer who was biased against whites.

7.3.2 Firing Lenient or Harsh Officers

Following this argument, we consider how the aggregate treatment disparity is changed with a policy targeting lenience rather than bias. We carry out two counterfactuals in this spirit, one reducing lenience and the other increasing it. Figure 15 plots what happens when the least lenient officers are removed. The x-axis is the minimum percentage of drivers at 10MPH over who must be discounted for an officer to be kept in the sample[15]. As the threshold increases, the aggregate disparity decreases substantially. In the limit, where all drivers are given a reduced ticket, the gap in treatment mechanically goes to zero because all drivers are discounted.

The opposite counterfactual is to consider removing the most lenient officers, which we present in Figure 16. The x-axis shows the threshold share of drivers discounted above which officers are removed from the sample. The reduction in disparity occurs faster than removing strict officers, due to the fact that bias can only arise in officers who exhibit lenience to at least one group. Similar to removing strict officers, the limiting case removes all of the aggregate gap.

While not the same exercise, we think of these counterfactuals as reflecting a potential policy where the highway patrol imposes a minimum or maximum standard that all officers must satisfy. Though not shown, we have performed such an analysis, and the results look nearly identical. Perhaps surprisingly, these policies targeting lenience appear to be more effective than simply removing biased officers.

[15]We look separately at the officers share of white drivers discounted and share minority. Both must satisfy the threshold.

7.3.3 Increasing Minority and Female Shares

We next consider increasing the share of minority or female officers. Given our earlier finding that minority and female officers exhibit lower levels of bias, we should expect that increasing their presence might lead to lower levels of aggregate bias. We do so by re-simulating which officer each driver draws, taken from within his county, where the probability of drawing a minority or female officer is exogenously changed. Results are presented in Figures 17. Consistent with our intuition, the gap in probability of discount declines, though very modestly. An increase in minority officers from the empirical share of 34% to 60% reduces the gap from .093 to .088. Increasing the share of women leads to a similar effect, with a change from 9% to 25% leading to a reduction from .093 to .088 as well.

Demographic policies have been suggested in the past as a possibility for systemically changing police behavior, particularly towards poor and minority communities. Donohue III and Levitt (2001) find that an increase in minority officers leads to an increase in arrests of white offenders, no effect on non-white offenders, and vice versa for an increase in white officers. Our results, though only counterfactuals, are consistent with their findings.

7.3.4 Assigning Officers by Bias and Lenience

The final counterfactuals we consider are to reassign officers to specific areas based on their behavior and the share of minorities in each county. Officers are assigned to Troops, which patrol 6-10 counties. Within the troops, officers regularly vary in which locations they patrol. It may be potentially feasible for a senior officer to, for example, change the assignment of officers such that minorities face less biased officers. Table 12 presents the results of such a policy. The first column is the baseline simulation of the model to match the true data. The second column sorts officers within a troop (which comprise 6-10 counties) such that the least biased officers are in counties with the most minorities. The third column sorts officers within a troop such that the most lenient officers are in counties with the most minorities.

Surprisingly, sorting officers to expose minorities to the least biased has a deleterious effect on the treatment gap. The least biased officers are also the least lenient on average, leading minorities to be treated poorly relative to whites, now exposed to lenient and biased officers. The gap in probability of discount increases from .097 to .102. Much more effective in reducing the gap in treatment is assigning the most *lenient* officers to minority counties. This policy reduces the treatment gap from .096 to .0047. The gap in speeds declines from -2.11 to -1.21, nearly identical to the true speed gap, -1.11.

In short, the counterfactual analyses highlight the importance of absolute lenience as a consideration separate from bias. The policy aimed at exposing minorities to lenience is much more effective, as are policies making lenience uniform, than removing overall bias through firing biased officers or hiring minority and female officers.

8 Conclusion

The glaring racial disparities in the criminal justice system have led many to claim bias as the root cause. Proving so is surprisingly hard, as it is difficult to truly control for differences in criminality and show that individuals are treated differently by race for the same offense. This study explores the question of bias in the criminal justice system and the extent to which it explains aggregate racial disparities. Specifically, we study speeding tickets and the choice of officers to discount drivers to a speed right below an onerous punishment. By using a bunching estimator approach that allows for officer-by-race measures of lenience in tickets, we can explore the entire distribution of both lenience and bias on the part of officers. We find that about 20% of officers explain all the aggregate bias, and 46% of the gap in charged speeds can be explained by differential exposure to lenience. The rest of the gap is due to underlying differences in driving speeds across races.

We explore whether bias is predictable by regressing individual officers' bias on demographic and personnel characteristics. We find that officers tend to favor their own race, older officers are more racially biased, and women and college educated officers are less biased on average. Personnel information, such as failing an entry exam, receiving civilian complaints, and seeking a promotion, are not strongly informative about bias.

Using a model of driver speeding and officer decision-making, we confirm that, while minorities drive faster on average, our officer-level estimates of bias are not confounded by differences in speeding across groups. We find that setting bias to be zero across officers fails to remove the majority of the treatment gap due to the fact that minorities tend to live in regions where officers are less lenient towards all drivers. Because of this fact, we find that policies directed at reducing bias directly have a significant but modest effect on the treatment gap. Policies that instead target officers' lenience, either by firing overly harsh or lenient officers or by re-assigning lenient officers to minority neighborhoods, are much more effective at reducing the aggregate treatment disparity.

Popular debate over police misconduct tends to revolve around a discussion of whether misbehavior is systemic or the product of a few bad apples. We make progress on this question by providing a specific answer: racial bias is due to 20% of officers, and there are effective policies for mitigating their harm.

References

Aaronson, D., Barrow, L. and Sander, W. (2007). Teachers and Student Achievement in the Chicago Public High Schools. *Journal of Labor Economics*, **25** (1), 95–135.

Agan, A. Y. and Starr, S. B. (2016). Ban the box, criminal records, and statistical discrimination: A field experiment.

Alesina, A. and La Ferrara, E. (2014). A Test of Racial Bias in Capital Sentencing . *American Economic Review*, **104** (11), 3397–3433.

Anbarci, N. and Lee, J. (2014). Detecting Racial Bias in Speed Discounting: Evidence from Speeding Tickets in Boston. *International Review of Law and Economics*, **38**, 11–24.

Antonovics, K. and Knight, B. (2009). A New Look at Racial Profiling: Evidence from the Boston Police Department. *Review of Economics and Statistics*, **91** (1), 163–177.

Anwar, S. and Fang, H. (2006). An alternative test of racial prejudice in motor vehicle searches: Theory and evidence. *The American Economic Review*, **96** (1), 127–151.

Ashenfelter, O. and Greenstone, M. (2004). Using Mandated Speed Limits to Measure the Value of a Statistical Life. *Journal of Political Economy*, **112** (S1), S226–S267.

Becker, G. S. (1957). *The economics of discrimination: an economic view of racial discrimination*. University of Chicago.

Brown, M. K. (1988). *Working the street: Police discretion and the dilemmas of reform*. Russell Sage Foundation.

Chalfin, A., Danieli, O., Hillis, A., Jelveh, Z., Luca, M., Ludwig, J. and Mullainathan, S. (2016). Productivity and Selection of Human Capital with Machine Learning. *American Economic Review*, **106** (5), 124–127.

Coate, S. and Loury, G. (1993). Will Affirmative-Action Policies Eliminate Negative Stereotypes? *American Economic Review*, **83** (5), 1220–1240.

Coviello, D. and Persico, N. (2013). An Economic Analysis of Black-White Disparities in NYPD's Stop and Frisk Program. *NBER Working Paper*, pp. 1–28.

Day, M. V. and Ross, M. (2011). The Value of Remorse: How Drivers' Responses to Police Predict Fines for Speeding. *Law and Human Behavior*, **35** (3), 221–234.

De Paola, M., Scoppa, V. and Falcone, M. (2010). The Deterrent Effects of Penalty Point System in Driving Licenses: A Regression Discontinuity Approach. *Universita Della Calabria Economics Dept Working Paper*, pp. 1–21.

DeAngelo, G. and Hansen, B. (2014). Life and Death in the Fast Lane: Police Enforcement and Traffic Fatalities †. *American Economic Journal: Economic Policy*, **6** (2), 231–257.

— and Owens, E. G. (2014). Learning the Ropes: Task Specific Experience and the Output of Idaho State Troopers. *Working Paper*, pp. 1–50.

Dee, T., Dobbie, W., Jacob, B. and Rockoff, J. (2016). The Causes and Consequences of Test Score Manipulation: Evidence from the New York Regents Examination. *NBER Working Paper*, pp. 1–67.

Delaigle, A., Hall, P. and Meister, A. (2008). On deconvolution with repeated measurements. *The Annals of Statistics*, pp. 665–685.

Diamond, R. and Persson, P. (2016). The Long-Term Consequences of Teacher Discretion in Grading of High-Stakes Tests. *SIEPR Discussion Paper*, **16-003**, 1–70.

Donohue, J. (2014). An Empirical Evaluation of the Connecticut Death Penalty System Since 1973: Are There Unlawful Racial, Gender, and Geographic Disparities? *Journal of Empirical Legal Studies*, **11** (4), 637–696.

Donohue III, J. J. and Levitt, S. D. (2001). The impact of race on policing and arrests. *The Journal of Law and Economics*, **44** (2), 367–394.

Fryer, R. G. (2016). An Empirical Analysis of Racial Differences in Police Use of Force. *NBER Working Paper*, **22399**, 1–63.

Garrett, T. A. and Wagner, G. A. (2009). Red Ink in the Rearview Mirror: Local Fiscal Conditions and the Issuance of Traffic Tickets. *The Journal of Law and Economics*, **52** (1), 71–90.

Gehrsitz, M. (2015). Speeding, Punishment, and Recidivism Evidence from a Regression Discontinuity Design. *Unpublished Manuscript*, pp. 1–45.

Gelman, A., Fagan, J. and Kiss, A. (2007). An Analysis of the New York City Police Department's "Stop-and-Frisk" Policy in the Context of Claims of Racial Bias. *Journal of the American Statistical Association*, **102** (479), 813–823.

Goldin, C. and Rouse, C. (2000). Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians. *American Economic Review*, **90** (4), 715–741.

Goldstein, J. (1960). Police discretion not to invoke the criminal process: Low-visibility decisions in the administration of justice. *The Yale Law Journal*, **69** (4), 543–594.

Grogger, J. and Ridgeway, G. (2006). Testing for Racial Profiling in Traffic Stops From Behind a Veil of Darkness. *Journal of the American Statistical Association*, **101** (475), 878–887.

Hansen, B. (2015). Punishment and Deterrence: Evidence from Drunk Driving †. *American Economic Review*, **105** (4), 1581–1617.

Horrace, W. C. and Rohlin, S. M. (2016). How Dark Is Dark? Bright Lights, Big City, Racial Profiling. *Review of Economics and Statistics*, **98** (2), 226–232.

Kessler, D. P. and Piehl, A. M. (1998). The role of discretion in the criminal justice system. *Journal of Law, Economics, & Organization*, pp. 256–276.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, **8**, 435–464.

Knowles, J., Persico, N. and Todd, P. (2001). Racial Bias in Motor Vehicle Searches: Theory and Evidence. *Journal of Political Economy*, **109** (1), 203–229.

Koedel, C., Mihaly, K. and Rockoff, J. E. (2015). Value-Added Modeling: A Review. *Economics of Education Review*, **47** (C), 180–195.

Lange, J. E., Johnson, M. B. and Voas, R. B. (2005). Testing the racial profiling hypothesis for seemingly disparate traffic stops on the new jersey turnpike. *Justice Quarterly*, **22** (2), 193–223.

Lee, D. (2012). Do Traffic Tickets Reduce Motor Vehicle Accidents? Evidence from a Natural Experiment. *Unpublished Manuscript*, pp. 1–32.

Makowsky, M. D. and Stratmann, T. (2009). Political Economy at Any Speed: What Determines Traffic Citations? *American Economic Review*, **99** (1), 509–527.

— and — (2011). More Tickets, Fewer Accidents: How Cash-Strapped Towns Make for Safer Roads. *The Journal of Law and Economics*, **54** (4), 863–888.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, **142** (2), 698–714.

Morris, C. N. (1983). Parametric Empirical Bayes Inference: Theory and Applications. *Journal of the American Statistical Association*, **78** (381), 47–55.

Muir, W. K. (1979). *Police: streetcorner politicians*. University of Chicago Press.

Persico, N. (2009). Racial Profiling? Detecting Bias Using Statistical Evidence. *Annual Review of Economics*, **1** (1), 229–254.

Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, **70** (6), 2295–2350.

Price, J. and Wolfers, J. (2010). Racial Discrimination Among NBA Referees. *Quarterly Journal of Economics*, **125** (4), 1859–1887.

Quintanar, S. (2011). Do Driver Decisions in Traffic Court Motivate Police Discrimination in Issuing Speeding Tickets? *LSU Department of Economics Working Paper*, pp. 1–42.

Rehavi, M. M. and Starr, S. B. (2014). Racial disparity in federal criminal sentences. *Journal of Political Economy*, **122** (6), 1320–1354.

Reiss, A. J. (1973). *The police and the public*, vol. 39. Yale University Press.

Ridgeway, G. and MacDonald, J. (2010). Methods for Assessing Racially Biased Policing. In S. Rice and M. White (eds.), *Race, Ethnicity, and Policing New and Essential Readings*, pp. 180–204.

Roach, M. (2015). Is the Highway Patrol Really Tougher on Out-of-State Drivers? An Empirical Analysis. *The B.E. Journal of Economic Analysis & Policy*, **15** (2), 769–796.

Rowe, B. (2008). Gender Bias in the Enforcement of Traffic Laws: Evidence Based on a New Test. *American Law and Economics Association Annual Meetings*, pp. 1–36.

— (2010). Discretion and Ulterior Motives in Traffic Stops: The Detection of Other Crimes and the Revenue from Tickets. *Unpublished Manuscript*, pp. 1–27.

Smith, D. A., Visher, C. A. and Davidson, L. A. (1984). Equity and discretionary justice: The influence of race on police arrest decisions. *The Journal of Criminal Law and Criminology (1973-)*, **75** (1), 234–249.

Smith, W., Tomaskovic-Devey, D., Zingraff, M., Mason, H. M., Warren, P. and Wright, C. (2004). The North Carolina Highway Traffic Study. *NCJRS Grant Report*, pp. 1–407.

Walker, S., Spohn, C. and DeLone, M. (2012). *The color of justice: Race, ethnicity, and crime in America*. Cengage Learning.

West, J. (2015). Racial Bias in Police Investigations. *Working Paper*, pp. 1–29.

Wozniak, A. (2015). Discrimination and the effects of drug testing on black employment. *Review of Economics and Statistics*, **97** (3), 548–566.

Technical Appendix

Calculating Weights for Bias Prediction Regressions

In Section 6.2, we estimate regressions predicting officer bias, where each officer's left-hand side variable is measured with noise:

$$\begin{aligned} \hat{\theta} &= X_i\beta + \epsilon_i \\ \theta &= X_i\beta + \epsilon_i + u_i \end{aligned}$$

where $V_u = V(u_i)$ is the variance of the parameter estimate, and $V_e = V(\epsilon_i)$ is the true variation in bias. We therefore use a weighted least squares regression approach to deal with heteroskedasticity.

While we know V_u from our estimation procedure for θ, we have to calculate V_e. We do so using the following iterative procedure, from Morris (1983) :

$$\begin{aligned}
&(1) \quad \text{guess starting } \mu,\ \mu^{(0)} \equiv \frac{1}{N}\sum_i \hat{\theta}_i \\
&(2) \quad \text{guess starting V}(\theta_i),\ V^{(0)} \equiv \frac{1}{N}\sum_i^N \left[(\hat{\theta}_i - \mu^{(0)})^2 - V(\epsilon_i)\right] \\
&(3) \quad \text{update guess of } \mu,\ \mu^{(1)} = \frac{\sum_i^N \hat{W}_i \cdot \hat{\theta}}{\sum_i^N \hat{W}_i}, \quad \text{where } \hat{W}_i = \frac{1}{V^{(0)} + V(\epsilon_i)} \\
&(4) \quad \text{update guess of V}(\theta_i),\ V^{(1)} = \frac{\sum_i^N \hat{W}_i\left[\frac{N}{N-1}(\hat{\theta} - \mu^{(1)})^2 - V(\epsilon_i)\right]}{\sum_i^N \hat{W}_i} \\
&(5) \quad \text{iterate until } \mu^{(k)} \approx \mu^{(k+1)}, V^{(k)} \approx V^{(k+1)}
\end{aligned}$$

The use of weights $\hat{W}_i$ reflects the fact that some officers' bias parameters are more precisely estimated than others and should thus be given more weight in estimating the distribution of θ_i. In practice, the final optimal estimates of μ and $V(\theta_i)$ are very close to the initial unweighted estimates, and some papers use these initial guesses as their values for the distribution of θ_i (Aaronson *et al.*, 2007). We then weight each observation by $W = \frac{1}{V_e + V_{ui}}$.

8.1 DeConvolution Procedure

While the Morris (1983) procedure gives us an estimate of the true variance in bias, we are also interested in getting the full distribution of bias, after accounting for the estimation error in the parameters. We do so by using the deconvolution method from Delaigle and Meister (2008), which provides the density estimator

$$\hat{f}_n(\theta) = \frac{1}{2\pi} \int \exp(-it\theta) K^{\text{ft}}(t/\omega_n) \Phi_n(t) dt,$$

where

$$\Phi_n(t) = \sum_{j=1}^{n} f_{\epsilon_j}^{ft}(-t) \exp(itY_j) / \Big(\sum_{k=1}^{n} |f_{\epsilon_k}^{ft}(t)|^2 \Big),$$

K is a square-integrable kernel function, ω_n is a smoothing parameter, and f_{ϵ_k} are the distributions of the error terms. We use the stndard normal distribution for the kernel, and distribution of the error terms are normal with standard deviation equal to the standard error of the parameter estimate.

The deconvolution takes advantage of the fact that when a random variable is a sum of two other random variables, its characteristic function is the product of the variables' characteristic functions. In our case, where we $\hat{\theta} = \theta + \epsilon$, we solve for the distribution of θ by taking the quotient of the characteristic functions of $\hat{\theta}$ and ϵ and applying the characteristic function inversion.

Figure 1: Distribution of Charged Speeds and Fine Schedule

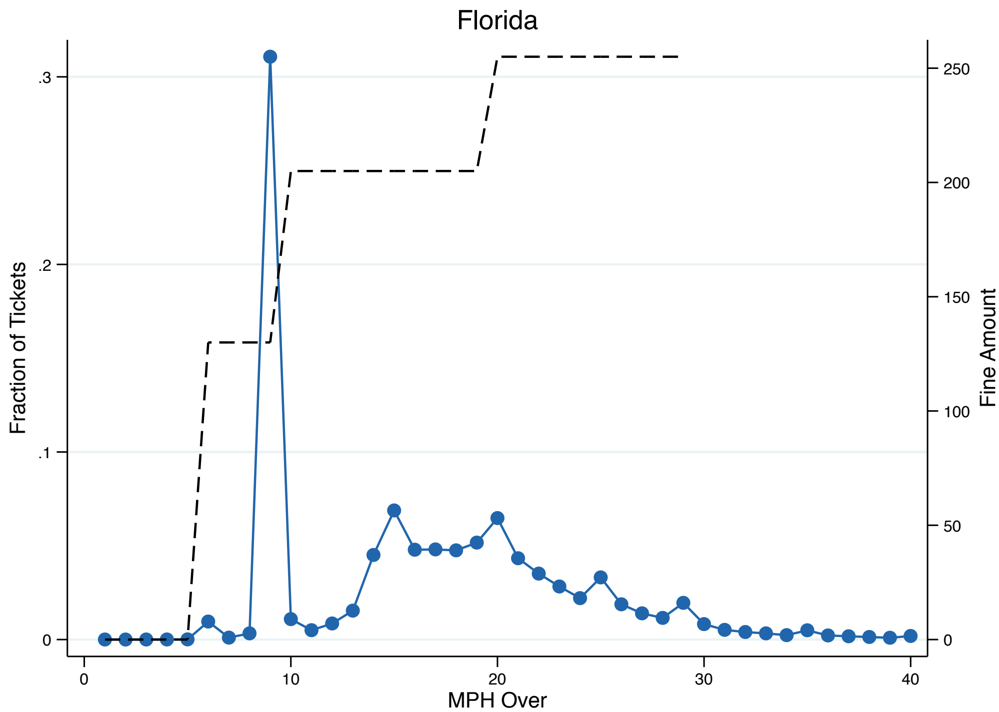

Notes: Connected line shows histogram of tickets. Dashed line plots fine schedule for Broward County. 30 MPH over is felony speeding and carries a fine to be determined following a court appearance.

Figure 2: Evidence of Officer Lenience

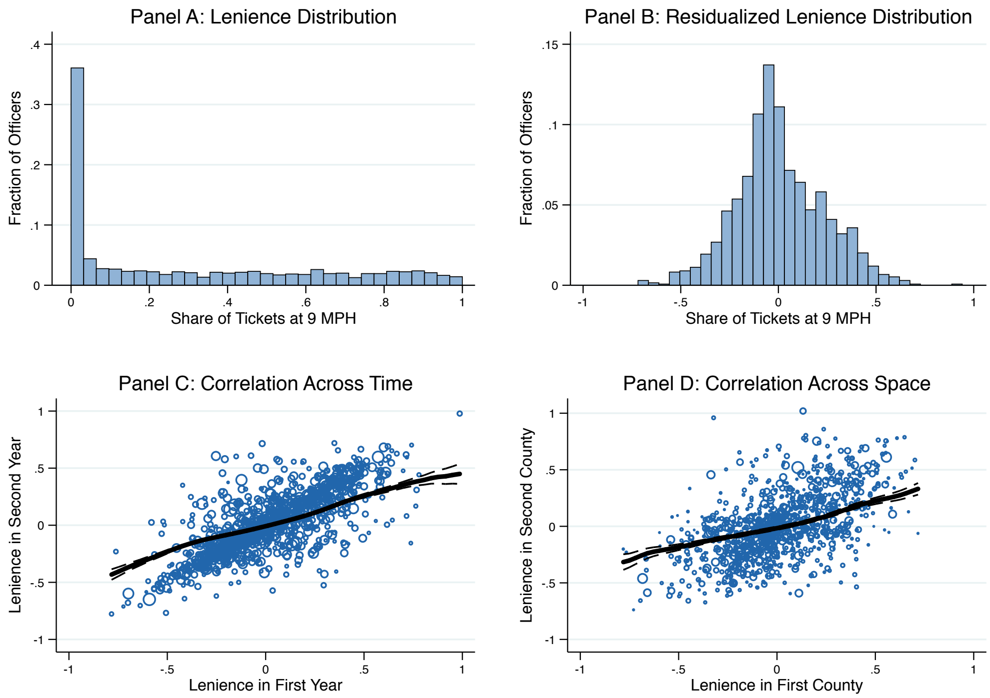

Notes: Panel A plots the across-officer distribution of lenience. Panel B plots the across-officer distribution of residualized lenience. Panel B plots officers' residualized lenience in the years with the most and second most citations. Panel D plots the residualized lenience in the county with the most and second most citations. Estimates residualized by conditioning on county fixed effects, speed zone fixed effects, year and month fixed effects, and day of week fixed effects. See text for additional details.

Figure 3: Charged Speed Distributions by Driver Race

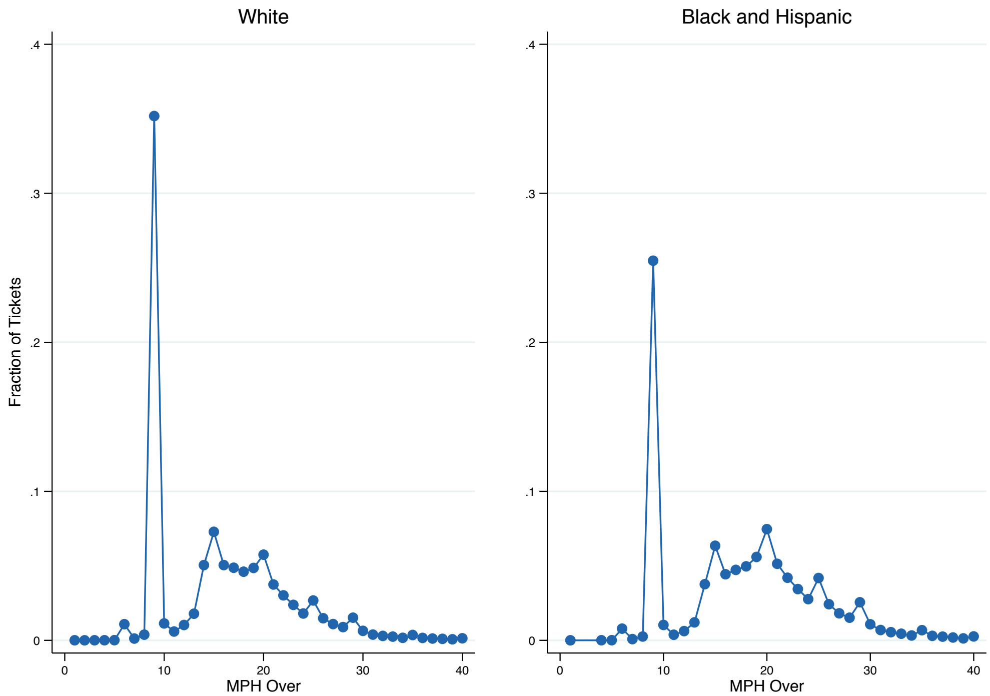

Figure 4: Distributions of Officer Bias Estimates

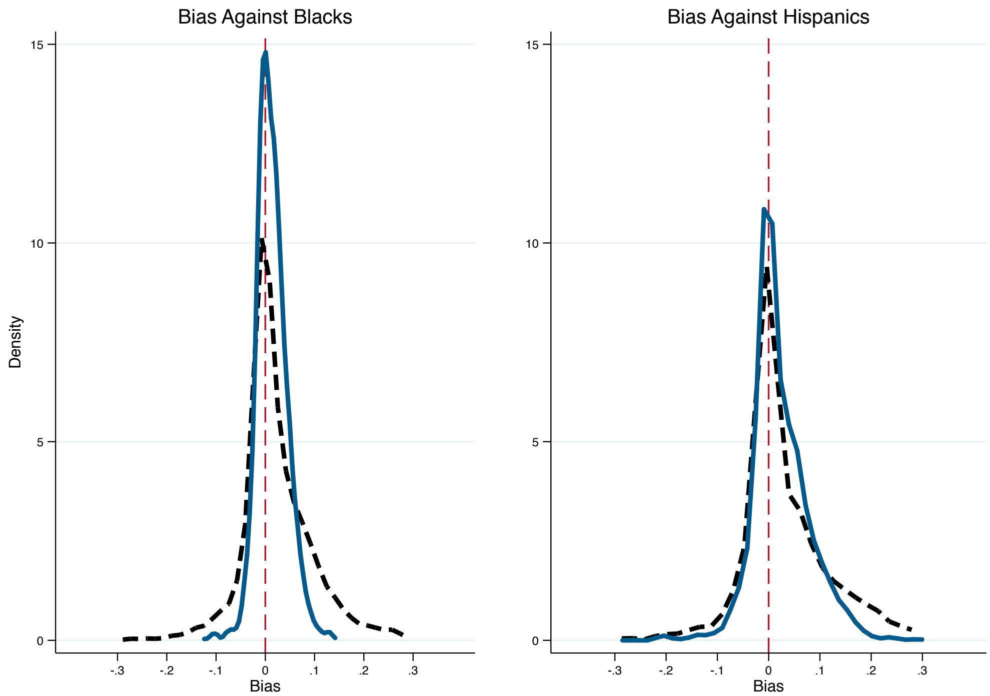

Notes: Dashed line plots kernel density estimate of the distribution of officer bias estimates. Solid line plots corresponding kernel density estimate of distribution of Bayes shrunk officer bias.

Figure 5: Estimated Distributions of Bias

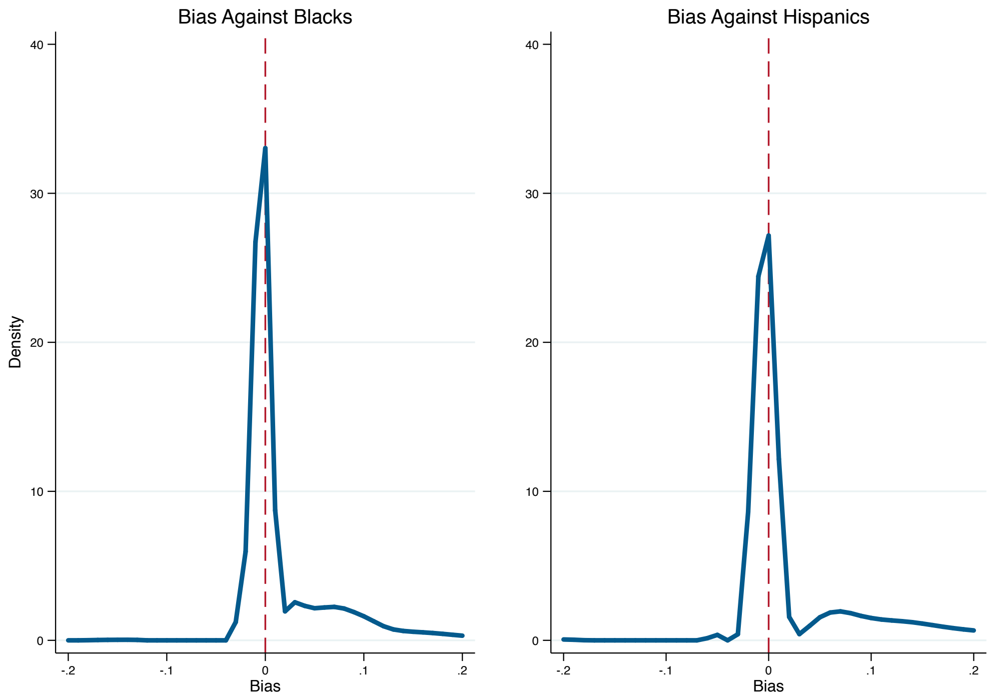

Notes: Figure plots estimated distributions of bias computed using the deconvolution technique from Delaigle and Meister (2008).

Figure 6: Estimated CDFs of Bias

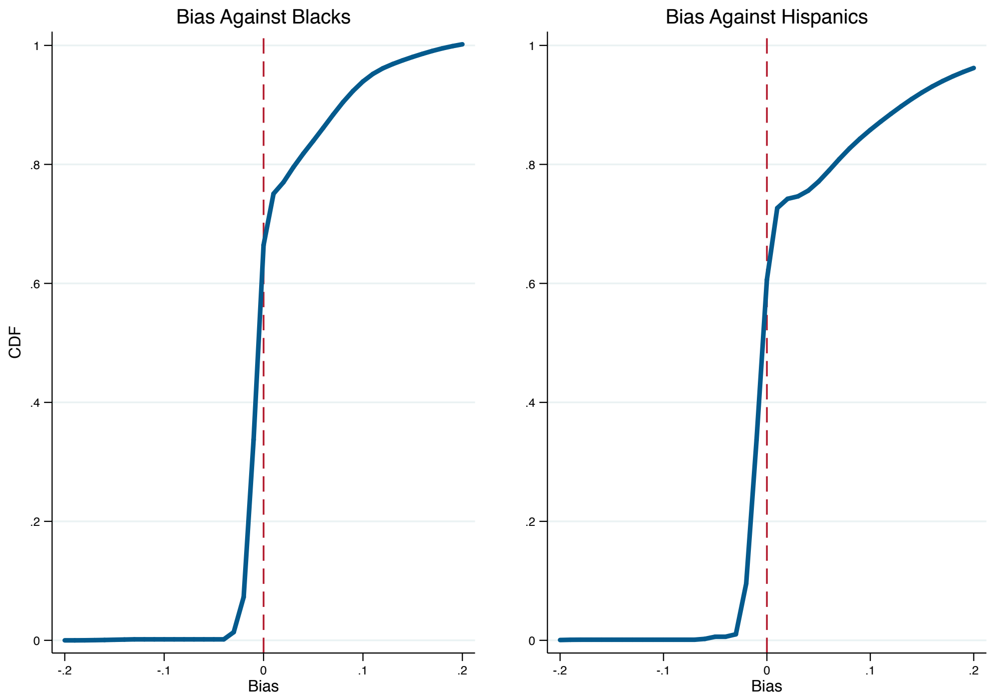

Notes: Figure plots estimated cumulative density functions of bias computed using the deconvolution technique from Delaigle and Meister (2008).

Figure 7: Distributions of Bias by Officer Race

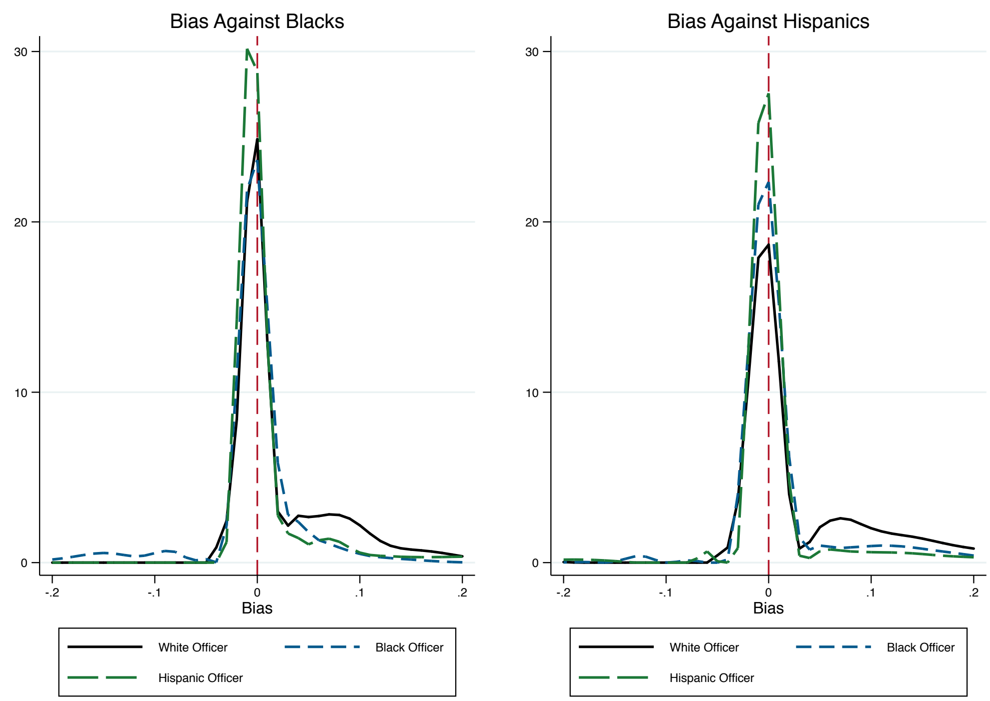

Notes: Figure plots estimated distributions of bias computed using the deconvolution technique from Delaigle and Meister (2008) and computed separately by officer race.

Figure 8: Driver Race Shares by Officer Bias

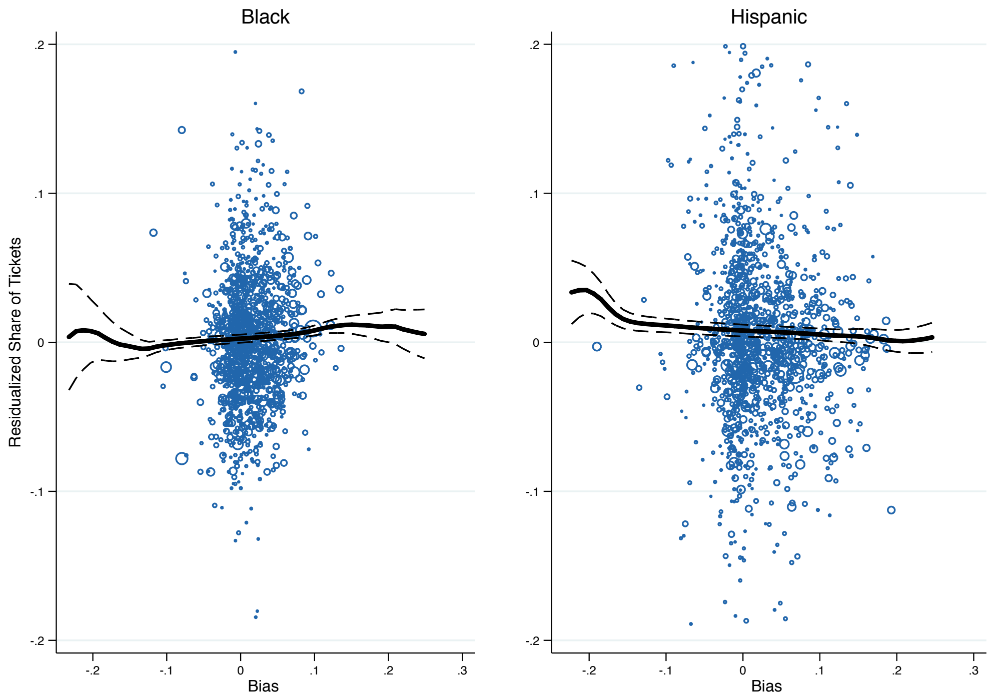

Notes: Figure plots an officer's residualized share of stops which are black (Hispanic) against his estimate bias against that racial group. Residualized shares are computed by regressing a black (Hispanic) indicator on year, month, day-of-week, speed zone, and county fixed effects and averaging the residuals at the officer level.

Table 1: Summary Statistics

	White	Black	Hispanic	Total
Female	0.362	0.397	0.301	0.354
	(0.481)	(0.489)	(0.459)	(0.478)
Age	37.39	34.15	34.20	36.03
	(14.89)	(12.10)	(11.93)	(13.83)
Florida License	0.818	0.853	0.893	0.842
	(0.386)	(0.355)	(0.309)	(0.365)
Zip Code Income	56.96	39.70	46.49	51.13
	(53.20)	(31.18)	(42.81)	(47.79)
Citations in Past Year	0.0391	0.0482	0.0459	0.0424
	(0.209)	(0.231)	(0.232)	(0.219)
MPH Over	15.49	16.67	18.32	16.38
	(6.518)	(7.046)	(6.972)	(6.829)
Discount	0.352	0.316	0.206	0.311
	(0.478)	(0.465)	(0.405)	(0.463)
Fine Amount	182.3	190.2	200.7	188.0
	(76.40)	(80.50)	(79.33)	(78.23)
White	-	-	-	0.577
				(0.494)
Black	-	-	-	0.186
				(0.389)
Hispanic	-	-	-	0.237
				(0.425)

Notes: Number of observations is 571,751 (White); 184,567 (Black); 234,323 (Hispanic); 990,641 (Total). Standard deviations in parentheses. Zip code income is missing for 42% of White stops, 40% of Black stops, 37% of Hispanic stops. To account for the fact that a large share of fine amounts are missing or zero in our data, we impute the fine amount with the modal non-zero fine for each county $\times$ speed over the limit cell.

Table 2: Characteristics of Cited Drivers Relative to Other Data Sources

	Citations	ACS - Any	ACS - Drivers	Crash - Any	Crash - Injury
Female	0.356	0.515	0.474	0.424	0.441
	(0.479)	(0.500)	(0.499)	(0.494)	(0.497)
Age	34.90	47.46	41.70	39.65	39.77
	(13.45)	(19.39)	(13.72)	(16.78)	(17.11)
White	0.578	0.625	0.606	0.556	0.576
	(0.494)	(0.484)	(0.489)	(0.497)	(0.494)
Black	0.181	0.138	0.136	0.189	0.193
	(0.385)	(0.345)	(0.343)	(0.391)	(0.394)
Hispanic	0.241	0.200	0.217	0.233	0.211
	(0.428)	(0.400)	(0.412)	(0.423)	(0.408)

Notes: Standard deviations in parentheses. ACS data include individuals aged 16 or older and sampling weights are used.

Table 3: White-Minority Speed Charged Gap

	(1) FL	(2) FL	(3) FL
Black	1.180***	1.055***	0.721***
	(0.0185)	(0.0185)	(0.0167)
Hispanic	2.857***	2.601***	0.694***
	(0.0169)	(0.0170)	(0.0162)
White Mean	15.49	15.49	15.49
Controls	No	Yes	Yes
FE	No	No	Yes
Obs	988096	988096	988096

Notes: Dependent variable is the ticketed speed minus the speed limit. Robust standard errors in parentheses. Dependent variable is speed charged. Age divided by 1000. Controls include gender, age, age squared, and whether driver is in-state. Fixed effects include year, month, day of week, speed zone, and county.

Table 4: Linear Probability Estimates

	(1) Discount	(2) Discount	(3) Discount	(4) Discount	(5) Discount
Black	-0.0376***	-0.0331***	-0.0228***	-0.0243***	-0.0242***
	(0.00127)	(0.00127)	(0.00113)	(0.00112)	(0.00116)
Hispanic	-0.149***	-0.137***	-0.0326***	-0.0335***	-0.0333***
	(0.00106)	(0.00107)	(0.00103)	(0.00102)	(0.00106)
Female		0.0547***	0.0354***	0.0314***	0.0302***
		(0.000992)	(0.000875)	(0.000869)	(0.000896)
Florida License		-0.0625***	-0.000213	-0.000155	0.000592
		(0.00133)	(0.00125)	(0.00124)	(0.00127)
Age		1.453***	2.917***	2.748***	2.644***
		(0.179)	(0.158)	(0.157)	(0.162)
Age Squared		-1.687	-16.05***	-17.11***	-16.79***
		(2.149)	(1.890)	(1.878)	(1.935)
White Mean	.36	.36	.36	.36	.36
Zone FE	No	No	Yes	Yes	Yes
County FE	No	No	Yes	Yes	Yes
Year + Month FE	No	No	No	Yes	Yes
DOW FE	No	No	No	Yes	Yes
Hour FE	No	No	No	No	Yes
Observations	976821	976821	976821	976821	926622

Notes: Robust standard errors in parentheses. Dependent variable is an indicator for charged speed equals 9 MPH above the limit. Age divided by 1000.

Table 5: Linear Probability Estimates with Additional Controls

	(1) Discount	(2) Discount	(3) Discount	(4) Discount	(5) Discount	(6) Discount
Black	-0.0368***	-0.0374***	-0.0188***	-0.0213***	-0.0251***	-0.0249***
	(0.00197)	(0.00198)	(0.00146)	(0.00152)	(0.00124)	(0.00124)
Hispanic	-0.0482***	-0.0490***	-0.0276***	-0.0293***	-0.0311***	-0.0311***
	(0.00189)	(0.00189)	(0.00135)	(0.00137)	(0.00114)	(0.00114)
Log Zip Code Income				-0.00749***		
				(0.00117)		
Tickets Past Three Years						-0.0199***
						(0.00233)
Tickets Squared						0.000141
						(0.00115)
White Mean	.38	.38	.36	.36	.36	.36
Vehicle Chars	No	Yes	No	No	No	No
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Sample	Vehicle	Vehicle	ZIP Code	ZIP Code	History	History
Observations	297963	297963	581709	581709	798668	798668

Notes: Robust standard errors in parentheses. Dependent variable is an indicator for charged speed equals 9 MPH above the limit. Odd columns repeat the specification from Column (4) of Table 4 but restricting the sample to stops where the relevant variable is non-missing, and even columns add those controls. All regressions include controls, year and month fixed effects, day-of-week fixed effects county fixed effects, and speed zone fixed effects

Table 6: Characteristics of Distribution of Estimated Bias

	Bias Against Blacks	Bias Against Hispanics
Mean	.013	.022
Standard Deviation	.031	.058
10th Percentile	-.02	-.03
25th Percentile	-.007	-.009
50th Percentile	.01	.011
75th Percentile	.02	.051
90th Percentile	.053	.096
Officers <0	101	160
Officers >0	161	236
Officers	1337	1339

Notes: This table presents the distributional statistics of shrunken Officer-level estimated racial bias (graphed in Figure 4).

Table 7: External Benchmark Estimates

	Discount
Black Driver	-0.0297***
	(0.00123)
Hispanic Driver	-0.0369***
	(0.00117)
Black Officer	-0.00415***
	(0.00129)
Hispanic Officer	-0.0123***
	(0.00136)
Black Driver x Black Officer	0.0304***
	(0.00272)
Hispanic Driver x Hispanic Officer	0.0163***
	(0.00199)
White Mean	.36
Controls	Yes
Fixed Effects	Yes
Observations	976821

Notes: Robust standard errors in parentheses. Dependent variable is an indicator for charged speed equals 9 MPH above the limit.

Table 8: Predicting Officer Bias

	(1) Black	(2) Black	(3) Black	(4) Hispanic	(5) Hispanic	(6) Hispanic
Black	-0.0213***	-0.128***	-0.0650**	-0.0202***	-0.0973**	-0.0586**
	(0.00359)	(0.0426)	(0.0261)	(0.00520)	(0.0403)	(0.0298)
Hispanic	-0.00959***	-0.122***	-0.0720***	-0.0231***	-0.0862**	-0.0925***
	(0.00315)	(0.0393)	(0.0248)	(0.00517)	(0.0378)	(0.0256)
Other	0.0135	0.161	0.0546	0.000741	0.129	-0.109
	(0.0111)	(0.100)	(0.0990)	(0.0127)	(0.101)	(0.0713)
Female	-0.0154***	-0.139***	-0.0680**	-0.0150***	-0.130***	-0.0609*
	(0.00375)	(0.0522)	(0.0279)	(0.00543)	(0.0496)	(0.0326)
Age	0.0155*	0.298***	0.138**	0.0156	0.189**	0.0668
	(0.00854)	(0.0912)	(0.0586)	(0.0126)	(0.0879)	(0.0688)
Age Squared	-0.00263**	-0.0470***	-0.0198**	-0.00246	-0.0269**	-0.00421
	(0.00129)	(0.0136)	(0.00883)	(0.00198)	(0.0133)	(0.0107)
Experience	-0.00623	0.0155	-0.0690	-0.0118	-0.0535	-0.130***
	(0.00584)	(0.0626)	(0.0452)	(0.0102)	(0.0562)	(0.0488)
Exp Squared	0.00406*	0.0183	0.0306*	0.00864**	0.0447**	0.0448**
	(0.00243)	(0.0241)	(0.0177)	(0.00406)	(0.0211)	(0.0198)
Failed Entrance Exam	-0.00558	-0.0881	-0.0464	-0.000403	-0.103**	-0.0351
	(0.00519)	(0.0537)	(0.0314)	(0.00651)	(0.0521)	(0.0355)
Any College	-0.00443	-0.0876***	-0.0422*	0.000194	0.0112	0.0166
	(0.00297)	(0.0330)	(0.0220)	(0.00450)	(0.0317)	(0.0249)
Any Complaints	-0.00267	-0.127***	0.0337	0.00117	-0.0353	0.0142
	(0.00428)	(0.0451)	(0.0340)	(0.00545)	(0.0452)	(0.0346)
Sought Promotion	-0.00143	-0.0138	0.0162	-0.00319	0.0133	0.0102
	(0.00282)	(0.0315)	(0.0216)	(0.00459)	(0.0303)	(0.0234)
Dep Var	Bias	Bias<0	Bias<0 (Sig)	Bias	Bias<0	Bias<0 (Sig)
Mean	.023	.582	.124	.029	.59	.178
Observations	1332	1332	1332	1334	1334	1334

Notes: Robust standard errors in parentheses. In Columns (1) and (4), dependent variable is our measure of the officer's bias. In Columns (2) and (5), it is an indicator for whether our bias estimate is positive. In Columns (3) and (6), it is an indicator for whether the estimate is positive and statistically significant. Columns (1)-(3) refer to bias against blacks while Columns (4)-(6) refer to bias against Hispanics.

Table 9: Predicting Officer Bias, Sergeant's Exam Takers

	(1) Black	(2) Black	(3) Black	(4) Hispanic	(5) Hispanic	(6) Hispanic
Black	-0.0269***	-0.176**	-0.147***	-0.0162*	-0.0537	-0.0945*
	(0.00606)	(0.0699)	(0.0342)	(0.00928)	(0.0693)	(0.0489)
Hispanic	-0.00978**	-0.0407	-0.114***	-0.0230***	-0.0539	-0.149***
	(0.00486)	(0.0620)	(0.0407)	(0.00747)	(0.0602)	(0.0409)
Other	0.0269	0.291*	0.196	-0.00478	-0.0132	-0.267***
	(0.0297)	(0.151)	(0.207)	(0.0226)	(0.202)	(0.0416)
Age	0.0128	0.216	0.111	0.0194	0.0569	-0.0747
	(0.0176)	(0.190)	(0.131)	(0.0259)	(0.189)	(0.144)
Age Squared	-0.00285	-0.0365	-0.0217	-0.00218	-0.0125	0.0160
	(0.00252)	(0.0267)	(0.0188)	(0.00360)	(0.0268)	(0.0209)
Experience	-0.00356	0.173	-0.114	-0.00357	-0.0183	-0.137
	(0.0117)	(0.114)	(0.0843)	(0.0178)	(0.112)	(0.0843)
Exp Squared	0.00521	-0.0324	0.0612*	0.00157	0.0267	0.0326
	(0.00526)	(0.0474)	(0.0360)	(0.00806)	(0.0462)	(0.0369)
Failed Entrance Exam	-0.00526	-0.0176	-0.115***	0.00458	-0.0391	-0.0184
	(0.00923)	(0.0874)	(0.0378)	(0.00888)	(0.0839)	(0.0598)
Any College	-0.00697	-0.0914*	-0.0149	-0.0140*	-0.0179	-0.0318
	(0.00471)	(0.0493)	(0.0335)	(0.00716)	(0.0478)	(0.0370)
Any Complaints	-0.00703	-0.204***	0.00417	0.00161	-0.0691	0.0526
	(0.00754)	(0.0717)	(0.0475)	(0.00828)	(0.0720)	(0.0567)
Sgt Exam Score	-0.0000316	-0.00102	-0.00325**	0.000498	0.00376	0.00275
	(0.000223)	(0.00252)	(0.00159)	(0.000321)	(0.00246)	(0.00185)
Dep Var	Bias	Bias<0	Bias<0 (Sig)	Bias	Bias<0	Bias<0 (Sig)
Mean	.023	.582	.124	.029	.59	.178
Observations	513	513	513	514	514	514

Notes: Same as Table 8 except that only the subset of officers ever to take the Sergeant's Promotional exam are included.

Table 10: Model Parameter Estimates

	White			Minority			
	μ	σ^2	# Param	μ	σ^2	# Param	Mean Diff
b	.0228*** (9.38 X 10^{-5})	—	1	—	—	—	—
t	-2.41*** (.046)	20.12*** (0.781)	1327	-2.52*** (.057)	19.50*** (0.757)	1327	0.10*** (.013)
λ	18.387*** (0.002)	2.905*** (0.506)	67	19.107*** (0.003)	2.186*** (0.381)	67	-0.720*** (0.004)
Pr(Discount)	0.357*** (6.90 X 10^{-4})	0.125*** (0.005)	1327	0.325*** (7.98 X 10^{-4})	0.112*** (0.004)	1327	0.033*** (0.001)

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table presents estimates of the model introduced in section 7. *b* is the slope parameter for how officers weight the speed of drivers in choosing to discount, t is each officer's mean valuation of a racial group in choosing to discount, and λ is the poisson speed parameter for each race by county. Pr(Discount) = $\Phi(t - 10b)$, i.e. the probability of being discounted when driving right above the bunch point. Note that the discount probabilities are not technically parameters but rather are calculated using the estimated t's and b. The variances are empirical variances of the estimates, not adjusted for sampling error.

Table 11: Racial Gap Decomposition

	Baseline	No Bias	No Sorting	Neither
Probability Gap	0.0968	0.0653	0.0341	0
Speed Gap	-2.114	-1.783	-1.475	-1.1135

Notes: This table presents how the racial gap in discounting and speeds change without bias and sorting of officers across counties. The probability gap is the probability of being discounted if you are at the speed right above the jump in fine. Both gaps are the white drivers' outcome minus minority drivers' outcome. No bias is calculated by assigning each officer's preferences towards minorities to be the same as his preference to whites. No sorting is calculated by simulating a new draw of officers for each driver, where the draw is done at the state level.

Figure 9: Model Discount Probability Across Officer Lenience

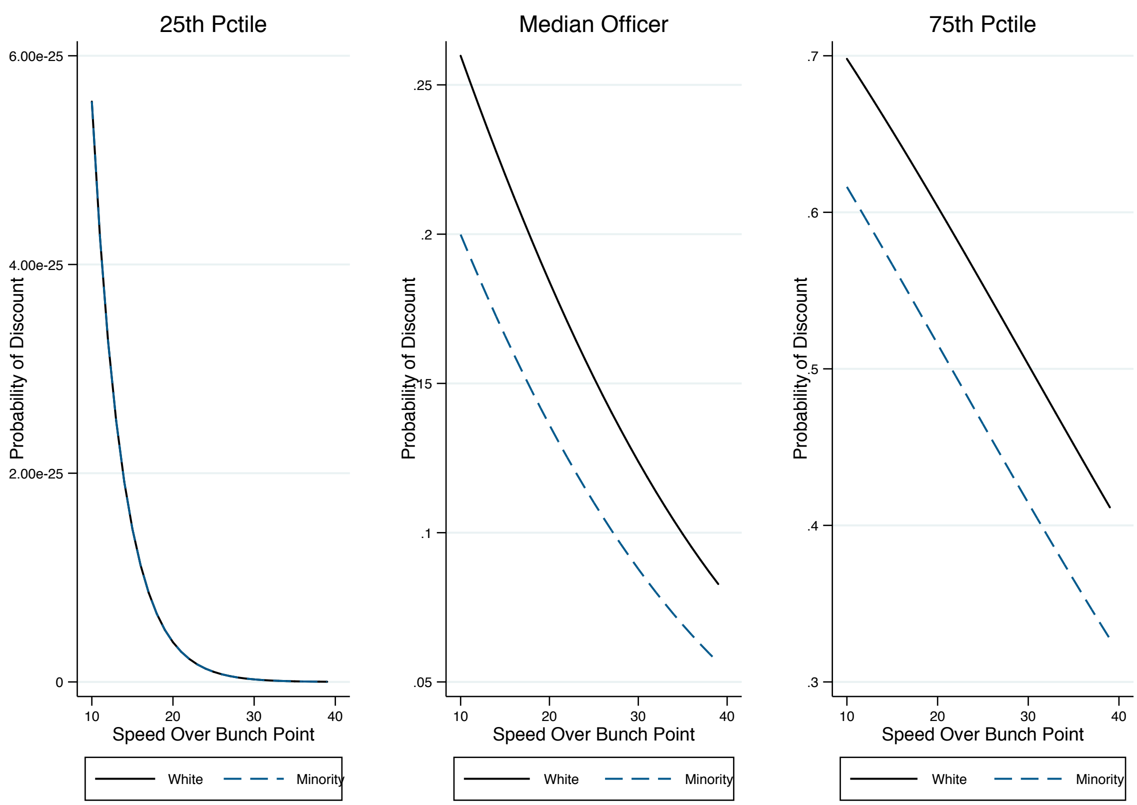

Notes: Each figure plots the probability of being discounted to 9MPH over for each speed above 9, separately by race. The left panel presents the discount probabilities for officers at the 25th percentile of *lenience*, the second panel for the median lenience, and the third panel for the 75th percentile of lenience. The most important fact to note is that bias only appears for officers that have some degree of lenience.

Figure 10: Discount Probability Across Officer Distribution

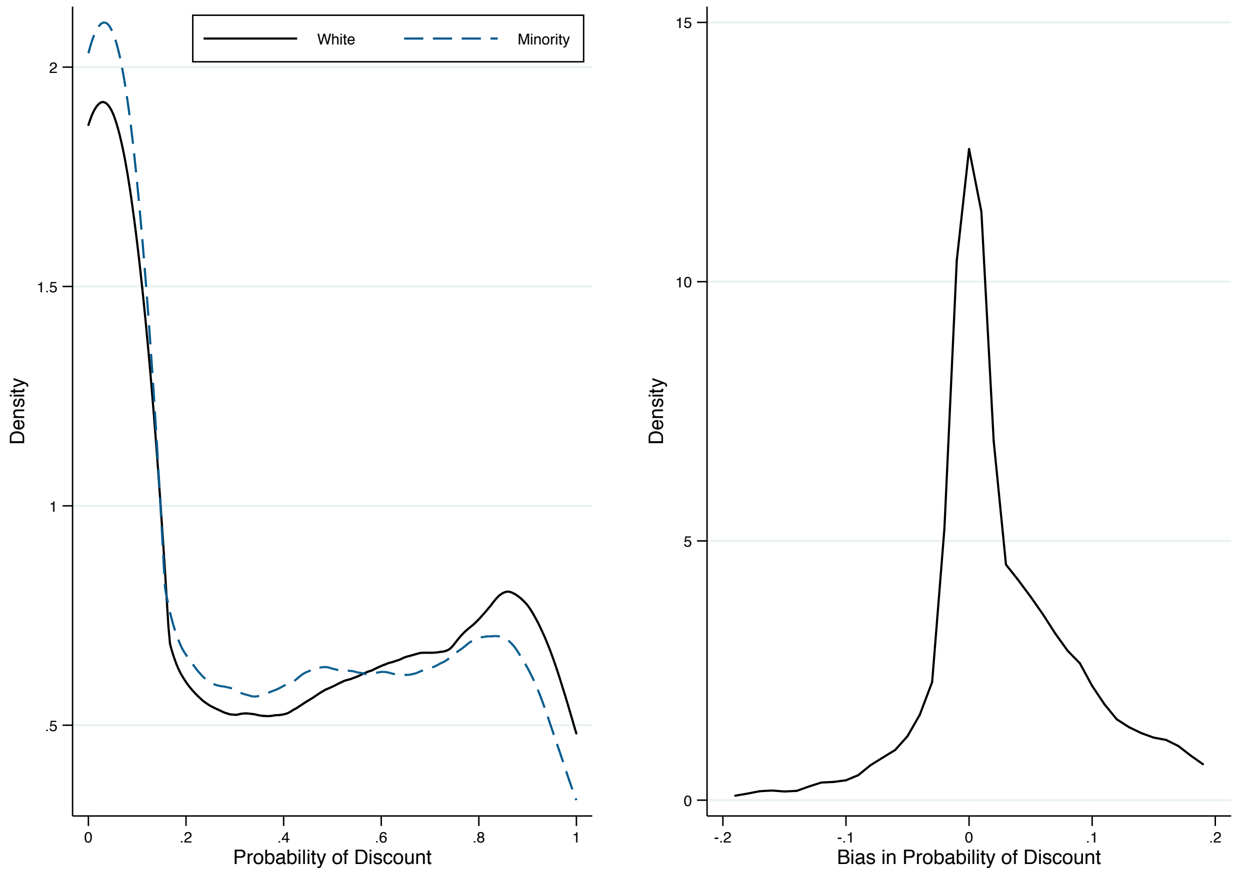

Notes: The left panel plots the kernel density estimate of the distribution of discount probabilities across officers, calculated separately for whites and minorities. The discount probability is calculated as Pr(Discount) = $\Phi(t - 10b)$, i.e. the probability of being discounted when driving right above the bunch point. The right panel plots the kernel density estimate of each officer's difference in discount probability between whites and minorities. The further right, the more biased an officer is towards whites.

Figure 11: Speeds By Race Across Counties

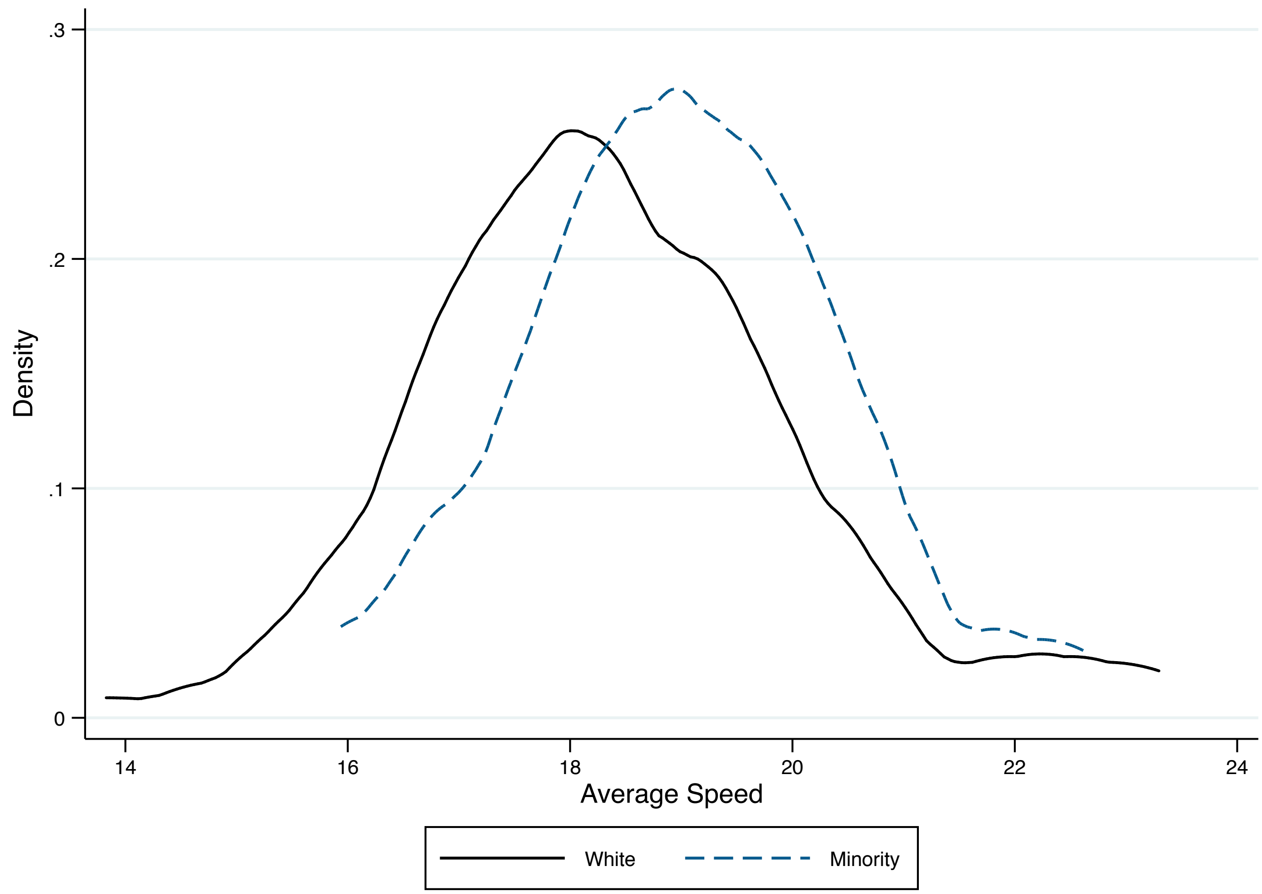

Notes: This figure plots the kernel density estimate of average speeds across counties, calculated separately by race. The average speeds correspond to the λ calculated in the model.

Figure 12: Speeds By Race Across Counties

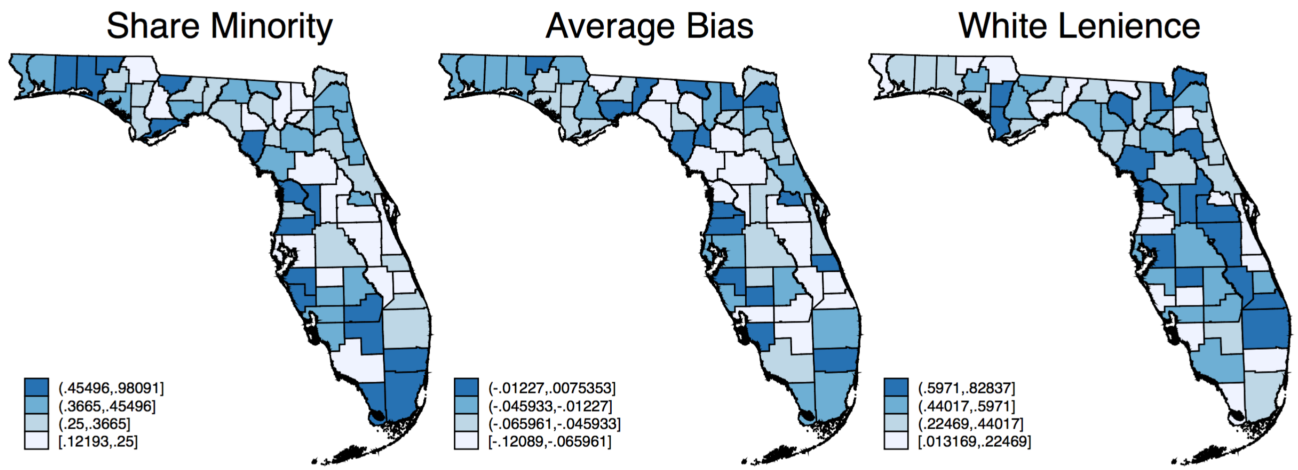

Notes: This figure plots the degree of minorities, bias, and lenience across counties of Florida. The darker the shade of blue, the more minorities, bias, and white lenience are present in the county, respectively across panels. The most important fact from these graphs is that white lenience is greatest in counties with few minorities.

Figure 13: Removing Most Biased Officers, Explanation

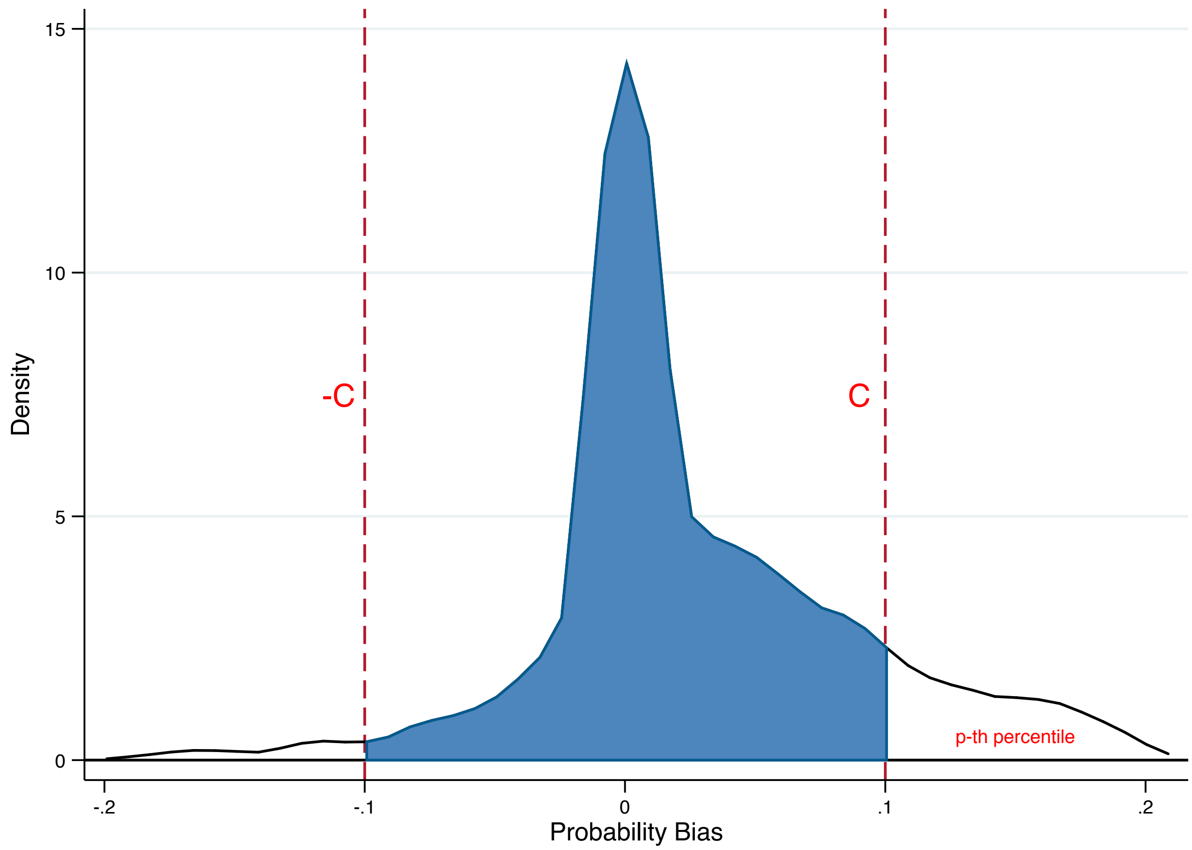

Notes: This plot describes how the first counterfactual is conducted. For a p-th percentile chosen, officers are removed for bias greater than C, the level of bias at the p-th percentile, and -C at the opposing end of bias. Individuals are all re-assigned officers from within their county who are not 'fired,' with the probability of encountering a certain officer proportional to the number of tickets he writes in that county.

Figure 14: Removing Most Biased Officers

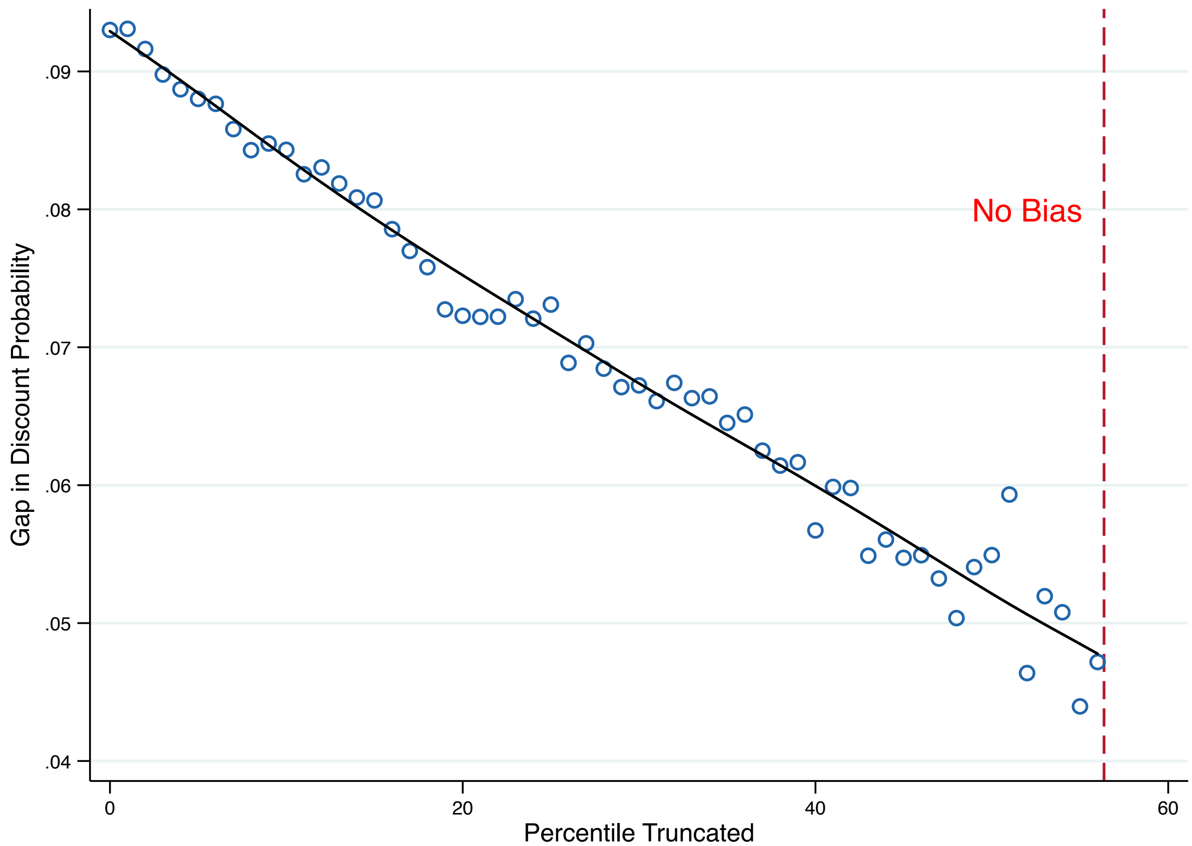

Notes: Plot of the counterfactual of removing the most biased officers. The x-axis is the p-th percentile being truncated, and the y-axis is the average gap in probability of discount across whites and minorities. The extreme of no bias is at the 57th percentile rather than 50 because the median officer is slightly biased against minorities.

Figure 15: Removing Least Lenient Officers

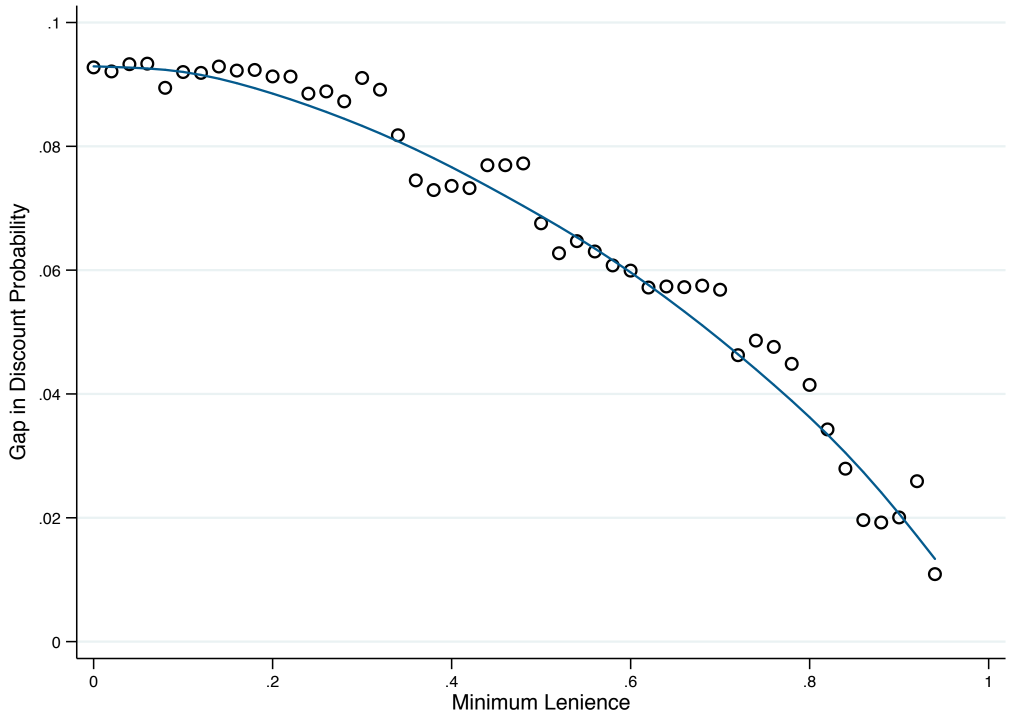

Figure 16: Removing Most Lenient Officers

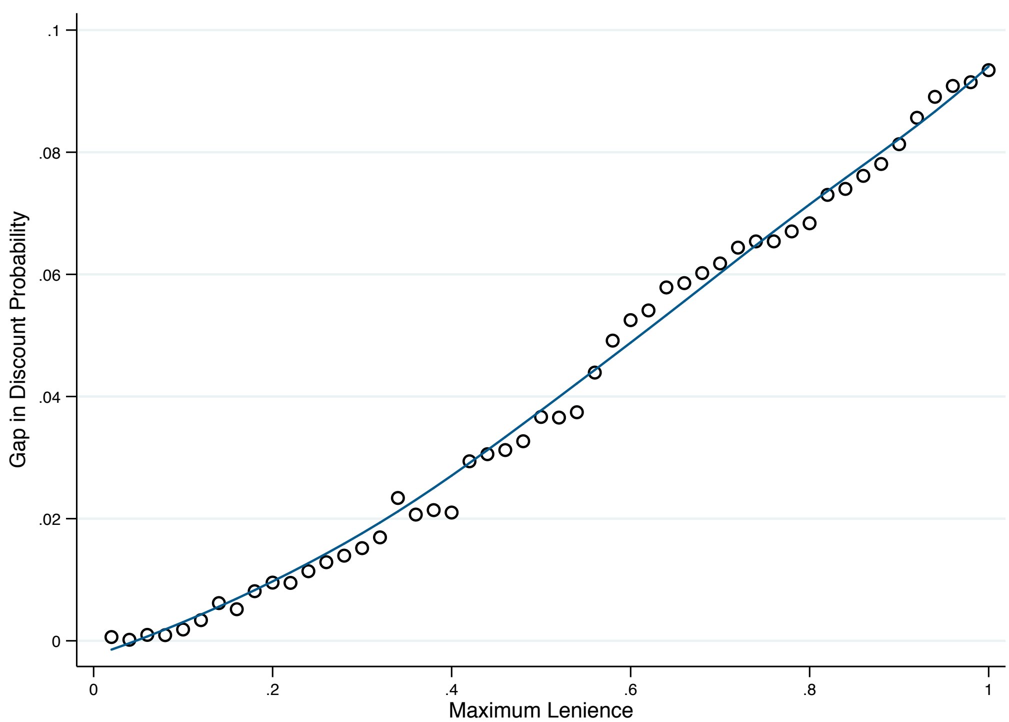

Notes: Plot of the counterfactuals of removing the most harsh and most lenient officers. The x-axis of the first figure is the minimum lenience of the officers allowed in the sample, below which officers are fired, and the y-axis is the gap in the probability of discount between whites and minorities. The x-axis of the second figure is the maximum lenience of the officers allowed in the sample. Lenience is defined as the probability of discounting drivers going at 10MPH over. As in the first counterfactual, individuals receive a new draw of officers within their county.

Figure 17: Changing Share of Minority and Female Officers

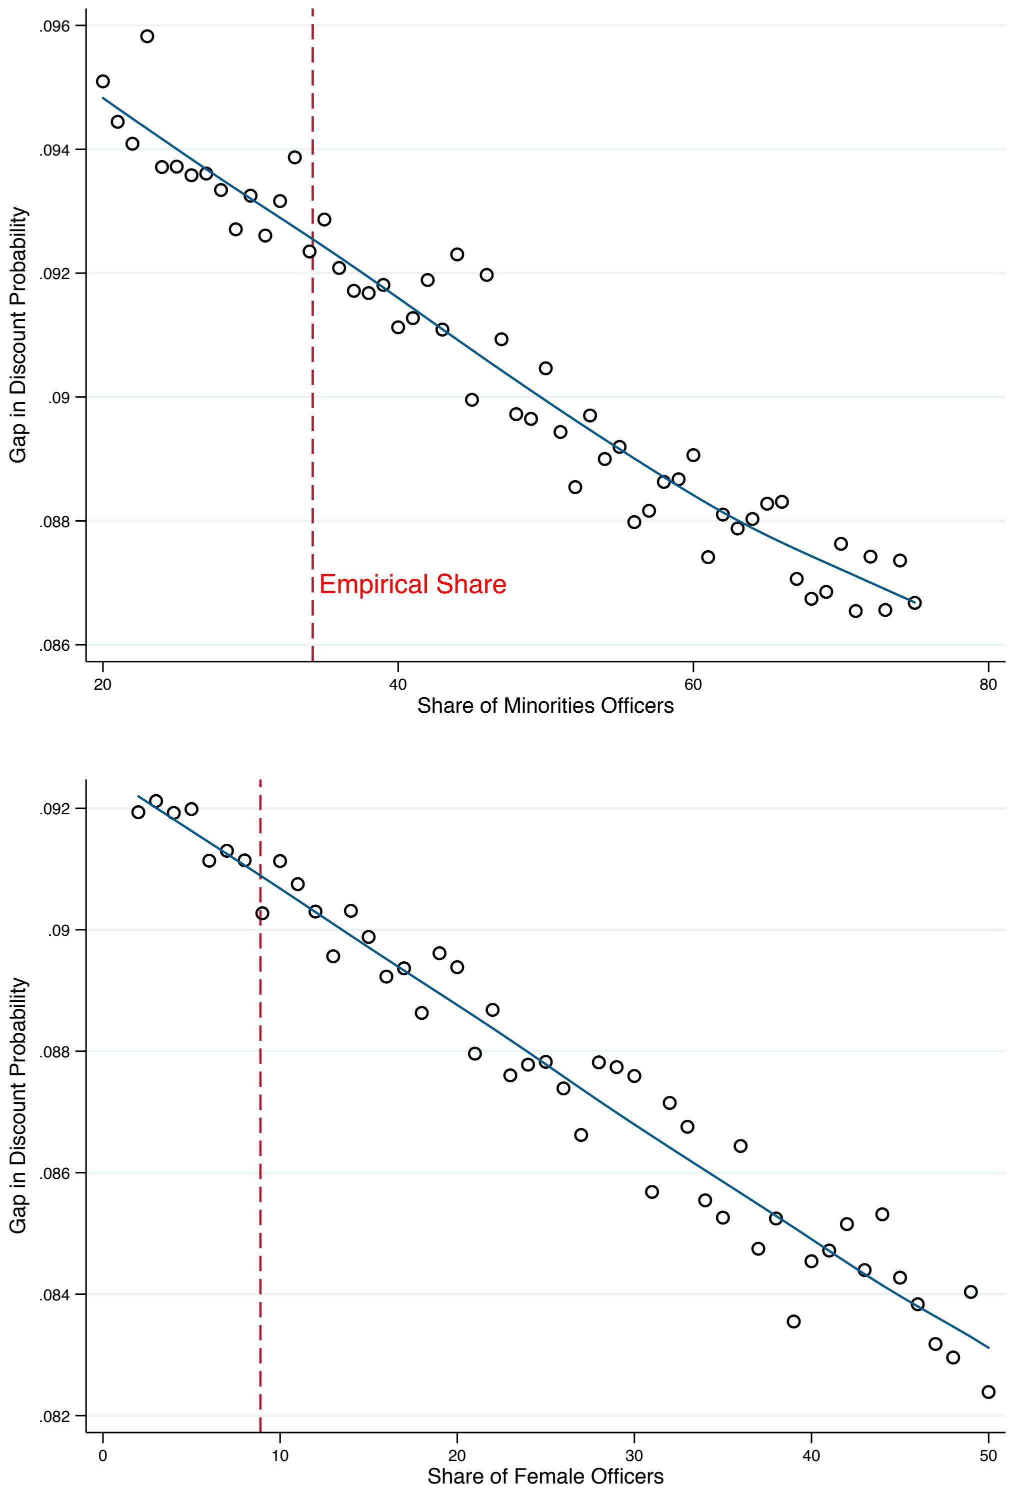

Notes: Plot of the counterfactuals of increasing the share of minority and female drivers in the sample. The x-axis of the first and second figures are the share minority and share female, respectively, and the y-axis is the gap in probability of discount between whites and minorities.

Table 12: Assigning Officers Based on Bias and Lenience

	Baseline	Assigning on Bias	Assigning on Lenience
Probability Gap	0.097	0.102	.0047
Speed Gap	-2.114	-2.116	-1.210

Notes: This table presents how the racial gap in discounting and speeds change when officers are assigned to minimize minorities' exposure to bias and officer harshness. The first column is the baseline simulation of the model to match the true data. The second column sorts officers within a troop (which comprise 6-10 counties) such that the least biased officers are in counties with the most minorities. The third column sorts officers within a troop such that the most lenient officers are in counties with the most minorities.

Appendix Figures and Tables

Figure A-1: Model Goodness of Fit

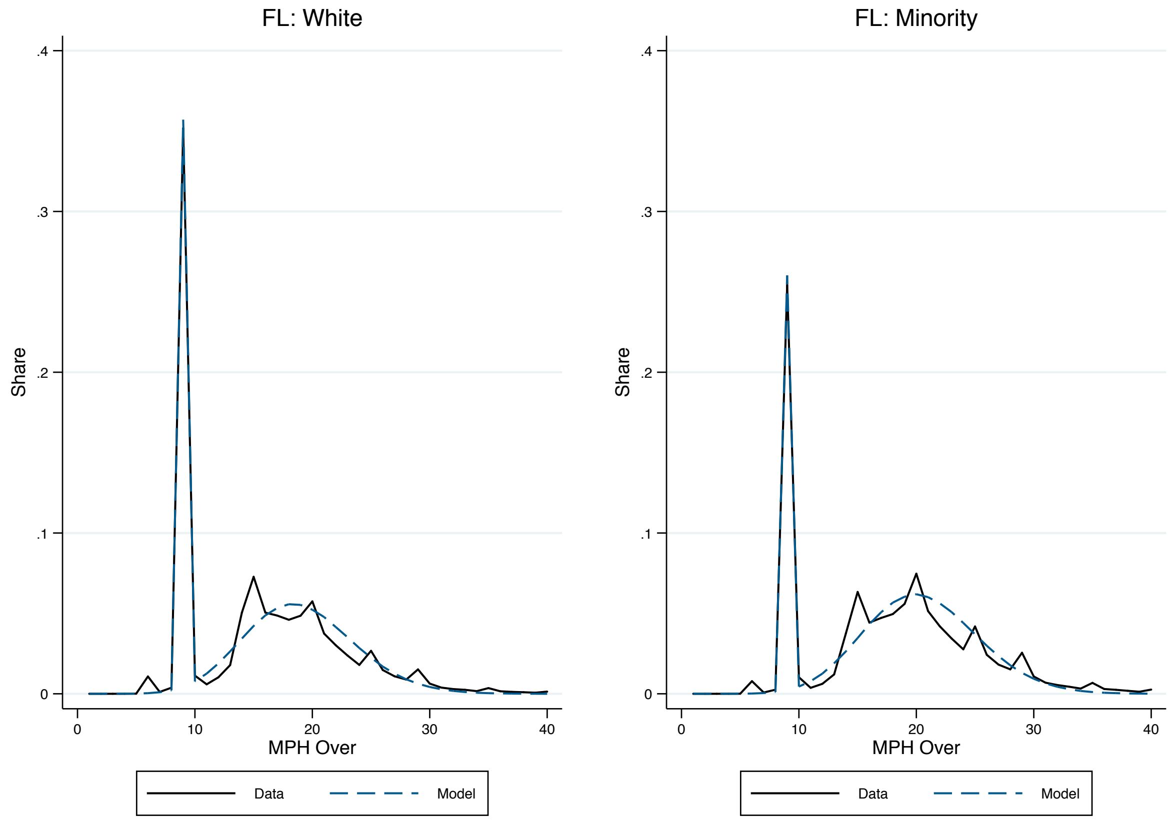

Notes: Figure plots observed versus model-estimated speed distributions for white and nonwhite drivers.

Figure A-2: Model Goodness of Fit

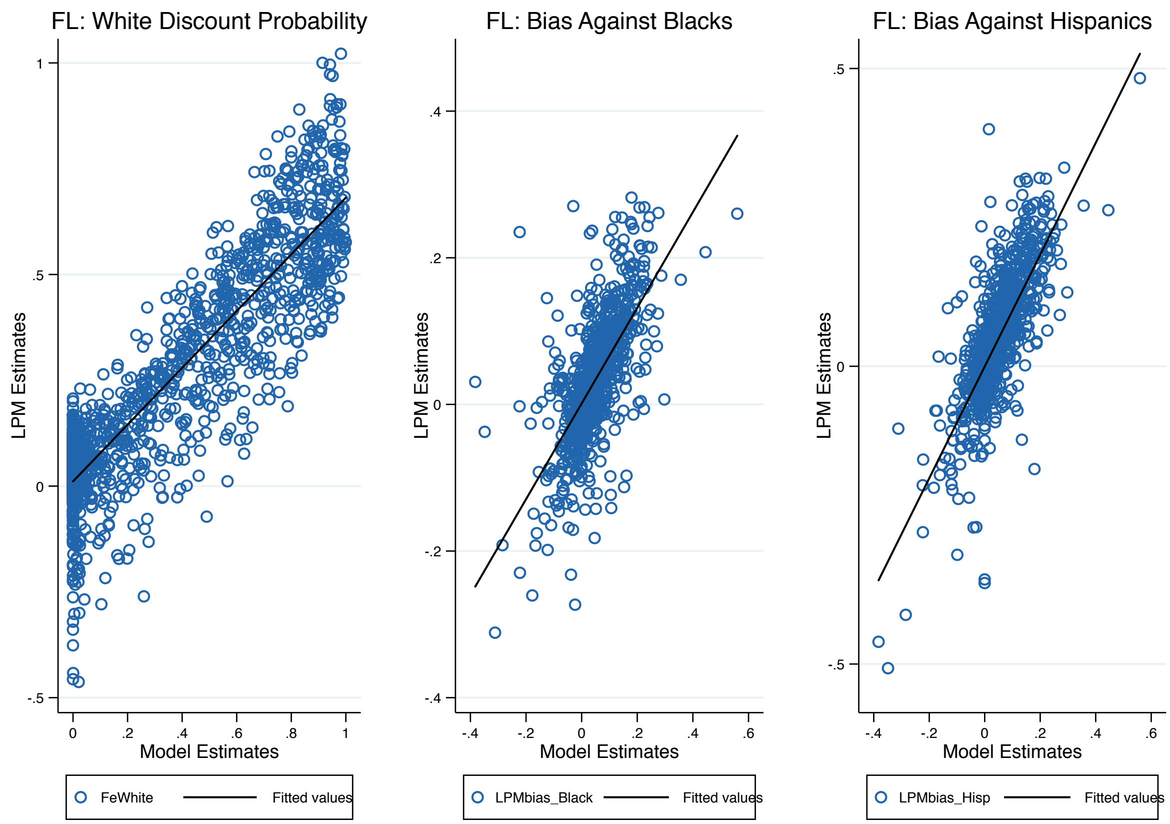

Notes: Figure plots the officer-level relationship between lenience towards whites and bias against nonwhites estimated via the simple linear probability model approach (y-axis) and the structural model approach (x-axis).