8-4-2008

Discretionary Behavior and Racial Bias in Issuing Tra±c Tickets: Theory and Evidence

Nejat Anbarci
Department of Economics, Florida International University

Jungmin Lee
Department of Economics, Florida International University

Discretionary Behavior and Racial Bias in Issuing Traffic Tickets: Theory and Evidence

Nejat Anbarci and Jungmin Lee*

May 27, 2008

Recently, police departments, legislators, media, and the public at large in the U.S. have increasingly been concerned about racial disparities in officers' issuing traffic tickets. Ascertaining the extent to which an observed disparity reflects racial bias is the crucial issue. First, we use a theoretical model which borrows features from the recent literature regarding racial bias in vehicle searches. In our model, motorists, picking the speed to travel at, take into account the probability of getting ticketed and the speed that the officer will cite, while officers maximize a benefit function generically increasing in the speed of ticketed drivers; this benefit function, however, is general enough to allow officers to give certain drivers a break by citing them at a lower speed than they were traveling. Empirically, we exploit the existence of a massive accumulation of speeding tickets at 10 m.p.h. over the speed limit to elicit officers' discretionary behavior and leniency. Surprisingly, about 30% of all ticketed drivers were cited for driving exactly at this particular speed. Using our novel measure of officers' leniency, we find that especially white and male officers are heavily engaged in discretionary behavior. We also find officers' discretion is racially biased; minority officers are less lenient to minority drivers. This is interesting in comparison with Antonovics and Knight (forthcoming) who, using the same data set, found evidence on own-race preferences in vehicle searches.

JEL classification numbers: J70, K42

Keywords: Discretionary behavior, strict behavior, leniency, racial bias, drivers' speeding decision, officers' ticketing and citation decision.

1 Introduction

Since the early 1990s, the national debate over racial profiling has mostly focused on countless accounts of unjustified searches, videotaped beatings, and so on. The national attention as well as numerous lawsuits brought about against police departments nationwide - alleging racially prejudiced and at times harsh law enforcement practices - have grown like an avalanche since. The

*We are very grateful of Bill Dedman, Brian Knight, and Nicola Persico for kindly providing the data. We would like to thank Kate Antonovics, Mustafa Caglayan, Hyunbae Chun, Bill Dedman, Monica Escaleras, Laura Giuliano, Gulcin Gumus, Daniel Hamermesh, Cem Karayalcin, Jinyoung Kim, Patrick Mason, Oscar Mitnik, Nicola Persico, Phil Robins, Peter Thompson, Mehmet Ulubasoglu, and participants in seminars at several universities for helpful comments.
Nejat Anbarci: Department of Economics, Florida International University, University Park, Miami, FL 33199, USA; email: anbarcin@fiu.edu; phone: +1.305.348.2735; fax: +1.305.348.1524.
Jungmin Lee: Department of Economics, Florida International University, University Park, Miami, FL 33199, USA; email: leej@fiu.edu; phone: +1.305.348.6639; fax: +1.305.348.1524.

resulting public outcry has later been recognized by politicians at the highest levels as well.[1]

In July 2003, Pulitzer Prize winning journalist Bill Dedman and his co-author, Francie Latour, reported the initial results of an analysis of 166,000 tickets and warnings from every police department in Massachusetts over a two-month period, April and May of 2001. Comparing speeding tickets and warnings, they drew serious conclusions such as minorities, men and young drivers were least likely to receive a warning, and thus less likely to get away with speeding. In this paper, using the same data set, we will take a closer look at these issues focusing on racial bias in issuing speeding tickets.

It is small wonder that Dedman and Latour's articles have attracted so much attention in the short period of time since their initial appearance in the Boston Globe. First, the economic impact of a traffic ticket is considerable. Speeding tickets in Massachusetts start at \$75 including \$25 surcharge for the Head Injury Trust Fund; for the first ten miles above the speed limit, the fine is \$75, and then it rises by ten dollars for each additional mile. But that typically is the least of the problems. The Boston Globe estimates that a typical Massachusetts driver will pay \$350 in higher insurance premiums for a single ticket and \$1400 for two tickets, over the six years a ticket stays on the driving record. Second, while its monetary costs are explicit and recognizable by everyone, disparate treatment of different races, ethnicities and genders has other serious implications: Individuals subjected to disparate traffic enforcement may in time experience a loss of respect and trust in their local police force.

Our main contribution in this paper is that we investigate police discretion especially regarding reported speeds on issued citations. Police discretion goes beyond deciding whether to ticket a speeding driver or just let drivers go with or without oral warnings. It also includes cases where the officers give them a "break" - and a smaller fine - by citing them at a lower speed than they were traveling. One of this paper's most important contributions will be identifying a measure of police discretion in terms of officers' underreporting especially at a particular lower speed very often. Indeed Figure 1, the histogram of reported speeds in our sample, shows that reported speeds are likely to be different from drivers' actual speeds.[2] There are outstanding spikes in some specific levels of speed, such as 10 and 15 m.p.h. over the speed limit. We will exploit this unique empirical

[1] On June 1999, President Clinton condemned racial profiling as "morally indefensible," and described it as "the opposite of good police work where actions are based on hard facts, not stereotypes." Finally, on June 18, 2003, the Bush administration ordered a broad ban on racial and ethnic profiling at all 70 federal law enforcement agencies.

[2] We will discuss Figure 1 in more detail in Section 5. In the process of providing us with the data, Bill Dedman was quick to notice and point out this curious nature of the citation distribution as well - in particular the huge heap of citations exactly at the speed of 10 m.p.h. over the limit.

feature to elicit officers' discretionary behavior and leniency.[3]

Discretion - at least for our purposes - is broadly defined as "latitude of choice within certain bounds imposed by law" (Merriam-Webster, 1996). A discretionary behavior by an officer may prevent him from reporting the actual speed of the driver for various plausible concerns. Strict behavior, in contrast, implies that the officer tickets all speeding drivers at the exact speed they drive. As one can easily predict, various factors such as drivers' age and financial situation - as much as the latter can be judged by officers - as well as the current high levels of the fines may play significant roles in officers' discretion. In this paper, while we look at the above issues, we examine particularly whether officers' discretionary behavior reflects their racial bias.

Our simple theoretical model borrows features from the recent literature regarding racial bias in police stops and searches (such as Knowles, Persico and Todd, 2001, and Antonovics and Knight, forthcoming; AK hereafter); in addition, our model incorporates several stylized facts pertaining to the institutional details in Massachusetts and to officers' and drivers' incentives. In our model, motorists, picking the speed to travel at, take into account the probability of getting ticketed as well as the speed that the officer will cite in deciding at what speed they will travel, while officers - net of the cost of ticketing motorists - maximize a benefit function generically increasing in the speed of ticketed drivers; this benefit function is then generalized to allow officers to give some drivers a break by citing them at a lower speed than they were traveling.[4] We then obtain results on discretionary behavior and differential bias among different races within different subgroups of the police.

Empirically, we examine the citation-officer matched data from the Boston Police Department between April 2001 and November 2002. As illustrated by the above histogram, we find that officers exhibit significant degrees of discretionary behavior. We verify that this discretionary behavior indeed indicates a major form of leniency by officers – underreporting. We find that white officers are the most lenient ones overall, followed by Hispanic and African-American officers, in terms of speed-discounting leniency. Female officers are relatively stricter than male officers. We

[3]The following anecdote seems to support our empirical strategy. "There are always mitigating circumstances in a stop," Officer Knecht said in an interview with the Boston Globe. "Anything could be said or could happen. Attitudes, people talking back to you. The circumstances change with each individual driver. But for most cops I know, race has nothing to do with it." He recalled that he did indeed show leniency to at least one African-American motorist on that day. "Although he wrote the man a ticket for only 10 m.p.h. over the 35 m.p.h. limit, he made a note in the top right-hand corner of the ticket: '64'." That meant that the driver was actually going 64 m.p.h., or 29 m.p.h. over the limit. Admittedly, Knecht would sometimes lower the speed on a ticket, to save a driver a high fine, and the notation was there in case the driver challenged the ticket in court.

[4]Another distinct source of discretionary behavior which is quite innocuous is to set cutoff speed levels in stopping drivers.

do not find evidence of systematic racial bias in the form of 'mutual or reciprocal' (e.g., white officers discriminating against Hispanic motorists, and in turn, Hispanic officers discriminating against white drivers) or 'monolithic' (e.g., all officers discriminating against African-American officers) racial bias. However, we find strong evidence that minority officers are less lenient to minority drivers. Interestingly, there is no such minority-on-minority bias in the case of female drivers or new-vehicle drivers. Our results, to a large extent, echo Dedman and Latour's findings about officers' discretionary behavior in issuing warnings.[5] There are also differences in how officers treat in-town vs. out-of-town motorists, and commercial vs. non-commercial motorists. Finally, we find that the degrees and forms of racial bias depend on its surroundings; there are significant variations in terms of time of the day as well as different neighborhoods. This conclusion calls for more empirical studies on racial bias across different contexts.

This paper is organized as follows. Next section summarizes the Related Literature. We then present the Theoretical Setup, which is followed by a section on the Empirical Predictions of the Theoretical Setup. Sections 5 and 6 are on Data and Empirical Strategy, respectively. In Section 7, we present our Empirical Findings, and finally Section 8 concludes.

2 Related Literature

Our work is related to two interrelated strands of recent research. The first strand concerns officers' decision-making regarding whether they issue tickets or warnings to a driver with certain characteristics. The main question is whether the decision-making is affected by the racial preferences of officers. The results are mixed. The state-sponsored Northeastern Study, conducted by criminologists Farrell and McDevitt (2004), used the same Massachusetts data set as that of Dedman and Latour.[6] Their general results reveal major disparities in ticketing behavior of officers as manifested by the distribution of tickets and warnings officers issued at different speeds to motorists of different races.[7]

[5] We also report results regarding warnings.

[6] The Northeastern study employs the standard benchmark test which has been traditionally used in the early literature studying racial bias in police stops and searches as well. This test compares the shares of racial minorities in the population to their shares in the sample of drivers ticketed (or stopped and searched by the police in the police search literature). In the use of this test, it is not clear what would be the right benchmark to compare the traffic citations. Ideally, it should be the racial composition of drivers on the road, but such information is typically unavailable.

[7] This naturally raises a red flag regarding the officers' intentions given that an early study by Lamberth (1996), who examined driving habits of African-American and white motorists on Maryland highways, found no difference in the rate at which these two segments of motorists engaged in speeding.

McConnell and Scheidegger (2001) compared tickets issued by air-patrol officers with citations issued by ground-patrol officers in order to overcome the traditional benchmark test's limitation. This is indeed a "blind" vs. "not-blind" comparison since the race of the driver is hard to be determined by the air-patrol officer. They found, surprisingly, that a smaller proportion of African-Americans received ground-patrol citations than air-patrol citations. Ridgeway (2006) studied the 7,607 stops recorded by the Oakland Police Department in 2003. Using the propensity-score matching method, he found that "black drivers are significantly less likely to be cited than non-black drivers, black drivers are slightly less likely to be cited than white drivers, and white and non-white drivers are not cited at significantly different rates" (p. 19).

Some studies also looked at other types of discretionary behavior. The above-mentioned Northeastern Study examined gender preferences of officers. Similarly, Blalock, DeVaro, Leventhal, and Simon (2007) compared the data from Bloomington and Highland Park in Illinois, Wichita, Boston, and the entire state of Tennessee and found out that female drivers are more likely to receive citations in three of the five locations, while male drivers are more likely to receive citations in the other two locations. Makowsky and Stratmann (forthcoming) took a different perspective and examined the political-economy determinants of speeding tickets and traffic fines. Using the first-two-month Massachusetts data excluding Boston, they found evidence that when local police officers issue tickets, they pursue additional objectives as well - apart from strict law enforcement -, such as maximizing their own utility and raising local government revenues from out-of-towners. Their paper is in the same spirit with and complementary to our paper in that both papers focus on police officers' discretionary behavior in issuing speeding tickets.

The second strand that our paper is related to is the quickly-growing literature on racial profiling in traffic stops and searches. Economists have recently joined the debate which was initially dominated by criminologists and statisticians employing the benchmark test. The interest of criminologists and statisticians in the subject started shortly after the 1993 civil case involving an African-American attorney, Robert L. Wilkins, as the plaintiff, who alleged that the Maryland State Police stopped and searched him simply because of his ethnicity. As a result of a consent decree, a researcher, John Lamberth, conducted extensive research as a part of criminal prosecution on Maryland highways; he later conducted similar research on New Jersey Turnpike. In both cases, he found major racial disparities in traffic stops and searches; his latter results were reported in State of New Jersey vs. Soto (1996). Later, Harris (1999) conducted interviews with motorists who were stopped, examined official records collected by several large departments in Ohio, and found similar results. More recently, Novak (2004) reported that, although disproportionate number of minorities is stopped and searched by the Kansas police, they are more likely to be stopped at night and to reside outside the city.

Grogger and Ridgeway (2006), on the other hand, had a different conclusion by examining the differences in stops when police are unable to observe the drivers' race at night versus in the presence of daylight. Using the data from Oakland, California, they find that African-American drivers are more likely to be stopped when it is dark outside and officers cannot easily observe the motorist's race. Consequently, they are unable to reject the null hypothesis of no racial profiling.

A number of recent papers in economics have attempted to determine whether the observed racial disparities in policing patterns can be explained better by models of statistical discrimination or by models of preference-based discrimination. In Knowles, Persico and Todd (2001), police decide which vehicles to search and motorists decide whether to carry contraband such as drugs or illegal weapons. Officers who are not racially biased maximize the number of successful searches, defined as uncovering contraband, net of conducting the cost of a search.[8] Racial bias is incorporated in the model as a reduction in the perceived cost of searching vehicles of certain types of motorists. Biased monitoring implies that the equilibrium rate at which contraband is seized (the "hit rate") is lower for the groups subject to bias. Using vehicle search data from Maryland, they found that the hit rates are indeed equalized across races.

Three other particularly relevant papers in that literature are Dharmapala and Ross (2004), Anwar and Fang (2006), and AK (forthcoming). Although they have their modeling differences, the first two papers assume that it is infeasible for the police to deter crime in a given subgroup of the population perfectly; given this assumption, they in turn show that the hit rate test is not necessarily valid. In addition, Anwar and Fang (2006) provide a test for "differential bias" within different subgroups of the police. Using the Florida State Highway Patrol data, they cannot reject the hypothesis of no differential bias. AK (forthcoming), whose analysis is similar in spirit to that of Anwar and Fang (2006), use the same Boston data we use in this paper; they find that officers are more likely to conduct a search if the race of the officer differs from the race of the driver. It is interesting to note that the same data show different types of racial bias; they find own-race preferences or racial bias against different races, while we find minority-on-minority racial bias.

[8] For in-depth discussion regarding the objective of police vehicle search, refer to Dominitz and Knowles (2006) and Close and Mason (2007). In addition, Close and Mason (2007), using the same data from the Florida State Highway Patrol used by Anwar and Fang (2006), compared consent versus non-consent searches and examined search rationale given by officers.

Our paper is in line with the latter two papers in that we focus on differential bias.

3 Theoretical Setup

3.1 Equilibrium without Underreporting and Racial Preferences

Suppose an officer is in charge of a particular segment of a route at a certain time of the day. We first consider the case of officers who neither have any racial preference nor engage in underreporting - the latter implies that, when these officers choose to ticket a speeding driver,[9] they report the actual speed of a driver.

Let c denote all *characteristics* of a motorist other than race that may affect the decision of the police officer as to whether to issue a traffic citation (interchangeably, ticket) or not. The variable c may be unobserved or only partially observed by third parties (including the econometrician). As in the racial bias literature, here too the variable c will be treated as a one-dimensional variable. Likewise, let C denote *all characteristics of an officer other than race* that may affect his decision as to whether to issue a traffic citation or not. The variable C too may be unobserved or only partially observed by third parties (including the econometrician) and will be treated as a one-dimensional variable. Experience and gender are among notable officer characteristics.

Officers compare the marginal cost and marginal benefit of *issuing a ticket* to a motorist who is *traveling at speed* S. Let $b(S, c|C) > 0$ denote the *benefit function of an officer of type* (C) from citing a driver of type (c) at their actual speed[10] $S > \underline{S}$, where $\underline{S} > 0$ is the *speed limit* on the portion of the road at which the driver was stopped.

Let $\tilde{S}(c) > 0$ denote the *maximum speed* a driver of type (c) may technically and safely find desirable in the absence of any fines. Let $\tilde{s}(c) \equiv \tilde{S}(c) - \underline{S}$ (and $s \equiv S - \underline{S}$). Observe that this re-scaling will allow us to measure every relevant speed level relative to the speed limit. To avoid trivial cases, we will assume that $\tilde{s}(c)$ for type (c) is high enough to accommodate real-life speeds (such as the ones, as data indicate, that can exceed the speed limit by more than 40-50 miles at times).

[9] Issuing warnings and not stopping a speeding driver are all equivalent in the eyes of the driver in terms of avoiding a hefty cost regardless of the particular action taken by officers.

[10] Although a driver may have an optimal speed in mind, it may be difficult to maintain that speed consistently; even when it is possible, a strict officer's radar gun as well may record it at a different speed level due to the inherent margin of error such electronic devices have. Thus, there is some degree of randomness involved regarding the speed of a driver. We will, however, deliberately abstract away from this aspect in our theoretical model in order not to complicate things further by incorporating motorists' risk preferences. As we will explain later, the randomness in fact strengthens our empirical findings.

Let $\tilde{s} \equiv \max\ \tilde{s}(c)$ over all driver types (c). We normalize the officer's benefit of a traffic ticket issued at that speed, $b(\tilde{s}, c|C)$, to equal one, so that the marginal cost of writing a ticket (or of stopping a driver - which we will consider in the Appendix) is scaled as a fraction of a well-defined maximum possible benefit to a police officer. For an officer, the *marginal cost of writing a ticket* to *any* motorist is denoted by t where $0 < t < 1$. Given the officer's cost of issuing a ticket to a driver and given the benefit function, we can define a *strict officer*: a strict officer of type (C) has $b(s, c|C) > t$ for any given s, c, and t; thus either $b(s, c|C)$ is very high for any $s > 0$ or t is very low or both. Whenever on duty, a strict officer will stop any speeding driver on the route he is in charge of, and report the actual speed the motorist was traveling at.

The *probability of a ticket* a driver of type (c) will receive while traveling at speed s will be denoted by $\gamma(s, c|C)$ (where $0 \leq \gamma(s, c|C) \leq 1$). Since any particular route may not be policed all the time, $\gamma(s, c|C)$ may be zero sometimes during a given day, and very low on certain days if police enforcement is absent for prolonged periods of time on specific days. Let $z > 0$ denote this probability that an officer will be present on a certain route at a given time. A motorist will take this probability into consideration and will consider $z \cdot \gamma(s, c|C)$ as the ticketing probability she faces.

Unlike the officers who encounter different driver types, many times every day, a motorist may hardly encounter different officer types every day or even every month. Therefore, a motorist's conjecture regarding the value of $\gamma(s, c|C)$ is first based on her observation of the frequency of stopped cars by police officers on her usual route on different days and times. Although the motorist may not be able to distinguish between the cases whether an officer is writing a ticket or a warning (or will be letting the stopped driver go with only an oral warning), such observations and the information that she gathers from other sources are her most important sources in reaching a conjecture on the value of $\gamma(s, c|C)$. Let $f(C)$ be the distribution of C in the officer population on that route. Let $E\gamma \equiv \int \gamma(s, c|C) f(C) dC$ denote the expectation of the driver of type (c) regarding $\gamma(s, c|C)$.

The *penalty function* is denoted by $p(s|c) > 0$. In Massachusetts, as mentioned before, for the first ten miles above the speed limit, the fine is \$75, and then it rises by ten dollars for each additional mile. But the major hit typically comes with the car insurance bill. As the Boston Globe has estimated, a Massachusetts driver will pay \$350 in higher insurance premia for a single ticket (and \$1,400 altogether for two tickets - in that sense as well, the penalty function depends on c), over the six years it stays on the driving record. Thus, the penalty function takes the form:

$$\begin{aligned} p(s|c) &= \alpha && \text{for} \quad 0 < s \leq 10 \\ p(s|c) &= \alpha + 10 \cdot (s - 10) && \text{for} \quad 10 < s \end{aligned} \tag{1}$$

where $\alpha > 0$ denotes the initial speeding fine \$75 plus the present discounted value of the \$350 insurance premium because of the first speeding ticket (or \$1,400 with a second ticket - and even more with yet another one) the driver will incur over the next six years.

For a driver of type (c), the function $v(s, c)$ will measure the *variable benefit from speeding* (such as arriving at a particular destination earlier, the pure joy of driving at a particular high speed, and so on), net of safety and gas mileage concerns. These variable benefits will be wiped out if there is a 100% ticketing probability that the driver faces at some speed; that is, this variable part of the benefit function will affect a driver's utility more as the ticketing probability decreases. Thus, it is reasonable to presume that a motorist's utility function has a non-negative *fixed benefit* portion that the driver does not lose even after getting ticketed with 100% probability. Let $D \geq 0$ denote this *fixed benefit* in the motorist's utility function, unaffected by ticketing probability she faces. Putting everything together, the *utility function of driver of type* (c), will have the canonical form (a somewhat similar version of which is also used in the police stop-and-search literature):

$$\begin{aligned} & u(s, c, z, \gamma) = D + (1 - z)v(s, c) + z(1 - E\gamma)v(s, c) - zE\gamma p(s|c) \\ & \Leftrightarrow u(s, c, z, \gamma) = D + (1 - zE\gamma) \cdot v(s, c) - zE\gamma p(s|c) \end{aligned} \tag{2}$$

A driver will maintain a speed s^* at which her expected utility is maximized. Any realistic variable benefit function should comply with the fact that many drivers find it optimal to drive beyond the speed limit as well as to speed more than ten miles above that limit. This implies the presence of certain types of variable benefit functions.[11] Further, given some parameter values, any such function should generate an optimal natural maximum speed in the absence of any fines ("natural" due to safety and technical reasons, as mentioned above). For that to happen, the variable benefit function should also exhibit a declining range of marginal benefit at very high speeds and eventually reach a value of zero at $s^* = \tilde{s}(c)$.

The simplest functional form that exhibits all of these properties is the quadratic form. Observe that the variable benefit function below exhibits increasing benefit in some range of speeds exceeding

[11] In the Appendix, we show that some simple linear functional forms - such as $v(s, c) = a_0$ and $v(s, c) = a_0 + a_1 \cdot s$ - should not be considered due to that concern.

the limit by 10 m.p.h. and declining benefit at some higher speeds (eventually reaching the peak at $s^* = \tilde{s}(c)$, at least for some coefficient values). Let $a_0 \geq 0$ and $a_1, a_2 > 0$:

$$v(s, c) = a_0 + a_1 \cdot \tilde{s}(c) \cdot s - a_2 \cdot s^2 \qquad \text{for } s \leq \tilde{s}(c) \tag{3}$$

The marginal benefit of speeding with this functional form is positive first and becomes negative for higher s. We will use the functional form above as the canonical form of a driver's variable benefit function with type (c). First, observe that, in order for a driver to speed above the limit, $u(s, c, z, \gamma) = D + (1 - zE\gamma)v(s, c) - zE\gamma p(s|c) > 0$ must hold. That is, if $u(s, c, z, \gamma) < 0$, the driver will not speed. If $u(s, c, z, \gamma) = 0$, the driver will be indifferent between speeding and not speeding.

Given Equations (1), (2) and (3), the optimizing behavior yields, for the driver of type (c), the optimal s^* miles per hour above the speed limit as follows:

$$s^*(c, z, E\gamma) = \frac{1}{2a_2}\left(a_1\tilde{s}(c) - p'\frac{zE\gamma}{(1 - zE\gamma)}\right) \equiv \arg\max u(). \tag{4}$$

where p' denote $dp(s|c)/ds$, which is 0 if $s^* < 10$ and 10 if $s^* \geq 10$. Observe that s^* decreases in both $E\gamma$ and z. In addition, for $\tilde{s}(c)$ to be the maximum speed, $a_1 = 2a_2$ must hold. Note that there is no restriction for the range of s^* except $s^* \leq \tilde{s}(c)$. Given the undifferentiable structure of the penalty function, one may erroneously expect $s^*(c, z, \gamma)$ to be either zero or greater than 10. The following figure, however, illustrates that $0 < s^*(c, z, \gamma) \leq 10$ is possible as well since $s^*(c, z, \gamma)$ can reach a maximum in that range.

An officer makes two main types of decisions: the *stopping decision* and the *ticketing decision*. The *stopping decision* (which is not fully observable to the econometrician) is rather complicated - and to some extent non-crucial to our analysis; a full-fledged version of it (incorporating under-reporting and racial preferences) will take place in the Appendix.

Suppose the officer chooses to stop a driver who is traveling at speed s; upon stopping the driver, the officer can find out about (c). Then an officer of type (C) will obtain the following net benefit from ticketing that driver

$$b(s, c|C) - t \tag{5}$$

Equation (5) implies the following. If $b(s, c|C) - t < 0$, then the optimizing behavior implies $\gamma(s, c|C) = 0$; that is, the officer's best response is to never ticket motorist type (c), who speeds at s. If $b(s,c|C)-t>0$, then the officer will be willing to issue a ticket to type (c); that is, the officer's best response is to always ticket motorist type c who speeds at s. If $b(s,c|C)-t=0$, then the officer will be willing to randomize over whether to issue a ticket to type c, who speeds at s or not ticketing the motorist at all.

Thus, for motorists of type (c), and officers of type (C), there is a unique equilibrium of this discretionary interaction (i.e., game) in which (1) drivers' correct conjecture regarding the expected ticketing probability of an officer, $zE\gamma$, on a given route and the penalty function, $p(s|c)>0$, renders motorists of type (c) driving on that route to choose a particular best response speed level, and (2) the behavior of motorists of type (c) and an officer type (C)'s maximum benefit associated with each $(s,c|C)$ render the officer indifferent between ticketing them after stopping them at some speed level (i.e., such an officer will surely be willing to ticket them beyond that speed level and not willing to ticket them below).

Figure 3a graphs the canonical best response functions for motorists and officers when $s^* \geq 10$. Then the step function represents the best response function of the officer on that route and the downward-sloping curve represents the best response function of motorists type (c) given the expectation about the officer's type. Recall that z denotes the probability that the officer will be present on that certain route at any given time that the driver travels. Figure 3b depicts the equilibrium in which the driver takes into consideration the expected value of the two different types of officers' ticketing best response functions (high cutoff speed, 'high s^*', and low cutoff speed, 'low s^*'). Taking that expected value into consideration, the driver's equilibrium speed is such that she would be ticketed by the stricter officer type while she would not be ticketed by the more lenient officer type.

3.2 Equilibrium with Underreporting and Racial Preferences

Let $r \in \{a,w\}$ denote a driver's *race*, and let $R \in \{A,W\}$ denote an officer's *race* which are observable by both parties upon meeting each other in person. Let $f(C|W)$ and $f(C|A)$ be the distribution of C in the white and African-American officer populations on that route, respectively. Let $g(R)>0$ be the probability that the officer on the route the driver is traveling is of race R. Then, in this subsection, let $E\gamma \equiv \sum_{A,W}\left(\int(\gamma(s,c,r|C,R)f(C|R)dC\right)g(R)$ denote the expectation of the driver of type (c,r) regarding the different ticketing probabilities,$\gamma(s,c,r|C,R)$, by different officer types (C,R).

Underreporting is such that the officer cites a driver of type (c,r) at speed $\sigma<s$; if $\sigma=0$,

observe that the motorist is not ticketed.[12] Let $b_r^R(\sigma, s, c|C) \in (0.1]$ denote the *benefit function of an officer of type* (C, R) from citing an additional driver of type (c, r) at speed σ, who was actually traveling at speed s.

There may be some drivers with type (c, r) assessed more favorably by some officer of type (C, R) - e.g., polite, old, quite, female - and some other characteristics assessed less favorably by him - e.g., rude, young, talking-back, male; likewise, non-minorities may be assessed more favorably than minorities. It is reasonable to think that such an officer will report a speed $\sigma^* < s$ in the former cases, and a speed $\sigma^* = s$ in the latter cases. Nevertheless, we will assume that σ will be non-decreasing in the actual speed s.

With underreporting, the penalty function takes the form:

$$\begin{aligned} p(\sigma|c) &= \alpha && \text{for} \quad 0 < \sigma \leq 10 \\ p(\sigma|c) &= \alpha + 10(\sigma - 10) && \text{for} \quad 10 < \sigma \end{aligned} \tag{1'}$$

Then let $E\sigma = \sum_{A,W} \left(\int (\sigma(s|c, r, C, R) f(C|R) dC \right) g(R)$ denote the expectation of the driver of type (c, r) regarding different underreporting possibilities,$\sigma(s|c, r, C, R)$, by different officer types (C, R). With underreporting and the different races of drivers, the driver of type (c, r)'s utility function will take the above penalty function into consideration:

$$u_r(s, c, r, z, \sigma, \gamma) = D + (1 - zE\gamma)v_r(s, c, r) - zE\gamma p(E\sigma) \tag{2'}$$

The variable benefit function of a driver type (c, r) is given by

$$v_r(s, c) = a_0 + a_1\tilde{s}(c, r)s - a_2s^2 \qquad \text{for} \ \ s \leq \tilde{s}(c, r) \tag{3'}$$

Since σ is assumed to be non-decreasing in the actual speed s that is clocked by the officer up to certain speed level and increasing at or beyond that, it implies that the driver also knows that $dp(E\sigma)/ds = k \geq 0$ up to some relatively high speed level and $dp(E\sigma)/ds = k > 0$ at or beyond that speed level.

Given Equations (1'), (2') and (3'), the optimizing behavior of the driver of type (c, r) yields

[12]For some obvious reasons for officers such as reputation concerns and court-time loss, we rule out overreporting as a theoretical possibility. We will check the validity of this assumption in Section 7.4.

speeding s_r^* miles above the speed limit as follows:

$$s_r^*(c,z,\gamma) = \frac{1}{2a_2}\left(a_1\tilde{s}(c,r) - k\frac{zE\gamma}{(1-zE\gamma)}\right). \tag{4'}$$

Figure 4a depicts the driver's best response function based on the expectation of reported speed when the driver is stopped and ticketed. Similar to Figure 3b, Figure 4b depicts the equilibrium in which the driver takes into consideration the expected value of the two different types of officers (high σ and low σ). When the driver believes that the officer will favor him (and cite at low σ), the driver will speed faster given a value of $E\gamma$.

We assume that an officer of race (R)'s cost of writing a ticket to a driver of race (r) has an additional component: t_r^R is the additional cost of underreporting one additional mile, where $0 < t_r^R < 1$; that is, if she drives at speed s and the officer decides to underreport the speed at $\sigma < s$, then the "integrity cost" of underreporting is $t_r^R(s-\sigma)/s$; if the officer has no racial preferences, then regardless of different driver races, the officer has the same t^R. Suppose the officer chooses to stop a driver. Then the officer of type (C,R), will choose the *probability of ticketing*, $\gamma^R(\sigma,s,c,r|C)$, each motorist of type (c,r) at speed σ, who is actually driving at speed s, by considering the following net benefit from ticketing that driver[13]

$$b_r^R(\sigma,s,c|C) - t_r^R\left(\frac{s-\sigma}{s}\right) - t \tag{5'}$$

Let b' denote $d\left(b_r^R(\sigma,s,c|C)\right)/d\sigma$; i.e., b' is the marginal benefit of reporting one more mile. The officer will determine his optimal underreporting level, if any, by equating

$$b' = -t_r^R/s.$$

Figure 5a depicts a benefit function regarding a favorite driver type (c,r) which reaches a maximum below the actual speed of the driver. In Figure 5b, we observe a benefit regarding a non-favorite driver type (c,r) which reaches a maximum at the actual speed of the driver. The two different levels of marginal costs imply the same level of citation, $\sigma = s$.

After the officer figures out the optimal citation, $\sigma^* \leq s$, he will next figure out whether to ticket the driver at that cited speed or not. If $b_r^R(\sigma^*,s,c|C) - t_r^R\left(\frac{s-\sigma^*}{s}\right) - t < 0$, then the optimizing behavior implies $\gamma = 0$; that is, the officer's best response is to never ticket motorist type (c,r) who

[13]There is also the issue of discounting a fine, although there is an explicit formula for the dollar amount set by state law. We will examine this issue in Section 7.6.

speeds at s. If $b_r^R(\sigma^*, s, c|C) - t_r^R\left(\frac{s-\sigma^*}{s}\right) - t > 0$, then the officer will be willing to issue a ticket to type (c, r); that is, the officer's best response is to always ticket motorist type (c, r) who speeds at s. If $b_r^R(\sigma^*, s, c|C) - t_r^R\left(\frac{s-\sigma^*}{s}\right) - t = 0$, then the officer will be willing to randomize over whether to issue a ticket to type (c, r), who speeds at s or not ticketing the motorist at all.

Thus, for motorists of type (c, r) and officer type (C, R), there is a unique equilibrium of this discretionary interaction in which (1) an officer of type (C, R), given his maximum benefit associated with each (σ, s, c, r), decides on the optimal speed he will cite, $\sigma^* \leq s$, (2) drivers' correct conjecture regarding the expected ticketing probability of officers, $zE\gamma$, on a given route and the expected cited speed render motorists of different types driving on that route to choose particular best response speed levels, and (3) the behavior of motorists and an officer's maximum benefit associated with each (σ, s, c, r) render the officer indifferent between ticketing them after stopping them at some speed level.[14]

4 Empirical Implications of the Theoretical Setup

An important goal of an officer, who underreports the speed of a driver with a favorable type (c or r or both), is that the driver faces a smaller penalty than her actual speed would imply. Suppose that there are two segments of drivers. Both segments take the penalty function (which depends on σ) into consideration, which leads to a higher speeding best-response for the segment (with characteristic (c)) that is favored by the officer up to the particular high speed beyond which the officer does not show any leniency to any type of drivers. Let (c') denote the characteristic of the other segment that is not favored by the officer. Then, in Figure 6, the officer's best-response ticketing probability step functions intersect the best-response functions of the (c) and (c') segments of drivers at two different levels of speed, e.g. $s = 11$ for type (c) and $s' = 14$ for type (c'). But suppose the case in which the (c) type's speed is cited as $\sigma = 0$ (i.e., the driver is let go with an oral warning) while the (c') type's speed is cited as $\sigma = 14$.

Consider another officer with somewhat similar benefit functions but with different ticketing costs such that his step functions intersect those best-response functions of two segments of drivers at speeds, for instance, $s = 12$ and $s' = 15$, and yet another officer at speeds, $s = 13$ and $s' = 16$; suppose that in both of these cases, the (c) type is cited at speed $\sigma = 10$ while the (c') type is cited at her actual speed. Then there will be one citation each at $\sigma = 14$, $\sigma = 15$, $\sigma = 16$, two

[14]The Appendix considers various further generalizations our theoretical model, including the incorporation of the officers' stopping decision.

citations at $\sigma = 10$, and one warning. As we add more such officers, their discretionary behavior will imply a rather disproportionate number of citations at $\sigma = 10$.[15] This example clearly shows that underreporting is one plausible cause for massive spikes on speed dispersion like Figure 1.[16]

Now suppose a certain race of officers (or perhaps a couple of races of officers) have the following racial preferences: $r \succ r'$. Then race (r) motorists will obtain disproportionately more citations at $\sigma = 10$ and warnings (and, compensatingly, relatively fewer citations at nearby higher speeds) from these officers, and that situation will be the least for race (r') motorists. Thus, discretionary behavior and racial bias by such officers lead to a disproportionate number of citations by these officers at a focal level of speed, such as $\sigma = 10$, and warnings. In other words, such disproportionate number of citations at a particular level of speed will readily indicate the presence of discretionary behavior or racial bias by some segments of officers. Therefore, average officer characteristics at the prominent speed are more likely to reflect observable characteristics of discretionary officers. Similarly, drivers who were cited at the speed are more likely to have some favorable characteristics (and they actually drove faster). Finally, we should observe more favorable combinations of officers' and drivers' characteristics at the prominent speed level.

Also, the "integrity marginal cost" parameter, t_r^R, of an officer can presumably be even higher in a neighborhood more populated with people of his own race. Then, officers can be expected to be stricter in such neighborhoods. Further, with the help of some stylized facts regarding which segments of drivers are more likely to be on the roads on certain days and times, we can make predictions regarding 'day of the week' and 'time of the day' as well. For example, on Sundays, the prevalence of elderly churchgoers may elicit more lenient behavior from officers in terms of underreporting. Likewise, at night, even their preferred-race drivers may be perceived less favorably by officers.

Lastly, but not leastly, our model shows the importance of controlling for actual speed when we compare the ticketing probability between two segments of drivers. If actual speed is unobservable or omitted, comparing the ticketing probability might as well be misleading because it is possible to find, like Figure 6, that favored drivers are more likely to be ticketed than non-favored drivers since favored drivers are on average more likely to speed. This omitted-variable problem is troublesome

[15] Albers and Albers (1983 noted that certain numbers and fractions are more prominent than others - such as round numbers and percentages. We do not attempt to explain why 10 is a prominent number in our context.

[16] Another possible cause is the kinked penalty function at 10 m.p.h. over the limit, which leads some drivers to choose that exact speed rather than its adjacent speeds. However, first of all, this does not explain the spikes at other speed levels. Also, since the kinked-penalty-function explanation totally depends on drivers' heterogeneity, we can distinguish this hypothesis and our discretionary-behavior hypothesis by testing whether tickets cited at those focal speed levels are related to certain characteristics matching between officers and drivers.

since reported speed is likely to be different from actual speed even when information on speed is available in the data.

5 Data

The main data set we use in this paper is based on the record of 2,001,562 traffic citations issued in Massachusetts between April 2001 and November 2002. The data set contains basic demographic information about drivers, such as race, gender, age, and home town. Our data also contain information on time and date when citation was issued as well as the neighborhood in which the motorist's vehicle was stopped. We are able to match the citation-level data with the officer-level data including individual officers' race, gender, and experience on the force. The officer-level data are available only for the Boston police department. All officers in our sample are, therefore, municipal police officers. There are 161,133 officer-citation matched citations issued by Boston police officers within Boston.

Because of the particular focus of our paper, we only consider speeding tickets and warnings. They account for 25.7% of all issued citations, the largest portion as any single category. Most citations in our sample are tickets because warnings were recorded for the first two months only, April and May in 2001, and then stopped being collected due to budgetary shortfalls. We deleted citations with missing information regarding any variable that we will use for our regression analysis; in particular, there are 2,041 citations without speed and 3,128 citations without drivers' race. To be comparable to the previous literature on vehicle searches, we focus on three races, white, African-American, and Hispanic. We deleted 1,875 citations where drivers are not white, African-American, or Hispanic and 1,031 citations issued by Asian officers. Consequently, the final sample consists of 25,738 tickets and 2,644 warnings. Table 1 presents sample selection criteria.

5.1 Histogram of Tickets

Recall that Figure 1 graphs the frequency histogram of vehicle speed in miles per hour over the posted speed limit as written on tickets, as denoted by σ in our theory section. Two notable features stand out. First, there exist sizable spikes at multiples of 5 m.p.h. above the limit and, more distinctly, a massive one at the speed of 10 m.p.h. above the limit. Given the size of the spikes, it does not seem likely to explain their presence mainly by drivers' heterogeneity, needless to say that it would be hard to believe that drivers can control their speed so delicately particularly considering the city traffic conditions in Boston.

The more likely possibility is that the spikes result from officers' discretionary reporting behavior. There are two possible explanations, which are not necessarily mutually exclusive: first, as the stopping decision - captured by Equation (A5) in the Appendix - implies, officers can set their speed gun to beep at those specific speeds. However, such a stopping decision can only provide a partial explanation for those spikes. Suppose that an officer sets 10 m.p.h. above the limit as the stopping threshold. Then he will not stop the motorists driving under 10 m.p.h. above the limit, which can explain as to why the histogram abruptly drops below 10 m.p.h. above the limit.

The second, and more comprehensive explanation lies in the ticketing decision of officers - captured by Equation (5') in the theory section -, that entails officers having leeway to give a break to certain types of drivers. When officers decide to give a motorist a break, they can do so by letting that motorist go without ticket or by reporting a lower speed. It is still hard to explain why officers prefer those distinct numbers just on the basis of the ticketing decision analysis provided by Equation (5'). But for various other reasons (such as officer using 'prominence levels' as well as officers 'not wanting to look too meticulous'), it may seem reasonable to imagine that once an officer decides to give a break to a driver who drove - somewhat but not much - over 10 m.p.h. above the limit, say, 16 m.p.h. above the limit, the officer would choose 10 rather than 9 or 11 m.p.h. above the limit. Nevertheless, these spikes, particularly the massive one at 10, will play a crucial role in our empirical identification in this paper.

Recall that Figure 3.a illustrates the equilibrium speed ($s^* \geq 10$) and ticketing probability for a given type of drivers and a given type of officers when there is neither underreporting nor racial bias. It is possible to derive equilibrium speed dispersion by allowing the ticketing cost, t, to vary across individual officers. There will be a spike, for example, at 10 m.p.h. over the limit if and only if the distribution of t also has a corresponding spike exactly at $t = b(10, c|C)$. In this case, if the officer's benefit function is continuous, drivers cited at immediately higher speeds (e.g. 11 m.p.h. over the limit) should be similar in characteristics with those cited at 10 m.p.h. over the limit since the spike results from officers' heterogeneity. This is a testable hypothesis. On the other hand, if we introduce the possibility of underreporting, it is not necessary for us to assume an unusual and unlikely distribution of t in order to explain the spike in the distribution of reported speed. Instead it is possible that, as shown in Figure 4, favored drivers who are stopped at higher speeds are cited at the speed of the spike (i.e. 10 m.p.h. over the limit). Since there are non-favored drivers who were actually traveling at 10 m.p.h. over the limit and were ticketed and cited at that speed, we should find a mix of favored and non-favored speeders. However, those drivers cited at immediately higher speeds (e.g. 11 or 12 m.p.h. over the limit) are not necessarily similar in characteristics with those cited at 10 m.p.h. over the limit because there remain relatively more non-favored drivers at those higher speeds.

The second notable feature that stands out is that there are very few tickets below 10 m.p.h. above the limit. Recall that it is reasonable to imagine that many motorists - considering the higher probability of getting caught at a higher speed and the higher probability of being involved in a fatal accident - may have benefit functions due to which they may find it optimal to drive below 10 m.p.h. above the limit once they decide to speed (see Figure 2 in the theory section). Thus we think that there are few tickets below 10 m.p.h. above the speed limit because officers forgave these slow speeders or did not even attempt to stop them - due to their stopping and ticketing decisions captured by Equation (A5) in the Appendix and Equation (5') in the theory section respectively.

In Figure 1, we overlay a hypothetical distribution of actual speeds. The distribution is graphed under some reasonable assumptions. We assume that: (1) the distribution should be smooth; (2) those citations accumulated at the prominent speeds are discounted from some speeds above; (3) there are few unstopped vehicles above 10.[17] We can see that, below 10 m.p.h. above the limit, there are a substantial number of speeding vehicles that are not even stopped. It seems likely that the number of completely forgiven or ignored speeding vehicles is larger than the total number of tickets in our sample.

5.2 Characteristics by Speed

Given our discussion in the previous subsection, in Tables 2 and 3, we examine drivers' and officers' characteristics disaggregated by reported speed on tickets. As motivated by our findings in the previous section, our discussion will focus on those citations at 10 m.p.h. above the limit and its adjacent speeds.[18] If a certain characteristic of drivers or officers is related to the manipulation of reported speed, we should find a significant discontinuity in that variable around the specific speed. If officers were strict, different speed levels should only reflect heterogeneity in drivers' propensity to speed and, if any, changes in the means over speed should be gradual.

[17] According to the 2002 National Survey of Speeding and Other Unsafe Driving Actions conducted by National Highway Traffic Safety Administration in U.S. Department of Transportation, about 51% of drivers say that they sometimes or often drive 10 m.p.h. over the speed limit on interstate highways. People believe that they can travel about 8 m.p.h. over the limit on interstate highways without getting a speeding ticket. Allowable speed margins over the limit for city or neighborhood streets are deemed slightly lower.

[18] One should be careful in interpreting statistics below 10 m.p.h. over the limit given that there are very few observations.

Table 2 shows the means and standard deviations of drivers' characteristics. The column for the prominent speed (10 m.p.h. above the limit) is highlighted. First, the proportion of in-town drivers is lower at 10 m.p.h. above the limit. Since there is no reason that in-town drivers drive less frequently at 10 m.p.h. above the limit than others, it seems reasonable to suppose that officers are less likely to cite those drivers at that particular speed. This may be due to the fact that officers tend to give a break to out-of-towners simply because of learning-effects; i.e., they may be more lenient to out-of-towners since these drivers may not be expected to know the road and driving conditions as much and consequently could be forgiven for paying more attention to these conditions than to their speedometers, particularly in the traffic condition of Boston. There is a similar but weaker pattern for in-state drivers. We also find that the proportion of male drivers and that of white drivers are higher at the speed, while the proportion of African-American drivers is lower. This finding suggests the possibility that drivers are treated differently depending on their gender and race.

Table 3 shows officers' characteristics over speed. We find significant differences in various aspects. We find that male and/or inexperienced (or younger) officers are more likely to give a break to drivers. The most striking difference is found in the proportion of white officers at 10 m.p.h. above the limit. About 70 percent of all the tickets cited at 10 m.p.h. are issued by white officers, while they account for 40 percent and 30 percent at 9 and 11 m.p.h. above the limit, respectively. Lastly, note that the proportion of African-American officers is high particularly at 11 and 12 m.p.h. above the limit. This seems to reflect the opposite side of the same coin as the high proportion of white officers at 10 m.p.h. above the limit. It is reasonable to imagine that there are relatively more African-American officers at 11 and 12 m.p.h. above the limit because those tickets issued by white officers, which are supposed to be at 11 or 12 m.p.h. above the limit, are moved to 10 m.p.h. above the limit.

6 Empirical Strategy

Now we consider a statistical model for an officer's choice of reported speed given that the officer has already stopped a vehicle and decided to cite the driver. As a starting point, suppose that we could observe whether the officer gives a driver a break. We can then specify the officer's propensity to give a break based on observable characteristics including the officer's and the driver's races. Let y_{ij} denote the variable that takes on the value of one if officer j gives a break to driver i and, otherwise, takes the value of zero.

$$y_{ij} = \begin{cases} 1 & if\ y^*_{ij} \geq 0 \\ 0 & if\ y^*_{ij} < 0 \end{cases} \tag{6}$$

where $y^*_{ij} = \beta_0 - \beta_1 s_i + c_i\beta_2 + C_j\beta_3 + L_{ij}\beta_4 + \{driver\ race\} + \{officer\ race\} + \{racial\ bias\} - \varepsilon_{ij}$. The latent variable y^*_{ij} represents the officer's propensity to give a break. The variable s_i represents the driver's actual speed over the limit that is observed by the officer when the officer stopped the vehicle. Consistent with an assumption we have made in the theory section, we assume here that the propensity will be decreasing in s_i. The vectors c_i and C_j include driver and officer characteristics other than race, respectively. Driver characteristics include age, gender, dummy for in-town drivers, dummy for in-state drivers, and indicator for commercial license. Officers' characteristics include gender and years on the police force. To address the possibility that officers' leniency varies by surroundings, the vector L_{ij} includes time, date, and location for tickets; there are three time dummies (morning, afternoon, and evening, with predawn excluded), six date dummies (Tuesday excluded), ten neighborhood dummies, and one continuous variable for the speed zone (the speed limit in m.p.h.). Lastly, the variable ε_{ij} is a disturbance variable representing what is unobservable to the econometrician but may influence the officer's ticketing decision.

To estimate the effects of racial bias on officers' leniency, we include three sets of dummy variables created by driver race and officer race. The first set, denoted as $\{driver\ race\}$, includes two dummy variables, $1[i = a]$ and $1[i = h]$, which represent African-American drivers and Hispanic drivers, respectively. Holding other things, particularly the actual speed s_i, constant, the coefficients for these two driver race dummies are intended to capture statistical discrimination or monolithic racial bias. The first emphasizes the 'schooling drivers' aspect: officers may be stricter with drivers of a specific race if they believe that those drivers will likely speed again when given a warning or treated leniently. The latter, monolithic racial bias, could be socially imposed: officers may feel obliged to enforce their chiefs' or their communities' racial bias against a particular race. For example, minority community leaders often call for harsh law enforcement because they are more easily blamed than whites.

The second set, denoted as $\{officer\ race\}$ in Equation (6), also includes two office race dummy variables $1[j = A]$ and $1[j = H]$, which represent African-American officers and Hispanic officers, respectively. The coefficients for these two dummies should capture officers' race-specific strictness (relative to white officers). For example, if African-American officers are strict or less lenient, then the corresponding coefficient should be negative.

The final set of dummy variables is denoted as $\{racial\ bias\}$. These variables are supposed to capture officers' preferences for drivers with different races, which also may change by officers' race. For this purpose, this set should include 6 interaction variables given three racial groups.

$$\begin{aligned} racial\ bias &= d_w^W 1[i=w, j=W] + d_a^W 1[i=a, j=W] + d_h^W 1[i=h, j=W] \\ &+ d_w^A 1[i=w, j=A] + d_a^A 1[i=a, j=A] + d_h^A 1[i=h, j=A] \\ &+ d_w^H 1[i=w, j=H] + d_a^H 1[i=a, j=H] + d_h^H 1[i=h, j=H] \\ \\ &= 1 + (d_a^W - d_w^W) 1[i=a, j=W] + (d_h^W - d_w^W) 1[i=h, j=W] \\ &+ 1 + (d_w^A - d_a^A) 1[i=w, j=A] + (d_h^A - d_a^A) 1[i=h, j=A] \\ &+ 1 + (d_w^H - d_h^H) 1[i=w, j=H] + (d_a^H - d_h^H) 1[i=a, j=H] \end{aligned} \tag{7}$$

Te coefficient $(d_i^j - d_j^j)$ represents the racial bias by officers with race j for drivers with race i against drivers with officers' own race j. However, it is not possible to estimate all the six coefficients due to perfect multicollinearity. It is necessary to impose some reasonable restrictions on parameters (we can only estimate 4 parameters). For example, AK (forthcoming) imposed the following restrictions:

$$d_a^W - d_w^W = d_h^W - d_w^W = d_w^A - d_a^A = d_h^A - d_a^A = d_w^H - d_h^H = d_a^H - d_h^H = d.$$

Therefore,

$$\begin{aligned} racial\ bias &= d\{1[i=a, j=W] + 1[i=h, j=W] + 1[i=w, j=A] + 1[i=h, j=A] \\ &+ 1[i=w, j=H] + 1[i=a, j=H]\} = dMismatch, \end{aligned} \tag{8}$$

where *Mismatch* is a dummy variable for whether the officer's race and the driver's race are different. This is the sum of the above six indicator variables.

The validity of this restriction is an empirical as well as conceptual question; that is, whether, if any, racial bias occurs in the form of own-race preferences, $d_i^j - d_j^j < 0$ for all i and j (or distaste about own-race drivers, $d_i^j - d_j^j > 0$ for all i and j). In this paper, we allow and test two more hypothetical forms of racial bias: (1) minority-on-minority bias and (2) African-American-white confrontation. Both hypotheses are motivated by our preliminary scrutiny of the data and the literature. The first hypothesis is that minority officers are stricter (less lenient) on minority drivers. Technically, we include the interaction term *Minority*×*Minority*, which is $1[i \neq w] \times 1[j \neq W]$.

Recall that Dedman and Latour, using the same data set, found that minority officers are tougher on minority drivers by issuing more tickets than warnings. The latter is that there exists two-way bias between white and African-American officers, which might be a type of racial bias usually perceived by the public. Note that these three forms of racial bias (own-race preferences, minority-on-minority bias, and African-American-white confrontation) may coexist. Indeed they are not mutually exclusive although they might as well compete with each other. Our full specification of racial bias is the following:

$$\begin{aligned} y^*_{ij} = \beta_0 - \beta_1 s_i + c_i\beta_2 + C_j\beta_3 + L_{ij}\beta_4 + \beta_5 1[i = a] + \beta_6 1[i = h] + \beta_7 1[j = A] + \beta_8 1[j = H] \\ + d_1 Mismatch + d_2 Minority \times Minority + d_3 1[i = a, j = W] + d_4 1[i = w, j = A] - \varepsilon_{ij} \end{aligned} \tag{9}$$

6.1 Warnings versus Tickets

An obvious measure of officers' leniency is whether they give warnings rather than tickets given a level of speed. In Table 4, we estimate the Probit model where the dependent variable is the indicator for whether the driver is warned rather than ticketed. Here we assume that the reported speed is the actual speed, which is not true and will be discussed later.[19] Note, however, that the main purpose of this subsection is to compare our analysis with the Boston Globe's analysis presented in Column (1). We try to replicate the regression analysis done by Professor Elaine I. Allen which was asked by the Boston Globe. Although the samples and specifications are not exactly identical, our findings are overall in harmony with the Boston Globe's: officers are stricter for faster drivers; white, older drivers, and/or in-town drivers are more likely to be warned; male drivers are less likely to be warned.

In Column (2), we include our racial bias terms as well as other control variables as specified in Equation (9). Some results become different from the Boston Globe's. We find that the driver age effect is nonlinear: officers are more lenient for younger and older drivers while they are strict for prime age drivers. We find no gender disparity.

More importantly for our purpose in this paper, we find that officer characteristics are significant: male and/or less experienced officers are more lenient. Given a cited speed and holding other things constant, they are more likely to issue warnings instead of tickets. The racial bias terms

[19] We can minimize this problem by using a series of dummy variables for speed ranges (such as below 10, 10-14, and 15 or above) instead of using the continuous variable. This solution is reasonable in that the speed is discounted to its nearest prominent level. Our results change little by using dummy variables for speed ranges.

are also significant. Our full specification in Column (2) shows that minority officers are stricter to minority drivers. In other words, minority drivers are more likely to be warned and not ticketed when they are stopped by white officers. This is consistent with Dedman and Latour's articles.

6.2 Reported Speed

Now we allow the possibility that officers manipulate the speed on tickets. The actual speed is not observable to the econometrician. Thus we assume that:

$$s_i = \mathrm{X}_i\alpha_1 + L_i\alpha_2 + \alpha_3 1[i = a] + \alpha_4 1[i = h] + \omega_i \tag{10}$$

where the vector X_i includes the driver's characteristics affecting speeding behavior. The vector L_i is the same as defined in Equation (6). The coefficient α_3 and α_4 are supposed to capture average racial differentials in speeding behavior. Depending on the driver's race, the maximum speed at which he or she can drive - $\tilde{s}(c, r)$ in our theoretical model - may differ.[20] In addition, drivers of different races may have different perceptions regarding their likelihood they will be ticketed - $E\gamma$ in our theoretical model. The variable ω_i represents unobservable individual heterogeneity such as, among others, risk attitude and time discount.

One noteworthy thing is that there is no officer variable included in Equation (7) because it is unknown to the driver which type of officer s/he will encounter when s/he is stopped. Also note that we include location variables controlling which neighborhood the motorist was driving through. The path the driver follows is clearly his/her choice, and if officers are assigned across districts in a systematic way based on their observable types (e.g., race) and if drivers know this assignment rule, the driver may be able to predict to some extent which type of officers he or she will likely face on their path. For example, people might expect more African-American officers in a neighborhood in which African-Americans are concentrated (AK, forthcoming). While this expectation should be expected to affect speeding behavior or route choice for criminals (or those who try to avoid police stops, e.g. joyriders), it should not for other types of motorists. First, it should be quite costly, in terms of both time and gasoline, to change their route given their origin and destination. Again we believe that those drivers within our focused range of speed, 10-14 m.p.h. above the limit, should not choose their travel route in order to avoid certain types of officers.

Drivers' characteristics that affect officers' leniency could affect drivers' speeding behavior. To

[20] Surprisingly there is little empirical research on drivers' race and speeding. But we believe that speeding behavior should not be significantly different by drivers' race within our narrow range of speed, 10-14 m.p.h. above the limit.

be general, we do not assume any exclusion restriction; $c_i = \mathrm{X}_i$. For example, the driver's age should not only determine the speeding behavior but also affect the likelihood that he or she gets a break from the officer who stops him/her. Old drivers tend to drive slowly while officers tend to give a break to them. Under the assumptions, substituting Equation (7), we have:

$$\begin{aligned}
y^*_{ij} &= \beta_0 - \beta_1(\alpha_0 + c_i\alpha_1 + \Lambda_i\alpha_2 + \alpha_3 1[i=a] + \alpha_4 1[i=h] + \omega_i) \\
&\quad + c_i\beta_2 + C_j\beta_3 + \Lambda_{ij}\beta_4 + \beta_5 1[i=a] + \beta_6 1[i=h] + \beta_7 1[j=A] + \beta_8 1[j=H] \\
&\quad + d_1 Mismatch + d_2 Minority \times Minority + d_3 1[i=a, j=W] + d_4 1[i=w, j=A] - \varepsilon_{ij} \\
\\
&= (\beta_0 - \beta_1\alpha_0) + c_i(\beta_2 - \beta_1\alpha_1) + C_j\beta_3 + \Lambda_{ij}(\beta_4 - \beta_1\alpha_2) \\
&\quad + (\beta_5 - \beta_1\alpha_3)1[i=a] + (\beta_6 - \beta_1\alpha_4)1[i=h] + \beta_7 1[j=A] + \beta_8 1[j=H] \\
&\quad + d_1 Mismatch + d_2 Minority \times Minority + d_3 1[i=a, j=W] + d_4 1[i=w, j=A] - (\varepsilon_{ij} + \beta_1\omega_i)
\end{aligned} \tag{11}$$

It is obvious in Equation (11) that we cannot identify the coefficients for officers' statistical discrimination or monolithic preferences because drivers' race also possibly matters in their speeding behavior. On the other hand, we can still identify officers' race-specific strictness and, more importantly, can test for different types of racial bias: own-race preferences, minority-on-minority bias, and African-American-white confrontation.

We do not directly observe officers' choice of whether or not to give a break. Thus, we use a proxy variable, an indicator as to whether a motorist is cited for driving exactly 10 m.p.h. above the limit. The use of the proxy variable is well rationalized from our discussion in previous sections. The variable is, however, subject to misclassification error. For example, those drivers whose actual speed is 10 m.p.h. above the limit and whose speed gets cited at exactly 10 m.p.h. above the limit are classified as those who are favored by officers although they actually are not.[21] Misclassification error is likely to lead to attenuation bias.

As mentioned before, we restrict our sample to tickets between 10 and 14 m.p.h. over the speed limit. Due to the massive spike at 10 m.p.h. above the limit, the restricted sample includes the majority of tickets (55%) in our whole sample. This sample restriction will reduce drivers' heterogeneity. Since our purpose is to identify officers' discretionary behavior and racial bias as distinctly as possible, we want to minimize drivers' heterogeneity and remove the potentially

[21] There is another misclassified group in which drivers get cited above 10 m.p.h. over the limit while they drive faster than that. We believe that there are only a negligible number of tickets in that category because officers would rather cite 10 m.p.h. above the limit (than 11 or 12) once they have decided to give a break.

confounding effects. It is also reasonable to assume that officers are less likely to give a break (or likely to give a smaller break) to motorists driving 15 m.p.h. or faster above the limit (recall our assumption in the theory section that "the reported speed will be increasing in s_i beyond a certain speed" is to that effect). Slow speeders driving below 10 m.p.h. over the limit are either not even stopped or strictly given a ticket since they are more likely to violate the speed limit in a restricted zone (such as school zone). In either case, officers' manipulation of speed on tickets should be expected to be insignificant. The range of speed between 10 and 14 m.p.h. over the limit seems to be appropriate for the study of leniency in terms of speed discounting.

The restriction potentially raises sample selection bias because we drop those tickets whose actual speed is between 10 and 14 m.p.h. above the limit but it is reported below 10 m.p.h. The bias is likely to be ignorable given that there are very few tickets cited under 10 m.p.h. over the limit; in addition, - as mentioned before - citing a speed at or below 10 m.p.h. above the speed limit does not matter for the motorist or the officer in Massachusetts in any way. Finally, although there are also some observable spikes at multiples of 5 above 14 m.p.h. over the limit, the magnitude of underreporting should be weak for drivers who are driving above 15 m.p.h. over the limit.

7 Empirical Findings

Table 5 presents our main results from the Probit model.[22] The dependent variable is the indicator for whether drivers get cited for driving exactly at 10 m.p.h. over the limit. We include all three sets of variables for the racial bias, *Mismatch*, $Minority \times Minority$ ($1[i \neq w]1[j \neq W]$), $1[i = a, j = W]$, and $1[i = w, j = A]$. In Column (1), as discussed before, we use our sample of tickets between 10 and 14 m.p.h. over the limit. In Column (2), we restrict the sample further to those between 10 and 11 m.p.h. above the limit. In this extremely limited sample, it seems to be true that whether they drive at 10 or 11 m.p.h. above the limit is not systematically determined by their characteristics. Therefore, any significant effect we find in this sample should be attributed to officers' manipulation - due to discretionary or racial reasons. The results are strikingly similar between the two samples.

[22] An alternative specification is a zero-inflated Poisson model, which can allow for different data-generating processes, one for exactly 10 m.p.h. over the limit and another for higher speeds. The results are similar.

7.1 Driver Characteristics

As determinants of whether drivers get cited for driving exactly 10 m.p.h. over the limit, we include driver characteristics such as age, gender, whether the driver resides in the same town where he/she is stopped, whether he/she resides in Massachusetts, and whether he/she holds a commercial driver's license. We find that no such driver characteristic variable is significant, except for race. As shown in Equation (11), we cannot differentiate the direct influence of any driver characteristic on the officer's ticketing decision from its effect on the actual speed and its subsequent impact on the officer's decision. The two effects might be opposite, which is perhaps why our driver characteristics turn out to be insignificant. In addition, the insignificance of driver characteristics variables might result from our limited range of speed. It is not surprising that drivers' characteristics do not determine the speeding behavior in a range between 10 and 14 m.p.h. above the limit (needless to say, the range between 10 and 11 m.p.h. above the limit).

One exception is the number of violations which significantly decreases the likelihood in which drivers get cited for driving exactly at 10 m.p.h. over the limit. First, officers could be tougher on those speeders who also violates the baby car seat rule. This variable may be also related to the driver's behavior. The drivers without appropriate registration would refrain from speeding. On the contrary, it might be equally true that drunken drivers would be more likely to speed due to the influence of alcohol.

7.2 Officer Characteristics

Unlike driver characteristics, most of officer characteristics variables are significant in explaining why some drivers get cited at 10 m.p.h. above the limit while others do not. First of all, we find that male officers are significantly more likely to issue tickets at 10 m.p.h. above the limit. This suggests that male officers are more lenient than female officers. At the sample averages of the other variables, male officers are 33% (or 19% in Column (2)) more likely to give a break than female officers. Second, less experienced or young officers are more likely to give a break to drivers than experienced officers. It seems reasonable that new officers are more lenient since - especially during this crucial learning-by-doing process of theirs in which they pay attention to all aspects of becoming a full-fledged officer - they may be more vulnerable to many things including drivers' complaints. Also they might be not as deft at handling speeders as their seniors.

7.3 Interactions Between Driver Races and Officer Races

Like other driver characteristics, we find no significant effect for driver race. However, keep in mind that we cannot differentiate between the direct effect of driver race on the officer's ticketing decision or its effect on the actual speed and its subsequent impact on the officer's decision. On the other hand, officer race variables are significant. We find that relative to white officers, African-American and Hispanic officers are significantly less likely to give a break. On average, holding other things constant, African-American (or Hispanic officers) are about 13% (or 9%) less likely to give a break.

The variable *Mismatch* is significant when it is included alone. However, as soon as we include *Minority*×*Minority*, *Mismatch* turns out to be insignificant. This means that the significant result for *Mismatch* is driven by minority officers being tough on minority drivers. We find that *Minority*×*Minority* is significant. Minority officers are 16% (or 8% in Column (2)) less likely to be lenient to minority drivers. We additionally find that African-American officers are less lenient to white drivers. Combined with the previous finding that minority officers are tougher on minority drivers, the last finding shows that African-American officers are stricter to all races of drivers than other officers while they are slightly more stricter to minority drivers.

7.4 Validity Check of Identification Strategy and Robustness Check

In Table 6, in Column (1), we restrict our sample to those between 11 and 14 m.p.h. above the limit. The dependent variable is the indicator for whether drivers get cited for driving exactly at 11 m.p.h. above the limit. We call this kind of dependent variable "fake" dependent variable since it seems unlikely from the data that officers cite drivers at 11 m.p.h. rather than 12, 13, or 14 m.p.h. above the limit. In fact, an officer might look too meticulous when he decides to give a break to a driver and lowers the reported speed from 14 to 11 m.p.h. over the limit. The purpose of this subsection is to check the validity of our identification strategy of using 10 m.p.h. over the limit as a proxy variable for officers' leniency.

The results support our empirical strategy. Using fake dependent variables, all the racial bias variables except one become insignificant. Only the variable for African-American officers is significant and positive. This is in fact the indirect consequence of the racial bias found at 10 m.p.h. above the limit. Recall that in Table 5 we find that African-American officers are the least lenient. As white and Hispanic officers discount tickets that should be cited at a speed higher than 10 m.p.h. over the limit and less than 15 m.p.h. over the limit and, in fact, cite 10 m.p.h. above the limit, it is a natural consequence of it to find out that there are relatively more African-American officers citing speeds at 11 and 12 m.p.h. over the limit.

Also note that the effect of speed limit is significantly positive in Table 6 while it is significantly negative in Table 5. Given that our results in Table 6 are driven by drivers' behavior rather than officers' behavior, this contrasting finding suggests that the negative effect we found in Table 5 is a result of officers' discretionary behavior of speed reporting. We conclude that holding the actual speed constant, officers are less likely to give a break to those who speed in a high speed zone. This has a reasonable explanation: the officers in general deem that high speed in itself is a dangerous and risky act and should be curbed more as the speed the motorists are allowed to travel at increases.

Likewise we can also check if our findings in Table 5 regarding driver characteristics and time and location variables result from drivers' speeding behavior or officers' discretionary manipulation of the speed. First, in Table 6, there is no significant effect of the number of violations, which means that what we found in Table 5 is the consequence of officers' discretionary behavior. Second, with the fake dependent variables, we find no systematic effect of dates or time of the day. When significant, they are positive. On the contrary, we found in Table 5 the negative effects of those variables. This also suggests that officers do manipulate the reported speed differently according to time and date.

For further robustness check, in Table 7, in Columns (1) and (2), we restrict our sample to those speeding tickets issued while there was no vehicle search. One might think that officers would behave differently when drivers look suspicious. In addition, on the driver's side, speeding behavior and criminal activities could be correlated. The results change little. There is a significant effect of $Minority \times Minority$.

In Columns (3) and (4), we disaggregate the minority group since one may think that African-American and Hispanic officers have different incentives and preferences. We include four interaction terms between African-American/Hispanic officers and drivers instead of $Minority \times Minority$. Unfortunately, we cannot test for own-race preferences and African-American and white reciprocal bias because we can only estimate up to four parameters. This should not be a major issue since the variables except $1[i = w, j = A]$ were insignificant in Table 5. The results show that African-American and Hispanic officers do have different preferences; the magnitude of bias against minority drivers is larger among Hispanic officers. However, minority officers do not differentially treat African-American and Hispanic drivers.

Finally, we check the possibility of overreporting, although we believe that overreporting should not be very common because it could irritate drivers for no explicit benefit to officers. We replicate our main analysis in Table 5 using the new sample of those tickets cited at 10 m.p.h. over the limit and below. The idea is that if our previous results were driven by officers' overreporting rather than underreporting, we should find the same and even stronger results with the correct sample and specification. The results in Column (5) show that those who are cited at 10 m.p.h. over the limit are more likely to be white drivers and cited by white officers compared to those cited below the speed. It is not surprising given the racial composition of drivers and officers over different speed levels in Tables 2 and 3. The most important finding here is that there is no evidence for racial bias; none of racial interaction terms is significant. We conclude that although we still cannot fully exclude the possibility of overreporting, if any, it should not be motivated by racial bias. In other words, overreporting cannot explain our previous finding about minority-on-minority bias.

7.5 Differences across Subsamples

Table 8 compares the results across different subsamples: by time of day in Column (1), by driver's gender in Column (2), by driver's age in Column (3), and by vehicle age in Column (4). The variable *Minority*×*Minority* is significant for both day and night. Interestingly, the effect is stronger at night. And, at night, the variable for African-American drivers turns out to be positive and significant. In combination with the stop-and-search literature's findings, this may reflect an after-the-fact leniency shown to these drivers who may initially be considered as a higher statistical criminal threat by officers, right after they stop such drivers; once officers figure out - one way or another – that many of such drivers are not carrying any contraband, the leniency may be a form of implicit reward to such non-criminal African-American drivers. Also, only at night, African-American officers are less lenient to white drivers; this may be either due to the fact that white drivers that drive at night have different characteristics that are unobservable to the econometrician - such as fewer older looking white drivers might be driving at night.

In addition, we find that minority officers are less lenient to minority drivers in the context of male drivers while there is no evidence for minority-on-minority bias in the case of female drivers. The bias exists regardless of driver age, while it does not exist for relatively new vehicles.

Table 9 shows the results for different racial neighborhoods. It may matter in what kind of neighborhood an officer stops and tickets a driver. This additional feature can easily be incorporated into our theoretical setup by employing an additional neighborhood notation, $N \in \{AA,\ WH\}$ where

AA and WH stand for predominantly African-American and white neighborhoods respectively, and so on). As mentioned in the empirical predictions section, one can imagine that an officer's integrity marginal cost is higher in a neighborhood more populated with people of his own race;

$$t_r^A(AA) > t_r^A(WH) \qquad \text{and} \qquad t_r^W(WH) > t_r^W(AA)$$

In Column (1) we focus on white neighborhoods (60% or more white) and find no evidence for the racial bias. We still find that African-American officers are the least lenient followed by Hispanic officers. In Column (2) we look at African-American neighborhoods where 20% or more population is African-American. We find that, in this area, minority officers are less lenient to minority drivers. In addition we find that African-American officers are less lenient to white drivers. Lastly, in Column (3) where we examine Hispanic neighborhoods, we find no significant effect of any race variables. The insignificance might result from small sample size.

7.6 Traffic Fines

As a supplementary study, in this subsection, we examine whether police officers also manipulate the dollar amount of a fine directly. The question is motivated by Makowsky and Stratmann (forthcoming) which shows that fines are somewhat arbitrarily determined by officers according to their own objectives or local public interests. Information on the fine amount is also available in our data since officers were required to record the exact dollar amount as well as vehicle speed.

$$Gap = \begin{cases} p_{ij} - p(\sigma_{ij}) & \text{if } p_{ij} \text{ is reported} \\ 0 & \text{if } p_{ij} \text{ is missing} \end{cases} \tag{12}$$

By using Equation (12), we compute the gap between the fine amount written on a ticket (p_{ij}) and that calculated by the formula and the reported speed - $p(\sigma_{ij}) = 75 + 10(\sigma_{ij} - 10)$ when the reported speed exceeds 10 m.p.h. above the speed limit and \$75 otherwise. Surprisingly, the fine amount is missing for about 44 percent of 25,738 tickets. Instead of dropping these observations, we assume that the amount should have been calculated by the formula and impute it by the formula and the reported speed. This assumption is reasonable in that there is no particular reason why something other than the amount implied by the formula should be imposed when a ticket with missing fine is sent to the collection office. When the fine amount is missing, by construction, the gap is zero.

Table 10 presents the distribution of the fine gap. We continue to use our restricted sample of tickets between 10 and 14 m.p.h. above the speed limit.[23] There are three noteworthy things in Table 10. First, it is notable that the gap is zero for 89 percent of tickets. This means that officers rarely lower fines and tend to apply the formula strictly. This is in contrast with our previous finding that officers are lenient in terms of speed reporting. This makes sense in that officers do not want to look "inconsistent" by reporting a penalty different from what the reported speed implies. Second, there is an accumulation (3%) of tickets at $25. This is probably due to the fact that officers omit the surcharge of $25 for the Head Injury Trust Fund. In this case, we should not interpret the gap as evidence that officers lower fines. Third, a tiny number of tickets have negative gaps; officers impose larger fines than the recommended fine according to the statute. This is possible in special zones like construction site.

Table 11 presents the results from Tobit models where the dependent variable is the fine gap. In Column (2) we exclude those tickets whose fine gap is exactly $25 because of the above mentioned concern. There are some significant estimates. First, the higher the speed limit is, the larger the fine gap is. The size of the effect is, however, small. A 10 m.p.h. increase in the speed limit on average increases the gap by 50 to 80 cents. Second, the gap is larger for younger drivers. Again the effect is not substantial. The negative effect of driver age seems to reflect the notion that officers are less lenient to prime-age adults with higher earning potential. It might indicate that officers' penalty policy is overall progressive or simply that officers have certain age preferences. Third, we find a similar negative (and small) effect for commercial drivers. This too may be due to the fact that officers may tend give a break to less experienced drivers such as non-commercial drivers. Fourth, officers' experience increases the fine gap. This is opposite to our previous finding that inexperienced officers are more lenient. A plausible explanation is that these younger and newer officers are more concerned about inconsistency between the fine and reported speed. We also find significant effects of time of the day, but the estimates between Columns (1) and (2) are too different to interpret appropriately.

We also find that minority officers give smaller fines. Compared to white officers, the gap is larger by $1.8 for African-American officers and by $4 for Hispanic officers. Like other significant estimates, the size of the effect is not substantial. Furthermore, these differences disappear after we drop those tickets with the $25 gap in Column (2). More importantly, we find that none of the

[23] In addition, we focus on tickets without additional violation other than speeding in order to impute the fine amount accurately where it is missing.

racial bias variables is significant in both samples. The conclusion is that racial bias occurs mostly when officers decide which speed to report. Once they decide on the speed, there seems no further racial consideration in deciding fines.

8 Concluding Remarks

Our theoretical section considered motorists who take into account the probability of getting ticketed and the speed that the officer will cite in deciding at what speed they will travel and officers who - net of the cost of ticketing motorists - maximize a benefit function which generically increases in the speed of ticketed drivers; this framework is general enough to allow officers to give some drivers a break by citing them at a lower speed than they were traveling.

In our empirical section, we exploit the existence of a massive accumulation of speeding tickets exactly at 10 m.p.h. over the speed limit to elicit officers' discretionary behavior and leniency. We show that the accumulation of tickets at the specific speed level is likely to result from officers' manipulation of reported speed – underreporting. Using our novel measure of officers' leniency, we find that white officers are the most lenient ones overall. Female officers are the least lenient group of officers. We find strong evidence that minority officers are less lenient to minority drivers. There is no minority-on-minority bias when vehicles are new and/or when drivers are female. The bias, on the other hand, gets stronger at night and/or in minority concentration neighborhoods. Although we find evidence of racial bias, we find no systematic racial bias in the form of mutual or monolithic racial bias. Our findings about minority-on-minority bias are interesting particularly in that AK (forthcoming), using the same data set, found evidence on own-race preferences in vehicle searches.

In the next few paragraphs we will attempt to reconcile well-established sociological - and other - perspectives and our two findings that (1) there is minority-to-minority bias and (2) female officers are relatively stricter. On the sociological front, we first note an observation by Weber (1968) that a social group's superior material resources relative to another group can give rise to the development of status beliefs favoring that group over the other. Further, a wide variety of historical contingencies can also help such status beliefs. Once such beliefs form, a member of the materially-disadvantaged group suffers a social disadvantage even vis-à-vis those members of the other group who are, in fact, his or her material equals.

In addition, recently there have been attempts to explain how bias against minorities may arise in the context of network structure of social interactions and categorizations, which influence the formation of status beliefs. Fiske (1993) has shown that people tend to more finely categorize groups who are above them in a hierarchy and more coarsely categorize groups who are below them in a hierarchy. Furthermore, types of experiences and groups of people that are less frequent in the population are more coarsely categorized and more often lumped together. As a result, this can give rise to bias against minority groups even when there is no malevolent taste for bias (Fryer and Jackson, 2008).

When, along with the above sociological perspectives, one uses further research by sociologists - as well as by anthropologists and psychologists - observing American children at play, one may be able to shed even more light on our two findings emphasized above. In that strand of research, it is found that girls tend to play in small groups (or with a single best friend) and "learn to downplay ways in which one is better than the others and to emphasize ways in which they are all the same." Boys, on the other hand, tend to play in larger groups in which they are not treated as equals. "Boys generally don't accuse one another of being bossy, because the leader is expected to tell lower-status boys what to do" (Tannen, 2001).

In the light of the contents of the last few paragraphs, one can imagine a typical police department which is populated by mostly white officers. Accordingly, in such a police department, a status- or bias-formation process in the eyes of (especially newly-hired) male-minority officers may develop more or less along the lines of the above-mentioned Weberian and network-based sociological theories - many of these officers may perceive that their chances of surviving in that department could be increased if they reflected these status-differential beliefs in their behavior, while groups of male-minority officers may also be mutually guilty of coarsely categorizing the other minority groups that they encounter less frequently and do not particularly categorize above themselves in a hierarchy. While white-male officers may not share these perspectives of the male-minority officers, they could be involved in various types of non-racial discretionary behavior simply due to their unwillingness to perceive all drivers as equals. Female officers, on the other hand, may downplay any such status inequalities and treat all drivers as equals - as explained by the above gender differences that start developing at childhood. Surely, elaborate - and inter-disciplinary - future research would be very useful in substantiating this section's attempts to reconcile our two above-mentioned empirical findings and various prominent behavioral perspectives.

References:

1. Albers, Wulf, and Gisela Albers. 1983. "On the Prominence Structure of the Decimal System," In Roland Scholz, ed., *Decision Making under Uncertainty*, Amsterdam, Holland: North-Holland, pp. 271-287.

2. Antonovics, Kate L. and Brian G. Knight. forthcoming. "A New Look at Racial Profiling: Evidence from the Boston Police Department." *Review of Economics and Statistics.*

3. Anwar, Shamena and Hanming Fang. 2006. "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence." *American Economic Review*, 96(1), pp. 127-151.

4. Blalock, Garrick, Jed DeVaro, Stephanie Leventhal, and Daniel H. Simon. 2007. "Gender Bias in Power Relationships: Evidence from Police Traffic Stops" ILR Collection Working Papers, Cornell University.

5. Close, Billy R. and Patrick L. Mason. 2006. "After the Traffic Stops: Officer Characteristics and Enforcement Actions." *Topics in Economic Analysis & Policy*, Vol. 6, Issue 1, Article 24. (http://www.bepress.com/bejeap/topics/vol6/iss1/art24)

6. Close, Billy R. and Patrick L. Mason. 2007. "Searching for Efficient Enforcement: Officer Characteristics and Racially Biased Policing." *Review of Law & Economics*, Vol. 3, No. 2, Article 5. (http://www.bepress.com/rle/vol3/iss2/art5)

7. Dedman, Bill and Francie Latour. 2003. "Speed Trap: Who Gets a Ticket and Who Gets a Break?" *Boston Globe.*

   (http://www.boston.com/globe/metro/packages/tickets/072003.shtml)

   (http://www.boston.com/globe/metro/packages/tickets/072003b.shtml)

   (http://www.boston.com/globe/metro/packages/tickets/072003c.shtml).

8. Dharmapala, Dhammika and Stephen Ross. 2004. "Racial Bias in Motor Vehicle Searches: Additional Theory and Evidence." *Contributions to Economic Analysis & Policy*, 3(1).

9. Dominitz, Jeff and John Knowles. 2005. "Crime Minimization and Racial Bias: What Can We Learn from Police Search Data." *Economic Journal*, 116(515), pp. F368-F384.

10. Farrell, J., McDevitt, A., Bailey, L., Andresen, C. and E. Pierce. 2004. "Massachusetts Racial and Gender Profiling Study," Northeastern University, *Institute on Race and Justice* (http://www.boston.com/globe/metro/packages/tickets/profiling_finalreport.pdf).

11. Fiske, S.T. 1993. "Controlling Other People: the Impact of Power on Stereotyping," *American Psychologist*, 48, 621-638.

12. Fryer, Roland and Jackson, Matthew O. (2008) "A Categorical Model of Cognition and Biased Decision Making," *The B.E. Journal of Theoretical Economics*: Vol. 8 : Issue 1 (Contributions), Article 6. Available at: http://www.bepress.com/bejte/vol8/iss1/art6.

13. Grogger, Jeffrey and Greg Ridgeway. 2006. "Testing for Racial Profiling in Traffic Stops from behind a Veil of Darkness." *Journal of the American Statistical Association*, 101(475), pp. 878-887.

14. Harris, D. A. 1999. *Driving while Black: Racial Profiling on our Nation's Highways.* New York: American Civil Liberties Union. (http://www.aclu.org/profiling/report).

15. Knowles, John, Nicola Persico, and Petra Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy*, 109(1), pp. 203-232.

16. Lamberth, J. 1996. *A Report to ACLU.* New York: American Civil Liberties Union. (http://www.aclu.org/court/lamberth.html.)

17. Makowsky, Michael and Thomas Stratmann. forthcoming. "Determinants of Traffic Citations: Political Economy at Any Speed." *American Economic Review.*

18. McConnell, E. H. and A. R. Scheidegger. 2001. "Race and Speeding Citations: Comparing Speeding Citations Issued by Air Traffic Officers with Those Issued by Ground Traffic Officers." *Presented at the annual meeting of the Academy of Criminal Justice Sciences*, Washington D.C.

19. National Highway Traffic Safety Administration. 2002. *National Survey of Speeding and Other Unsafe Driving Actions, Volume II: Driver Attitudes and Behavior*, U.S. Department of Transportation.

20. Novak, Kenneth J. 2004. "Disparity and Racial Profiling in Traffic Enforcement." Police *Quarterly*, 7(1), pp. 65-96

21. Ridgeway, Greg. 2006. "Assessing the Effect of Race Bias in Post-traffic Stop Outcomes Using Propensity Scores." *Journal of Quantitative Criminology*, 22(1), pp. 1-29.

22. State of New Jersey v. Soto. 1996. A734A. 2d 350 (N.J. Super. Ct. Law Div. 1996).

23. Tannen, Deborah. 2001. *You Just Don't Understand: Women and Men in Conversation.* New York: Quill.

24. Weber, Max. 1968. *Economy and Society.* Edited by G. Roth and C. Wittich (translated by E. Frischoff, et al.). New York: Bedminster.

A Appendix

A.1 Ruling out some simple driver benefit functions

Suppose the variable benefit function is such that the marginal benefit is declining beyond the limit. Then it is easy to observe that the driver will only drive at speeds less than 10 m.p.h. above the limit since the penalty function will be increasing beyond 10 m.p.h. above the limit. Next, consider the variable benefit function, $v(s,c) = a_0$, where $a_0 > 0$. Then note that, even if $E\gamma = 0$, (i) there will be no speeding if $\alpha > a_0$, (ii) any speed between the limit and 10 m.p.h. above the limit would be possible if $\alpha = a_0$, and (iii) the driver would never exceed a speed which is 10 m.p.h. above the limit if $\alpha < a_0$.

Thus, to generate any speeds especially beyond 10 m.p.h. above the limit, we need to assume that at least a group of drivers will have a variable benefit function with increasing marginal benefit in some range of speeds exceeding the limit by 10 m.p.h. A possible tractable - though unrealistic - candidate for such a benefit function would be the linear form

$$v(s,c) = a_0 + a_1 s, \tag{A1}$$

$a_0 > 0$ is the level of benefit at the speed limit and $a_1 > 0$ is the marginal benefit of speeding beyond the limit Note that if $a_0 > \alpha$ and $a_1 > 10$, then the optimal speed is ∞ at low levels of $E\gamma$. With $a_0 > \alpha$ and $a_1 < 10$, optimizing behavior yields

$$(1 - E\gamma)/E\gamma = a_1/10. \tag{A2}$$

Note that this condition does not allow the driver to pick a particular optimal speed. Examining this condition a little closer reveals that the driver will definitely be speeding but will be indifferent among all speeds exceeding the limit by 10 m.p.h. as long as

$$(1 - E\gamma)/E\gamma > a_1/10. \tag{A3}$$

Further, the driver is indifferent between any speed above the limit by 10 miles as well as not speeding at all as long as

$$(1 - E\gamma)/E\gamma = a_1/10, \tag{A4}$$

Thus, the driver will randomize over any speed exceeding the limit by ten miles as well as speeding at or below the limit. In addition, the marginal benefit of speeding staying constant at each speed is surely far from describing a plausible and realistic situation.

A.2 Personal variations on the driver benefit function

Assume that there are $n > 1$ different segments of motorists with different levels of benefits at speeds exceeding the speed limit. Using the following driver variable benefit function will allow such personal variations regarding the drivers' taste for speeding, where $\theta_i,\ \lambda_i > 0$ with: $v_i(s,c) = a_0 + \theta_i \tilde{s}(c,r)s - \lambda_i s^2$ such that $\theta_n > \theta_{n-1} > ... > \theta_1 > 0$ and $0 < \lambda_n < \lambda_{n-1} < ... < \lambda_1$, i= 1, 2, $\ldots$, n. Given any utility function, if the driver chooses to speed, the optimizing behavior yields speeding s^* miles above the speed limit: $s_r^*(c,z,\gamma) = \frac{1}{2\lambda_i}\left(\theta_i \tilde{s}(c,r) - k\frac{zE\gamma}{(1-zE\gamma)}\right)$. Thus, $\theta_n/\lambda_n \leq 2$ is implied so that the "maximum speed s^* chosen by a driver when $\gamma = 0$" cannot possibly exceed $\tilde{s}(c,r)$. Note that a very low θ_1 and a very high λ_1 may allow the presence of a segment of drivers who would not speed even when γ is very low (especially if a_0 is also sufficiently low or zero).

A.3 Personal variations on officers' types

Again, here too, to allow personal variations one can assume that there are $m > 1$ different segments of officers by re-scaling t_r^R via a parameter $\tau_i > 0$, where $i = 1,\ 2,\ ...,\ m$. This will allow us - among other possibilities - to consider real-life cases where different segments of officers using prominent speed-cutoff probabilities such as 10 or 15 miles above the limit before they would consider issuing a ticket, and so on.

A.4 The stopping decision

The *stopping decision* (which is unobservable to the econometrician) is typically made in the absence of any information about c as well as about r. There is the time cost of stopping a driver, , separate from the ticketing cost t. Let the set below, $(c^*, r^*|s)$, denote the set of driver types whose ticketing - when stopped at speed (s) - would yield a non-negative payoff to the officer of type (C,R).

$$(c^*, r^*|s) = \left\{(c,r) : b_r^R(\sigma^*, s, c|C) - t_r^R\left(\frac{s-\sigma^*}{s}\right) - t \geq 0\right\}.$$

That is, a police officer, upon stopping a driver, will ticket her if and only if his ticketing net benefit is non-negative. Note that (c,r) cannot be observed by the officer before he makes the stop. Let $F(c|w,s)$ and $F(c|a,s)$ be the distribution of c in the white and African-American populations, respectively, conditional on observed speed s. Let $G(r|s)$ be the probability that the driver who is stopped turns out to be of race $r$conditional on s. Let $B(c^*, r^*|C, R, s)$, denote the expected net benefit of an officer of type (C, R) from ticketing a driver of type $(c^*, r^*|s)$ - who is stopped at speed (s) - by citing her speed at (σ^*):

$$B(c^*, r^*|C, R, s) = \sum_{r=w,a} \left[\int \left(b_r^R(\sigma^*, s, c|C) - t_r^R \left(\frac{s - \sigma^*}{s} \right) \right) F(c|r, s) dc \right] G(r|s) - t,$$

such that $(c, r) \in (c^*, r^*|s)$. Once the officer observes a vehicle speeding at speed (s), he compares his stopping cost T to it to choose the probability β of stopping that motorist.

$$B(c^*, r^*|C, R, s) - T. \tag{A5}$$

Equation (A5) implies the following. If the term in (A5) is positive, the optimizing behavior implies $\beta(C, R, s) = 1$. If that term is negative, then the optimizing behavior implies $\beta(C, R, s) = 0$. If that term is zero, then the officer will be willing to randomize over whether or not to stop a motorist traveling at speed s.

T may be expected to be fully idiosyncratic. However, there is a lot of anecdotal evidence that many officers set their radar guns to beep at certain focal speed levels such as 10 m.p.h. above the speed limit. The generic explanation by Albers and Albers (1983) seems to be relevant here as well.