A Collection of Methods for Racial Profiling Analysis

Greg Ridgeway
RAND Safety & Justice Program
Santa Monica, CA

January 22, 2007

Racial profiling is a growing concern

- I-95 "turnpike" studies in the mid-1990s raised public concern about racial profiling
- Public concern has led to state and local-level action
  - At least 26 states have passed legislation
  - Hundreds of other localities collect data; some compelled by the Justice Department
- Congress considering the End of Racial Profiling Act mandating data collection to receive Federal funds
- Should officers use racial profiling?
  - Tenth Circuit: "unequal application of criminal law to white and black persons was one of the central evils addressed by the framers of the Fourteenth Amendment"

Analytic quality is weak

- A growing number of studies claim racial profiling based on analysis of data collected
  - Texas: Concluded that "75% of agencies stop more black and Latino drivers than white drivers"
- And some studies hastily conclude no profiling occurs based on analyzed data
  - Sacramento:
    % black drivers stopped =
    % black crime suspect descriptions

Why is testing for racial profiling so hard?

Racial Distribution of People Stopped

Racial Distribution of People at Risk of Being Stopped

Why is testing for racial profiling so hard?

Racial Distribution of People Stopped

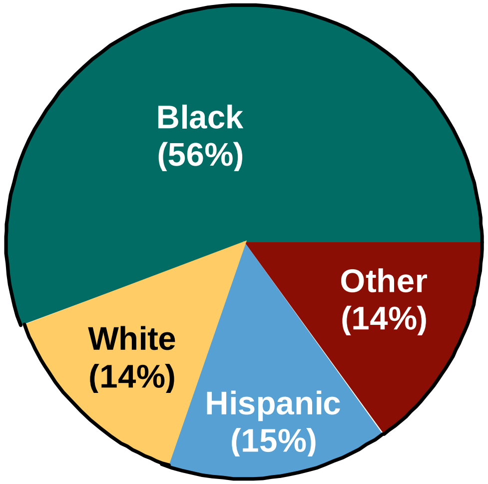

Racial Distribution of People at Risk of Being Stopped

Why is testing for racial profiling so hard?

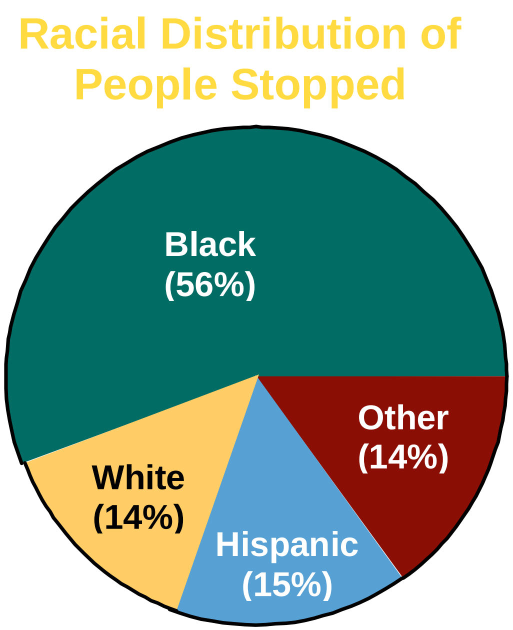

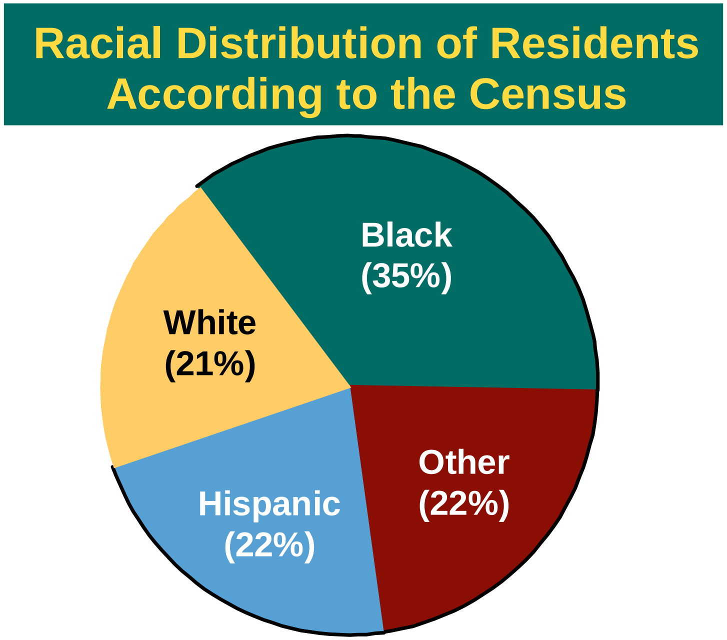

- The difference may result from:
  - A race bias
  - Car ownership, time on the road, and care
  - Exposure to police

A new approach

- Gauge department wide racial bias in the decision to stop
- Identify potential problem officers with internal benchmarking
- Assess racial bias in post-stop activity with propensity scores

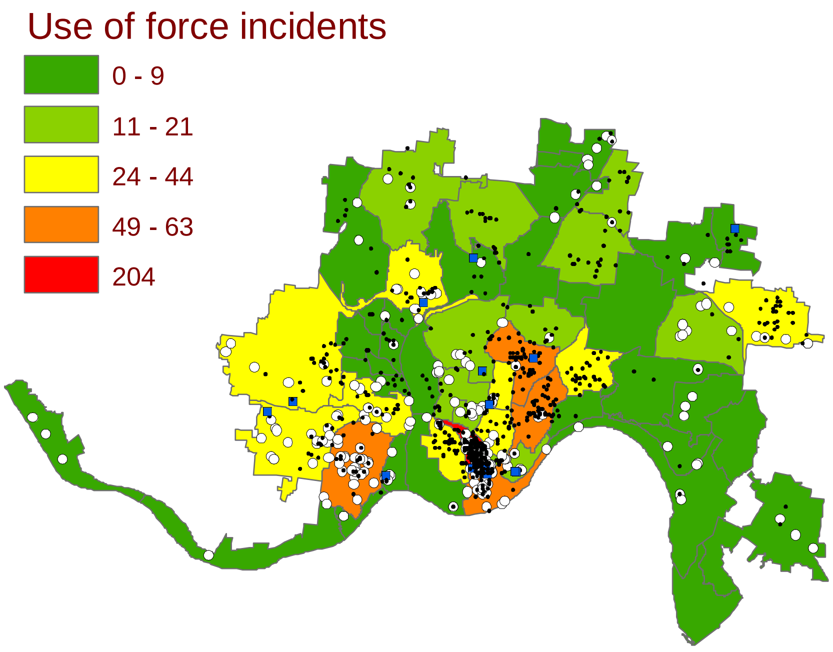

Step #1: Bias in the decision to stop

Grogger & Ridgeway (2006). "Testing for Racial Profiling in Traffic Stops from Behind a Veil of Darkness," *JASA* 101(475):878-887.

Central question: Does an officer's ability to identify race of driver in advance influence which drivers he stops?

- The ability to discriminate requires officers identifying the race in advance (e.g. Goldin & Rouse, bias in orchestra auditions)
- The ability to identify race in advance of the stop decreases as it becomes dark
- We directly test whether the ability to identify the race affects the race distribution of the stopped drivers

Simple veil of darkness test

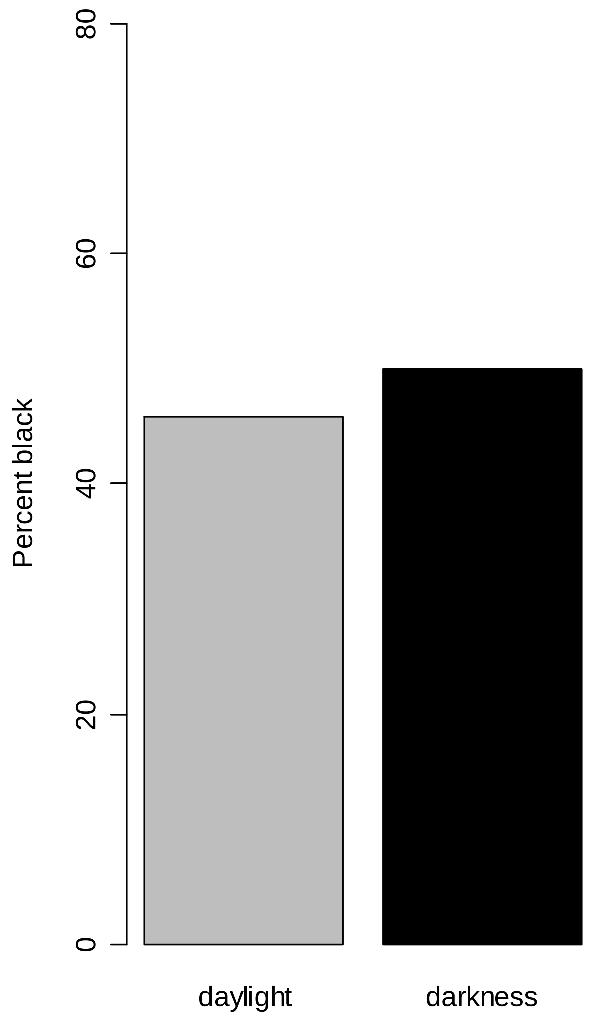

- CPD officers stop a greater proportion of black drivers at night than during the day
- This is counter to the racial profiling hypothesis

Adjusting for “clock time”

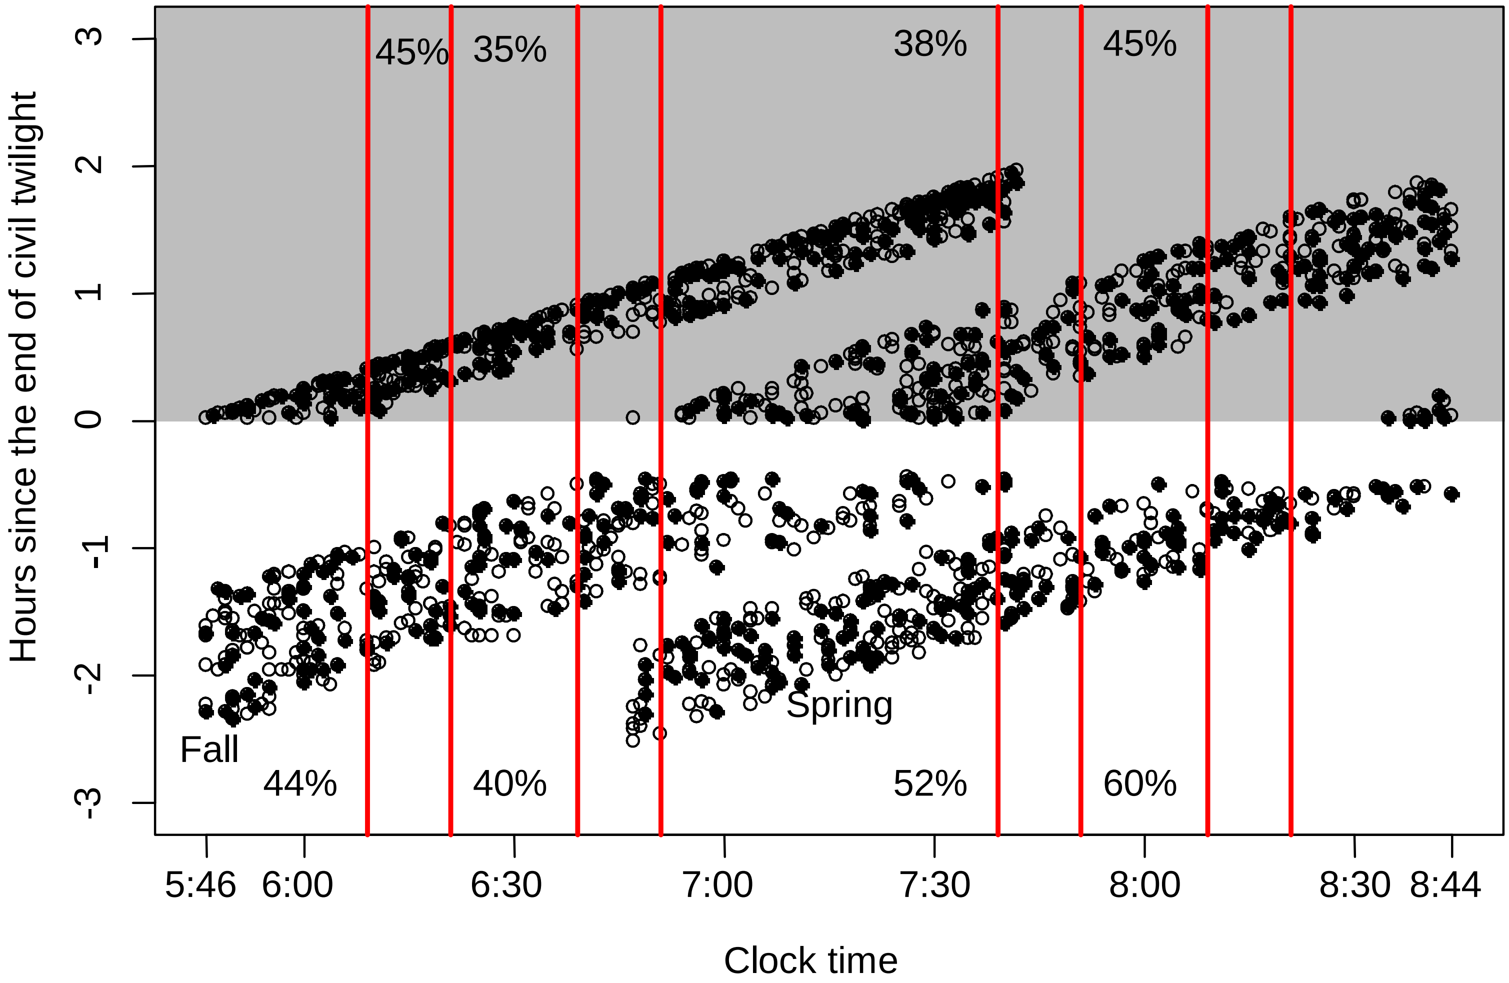

Development of the test

- In the absence of a race bias $K(t) = 1$

$$\frac{P(S|B,t,d=0)}{P(S|\bar{B},t,d=0)} = K(t)\frac{P(S|B,t,d=1)}{P(S|\bar{B},t,d=1)}$$

- Bayes' Theorem and some algebra yield

$$K(t) = \frac{P(B|S,t,d=0)}{P(\bar{B}|S,t,d=0)}\frac{P(\bar{B}|S,t,d=1)}{P(B|S,t,d=1)}\frac{P(\bar{B}|t,d=0)}{P(B|t,d=0)}\frac{P(B|t,d=1)}{P(\bar{B}|t,d=1)}$$

Accommodate underreporting

- There is some potential underreporting

$$P(B|S,t,d) = \frac{P(B|R,S,t,d)P(R|S,t,d)}{P(R|B,S,t,d)}$$

$$\begin{aligned}\log K(t) = \\ &\log \frac{P(B|R,S,t,d=0)}{1-P(B|R,S,t,d=0)} - \log \frac{P(B|R,S,t,d=1)}{1-P(B|R,S,t,d=1)} + \\ &\log \frac{P(\bar{B}|t,d=0)}{P(B|t,d=0)} \frac{P(B|t,d=1)}{P(\bar{B}|t,d=1)} + \\ &\log \frac{P(R|\bar{B},S,t,d=0)}{P(R|\bar{B},S,t,d=1)} \frac{P(R|B,S,t,d=1)}{P(R|B,S,t,d=0)}\end{aligned}$$

Decomposition of the race effect

$$\log K(t) = \text{stop distribution} + \text{exposure} + \text{reporting}$$

- We can estimate the stop ratio using logistic regression

$$\log \frac{P(B|R,S,d,t)}{1-P(B|R,S,d,t)} = \beta_0 + \beta_1 d + g(t)$$

- $g(t)$ is some flexible function of t (e.g. $t + t^2 + t^3$)
- Assume exposure term is 0
- Assume reporting term is 0
- $\log K(t) = -\beta_1$

Results: VoD estimates of bias, all months

Year	$K(t)$	95% interval	N
2003	1.01	(0.88,1.16)	4,013
2004	0.98	(0.86,1.12)	4,589
2005	1.07	(0.98,1.16)	10,890
Combined	1.02	(0.95,1.09)	19,492

- Includes all stops during the evening intertwilight period

Results: VoD estimates of bias, Daylight Savings Time

Year	$K(t)$	95% interval	N
2003	1.15	(0.79,1.68)	470
2004	1.19	(0.79,1.80)	403
2005	1.11	(0.81,1.52)	764
Combined	1.10	(0.91,1.33)	1,637

- Includes all stops occurring within four weeks of the spring or fall Daylight Saving Time change during the evening intertwilight period

Step #2: Internal benchmarking

- Consider a particular officer #534
- 71% of this officer’s stops involve a black driver

		Percentage
Time	(12-4pm]	9
	(4-8pm]	57
	(8pm-12am]	34
Day	Mon	20
	Tue	12
	Wed	12
	⋮	⋮
Month	Jan	12
	Feb	14
	Mar	7
	Apr	6
	May	8
	⋮	⋮
Area	J	49
	K	33
	L	5
	M	11

Internal benchmark

- 46% of similarly situated stops made by other officers involved black drivers

		Percentage	Comparison
Time	(12-4pm]	9	9
	(4-8pm]	57	56
	(8pm-12am]	34	35
Day	Mon	20	20
	Tue	12	11
	Wed	12	12
	⋮	⋮	⋮
Month	Jan	12	12
	Feb	14	15
	Mar	7	7
	Apr	6	6
	May	8	7
	⋮	⋮	⋮
Area	J	49	48
	K	33	34
	L	5	5
	M	11	11

Propensity score weighting

- Reweight stops that other officers made so that they have the same distribution of features

$$f(\mathbf{x}|t=1) = w(\mathbf{x})f(\mathbf{x}|t=0)$$

- Solving for $w(\mathbf{x})$ yields the propensity score weight

$$w(\mathbf{x}) = \frac{f(t=1|\mathbf{x})}{f(t=0|\mathbf{x})}K = \frac{p(\mathbf{x})}{1-p(\mathbf{x})}K$$

  where $p(\mathbf{x})$ is the probability that a stop with features $\mathbf{x}$ involves the officer in question

- Estimate $p(\mathbf{x})$ using a flexible, non-parametric version of logistic regression
- Compare the percentage of black drivers among the officer's stops with the weighted percentage of black drivers among other stops using weights $w_i = p(\mathbf{x}_i)/(1-p(\mathbf{x}_i))$

Common approach

- A common approach is to compute z-statistics for each officer

$$z = \frac{p_t - p_c}{\sqrt{\frac{p_t(1-p_t)}{n_t} + \frac{p_c(1-p_c)}{ESS}}}$$

- In the absence of racial bias this would be distributed N(0,1) and a cutoff of 2.0 would be reasonable
- With 133 officers and 133 correlated zs an appropriate reference distribution can be much wider (Efron 2006).

Estimating the false discovery rate

- Estimate $f_0(z)$ and $f(z)$ from the observed zs
- Right tail consists of 5 officers with “problem officer” probabilities ranging from 70% to 86%

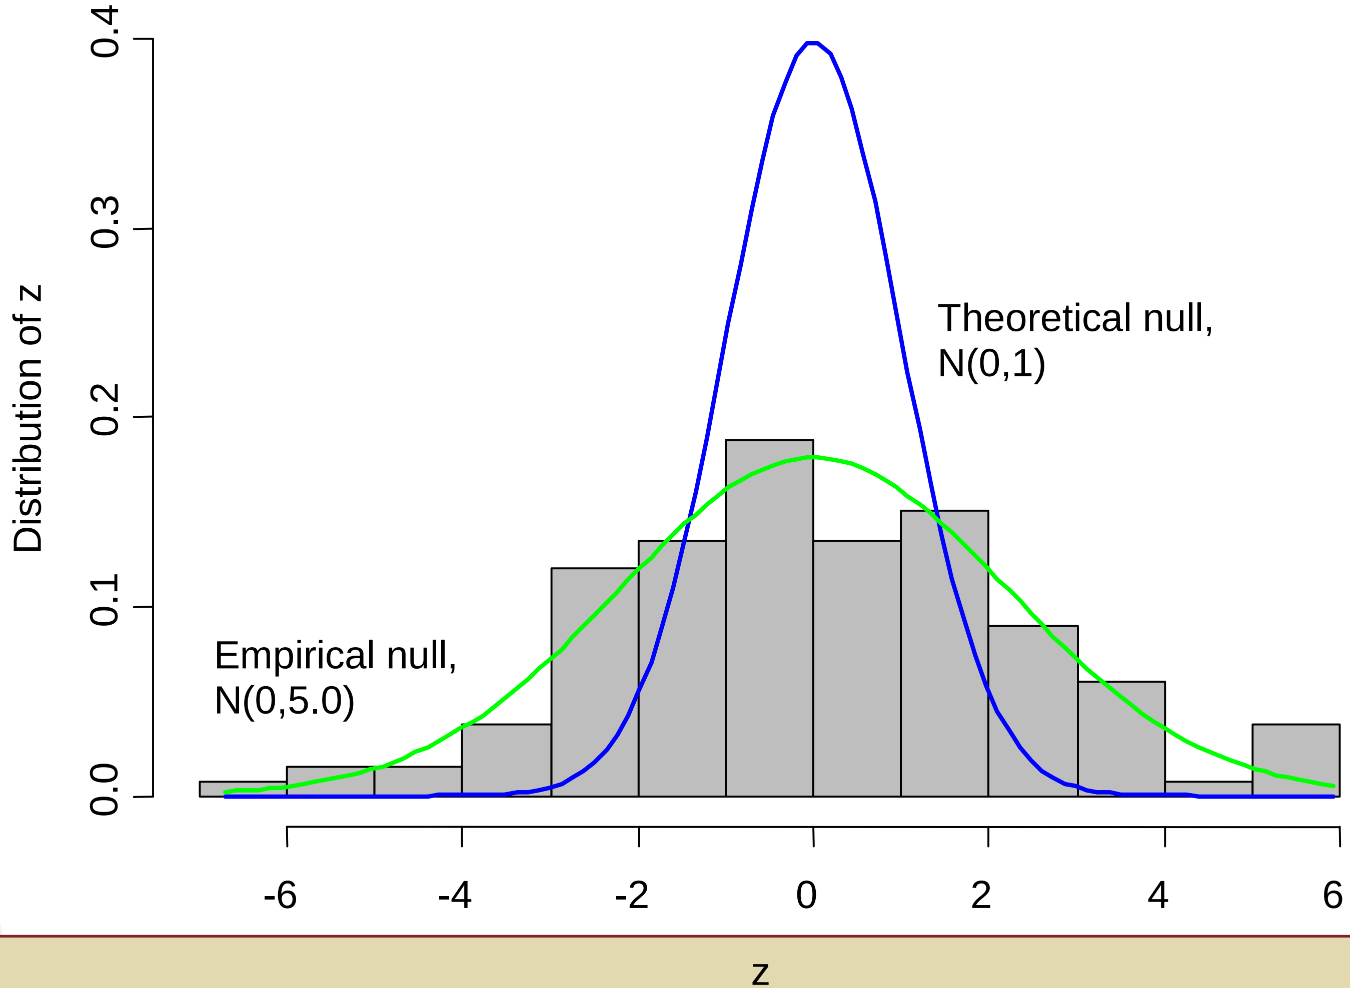

Step #3: Assessing race bias post-stop

G. Ridgeway (2006). "Assessing the effect of race bias in post-traffic stop outcomes using propensity scores," *JQC* 22(1):1-29.

- Central question: Are black drivers more/less likely to be cited, have long stop durations, or be searched?

Stop feature	% Black drivers (N=3,703)	% Nonblack drivers (N=3,033)
Region		
A	32%	14%
Time of day		
12am-4am	16%	8%
Resident	76%	64%
Age		
18-29	47%	38%
Reason		
Mechanical/ Registration	26%	16%
Male	75%	74%

Reweighting balances the group

- $w(\mathbf{x}) = \frac{P(\text{black}|\mathbf{x})}{1-P(\text{black}|\mathbf{x})}$

Stop feature	% Black drivers (N=3,703)	% Nonblack drivers weighted (ESS=1,689.2)	% Nonblack drivers (N=3,033)
Region			
A	32%	33%	14%
Time of day			
12am-4am	16%	16%	8%
Resident	76%	76%	64%
Age			
18-29	47%	48%	38%
Reason			
Mechanical/ Registration	26%	26%	16%
Male	75%	76%	74%

Results: Cincinnati stop duration

Year	Stop Duration (Minutes)	Black Drivers	Nonblack (reweighted)	Nonblack (unweighted)
2003	n =	16,708	4,881	18,548
	(0,10)	40%	43%	56%
2004	n =	18,721	5,190	20,390
	(0,10)	40%	44%	59%
2005	n =	15,571	4,965	20,431
	(0,10)	45%	47%	60%

- Black drivers in 2005 were three times more likely to have invalid licenses than white drivers (23% vs. 7%)

Results: Cincinnati search rates

Year	Discretion (Minutes)	Black Drivers	Nonblack (reweighted)	Nonblack (unweighted)
2003	n =	16,708	4,881	18,548
	High	5.9%	5.4%	2.8%
	Low	8.1%	5.5%	2.7%
2004	n =	18,721	5,190	20,390
	High	6.7%	6.2%	3.2%
	Low	10.7%	7.0%	3.9%
2005	n =	19,375	6,141	25,163
	High	6.1%	5.2%	2.8%
	Low	4.4%	3.5%	1.6%

- Hit rates for black and white drivers are about 28% for high discretion searches.

Summary

- Racial profiling analyses have generally confused the issue by studying irrelevant comparisons
- Credible and relevant comparisons are not difficult
  - Assess whether the ability to identify race in advance influences who gets stopped
  - Compare similarly situated officers
  - Equalize race groups on the obvious features on which they might legitimately differ

For more information

- Oakland 2003 report endorsed by OPD, the ACLU, the NAACP, and the Oakland CPRB
- Oakland Tribune reported “blacks are more likely than other races to be pulled over by police”
- Cincinnati Enquirer “Study: No bias in traffic stops, But many perceive discrimination based on race”

More available at http://www.i-pensieri.com/gregr/rp.shtml or Google “racial profiling analysis” or “Greg Ridgeway”