RAND Center on Quality Policing

A RAND INFRASTRUCTURE, SAFETY, AND ENVIRONMENT PROGRAM

Methods for Assessing Racially Biased Policing

Greg Ridgeway

Director, RAND Safety & Justice

Racial Profiling Continues to Be a 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

- Arrest of Henry Louis Gates in July 2009 and the resulting “beer summit” renewed interest

Unfortunately, the Quality of the Analysis Using Collected Data Is Weak

- A large 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: the percentage of black drivers stopped matched the percentage of blacks among crime suspect descriptions

Outline

- Assessing disparities in the decision to stop
- Internal benchmarking and early warning systems
- Assessing disparities in post-stop outcomes

Outline

- **Assessing disparities in the decision to stop**
- **Internal benchmarking and early warning systems**
- **Assessing disparities in post-stop outcomes**

Why Is Testing for Racial Profiling So Hard?

Difference Between

Racial Distribution of People Stopped

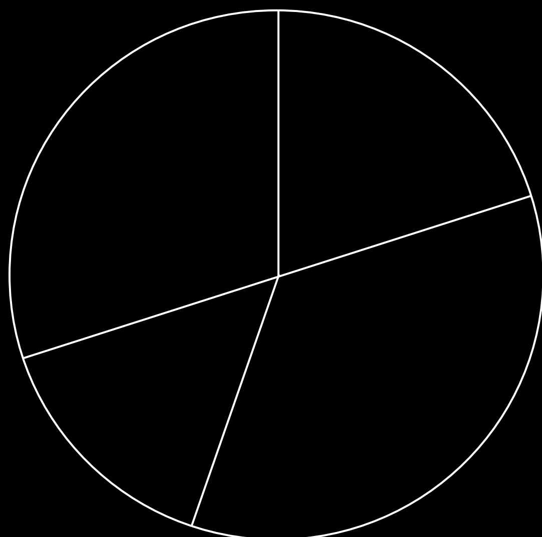

And

Racial Distribution of People at Risk of Being Stopped

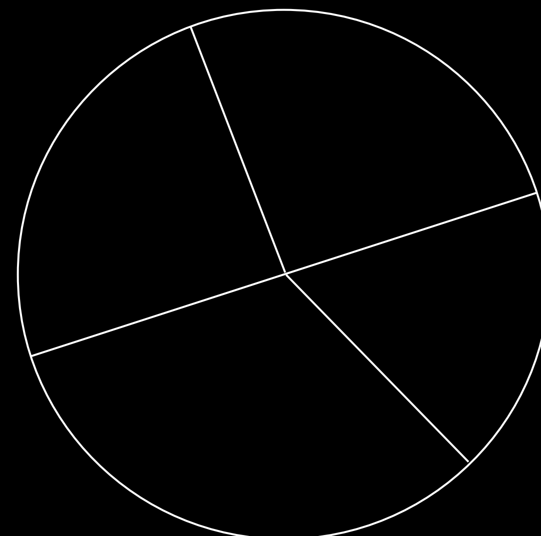

= Racial Profiling

Why Is Testing for Racial Profiling So Hard?

Racial Distribution of People Stopped

Racial Distribution of People at Risk of Being Stopped

Difference Between

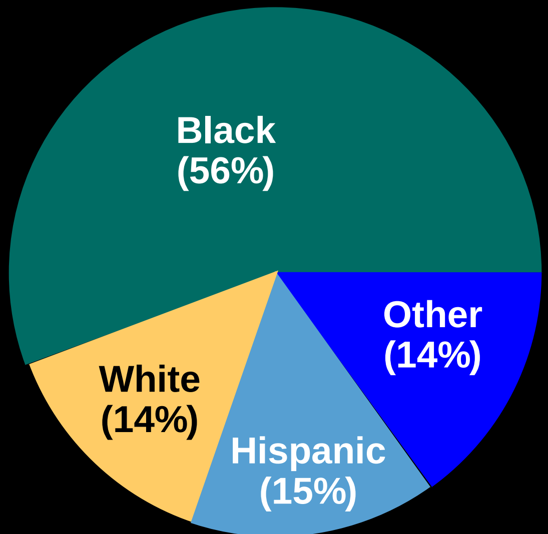

And

?

= **Racial Profiling**

Source: Oakland Police Department, 2003

Why Is Testing for Racial Profiling So Hard?

Racial Distribution of People Stopped

Racial Distribution of Residents According to the Census

Difference Between

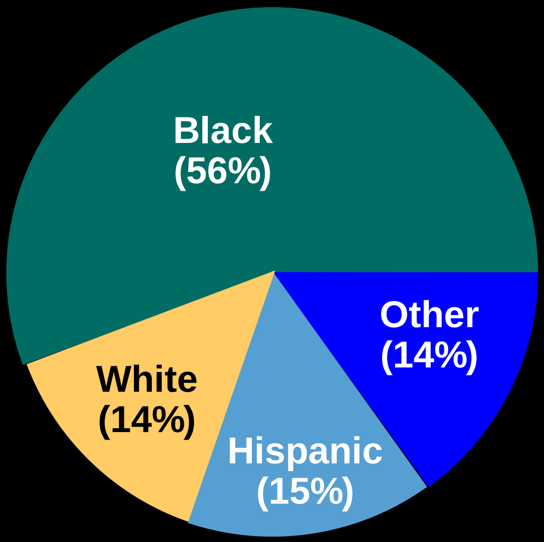

And

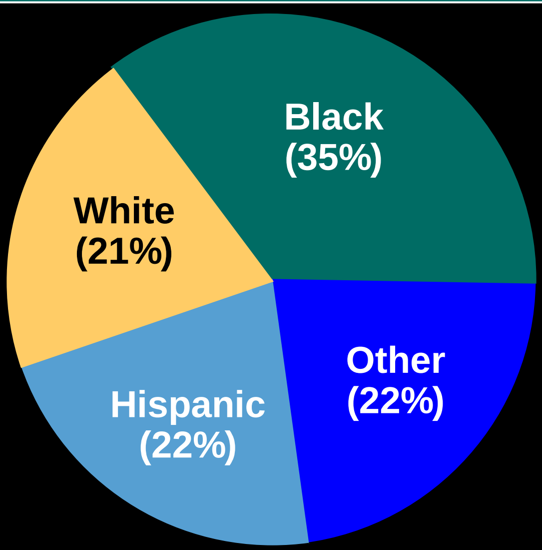

= 1.6

Source: Oakland Police Department, 2003

Source: U.S. Census, 2000

Why Is Testing for Racial Profiling So Hard?

Racial Distribution of People Stopped

Racial Distribution of Residents According to the Census

Difference Between

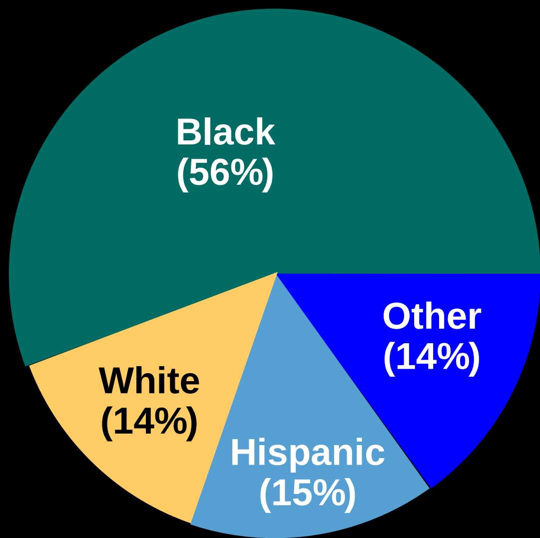

Source: Oakland Police Department, 2003

And

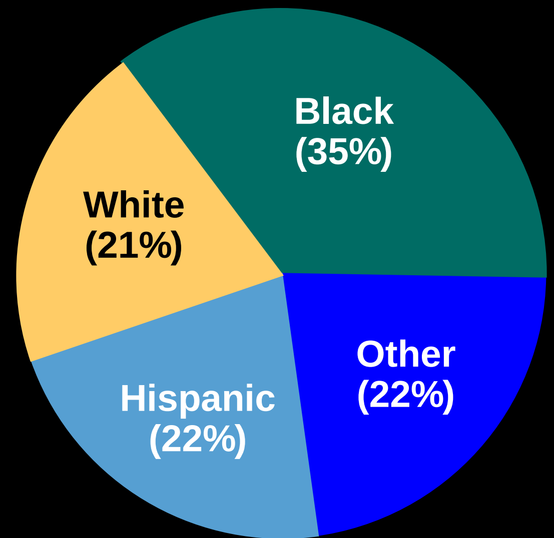

Source: U.S. Census, 2000

= 1.6

- **The 1.6 disparity between the racial distributions may result from:**
  - **A race bias**
  - **Driving behavior: car ownership, time on the road, and care**
  - **Exposure to police by area of city, neighborhood characteristics, etc.**

Does the Ability to See the Driver Influence Which Drivers Are Stopped?

1. **The ability to discriminate requires officers to identify the race in advance**

2. **The ability to identify race in advance of the stop decreases as it becomes dark**

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

Simple "Veil of Darkness" Test Shows No Evidence of Racial Bias in the Decision to Stop

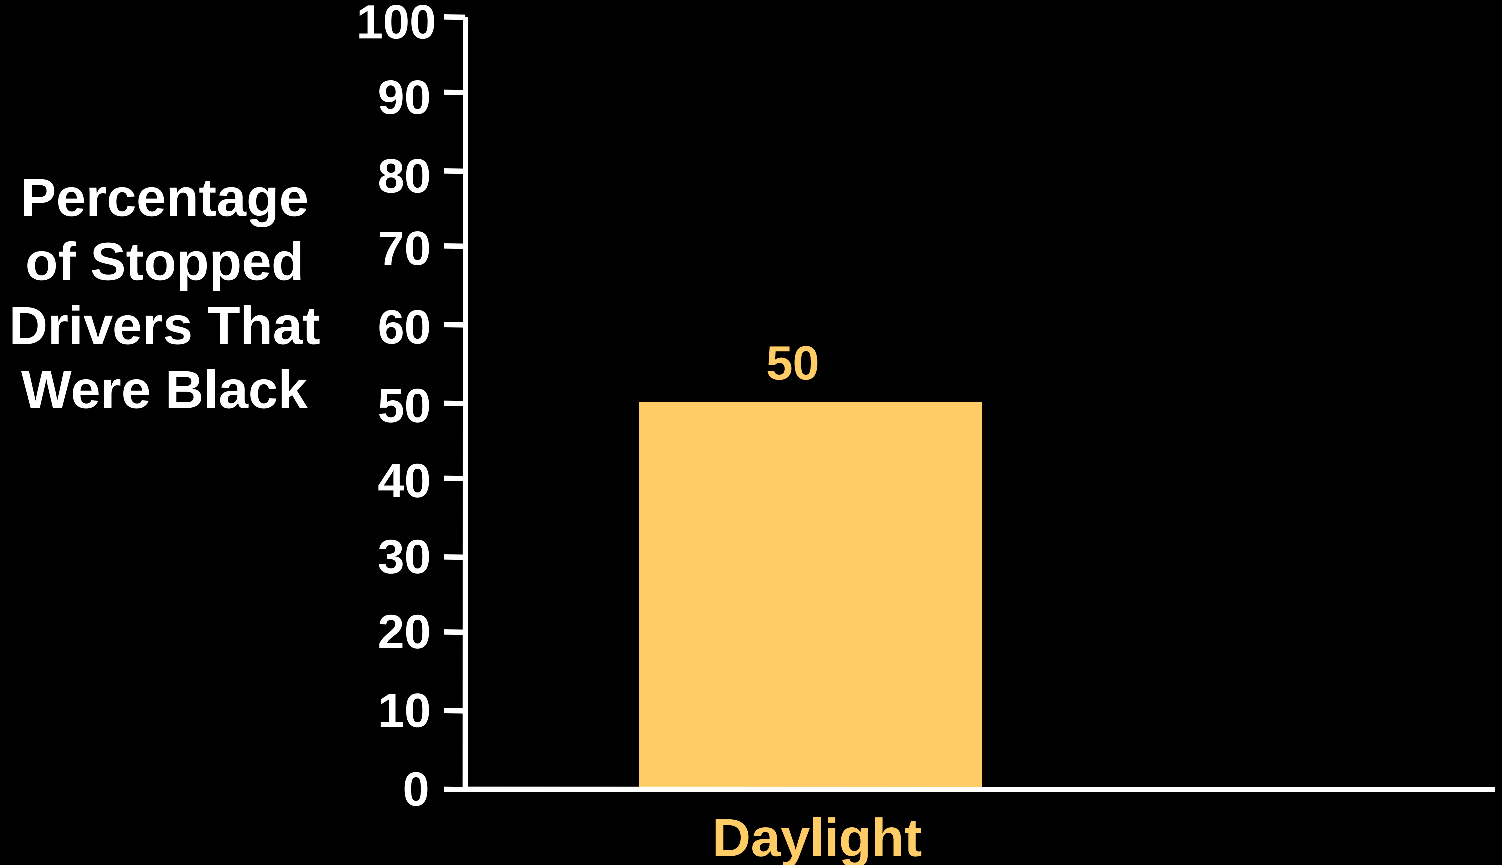

Simple "Veil of Darkness" Test Shows No Evidence of Racial Bias in the Decision to Stop

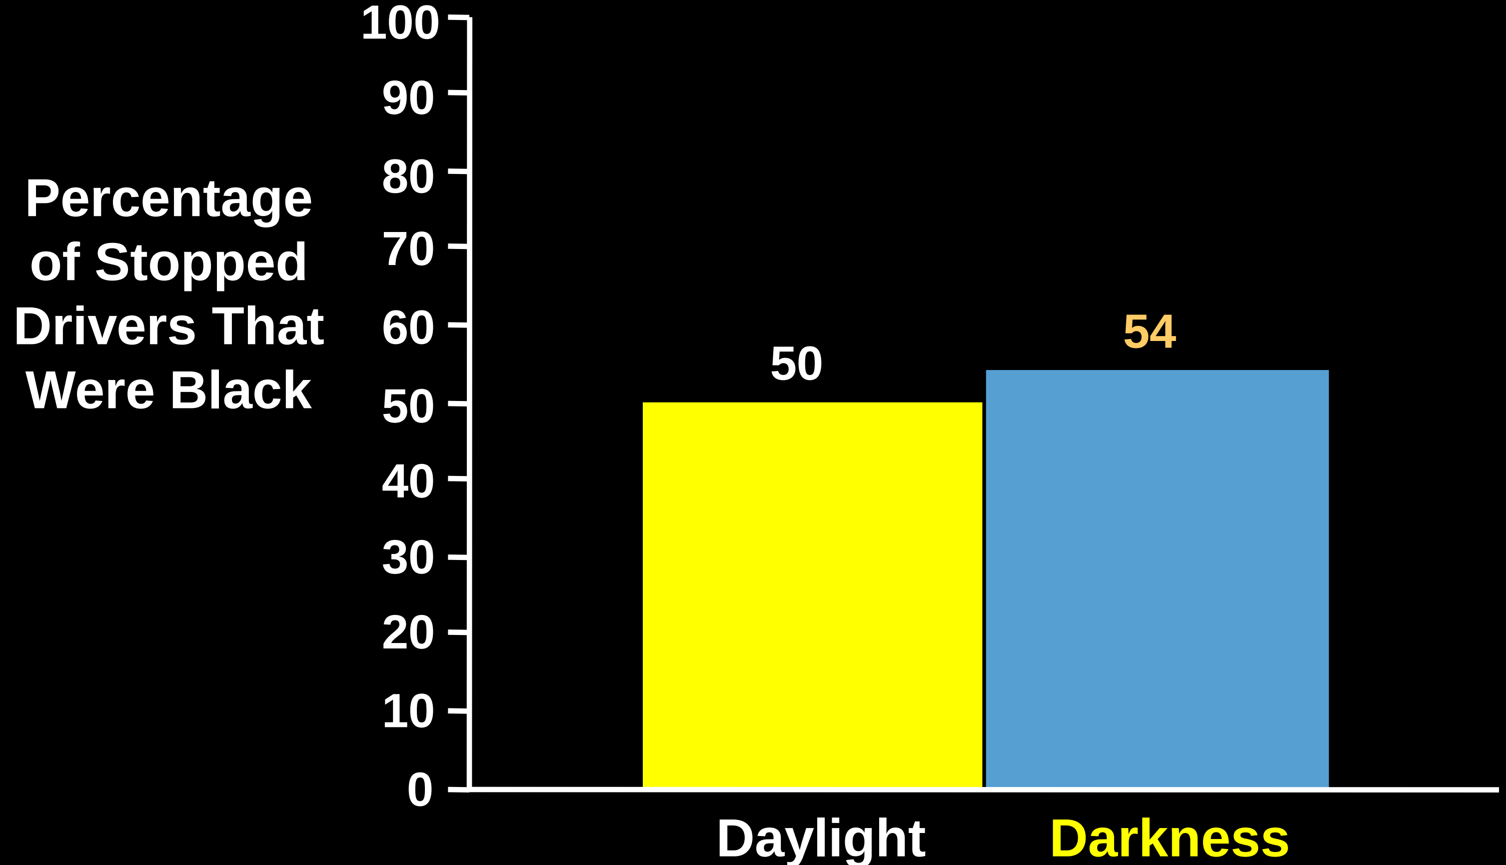

An Approach That Involved Adjusting for "Clock Time"

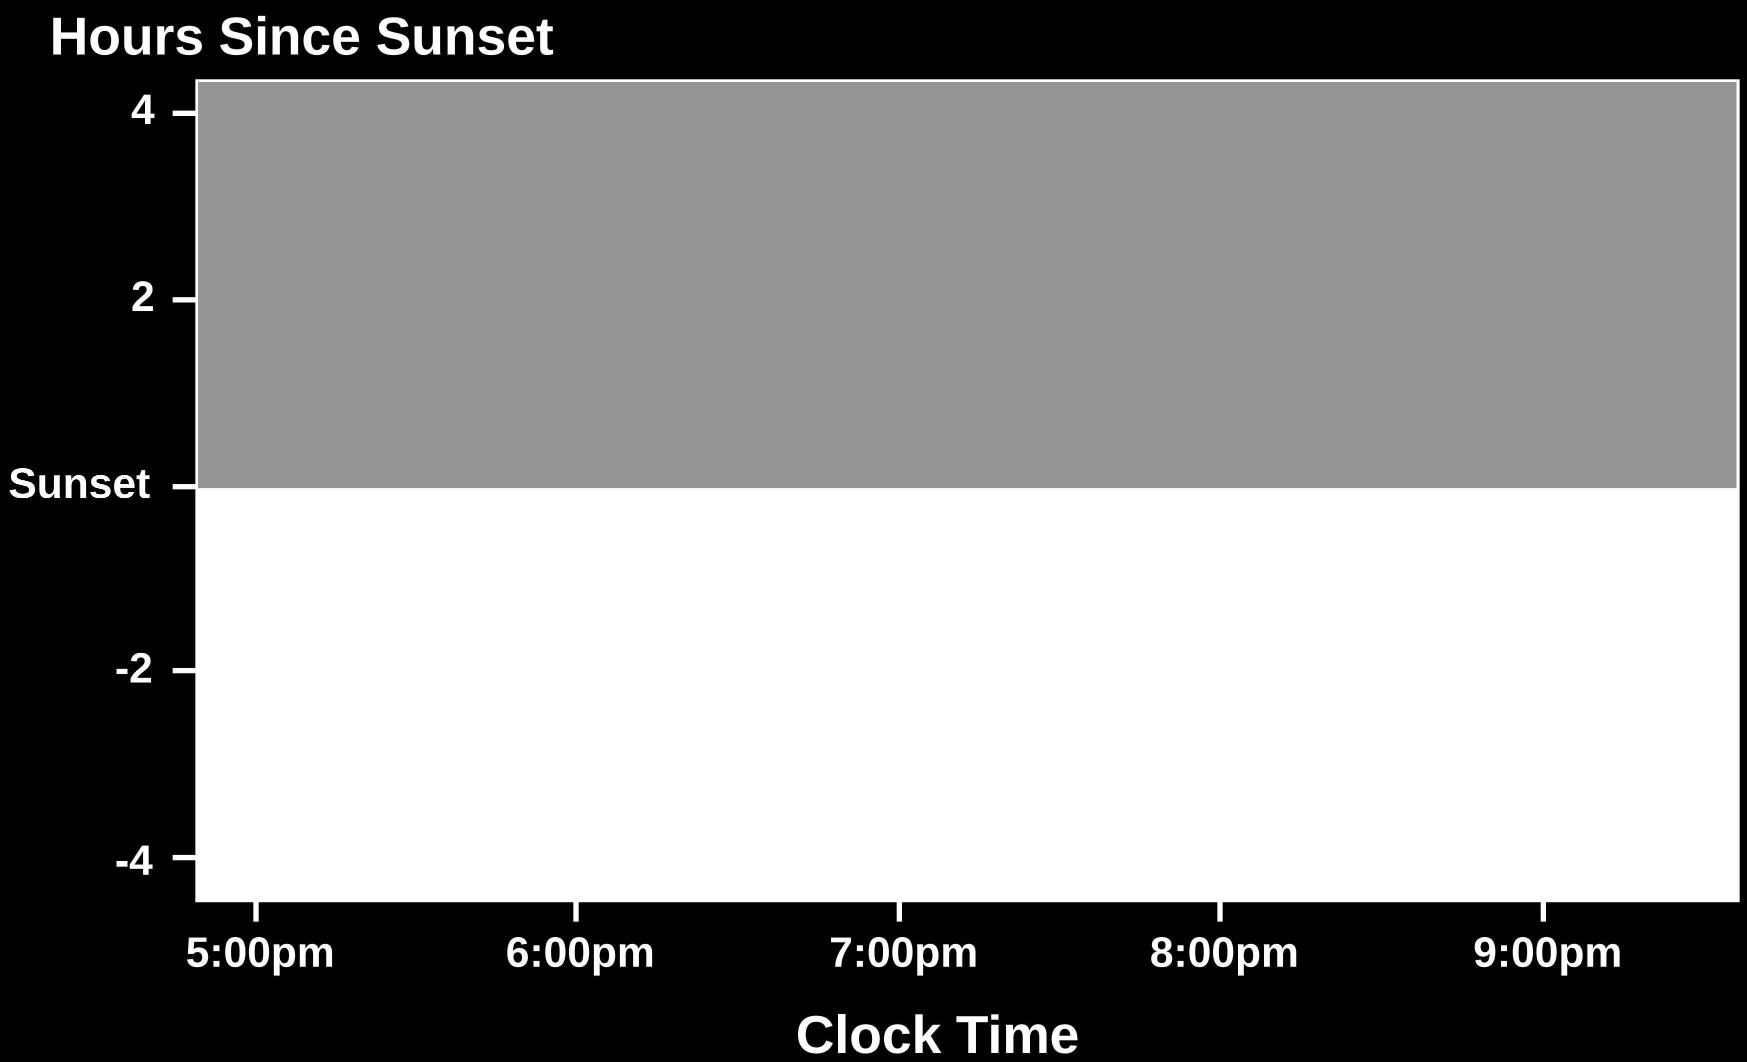

Compare Stops During Daylight with Stops in Darkness

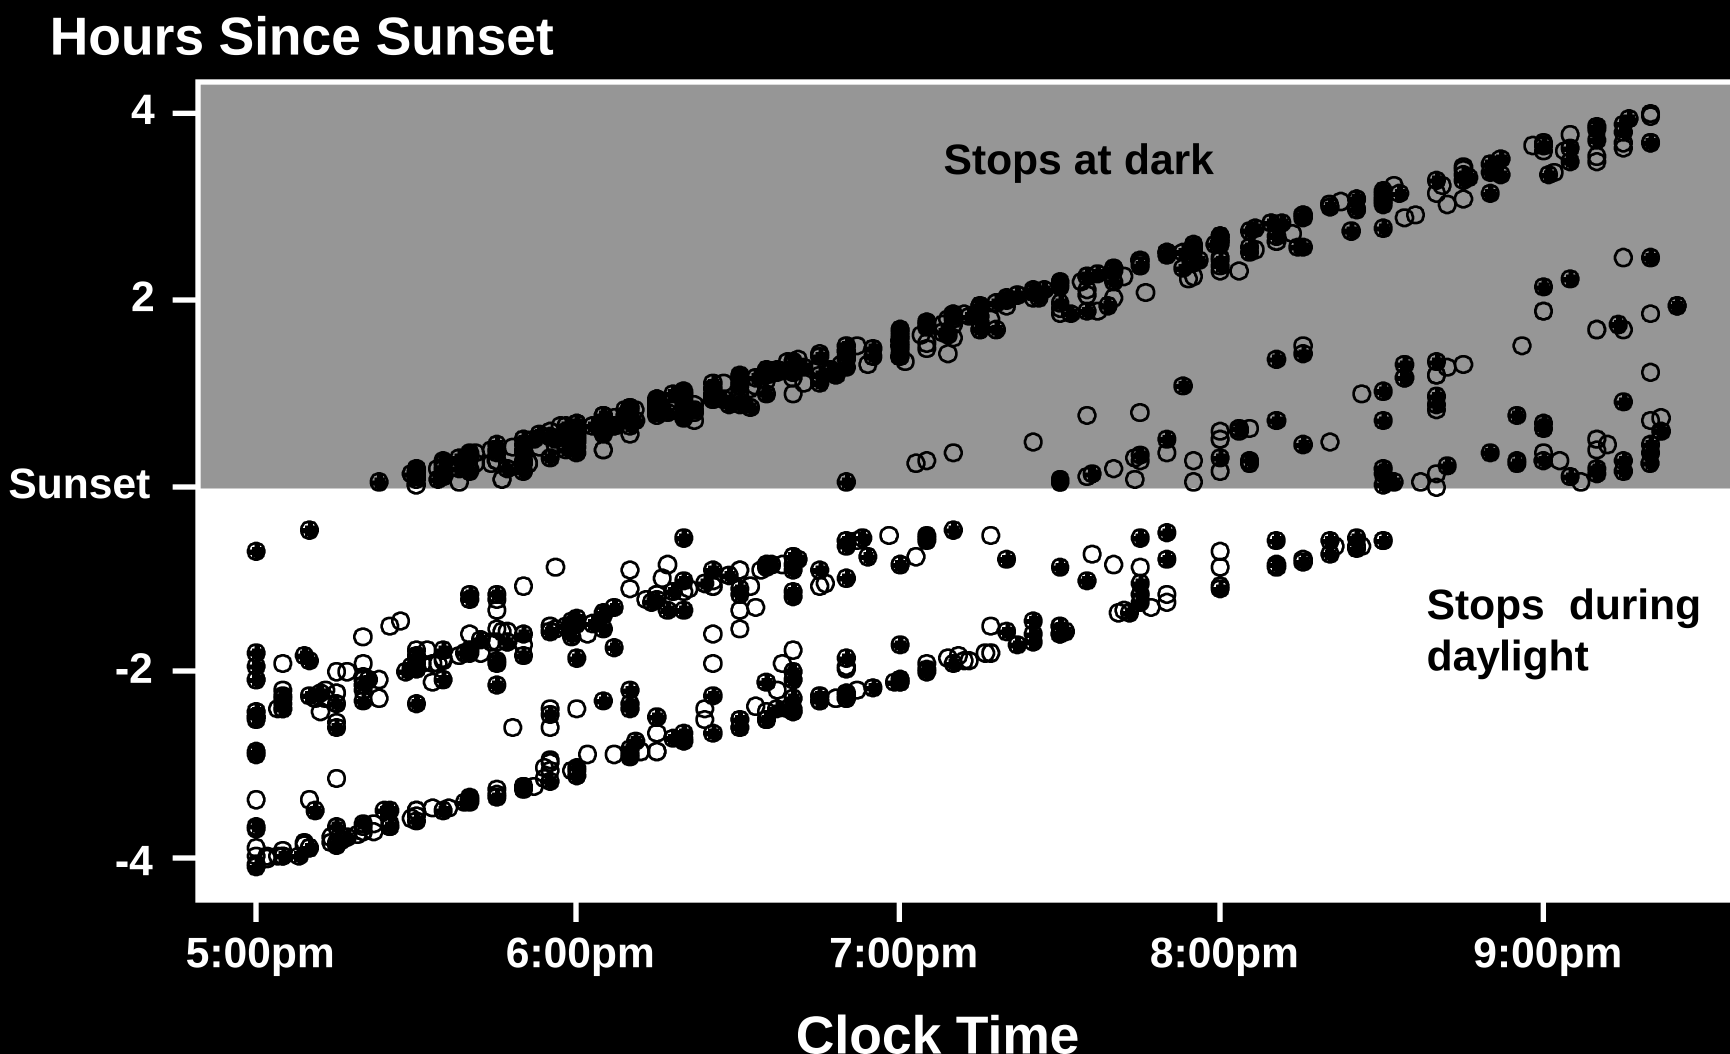

There Is No Difference in the Rate that Black Drivers Are Stopped

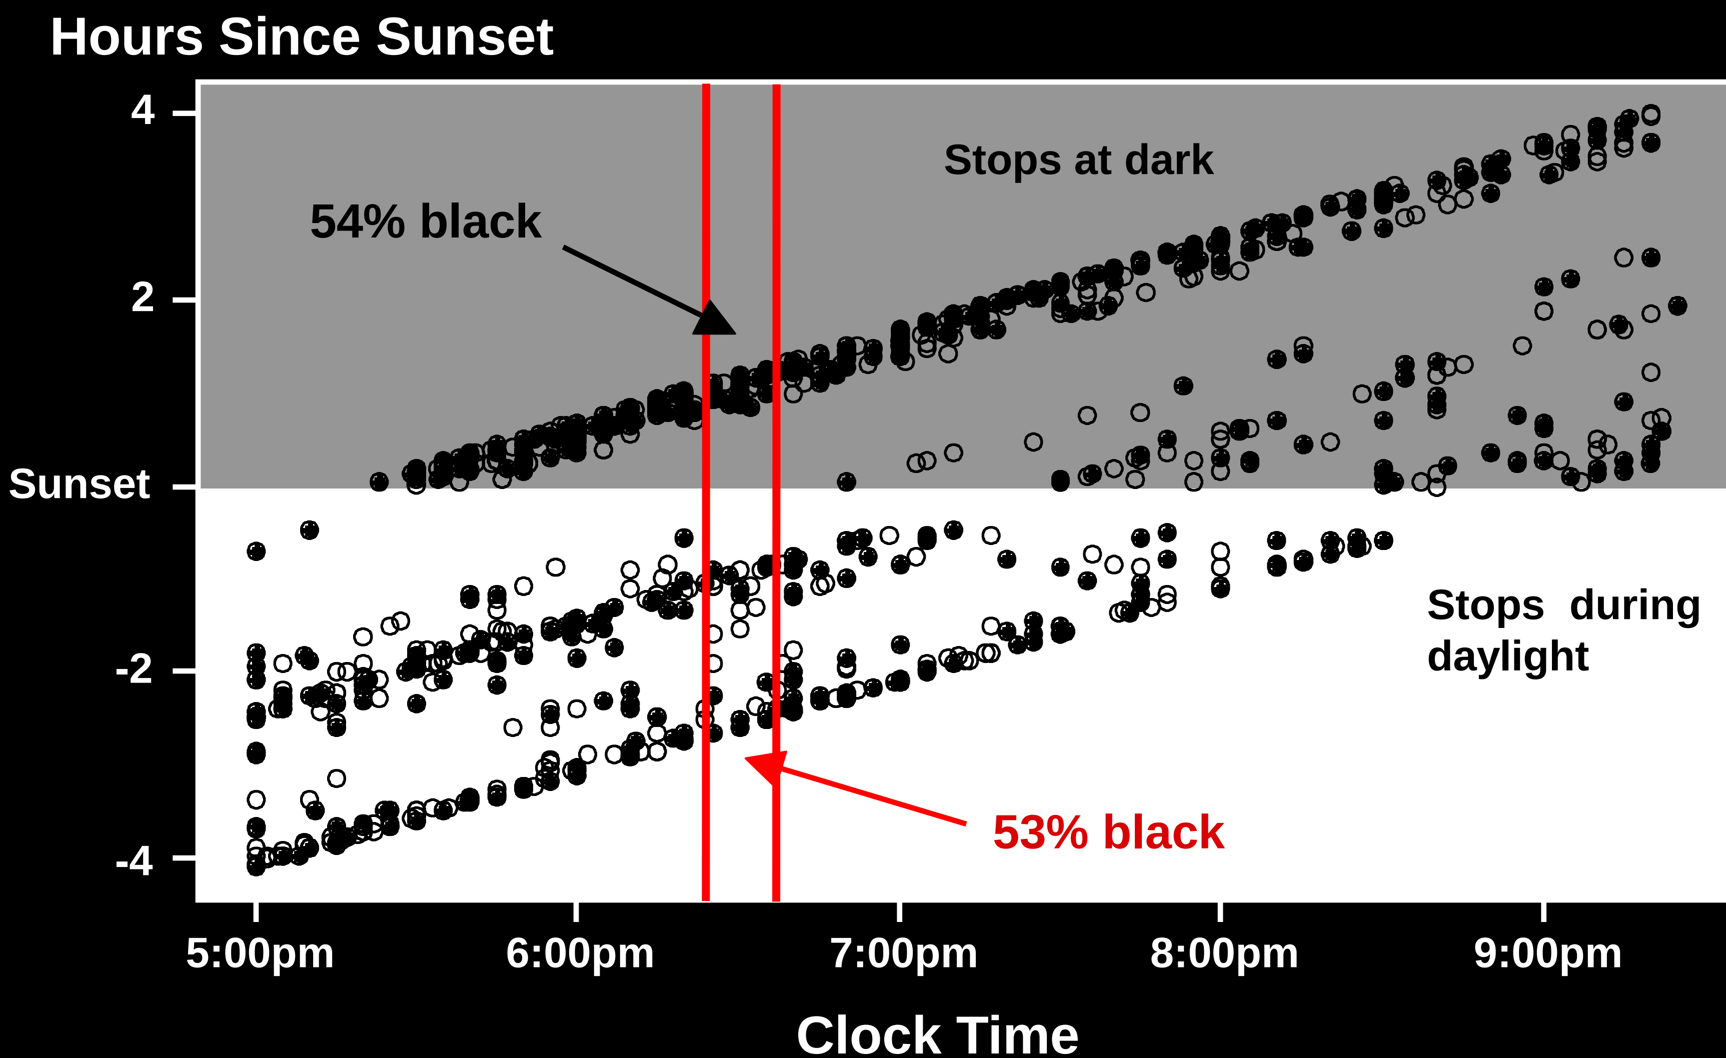

Decomposition of the VoD Estimator

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

- **S – Stop**
- **B – Black driver**
- **t – Clock time**
- **d – Darkness**

- ***K* > 1 suggests officers are more likely to stop black drivers when their race is visible**

Decomposition of the VoD Estimator

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

Decomposition of the VoD Estimator

$$K =$$

$$\frac{P(B|R,S,t,d=0)}{1-P(B|R,S,t,d=0)} \frac{1-P(B|R,S,t,d=1)}{P(B|R,S,t,d=1)} \times$$

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

Decomposition of the VoD Estimator

$$K =$$

$$\frac{P(B|R,S,t,d=0)}{1-P(B|R,S,t,d=0)}\,\frac{1-P(B|R,S,t,d=1)}{P(B|R,S,t,d=1)} \times$$

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

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

VoD is Easily Implemented with Logistic Regression

- **For each stop record race of driver, darkness indicator, and clock time**
- **Regress**

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

- **Report VoD estimate as** $K = -\beta_1$
- **Oakland 2003: *K* = 0.88**
- **Cincinnati 2003-2008: *K* = 0.96**

Outline

- Assessing disparities in the decision to stop
- **Internal benchmarking and early warning systems**
- Assessing disparities in post-stop outcomes

Example Internal Benchmark for an NYPD Officer

Stop Characteristic		Example Officer (%) *n* = 392
Month	January	3
	February	4
	March	8
Day of the week	Monday	13
	Tuesday	11
	Wednesday	14
Time of day	(4-6 p.m.]	9
	(6-8 p.m.]	8
	(8-10 p.m.]	23
	(10 p.m. -12 a.m.]	17
Patrol borough	Brooklyn North	100
Precinct	B	98
	C	1
Outside		96
In uniform	Yes	99
Radio run	Yes	1

Example Internal Benchmark for an NYPD Officer

Stop Characteristic		**Example Officer (%)** *n* **= 392**	**Internal Benchmark (%)** **ESS = 3,676**
Month	January	3	3
	February	4	4
	March	8	9
Day of the week	Monday	13	13
	Tuesday	11	10
	Wednesday	14	15
Time of day	(4-6 p.m.]	9	10
	(6-8 p.m.]	8	8
	(8-10 p.m.]	23	23
	(10 p.m. -12 a.m.]	17	17
Patrol borough	Brooklyn North	100	100
Precinct	B	98	98
	C	1	1
Outside		96	94
In uniform	Yes	99	97
Radio run	Yes	1	3

Benchmark Also Matches on Fine Location Data

86% of the Officer's Stops Were Black…

…Compared with 55% for the Benchmark

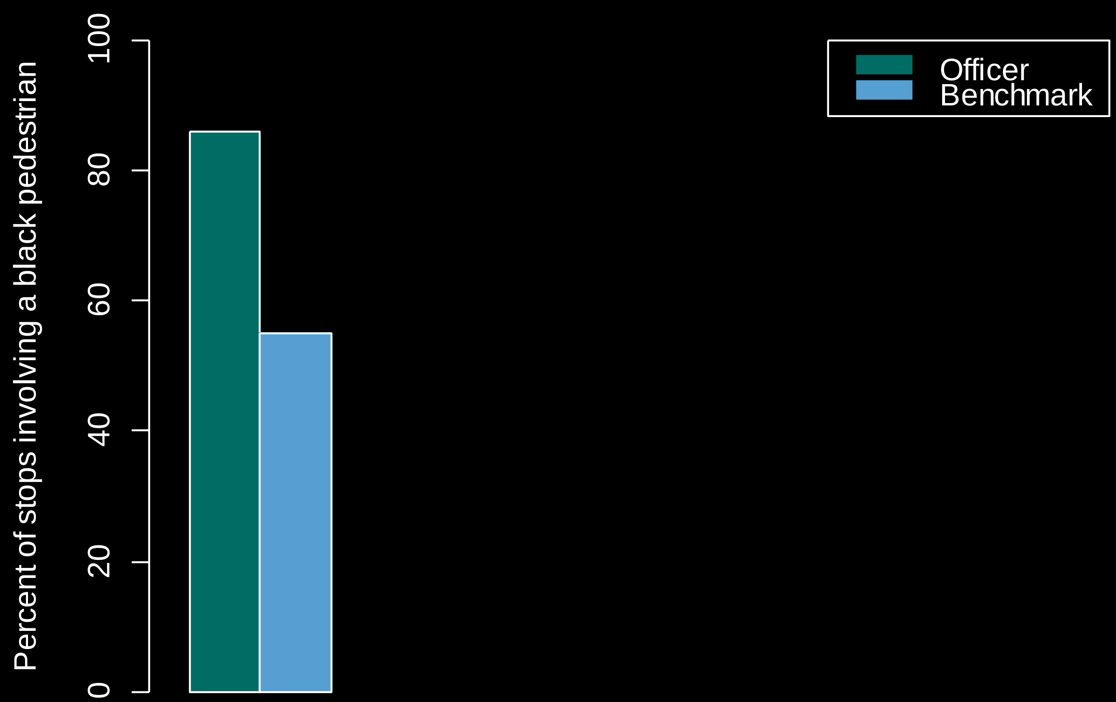

Analysis in NYPD Flagged Five Officers

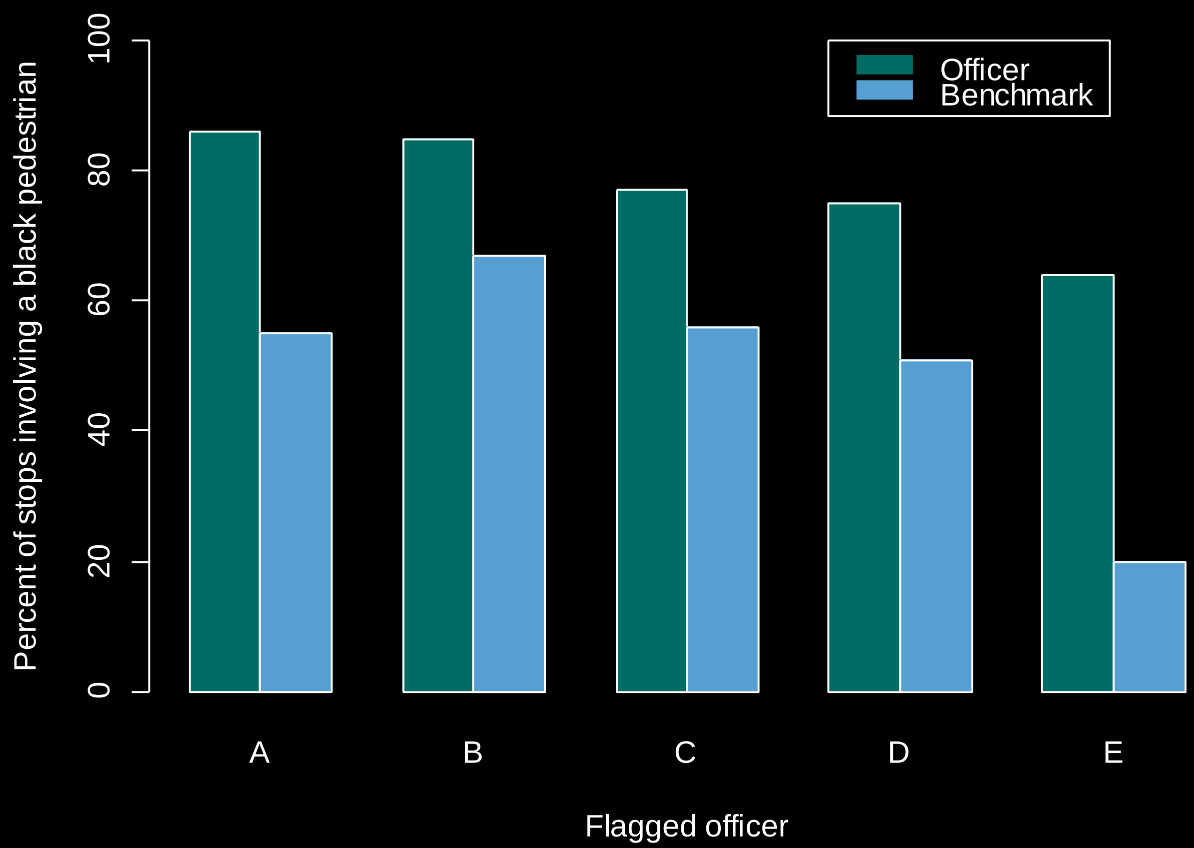

Benchmark Combines Three Modern Statistical Techniques

- **Propensity score weighting**
- **Double robust estimation**
- **False discovery rate**

G. Ridgeway and J.M. MacDonald (2009). "Doubly Robust Internal Benchmarking and False Discovery Rates for Detecting Racial Bias in Police Stops." JASA 104:661–668.

Propensity Score Weighting

- **Propensity scores reweight the other officer's stops to resemble the target officer's stops**

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

- **Estimate *p*(x) using a flexible, non-parametric version of logistic regression**

Double Robust Estimation

- **Propensity score weighted logistic regression removes remaining observed confounding**

$$\ell = \sum_{i=1}^{N} w_i(y_i f(t_i, \mathbf{x}_i) - \log(1 + \exp(f(t_i, \mathbf{x}_i))))$$

$$f(t, \mathbf{x}) = \alpha + \gamma t + \beta' \mathbf{x}$$

- **The z-test for γ = 0 will be consistent if *either* the propensity score *or* regression model is correct**

False Discovery Rates

- In the absence of racial bias the $z \sim N(0,1)$
- For 2,756 correlated zs, an appropriate reference distribution can be much wider (Efron 2006)

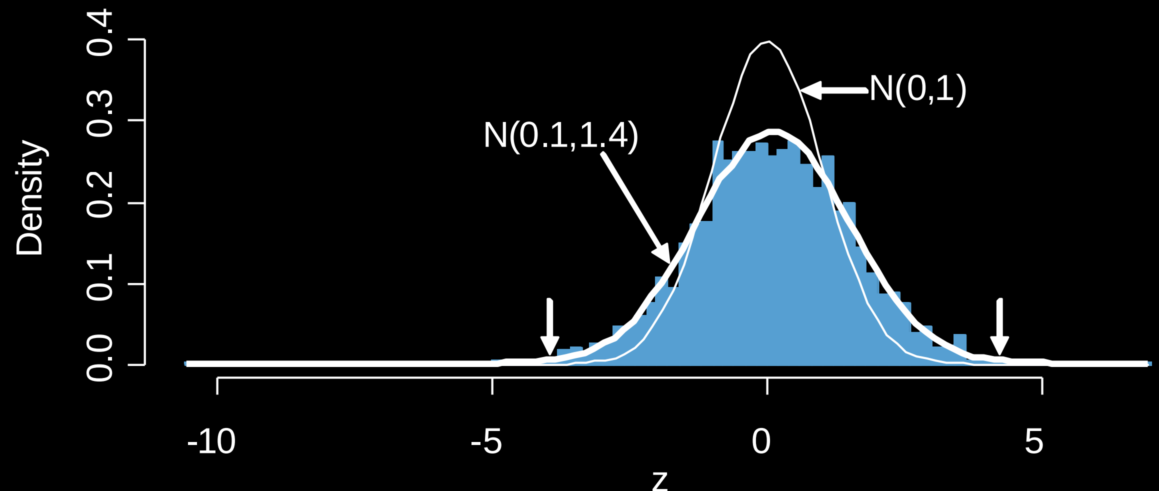

- $P(\text{problem} \mid z) \geq 1 - f_0(z) / f(z)$
- Standard cutoff of z > 2.0 flags 242 officers, 90% of which have fdr estimated to be greater than 0.999

Outline

- Assessing disparities in the decision to stop
- Internal benchmarking and early warning systems
- **Assessing disparities in post-stop outcomes**

Post Stop Outcomes

- Auditing police-citizen interactions
  - Video taped analysis
- Hit Rates
  - Comparing yields from contraband searches
- Matching on characteristics of stopped citizens
  - Comparing race groups who are similarly situated
  - Use the same methodology for matching officers' stops

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

Cincinnati Reported Large Disparities in Stop Duration

Stop feature	% Black drivers N=26,941	% Nonblack drivers (unadjusted) N=25,149
Stop < 10 minutes	55	65

Black and Nonblack Drivers Differ in Numerous Ways

Stop feature	% Black drivers N=26,941	% Nonblack drivers (unadjusted) N=25,149
Stop < 10 minutes	55	65
Invalid license	22	7
Male	65	66
Neighborhood		
Over-the-Rhine	9	5
Avondale	5	1
I-75	4	11
Residence		
Cincinnati	93	61
Date\Time		
12am-4am	16	8
Monday	15	14
August	9	11
Age		
18-25	33	29
Reason		
Equipment violation	27	16
Moving violation	51	76

There Are Similarly Situated Nonblack Drivers

Stop feature	% Black drivers N=26,941	% Nonblack drivers (weighted) ESS=4,952	% Nonblack drivers (unadjusted) N=25,149
Stop < 10 minutes	55		65
Invalid license	22	20	7
Male	65	65	66
Neighborhood			
Over-the-Rhine	9	10	5
Avondale	5	5	1
I-75	4	5	11
Residence			
Cincinnati	93	92	61
Date\Time			
12am-4am	16	16	8
Monday	15	15	14
August	9	9	11
Age			
18-25	33	32	29
Reason			
Equipment violation	27	28	16
Moving violation	51	52	76

No Significant Difference in Stop Duration

Stop feature	% Black drivers N=26,941	% Nonblack drivers (weighted) ESS=4,952	% Nonblack drivers (unadjusted) N=25,149
Stop < 10 minutes	55	57	65
Invalid license	22	20	7
Male	65	65	66
Neighborhood			
Over-the-Rhine	9	10	5
Avondale	5	5	1
I-75	4	5	11
Residence			
Cincinnati	93	92	61
Date\Time			
12am-4am	16	16	8
Monday	15	15	14
August	9	9	11
Age			
18-25	33	32	29
Reason			
Equipment violation	27	28	16
Moving violation	51	52	76

RAND

Most of the Search Rates Disparity Is Also Due to Non-Racial Factors

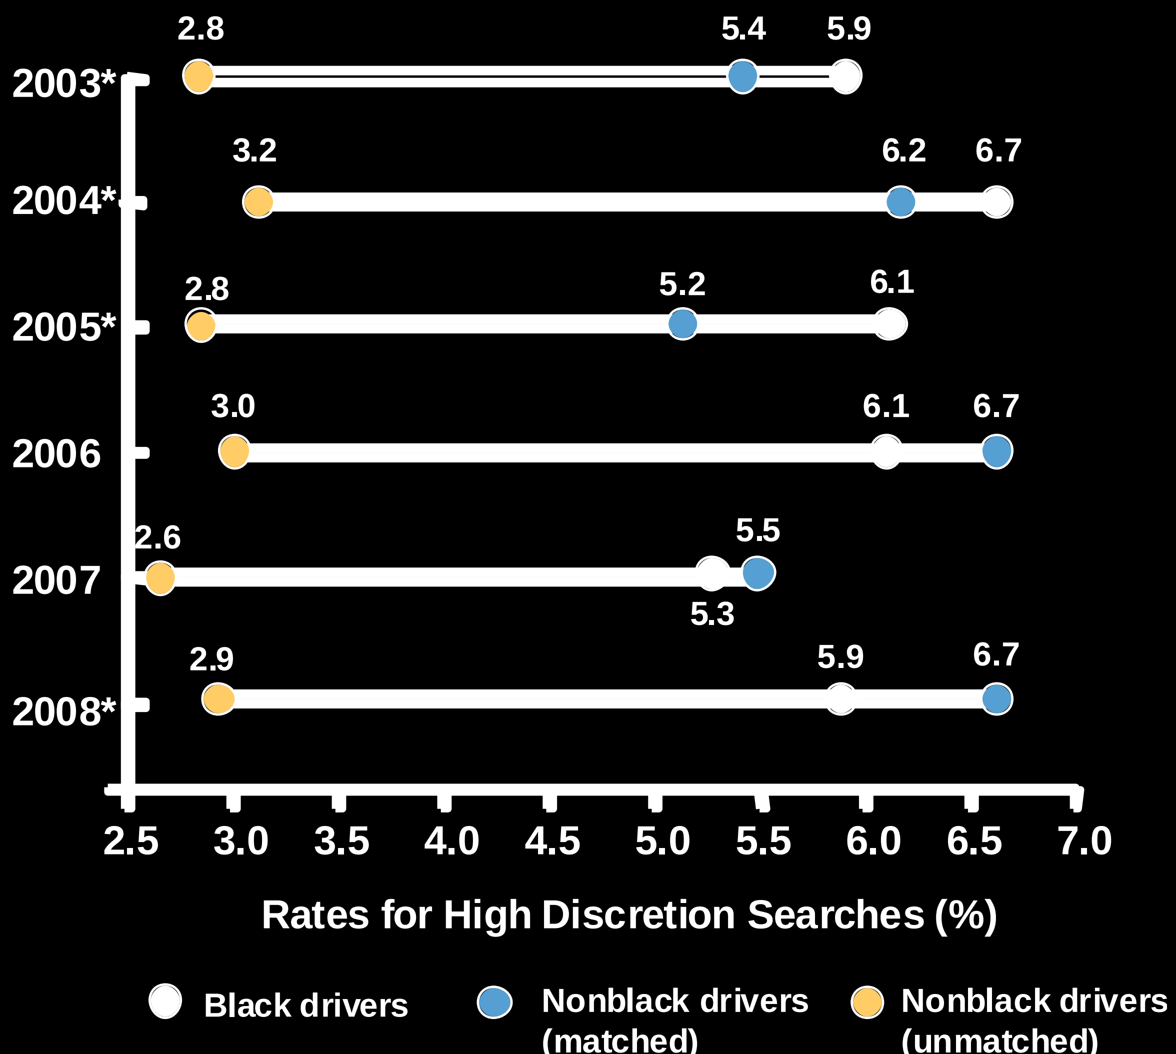

In Closing

- Naïve analysis methods can exaggerate or understate the effect of racial bias
- Three transparent methods for assessing racially biased policing
  - VoD natural experiment
  - Internal benchmarking
  - Comparisons of similarly situated drivers
- cqp.rand.org for the latest

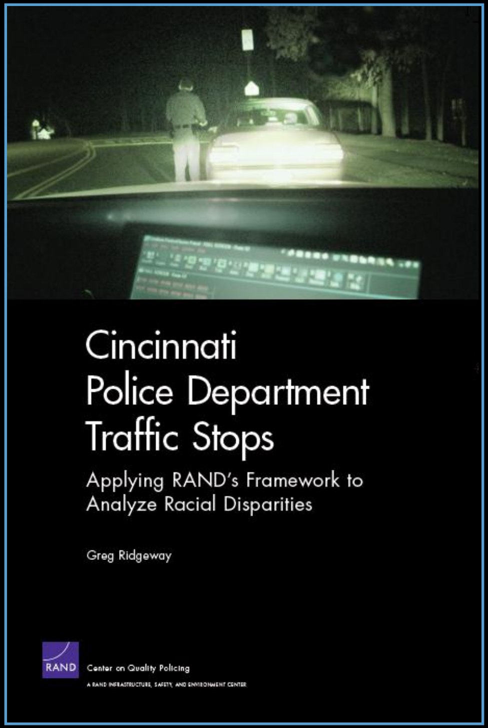

RAND

Center on Quality Policing

A RAND INFRASTRUCTURE, SAFETY, AND ENVIRONMENT PROGRAM