

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

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

July 13, 2005

Abstract

This paper provides new evidence on the role of preference-based versus statistical discrimination in racial profiling using a unique data set that includes the race of both the motorist and the officer. We build upon the model presented in Knowles, Persico and Todd (2001) and show that their test is not robust to alternative modelling assumptions. However, we also show that if statistical discrimination alone explains differences in the rate at which the vehicles of drivers of different races are searched, then, all else equal, search decisions should be independent of officer race. We then test this prediction using data from the Boston Police Department. Consistent with preference-based discrimination, our baseline results demonstrate that officers are 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: officers are better at searching members of their own racial group and the non-random 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, Hanming Fang, Arthur Goldberger, Roger Gordon, Nora Gordon, Winfried Koeniger and Finis Welch 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. The courts have tended to uphold racially biased policing patterns when they can be reasonably justified by racial differences in crime rates, but 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 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 by models of preference-based discrimination (see, for example, Knowles, Persico and Todd (2001), Hernández-Murillo and Knowles (2003), Anwar and Fang (2005) and Dharmapala and Ross (2004)).

In models of statistical discrimination, discrimination arises because law enforcement officials are uncertain about whether a suspect has committed a particular crime. 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 raises the benefit (or, equivalently, lowers the cost) of searching motorists from one group relative to those from some other group.¹

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.

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. We build upon the model of police search developed in Knowles, Persico and Todd (2001) (hereafter, often, KPT) and show that the test that they use to distin-

¹For an extended discussion of models of statistical discrimination and of preference-based discrimination see Arrow (1973) and Becker (1954), respectively.

guish between statistical discrimination and preference-based discrimination is not robust to alternative modelling assumptions.

According to KPT's model, in the absence of preference-based discrimination, the probability of guilt conditional on search will be the same for all identifiable groups of motorists. The logic is that, since the police will always search motorists for whom the likelihood of finding contraband exceeds the cost of search, in equilibrium, the drug trafficking behavior of motorists must adjust to equate the probability of guilt conditional on search for all individuals. The power of this insight is that, even if black and white motorists differ along dimensions other than race, the probability of guilt conditional on search will still be the same for all groups. This is critical since it is generally impossible to distinguish between statistical discrimination on the basis of race and statistical discrimination on the basis of characteristics that are correlated with race but that are unobserved to the econometrician (see for example, Altonji and Pierret (2001) and Dharmapala and Ross (2004)).

On the other hand, according to the model in KPT, if officers have discriminatory preferences, then the probability of guilt conditional on search will differ across racial groups. KPT then examine traffic stop data from the State of Maryland. They find that the probability of guilt conditional on search does not differ (in a statistical sense) across racial groups and interpret this as evidence that officers do not have discriminatory preferences.

As we show below, however, the validity of the test that KPT employ relies on a step-shaped best response function for officers. Any alteration that smooths out this officer best response function will permit equilibrium differences in guilt rates between groups of drivers, even in the absence of preference-based discrimination. In this paper, we focus on heterogeneity in officer search costs as the source of smoothness in the officer best response, but other models, such as one in which officers observe a noisy signal of motorist guilt, will also generate a smooth best response function.

To circumvent these problems, we develop an alternative mechanism for distinguishing between these two forms of discrimination that does not rely upon the probability of guilt conditional on search.² In particular, we show that if statistical discrimination alone explains differences in the rate at which African-American and white drivers are searched, then, all else equal, search decisions should be independent of the race of the police officer. Thus, we argue that if searches are more likely to occur when the race of the officer differs from the race of the driver, then this provides evidence of preference-based discrimination.

We then apply our test to a unique data set in which we are able to match the race of the

²Anwar and Fang (2005) develop a similar test to ours. We relate their test to ours in greater detail below.

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 2001.³ 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.

We find that if the race of the officer differs from the race of the motorist, then the officer is more likely to conduct a search than otherwise. We argue that our results cannot be explained by standard models of statistical discrimination and, instead, are consistent with preference-based discrimination. In addition, we rule out the possibility that our findings are driven by officers being better able to search members of their own racial group and by the way in which officers are assigned to neighborhoods.

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 is 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 0.40 percent if the officer is white and 0.62 percent if the officer is black. Similarly, the probability that a black motorist is searched is 0.82 percent if the officer is black but 0.97 percent 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

³For an alternative discussion of these data, see the series of articles by Bill Dedman and Francie Latour (2003).

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 when the officer making the stop is a member of the motorist's own racial group than they are when there is a mismatch between the race of the officer and the race of the motorist.

Abstracting at this stage from issues of statistical significance and other possible concerns, 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. Assuming that 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, officers may be better able to search motorists who are members of their own racial group. Second, officers may not be randomly assigned to neighborhoods. 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.

The Model

In this section, we re-examine the model of police search presented in Knowles, Persico and Todd (2001). In their model, motorists are assumed to be either African-American or white, denoted by a and w , respectively. Motorists are also distinguished by some characteristic, c , that is potentially useful to the police in determining whether or not to search a motorist's car. For now, we assume that both the police and the econometrician observe these driver characteristics (c); the case in which the econometrician does not observe driver characteristics will be investigated in the next section.

In deciding whether or not to carry contraband, motorists weigh the benefit of carrying

contraband against the penalty of being caught. If a driver does not carry contraband, then his payoff is zero regardless of whether or not his car is searched. However, if a motorist of type (c, r) carries contraband, then he is “guilty”, and faces cost $-j(c, r)$ if his car is searched and benefit $\nu(c, r)$ if his car is not searched.

Let $\gamma(c, r)$ be the probability the police search motorists of type (c, r) . Thus, the expected payoff to carrying contraband for motorists of type (c, r) is given by

$$-\gamma(c, r)j(c, r) + [1 - \gamma(c, r)]\nu(c, r).$$

Motorists are playing a best response to the search behavior of police if they carry contraband whenever the above expression is greater than zero. This occurs when

$$\gamma(c, r) < \frac{\nu(c, r)}{j(c, r) + \nu(c, r)}.$$

The police cannot perfectly observe whether a motorist of type (c, r) is guilty of carrying contraband. Instead, it is assumed that police maximize the expected payoff from making an arrest minus the cost of search. The benefit to making an arrest is normalized to one so that the cost of search is relative to the benefit. Let t_r be cost of searching a motorist from group r , where it is assumed that $0 < t_r < 1$ to rule out trivial equilibria. We do not incorporate a resource constraint on the total time that police spend searching; that is, officers can search all drivers if they so choose. The implications of this assumption, which is also employed in the baseline model of KPT, will be discussed at the end of this section.⁴

Let $\pi(c, r)$ denote the probability that a motorist of type (c, r) is guilty of carrying drugs or other contraband. Thus, the expected payoff to officers from searching motorists of type (c, r) is given by

$$\pi(c, r) - t_r.$$

Officers are playing a best response to the behavior of motorists if they search whenever the above expression is greater than zero. Figure 1 graphs the best response function for officers and motorists. For motorists of type (c, r) , the line labelled $\gamma(c, r)$ represents the best response function of police and the line labelled $\pi(c, r)$ represents the best response function of motorists.

As the figure reveals, the best response functions in this game are equivalent to those in a standard matching pennies game. Thus, for motorists of type (c, r) , there is a unique mixed strategy equilibrium of this game in which the search behavior of the police renders motorists indifferent between carrying and not carrying contraband, and the behavior of

⁴KPT describe this extension to their baseline model on page 215 of their paper.

motorists renders the police indifferent between searching and not searching. That is, in equilibrium the probability that the police search motorists of type c, r is given by

$$\gamma^*(c, r) = \frac{\nu(c, r)}{j(c, r) + \nu(c, r)}.$$

and the probability that motorists carry contraband is given by

$$\pi^*(c, r) = t_r.$$

Since the police randomly select motorists from this group, the probability of guilt conditional on search is exactly equal to $\pi^*(c, r)$, the probability that motorists carry contraband.

Before analyzing how statistical discrimination and preference-based discrimination manifest themselves in this model, it is worth enriching the model to allow for heterogeneity in motorists' utility from carrying contraband.⁵ In particular, we assume that the benefit to carrying contraband is given by

$$-\gamma(c, r)j(c, r) + [1 - \gamma(c, r)]\nu(c, r) - Z,$$

where Z represents an idiosyncratic cost of carrying contraband. Thus, letting $G(\cdot)$ denote the distribution of Z in the population, the probability that motorists of type (c, r) carry contraband is given by

$$G(-\gamma(c, r)j(c, r) + [1 - \gamma(c, r)]\nu(c, r)).$$

Figure 2 plots the corresponding best response functions for motorists and officers.

As in Knowles, Persico and Todd, we now examine the implications of statistical discrimination and preference-based discrimination for the equilibrium of this model. Police officers in this model are defined to have discriminatory tastes if the cost of search varies by the race of the motorist, so that $t_a \neq t_w$. For example, Figure 3 displays the best response functions when officers discriminate against African-Americans so that $t_a < t_w$. As Figure 3 reveals, if $t_a < t_w$, then, among motorists with characteristic c , the police will be more likely to search African-Americans than whites ($\gamma^*(c, a) > \gamma^*(c, w)$), and African-Americans will be less likely than whites to carry contraband ($\pi^*(c, a) < \pi^*(c, w)$).

In this model, police statistically discriminate on the basis of race if, among motorists with the same characteristic, c , the probability of search differs by the race of the motorist because of racial disparities in the propensity to carry contraband. This occurs whenever the net benefit of carrying contraband differs by race. That is, if $-j(c, r)$ and $\nu(c, r)$ vary by

⁵Knowles, Persico and Todd also discuss this extension on page 214 of their paper.

r , then so will the equilibrium probability of search. Figure 4 shows an example in which, even among motorists with the same observable characteristic, c , the net benefit of carrying contraband is higher for African Americans than it is for whites, but the cost of search is assumed to be the same for all motorists, so that $t_a = t_w = t$. In this case, for any given search probability γ , African-Americans are more likely than whites to carry contraband, and, as the figure reveals, in equilibrium, the police are more likely to search African-American motorists than white motorists ($\gamma^*(c, a) > \gamma^*(c, w)$). Importantly, however, the probability of guilt conditional on search is independent of motorist race ($\pi^*(c, a) = \pi^*(c, w) = t$).

Thus, in KPT, statistical discrimination has no effect on the probability of guilt conditional on search whereas preference-based discrimination does. In the absence of preference-based discrimination, black and white drivers must carry at the same rate [$\pi^*(c, a) = \pi^*(c, w) = t$] in order to make officers indifferent between searching and not searching. If black drivers, for example, carry at a rate higher than t , then officers will search black drivers with probability 1. However, if black drivers are searched with probability 1, then it is optimal for them to carry at rate 0 and thus black drivers carrying at a rate higher than t cannot be an equilibrium. Similarly, if black drivers carry at rate lower than t , they will be searched with probability 0, and it is now optimal for them to carry with probability 1. Thus, in the absence of preference-based discrimination, all drivers must carry at the same rate in equilibrium.

However, Figure 4 also makes clear the fact that the test that KPT adopt depends on the specialized step-shaped best response function for police officers. Any alteration to the model that smooths out the best response function for police officers will allow drivers to carry at different rates even in the absence of preference-based discrimination. One such alteration is to allow the police to be heterogenous in their preferences for search. Indeed, this type of heterogeneity is often used to justify the randomization that characterizes most mixed-strategy equilibria (for example, see Harsanyi (1973)). In addition, our data from Boston suggest that officers do vary in their preferences for search in the sense that we observe substantial variation across officers in the likelihood that they search motorists whom they have pulled over.

Formally, we assume that the cost of searching motorists from group r is given by $t_r + U$, where U captures a mean-zero, idiosyncratic search cost. With this cost, an officer is playing a best response to motorists of type (c, r) if they search whenever

$$\pi(c, r) - t_r - U > 0.$$

Letting $H(\cdot)$ denote the distribution of U among officers, the probability that optimizing

officers search is given by

$$H(\pi(c, r) - t_r),$$

which is clearly increasing in $\pi(c, r)$. Figure 5 plots best response functions with heterogeneity in both motorist preferences for carrying contraband and officer search costs. In this figure, we plot the best response function for both African-American and white motorists of type c under the assumption that the net benefit to carrying contraband is higher for African-Americans than it is for whites and under the assumption that $t_a = t_w = t$.

As shown, drivers now carry at different rates even in the absence of preference-based discrimination. In the model without officer heterogeneity, if one group carries at a higher rate than another, then officers will search at least one of the groups with probability 0 or 1 in the absence of preference-based discrimination; as described above, however, such extreme search rates cannot be supported in equilibrium. In a model with officer heterogeneity, by contrast, one group carrying at a higher rate can be supported in equilibrium even in the absence of preference-based discrimination due to the smoothness of the officer best response function. Thus, in this generalized model, the test that KPT employ no longer distinguishes between preference-based discrimination and statistical discrimination.

Another alteration to the model presented in KPT that leads to smooth, upward-sloping best response functions for police officers is to assume, like Anwar and Fang (2005) and Bjerk (2005), that officers can observe some characteristic, θ , that is correlated with the likelihood that the motorist is guilty, but that (unlike c) is unknown to the motorist at the time that he or she decides to carry contraband. It can be shown that, in this setting, the best response function for officers will have the same upward-sloping shape as the best response function for officers in Figure 5.⁶ In addition, if officers are not able to perfectly observe some groups of criminals (as suggested by Bjerk (2005) and Dharmapala and Ross (2004)), then this will also affect the step-shaped best response function of police officers.

It should now be clear that in a model with a smooth officer best response function, whatever the source of this smoothness, guilt rates may differ across groups even in the absence of preference-based discrimination. While homogenous search costs lead to step-shaped officer best response functions in KPT's model, a binding constraint on officers' search capacity may provide an alternative source of such discontinuities. In such a model, search decisions are no longer independent across drivers, and officers optimally target their constrained resources towards the group with the highest net returns from searching. An

⁶The models in Anwar and Fang (2005) and Bjerk (2005) highlight the added problem that the guilt rate for the marginal motorist will not equal the guilt rate for the average motorist. A complete analysis of how these models affect the shape of the best response functions is available upon request.

important question is whether the officer best response functions are smooth or step-shaped in a model with both heterogeneity in officer search costs and binding resource constraints. Persico and Todd (2004) analyze such a model and find that their hit rate test is valid. Thus, if officers have heterogeneous preferences, the test that KPT employ will only be valid if officers face binding resource constraints.

Whether or not resource constraints are binding in practice is an open question and may vary from application to application. Like us, Anwar and Fang (2005) and Bjerck (2005) assume that officers do not face binding capacity constraints, and we feel that this assumption is appropriate for our analysis of traffic stops by the Boston Police Department. Searches are relatively rare in practice: as shown in Table 1, under 1,000 searches occurred in Boston over the period April 1, 2001-January 31, 2003. Moreover, search rates vary significantly across officers, suggesting that resource constraints are unlikely to be binding for the many officers who search at low intensity in our dataset.

An Alternative Test

In this section we propose a simple alternative test to distinguish between preference-based discrimination and statistical discrimination. We start from the supposition that if officers have discriminatory preferences, then those prejudices will be directed towards motorists who are not members of that officer's racial group. That is, the cost of search, t , will depend on the match between the race of the officer and the race of the motorist. We assume that, like motorists, officers belong to one of two racial groups: African-Americans and whites. Let t_r^j denote the cost to officers from group j of searching motorists from group r .

Thus, the payoff to officers from group j of searching motorists from group r becomes

$$\pi(c, r) - t_r^j - U.$$

Let $\gamma^j(c, r)$ denote the probability that officers from group j search motorists from group r and let ρ denote the proportion of officers from who are African-American. Thus, the payoff to motorists from group r with idiosyncratic cost Z from carrying contraband is given by

$$-\gamma(c, r)j(c, r) + [1 - \gamma(c, r)]\nu(c, r) - Z,$$

where $\gamma(c, r) = \rho\gamma^a(c, r) + (1 - \rho)\gamma^w(c, r)$.

An equilibrium for motorists of type (c, r) occurs at any $\pi^*(c, r)$, $\gamma^{a*}(c, r)$, and $\gamma^{w*}(c, r)$ such that

$$\gamma^{a*}(c, r) = H(\pi^*(c, r) - t_r^a)$$

$$\gamma^{w*}(c, r) = H(\pi^*(c, r) - t_r^w)$$

and

$$\pi^*(c, r) = G(-\gamma^*(c, r)j(c, r) + [1 - \gamma^*(c, r)]\nu(c, r))$$

where $\gamma^*(c, r) = \rho\gamma^{a*}(c, r) + (1 - \rho)\gamma^{w*}(c, r)$.

Figure 6 graphs the best response functions and the equilibrium outcome of this model for motorists of type (c, r) under the assumption that $t_r^w < t_r^a$ so that white officers find it less costly to search motorists from group r than do African-American officers. The notion here is that white officers have discriminatory preferences against group r .⁷ As one might expect, white officers will be more likely to search motorists from group r than will African-American officers in equilibrium. By contrast, if $t_r^a = t_r^w = t$, then officers from both groups will be equally likely to search motorists of type (c, r) . Thus, one clear implication of the model is that, in the absence of preference-based discrimination, there should be no difference in the rate at which officers from different racial group search drivers of type (c, r) . This insight forms the basis of the empirical strategy that we employ.

In independent work, Anwar and Fang (2005), whose paper we became aware of after developing the first draft of our paper, employ a similar test to distinguish between statistical discrimination and preference-based discrimination. Their test is like ours in that it requires information on both the race of the police officer and the race of the motorist and in that it focuses on search rates rather than on conditional guilt probabilities. However, since their model emphasizes the role of motorist characteristics that serve as noisy signals of motorist guilt, their test differs somewhat from ours. Interestingly, when applied to our data, both tests provide support for the existence of discriminatory preferences among police.

Empirical Strategy

In this section, we discuss how we test our model's prediction. Recall that our model predicts that, in the absence of preference-based discrimination, there should be no difference in the rate at which officers from different racial groups search motorists of type (c, r) . Thus, assuming that the officer's race, the motorist's race and c are known, this implication can be tested. Below we discuss what happens if c is unobserved. However, in order to establish the link between our model and our empirical strategy, it is useful to start with the case in which c is observed.

⁷Or, equivalently, that African-American officers exhibit favoritism towards group r .

In order to motivate the probit model that we employ, recall that officers of race j search drivers of race r and characteristics c with the following probability:

$$\Pr(\text{search}|j, c, r) = H [\pi(c, r) - t_r^j].$$

Note that equilibrium guilty probabilities $[\pi(c, r)]$ are independent of officer race, which is revealed to drivers in the model only after the decision over whether or not to carry contraband has been made; this independence is key to our identification strategy.⁸

In order to identify racial prejudice separately by officer race, we would ideally estimate the following fully-specified probit model:

$$\Pr(\text{search}|j, c, r) = H (\beta_0 + \beta_1 c + \beta_2 1[j = a] + \beta_3 1[r = a] + \beta_4 1[j = a, r = w] + \beta_5 1[j = w, r = a]),$$

Unfortunately, the parameters of this fully specified model cannot be estimated as the model is perfectly collinear. To see this, take the difference between the final two regressors:

$$\Delta = 1[j = a] \times 1[r = w] - 1[j = w] \times 1[r = a] = 1[j = a] - 1[r = a]$$

Thus, this difference (Δ) equals a linear combination of the first two regressors. Our inability to estimate this fully specified model is not surprising since, even if c is a constant, there are only four possible cases of driver / officer interactions but five parameters. We can, however, feasibly estimate the following restricted model:

$$\Pr(\text{search}|j, c, r) = H (\beta_0 + \beta_1 c + \beta_2 1[j = a] + \beta_3 1[r = a] + \beta_4 \text{mismatch}),$$

where $\text{mismatch} = 1[j = a, r = w] + 1[j = w, r = a]$ indicates a traffic stop in which the race of the officer differs from the race of the driver. Given that we cannot identify racial prejudice separately for African-American and white officers, we assume that they are equally prejudiced ($t_a^a - t_w^a = t_w^w - t_a^w$).⁹ Under this assumption, we can write the following relationships between the theoretical and empirical specifications for this restricted model:¹⁰

⁸Drivers do, however, know the distribution of officer race, which is assumed to be the same for all motorists.

⁹If blacks and whites are not equally prejudiced, then our estimates will uncover the average level of prejudice across black and white officers.

¹⁰In order to derive these relationships, consider the following four possible cases of driver/officer interactions for both the theoretical and empirical models: 1) $j = w, r = w$, 2) $j = a, r = a$, 3) $j = w, r = a$, 4) $j = a, r = w$. One can then show that $\beta_0 + \beta_1 c = \pi(c, w) + t_w^w$. Using this relationship, the three key parameters can then be solved for.

RELATIONSHIP	INTERPRETATION
$\beta_2 = t_w^w - t_a^a$	Cost differences by officer race
$\beta_3 = \pi(c, a) - \pi(c, w)$	Statistical discrimination
$\beta_4 = t_a^a - t_w^a = t_w^w - t_a^w$	Racial prejudice

With data on the race of both the driver and the officer, we can thus distinguish between racial profiling based upon statistical discrimination, which is captured by the coefficient on driver race (β_3), and racial profiling based upon prejudice, which is captured by the coefficient on mismatch (β_4).

An implicit assumption underlying this Probit formulation is that search rates are separable in driver characteristics (c) and driver race ($1[r = a]$). That is, officers do not condition on driver characteristics in a manner that differs between black and white drivers. While this formulation appears to be restrictive, it is straightforward to incorporate an interaction between driver characteristics and driver race ($c \times 1[r = a]$) into the econometric specification. Let β_5 be the coefficient on this interaction term. Then, the above relationships are identical except for the expression for statistical discrimination [$\pi(c, a) - \pi(c, w)$], which was previously equal to β_3 and now equals $\beta_3 + \beta_5 c$. Intuitively, any conditioning on driver characteristics that differs between African and white drivers should be considered statistical discrimination. Importantly, however, the interpretation of the coefficient on mismatch is unchanged [$\beta_4 = t_a^a - t_w^a = t_w^w - t_a^w$] and still captures racial prejudice. As noted below, the coefficient on mismatch is positive in our empirical application even after including these interactions between driver race and driver characteristics.

Consider next the case in which driver characteristics (c) are unobserved to the econometrician. We show below that, under assumptions of normality and random matching of officers and drivers, our approach retains the ability to distinguish between racial prejudice and statistical discrimination *even if unobserved driver characteristics are correlated with driver race*. Intuitively, the coefficient on driver race absorbs any unobserved differences between black and white drivers, and the coefficient on mismatch is thus not contaminated by the presence of these unobserved characteristics.

Recall that, according to the probit model, officers search if the following expression holds:

$$\beta_0 + \beta_1 c + \beta_2 1[j = a] + \beta_3 1[r = a] + \beta_4 \text{mismatch} - U > 0$$

where $U \sim N(0, 1)$. Assume next that unobserved driver characteristics are normally distributed with a mean that varies by race:

$$c = c_r - \sigma\varepsilon, \quad r = a, w$$

where $\varepsilon \sim N(0, 1)$ and is assumed to be independent of both driver race (r) and officer characteristics (U, j).¹¹ As shown in Yatchew and Griliches (1985), without the normality assumption, which is made here for reasons of tractability, the presence of unobserved characteristics leads to complicated asymptotic bias formulas in probit models.¹² We refer to the assumption of independence between unobserved driver characteristics and mismatch as random matching. This random matching assumption will be satisfied if ε is uncorrelated with j and r . We address the validity of this assumption in the empirical analysis to follow by studying how police officers are assigned to neighborhoods in Boston.¹³

Substituting in the above expression for unobserved driver characteristics, officers of race j search drivers of race r if:

$$\beta_0 + \beta_1 c_w + \beta_2 1[j = a] + [\beta_3 + \beta_1(c_a - c_w)]1[r = a] + \beta_4 \text{mismatch} - U - \sigma\beta_1\varepsilon > 0$$

Under the assumption that U and ε are independently distributed, $U - \sigma\beta_1\varepsilon \sim N(0, 1 + \beta_1^2\sigma^2)$ and the probability of search, unconditional on driver characteristics, is given as follows:

$$\Pr(\text{search}|j, r) = H \left[\frac{\beta_0 + \beta_1 c_w + \beta_2 1[j = a] + [\beta_3 + \beta_1(c_a - c_w)]1[r = a] + \beta_4 \text{mismatch}}{\sqrt{1 + \beta_1^2\sigma^2}} \right],$$

We can thus define the unconditional probit parameters ($\gamma_0, \gamma_2, \gamma_3, \gamma_4$) as follows:

¹¹The assumption that c is a scalar is not crucial and can be generalized. In particular, allow an $N \times 1$ vector of unobserved driver characteristics (C) to vary according to driver race and a random vector: $C = C_r - E$, where C_r and E are both $N \times 1$ vectors, and the components of E are assumed to be distributed jointly normal with covariance matrix Σ . In this case, the unconditional probit can be written as follows:

$$\Pr(\text{search}|j, r) = H \left[\frac{\beta_0 + \beta_1 C_w + \beta_2 1[j = a] + [\beta_3 + \beta_1(C_a - C_w)]1[r = a] + \beta_4 \text{mismatch}}{\sqrt{1 + \beta_1 \Sigma \beta_1'}} \right]$$

where β_1 is now a $1 \times N$ vector.

¹²In particular, the asymptotic bias formulas depend on the cumulative distribution function for unobserved characteristics. Applying this lesson to our analysis, if traffic stops in which the race of the driver differs from the race of the officer are also stops in which drivers disproportionately carry contraband, then the coefficient on mismatch could be asymptotically biased in either direction.

¹³Without the assumption of random matching, our empirical strategy no longer directly measures racial prejudice. Suppose, for example, that unobserved characteristics are correlated with mismatch so that $c = c_r + \eta \text{mismatch} - \sigma\varepsilon$. In this case, the probit specification is given as follows:

$$\Pr(\text{search}|j, r) = H \left[\frac{\beta_0 + \beta_1 c_w + \beta_2 1[j = a] + [\beta_3 + \beta_1(c_a - c_w)]1[r = a] + (\beta_4 + \beta_1\eta) \text{mismatch}}{\sqrt{1 + \beta_1^2\sigma^2}} \right]$$

Thus, the coefficient on mismatch will capture both racial prejudice (β_4) and non-random matching ($\beta_1\eta$).

$$\begin{aligned}\gamma_0 &= \frac{\beta_0 + \beta_1 c_w}{\sqrt{1 + \beta_1^2 \sigma^2}} \\ \gamma_2 &= \frac{\beta_2}{\sqrt{1 + \beta_1^2 \sigma^2}} \\ \gamma_3 &= \frac{\beta_3 + \beta_1(c_a - c_w)}{\sqrt{1 + \beta_1^2 \sigma^2}} \\ \gamma_4 &= \frac{\beta_4}{\sqrt{1 + \beta_1^2 \sigma^2}}\end{aligned}$$

Using these definitions and the relationships listed above between the theoretical parameters and the probit parameters conditional on driver characteristics, we can thus relate the probit parameters unconditional on driver characteristics to the theoretical parameters as follows:

RELATIONSHIP	INTERPRETATION
$\gamma_2 = (t_w^w - t_a^a)/\sqrt{1 + \beta_1^2 \sigma^2}$	Cost differences by officer race
$\gamma_3 = [\pi(c, a) - \pi(c, w) + \beta_1(c_a - c_w)]/\sqrt{1 + \beta_1^2 \sigma^2}$	Statistical discrimination
$\gamma_4 = (t_a^a - t_w^w)/\sqrt{1 + \beta_1^2 \sigma^2} = (t_w^w - t_a^a)/\sqrt{1 + \beta_1^2 \sigma^2}$	Racial prejudice

These relationships yield several key insights. First, results from the case in which the econometrician observes and does not observe driver characteristics are identical if officers do not rely on driver characteristics in their search decisions ($\beta_1 = 0$). In addition, if there is no heterogeneity other than race in unobserved characteristics ($\sigma = 0$), then the coefficients on officer race and mismatch are unchanged. The coefficient on driver race (γ_3), however, is altered and now captures both statistical discrimination based purely upon race $[\pi(c, a) - \pi(c, w)]$ and statistical discrimination based upon driver characteristics that vary according to race ($\beta_1(c_a - c_w)$); without further information, we cannot distinguish between these two forms of statistical discrimination. However, even if $\beta_1 \neq 0$ and $\sigma \neq 0$, our approach retains the ability to distinguish between statistical discrimination, in whatever form it may take, and racial prejudice ($\gamma_4 = (t_a^a - t_w^w)/\sqrt{1 + \beta_1^2 \sigma^2}$). In fact, the presence of unobserved driver characteristics only serves to bias our analysis away from measuring racial prejudice due to the scaling factor ($\sqrt{1 + \beta_1^2 \sigma^2}$), which exceeds one.

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 required 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 merely issued a written or 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 carrying contraband. If so, then our estimates will tend to *understate* the extent of the racial bias in search patterns. We will address this data limitation and related sample selection issues later in the analysis by restricting our sample to the two-month period that includes data on warnings and by separately examining stops conducted at night when officers are less likely to know a driver's race prior to pulling them over.

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 percent of the citations issued by officers in the Boston Police Department in our data. That is, for approximately 20 percent 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 the 6 citations for which contradictory race information was recorded. In addition, we drop citations issued by Asian officers (23 officers in total), and 7,732 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. We also drop the ten citations that were issued to motorists outside the city of Boston. This may have happened, for example, if an officer followed a speeding driver outside of the city limits. Finally, we drop about 4,500 observations with missing information

on the race, age and residence of the driver and whether an accident occurred. Once these restrictions have been made we are left with 95,855 citations issued by 1,317 officers.

Of considerable concern is the fact that the search variable is missing for over 18 percent 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. There is substantial 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 see if 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 when the search variable is missing, no search was conducted. Then, we drop all officers for whom search is missing for more than 10 percent of the citations that those officers issue. Doing so eliminates approximately 24 percent of the citations (and 48 percent of the officers) in our data.¹⁴ For the remaining 685 officers, we drop observations for which search is missing, and are left with a sample comprising 70,652 citations. Tables 1 and 2 are calculated using this sample.

Table 3 presents some basic summary statistics. The first column includes all of the citations for which our baseline search measure is missing, whereas the second column includes all of the 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, citations for which the search variable is missing appear to be quite similar to those for which it is not.¹⁵ The last three columns of Table 3 show the average characteristics

¹⁴In calculating the percentage of citations for which search is missing, we do not include citations in which the race of the driver is missing.

¹⁵We 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

of the citations in our sample broken down by the race of the officer issuing the citation. We see that drivers are disproportionately issued citations by officers 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.

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 26, 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 indicator variables that control for the district in which the stop occurred. In Table 4 (and in all remaining relevant tables) we report the coefficients of our probit model. 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 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.

coefficient were statistically significant, this result 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.

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 also present robustness checks using the unweighted probits. First, describing search outcomes for the average motorist is interesting in its own right. Second, 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 Table 4 indicates, black drivers are more likely to have their cars searched than are white drivers. This result also holds for the weighted probit in column 3. Like Knowles, Persico and Todd, we find no evidence that the probability of guilt conditional on search differs by the race of the driver. In particular, in both columns 2 and 4, the coefficient of the indicator variable for whether the driver is 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 is Hispanic. The results are very similar to those in Table 4 though there is some evidence that black motorists are less likely to be found with contraband than are white motorists.

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 model of KPT is generalized to allow for heterogeneity in officer search costs, this prediction no longer holds.

However, our model delivers an alternative method for distinguishing between preference-based discrimination and statistical 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 there should be no difference in the rate at which officers from different racial groups search drivers from any given racial group. To examine

this, 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. 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 significance 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. Our results also suggest that Hispanic officers are more likely to conduct searches than are white officers, and the second and third columns suggest that officers are more likely to search motorists who 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 on mismatch is positive and significantly different from zero at standard confidence levels.¹⁶

As mentioned previously, our estimates may be biased if the mismatch between the race of the officer and the race of the motorist is correlated with motorist characteristics (c) that are not included in our regressions. Thus, it is comforting that the point estimate on mismatch changes very little as we add more regressors, suggesting that mismatch tends not to be correlated with unobserved motorist characteristics. We investigate this potential bias more fully below by empirically analyzing the assignment of officers to neighborhoods in Boston.

A positive coefficient on our mismatch variable 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. As noted above, however, we cannot identify differences in racial prejudice by officer race. 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. Rather, our results simply indicate that the interaction between the race of the motorist and the race of the officer is positively related to the probability that the motorist is searched, a pattern that is consistent with preference-based discrimination.

¹⁶As noted in the empirical strategy section, we have also estimated models in which the effect of driver characteristics is allowed to vary between black and white drivers. These results, not reported here, are similar to those in the baseline analysis. That is, after controlling for the interaction between driver race and driver characteristics, the coefficient on mismatch remains positive and statistically significant.

While the coefficient on mismatch is statistically significant in all of the specifications of Table 7, we should also evaluate the magnitude of these results. Conventional computation of the marginal effect of the variable mismatch is somewhat nonsensical as mismatch is an interaction term and thus changes in mismatch also require changes in either officer or driver race. Thus, in order to gauge the magnitude of these results, we calculate the change in the probability of search associated with a change in officer race, holding driver race constant. For example, using the coefficients in column 6 of Table 7, we calculate that black drivers are 3.74 percentage points more likely to be searched by a white officer, relative to a black officer, evaluated at the average characteristics of black drivers. These effects are large relative to the baseline weighted search rate of 3.10, but are similar in magnitude to the raw averages, in which blacks are searched by white officers at rate 5.85 percent but by black officers at rate 1.95 percent, as reported in Table 2. We calculate that white drivers are also more likely to be searched, an increase of 0.88 percentage points, by a black officer, relative to a white officer, where this effect is evaluated at the average characteristics of white drivers.¹⁷

As mentioned earlier, the search variable is missing for over 18 percent of the citations in our data. To see whether our results are 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 officers for whom the search variable is missing in more than 10 percent of the citations they issue. In the second column, we repeat the analysis in column 1 but 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 3, 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 in both column 2 and column 3. Table 9 repeats the analysis in Table 8, except that it does not include stops that involve either Hispanic officers or Hispanic motorists. Results are very similar to those from Table 8.

¹⁷Consistent with the raw averages in Table 1, the effect of changes in officer race associated with the unweighted results, those of column 3 of Table 7, are smaller in magnitude. Black drivers are 0.15 percentage points more likely to be searched by white officers, relative to black officers, while white driver are 0.21 percentage points more likely to be searched by black officers, relative to white officers. For comparison purposes, the baseline unweighted search rate is 0.65 percent, as reported in Table 1.

Recall that in our baseline search measure we drop officers for whom the search variable is missing for more than 10 percent of the citations issued by that officer. We have also experimented with changing that 10 percent cutoff. Lowering the cutoff (to say 5 percent or 3 percent), 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 percent (all officers are included).

Approximately 82 percent 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). In Table 10, we repeat the analysis in Table 6, but restrict attention to citations that are issued by Patrol Officers. While the coefficients remain positive, they are now statistically insignificant, suggesting that non-patrol officers may be more important in explaining the baseline results. One possible explanation for this result is that non-patrol officers tend to be more experienced and may thus have more previous exposure to drivers of different races; we examine in more detail these potential differences between experienced and non-experienced officers below.

Sample Selection Bias

One concern is that the set of drivers whom officers pull over and issue citations to is non-random. For example, as mentioned above, 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 carrying contraband. If so, then our estimates in Table 6 will tend to *understate* the extent of the racial bias in search patterns.

We empirically address the possibility that our results are driven by selection on officers' decisions to issue citations in two ways. First, as noted above, our data include written warnings for the two-month period of April-May 2001. If officers tend to issue more warnings to drivers of their own race, then our data after this two-month period may include only the own-race interactions in which drivers had committed the most severe infraction.¹⁸ To address this concern, in Table 11, we restrict our sample to stops that occur within this

¹⁸We 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 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.

two-month period. As expected, the coefficients on mismatch are 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 conventional levels.

Second, as suggested by Grogger and Ridgeway (2004), we examine citations issued at night. The idea is that officers are less likely to know a motorist’s race prior to pulling them over when it is dark outside and thus selection into stops may be less problematic during these stops. These results are presentend in Table 12. As shown, for the sample of stops occuring only at night, the coefficient on mismatch is larger than in the baseline analysis in Table 6 and is significant at the 99 percent level.¹⁹

Taken together, these robustness checks demonstrate that our findings of preference-based discrimination are not driven by selection on the officer’s decision to issue warnings or tickets or to deliberately stop certain types of drivers. In the remainder of the paper, we address other alternative explanations for our findings: informational asymmetries and the non-random assignment of officers to neighborhoods.

Asymmetric Search Ability

One concern is that our results may be driven by the fact that officers may be more successful at finding contraband in cars that are driven by motorists who are in their own racial group. For example, Donohue and Levitt (2001) find evidence that own-race policing may be more efficient than cross-race policing. To see how this would affect our model, let ϕ_r^j denote the probability that an officer from group j is successful of searching a motorist from group r ; the baseline model is one in which this probability equals one for all officers and drivers. In this generalized model, the payoff to an officer from group j to searching a motorist from group r is given by

$$\phi_r^j P(G|c, r) - t_r - U$$

The higher is ϕ_r^j , the *higher* will be the benefit to officers from group j of searching a motorist from group r . Thus, if officers are better at finding contraband when the motorist is a member of the officer’s own racial group, then we would expect officers to be *more* likely to search motorists from their own racial group and our estimates will tend to *understate* the extent of preference-based discrimination.²⁰ This result is shown graphically in Figure 7.

¹⁹We also attempted to examine drivers who were pulled over for going more than 15mph and more than 20mph over the posted speed limit, since the police arguably have less discretion in whether or not to pull over these motorists. However, the results were sensitive to how we handled the large number of observations with missing information on mph and so we do not report them here.

²⁰Like Anwar and Fang (2005) and Bjerck (2005) we have also considered a model in which officers observe a noisy signal of a motorist’s guilt that is unknown to motorists at the time they make their decision about

To empirically address the possibility that asymmetric information drives our results, we examine whether our results hold among officers with more than 10 years of experience. The idea is that if officers become better at searching motorists from a particular group as their exposure to that group increases, then assuming there are decreasing returns to experience, officers with substantial experience should be equally able to search the cars of motorists from different racial groups.

Table 13 presents the results; 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 10 or more years of experience. We chose 10 years 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 shown, 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 suggest that our findings of preference-based discrimination are *not* driven by differences in the ability of officers to accurately search the cars of motorists from their own racial group. Rather, this analysis suggests that our results are stronger when we examine only experienced officers, for whom we would expect the likelihood of a successful search to be independent of the match between the officer's race and the driver's race.

Unobserved Driver Characteristics and the Assignment of Officers to Neighborhoods

As discussed in the section on the econometric specification, our approach assumes random matching of motorists and officers. If there is some relevant characteristic, c , that is not included in our regressions and that is correlated with mismatch, then this assumption may be violated and our estimate of β_4 , the coefficient on mismatch, may be biased. One of the most plausible explanations for the source of this bias is that officers may not be randomly assigned to different neighborhoods within the city.

Suppose, for example, that white officers are disproportionately assigned to neighborhoods in which blacks commit a large fraction of the drug trafficking offenses. We would expect white officers to be more likely than black officers to search the cars of black motorists, even in the absence of preference-based discrimination. It seems unlikely that officers would be

whether to carry contraband, but that is correlated with the likelihood that they carry contraband. One might expect officers to receive more informative signals from motorists who are in the officer's racial group than from those who are not. In an appendix available from the authors upon request, we show that changes in the information content of the signal deliver ambiguous predictions about search behavior.

assigned to neighborhoods in this fashion, but 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 8 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 percent 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. 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.

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 district 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 pattern holds for Hispanics. For whites, the correlation between the fraction of the population aged 18 and older in each district who is white and the fraction of officers in that district who is white is 0.751. For blacks, Asians and Hispanics the analogous correlations are 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.

However, even if officers tend to be assigned to districts in which the majority of residents are members of their own racial group, mismatch may be inherently correlated with the likelihood that the motorist is guilty. For example, if African-Americans go to predominantly white neighborhoods to sell drugs and if whites go to predominantly African-American neighborhoods to buy drugs, then officers in predominantly white neighborhoods may deliberately target African-American motorists and officers in black predominantly black neighborhoods may deliberately target white motorists. As a result, a mismatch between the race of the motorist and the race of the officer may be correlated with the likelihood that a motorist is carrying contraband.

In order to address this concern, Table 15 provides results from only those traffic stops in which the motorist was a resident of the district in which he or she was pulled over. Presumably, when an officer stops a motorist, the officer can observe (by looking at the motorist’s driver’s license) whether the motorist is a resident of the district and, thus, whether he or she is a “suspicious outsider”. As the Table reveals, the coefficient on mismatch remains positive and is statistically different from zero in all three specifications.²¹ Thus, even if we focus on stops that take place in the district in which the motorist resides, we find that officers are more likely to conduct a search if the motorist is not a member of the officer’s racial group.

In addition, in Table 16, we break districts down into two categories: those that are racially diverse and those that are racially homogeneous. The idea is that in racially diverse neighborhoods, a driver’s race is less likely to signal that the driver is out of place. Using Table 14, we categorize East Boston, Roxbury/Mission Hill, Jamaica Plain and Hyde Park as diverse, while we categorize Boston Central (A-1 and D-4), South Boston, Allston/Brighton and West Roxbury/Roslindale as homogeneous.²² We drop citations that were issued in Dorchester and North Dorchester because these districts differ substantially in their racial composition and our data do not allow us to distinguish between the two. Interestingly, we find that the results are somewhat stronger in diverse neighborhoods, suggesting that our baseline results are not driven by the possibility that police target drivers whose race differs from that of the local population.²³

Thus, Tables 13-16 provide evidence that our results are not driven by the possibility that officers are better able to search members of their own racial group or by the non-random assignment of officers to neighborhoods.

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 than 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 model of police search developed in Knowles, Persico and

²¹Our results are qualitatively similar when we only examine blacks and whites.

²²Charlestown is included in East Boston.

²³Our results are qualitatively similar when we only examine blacks and whites.

Todd (2001) to include heterogeneity in officer search costs; under this generalized model, the fundamental insight that allows KPT to distinguish between preference-based discrimination and statistical discrimination falls away. 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 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 to search members of their own racial group. We find no evidence that this sort of search asymmetry explains our results. We also show that the manner by which officers are assigned to neighborhoods within the city does not account for our empirical findings. Rather, our results suggest that preference-based discrimination plays a substantial role in explaining differences in the rate at which motorists from different racial groups are searched during traffic stops.

References

- [1] Altonji, Joseph G. and Pierret, Charles R. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics*, February 2001, 116(1), 313-350.
- [2] Anwar, Shamena and Fang, Hanming "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence." forthcoming, *American Economic Review*.
- [3] Arrow, Kenneth "The Theory of Discrimination," in O. Ashenfelter and A. Rees, eds., *Discrimination in Labor Markets*. Princeton, NJ: Princeton University Press, 1973, 3-33.
- [4] Becker, Gary. *The Economics of Discrimination*. Chicago, IL: The University Chicago Press, 1957.
- [5] Bjerck, David. "Racial Profiling, Statistical Discrimination, and the Effect of a Colorblind Policy on the Crime Rate." *Mimeo, McMaster University*, March 2005.
- [6] Dedman, Bill and Latour, Francie. "Traffic Citations Reveal Disparity in Police Searches." *The Boston Globe*, January 6, 2003, pp.A1.
- [7] Dedman, Bill and Latour, Francie. "Race, Sex and Age Drive Ticketing." *The Boston Globe*, July 20, 2003.
- [8] Dedman, Bill and Latour, Francie. "Punishment Varies by Town and Officer." *The Boston Globe*, July 21, 2003.
- [9] Dedman, Bill and Latour, Francie. "Troopers Fair, Tough in Traffic Encounters." *The Boston Globe*, July 22, 2003.
- [10] Dharmapala, Dhammika and Ross, Stephen L. "Racial Bias in Motor Vehicle Searches: Additional Theory and Evidence." Forthcoming, *Contributions to Economic Analysis and Policy*, 2004.
- [11] Dominitz, Jeff. "How Do the Laws of Probability Constrain Legislative and Judicial Efforts to Stop Racial Profiling?" *American Law and Economic Review*, August 2003, 5(2), 412-432.
- [12] Donohue, John J. and Levitt, Steven D. "The Impact of Race on Policing and Arrests." *Journal of Law and Economics*, October 2001, 44, 367-394.
- [13] Grogger, Jeffrey and Ridgeway, Greg. "Testing for Racial Profiling in Traffic Stops from Behind a Veil of Darkness." *Mimeo*, 2004.

- [14] Harsanyi, J. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points." *International Journal of Game Theory* 1973, 2, 1-23.
- [15] Hernández-Murillo, Ruben and Knowles, John. "Racial Profiling or Racist Policing?: Testing in Aggregated Data." Forthcoming, *International Economic Review*, 2004.
- [16] Knowles, John; Persico Nicola; and Todd, Petra. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy*, February 2001, 109(11), 203-229.
- [17] Persico, Nicola and Todd, Petra. "Using Hit Rate Tests to Test for Racial Bias in Law Enforcement: Vehicle Searches in Wichita," NBER Working Paper 10947, December 2004.
- [18] Yatchew, Adonis and Griliches, Zvi. "Specification Error in Probit Models." *Review of Economic Studies*, February 1985, 67(1), 134-39.

Figure 1: Best Response Functions in KPT Model




Figure 2: Best Response Functions with Motorist Heterogeneity




Figure 3: Preference-Based Discrimination in KPT Model

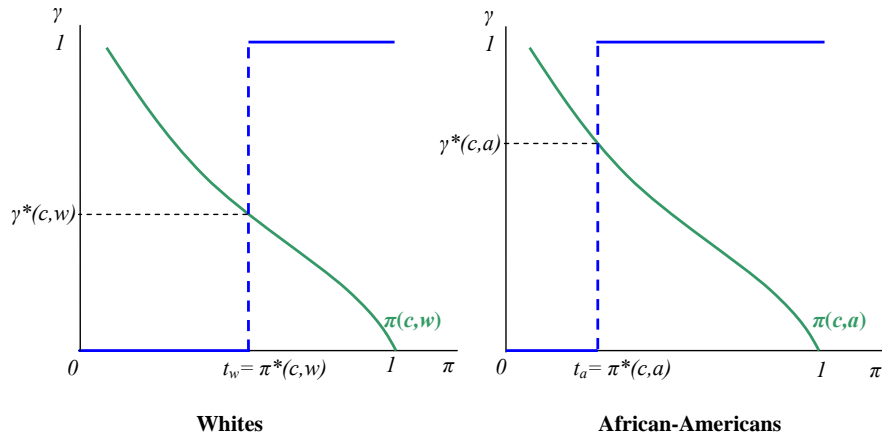


Figure 4: Statistical Discrimination in KPT Model




Figure 5: Statistical Discrimination with Officer Heterogeneity




Figure 6: Preference-Based Discrimination with Officer Heterogeneity




Figure 7: Asymmetric Search Ability

(Assumes $t_r^w = t_r^a$, but that $\phi_w^w > \phi_w^a$)


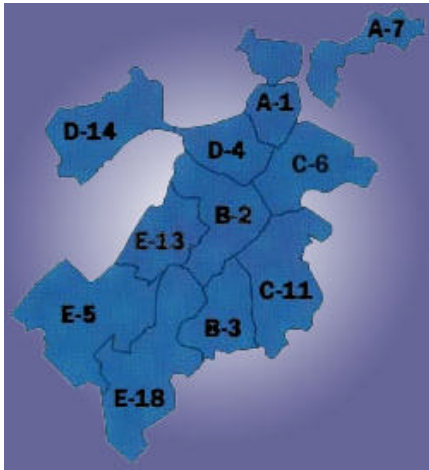


Figure 8: City of Boston, Police Districts



- 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 1: Probability of Search by Officer Race and Driver Race
(Standard Deviation of Sample Mean in Parentheses)

<i>Driver Race</i>	<i>Officer Race</i>			
	White	Black	Hispanic	All
White	0.40% (0.04%) n=22471	0.62% (0.07%) n=11132	0.25% (0.09%) n=3256	0.46% (0.04%) n=36859
Black	0.97% (0.09%) n=13131	0.82% (0.09%) n=9116	0.49% (0.15%) n=2258	0.87% (0.06%) n=24505
Hispanic	0.97% (0.14%) n=5058	0.82% (0.16%) n=3164	0.38% (0.19%) n=1066	0.85% (0.10%) n=9288
All	0.65% (0.04%) n=40660	0.73% (0.06%) n=23412	0.35% (0.07%) n=6580	0.65% (0.03%) n= 70652

Note: 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 not included.

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)

<i>Driver Race</i>	<i>Officer Race</i>			
	White	Black	Hispanic	All
White	1.18% (0.47%) n=404	2.59% (0.66%) n=138	2.31% (1.91%) n=46	1.99% (0.39%) n=588
Black	5.85% (1.18%) n=361	1.95% (0.69%) n=135	0.49% (0.21%) n=42	4.46% (0.82%) n=538
Hispanic	4.05% (1.41%) n=265	4.74% (2.42%) n=110	0.30% (0.17%) n=37	3.86% (1.11%) n=412
All	3.45% (0.52%) n=470	2.64% (0.53%) n=163	1.37% (0.96%) n=52	3.10% (0.38%) n=685

Note: 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 not included. 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		Primary Sample		
	Missing	All Officers	All Officers	White Officers	Black Officers
White Driver	57.3% (49.5%)	52.2% (50.0%)	55.3% (49.7%)	47.5% (49.9%)	49.5% (50.0%)
Black Driver	30.8% (46.2%)	34.7% (47.6%)	32.3% (46.8%)	38.9% (48.8%)	34.3% (47.5%)
Hispanic Driver	11.9% (32.4%)	13.1% (33.8%)	12.4% (33.0%)	13.5% (34.2%)	16.2% (36.8%)
Mismatch	49.1% (50.0%)	53.8% (49.9%)	49.7% (49.7%)	61.1% (48.8%)	83.8% (36.8%)
Baseline Search	-	0.7% (8.0%)	0.7% (8.1%)	0.7% (8.5%)	0.3% (5.9%)
Stop at Night	30.4% (46.0%)	30.4% (46.0%)	26.6% (44.2%)	36.9% (48.3%)	30.4% (46.0%)
Young Driver (Age<26)	24.7% (43.1%)	24.2% (42.8%)	24.2% (42.8%)	23.8% (42.6%)	26.2% (44.0%)
Male Driver	71.8% (45.0%)	68.1% (46.6%)	69.4% (46.1%)	65.5% (47.5%)	69.8% (45.9%)
In-State Driver	93.8% (24.1%)	93.6% (24.6%)	93.3% (25.1%)	94.1% (23.6%)	93.4% (24.7%)
In-Town Driver	48.7% (50.0%)	51.1% (50.0%)	49.0% (50.0%)	54.3% (49.8%)	52.4% (49.9%)
Accident	2.5% (15.7%)	1.3% (11.3%)	1.5% (12.0%)	0.9% (9.7%)	1.3% (11.5%)
Allston-Brighton	7.5% (26.4%)	6.7% (25.0%)	8.2% (27.5%)	4.6% (21.0%)	4.6% (21.0%)
Boston Central	20.4% (40.3%)	13.1% (33.7%)	12.9% (33.5%)	12.1% (32.6%)	18.0% (38.4%)
Charlestown-East Boston	10.4% (30.5%)	6.3% (24.2%)	8.5% (27.9%)	3.2% (17.5%)	3.2% (17.7%)
Dorchester-Mattapan	21.8% (41.3%)	20.1% (40.0%)	20.6% (40.4%)	20.2% (40.1%)	16.5% (37.1%)
Hyde Park	0.7% (8.6%)	0.9% (9.3%)	0.7% (8.3%)	1.3% (11.3%)	0.4% (6.4%)
Jamaica Plain	2.4% (15.3%)	2.5% (15.7%)	3.2% (17.5%)	0.4% (6.2%)	6.1% (23.9%)
Roslindale	0.5% (7.4%)	1.1% (10.6%)	1.3% (11.2%)	1.0% (10.1%)	0.7% (8.2%)
Roxbury	13.0% (33.6%)	17.6% (38.1%)	18.6% (38.9%)	15.1% (35.8%)	20.0% (40.0%)
South Boston	6.0% (23.7%)	4.0% (19.7%)	4.7% (21.1%)	3.5% (18.4%)	1.8% (13.2%)
Number of Officers	922	685	470	163	52
Number of Citations	25,203	70,652	40,660	23,412	6,580

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

	Unweighted Probits		Weighted Probits	
	Search	Guilt	Search	Guilt
Black Driver	0.213*** (0.059)	-0.472 (0.388)	0.387*** (0.144)	-0.622 (0.464)
Hispanic Driver	0.144 (0.108)	-0.228 (0.409)	0.219 (0.163)	0.262 (0.452)
Stop at Night	0.154 (0.101)	0.012 (0.329)	0.201* (0.116)	-0.487 (0.349)
Young Driver (Age<26)	0.087** (0.038)	-0.314 (0.236)	0.110 (0.129)	-0.413 (0.367)
Male Driver	0.064 (0.046)	-0.188 (0.261)	0.096 (0.123)	-0.062 (0.365)
In-State Driver	0.084 (0.092)		0.246 (0.194)	
In-Town Driver	0.028 (0.036)	-0.030 (0.335)	0.032 (0.105)	0.045 (0.402)
Accident	0.854*** (0.153)	-0.138 (0.433)	0.022 (0.188)	0.481 (0.531)
Neighborhood Controls	YES	YES	YES	YES
Observations	70652	369	70652	369

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”. 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	0.181*** (0.057)	-0.460 (0.365)	0.378*** (0.140)	-0.789* (0.414)
Stop at Night	0.115 (0.106)	0.100 (0.359)	0.172 (0.127)	0.568 (0.430)
Young Driver (Age<26)	0.066 (0.046)	-0.445 (0.372)	0.027 (0.141)	-1.055* (0.564)
Male Driver	0.055 (0.049)	-0.414* (0.237)	0.052 (0.131)	-1.323*** (0.374)
In-State Driver	0.001 (0.091)		0.120 (0.194)	
In-Town Driver	0.077* (0.040)	-0.041 (0.325)	0.131 (0.121)	0.116 (0.359)
Accident	0.896*** (0.173)	-0.077 (0.434)	0.020 (0.209)	-0.797 (0.542)
Neighborhood Controls	YES	YES	YES	YES
Observations	55850	334	55850	334

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”. 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	0.207*** (0.059)	0.175*** (0.059)	0.190*** (0.057)	0.167 (0.126)	0.144 (0.126)	0.204 (0.142)
Hispanic Driver	0.173 (0.115)	0.120 (0.118)	0.083 (0.106)	0.061 (0.166)	0.023 (0.174)	-0.006 (0.176)
Black Officer	0.027 (0.185)	0.044 (0.188)	0.058 (0.165)	-0.134 (0.134)	-0.115 (0.134)	-0.085 (0.135)
Hispanic Officer	-0.251 (0.182)	-0.260 (0.180)	-0.225 (0.176)	-0.487* (0.279)	-0.511* (0.269)	-0.501** (0.249)
Mismatch	0.099* (0.057)	0.110* (0.057)	0.124** (0.059)	0.354*** (0.126)	0.355*** (0.125)	0.345*** (0.121)
Stop at Night		0.144 (0.104)	0.156 (0.107)		0.207* (0.123)	0.208* (0.117)
Young Driver (Age<26)		0.087** (0.038)	0.092** (0.038)		0.101 (0.128)	0.106 (0.126)
Male Driver		0.077 (0.054)	0.069 (0.044)		0.100 (0.128)	0.088 (0.122)
In-State Driver		0.116 (0.096)	0.084 (0.093)		0.255 (0.182)	0.254 (0.185)
In-Town Driver		0.013 (0.046)	0.029 (0.036)		-0.015 (0.099)	0.025 (0.105)
Accident		0.867*** (0.144)	0.863*** (0.151)		0.036 (0.179)	0.018 (0.188)
Neighborhood Controls	NO	NO	YES	NO	NO	YES
Observations	70652	70652	70652	70652	70652	70652

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	0.206*** (0.061)	0.159** (0.064)	0.159*** (0.057)	0.240* (0.133)	0.187 (0.133)	0.194 (0.149)
Black Officer	0.044 (0.182)	0.061 (0.186)	0.049 (0.166)	-0.167 (0.142)	-0.155 (0.139)	-0.146 (0.139)
Mismatch	0.104* (0.061)	0.112* (0.060)	0.118* (0.060)	0.335** (0.133)	0.338*** (0.131)	0.339*** (0.126)
Stop at Night		0.109 (0.112)	0.115 (0.111)		0.175 (0.134)	0.174 (0.126)
Young Driver (Age<26)		0.058 (0.048)	0.065 (0.046)		0.006 (0.142)	0.013 (0.139)
Male Driver		0.052 (0.053)	0.057 (0.047)		0.041 (0.137)	0.036 (0.129)
In-State Driver		0.019 (0.094)	-0.000 (0.092)		0.098 (0.184)	0.110 (0.184)
In-Town Driver		0.082** (0.041)	0.077* (0.040)		0.113 (0.112)	0.131 (0.120)
Accident		0.896*** (0.169)	0.906*** (0.171)		0.033 (0.207)	0.039 (0.209)
Neighborhood Controls	NO	NO	YES	NO	YES	YES
Observations	55850	55850	55850	55850	55850	55850

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

Table 8: Probability of Search, Robustness Checks

	Search 1	Search 2	Search 3
Black Driver	0.284** (0.111)	0.192* (0.106)	0.198* (0.104)
Hispanic Driver	0.068 (0.128)	0.035 (0.130)	0.021 (0.129)
Black Officer	0.022 (0.119)	-0.073 (0.106)	-0.052 (0.106)
Hispanic Officer	-0.111 (0.161)	-0.289* (0.157)	-0.128 (0.156)
Mismatch	0.109 (0.106)	0.259*** (0.091)	0.189** (0.091)
Stop at Night	0.239*** (0.090)	0.195** (0.089)	0.217** (0.084)
Young Driver (Age<26)	0.183** (0.092)	0.162* (0.091)	0.242*** (0.088)
Male Driver	0.165 (0.102)	0.128 (0.101)	0.115 (0.098)
In-State Driver	0.099 (0.172)	0.156 (0.140)	0.075 (0.168)
In-Town Driver	0.110 (0.093)	0.033 (0.084)	0.118 (0.085)
Accident	0.079 (0.162)	-0.028 (0.133)	-0.045 (0.141)
Neighborhood Controls	YES	YES	YES
Observations	79337	95855	79369

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**

	Search 1	Search 2	Search 3
Black Driver	0.267** (0.117)	0.183 (0.118)	0.183 (0.113)
Black Officer	-0.047 (0.121)	-0.149 (0.111)	-0.132 (0.109)
Mismatch	0.093 (0.112)	0.260*** (0.099)	0.191* (0.099)
Stop at Night	0.210** (0.101)	0.135 (0.100)	0.163* (0.094)
Young Driver (Age<26)	0.093 (0.105)	0.078 (0.104)	0.165* (0.100)
Male Driver	0.105 (0.109)	0.083 (0.109)	0.057 (0.106)
In-State Driver	0.014 (0.166)	0.030 (0.140)	-0.018 (0.163)
In-Town Driver	0.125 (0.102)	0.087 (0.094)	0.143 (0.095)
Accident	0.034 (0.187)	-0.077 (0.151)	-0.079 (0.157)
Neighborhood Controls	YES	YES	YES
Observations	62539	74837	62560

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	0.038 (0.135)	0.034 (0.137)	0.142 (0.152)
Hispanic Driver	0.134 (0.163)	0.109 (0.168)	0.062 (0.158)
Black Officer	-0.114 (0.144)	-0.086 (0.144)	-0.042 (0.141)
Hispanic Officer	-0.372 (0.278)	-0.359 (0.267)	-0.341 (0.246)
Mismatch	0.185 (0.133)	0.195 (0.134)	0.178 (0.128)
Stop at Night		0.147 (0.131)	0.139 (0.124)
Young Driver (Age<26)		-0.057 (0.115)	-0.053 (0.108)
Male Driver		0.169 (0.112)	0.166 (0.104)
In-State Driver		0.181 (0.181)	0.164 (0.180)
In-Town Driver		-0.049 (0.097)	-0.021 (0.111)
Accident		0.298 (0.187)	0.301 (0.197)
Neighborhood Controls	NO	NO	YES
Observations	69379	69379	69379

Robust standard errors in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

Table 11: Probability of Search, April-May 2001

	Weighted Probits		
	(1)	(2)	(3)
Black Driver	-0.319 (0.200)	-0.434* (0.234)	-0.479** (0.237)
Hispanic Driver	0.219 (0.264)	0.095 (0.296)	0.060 (0.296)
Black Officer	-0.032 (0.222)	0.009 (0.218)	0.019 (0.219)
Hispanic Officer	-0.191 (0.412)	-0.191 (0.379)	-0.186 (0.379)
Mismatch	0.369* (0.200)	0.411** (0.205)	0.412** (0.206)
Stop at Night		0.372* (0.196)	0.384* (0.199)
Young Driver (Age<26)		-0.146 (0.181)	-0.146 (0.181)
Male Driver		0.380* (0.202)	0.381* (0.202)
In-State Driver		0.018 (0.326)	0.046 (0.331)
In-Town Driver		0.065 (0.205)	0.058 (0.205)
Accident		0.540* (0.315)	0.535* (0.315)
Neighborhood Controls	NO	NO	YES
Observations	15413	15413	13689

Robust standard errors in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

Note: The sample size is lower in the last column because no search was conducted in some neighborhoods during this time period.

Table 12: Probability of Search, Stops at Night

	Weighted Probits		
	(1)	(2)	(3)
Black Driver	-0.001 (0.158)	0.050 (0.173)	0.000 (0.176)
Hispanic Driver	0.080 (0.226)	0.090 (0.246)	-0.002 (0.260)
Black Officer	0.042 (0.176)	0.050 (0.173)	0.029 (0.168)
Hispanic Officer	-0.347 (0.254)	-0.349 (0.250)	-0.395 (0.247)
Mismatch	0.425*** (0.155)	0.447*** (0.148)	0.489*** (0.149)
Young Driver (Age<26)		0.079 (0.146)	0.105 (0.140)
Male Driver		0.050 (0.171)	0.033 (0.166)
In-State Driver		0.377 (0.251)	0.377 (0.253)
In-Town Driver		-0.248 (0.155)	-0.244* (0.145)
Accident		0.309 (0.205)	0.283 (0.218)
Neighborhood Controls	NO	YES	YES
Observations	21467	21467	21467

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	-0.127 (0.130)	-0.087 (0.138)	-0.028 (0.151)	0.416** (0.201)	0.348* (0.201)	0.468** (0.219)
Hispanic Driver	0.090 (0.172)	0.100 (0.168)	0.106 (0.184)	0.073 (0.281)	-0.057 (0.286)	-0.108 (0.279)
Black Officer	-0.043 (0.174)	-0.013 (0.170)	-0.062 (0.157)	-0.172 (0.190)	-0.172 (0.189)	-0.085 (0.196)
Hispanic Officer	-0.755*** (0.170)	-0.780*** (0.170)	-0.758*** (0.180)	-0.109 (0.430)	-0.153 (0.425)	-0.152 (0.392)
Mismatch	0.165 (0.134)	0.210 (0.131)	0.200 (0.127)	0.494** (0.197)	0.486** (0.192)	0.474** (0.184)
Stop at Night		0.463*** (0.171)	0.460*** (0.157)		0.023 (0.161)	0.017 (0.162)
Young Driver (Age<26)		-0.220* (0.118)	-0.186 (0.115)		0.354** (0.176)	0.347** (0.176)
Male Driver		0.006 (0.119)	0.012 (0.111)		0.191 (0.194)	0.181 (0.185)
In-State Driver		0.545** (0.219)	0.560** (0.229)		-0.079 (0.257)	-0.082 (0.259)
In-Town Driver		-0.261** (0.113)	-0.277** (0.125)		0.220 (0.153)	0.271* (0.160)
Accident		0.163 (0.168)	0.133 (0.163)		0.037 (0.276)	0.058 (0.291)
Neighborhood Controls	NO	NO	YES	NO	NO	YES
Observations	33249	33249	33249	37403	37403	37094

Robust standard errors in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

Note: There are fewer observations in the last column because there is one neighborhood (Jamaica Plain) in which no search was conducted by an officer with more than 10 years of experience.

Table 14: Racial Composition of Police Districts

Census Benchmark					Citation-Level Data			
Population 18 and Older					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%

Source: The Census numbers are taken from the "Massachusetts Racial and Gender Profiling Project: Preliminary Tabulations," prepared by Northeastern University, Institute on Race and Justice. The citation-level numbers were derived from our data.

Table 15: Probability of Search, Drivers Stopped in Their Own Neighborhood Weighted Probits

	(1)	(2)	(3)
Black Driver	-0.017 (0.190)	-0.050 (0.191)	0.242 (0.259)
Hispanic Driver	0.251 (0.242)	0.198 (0.244)	0.259 (0.256)
Black Officer	-0.219 (0.159)	-0.210 (0.159)	-0.146 (0.154)
Hispanic Officer	-1.161*** (0.292)	-1.136*** (0.296)	-1.119*** (0.318)
Mismatch	0.499*** (0.189)	0.499*** (0.184)	0.491*** (0.173)
Stop at Night		-0.065 (0.161)	-0.039 (0.156)
Young Driver (Age<26)		0.244 (0.174)	0.275* (0.165)
Male Driver		-0.014 (0.179)	-0.046 (0.170)
Accident		0.290 (0.280)	0.255 (0.279)
Neighborhood Controls	NO	NO	YES
Observations	11234	11234	10944

Robust standard errors in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

Note: There are fewer observations in the last column because there is one neighborhood (Jamaica Plain) in which residents of that neighborhood were never searched.

**Table 16: Probability of Search, Diverse vs. Homogeneous Neighborhoods
Weighted Probits**

	Diverse Neighborhoods			Homogeneous Neighborhoods		
	(1)	(2)	(3)	(4)	(5)	(6)
Black Driver	-0.364* (0.206)	-0.399* (0.226)	-0.313 (0.220)	0.445*** (0.159)	0.450*** (0.161)	0.450*** (0.160)
Hispanic Driver	-0.147 (0.226)	-0.157 (0.234)	-0.172 (0.232)	-0.083 (0.232)	-0.055 (0.241)	-0.042 (0.239)
Black Officer	-0.382* (0.226)	-0.356 (0.219)	-0.327 (0.211)	0.101 (0.167)	0.130 (0.168)	0.142 (0.166)
Hispanic Officer	-0.360 (0.293)	-0.355 (0.274)	-0.313 (0.256)	-0.918*** (0.245)	-0.908*** (0.246)	-0.873*** (0.242)
Mismatch	0.386* (0.203)	0.360* (0.196)	0.355* (0.192)	0.233 (0.161)	0.261 (0.167)	0.249 (0.163)
Stop at Night		0.415** (0.168)	0.469*** (0.158)		0.123 (0.176)	0.098 (0.169)
Young Driver (Age<26)		-0.155 (0.169)	-0.152 (0.174)		0.094 (0.186)	0.112 (0.181)
Male Driver		0.123 (0.184)	0.139 (0.171)		-0.209 (0.162)	-0.223 (0.157)
In-State Driver		0.272 (0.290)	0.221 (0.289)		0.410* (0.241)	0.428* (0.251)
In-Town Driver		-0.017 (0.193)	-0.001 (0.191)		-0.101 (0.145)	-0.076 (0.147)
Accident		0.035 (0.223)	0.002 (0.253)		0.443 (0.300)	0.434 (0.298)
Neighborhood Controls	NO	NO	YES	NO	NO	YES
Observations	19247	19247	19247	37234	37234	37234

Robust standard errors in parentheses

* significant at 10%; ** significant at 5%; *** significant at 1%

The neighborhoods categorized as diverse are East Boston (including Charlestown), Roxbury/Mission Hill, Jamaica Plain and Hyde Park. The neighborhoods characterized as homogeneous are Boston Central, South Boston, Allston/Brighton and West Roxbury/Roslindale.