HARRIS SCHOOL WORKING PAPER
SERIES 05.7

TESTING FOR RACIAL PROFILING IN TRAFFIC STOPS FROM BEHIND A VEIL OF DARKNESS

Jeffrey Grogger and Greg Ridgeway

Testing for Racial Profiling in Traffic Stops from Behind a Veil of Darkness

Jeffrey Grogger
Harris School
University of Chicago
1155 E. 60th St.
Chicago, IL 60637
jgrogger@uchicago.edu

Greg Ridgeway
RAND
gregr@rand.org

Acknowledgements go here.

Abstract

The key problem in testing for racial profiling in traffic stops is estimating the risk set, or "benchmark," against which to compare the race distribution of stopped drivers. To date, the two most common approaches have been to employ Census-based residential population data or to conduct traffic surveys in which observers tally the race distribution of drivers at a certain location. It is widely recognized that residential population data may provide poor estimates of the population at risk of a traffic stop; at the same time, traffic surveys have limitations and may be too costly to carry out on the ongoing basis required by recent legislation and litigation. In this paper, we propose a test for racial profiling that does not require explicit, external estimates of the risk set. Rather, our approach makes use of what we refer to as the "veil of darkness" hypothesis, which asserts that at night, police cannot determine the race of a motorist until they actually make a stop. The implication is that the race distribution of drivers stopped at night should equal the race distribution of drivers at risk of being stopped at night. If we further assume that racial differences in traffic patterns, driving behavior, and exposure to law enforcement do not vary between day and night, we can test for racial profiling by comparing the race distribution of stops made during daylight to the race distribution of stops made at night. We propose a means of weakening this assumption by restricting the sample to stops made during the evening hours and controlling for clock time while estimating day/night contrasts in the race distribution of stopped drivers. We provide conditions under which our estimates are robust to a substantial non-reporting problem present in our data and in many other studies of racial profiling. We propose an approach to assess the sensitivity of our results to departures from our maintained assumptions. Finally, we apply our method to data from Oakland, California. In this example, the data yield little evidence of racial profiling in traffic stops.

I. INTRODUCTION

Racial profiling is a significant social problem. Forty-two percent of African-Americans say police have stopped them just because of their race. Fifty-nine percent of the U.S. public believes that the practice is widespread. Eight-one percent disapprove of the practice (Gallup 1999).

Public concern over racial profiling recently has resulted in massive amounts of costly data collection. At least 14 states have passed legislation to deal with racial profiling (Northeastern University, 2004). More than 400 police agencies now compile data on the race distribution of stopped motorists (McMahon et al., 2002). Many localities collect such data voluntarily. Some, such as the Cincinnati and Los Angeles Police Departments, have entered into consent decrees with the U.S. Justice Department that require such data collection on an ongoing basis.

Despite all the data collection, there remains considerable uncertainty as to how those data should be used to test for racial profiling. Many researchers suggest that a difference between the racial distribution of persons stopped by police and the racial distribution of the population at risk of being stopped would constitute evidence of its existence (San Jose Police Department 2002, Kadane and Terrin 1997, Smith and Alpert 2002, McDonald 2001, Dominitz 2003, General Accounting Office 2000, Zingraff et al. 2000). This implicit definition reveals the key empirical problem in testing for racial profiling: measuring the risk set, or the "benchmark," against which to compare the racial distribution of traffic stops.

Measuring the risk set explicitly poses a number of problems. First, the race distribution of drivers within a jurisdiction may differ from the race distribution of the residential population because car ownership and travel patterns may vary by race. They may differ also because part of the driving population originates outside of the jurisdiction, driving to or passing through it on the way to work, to shop, or the like.

Furthermore, the race distribution of the at-risk population may differ even from that of the driving population if drivers of different races exhibit different levels of care in their driving. Police may legally stop drivers who violate traffic laws, even those driving just one mile per hour over the speed limit. In the absence of a traffic violation, police must have probable cause or reasonable suspicion specific to the driver or a passenger in order to legally conduct a stop (Harris 1997). If there are racial differences in driving behavior, that is, in the care with which drivers drive, then the racial distribution of the at-risk population differs even from the racial distribution of the driving population.

Finally, the at-risk population may vary due to differences in exposure to police, even controlling for driving behavior. Police argue that they deploy patrols among neighborhoods in proportion to calls for service. Since calls for service come disproportionately from minority neighborhoods, minority neighborhoods have more patrols, so police may observe minority drivers more frequently (McMahon et al. 2002; San Jose Police Department 2002).

The benchmarking problem has generally been dealt with in one of three ways: analysts have used benchmarks based on residential populations or driver's license records, despite their limitations; researchers have conducted expensive traffic surveys, employing observers to tally the race distribution of drivers or traffic violators at a certain location; or they have ignored data on stops altogether, looking for racial disparities in other measures of police behavior. We discuss these approaches in more detail below (see Fridell 2004 for a full description of current methods).

Our main goal in this paper is to propose an alternative approach to testing for racial profiling in traffic stops that does not require explicit external estimates of the race distribution in the population at risk of being stopped. An important advantage of our approach is that it is inexpensive to implement, even on an ongoing basis, because the benchmark we propose can be constructed from traffic stop data themselves. We present the assumptions under which our approach yields a valid test, discuss how some of those assumptions may be relaxed, and provide some calculations to assess the sensitivity of the test to violations of those assumptions.

Our approach is based on a simple assumption: that during the night, police have greater difficulty observing the race of a suspect before they actually make a stop. This is a common assertion police make to us and to the press: "[m]any [officers] say … that they often cannot tell the race of people they are tailing, particularly at night." (Daunt and Leovy 2003). We refer to this as the "veil of darkness" hypothesis. The implication of the veil of darkness hypothesis is that the race distribution of drivers stopped at night should be equal to the race distribution of drivers at risk of being stopped at night. If travel patterns, driving behavior, and exposure to police are similar between night and day, then we can test for racial profiling by comparing the race distribution of drivers stopped during the day to the race distribution of drivers stopped during the night.

The assumption that travel patterns are similar day and night may be restrictive because the time of employment is known to vary by race (Hamermesh 1996). As a result, there may be racial differences in travel patterns by time of day. To deal with this issue, we make use of natural variation in hours of daylight over the year. In the winter, it is dark by early evening, whereas in the summer it stays light much later. Limiting much of our analysis to stops that take place during the hours between roughly 5 p.m. and 9 p.m., we can test for differences in the race distribution of traffic stops between night and day, while controlling implicitly for racial variation in travel patterns by time of day. As we argue below, limiting the sample period and employing time-of-day controls may also equalize differences in risk that arise due to differences in driving behavior or police exposure. Neighborhood controls may equalize any differences that remain.

In the next section of the paper, we provide more detail on previous analyses of racial profiling. In section 3, we discuss our data. In section 4, we formalize and extend our analytical approach. One important extension deals with a serious non-reporting problem that is common in the literature. We present the assumptions under which our approach yields valid tests. In section 5, we present our main results based on traffic-stop data from Oakland, California. We follow our main results with a sensitivity analysis that helps to quantify the extent by which our assumptions would have to fail in order to reverse our conclusions. We conclude with section 6, in which we discuss some of the limitations of our approach.

2. PREVIOUS RESEARCH ON RACIAL PROFILING

A number of studies have used Census-based estimates of the race distribution of residential populations to estimate the race distribution of the at-risk population. For the reasons discussed above, this approach has serious limitations, which have been recognized by both researchers and the courts (San Jose Police Department 2002; Dominitz 2003; Smith and Alpert 2002; Chavez v. Illinois State Police). Zingraff et al.

(2000) have refined this approach by using the race distribution of licensed drivers, rather than the residential population, to estimate the race distribution of drivers at risk of being stopped. Although this accounts for racial differences in the rate at which the population holds driver's licenses, it does not account for out-of-jurisdiction drivers, nor does it account for potential racial differences in travel patterns or driving behavior. Alpert, Smith, and Dunham (2003) have proposed using data on traffic accidents to estimate the race distribution of the at-risk population; although this approach potentially accounts for traffic patterns and driving behavior, it does not account for racial differences in police exposure. Some studies have experimented with "race blind" methods, comparing the race of drivers stopped by officers with the race of drivers flagged by photographic stoplight enforcement (Montgomery County Police Department, 2002) or aerial patrols (McConnell and Scheidegger, 2001). These race blind methods are closest to the line of reasoning behind the method presented in this paper; in the absence of racial profiling the ability to identify race in advance should not influence which drivers officers stop. In any case, the main appeal in using these types of data is their cost: they provide an inexpensive, if problematic, means of estimating the race distribution of the at-risk population.

At the other end of the cost spectrum are traffic surveys. Such surveys employ observers to tally the race distribution of drivers, and in some cases, the race distribution of drivers failing to take some specified level of care. For example, Lamberth (1994) employed observers to estimate the race distribution of all drivers, and of drivers exceeding the speed limit by at least 5 miles per hour, on a stretch of the New Jersey turnpike where motorists had lodged allegations of racial profiling against police.

The advantage of traffic surveys is that they provide plausibly valid estimates of the race distribution of drivers at a specific set of locations. However, traffic surveys have disadvantages as well. The first is their expense. By one estimate, such surveys require 800 person-hours of labor to carry out (Pritchard 2001). This may make them particularly unsuited for the ongoing monitoring of traffic stops required by much recent legislation and litigation. Another problem with the approach is that its validity may fail in multi-ethnic environments, where the ethnicity of a driver is difficult to discern with precision during an observation period that may last only a few seconds. Finally, the approach generally provides only limited measures of driver care. Although all drivers exceeding the speed limit by any amount are technically at legal risk of being stopped, the threshold of care observed in a traffic study may not reflect the threshold police typically employ when deciding whom to stop (General Accounting Office 2000). In Lamberth's (1994) analysis, which employed a low threshold for violations, virtually all drivers failed the care threshold (i.e., drove in excess of 5 miles/hour over the speed limit), independent of race. However, in a separate traffic study which employed a higher threshold over the highway, Lange, Blackman, and Johnson (2001) found that black drivers were more likely than non-blacks to exceed speeds of 80 miles an hour. If police are more likely to stop drivers, the faster they drive, then differences in driving behavior not accounted for in a low-threshold traffic study could potentially explain racial differences in traffic stops.

A final vein of research has ignored traffic stop data altogether, focusing on other measures of police behavior, such as the rate at which stopped drivers are searched, or the rate at which searches yield contraband, referred to as the "hit rate." A practical virtue of this approach is that the risk sets are readily measured: the population at risk of being searched consists of drivers who are stopped, and the population at risk of being found with contraband consists of drivers who are searched. Beyond mere practicality, the emphasis on hit rates stems from an economic model of police behavior. Knowles, Persico, and Todd (2001) show that in an environment where police seek to maximize arrests, equality of hit rates by race implies that police do not intentionally discriminate. However, the model implicitly assumes that police place no weight on the rate at which innocent motorists are detained. In contrast, much of American criminal law (starting with the Fourth Amendment) stresses the protection of the rights of the innocent. Since the rate at which innocents are wrongly detained is a function of the stop rate (Dominitz 2003), analyses that exclude stop rates omit this important consideration. In the next section, we discuss the stop data to which we apply the approach that we spell out below.

3. DATA

The genesis for the data we analyze were complaints by motorists and advocates that the Oakland Police Department (OPD) had engaged in racial profiling, discriminating in particular against black drivers (Oakland Police Department, 2004). An early analysis of stop data indicated that over half of drivers stopped by OPD were black, whereas blacks composed only 35 percent of the city's residential population. Although OPD started collecting stop data voluntarily, it later entered a settlement agreement requiring that such data be collected on an ongoing basis (Allen et al. v. City of Oakland et al. 2003, Section VI.B). Similar to the consent decrees involving other police departments, the Oakland litigation required regular monitoring of the stop data so as to detect trends in potentially discriminatory police behavior.

Under the terms of the agreement, police must record information on every stop that they initiate. Police officers must complete a report including items such as the reason for the stop, the time and location of the stop, and the race/ethnicity of person stopped. These stop-by-stop data are then entered into an electronic database, which the OPD made available for our analysis. Here we focus on motor vehicle stops, which are the type of stop that generate the most racial profiling complaints. Our sample data include stops carried out between June 15 and November 30, 2003.

Despite the terms of the decree, there is evidence of a substantial non-reporting problem in the data. An audit of the stop reports led OPD's Independent Monitoring Team to estimate that as many as 70 percent of all motor vehicle stops were not reported in the early phases of this study (Burges et al., 2004, pg 41). Close court-ordered oversight and increases in sanctions against noncompliant has resulted in increases in the number of completed stop forms, especially in October and November. Such sizeable non-reporting problems seem fairly common in the literature. Kadane and Terrin (1997) note that either race data were missing, or no report was available, for about 69 percent of the drivers who were stopped during the course of data collection for Lamberth's (1994) New Jersey Turnpike study. GAO (2000) report that the driver's race was missing from about half the stops carried out during a racial profiling study in Philadelphia. Smith and Alpert (2002) report that data were missing for 36 percent of the stops made in the course of a Richmond, Virginia study. Steward (2004) reports that 34 percent of Texas law enforcement agencies failed to collect stop data mandated by recent state legislation.

Clearly, non-reporting problems are an issue that must be considered in testing for racial profiling. In the next section, we provide conditions under which the veil of darkness approach yields valid tests despite the presence of substantial non-reporting. These conditions are weaker than one might expect. For example, we do not need to assume that the rate of non-reporting is independent of race. After we present our main analyses, we return to the non-reporting issue by assessing the extent to which the assumptions that we do require would have to be violated in order to overturn our qualitative conclusions.

4. METHODS

We begin with a simple test based on the veil of darkness hypothesis that is valid under relatively restrictive assumptions. We then analyze the extent to which we can relax those restrictions. We argue that conducting the test by use of a regression model that includes time-of-day covariates allows us to replace some strong assumptions with weaker assumptions. Other assumptions, such as those involving non-reporting, cannot be dealt with explicitly via the regression model. In those cases, we state the necessary conditions for our approach to yield valid tests, then provide a sensitivity analysis in the next section to assess the extent to which those conditions would have to fail in order to reverse our qualitative conclusions.

4.1 A Test for Racial Profiling Based on the Veil of Darkness Hypothesis

To formalize our approach in its simplest yet most restrictive form, let S be a binary random variable indicating whether officers stop a vehicle. Let the binary random variable B denote the event that a person is black and at risk of being stopped. To be at-risk, the person must be driving a vehicle, driving with a level of care that would normally lead police to stop the vehicle if observed, and be exposed to police. Let the binary random variable $\overline{B}$ denote the event that a person is non-black and at risk of being stopped. In the text we will often use the terms "black driver" and "non-black driver" as shorthand to refer to drivers in the at-risk population who are black and non-black, respectively.

In the absence of racial profiling, black and non-black drivers should be equally at risk of being stopped, that is, we should have

$$\frac{P(S \mid B)}{P(S \mid \overline{B})} = 1. \tag{1}$$

Neither quantity on the left-hand side of (1) can be estimated with stop data, but by Bayes rule the expression can be re-written and rearranged to yield

$$\frac{P(B \mid S)}{P(\overline{B} \mid S)} = \frac{P(B)}{P(\overline{B})}. \tag{2}$$

The left-hand side of (2) is the population *stop ratio*, that is, the ratio of the proportion of blacks among drivers who are stopped to the proportion of non-blacks among drivers who are stopped. This population stop ratio can be estimated from the traffic stop data. The right-hand side of (2) is the population *risk ratio*, that is, the ratio of the proportion of blacks among the at-risk population to the proportion of non-blacks among the at-risk population. It is the benchmark against which the race distribution of stop data should be compared in order to test for racial profiling. As discussed above, it is difficult to estimate the population risk ratio explicitly.

Rather than estimating the risk ratio explicitly, we propose an alternative test based on the veil of darkness hypothesis. Under the veil of darkness hypothesis, police are unable to observe the race of a driver when it is dark before they make the stop. Therefore when it is dark, the stop ratio should equal the risk ratio. Algebraically, let *d* be an indicator variable equal to one for a stop that occurs when it is dark and equal to zero for a stop that occurs when it is light. The veil of darkness hypothesis asserts that

$$\frac{P(B \mid S, d=1)}{P(\overline{B} \mid S, d=1)} = \frac{P(B \mid d=1)}{P(\overline{B} \mid d=1)},$$

since race cannot inform the decision to stop when it is not visible. Under the assumption that the race distribution of at-risk drivers is independent of daylight, this becomes

$$\frac{P(B \mid S, d=1)}{P(\overline{B} \mid S, d=1)} = \frac{P(B)}{P(\overline{B})}. \qquad (3)$$

In conjunction with equation (2), this suggests a test for racial profiling based on a test of whether

$$\frac{P(B \mid S, d=0)}{P(\overline{B} \mid S, d=0)} = \frac{P(B \mid S, d=1)}{P(\overline{B} \mid S, d=1)}. \qquad (4)$$

That is, we test whether the stop ratio during daylight, when $d = 0$ and police are more likely to see the driver before they make the stop, is equal to the stop ratio during the night, when $d = 1$ and police are less likely to be able to see the driver until after the stop is made.

4.2 Generalizing the Test

For a number of reasons, the assumptions needed to base a test on equation (4) may be restrictive. As discussed above, temporal travel patterns may vary by race due to differences in hours of work. If so, then the race distribution of the at-risk population may vary by time of day, violating equation (3). Another drawback is that the analysis above neglects potential differences in driving behavior and exposure to police. A further limitation is that it sheds no light on how the substantial non-reporting problem discussed in Section 3 may affect our results.

In this section we generalize the approach above. We develop a regression model that allows for a test based on weaker assumptions about temporal variation in the race distribution of drivers at risk. The model potentially controls for differential travel patterns, care, and exposure. Our analysis provides conditions under which our regression-based test is valid in the presence of non-reporting.

Much of our concern that the risk ratio may vary between dark and light stems from differences in travel patterns by time of day. To address the possible dependence of the risk ratio on daylight, we introduce clock time t into the analysis. As we argue shortly, conditioning on t may control implicitly for racial differences in travel patterns, driving behavior, and police exposure.

We generalize the simple test from the previous subsection by basing our test for racial profiling on a test of $K(t)$ in the relation

$$\frac{P(S \mid B, t, d=0)}{P(S \mid \overline{B}, t, d=0)} = K(t) \frac{P(S \mid B, t, d=1)}{P(S \mid \overline{B}, t, d=1)}. \tag{5}$$

In the absence of racial profiling, we should find that $K(t) = 1$ for all t. In the presence of racial profiling, we should find $K(t) > 1$. That is, we should find that blacks are at greater relative risk of being stopped during the daylight than during the dark, when officers are unable to engage in racial profiling by hypothesis.

To deal with non-reporting, we first apply Bayes rule to each of the four probability terms in equation (5), then solve for the logarithm of $K(t)$ to obtain

$$\begin{aligned} \log K(t) &= \log \frac{P(S \mid B, t, d=0)}{P(S \mid \overline{B}, t, d=0)} \frac{P(S \mid \overline{B}, t, d=1)}{P(S \mid B, t, d=1)} \\ &= \log \frac{P(B \mid S, t, d=0)}{P(\overline{B} \mid S, t, d=0)} \frac{P(\overline{B} \mid S, t, d=1)}{P(B \mid S, t, d=1)} \frac{P(\overline{B} \mid t, d=0)}{P(B \mid t, d=0)} \frac{P(B \mid t, d=1)}{P(\overline{B} \mid t, d=1)}. \end{aligned} \tag{6}$$

Letting *R* be a binary random variable indicating whether the officer reported the stop, we can introduce non-reporting in the expression for log *K*(*t*) with the probability relation

$$P(B \mid S,t,d) = \frac{P(B \mid R,S,t,d)P(R \mid S,t,d)}{P(R \mid B,S,t,d)}. \quad (7)$$

Substituting (7) into (6), collecting similar terms, and making use of the fact that $P(\overline{B} \mid R,S,t,d) = 1 - P(B \mid R,S,t,d)$, we obtain

$$\begin{aligned}\log K(t) = {} & \log\frac{P(B \mid R,S,t,d=0)}{1-P(B \mid R,S,t,d=0)} - \log\frac{P(B \mid R,S,t,d=1)}{1-P(B \mid R,S,t,d=1)} + \\ & \log\frac{P(\overline{B} \mid t,d=0)}{P(B \mid t,d=0)}\frac{P(B \mid t,d=1)}{P(\overline{B} \mid t,d=1)} + \\ & \log\frac{P(R \mid \overline{B},S,t,d=0)}{P(R \mid \overline{B},S,t,d=1)}\frac{P(R \mid B,S,t,d=1)}{P(R \mid B,S,t,d=0)}.\end{aligned} \quad (8)$$

Equation (8) is the key to the analysis that follows. The probabilities in the first line condition only on reported stops, exactly the data that we observe. We can estimate this line from the observed data using a logistic regression in which the dependent variable is a race indicator (black/non-black) with *d* (the darkness indicator) and *t* (clock time) as covariates. The logistic regression model estimates the regression *f*(*d*, *t*) from the observed data as

$$\log\frac{P(B \mid R,S,t,d)}{1-P(B \mid R,S,t,d)} = f(t,d). \quad (9)$$

The first line of equation (8) is then simply *f*(*t*, 0) – *f*(*t*, 1). If the effect of darkness is additive then this difference is simply the coefficient on the darkness variable times –1.

The second line of equation (8) measures how the mix of black and white drivers in the at-risk population changes depending on darkness and clock time. If the race distribution of the at-risk population is independent of darkness, conditional on clock time, this term vanishes. This is a weaker assumption than that reflected in equation (3) above, which required the risk ratio to be independent of darkness. Here we discuss the circumstances that may satisfy this weaker condition.

First note that in order to condition on clock time while estimating dark/light contrasts in the race distribution of stopped drivers, we must limit the sample to stops made at times when it is daylight during certain times of year and dark at other times. In Oakland, the latest occurrence of the end of civil twilight, which we use to define "dark," falls on June 22 at 9:06 pm. The earliest occurrence falls on November 30 at 5:19 pm. For the remainder of the analysis, we limit the sample to stops occurring between 5:19 and 9:06 pm, which we refer to as the inter-twilight period. Restricting the sample in this way allows us to construct contrasts by dark and light while controlling for clock time.

Figure 1 represents this idea visually. The horizontal axis indicates the clock time and the vertical axis indicates hours since dark. Throughout the analysis, we omit stops carried out during the roughly 30-minute period between sunset and civil twilight, since that period is difficult to classify as either light or dark. The solid points indicate stops of black drivers, whereas open circles represent stops of non-black drivers. At any time between 5:19 and 9:06 p.m. some stops are carried out when it is dark (gray shading) and some are carried out when it is light (no shading). The diagonal bands are a result of the natural variation in daylight hours over the course of the study period. In particular, the large diagonal gap is a result of the shift from Pacific Daylight Time to Pacific Standard Time at the end of October. This shift is especially useful for our comparison since it creates extremes in driver's race visibility for fixed clock times.

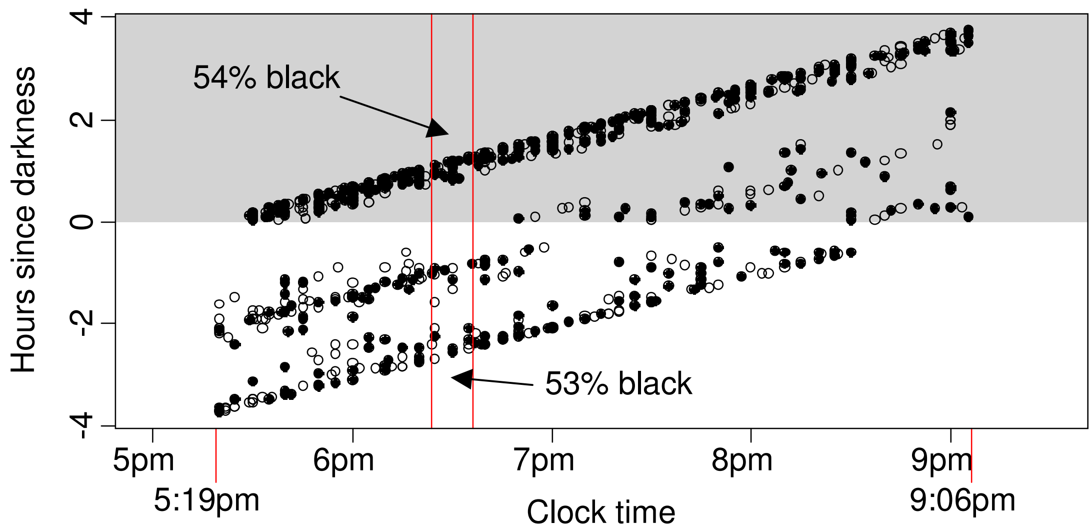

Figure 1: Plot of stops by clock time and darkness. The solid points indicate black drivers and the open circles represent non-black drivers. The shaded region indicates those stops occurring after the end of civil twilight. The large diagonal gap is a result of the shift from daylight savings time to standard time. The figure excludes stops occurring between sunset and the end of civil twilight. The vertical lines near 6:30pm mark the example region discussed in the text.

Thus within the inter-twilight period, we can construct contrasts by daylight and darkness in the fraction of stopped drivers who are black, controlling flexibly for time of day. For example, the vertical lines mark a period around 6:30pm within which we can assess whether darkness influences the race of drivers stopped. During daylight hours 53% of the stops involved black drivers while stops after dark involved black drivers 54% of the stops. The full regression analysis will combine such comparisons across the inter-twilight period. Note that, although we could potentially include stops carried out during the morning inter-twilight period as well as during the evening inter-twilight depicted, we exclude the morning stops simply because they are rare.

To the extent that differences in travel patterns result from differences in work hours and other routine activities, they may have more to do with clock time than with daylight per se. That is, if travel patterns vary between the races due to variation in commuting times, and commuting times are determined by work hours, it may be reasonable to assume that travel patterns are independent of daylight, conditional on time of day.

Time-of-day controls may also help equalize any racial differentials in driving behavior that may vary by time of day. Such differences may arise due to composition effects: drivers on the road at 8 p.m. may differ on average from those on the road at 6 p.m. because the former include a higher proportion of drivers en route to entertainment venues, whereas the latter include a higher proportion of those on their way home from work. If such composition effects vary by race, controlling for time-of-day should account for them.

To the extent that many of the motorists driving during our sample period are making routine trips home from work, controlling for time-of-day may also help control for police exposure. In driving home from work, many motorists travel at regular times along regular routes. This results in correlation between location and time-of-day, which time-of-day controls equalize by race. To provide further controls for differences in police exposure that arise due to differences in patrol intensity by location, we include neighborhood controls in one of the models we report below. More generally, the sensitivity test we carry out below will help to assess the extent to which our key assumption, that risk ratios are independent of daylight conditional on time-of-day, would have to be violated in order to reverse our conclusions.

The third line of equation (8) reveals the condition that reporting rates must satisfy in order for the regression to yield a valid test. The two ratios in this term measure how much reporting rates change between daylight versus darkness by race, given clock time. If reporting rates vary by race, but race-specific reporting rates do not vary between day and night (conditional on clock time), then these two terms vanish. It is important to note that equal reporting rates by race are not needed. As compared to the New Jersey traffic study, where equal reporting rates by race would have been necessary to identify the extent of racial profiling (Kadane and Terrin 1997), our requirement is weaker. Note, however, if there are a substantial number of officers that are not reporting stops and are engaging in racial profiling, then the reporting rate for black drivers during the day is likely to be smaller than the reporting rate for black drivers at night. After presenting our main results in the next section, we return to the non-reporting issue, asking to what extent racial reporting ratios would have to differ between day and night in order for the conclusions from our main analysis to be reversed.

5. RESULTS

5.1 Comparing Stops during Daylight and Dark

The simple approach described in Section 4.1 can be implemented with the full sample of data. In the full sample, we define daylight as extending from sunrise to sunset. We define dark as extending from the end of civil twilight in the evening until the beginning of civil twilight the following morning. We only used stops for moving violations, which comprise 75 percent of all stops. We did not consider mechanical and registration violations, which comprise most of the remaining stops (20 percent), since detection of these stops differs by darkness (e.g. headlight violations, noticing expired registration) and tend to be associated with non-white and non-Asian drivers. This may be due to racial profiling or differences in vehicle maintenance.

The first column of Table 1 presents the fraction of blacks among drivers stopped in the full sample. Among drivers stopped during daylight, 47 percent were black. Among drivers stopped when it was dark, 63 percent were black. Under the restrictive conditions discussed in Section 4.1, we can test for racial profiling by comparing these two numbers. They provide little evidence of racial profiling. If anything, they suggest that police are less likely to stop black drivers during the day, when the driver's race can be known in advance of the stop, than at night, when police cannot observe the driver's race until after the stop is made.

Table 1
Percent Black among Stopped Drivers, by Daylight

	Full sample	Inter-twilight sample
Total	0.52 (n=5144)	0.53 (n=862)
Light (d=0)	0.47 (n=3432)	0.50 (n=330)
Dark (d=1)	0.63 (n=1712)	0.55 (n=532)

The second column of Table 1 presents the fraction of blacks among drivers stopped in the inter-twilight sample. Among drivers stopped during daylight, 50 percent were black. Among drivers stopped when it was dark, 55 percent were black. Restricting the sample to the inter-twilight period has reduced the contrast between day and night, but the data still provide little evidence of racial profiling.

4.2 Regression Results

We first consider a simple model that assumes that the racial profiling effect is constant over time. It takes the form

$$\log \frac{P(B \mid d,t)}{1-P(B \mid d,t)} = \beta_0 + \beta_1 d + \gamma_1^T ns_6(t)\,. \tag{9}$$

where $ns_6(t)$ denotes a natural spline basis in clock time with 6 degrees of freedom, γ_1 is a column vector of 6 parameters, and the T superscript denotes transposition. For this model the racial profiling effect is a constant, $\log K(t) = -\beta_1$.

Table 2 presents the estimates of log K. The estimate in the first row uses no adjustment for clock time and essentially uses only the numbers presented in Table 1. The estimate in the second row adjusts for clock time. The estimate is negative, which constitutes evidence against racial profiling and is consistent with officers stopping black drivers less frequently during daylight than during the darkness. This is the same result observed in Table 1. Estimation of log K is imprecise, since the coefficient is smaller in absolute value than its standard error. Adding time of day controls does not change the conclusion from Table 1 that the data provide little evidence of racial profiling.

Table 2

Regression Estimates of the Racial Profiling Effect

Adjustments	log K	SE
None	-0.193	(0.1404)
Clock time	–0.092	(0.1510)
Clock time and neighborhood	–0.164	(0.1630)

Notes: Standard errors in parentheses. In addition to the indicator variable for darkness, the clock time adjusted models include a natural spline in clock time with 6 degrees of freedom. The third model also includes a set of patrol-area indictors.

We also estimated a model that allows for the extent of racial profiling to vary with clock time. It takes the form

$$\log \frac{P(B \mid t,d)}{1-P(B \mid t,d)} = \beta_0 + \beta_1 d + \gamma_1^T ns_6(t) + \gamma_2^T d * ns_6(t)\,. \tag{10}$$

For this model $\log K(t) = -\beta_1 - \gamma_2^T ns_6(t)$. Figure 2 plots the estimate by clock time. The shaded area indicates ±2 pointwise standard errors. Like the simpler model above, this model yields little evidence of racial profiling. Just before 7 p.m. log *K*(*t*) is at its highest, but it is still well within sampling variability of the horizontal line at 0.

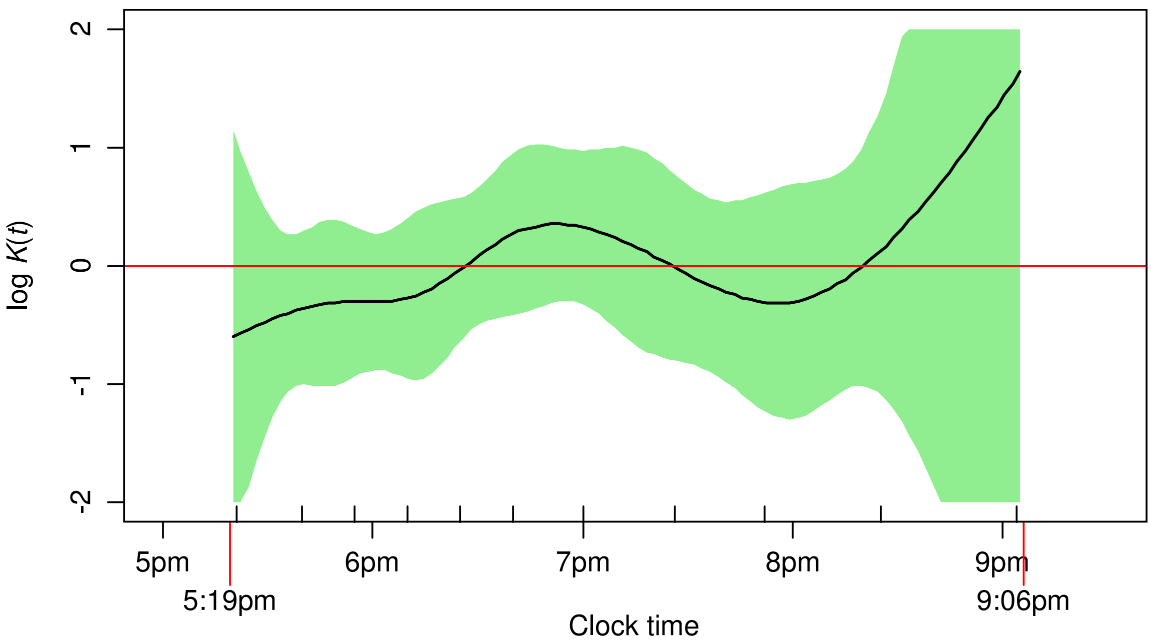

Figure 2: Estimate of log K(t). The curve is the best estimate of K(t) with the shaded area indicating ±2 pointwise standard errors. The horizontal line indicates the K(t) that we would expect under no racial profiling. The inward tickmarks along the x-axis indicate the deciles of the observed stop times.

Finally, we estimated a version of equation (9) to which we added an indicator variable for each patrol area in the city. These indicator variables provide additional controls for differential exposure to law enforcement between blacks and non-blacks that arises from differences in patrol intensity across neighborhoods. The Oakland police has divided the city into 35 community policing beats that we aggregated into six regions.

The second row of Table 2 reports the resulting estimate of log *K*(*t*). Controlling for neighborhood with an additive model still yields no evidence of racial profiling at the citywide level. We can refine this one step further with the data available to us by including a darkness*neighborhood interaction term. This will allow us to estimate a racial profiling effect for each neighborhood, as shown in Table 3. We continue to estimate an additive effect of clock time that does not vary by neighborhood.

Table 3

Regression Estimates of the Racial Profiling Effect by Neighborhood

Neighborhood	log *K*(neighborhood)	SE
Downtown	0.068	0.3010
East Oakland	-0.163	0.4433
Midtown Oakland	-0.081	0.4553
West Oakland	-0.678	0.4380
North Oakland	0.857	1.1693
North and South Hills	-0.972	0.9742

In all neighborhoods with the exception of West Oakland, the standard errors exceed the estimate of log *K*. West Oakland lies just south of the city's downtown core and is more than 80% non-white. However, log *K* is negative, again implying that, if anything, officers are less likely to stop black drivers when their race is visible.

4.3 Sensitivity analysis

Although the results above suggest that there is no racial profiling in traffic stops, those results hinge on assumptions concerning risk ratios and reporting rates. In the first column of Table 2, we estimated $-\beta_1$ to be –0.092. Under the assumptions maintained above, namely that differences in reporting ratios do not vary between day and night, and likewise that risk ratios do not vary between day and night in a manner that is independent of clock time, $\log K = -\beta_1$ since the last two lines of equation (8) equal 0. However, if our assumptions are violated, then the nuisance terms in the last two lines of equation (8) may be different from zero, in which case log *K* would differ from $-\beta_1$. If the sum of those nuisance terms differed from zero by enough that the lower end of the confidence interval exceeded zero, then we would question our conclusion regarding the absence of racial profiling. Although we cannot estimate the nuisance terms directly, in this section we illustrate the magnitude they would have to achieve to overturn our main conclusion.

The lower bound for a 95 percent confidence interval for $-\beta_1$ is –0.3880. This implies that if the sum of the nuisance terms exceeds 0.3880, then this would shift the estimate for log K enough for the data to suggest a racial profiling effect. We focus first on the risk ratio term (the second line in equation (8)), assuming for the moment that the reporting ratio term (the third line in equation (8)) equals zero.

We consider the circumstances under which

$$\frac{\dfrac{P(\overline{B} \mid t, d=0)}{P(B \mid t, d=0)}}{\dfrac{P(\overline{B} \mid t, d=1)}{P(B \mid t, d=1)}} = \exp(0.3880) = 1.40\,. \tag{11}$$

To assess this magnitude, assume that at 6:30 p.m. on days when 6:30 p.m. occurs during daylight black and non-black drivers are equally at risk for being stopped, that is, $P(B \mid t, d = 0) = 50$ percent. In this case an odds ratio of 1.40 implies that at 6:30 p.m. on dark days, black drivers compose 58 percent of the at-risk population. The proportion of black drivers would have to increase 16 percent between days that it was light at 6:30 and days that it was dark at 6:30.

Focusing next on the reporting term, and assuming the risk ratio term is 0, if the odds ratio for darkness exceeds 1.40 then we likewise have evidence for racial profiling:

$$\frac{\dfrac{P(R \mid B, S, t, d=1)}{P(R \mid B, S, t, d=0)}}{\dfrac{P(R \mid \overline{B}, S, t, d=1)}{P(R \mid \overline{B}, S, t, d=0)}} = \exp(0.3880) = 1.40\,. \qquad (12)$$

Assume that reporting rates for non-black drivers vary by *t* but not by *d* so that the denominator of (12) is 1. For the reporting term to exceed 0.3880, stops involving black drivers would have to be 40 percent more likely to be reported at night than during the day (e.g. 30 percent during daylight and 42 percent in darkness), requiring a substantial fraction of the non-reporting police force to be engaging in racial profiling. We can rearrange the left side of (12) to consider another black/non-black comparison. If stops involving black drivers were twice as likely to be reported during the day as stops involving non-black drivers, then officers would have to report black drivers more than 2.8 times as often as non-black drivers at night before it invalidates the "no racial profiling" conclusion.

The sensitivity analysis has considered deviating from the assumptions about the exposure term being 0 and the reporting term being 0, but has not considered both violations simultaneously. If the risk ratio in (11) is 1.18 and simultaneously the reporting ratio in (12) is 1.18, then we begin to have evidence of racial profiling.

6. CONCLUSIONS

The key problem in testing for racial profiling in traffic stops is estimating the risk set against which to compare the race distribution of stopped drivers. Previous analyses have relied on external estimates of the risk set constructed from either secondary data or traffic surveys. The validity of estimates from secondary data has been questioned; traffic surveys are expensive to carry out. The approach we have proposed here does not require external estimates of the risk set. It does require certain identifying assumptions. In the case of the Oakland data, our sensitivity analysis suggests that the departures from those assumptions would have to be substantial in order to overturn our conclusions.

A few points concerning limitations are in order. We have noted that our estimates are valid if, controlling for clock time, racial differences in risk sets do not vary between day and night. Implicitly, we have assumed that there is no seasonality in day-night risk differentials. In areas with substantial tourist inflows, this assumption may be violated.

A further caveat is that the results are limited to the inter-twilight period. Our approach cannot speak directly to the question of racial profiling during other hours. Other types of studies have similar limitations. Like our approach, studies based on traffic surveys, which require visual identification of the driver's race, cannot assess the extent of racial profiling at night.

Finally, our approach is designed only to assess the extent of racial profiling in traffic stops. Other studies have noted racial disparities in post-stop outcomes, such as search and arrest rates. Data on a full set of post-stop outcomes are needed to provide a comprehensive assessment of racial profiling.

References

Alpert, Geoffrey P., Smith, Michael R., and Dunham, R.G. (2003), "Toward a Better Benchmark: Assessing the Utility of Not-at-Fault Traffic Crash Data in Racial Profiling Research," Confronting Racial Profiling in the 21st Century: Implications for Racial Justice, Boston.

Burges, Rachel, Kelli Evans, Charles Gruber, and Christy Lopez (2004). Second Quarterly Report of the Independent Monitor. Available at http://www.oaklandpolice.com/agree/2qtr.pdf.

Daunt, Tina and Jill Leovy (January 7, 2003). "LAPD Offers 1st Data on Traffic Stops," Los Angeles Times.

Dominitz, Jeff. (2003), "How Do the Laws of Probability Constrain Legislative and Judicial Efforts to Stop Racial Profiling?" American Law and Economics Review, 5, 412-432.

Fridell, Lorie A. (2004), By the Numbers: A Guide for Analyzing Race Data from Vehicle Stops, Police Executive Research Forum, Washington, DC.

Gallup, George Sr. (1999), The Gallup Poll: Public Opinion 1999, Wilmington, DE: Scholarly Resources, Inc.

General Accounting Office. (2000), Racial Profiling: Limited Data Available on Motorist Stops, Washington, DC: General Accounting Office.

Hamermesh, Daniel S., (1996), Workdays, Workhours, and Work Schedules : Evidence for the United States and Germany, Kalamazoo, Mich.: W.E. Upjohn Institute for Employment Research.

Harris, David A. (1997), "'Driving While Black' and All Other Traffic Offenses: The Supreme Court and Pretextual Traffic Stops," Journal of Criminal Law and Criminology, 87, 544-581.

Harris, David A. (1999), "The Stories, the Statistics, and the Law: Whey 'Driving While Black' Matters," Minnesota Law Review, 84, 265-326.

Kadane, Joseph B., and Terrin, Norma. (1997), "Missing Data in the Forensic Context," Journal of the Royal Statistical Society. Series A, 160, 351-357.

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

Lamberth, John. (1994), "Revised Statistical Analysis of the Incidence of Police Stops and Arrests of Black Drivers/Travelers on the New Jersey Turnpike Between

Exits or Interchanges 1 and 3 from the Years 1988 through 1991," report, Temple University, Department of Psychology.

Lange, James E., Blackman, Kenneth O., and Johnson, Mark B. (2001), "Speed Violation Survey of the New Jersey Turnpike: Final Report," report to New Jersey Attorney General's Office, Public Services Research Institute.

MacDonald, Heather. (2001), "The Myth of Racial Profiling," City Journal, http://www.city-journal.org/html/11_2_the_myth_.html.

McConnell, Elizabeth H. and Amie R. Scheidegger. (2001), "Race and Speeding Citations: Comparing Speeding Citations Issued by Air Traffic Officers with Those Issued by Ground Traffic Officers." Paper presented at the annual meeting of the Academy of Criminal Justice Sciences, Washington, D.C., April 4-8.

McMahon, Joyce, Garner, Joel, Davis, Ronald, and Kraus, Amanda. (2002), How to Correctly Collect and Analyze Racial Profiling Data: Your Reputation Depends on It!, Washington, DC: Government Printing Office.

Montgomery County Department of Police. (2002), Traffic Stop Data Collection Analysis, 3rd report.

Northeastern University, (2004), Data Collection Resource Center, http://www.racialprofilinganalysis.neu.edu/legislation.php

Oakland Police Department (2004), Promoting Cooperative Strategies to Reduce Racial Profiling: A Technical Guide, Oakland, CA: Oakland Police Department.

Pritchard, Justin. (2001), "Racial Profiling a Conundrum for Police," Los Angeles Times, January 21, A22.

Ramirez, Deborah, McDevitt, Jack, and Farrell, Amy. (2000), A Resource Guide on Racial Profiling Data Collection Systems: Promising Practices and Lessons Learned, Washington, DC: U.S. Department of Justice.

San Jose, California Police Department. (2002), Vehicle Stop Demographic Study, San Jose, CA: San Jose Police Department.

Smith, Michael R., and Alpert, Geoffrey P. (2002), "Searching for Direction: Courts, Social Science, and the Adjudication of Racial Profiling Claims," Justice Quarterly, 19, 673-703.

Steward, Dwight. (2004), Racial Profiling: Texas Traffic Stops and Searches, Austin, TX: Steward Research Group.

Zingraff, Matthew T., Mason, Marcinda, Smith, William, Tomaskovic-Devey, Donald, Warrent, Patricia, McMurray, Harvey L., and Fenlon, Robert C. (2000), Evaluating North Carolina State Highway Patrol Data: Citation, Warnings, and Searches in 1998, report submitted to North Carolina Department of Crime Control and Public Safety and North Carolina State Highway Patrol.