A New Look at Racial Profiling: Evidence from the Boston Police Department

Kate L. Antonovics and Brian G. Knight[*†]

PRELIMINARY AND INCOMPLETE

DO NOT CITE WITHOUT AUTHORS' PERMISSION

April 28, 2004

Abstract

This paper provides new evidence on the role of preference-based versus statistical discrimination in racial profiling using a unique dataset that includes the race of both the driver and the officer. We first generalize the model of Knowles, Persico and Todd (2001) and show that when the police observe a noisy signal of a motorist's guilt, the insight that allows them to empirically distinguish between preference-based and statistical discrimination disappears. However, our model also predicts that if statistical discrimination alone explains differences in the rate at which the vehicles of drivers of different races are searched, then search decisions should be independent of officer race. Consistent with preference-based discrimination, our baseline results demonstrate that the officer is more likely to conduct a search if the race of the officer differs from the race of the driver. We then investigate and rule out two alternative explanations for our findings: race-based informational asymmetries between officers and the assignment of officers to neighborhoods.

[*]Respectively, Department of Economics, University of California, San Diego, email: kantonov@ucsd.edu, and Department of Economics, Brown University, email: Brian_Knight@brown.edu.

[†] We thank Peter Arcidiacono, Eli Berman, Richard Carson, Kim Sau Chung, Roger Gordon and Nora Gordon for their comments. We have also benefitted from discussions with Amy Farrall at Northeastern University and Carl Walter at the Boston Police Department. Finally, we are indebted to Bill Dedman at *The Boston Globe* for providing us with our data.

A New Look at Racial Profiling: Evidence from the Boston Police Department

To date, there have been over 200 court cases involving allegations of racial and ethnic profiling against law enforcement agencies in the United States. Typically, the focus in these cases has been on uncovering why law enforcement officials treat individuals from different racial groups differently. On the one hand, the courts have tended to uphold racially biased policing patterns when they can be reasonably justified by racial differences in crime rates. On the other hand, the courts have consistently ruled against what appear to be purely racist policing practices. The problem, of course, is that it is not easy to empirically distinguish between these two possibilities.

Economists, who have long struggled with explaining racial disparities in labor market outcomes, have now joined the debate over racial profiling, and a number of recent papers have attempted to determine whether the observed racial disparities in policing patterns are best explained by models of statistical discrimination or models of preference-based discrimination (see, for example, Knowles, Persico and Todd (2001) and Hernández-Murillo and Knowles (2003)).

In models of statistical discrimination, discrimination arises because law enforcement officials are uncertain about whether a suspect has committed a particular crime. Thus, if there are racial differences in the propensity to commit that crime, then the police may rationally treat individuals from different racial groups differently. On the other hand, in models of preference-based discrimination, discrimination arises because the police have discriminatory preferences against members of a particular group and act as if there is some non-monetary benefit associated with arresting or detaining members of that group. Thus, preference-based discrimination refers to anything that raises the benefit (or, equivalently, lowers the cost) of searching motorists from one group relative to those from some other group.

This debate among economists over the sources of racial disparities in policing patterns roughly parallels the debate over racial profiling within the court system. That is, statistical discrimination approximately corresponds to the type of behavior that the courts have tended to uphold, while preference-based discrimination approximately corresponds to the type of behavior that the courts have tended to condemn. For this reason, economic theory and economic analysis may lead to insights that are useful in litigating these hotly-contested court cases.

In this paper, we attempt to understand the reasons for observed racial differences in the rate at which the vehicles of African-American, Hispanic and white motorists are searched during traffic stops. In doing so, we contribute to the literature on racial profiling in a number of ways. From a theoretical standpoint, we show that if the information structure of the canonical model of police search behavior developed in Knowles, Persico and Todd (2001) (hereafter, often, KPT) is generalized, then the fundamental insight that allows KPT to empirically distinguish between preference-based discrimination and statistical discrimination falls away. Our modified model, however, provides an alternative mechanism for testing between these two forms of discrimination; in particular, our model predicts that if statistical discrimination alone explains differences in the rate at which African-American and white drivers are pulled over, then search decisions and search outcomes should be independent of the race of the police officer.

We then test these theoretical predictions using a unique data set in which we are able to match the race of the officer to the race of the driver for every traffic stop made by officers in the Boston Police Department for the two-year period starting in April 2001. Thus, in addition to being able to discern differences in the likelihood that motorists from different racial groups are subject to search, we are also able to determine whether these patterns differ depending on the race of the officer. Previous studies of racial profiling lacked the officer-level data required for this type of analysis.

We find that, even after controlling for a broad set of covariates including the location of the stop, if the race of the officer differs from the race of the driver, then the officer is more likely to conduct a search. These results cannot be explained by standard models of statistical discrimination and are consistent with preference-based discrimination. We then investigate two alternative explanations for these empirical findings. First, we examine whether these patterns could arise because of differences in the ability of African-American and white police officers to accurately assess the guilt of motorists from different racial groups. Although our model delivers ambiguous predictions about the effect of this type of informational asymmetry, we find that our results hold even among officers with greater than 10 years of experience, amongst whom informational asymmetries should be less severe. Finally, we investigate (and rule out) the possibility that our findings could be explained by the way in which officers are assigned to various neighborhoods within the city.

Some Initial Trends in the Data

In order to motivate our model and the analysis that follows, it is worthwhile to first highlight a few patterns in our data. For now, these patterns are merely meant to be suggestive, and we will discuss the data in greater detail below.

Table 1 presents, by officer race and motorist race, the probability that a motorist's car is searched during a traffic stop. Looking at the last column, we see that both Hispanics and blacks are almost twice as likely as are whites to have their cars searched. This differential search pattern could be the result of preference-based discrimination. However, it is also consistent with statistical discrimination. That is, if blacks and Hispanics are more likely to carry drugs or other contraband than are whites, then it's also possible that they are also more likely than whites to raise the suspicion of the police. Thus, the last column of Table 1 simply reiterates the well-known fact that racial disparities in search rates exist, but does not offer any insight into why those disparities might arise.

Columns 2-4, however, are more revealing; motorists are, in general, more likely to be searched if the officer making the stop is from a different racial group from that of the motorist. For example, the probability that a white motorist is searched is .41% if the officer is white and .67% if the officer is black. Similarly, the probability that a black motorist is searched is .81% if the officer is black but 1.0% if the officer is white. In order to insure that the patterns in Table 1 are not driven by a small number of officers who issue an unusually large number of tickets, Table 2 weights each citation by the inverse of the number of citations given by the officer issuing the citation. Since officers who issue a large number of tickets are less likely to conduct searches than officers who issue a small number of tickets, the search probabilities are generally larger in Table 2 than in Table 1. However, as in Table 1, we see that motorists are consistently less likely to be searched if the officer making the stop is a member of the motorist's own racial group.

Abstracting at this stage from issues of statistical significance and other possible concerns, we merely wish to point out that the patterns in Tables 1 and 2 are inconsistent with standard models of statistical discrimination in which racial differences in the rate at which motorists are searched arise because the police believe that motorists from some racial groups are more likely to have contraband than are motorists from other groups. Since these beliefs must be correct in equilibrium, there should be no difference in the rate at which officers from different racial groups search the vehicles of motorists from a particular racial group.

On the other hand, preference-based discrimination could explain these patterns. In particular, if officers favor members of their own racial group, then we would expect search rates to be lower when there is a match between the race of the officer and the race of the motorist.

However, two alternative explanations also come to mind. First, there may be racial differences in the ability of officers to accurately discern the likelihood that motorists from different racial groups are guilty. For example, it is natural to think that officers may be better able to assess the guilt of motorists who are members of their own racial group. A second explanation for the differential search rates in Tables 1 and 2 involves the mechanism through which officers are assigned to various neighborhoods within the city. For example, if white officers are assigned to neighborhoods in which crimes are more likely to be committed by blacks than whites, and if black officers are assigned to neighborhoods in which crimes are more likely to be committed by whites than blacks, then we might expect that, for the city as a whole, white officers would be more likely than black officers to search the cars of black motorists. We address both of these alternative explanations in the final sections of the paper.

Model

In this section, we generalize the information structure of the model of police search presented in Knowles, Persico and Todd (2001), and show that once this more general information structure is adopted, the fundamental insight that allows KPT to empirically distinguish between preference-based discrimination and statistical discrimination falls away.

In their model, the police decide whether or not to search motorists, motorists decide whether or not to carry drugs, and the optimal actions of both parties depend upon the behavior of the other. As in KPT, we assume that, when making their decisions, the police are able to observe each motorist's race and a characteristic that is related to the motorist's propensity to carry contraband. However, in contrast to KPT, we assume that this characteristic (or signal) is not known to the motorist at the time he or she decides whether or not to carry drugs. For example, even though motorists who traffic drugs may be more likely to appear nervous than those who do not, and even though the police observe how nervous motorists appear, motorists may not be able to perfectly predict their outward appearances in the event that they are pulled over. In this setting, the police will use these behavioral cues as signals of the motorist's guilt, and, more importantly, the police will weigh these signals against their prior beliefs about the likelihood that the motorist is carrying contraband.

To ease comparison, we adopt the notation in KPT whenever possible. We assume that both police officers and drivers are either African-American or white, which we denote a and w, respectively. In what follows, superscripts refer to the race of the officer and subscripts refer to the race of the motorist.

Motorists

In deciding whether or not to carry drugs, motorists weigh the benefit of carrying drugs against the penalty of being caught. If a driver does not carry drugs, then his payoff is assumed to be zero regardless of whether or not his car is searched. However, if a motorist from group r carries drugs, then he faces cost $-j_r$ if his car is searched and benefit ν_r if his car is not searched. In addition, we allow drivers to be heterogeneous in their preferences for carrying drugs. To capture this heterogeneity we assume that drivers who carry drugs face an additional cost, ϵ, that does not depend on whether the motorist's car is searched. Let $H(\cdot|r)$ denote the cdf of ϵ for drivers from group r. Differences in the distribution of ϵ across racial groups may stem, for example, from racial differences in labor market opportunities, where better labor market opportunities translate into higher values of ϵ.

Given that the payoff to carrying drugs depends on whether a motorist's car is searched, motorists must consider the likelihood that they will be searched. If the motorist decides to carry drugs (if he is guilty), then the police observe the signal $\theta \in [0,1]$ drawn from the pdf f_G (cdf F_G), and if the motorist does not carry drugs (if he is innocent), then the police observe the signal $\theta \in [0,1]$ drawn from the pdf f_N (cdf F_N). It is assumed that f_G and f_N satisfy the strict monotone likelihood ratio property so that $\rho(\theta) = \frac{f_N(\theta)}{f_G(\theta)}$ is strictly decreasing in θ. This assumption implies that higher values of θ are more likely if the driver carries drugs. Conceptually, θ corresponds to characteristics that are correlated with the likelihood that a motorist carries drugs, but that are not perfectly known to the motorist at the time he or she decides whether to carry drugs. For the time being, we ignore the possibility that θ may provide a more accurate signal to police officers from certain racial groups.

Let γ_r^a and γ_r^w denote the probability that African-American and white officers, respectively, search motorists from group r. Then, if the fraction of African-American officers is given by q, the probability that a motorist from group r is searched is given by

$$\gamma_r = q\gamma_r^a + (1-q)\gamma_r^w,$$

and the expected payoff to carrying drugs for motorists from group r is given by

$$-\gamma_r j_r + (1-\gamma_r)\nu_r - \epsilon.$$

Police

The police cannot perfectly observe whether a driver is carrying drugs. However, officers do observe each driver's race, r, and the noisy signal, θ. In addition, the police have prior beliefs about the likelihood that drivers from group r carry drugs. Let these beliefs be denoted by

π_r. Based on θ and π_r, the police form posterior beliefs about the likelihood that drivers from group r are carrying drugs or other contraband. Let G denote the event that a driver who is searched is caught carrying drugs and let $P(G|r, \theta)$ denote officers' posterior beliefs about the likelihood that drivers from group r with signal θ will be found carrying drugs, where

$$P(G|r, \theta) = \frac{\pi_r f_G(\theta)}{\pi_r f_G(\theta) + (1 - \pi_r) f_N(\theta)}.$$

It is assumed that when deciding whether or not to search a car, police seek to maximize the number of convictions less the cost of search. Let t_r^i be the cost to officers from group i of searching drivers from group r. To make the model viable, we assume that $0 < t_r^i < 1$. Thus, the expected payoff to officers from group i of searching motorists from group r with signal θ is given by:

$$P(G|r, \theta) - t_r^i.$$

Equilibrium

The timing of the game is as follow:

Stage 1: Drivers observe ϵ and decide whether to carry contraband. Thus, a best response function for motorists from group r is a decision rule $c_r : \Re \rightarrow \{0, 1\}$.

Stage 2: Motorists are randomly pulled over. Police observe the motorist's race and θ and decide whether or not to search the motorist's car. If the driver is guilty of carrying contraband, θ is drawn from the pdf f_G, otherwise θ is drawn from the pdf f_N. A best response function for officers from group i who have pulled over a motorist from group r is a decision rule $s_r^i : [0, 1] \rightarrow \{0, 1\}$.

Stage 3: Payoffs happen.

We now describe the equilibrium of the above model, and start by examining the best response function of police officers. Officers from group i will optimally search a motorist from group r with signal θ if $P(G|r, \theta) \geq t_r^i$. Since higher values of θ are more likely if the driver is guilty, this implies a cutoff value of θ above which police search and below which they do not. Let the cutoff value used by officers from group i for drivers from group r be denoted $\tilde{\theta}_r^i = \tilde{\theta}(\pi_r, t_r^i)$ where

$$\tilde{\theta}(\pi_r, t_r^i) = \rho^{-1} \left[\left(\frac{\pi_r}{1 - \pi_r} \right) \left(\frac{1 - t_r^i}{t_r^i} \right) \right] \tag{1}$$

and where $\rho(\theta) = \frac{f_N(\theta)}{f_G(\theta)}$. Given $\rho(\theta)$ is decreasing in θ, it is easy to show that $\tilde{\theta}(\pi_r, t_r^i)$ is decreasing in π_r and increasing in t_r^i. Thus, the higher is a police officer's prior belief and the lower is the cost of search, the lower will be the cutoff value of θ (and the more likely the police will be to search drivers from that group).

Drivers optimally carry drugs if the expected net benefit of carrying drugs is greater than zero. Thus, motorists from group r optimally carry drugs if and only if $-\gamma_r j_r + (1-\gamma_r)\nu_r \geq \epsilon$, so the probability that they traffic drugs is given by

$$H(\gamma_r|r) = H(-\gamma_r j_r + (1 - \gamma_r)\nu_r|r) \tag{2}$$

Note that drivers with values of θ that exceed $\tilde{\theta}_r^i$ and who are pulled over by officers from group i will experience ex post regret if they are carrying drugs, even though their decision was optimal ex ante.

A Bayesian Nash equilibrium of this model occurs where, based on ϵ, motorists are playing a best response to the search behavior of the police and where the police are playing a best response to the distribution of drug trafficking strategies played by motorists. This occurs at any belief π_r^* such that African-American and white officers, respectively, set $\tilde{\theta}_r^a$ and $\tilde{\theta}_r^w$ such that, in response to the implied search probabilities, motorists from group r carry drugs at the exact rate postulated by the police.

To see this more precisely, note that if motorists from group r traffic drugs at the exact rate postulated by the police, then the distribution of θ for motorists from group r is given by

$$F(\theta|\pi_r) = \pi_r F_G(\theta) + (1 - \pi_r)F_N(\theta).$$

Thus, given the optimal cutoff search rule for police, the probability that motorists from group r are searched by officers from group i is given by

$$\gamma^i(\pi_r) = 1 - F(\tilde{\theta}(\pi_r, t_r^i)|\pi_r). \tag{3}$$

and the probability that motorists from group r are searched by officers in general is given by

$$\gamma(\pi_r) = q\gamma^a(\pi_r) + (1 - q)\gamma^w(\pi_r) \tag{4}$$

where, as before, q is the fraction of African-American police officers. Thus, combining equations (2) and (4), we see that a Bayesian Nash equilibrium is any π_r^* such that

$$\pi_r^* = H(\gamma(\pi_r^*)|r).$$

Discrimination and Equilibrium

Given our interest in testing our model's predictions, we focus on examining how both statistical discrimination and preference-based discrimination affect the model's observable implications. In particular, we focus on how statistical discrimination and preference-based discrimination affect the probability of guilt conditional on search, $P(G|r, \theta > \tilde{\theta}_r^i)$, and the probability of search, $\gamma^i(\pi_r)$.

Given π_r^*, the equilibrium probability that officers from group i find that motorists from group r are carrying drugs conditional on search is given by:

$$P(G|r, \theta > \tilde{\theta}(\pi_r^*, t_r^i)) = \frac{\pi_r^*(1 - F_G(\tilde{\theta}(\pi_r^*, t_r^i)))}{\pi_r^*(1 - F_G(\tilde{\theta}(\pi_r^*, t_r^i))) + (1 - \pi_r^*)(1 - F_N(\tilde{\theta}(\pi_r^*, t_r^i)))} \quad (5)$$

It should be clear from above that $P(G|\pi_r^*, \theta > \tilde{\theta}(\pi_r^*, t_r^i)) > \pi_r^*$. That is, conditional on being searched, the probability that drivers from group r are found carrying drugs is higher than the proportion of drivers from that group who carry drugs overall. This makes sense; the police only search drivers who have a relatively high likelihood of being guilty.

Statistical Discrimination

In this model, statistical discrimination arises if there are equilibrium differences in officers' prior beliefs about the likelihood that drivers from different racial groups carry drugs, so that $\pi_a^* \neq \pi_w^*$. Since these prior beliefs must be self-confirming in equilibrium, this can only occur if the cost of carrying drugs differs for African-Americans and whites. For example, Figure 1 depicts the equilibrium of this model under the assumption that $H(\cdot|a)$ stochastically dominates $H(\cdot|w)$ so that the mean value of ϵ is lower for African-American than whites. This might occur, for example, if the opportunity cost of carrying drugs in terms of forgone labor market earnings is lower for African-Americans than whites. On the horizontal axis we plot π_r and on the horizontal axis we plot γ_r^i. As the figure reveals, for every search probability, γ, African-American motorists are more likely to traffic drugs than white motorists. As a result, the equilibrium probability that African-Americans carry drugs is higher than it is for whites ($\pi_a^* > \pi_w^*$), and the police are more likely to search African-American motorists than white motorists ($\gamma_a^* > \gamma_w^*$).

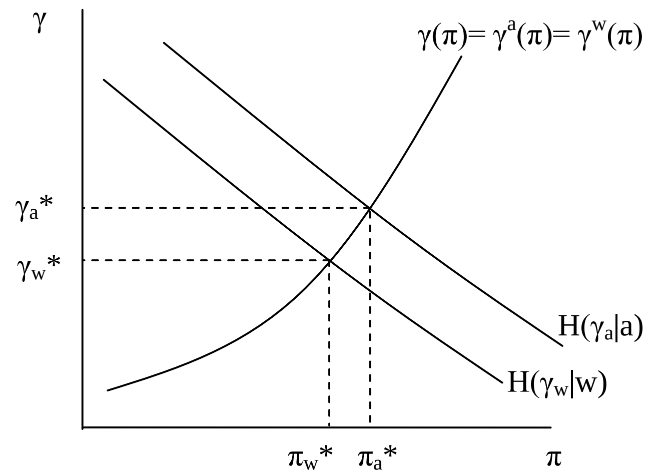

Figure 1: Equilibrium with Statistical Discrimination

However, note that Figure 1 does *not* imply that there will be differences in the rate at which African-American and white officers search motorists from group r. That is, as long as $t_r^a = t_r^w$, then it is clear from equation (3) that in equilibrium $\gamma_r^{a*} = \gamma_r^{w*}$.

We can also examine the effect of statistical discrimination on $P(G|r, \theta > \tilde{\theta_r^i})$, the probability of guilt conditional on search. It should be clear from equation (5) that even if the cost of search is the same across racial groups, if $\pi_a^* \neq \pi_w^*$, then there is no reason to expect that $P(G|a, \theta > \tilde{\theta}(\pi_a^*, t)) = P(G|w, \theta > \tilde{\theta}(\pi_w^*, t))$.

This result lies in sharp contrast to the central prediction of the model presented in KPT. In that model, if statistical discrimination alone explains the differential search rates of white and African-American motorists, then the probability of guilt conditional on search should be the same for motorists from all racial groups. The intuition for this prediction is that officers will search motorists until the marginal benefit of searching is equal to the marginal cost. Since the benefit of searching African-American and white motorists is exactly equal to the probability of guilt, then, as long as the marginal cost of search is the same across racial groups, so too should be the probability of guilt. That is, *all* motorists, regardless of their observable characteristics, will traffic drugs at the exact same rate. This prediction is powerful because the econometrician need not observe the same set of characteristics that is observed by the police in order to test whether it is true.

The crucial difference between our model and that in KPT is that we allow the police to observe a characteristic that is related to a motorist's propensity to carry contraband, *but that is not perfectly known to the driver at the time he or she decides to traffic drugs.*

As mentioned earlier, a prime example of one such characteristic is nervous behavior by the motorist. In KPT, in contrast, when motorists make their decision about whether or not to traffic drugs, they are able to perfectly predict the information that officers will be able to observe. It turns out that this seemingly subtle distinction greatly affects the equilibrium outcome. In particular, in our model, motorists differ in the likelihood that they traffic drugs. That is, θ is not known to motorists at the time they make their decision about whether or not to traffic drugs. Thus, it serves as an additional piece of information that the police can use in determining whether or not to search, and the police use this piece of information to update (in a Bayesian fashion) their prior beliefs about the likelihood that the motorist traffics drugs. Thus, if the police have different prior beliefs about the likelihood that drivers from different racial groups are trafficking drugs, then the police will treat drivers with the same θ but who are from different racial groups differently.

Note that in our model the cutoff value of θ, $\tilde{\theta}_r$ is exactly that value of θ such that $P(G|r, \tilde{\theta}_r) = t$. Thus, if the cost of search is the same across drivers from different racial groups, then the "marginal" motorist from each group will be equally likely to be found trafficking drugs if his or her car is searched. The problem, of course, is that, since we do not observe θ, we do not observe these "marginal" motorists, and so this prediction is impossible to test. In addition, there is no reason to think that the probability of guilt conditional on search will be the same for the *average* motorist who is searched. That is, as discussed above, if $\pi_a^* \neq \pi_w^*$, then there is no reason to expect that $P(G|a, \theta > \tilde{\theta}(\pi_a^*, t)) = P(G|w, \theta > \tilde{\theta}(\pi_w^*, t))$. Thus, once the information structure of KPT is expanded to allow the police to observe a noisy signal of a motorist's guilt, then, even in the absence of preference-based discrimination, there is no reason to expect the probability of guilt conditional on search to be the same across racial groups.

Preference-Based Discrimination

The figure below depicts an equilibrium for drivers from group r under the assumption that white police officers have discriminatory preferences against drivers from group r. That is, we assume that white officers find it less costly than African-American police officers to search drivers from group r so that $t_r^w < t_r^a$. As the figure reveals, regardless of the likelihood that drivers from group r carry drugs, white police officers are more likely to search drivers from group r than are African-American officers. As one would expect, this implies that the equilibrium outcome depends on the proportion of police officers who are white. In this case, the larger the proportion of white police officers, the more likely are drivers from group r to be searched and the less likely are they to carry drugs.

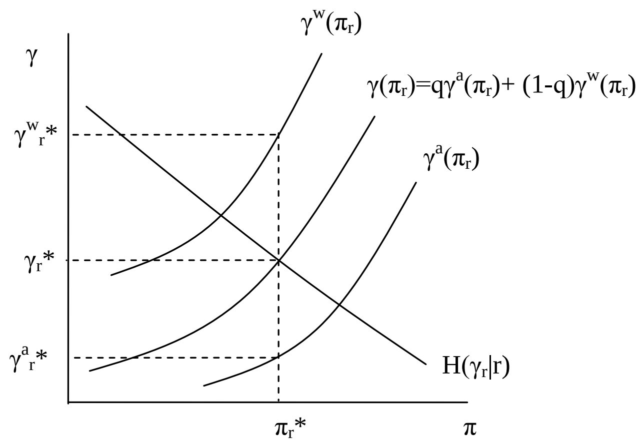

Figure 2: Equilibrium with Preference-Based Discrimination

In the model above, it is assumed that officers do not differ in their ability to accurately assess the likelihood that motorists from different racial groups are carrying drugs and that officers are randomly matched with drivers. Provided these assumptions hold, our model predicts that, in the absence of preference-based discrimination, the probability that officers from different racial groups search motorists from any one racial group should be the same. The next section attempts to test these predictions using data from Boston on both the race of the officer and the race of the driver.

Data

In July 2000, the Massachusetts legislature passed Chapter 228 of the Acts of 2000, *An Act Providing for the Collection of Data Relative to Traffic Stops*. Among other things, this statute requires that, effective April 1, 2001, the Registry of Motor Vehicles collect data on the identifying characteristics of all individuals who receive a citation or who are arrested. The data collected by the State contain a wide variety of information including: the age, race and gender of the driver, the year, make and model of the car, the time, date and location of the stop, the alleged traffic infraction, whether a search was initiated and whether the stop resulted in an arrest.

The statute also required the Registry of Motor Vehicles to collect data on warnings. However, citing budgetary shortfalls, the Registry only compiled data on warnings for two months. Thus, for most of the time period under investigation, we do not observe stops for which an officer issued either a written or a verbal warning. That is, unless an officer issued a citation, the stop does not appear in our data outside of the two-month period. If officers favor members of their own racial group, then we might expect officers to issue citations to members of their own racial group only if they have committed relatively serious traffic infractions or if the officer strongly suspects that the driver is trafficking drugs. If so, then our estimates will tend to *understate* the extent of the racial bias in search patterns. We will address this data limitation later in the analysis by restricting our sample to the two-month period that includes data on warnings.

In addition to the citation-level data collected by the State, we were also able to obtain officer-level data from the Boston Police Department. These data contain, among other things, information on the officer's race, gender, rank and number of years on the force. For the subset of citations issued by officers in the Boston Police Department, we are then able to match the officer-level data to the citation-level data collected by the state. In total, we are able to match officer-level data to over 112,473 citations issued by 1,369 officers, representing just over 80% of the citations issued by officers in the Boston Police Department in our data. That is, for approximately 20% of the citations issued by an officer in the Boston Police Department in our data, we were unable to identify the officer who issued the citation.

We restrict our sample in a number of ways. First, we drop a small number of citations (6) for which contradictory race information was recorded. In addition, we drop citations issued by Asian officers (23 officers in total), and 8,051 citations issued to Asian, Native American and Middle Eastern motorists. As a result, all of the motorists and officers in our data are either black, white or Hispanic. Finally, we drop a small number of citations (10) that were issued to motorists outside the City of Boston in one of the surrounding suburbs. This may have happened, for example, if an officer followed a speeding driver outside of the City limits. Once these restrictions have been made we are left with 100,408 citations issued by 1,335 officers.

Of considerable concern is the fact that the search variable is missing for over 18% of the citations in our data. When filling out a citation, officers are required to check either "yes" or "no" to indicate whether a search was conducted. If an officer neglected to check either box, then the search variable is missing in our data. We do not know why officers failed to check this box. One possibility is that they were careless. Another is that they did not fully understand how to fill out the citation and generally only checked the "yes" box if they conducted a search but otherwise left the question blank. Interestingly, there is significant variation across officers in the proportion of citations for which the search variable is left missing; some officers appear to have been better at accurately filling out the citation than others. There are a number of ways of dealing with these missing values. We pick the method that we think is the best and then check to make sure that our results are robust to alternative procedures. In our baseline specification, if the officer indicated that a search was conducted for all citations in which search was non-missing, then we assume that if the search variable is missing, then no search was conducted. Then, we drop all officers for whom search is missing for more than 10% of the citations that those officers issue. Doing so eliminates approximately 25% of the citations (and 48% of the officers) in our data. For the remaining 684 officers, we drop observations for which search is missing, and are left with a sample comprising 72,903 citations. Tables 1 and 2 were calculated using our baseline search measure. In the next section, we discuss our robustness checks in greater detail.

Table 3 presents some basic summary statistics. The first column includes only those citations for which our baseline search measure is missing, whereas the second column includes only those citations for which our baseline search measure is available. Thus, comparing these first two columns provides some idea as to whether the citations for which search is missing differ systematically from those where it is not. Among citations for which search is missing, accidents are about twice as likely to have occurred as among citations for which search is not missing. There is also some variation across the first two columns in the percentage of citations that are issued in each neighborhood, reflecting the fact that officers in some districts were less likely to leave the search question blank than were officers in other districts. Otherwise, however, citations for which the search variable is missing appear to be quite similar to those for which it is not.[1] The last three columns of Table 3 show the average characteristics of the citations in our sample broken down by the race of the officer issuing the citation. Interestingly, we see that officers disproportionately issue citations to motorists from their own racial group. As we will see below, this may reflect the fact that officers are more likely to issue tickets in districts in which a large portion of the population (and so, presumably, the drivers) are in the same racial group as the officer. Indeed, this is also reflected in the fact that there is variation across the last three columns in the proportion of citations issued in different neighborhoods. Finally, we see that black officers are more likely to issue citations at night and less likely to issue citations at which an accident has occurred than either white or Hispanic officers.

[1] We also estimated Probit models for whether or not the search variable was missing as a function of officer and driver characteristics. The mismatch coefficient turns out to be negative but statistically insignificant. This insignificance suggests that the omission of missing observations is not driving our results. Even if the coefficient were, this results would only serve to bias us against finding preference-based discrimination under the assumption that non-searches were more likely to be coded as missing observations. That is, our data are missing non-searches in which the race of the officer and driver were likely to match.

Search Patterns in the Boston Police Department

In this section we test our model's theoretical predictions. For the time being we abstract from the possibility that there exist racial differences in officers' abilities to assess the guilt of motorists from different racial groups and the possibility that officers may be non-randomly matched with motorists from different racial groups.

We start by replicating the results presented in KPT. To do so, we use a probit model to study the probability of search and the probability of guilt conditional on search. In order to determine how the probability of search and the probability of guilt conditional on search differ depending on the race of the driver, we include indicator variables for whether the driver is black or Hispanic (so that white drivers are our omitted category). We also include as controls indicator variables for whether the stop occurred at night (6pm-5am), whether the driver was below the age of 25, whether the driver was male, whether the driver was from in state, whether the driver was from in town and whether an accident had occurred. In addition, we include dummy variables that control for the district in which the stop occurred. In Table 4 (and in all remaining relevant tables) we report the estimated marginal impact of each variable on the probability of search. Column 1 presents the results from the probit model of the probability of search, and column 2 presents the results for the probability of guilt conditional on search. In these first two columns, each citation receives equal weight. However, concern that these results are driven by a small number of officers who issue an unusually large number of citations prompted us to repeat the analysis in columns 1 and 2, but instead weight each citation by one over the number of citations given by the officer issuing that citation. The last two columns of Table 4 present the results of these weighted probits.

As will be seen, our results are sometimes sensitive to whether or not we weight citations in this fashion. In fact, the merits of weighting depend upon the question that you wish to answer. If you are interested in understanding the behavior of the average officer, the weighted probits provide a better description of the data since officers who issue a large number of tickets do not exert a disproportionate impact on the estimates. On the other hand, if you are interested in understanding search outcomes for the average motorist who receives a citation, then the unweighted probits are more appropriate. In this paper, we are interested in understanding the search decisions of officers and, in particular, whether their behavior is consistent with preference-based discrimination. Thus, we believe that the results of the weighted probits are appropriate. For several reasons, however, we do present robustness checks using the unweighted probits. First, describing search outcomes for the average motorist is interesting in its own right. Second, we are privy in this paper to an immensely rich data set. Future economists, legal analysts and other researchers may or may not have the type of officer-level data that are available to us. The differences between the weighted and unweighted probits, and the concomitant differences in the interpretation of the results, highlight the fact that citation-level data (even if officer race is available) may lead to misleading results if it is not possible to account for the fact that officers who issue a large number of tickets will be over-represented in the sample.

As the first column of the table indicates, black drivers are significantly more likely to have their cars searched than are white drivers. This result also holds for the weighted probit in column three. In addition, like Knowles, Persico and Todd, we find no evidence that the probability of guilt conditional on search differs depending on the race of the driver. In particular, in both columns 2 and 4, the coefficient of the indicator variable for whether the driver is black is close to zero and not statistically different from zero. Table 5 is identical to Table 4 except that it drops citations for which either the officer or the driver was Hispanic. The results are very similar to those in Table 4.

KPT interpret the finding that the probability of guilt conditional on search is identical across racial groups as evidence against preference-based discrimination. However, as the discussion in the preceding section highlights, once the information structure of KPT is expanded to allow the police to observe a noisy signal of motorist guilt that it not perfectly known to motorists at the time they make their decision to carry drugs, this prediction no longer holds.

However, our model delivers an alternative method for distinguishing between preference-based discrimination and statistical discrimination. In particular, as discussed above, the model predicts that if statistical discrimination alone explains differences in the rate at which African-American and white drivers are pulled over, then there should be no difference in the rate at which officers from different racial groups search drivers from any given racial group. In order to determine how search patterns depend on the interaction between the race of the driver and the race of the officer, we again use a probit model to analyze the probability of search. Here, in addition to controlling for the race of the driver, we also include indicator variables for the race of the officer as well as an indicator variable that is equal to 1 if the race of the officer differs from the race of the driver (we call this indicator "mismatch"). Table 6 presents our results. In the first three columns, each citation receives equal weight, and each column includes a progressively broader set of controls. The last three columns are identical to the first three, except that in the last three columns each citation is weighted by one over the number of citations given by the officer issuing the citation. In all six columns, the coefficient on our mismatch indicator is positive and statistically different from zero at standard confidence levels. Thus, our results indicate that officers are more likely to search motorists who are not members of the officer's racial group. As mentioned before, this finding is inconsistent with standard models of statistical discrimination. In addition, our results also suggest that Hispanic officers are more likely to conduct searches than white officers, and the second and third columns suggest that officers are more likely to search motorists if they are black, young or involved in an accident. Table 7 presents the results from the same analysis as in Table 6 except that stops involving either Hispanic officers or Hispanic motorists are excluded from the sample. Again, in all six columns the coefficient estimate on the mismatch parameter is positive and significantly different from zero at standard confidence levels.

Note that a positive coefficient on our mismatch parameter could be driven, for example, by discrimination on the part of white officers against black drivers or by discrimination on the part of black officers against white drivers. The problem, of course, is that our data do not allow us to distinguish between these two possibilities since we have no known non-discriminatory group of officers against which to compare our results. Thus, for example, our results should not be taken as evidence that black motorists in the Boston area are the subject to discrimination by officers in the Boston Police Department. This may be true, but our results do not shed light directly on this issue. Rather, our results simply indicate that the interaction between the race of the motorist and the race of the officer is significantly related to the probability that the motorist is searched, and we argue that this pattern is consistent with preference-based discrimination.

As mentioned earlier, the search variable is missing for over 18% of the citations in our data. In order to make sure that our results are not sensitive to the way in which we treat these missing values, we conduct a number of robustness checks, the results of which are presented in Table 8. In the first column, we run the same basic specification as above with our full set of controls, but include in the analysis citations issued by officers for whom the search variable is missing in more than 10% of the citations issued by that officer. In the second column, we assume that if search was missing, then no search was conducted. The motivation for this assumption is the notion that officers may be more likely to leave the search question blank if no search was conducted. This obviously increases our sample size substantially. Finally, in column three, we repeat the analysis in column 1 except that if all of an officer's non-missing search citations indicate that a search was conducted, then we assume that no search was conducted for all of the missing observations. As shown, the point estimates drop in size relative to the comparable estimate using our baseline search measure. However, the mismatch coefficient remains statistically different from zero at the 99% confidence level.

Table 9 repeats the analysis in Table 8, except that it does not include stops that involve either Hispanic officers or Hispanic motorists. In all three columns, the coefficient on the mismatch indicator is positive and is statistically different from zero at standard confidence levels.

Recall that in our baseline search measure we drop officers for whom the search variable is missing for more than 10% of the citations issued by that officer. Although we do not present these results, we have also experimented with changing that 10% cutoff. Lowering the cutoff (to say 5% or 3%), tends to strengthen our results, while increasing the cutoff tends to weaken them. This is reflected in column 1 of Tables 8 and 9 where the cutoff is effectively 100% (all officers are included).

Approximately 82% of the officers in our data are patrol officers. The remaining officers are some manner of either Deputies, Detectives, Sergeants or Captains. We lack information on how an officer's duties vary according to his or her rank and, more importantly, we do not know how rank affects ticketing behavior (although it's clear that high-ranking officers issue fewer tickets). Thus, in Table 10, we repeat the analysis in Table 6, but restrict attention to citations that are issued by Patrol Officers. In all three columns, the coefficient on the mismatch indicator are positive, similar in magnitude to the comparable estimates in Table 6, and are statistically different from zero at standard levels. Table 11 repeats the analysis in Table 10 but excludes citations in which either the officer or the motorist is Hispanic. Again. the point estimates on the mismatch indicator are similar in magnitude to those in Table 7, but are statistically different from zero (at the 90% confidence level) in only two out of the three columns. Overall, we take the results in Tables 10 and 11 as evidence that our results are not driven by idiosyncratic policing practices among higher-ranking officers.

As noted above, our data only include warnings for the two-month period of April-May 2001. In order to determine whether selection on the officer's decision to issue tickets or warnings is driving our findings of preference-based discrimination, we next restrict our sample to those stops within this two-month period. As shown in Table 12, the coefficient on the mismatch variable is larger than those in the baseline analysis in Table 6. While the standard errors are also larger, probably reflecting the diminished sample size, the coefficients remain statistically significant at the 90 percent level.[2] Taken together, these robustness

[2]Although the results are not presented here, we also examined officer decisions over whether to issue warnings or tickets during this two-month period. Warnings were less likely to be issued if there was a

checks demonstrate that our findings of preference-based discrimination are not driven by observations with missing search variables, patrol officer versus non-patrol officer distinctions, or selection on the officer's decision to issue warnings or tickets. In the remainder of the paper, we address other alternative explanations for our findings: informational asymmetries and the assignment of officers to neighborhoods.

Informational Asymmetries

One possible explanation for our results is that officers may differ in their ability to assess the likelihood that motorists from different racial groups are carrying contraband. In particular, we would typically expect that officers are better able to assess the guilt of members of their own racial groups. For example, black officers may have better information than white officers about the likelihood that black motorists are carrying contraband, and, similarly, white officers might have better information than black officers about the likelihood that a white motorist is carrying contraband. Intuitively, however, it is not clear how this type of informational asymmetry will affect search patterns. For example, if one assumes that search costs are sufficiently high that officers search a small number of motorists relative to the number who traffic drugs, then one would expect officers with better information to be *more* likely to search since they can better avoid unnecessarily incurring the cost of search. On the other hand, if search costs are low so that officers search a large number of motorists relative to the number who traffic drugs, then one would expect officers with more information to be *less* likely to search since there is no point in searching drivers who are innocent. Since informational asymmetries do not have an intuitively obvious effect on search patterns, we rely upon our model for further insights. Thus, in this section, we allow the informational content of the signal, θ, to differ depending on the race of the motorist and the race of the officer.

In the framework of the model presented above, changes in the quality of information are reflected in F_G and F_N, the distribution of θ for guilty and innocent drivers, respectively. As information improves, $F_G(\theta) \to 1$ and $F_N(\theta) \to 0 \ \forall \ \theta$. On the other hand, as information declines, $F_G \to F_N \ \forall \ \theta$. Thus, as θ becomes more informative, the slope of $\rho(\theta)$ approaches infinity in absolute value, and as θ becomes less informative, the slope of $\rho(\theta)$ approaches zero. That is, when information is good, changes in θ have a strong effect on the officer's assessment of the motorist's guilt, but when information is bad, changes in θ will have almost

mismatch between the race of the officer and the race of the driver, although this coefficient was statistically insignificant. As discussed earlier, however, this finding would only serve to bias our results against finding preference-based discrimination during the post-warnings period.

no effect on the officer's assessment of the motorist's guilt.

We would like to know how better information will affect the equilibrium outcome. First, note that changes in information only affect the best response function of the police. Second, if the police have no information (so that $F_G = F_N$ and $f_G = f_N$), then our model is equivalent to the model in KPT (2001) in which police do not observe a noisy signal of whether the driver is carrying contraband. Thus, the model in KPT is a special case of the model presented in this paper in which $P(G|r, \theta) = P(G|r) = \pi_r$. Third, note that as information becomes complete (so that $F_G \to 0$ and $F_N \to 1$), then $\gamma(\pi) = \pi$. This makes sense; if you can tell who the criminals are, they are the only people you search.

Unfortunately, it is not generally possible to sign the effect of this sort of informational asymmetry on the probability of search. This is best illustrated in a simple example. Suppose that the cost of searching motorists from group r is the same for African-American and white officers so that $t_r^a = t_r^w = .3$. However, suppose that African-American officers are better able to assess the criminality of motorists from group r. In particular, for motorists from group r, $F_G(\theta) = \theta^\lambda$ and $F_N(\theta) = 1 - (1 - \theta)^\lambda$, where $\lambda = 2$ if the officer is African-American and $\lambda = 1.5$ if the officer is white.

Given these parameters, Figure 3 plots $\gamma_r^a(\pi_r)$ and $\gamma_r^w(\pi_r)$, the best response function for African-American and white police officers, respectively. As the figure reveals, if $\pi_r^* > .263$, then African-American officers will be less likely than white officers to search motorists from group r. Otherwise, African-American officers will be more likely than white officers to search motorists from group r. Thus, without making additional assumptions, we cannot determine the effect of information on the probability of search.

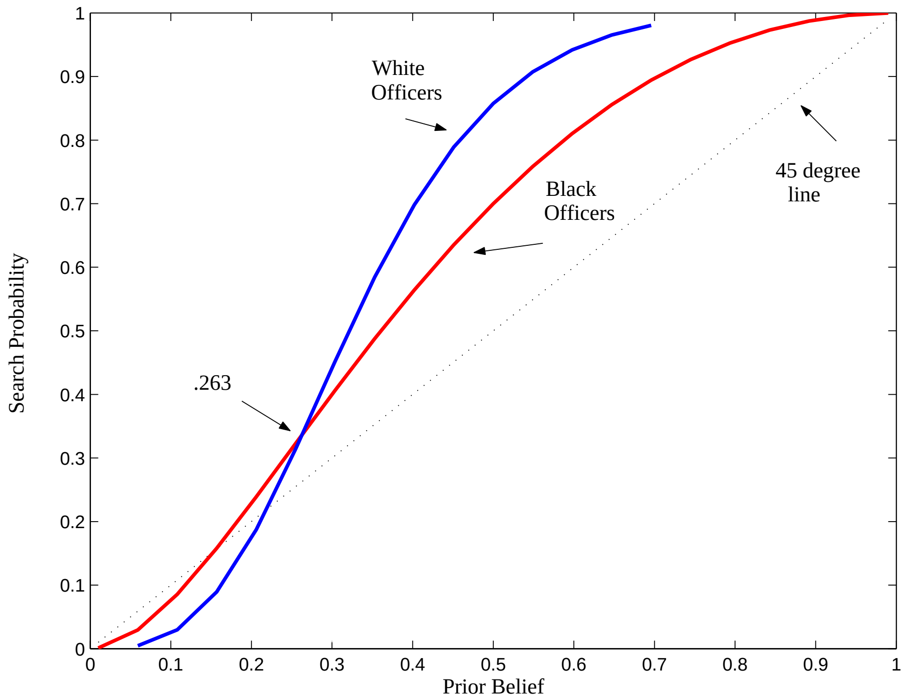

Figure 3: Model With Informational Asymmetries

Testing for Informational Asymmetries

As the above discussion highlights, our model delivers ambiguous predictions about the probability of guilt conditional on search. Thus, if officers are better at determining whether motorists from their own racial group are trafficking drugs, then this could either raise or lower the probability that they will search motorists from their own racial group relative to motorists from other groups.

As a result, the question of how improved information affects search probabilities is entirely an empirical one. In order to help answer this question, we examine whether our earlier results hold even among officers with more than 10 years of experience on the force. The hypothesis is that among these older officers, informational asymmetries should be less severe since these more experienced officers have had the opportunity to interact with a large number of drivers from different racial groups. Thus, if officers become better at determining whether motorists from a particular group are guilty as exposure to that group increases, then, provided these improvements happen at a decreasing rate, experienced officers should be similarly able to judge the likelihood that motorists from different racial groups are carrying drugs.

In other words, if informational asymmetries drive the results in Tables 4 and 5, then we would expect the coefficient on the mismatch indicator to be statistically indistinguishable from zero for more experienced officers. Thus, Table 12 presents the results of our weighted probits; the first three columns focus on citations issued by officers with less than 10 years of experience while the last three columns focus on citations issued by officers with more than 10 years of experience. We chose 10 years of experience as our cutoff because it is close to the average experience level of officers in our data, approximately 12 years. However, our results are not sensitive to the exact cutoff experience level that we employ.

As the Table shows, the coefficients on the mismatch indicator are small and statistically insignificant for inexperienced officers but large and statistically significant for experienced officers. Thus, these results suggests that our findings of preference-based discrimination are *not* driven by informational asymmetries.

How Are Officers Assigned to Neighborhoods?

One final explanation for our results is that officers may not be randomly assigned to different neighborhoods. For example, if black motorists are more likely to traffic drugs in neighborhoods that are disproportionately patrolled by white officers than in neighborhoods that are disproportionately patrolled by black officers, then, even in the absence of preference-based discrimination, we would expect white officers to be more likely than black officers to search the cars of black motorists. From a public relations perspective, it seems unlikely that officers would be assigned to neighborhoods in this fashion. Nonetheless, it is worth examining how the Department allocates officers across the City.

Officers in the Boston Police Department are assigned to one of 11 districts. These districts correspond to well-defined geographic areas within the City and are the primary organizational units for the Department. Figure 4 indicates both the name and location of these 11 districts. In addition, the Boston Police Department has a "Same Cop Same Neighborhood" (or "SC/SN") policing policy. Under SC/SN, patrol officers are assigned to a neighborhood beat within each district, and spend no less than 60% of their shift in that beat. The intent of SC/SN is to enable officers to become familiar with the local community to which they are assigned and, thus, be more effective at preventing crime. Unfortunately, while our data contain information on the district to which an officer was assigned at the time he or she issues a citation, we lack information on the officer's neighborhood beat.

Nonetheless, in Table 14, we compare the racial composition of the population aged 18 and over in each district with the racial composition of the officers who are assigned to that district. As the table shows, in districts in which a relatively large percentage of the population is white, a relatively large proportion of the officers assigned to that neighborhood are white. Similarly, in districts in which a relatively large proportion of the population is black, a relatively large proportion of the officers assigned to that district is black, and the same basic pattern holds for Hispanics. For whites the correlation between the percentage of the population aged 18 and older in each district and the fraction of officers in that district who are white is 0.751. For blacks, Asian and Hispanics the analogous correlation is 0.844, 0.575 and 0.885, respectively. To some extent, these patterns may reflect intentional assignment patterns on the part of officials at the Boston Police Department. However, officers also have some discretion about the district to which they are assigned. In any case, officers appear to patrol areas in which the majority of residents are members of the officer's racial group. Obviously, this table tells us little about the propensity of motorists from different racial groups to traffic drugs in each district.

Conclusion

Between April 2001 and January 2003, over 43 percent of all searches conducted by officers from the Boston Police Department were of cars driven by African-American motorists even though cars driven by African-Americans made up less then 33 percent of the cars that were pulled over. One possible explanation for this discrepancy is statistical discrimination. Another is preference-based discrimination. In this paper, we develop a test that allows us to distinguish between these two hypotheses.

We start by generalizing the information structure of the canonical model of police search developed in Knowles, Persico and Todd (2001) and show that if the police observe a noisy signal of the likelihood that a motorist is carrying contraband, then the fundamental insight that allows KPT to distinguish between preference-based discrimination and statistical discrimination falls apart. In particular, we show that even in the absence of preference-based discrimination, there is no reason to expect the probability of guilt conditional on search to be the same across racial groups. However, we suggest an alternative test for distinguishing between statistical discrimination and preference-based discrimination. Our model predicts that if statistical discrimination alone accounts for racial disparities in the rate at which motorists from different racial groups are subject to search, then there should be no difference in the rate at which *officers* from different racial groups search drivers from any given group.

We test this hypothesis using data from the Boston Police Department. Our results strongly suggest that officers are much more likely to conduct a search if the race of the motorist differs from the race of the officer. We then test whether this pattern could be explained by differences in the ability of officers from different racial groups to assess the guilt of motorists from a particular racial group. We find no evidence that this sort of informational asymmetry explains our results. We also discuss whether non-random assignment of officers to different neighborhoods within the city could account for our findings. They do not appear to do so. Rather, our results suggest that preference-based discrimination plays a significant role in explaining differences in the rate at which motorists from different racial groups are searched during traffic stops.

References

[1] Donohue, John J and Levitt, Steven D. "The Impact of Race on Policing and Arrests," *Journal of Law and Economics* 44 (October 2001), 367-394.

[2] Hernández-Murillo, Ruben and Knowles, John. "Racial Profiling or Racist Policing?: Testing in Aggregated Data," *Mimeo*, (April 2003).

[3] Knowles, John; Persico Nicola; and Todd, Petra. "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy* 109 (February 2001), 203-229.

Table 1: Probability of Search by Officer Race and Driver Race

(Standard Deviation of Sample Mean in Parentheses)

	Officer Race			
Driver Race	White	Black	Hispanic	All
White	0.41%	0.67%	0.24%	0.47%
	(0.04%)	(0.08%)	(0.08%)	(0.04%)
	n=23359	n=11399	n=3370	n=38128
Black	1.00%	0.81%	0.47%	0.88%
	(0.09%)	(0.09%)	(0.14%)	(0.06%)
	n=13533	n=9326	n=2339	n=25198
Hispanic	1.01%	0.80%	0.36%	0.87%
	(0.14%)	(0.16%)	(0.18%)	(0.09%)
	n=5233	n=3237	n=1105	n=9575
All	0.65%	0.72%	0.38%	0.65%
	(0.04%)	(0.05%)	(0.07%)	(0.03%)
	n=45755	n=25392	n=7181	n= 78328

Note: Stops made by Asian officers are not inlcuded. Includes only officers for whom the search variable is missing for at most 10% of all citations. Stops involving drivers from other racial groups are included in the "All" category.

Table 2: Probability of Search by Officer Race and Driver Race Weighted by Inverse of Number of Citations

(Standard Deviation of Sample Mean in Parentheses)

	Officer Race			
Driver Race	White	Black	Hispanic	All
White	1.91%	2.54%	2.50%	2.09%
	(0.53%)	(0.56%)	(2.10%)	(0.42%)
	n=404	n=139	n=46	n=589
Black	5.05%	2.04%	0.48%	3.95%
	(1.05%)	(0.91%)	(0.22%)	(0.74%)
	n=364	n=137	n=42	n=543
Hispanic	4.89%	4.55%	0.28%	4.34%
	(1.64%)	(2.43%)	(0.16%)	(1.23%)
	n=265	n=111	n=37	n=413
All	3.19%	2.78%	1.38%	2.96%
	(0.47%)	(0.56%)	(0.98%)	(0.36%)
	n=473	n=164	n=52	n=689

Note: Stops made by Asian officers are not inlcuded. Includes only officers for whom the search variable is missing for at most 10% of all citations. Stops involving drivers from other racial groups are included in the "All" category. For each officer, observations weighted by one over the number of citations given by that officer.

Table 3: Summary Statistics

(Standard Deviation in Parentheses)

Variable	Baseline Search Missing	Primary Sample			
	All Officers	All Officers	White Officers	Black Officers	Hispanic Officers
White Driver	56.7%	52.3%	55.5%	47.6%	49.5%
	(49.5%)	(49.9%)	(49.7%)	(49.9%)	(50.0%)
Black Driver	31.1%	34.6%	32.1%	38.9%	34.3%
	(46.3%)	(47.6%)	(46.7%)	(48.8%)	(47.5%)
Hispanic Driver	12.1%	13.1%	12.4%	13.5%	16.2%
	(32.7%)	(33.8%)	(33.0%)	(34.2%)	(36.9%)
Mismatch	49.4%	53.6%	44.5%	61.1%	83.8%
	(50.0%)	(49.9%)	(49.7%)	(48.8%)	(36.9%)
Baseline Search	-	0.7%	0.7%	0.7%	0.3%
		(8.1%)	(8.2%)	(8.6%)	(36.9%)
Stop at Night	32.1%	30.7%	27.0%	37.3%	30.9%
	(46.7%)	(46.1%)	(44.4%)	(48.4%)	(46.2%)
Young Driver (Age<26)	25.0%	24.3%	24.2%	23.8%	26.3%
	(43.3%)	(42.9%)	(42.8%)	(42.6%)	(44.0%)
Male Driver	71.8%	68.1%	69.3%	65.6%	69.9%
	(45.0%)	(46.6%)	(46.1%)	(47.5%)	(45.9%)
In-State Driver	93.3%	93.2%	92.9%	93.8%	93.0%
	(25.0%)	(25.2%)	(25.8%)	(24.2%)	(25.4%)
In-Town Driver	49.2%	51.1%	49.1%	54.3%	52.5%
	(50.0%)	(50.0%)	(50.0%)	(49.8%)	(49.9%)
Accident	2.6%	1.3%	1.5%	0.9%	1.3%
	(15.8%)	(11.3%)	(12.1%)	(9.7%)	(11.5%)
Allston-Brighton	7.6%	6.8%	8.3%	4.6%	4.7%
	(26.5%)	(25.1%)	(27.6%)	(21.0%)	(21.1%)
Boston Central	19.9%	13.1%	12.8%	12.2%	18.1%
	(39.9%)	(33.7%)	(33.4%)	(32.7%)	(38.5%)
Charlestown-East Boston	10.1%	6.2%	8.4%	3.2%	3.2%
	(30.2%)	(24.1%)	(27.8%)	(17.5%)	(17.5%)
Dorchester-Mattapan	21.4%	19.9%	20.3%	20.1%	16.4%
	(41.0%)	(39.9%)	(40.2%)	(40.1%)	(37.0%)
Hyde Park	0.7%	0.9%	0.7%	1.3%	0.5%
	(8.4%)	(9.3%)	(8.3%)	(11.4%)	(6.7%)
Jamaica Plain	2.3%	2.5%	3.1%	0.4%	6.1%
	(14.9%)	(15.6%)	(17.3%)	(6.2%)	(23.9%)
Roslindale	0.5%	1.1%	1.3%	1.1%	0.7%
	(7.2%)	(10.6%)	(11.2%)	(10.2%)	(8.2%)
Roxbury	13.3%	17.5%	18.5%	15.2%	19.7%
	(33.9%)	(38.0%)	(38.9%)	(35.9%)	(39.8%)
South Boston	6.0%	4.0%	4.6%	3.5%	1.8%
	(23.7%)	(19.6%)	(21.0%)	(18.4%)	(13.2%)
Number of Officers	946	684	469	163	52
Number of Citations	27,505	72,903	42,125	23,964	6,814

Table 4: Probability of Search and Guilt Conditional on Search Officer Race Excluded

	Unweighted Probits		Weighted Probits	
	Search	Guilt	Search	Guilt
Black Driver	0.003***	-0.001	0.019**	-0.001
	(0.001)	(0.001)	(0.010)	(0.002)
Hispanic Driver	0.003	-0.000	0.015	0.003
	(0.002)	(0.000)	(0.013)	(0.005)
Stop at Night	0.002	0.000	0.013	-0.001
	(0.002)	(0.000)	(0.008)	(0.001)
Young Driver (Age<26)	0.001**	-0.000	0.006	-0.002
	(0.001)	(0.000)	(0.008)	(0.001)
Male Driver	0.001	-0.000	0.006	-0.005
	(0.001)	(0.000)	(0.006)	(0.007)
In-State Driver	0.001		0.010	
	(0.001)		(0.007)	
In-Town Driver	0.000	-0.000	0.003	0.001
	(0.001)	(0.000)	(0.006)	(0.001)
Accident	0.037**	-0.000	-0.000	-0.001
	(0.015)	(0.000)	(0.011)	(0.001)
Neighborhood Control	YES	YES	YES	YES
Observations	70614	363	70614	363

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%
If the search outcomes for an officer are missing in all cases except those in which the officer indicates that a search was conducted, then we code the missing search outcomes as "no search" (this affects 33 citations). After this correction, officers with missing values for more than 10% of search outcomes are dropped. Missing search outcomes for the remaining officers are dropped.

Table 5: Probability of Search and Guilt Conditional on Search Officer Race Excluded, Blacks and Whites Only

	Unweighted Probits		Weighted Probits	
	Search	Guilt	Search	Guilt
Black Driver	0.003**	-0.014	0.018*	-0.013
	(0.001)	(0.014)	(0.010)	(0.012)
Stop at Night	0.002	0.003	0.012	0.010
	(0.002)	(0.011)	(0.009)	(0.007)
Young Driver (Age<26)	0.001	-0.011	0.005	-0.015*
	(0.001)	(0.009)	(0.009)	(0.008)
Male Driver	0.001	-0.014	0.003	-0.031
	(0.001)	(0.012)	(0.007)	(0.023)
In-State Driver	-0.000		0.007	
	(0.001)		(0.009)	
In-Town Driver	0.001*	-0.001	0.006	0.001
	(0.001)	(0.009)	(0.008)	(0.005)
Accident	0.043**	-0.002	0.002	-0.005*
	(0.018)	(0.011)	(0.014)	(0.003)
Neighborhood Controls	YES	YES	YES	YES
Observations	55824	329	55824	329

Robust standard errors in parentheses.
* significant at 10%; ** significant at 5%; *** significant at 1%
If the search outcomes for an officer are missing in all cases except those in which the officer indicates that a search was conducted, then we code the missing search outcomes as "no search" (this affects 33 citations). After this correction, officers with missing values for more than 10% of search outcomes are dropped. Missing search outcomes for the remaining officers are dropped.

Table 6: Probability of Search, Baseline Specification

	Unweighted Probits			Weighted Probits		
	(1)	(2)	(3)	(4)	(5)	(6)
Black Driver	0.004***	0.003**	0.003**	0.008	0.003	0.006
	(0.001)	(0.001)	(0.001)	(0.009)	(0.007)	(0.008)
Hispanic Driver	0.003	0.002	0.001	0.004	0.000	-0.001
	(0.003)	(0.002)	(0.002)	(0.012)	(0.011)	(0.010)
Black Officer	0.001	0.001	0.001	-0.007	-0.009	-0.007
	(0.003)	(0.003)	(0.003)	(0.008)	(0.007)	(0.007)
Hispanic Officer	-0.004*	-0.003*	-0.003	-0.019**	-0.018**	-0.016**
	(0.002)	(0.002)	(0.002)	(0.009)	(0.008)	(0.007)
Mismatch	0.002**	0.002*	0.002**	0.019**	0.023***	0.020***
	(0.001)	(0.001)	(0.001)	(0.009)	(0.007)	(0.006)
Stop at Night		0.002	0.002		0.014	0.013
		(0.002)	(0.002)		(0.009)	(0.008)
Young Driver (Age<26)		0.002**	0.002**		0.006	0.006
		(0.001)	(0.001)		(0.008)	(0.007)
Male Driver		0.001	0.001		0.006	0.005
		(0.001)	(0.001)		(0.006)	(0.006)
In-State Driver		0.002	0.001		0.011	0.010
		(0.001)	(0.001)		(0.007)	(0.007)
In-Town Driver		0.000	0.000		-0.000	0.002
		(0.001)	(0.001)		(0.006)	(0.006)
Accident		0.039***	0.037**		0.000	-0.001
		(0.014)	(0.015)		(0.011)	(0.011)
Neighborhood Controls	NO	NO	YES	NO	NO	YES
Observations	72903	70614	70614	72903	70614	70614

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 7: Probability of Search, Baseline Specification Blacks and Whites Only

	Unweighted Probits			Weighted Probits		
	(1)	(2)	(3)	(4)	(5)	(6)
Black Driver	0.004**	0.003*	0.002**	0.011	0.004	0.004
	(0.001)	(0.001)	(0.001)	(0.010)	(0.007)	(0.008)
Black Officer	0.001	0.001	0.001	-0.008	-0.012	-0.011*
	(0.003)	(0.003)	(0.003)	(0.009)	(0.007)	(0.007)
Mismatch	0.002*	0.002*	0.002*	0.017*	0.022***	0.021***
	(0.001)	(0.001)	(0.001)	(0.010)	(0.008)	(0.007)
Stop at Night		0.002	0.002		0.012	0.012
		(0.002)	(0.002)		(0.009)	(0.009)
Young Driver (Age<26)		0.001	0.001		0.004	0.004
		(0.001)	(0.001)		(0.009)	(0.008)
Male Driver		0.001	0.001		0.002	0.001
		(0.001)	(0.001)		(0.007)	(0.007)
In-State Driver		0.000	-0.000		0.006	0.007
		(0.001)	(0.001)		(0.009)	(0.008)
In-Town Driver		0.001*	0.001*		0.005	0.006
		(0.001)	(0.001)		(0.007)	(0.007)
Accident		0.043**	0.043**		0.003	0.002
		(0.018)	(0.019)		(0.014)	(0.013)
Neighborhood Controls	NO	NO	YES	NO	NO	YES
Observations	57621	55824	55824	57621	55824	55824

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 8: Probability of Search, Robustness Checks

	Weighted Probits		
	Search 1	Search 2	Search 3
Black Driver	0.011*	0.008	0.011*
	(0.007)	(0.005)	(0.007)
Hispanic Driver	0.003	0.002	0.003
	(0.008)	(0.006)	(0.008)
Black Officer	-0.005	-0.004	-0.005
	(0.006)	(0.004)	(0.006)
Hispanic Officer	-0.010	-0.009*	-0.010
	(0.007)	(0.005)	(0.007)
Mismatch	0.014***	0.011***	0.014***
	(0.005)	(0.004)	(0.005)
Stop at Night	0.014**	0.010**	0.014**
	(0.006)	(0.004)	(0.006)
Young Driver (Age<26)	0.008	0.006	0.008
	(0.006)	(0.004)	(0.006)
Male Driver	0.010**	0.007**	0.010**
	(0.005)	(0.003)	(0.005)
In-State Driver	0.007	0.006	0.008
	(0.007)	(0.005)	(0.006)
In-Town Driver	0.004	0.001	0.004
	(0.005)	(0.004)	(0.005)
Accident	0.001	-0.000	0.001
	(0.009)	(0.006)	(0.009)
Neighborhood Controls	YES	YES	YES
Observations	79,337	95,855	79,369

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 9: Probability of Search, Robustness Checks Blacks and Whites Only

	Weighted Probits		
	Search 1	Search 2	Search 3
Black Driver	0.010	0.007	0.009
	(0.007)	(0.005)	(0.007)
Black Officer	-0.010*	-0.007*	-0.010*
	(0.006)	(0.004)	(0.006)
Mismatch	0.014**	0.011**	0.014**
	(0.006)	(0.005)	(0.006)
Stop at Night	0.011*	0.008	0.011*
	(0.007)	(0.005)	(0.007)
Young Driver (Age<26)	0.006	0.004	0.006
	(0.006)	(0.005)	(0.006)
Male Driver	0.007	0.005	0.007
	(0.005)	(0.004)	(0.005)
In-State Driver	0.003	0.002	0.003
	(0.008)	(0.006)	(0.008)
In-Town Driver	0.005	0.002	0.005
	(0.006)	(0.004)	(0.006)
Accident	-0.002	-0.003	-0.002
	(0.010)	(0.006)	(0.010)
Neighborhood Controls	YES	YES	YES
Observations	62,539	74,837	62,560

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 10: Probability of Search, Patrol Officers

	Weighted Probits		
	(1)	(2)	(3)
Black Driver	-0.005	-0.002	0.003
	(0.008)	(0.008)	(0.008)
Hispanic Driver	0.001	0.003	0.001
	(0.011)	(0.011)	(0.009)
Black Officer	-0.012	-0.009	-0.007
	(0.008)	(0.008)	(0.007)
Hispanic Officer	-0.018**	-0.015*	-0.013*
	(0.009)	(0.009)	(0.008)
Mismatch	0.016**	0.016**	0.013**
	(0.008)	(0.008)	(0.006)
Stop at Night		0.009	0.007
		(0.009)	(0.007)
Young Driver (Age<26)		-0.004	-0.003
		(0.007)	(0.006)
Male Driver		0.008	0.007
		(0.006)	(0.005)
In-State Driver		0.008	0.007
		(0.008)	(0.007)
In-Town Driver		-0.002	-0.000
		(0.006)	(0.006)
Accident		0.018	0.017
		(0.018)	(0.018)
Neighborhood Controls	NO	NO	YES
Observations	71,560	69,346	69,346

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 11: Probability of Search, Patrol Officers Blacks and Whites Only

	Weighted Probits		
	(1)	(2)	(3)
Black Driver	-0.001	0.001	0.005
	(0.008)	(0.008)	(0.009)
Black Officer	-0.010	-0.007	-0.007
	(0.008)	(0.008)	(0.007)
Mismatch	0.014	0.013*	0.013*
	(0.009)	(0.008)	(0.007)
Stop at Night		0.009	0.008
		(0.010)	(0.009)
Young Driver (Age<26)		-0.010*	-0.009
		(0.006)	(0.006)
Male Driver		0.004	0.004
		(0.007)	(0.006)
In-State Driver		0.002	0.003
		(0.009)	(0.008)
In-Town Driver		0.000	0.000
		(0.007)	(0.008)
Accident		0.028	0.028
		(0.024)	(0.023)
Neighborhood Controls	NO	NO	YES
Observations	56,409	54,684	54,684

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 12: Probability of Search, April-May 2001

	Weighted Probits		
	(1)	(2)	(3)
Black Driver	-0.016	-0.014	-0.015*
	(0.015)	(0.011)	(0.009)
Hispanic Driver	0.017	0.011	-0.003
	(0.027)	(0.022)	(0.014)
Black Officer	0.005	0.003	0.009
	(0.017)	(0.014)	(0.015)
Hispanic Officer	-0.004	0.004	-0.003
	(0.026)	(0.024)	(0.016)
Mismatch	0.030*	0.028*	0.030**
	(0.018)	(0.016)	(0.013)
Stop at Night		0.030	0.031*
		(0.021)	(0.017)
Young Driver (Age<26)		-0.009	-0.012*
		(0.010)	(0.007)
Male Driver		0.022**	0.017**
		(0.011)	(0.008)
In-State Driver		0.012	0.008
		(0.011)	(0.010)
In-Town Driver		-0.004	-0.001
		(0.010)	(0.008)
Accident		0.010	0.010
		(0.023)	(0.024)
Neighborhood Controls	NO	NO	YES
Observations	16,162	15,407	13,683

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 13: Probability of Search, Unexperienced vs. Experienced Officers Weighted Probits

	Inexperienced Officers (<=10 years)			Experienced Officers (>10 years)		
	(1)	(2)	(3)	(4)	(5)	(6)
Black Driver	-0.008	-0.004	-0.002	0.024	0.007	0.014
	(0.008)	(0.007)	(0.007)	(0.018)	(0.012)	(0.012)
Hispanic Driver	0.002	0.004	0.004	0.009	-0.005	-0.005
	(0.011)	(0.010)	(0.010)	(0.022)	(0.013)	(0.010)
Black Officer	-0.004	-0.001	-0.003	-0.006	-0.014*	-0.008
	(0.009)	(0.008)	(0.006)	(0.012)	(0.008)	(0.007)
Hispanic Officer	-0.026***	-0.020***	-0.018***	-0.004	-0.005	-0.003
	(0.006)	(0.005)	(0.004)	(0.022)	(0.017)	(0.014)
Mismatch	0.008	0.010	0.008	0.026*	0.032**	0.024**
	(0.007)	(0.006)	(0.005)	(0.015)	(0.013)	(0.010)
Stop at Night		0.025**	0.022**		0.002	0.000
		(0.011)	(0.009)		(0.009)	(0.007)
Young Driver (Age<26)		-0.011**	-0.009*		0.024*	0.021*
		(0.005)	(0.005)		(0.013)	(0.011)
Male Driver		-0.000	0.000		0.009	0.008
		(0.006)	(0.005)		(0.008)	(0.006)
In-State Driver		0.015***	0.013***		-0.004	-0.003
		(0.005)	(0.004)		(0.015)	(0.013)
In-Town Driver		-0.012**	-0.011*		0.013	0.013*
		(0.006)	(0.006)		(0.008)	(0.007)
Accident		0.005	0.004		0.003	0.004
		(0.009)	(0.007)		(0.016)	(0.015)
Neighborhood Controls	NO	NO	YES	NO	NO	YES
Observations	34073	33232	33232	38830	37382	37073

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Figure 4: City of Boston, Police Districts

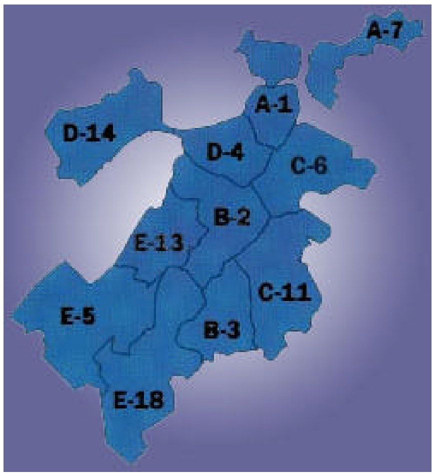

A-1 Downtown/Beacon Hill/Chinatown/Charlestown
A-7 East Boston
B-2 Roxbury/Mission Hill
B-3 Mattapan/North Dorchester
C-6 South Boston
C-11 Dorchester
D-4 Back Bay/Sound End/Fenway
D-14 Allston/Brighton
E-5 West Roxbury/Roslindale
E-13 Jamaica Plain
E-18 Hyde Park

Table 14: Racial Composition of Police Districts

	Census Benchmark Population 18 and Older				Citation-Level Data Racial Breakdown of Officers by District			
District	White	Black	Hispanic	Asian	White	Black	Hispanic	Asian
A-1	76.7%	3.3%	3.2%	15.1%	62.8%	24.8%	8.0%	4.4%
A-7	53.0%	2.4%	36.6%	3.7%	72.1%	16.4%	9.8%	1.6%
B-2	22.1%	47.8%	17.0%	4.7%	51.7%	35.7%	10.9%	1.7%
B-3	3.8%	78.9%	10.8%	1.1%	55.2%	37.3%	6.6%	0.9%
C-6	87.5%	1.8%	5.2%	4.0%	76.5%	14.8%	7.4%	1.3%
C-11	41.4%	28.7%	9.0%	12.5%	70.4%	17.3%	8.7%	3.6%
D-4	66.7%	9.9%	8.8%	11.5%	69.8%	21.2%	7.3%	1.7%
D-14	71.3%	3.9%	8.0%	13.1%	71.1%	16.3%	9.6%	3.0%
E-5	80.7%	6.2%	7.8%	3.0%	71.1%	22.2%	5.9%	0.7%
E-13	54.3%	15.2%	25.0%	2.3%	60.5%	21.6%	17.2%	0.8%
E-18	47.9%	33.7%	11.9%	3.0%	59.1%	30.7%	10.2%	0.0%