INFRASTRUCTURE, SAFETY,
AND ENVIRONMENT

Assessing the effect of race bias in post-traffic stop outcomes using propensity scores

Greg Ridgeway

Race might influence outcomes of the stop

- **Goal: Determine whether race affects the decision**
  - **to issue a citation versus a warning**
  - **to conduct a consent search**
  - **to pat search**
  - **to detain the vehicle for more than 10 minutes**

Race Groups Live in Different Oakland Neighborhoods

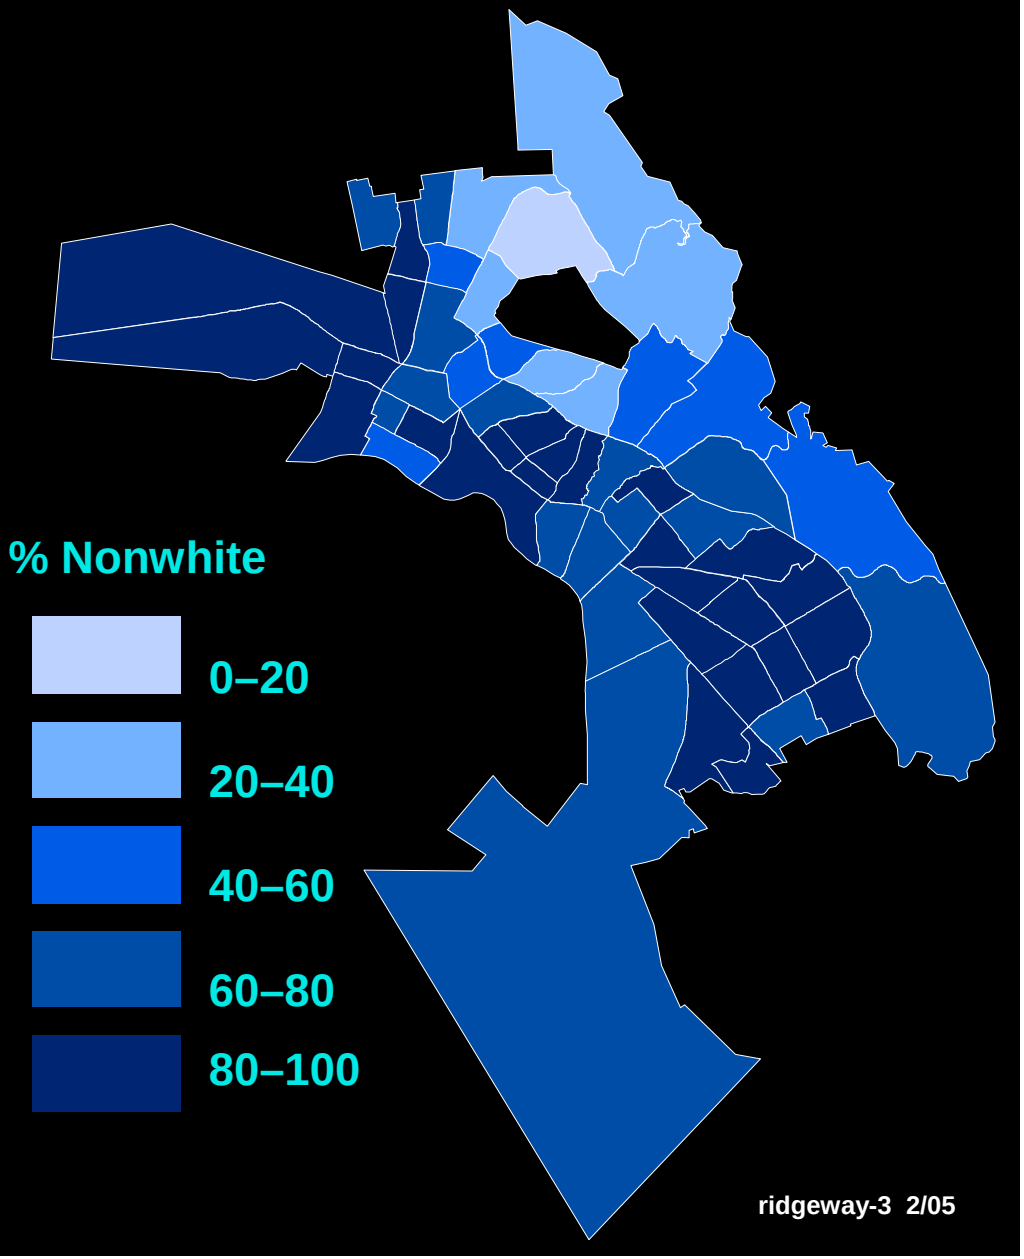

RAND

ridgeway-3 2/05

Black and Non-Black Drivers are Stopped in Different Parts of the City

% Stopped by Neighborhood

Region	Black	Non-Black
Downtown	31%	27%
East	32%	13%
Hills	1%	3%
Midtown	12%	21%
North	9%	8%
South Hills	3%	6%
West	14%	21%

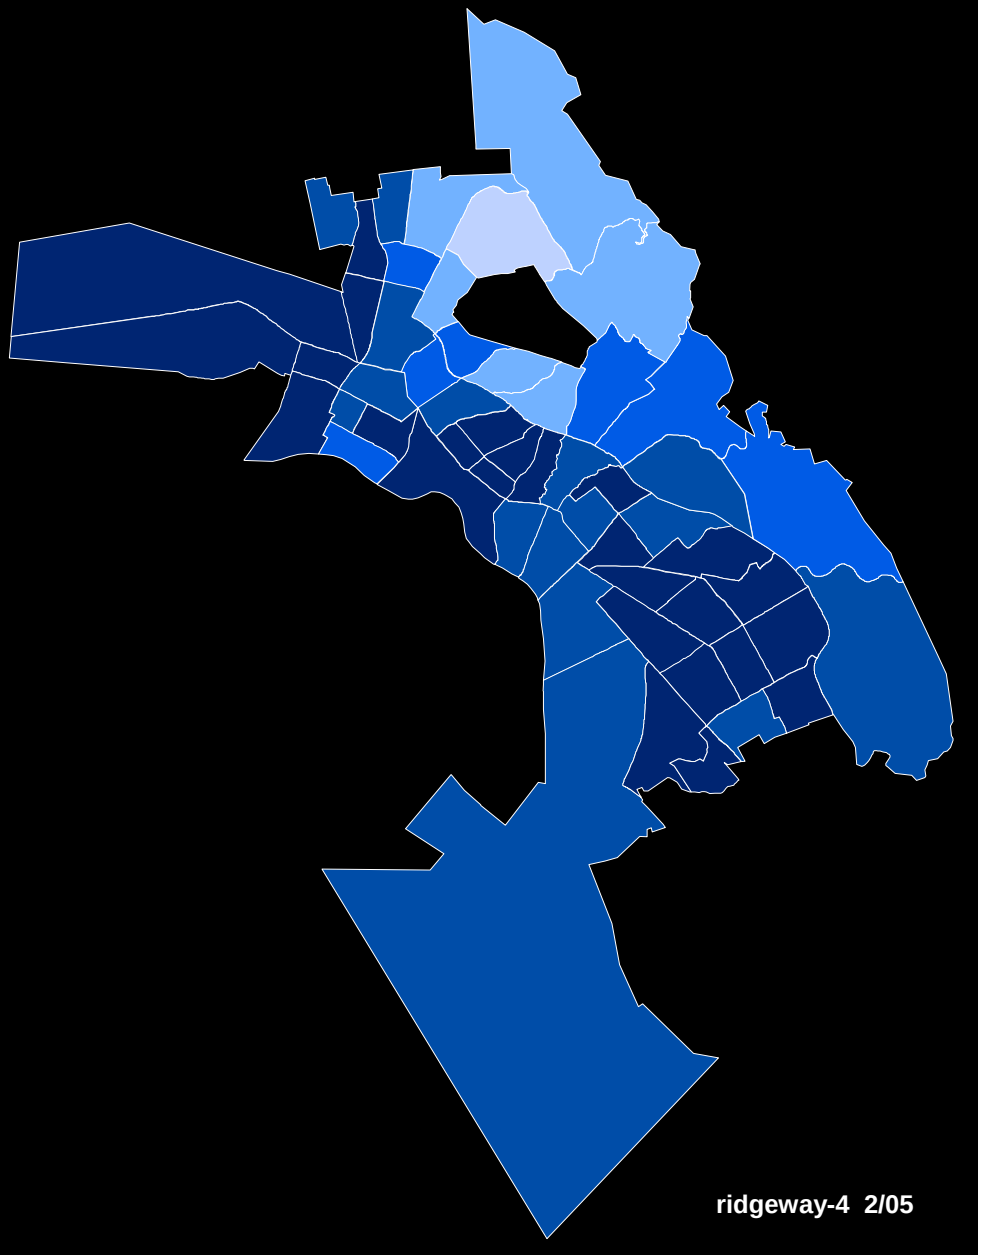

Black and Non-Black Drivers are Stopped in Different Parts of the City

% Stopped by Neighborhood

Region	Black	Non-Black
Downtown	31%	27%
East	32%	13%
Hills	1%	3%
Midtown	12%	21%
North	9%	8%
South Hills	3%	6%
West	14%	21%

Highest citation rates

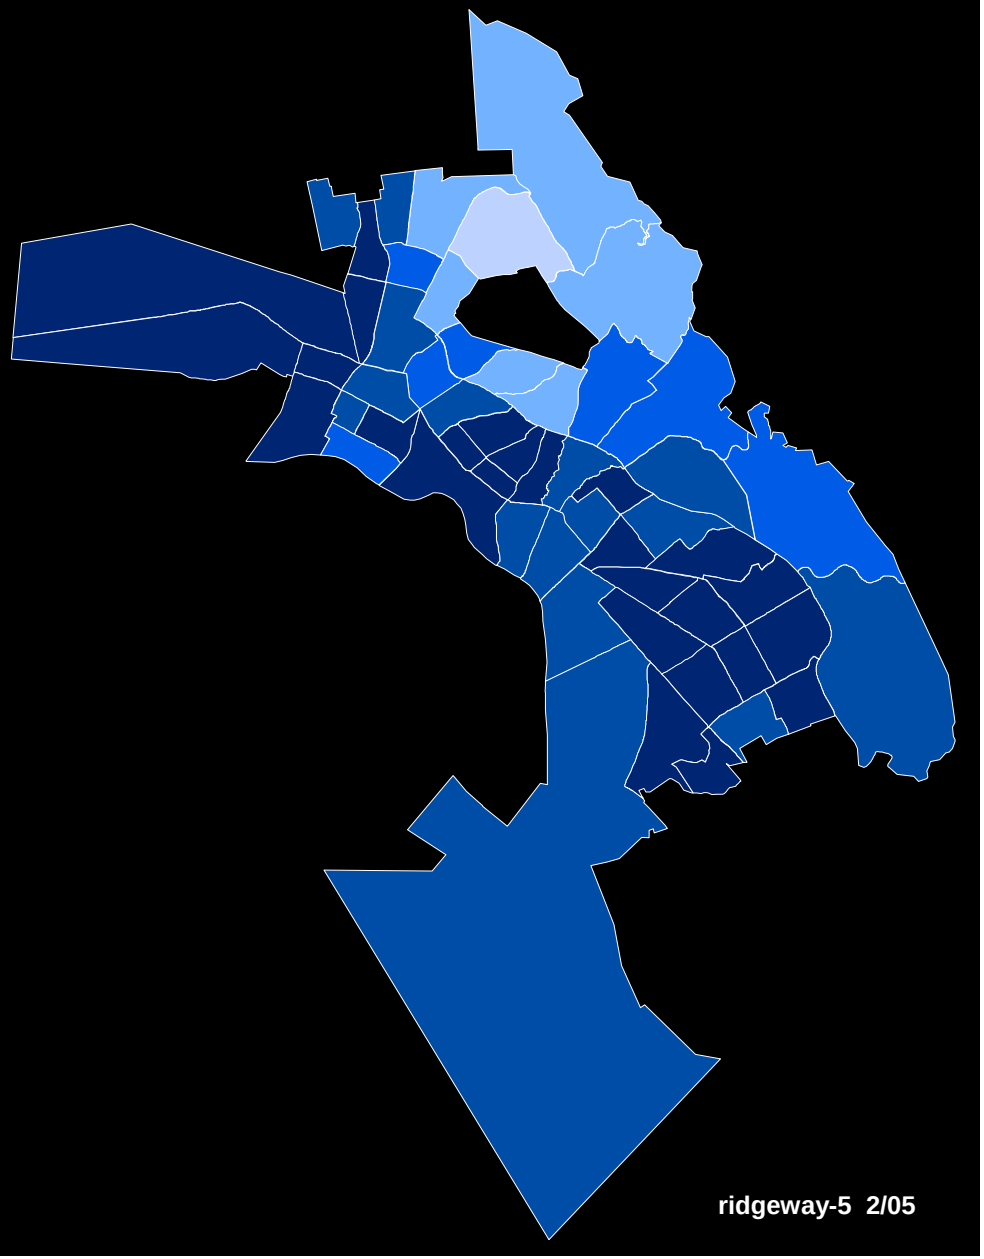

RAND

Is the Difference Due to Race or Neighborhood?

% Stopped by Neighborhood

Region	Black	Non-Black
Downtown	31%	27%
East	32%	14%
Hills	1%	3%
Midtown	12%	21%
North	9%	8%
South Hills	3%	6%
West	14%	21%

Highest citation rates

- **Citation rates**
  - **Black: 68%**
  - **Non-black: 79%**
- **Are black drivers stopped for no good reason?**
- **Are black drivers driving in neighborhoods where officers don't write tickets?**

Traffic Stops Differ on Several Features

	% Black drivers *N* = 3,703	% Non-black drivers *N* = 3,033
Region		
East	31.9%	14.1%
Hills	0.5%	2.7%
Time of day		
12:00am-4:00am	15.8%	7.5%
12:00pm-4:00pm	19.8%	28.2%
Reason		
Mechanical/Registration	25.7%	16.2%
Moving (dangerous)	22.5%	37.2%
Resident	75.8%	63.6%
Age		
18-29	46.5%	38.4%
Male	75.1%	73.6%

Regression approaches are not robust

- **The "first impulse" analysis is logistic regression**

$$\begin{aligned}\text{logit}(P(\text{citation}|\mathbf{x}, \text{black})) \\ = \beta_0 + \beta_1 \text{black} + \\ \beta_2 \text{neighborhood} + \\ \beta_3 \text{time of day} + \ldots\end{aligned}$$

- **When the distribution of stop features differ for black and white drivers, the estimate of β_1 is *not* robust to violations of the additivity assumption**
- **Regression is intended to adjust for *small* differences between the groups**

Propensity score weighting equalizes the two groups

- **Weight the feature distribution of the comparison group so that it matches that of the target group**

$$f(\mathbf{x}|\text{race} = \text{black}) = w(\mathbf{x})\, f(\mathbf{x}|\text{race} \neq \text{black})$$

Propensity score weighting equalizes the two groups

- **Weight the feature distribution of the comparison group so that it matches that of the target group**

$$f(\mathbf{x}|\text{race} = \text{black}) = w(\mathbf{x})\, f(\mathbf{x}|\text{race} \neq \text{black})$$

- ***w*(x) reduces to a propensity score weight**

$$w(\mathbf{x}) = K\frac{f(\text{race} = \text{black}|\mathbf{x})}{1 - f(\text{race} \neq \text{black}|\mathbf{x})}$$

- **Reweight each non-black driver by $p_i / (1 - p_i)$, estimating p_i using non-linear logistic regression**

Propensity score weighting equalizes the two groups

	% Black drivers *N* = 3,703	% Non-black drivers (weighted) ESS = 1,689.2	% Non-black drivers (unweighted) *N* = 3,033
Region			
East	31.9%	32.7%	14.1%
Hills	0.5%	0.6%	2.7%
Time of day			
12:00am-4:00am	15.8%	15.5%	7.5%
12:00pm-4:00pm	19.8%	19.7%	28.2%
Reason			
Mechanical/Registration	25.7%	25.8%	16.2%
Moving (dangerous)	22.5%	22.6%	37.2%
Resident	75.8%	75.5%	63.6%
Age			
18-29	46.5%	48.2%	38.4%
Male	75.1%	76.0%	73.6%

Analysis Shows That a Race Disparity in Citation Rates Might Exist

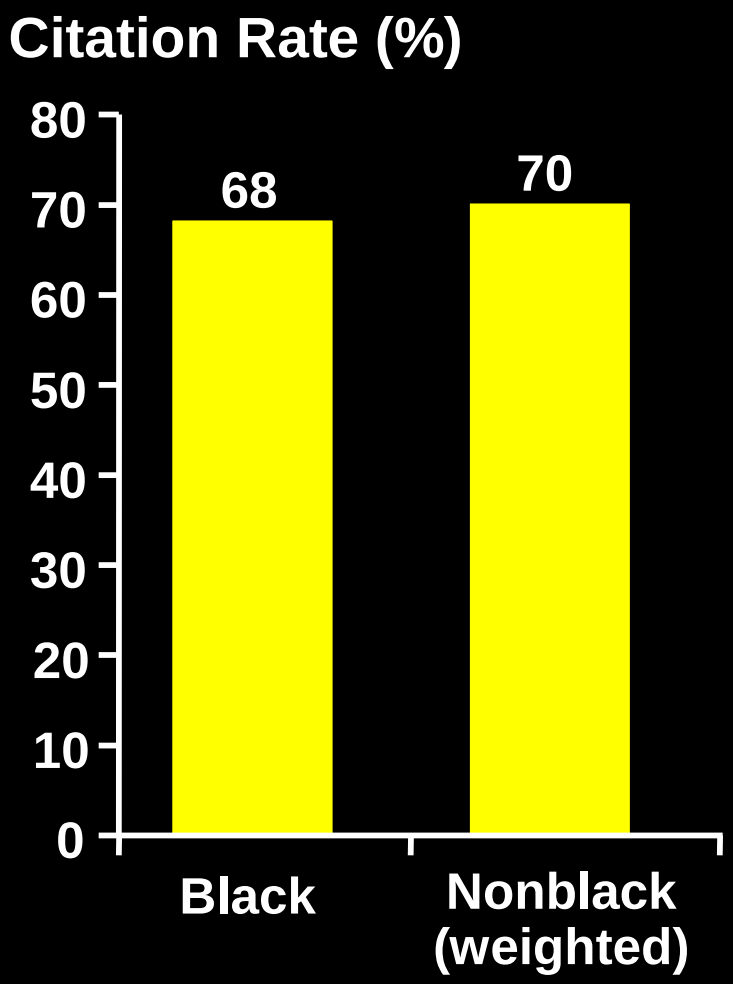

- **Citation rate for black drivers is 2% less than for comparable non-black drivers**

But the Analysis Also Shows the Danger of Making Naïve Comparisons

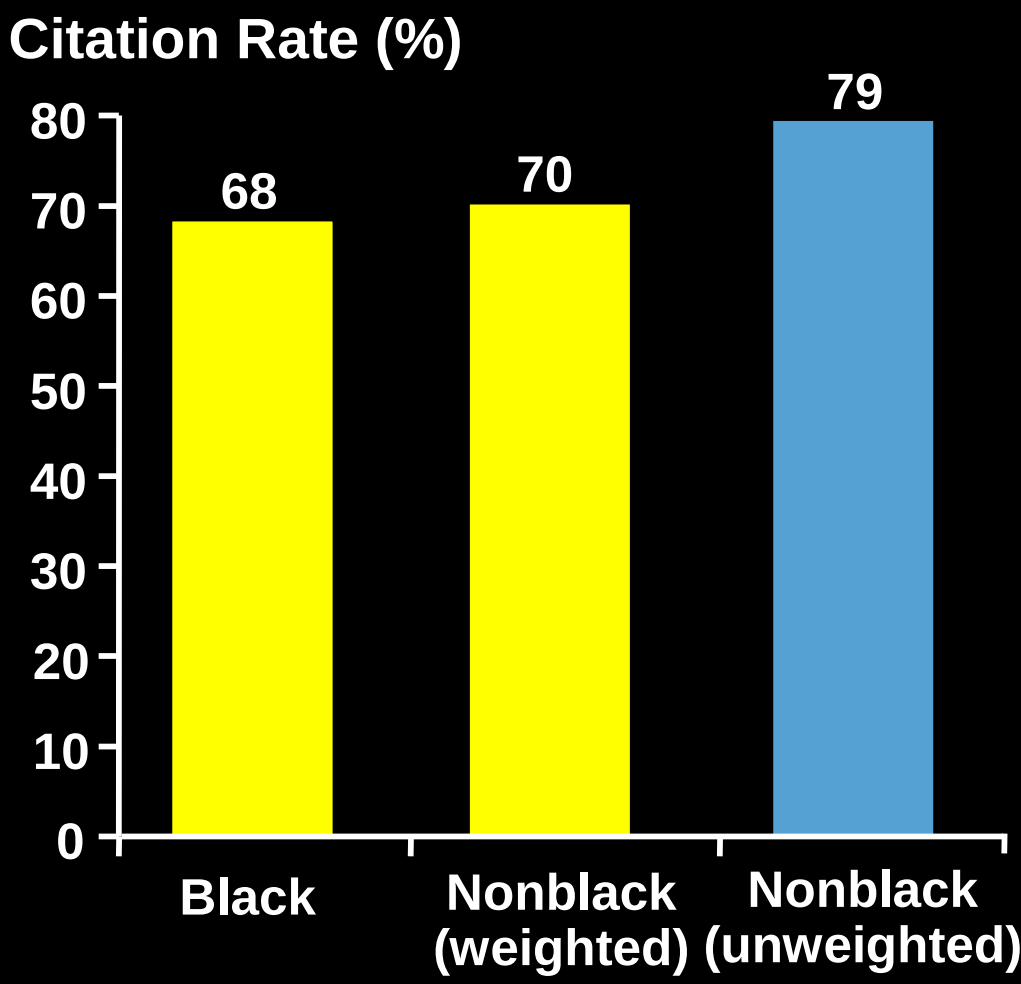

- When we compare black vs. nonblack (unweighted), difference is 11%
- Logistic regression gives an adjusted citation rate of 72%

Consent Searches Have Similar Rates

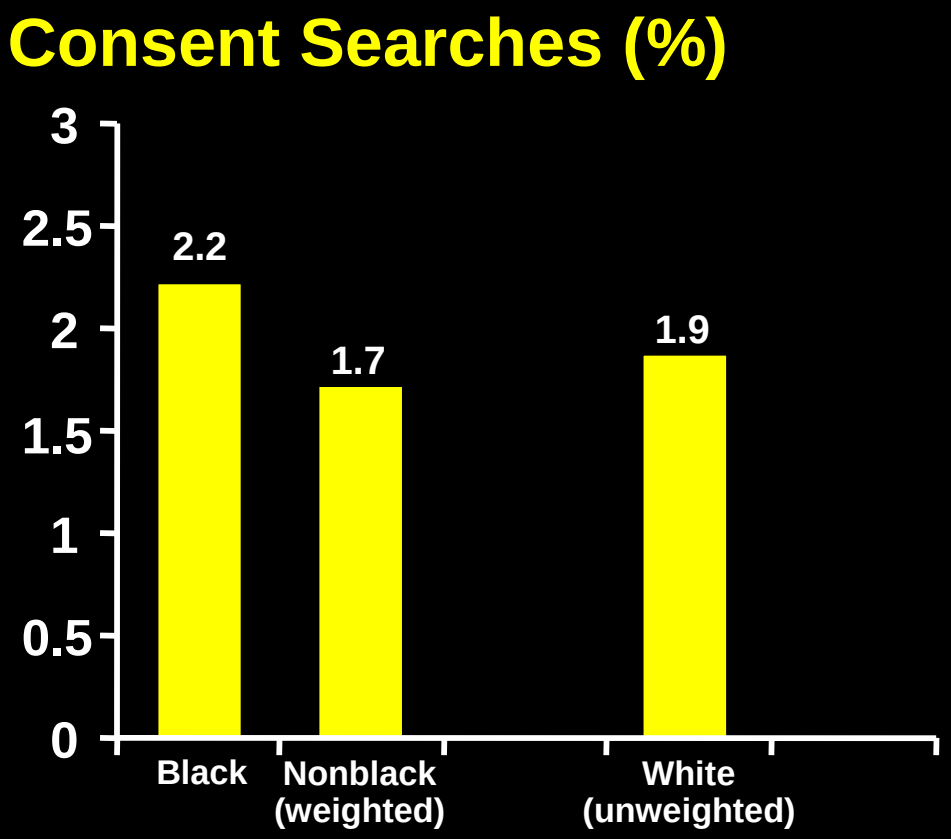

Consent Searches Have Similar Rates, but Pat Searches More Likely for Blacks Than Whites

Consent Searches (%)

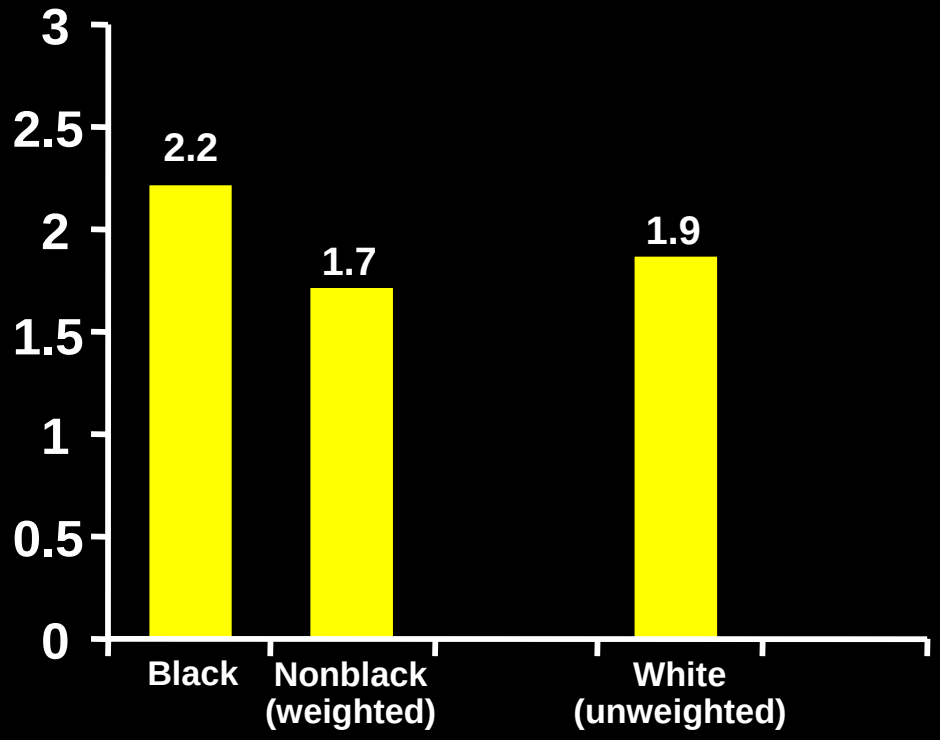

Pat Searches (%)

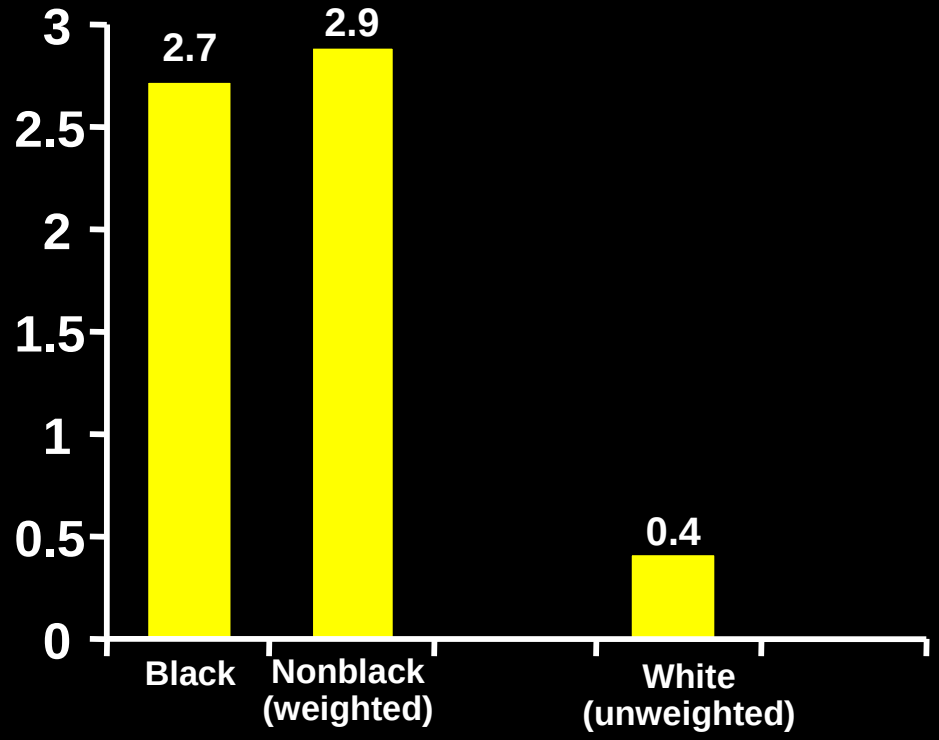

Once Again, Naïve Comparisons Can Distort the Findings

Consent Searches (%)

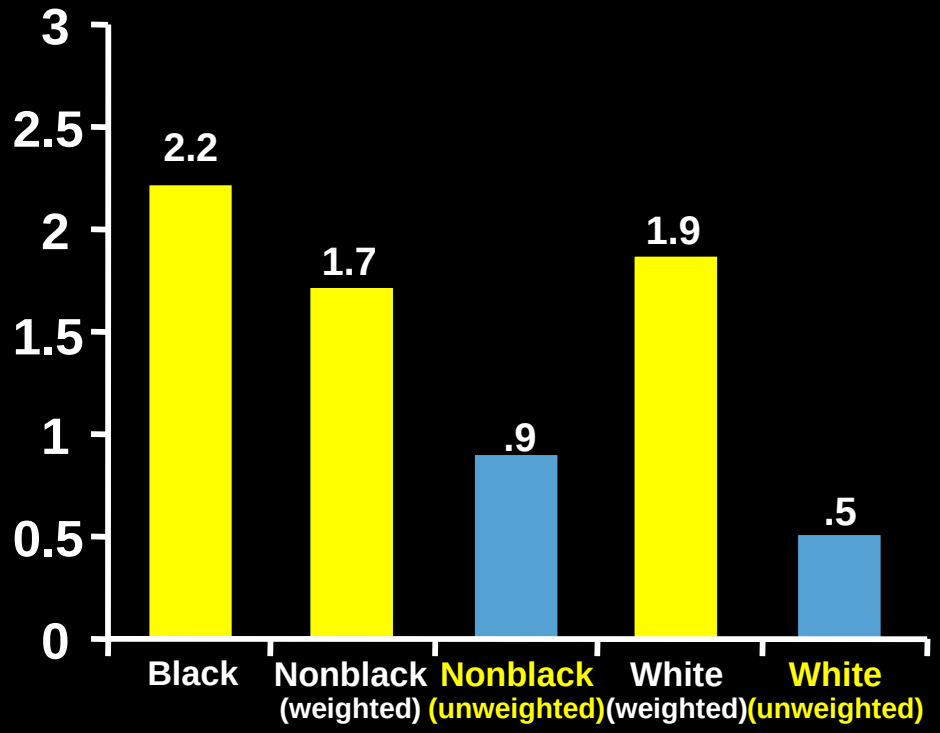

Pat Searches (%)

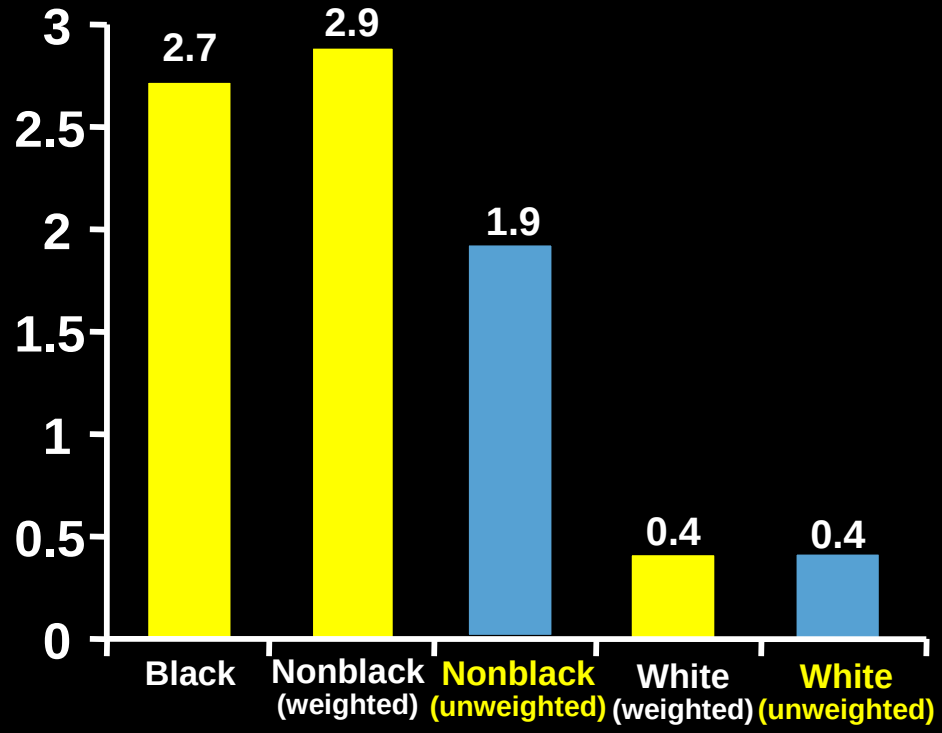

Black Drivers Seemed More Likely to Have Longer Stops Than Non-Black or White Drivers

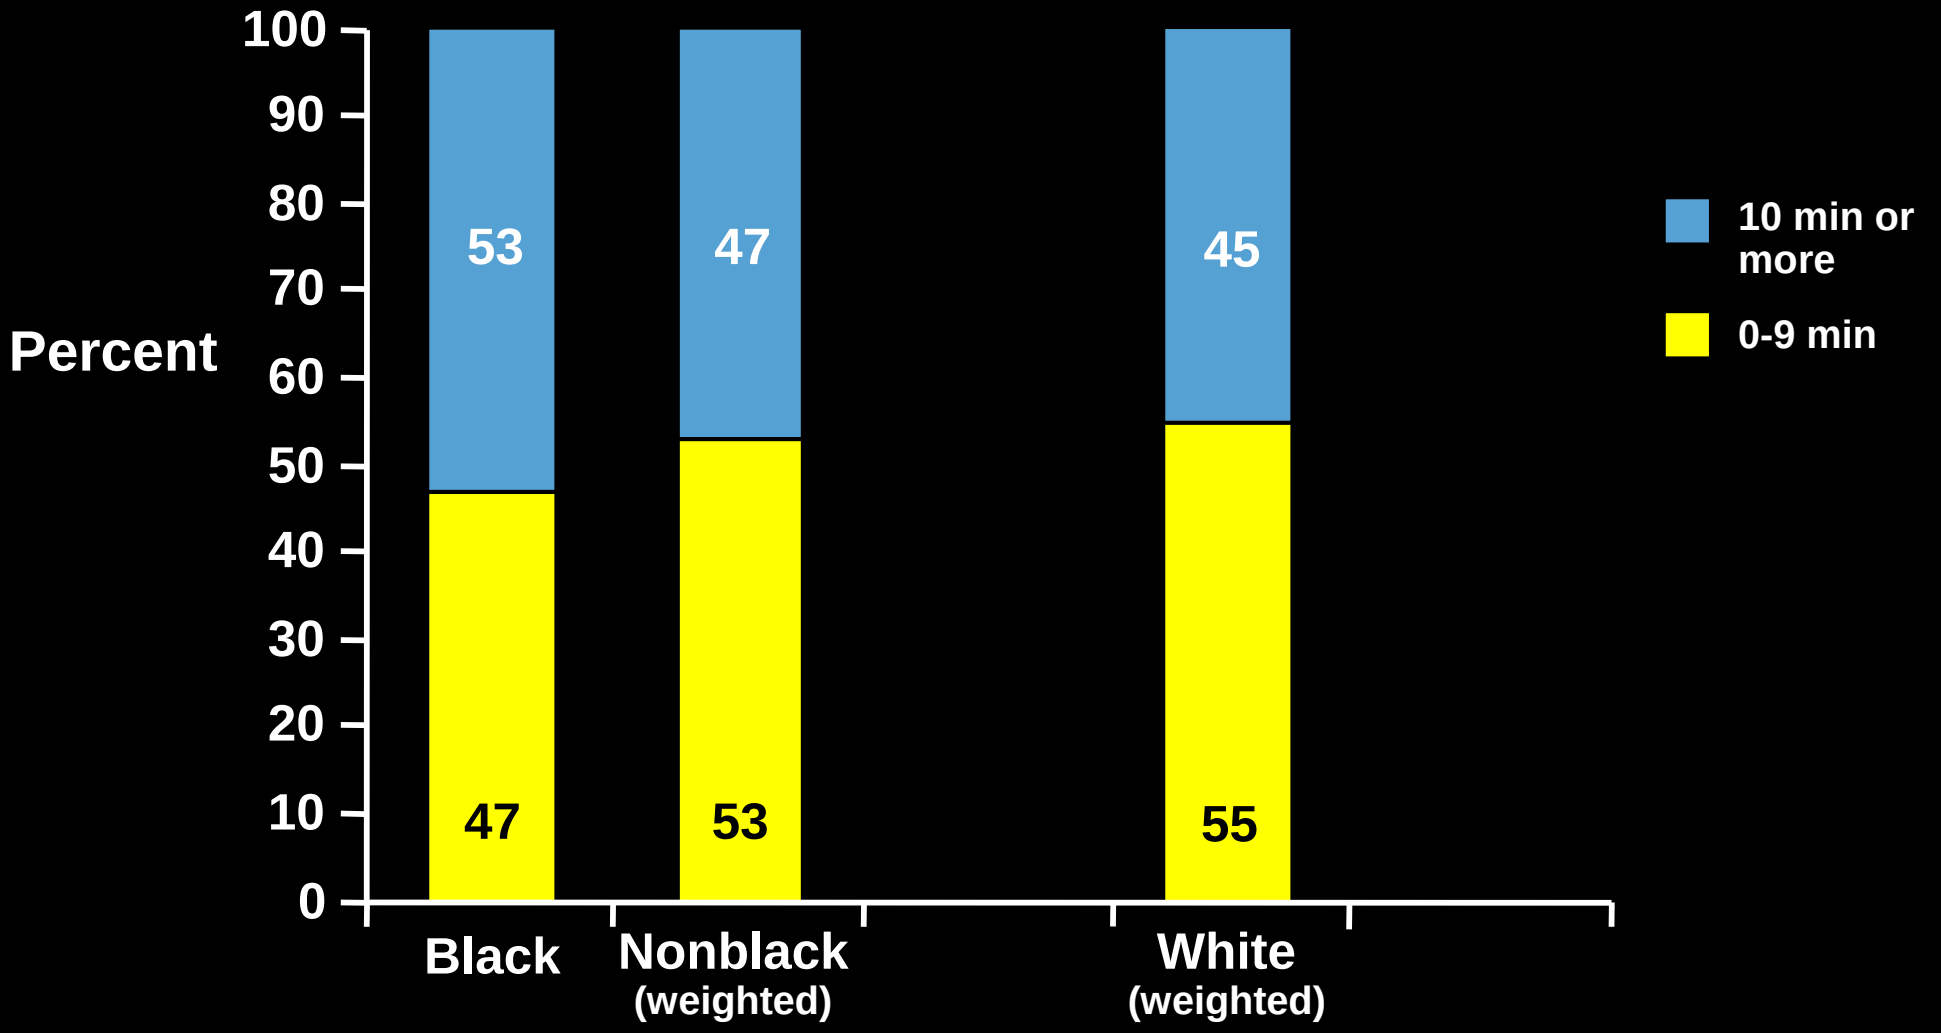

RAND

Naïve Comparisons Considerably Overstate the Problem

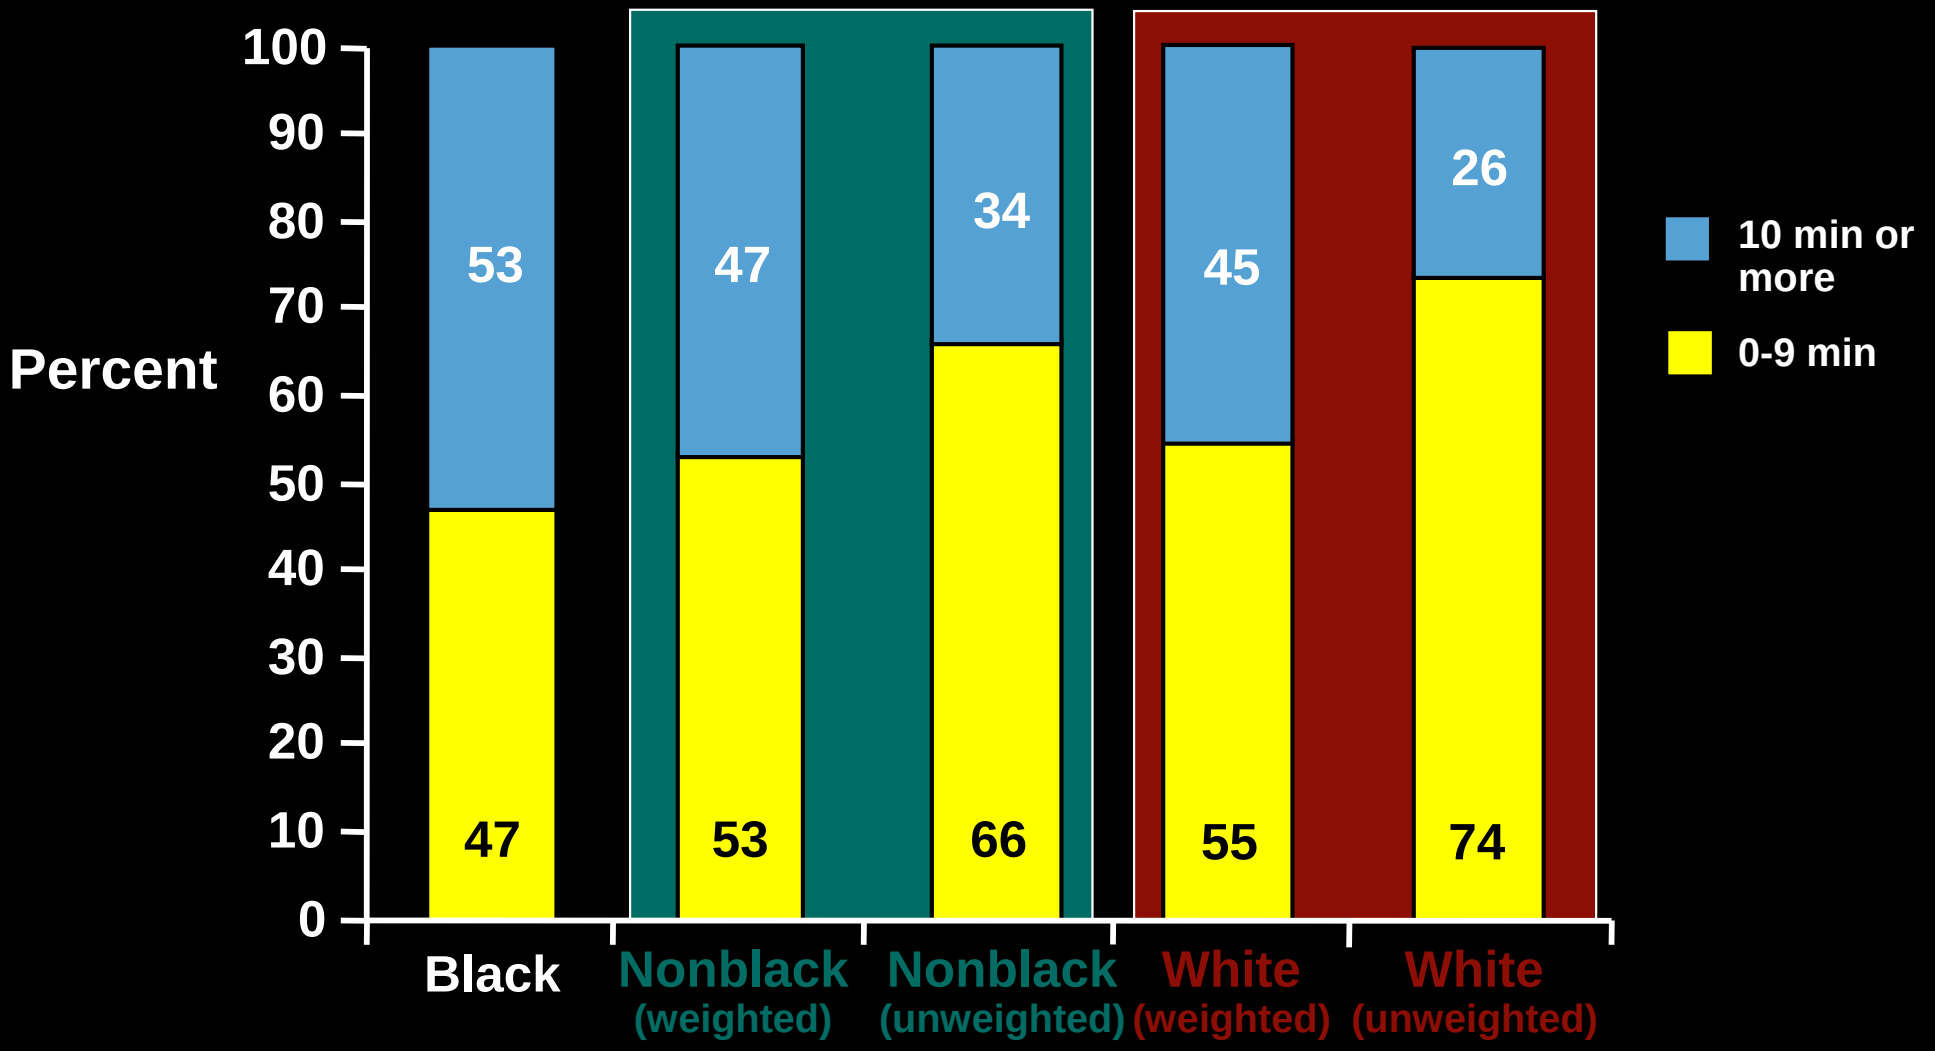

Summary: Is There a Race Bias in Oakland?

- **Citations**: Small differences in citation rates
- **Search**: Frequency of pat searches is greater among black drivers than against similarly situated white drivers
- **Duration**: Black drivers are more likely to be detained for more than 10 minutes than similarly situated drivers

Methodological Conclusions

- Naïve comparisons can exaggerate (or even understate) the effect of racial bias
- Multivariate regression does not work when the race groups differ greatly in stop features
- Propensity score weighting
  - balances the groups on all observed features
  - Model diagnosis is transparent
  - Results are easy to present to community members and police officers

INFRASTRUCTURE, SAFETY,
AND ENVIRONMENT

G. Ridgeway (2006, to appear). “Assessing the effect of race bias in post-traffic stop outcomes using propensity scores.” *Journal of Quantitative Criminology*.

Software available at www.i-pensieri.com/gregr

or google “Greg Ridgeway” and select the one that’s not a plumber in Florida