U.S. Department of Justice
Office of Community Oriented Policing Services

COPS
COMMUNITY ORIENTED POLICING SERVICES
U.S. DEPARTMENT OF JUSTICE

INNOVATIONS
A Closer Look

A Suggested Approach to Analyzing Racial Profiling:
Sample Templates for Analyzing Car-Stop Data

The CNA Corporation

Authors:
Joyce McMahon, Ph.D.,
Amanda Kraus, Ph.D.,
The CNA Corporation

A Suggested Approach to Analyzing Racial Profiling:
Sample Templates for Analyzing Car-Stop Data

By
Joyce McMahon, Ph.D.,
Amanda Kraus, Ph.D.,
The CNA Corporation

This project was supported by cooperative agreement #2002-HS-WX-0061 by the Office of Community Oriented Policing Services, U.S. Department of Justice. The opinions contained herein are those of the author(s) and do not necessarily represent the official position of the U.S. Department of Justice. References to specific companies, products, or services should not be considered an endorsement of the product by the author(s) or the U.S. Department of Justice. Rather, the references are illustrations to supplement discussion of the issues.

Introduction

Because of community concerns about the use of race as a significant factor in police decision-making, many police departments have collected and used data from car stops to investigate the issue of racial profiling, or bias-based policing. Many more departments are considering doing so. Most existing evaluations based on car-stop data have relied on aggregate comparisons between the racial composition of stops in a city and the racial composition of the city's total population. Unfortunately, this simple, aggregate approach can be dangerous. By failing to include important information on police operational procedures, specific city circumstances, and socio-demographic characteristics of neighborhoods, the aggregate approach is too simplistic and may generate misleading results.

Decisions regarding the merits of racial profiling concerns are important and should not be based on either anecdotal evidence or incomplete analysis. Evaluating the extent and nature of police profiling patterns may lead to decisions regarding proper training and appropriate police tactics. It is crucial that such evaluations rely on appropriate methodological approaches, objectively obtained data, and appropriate benchmarks or comparison guidelines.

The Justice Department's Office of Community Oriented Policing Services (COPS) asked The CNA Corporation (CNAC) to work with two police departments to address data collection and evaluation issues. This effort represents a follow-on to a previous CNAC study for the COPS Office, which produced the report, *How to Correctly Collect and Analyze Racial Profiling Data: Your Reputation Depends On It!* (available on the COPS web site).[1] The purpose of the new work was to apply the tools and methods recommended in the previous report by creating partnerships with two police departments and helping them use rigorous analytical methods that go beyond the standard practice. The approach relies on blending analytical research methods with operational police insights.

[1] *How to Correctly Collect and Analyze Racial Profiling Data,* 2002. Report available in hard copy from the COPS Office (800-421-6770) or at the COPS web site, http://www.cops.usdoj.gov/default.asp?Item=770

This pamphlet does not include any actual data or specific findings from the two departments. Instead, it describes the general approaches used, and illustrates them with sample templates of the analytical output. These templates represent examples of how to display and evaluate results from various methods of analysis. They were selected to demonstrate the range of analytical questions that can be addressed by these techniques.

The traditional (standard) approach to data analysis–and its limitations

The typical approach for using police administrative data to identify the presence of racial profiling practices is to compare the racial breakdown of vehicle stops to that of the city's population. Consider an example for a hypothetical city in which 33 percent of the people who are stopped belong to a minority race group and 67 percent belong to a majority race group. If the demographic data for the city show that the racial breakdown of the population is the same–33 percent minority and 67 percent non-minority–the standard conclusion that would be drawn from the fact of equal racial shares is that there is no racial profiling; all is well.

By using this approach, people are effectively allowing the demographic mix of a city's total population to act as an estimate of the demographic mix of the observed violator population. More specifically, they are assuming that members of each race group violate traffic laws at the same rates and that officers observe members of each race group committing these violations in the same proportions. Whether the first assumption holds has been a subject of much debate, and is virtually impossible to prove or disprove. The second assumption is generally not expected to hold because it is well understood that deployed police officers, drivers of different races, and problem traffic areas are all dispersed unevenly across a city's geography.

Finding the right comparison group

A more general characterization of the main conceptual problem with the standard approach is that city population data don't provide the right comparison groups. Figure 1 shows how members of a city's driving population–whether they are residents of the city or not–become part of the stopped and cited populations. The ideal approach to determining whether members of one race group are inappropriately stopped at higher rates than members of other groups would be (a) to compare the race breakdown of the total violator population with that of the observed violator population, (b) to compare the race breakdown of the observed violator population with that of the stopped population, and (c) to consider reasons why variations between the groups exist. Unfortunately, it is virtually impossible to determine the racial make-up of either the total violator population or the observed violator population, so researchers and police analysts have usually made comparisons between the race breakdowns of the stopped and the resident populations. The trouble is that the resident population is several steps removed from both violator populations and is likely to be quite different for reasons that are independent of police activity as well as for reasons that are due to police policies and behavior.

Reasons for differences in the racial compositions of the various populations in the figure can be considered in more detail. Starting from left to right in Figure 1, it can first be seen that the racial composition of the available driving population may be different from the racial composition of the city's population for two reasons. First, the city's own driving population may not have the same racial composition as the city's total population. This will be true if race breakdowns vary by driving age or car ownership. Second, if the non-resident driving population is sufficiently large, then differences between the racial compositions of non-resident drivers and the city's population will also cause the driving population to be racially different from the city's population. If the race mix of the driving population is different from that of the city's population for either of these reasons, then it is likely that all the populations that are "downstream" from the driving population will also be racially different from the city's population. Since these differences will occur independent of any action on the part of police officers, they cannot be the result of racial profiling.

Figure 1. How members of a city's total population become members of the city's driving and stopped populations

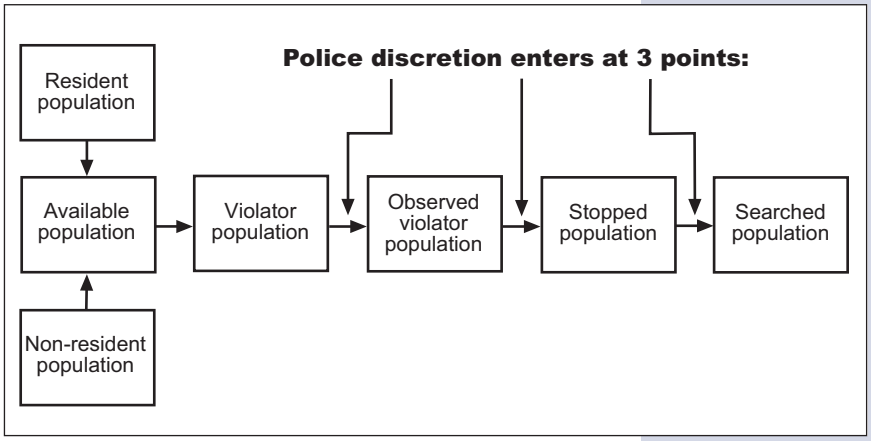

Next, consider sources of differences between the racial compositions of the driving population and the violator population. The main issue here is whether people of different races commit infractions at the same rates. If they do, then the race distributions of the two groups will match; if they don't then the race distributions will not match. Again, this difference will occur independent of police activity.

Moving from the violator population to the observed violator population is the first place where police department policies can contribute to differences between the stopped population and the city population. Specifically, the racial make-up of the observed violator population may differ from that of the total violator population due to variations in police deployments across geographic areas and times of day. Such differences are legitimate if deployment patterns are determined without respect to race. Next, differences in the race breakdowns of the observed violator and stopped populations can arise due to the interaction between racial differences in violation rates on the one hand, and police policies and procedures regarding which violations should be singled out for stops, on the other. Again, if such policies are determined without respect to race, then differences between populations can be considered legitimate. Therefore, an important corollary to recommendations on effectively using administrative data to show whether a department practices racial profiling is that departments need to be able to document the reasons behind existing deployment patterns and other operational decisions.

Finally, individual officers will use their own judgment and understanding of department policies to decide whom to stop, and then whom to cite or search among those they have stopped. Racial differences between the observed violator and stopped populations and racial differences between the stopped and cited (or searched) populations are the result of the behaviors of both the officer involved and the stopped citizen, therefore, they can be the result of racial profiling or not.

Following the flow from the city's total population to the stopped and searched or cited populations clearly shows the problems associated with using comparisons between the racial compositions of these groups to prove or disprove that a police department's officers are engaged in racial profiling. Therefore, such comparisons should be interpreted with extreme care.

Matching definitions of race across data sources

Another problem with making simple comparisons between the racial compositions of the stopped populations and racial compositions of cities' populations is the fact that measuring race is not always straightforward. The U.S. Census Bureau publishes self-reported race data by one race alone, two races, and three or more races. Census data also indicate self-reported Hispanic ethnicity. In contrast, police-collected information on a citizen's race typically reflects an officer's assessment of the individual's racial background based on how the person looks. Since it's difficult for police officers to capture the nuances of mixed races, stop or citation data usually just have one category for each race and don't capture multiple races. Therefore, in many cases, U.S. Census race definitions may not match race definitions in police data even when they have the same labels. In addition, in some instances, stop or citation data are recorded as Caucasian or African-American, omitting all other races and ethnic categories (e.g., Hispanic).

Controlling for many factors at one time

Finally, simplistically comparing race data for stops against race data for total populations can't capture multiple correlations between different factors associated with stops and violations. Therefore, a better approach is to use multivariate techniques that allow researchers to hold constant the effects of several factors all at the same time. For example, it is likely that more police officers are deployed and making stops in high-crime, low-income neighborhoods. Unfortunately, in the United States, such neighborhoods are also likely to be minority neighborhoods. Therefore, it is possible that disproportionately high stop rates for minorities may be the result of police deployment patterns rather than police bias.

One way to explore the combined impact of multiple factors, and investigate the direction of causality, is by the use of nested tables. Samples of these types of tables are provided later on (e.g., figures 3 and 4). For example, suppose that an investigator is examining stop rates by minority classification. Rather than just look at stop rates associated with various race/ethnic groups, it is important to look at stop rates for these groups within specific types of car stops (e.g., driving related and vehicle equipment), for different times of day (e.g., rush hour, daytime, evening), or within different neighborhoods (e.g., central city district, residential areas, shopping areas). Examining the basic measure of car stops in this way has the potential to provide substantially more information about the multiple influences that may affect car stops. The tables can be set up to explore these different influences by explicitly calculating rates for each pertinent cell defined (e.g., stop rate by race/ethnic group for driving related incidents at night in the central city).

In addition, other multivariate techniques (such as regression analysis) that include neighborhood-specific data on income, race, and the number of deployed police officers can measure the separate effects of each of these potentially correlated factors. One specific example relates to controlling for type of stop along with race. Research has shown that Black drivers are less likely than drivers of other races to wear seat belts.[2] Therefore, if a large percentage of stops were related to failure to wear a seat belt, Black drivers would be likely to have a disproportionately

[2] See: "Achieving a Credible Health and Safety Approach to Increasing Seat Belt Use Among African Americans," Department of Occupational and Preventive Medicine, Meharry Medical College, May 1999 and "Blue Ribbon Panel to Increase Seat Belt Use Among African Americans: A Report to the Nation," December 2000, p. 11, DOT HS 809 185.

high share of the total. Thus, it would be important to control for type of violation or reason for stop in a racial profiling study based on data from car stops. Regression analysis is well suited for this type of analytical investigation.

Another aspect of controlling for many factors is to consider which factors are likely to be important in a given jurisdiction. Different cities are likely to have different problems, geographic considerations, and histories with respect to citizen-police interactions over time. One approach is to survey the citizens of a jurisdiction to determine their perceptions (e.g., of the police force, crime, police initiatives) and to understand citizens' views concerning various issues (e.g., racial profiling, neighborhood crime concerns, traffic congestion, illegal parking). Using surveys is a productive way to determine the range of factors that are relevant for a particular city or jurisdiction.[3]

Data collection

What kind of data can be collected?

There are at least four general categories of data, as described on the next page. This should not be considered a definitive or exhaustive list of data elements. Depending on the questions that are considered to be important, and the particular characteristics of a department and the surrounding city or county, some of these measures may be irrelevant. There may also be data elements that would be important to consider that are not included.

Although this may appear to be a lot of information, most of it is already collected–often as part of standard practices. For example, officers making a traffic stop will usually record the year of the car and the date of stop. With this information, it is possible to create a variable to describe the age of the car. In addition, characteristics of the officer making the stop need not to be collected for each incident. Instead, with a coded officer identifier for each stop, stop data can be merged with data from personnel files. Finally, information about the stop location can be derived from U.S. Census data and supplementary operational data.

[3] See Deborah Weisel, Conducting Community Surveys: A Practical Guide for Law Enforcement Agencies, NCJ 178246, 1999, Washington, DC: U.S. Department of Justice, Office of Justice Programs. Weisel provides a practical guide for law enforcement agencies, which includes a description of how surveys can be applied to improve policing services, in addition to techniques for identifying survey goals and procedures for survey administration and analysis. The guide reviews many practical aspects of using surveys, including the use of telephone surveys, tailoring surveys to fit different needs, sample size, probability sampling, analytical techniques, and costs.

Categories of Data

Citizen characteristics

- Demographics
- Car age/type
- Address
- Behavior

Officer characteristics

- Demographics
- Career and assignment
- Behavior

Stop information

- Location
- Time of day
- Duration
- Reason
- Outcome

Site information

- Demographics
- Urban/rural, business/residential
- Socio-economic data
- Traffic flow
- Officer deployment
- Calls for service

Deciding what data to collect

Figure 2 provides a flow diagram of a process that can be used to guide data collection decisions. It is motivated by the most basic consideration–what do the department and the community want to learn? Or put another way, which questions are analysts attempting to answer? The ability to articulate what is to be investigated should precede the process of exploring data elements and selecting an analytical approach.

Following this process will allow departments to evaluate and categorize the data elements they are collecting in tandem with the methods available to them. One of the important aspects of this method is that it allows investigators to explicitly consider what questions can be answered using alternative approaches. For example, given a certain set of available data elements, there will be some questions that can be answered, but there are likely to also be some questions that can't be addressed by that data. The investigators need to know this and consider it in advance, so that they can determine whether it is important to obtain additional data, or alternatively so that they can explain that certain questions will not be answered using the current data elements.

Figure 2. Steps to Guide Data Collection

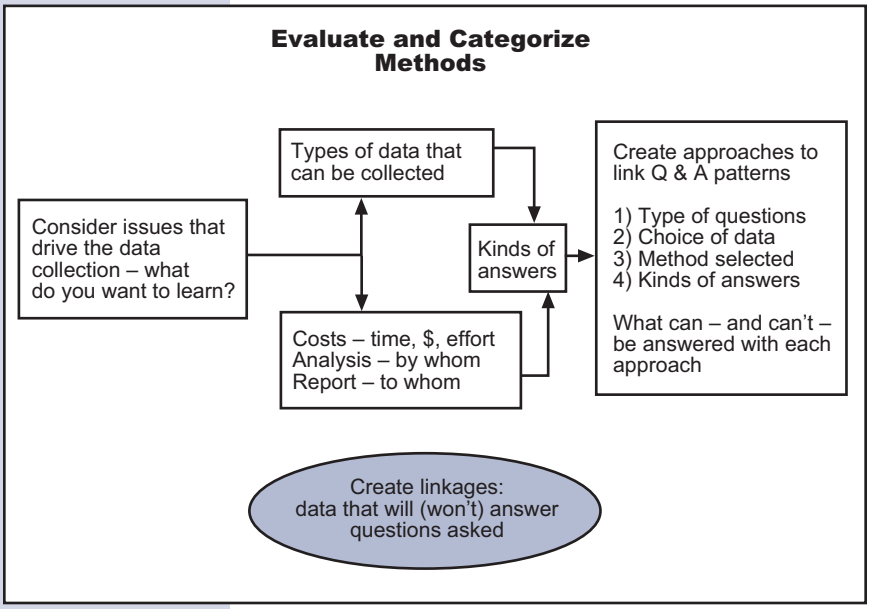

Depending on what analytical method is eventually selected, it may be advisable for police departments to consider the merits of a partnership with an analytically trained social scientist.

CNAC research partnerships results

General lessons learned

CNAC established research partnerships with two law enforcement agencies that differed substantially in terms of characteristics of the surrounding communities and the data being collected. In each case, there were separate challenges to evaluating the data; surprisingly, however, some lessons learned were the same for both departments in that similar patterns showed up in both data sets. Specifically, four findings held across both communities:

- Nonresident drivers affect outcomes.
- The race share of stops/citations varies significantly by time of day.
- Race and gender interact with type of stop/citation.
- The race mix of stops/citations varies with the race mix of the stop location's population.

General approach to interpreting data

To investigate the presence or absence of racial profiling within a force, data can be analyzed at three levels: department or citywide, by census tract, and by individual officer. Generally, our approach was to use the data to identify patterns in the race mix of stops and stop outcomes that might call for further investigation. More specifically, rather than using the standard technique of making simplistic comparisons between the race mix of stops and the race mix of the cities, we focused on within-sample comparisons across race groups, and used both bivariate and multivariate analysis techniques. This approach allowed us to avoid some of the pitfalls associated with the standard technique.

Although the focus of this project was on using data to analyze racial profiling issues, the same data can be used to help departments analyze operational efficiency and to learn more about what practices are more effective in achieving operational goals (and what practices are less effective). In the process of trying to understand why outliers exist, a department can learn about many aspects of its operations–not just racial profiling.

Sample standard templates to evaluate stop and citation data

In the following samples of standard templates, we present table structures that include a selection of important variables and a suggested structure of analysis. These are examples of how investigators can use templates, either in table or regression formats, to investigate stop and citation data. In addition, these approaches can be used to investigate associated search data, as well as officer efficiency metrics. In these tables, we do not present actual numbers–while we show the structure of the table, all numerical entries are represented by XX. We do not show actual numbers, in part because we do not want to imply that our examples support findings that might indicate the possibility of racial profiling, or appear to negate the existence of racial profiling.

If we were to use actual numbers in our tables, it would be tempting for readers to try to interpret the meaning of the results. This would be inappropriate for a number of reasons. First, we are not using full models – just template examples. In the real world, there might be other influences that should be considered, and these factors will vary for different jurisdictions. In addition, if we used numerical examples, we would need to describe how all of the variables are measured (their scale), discuss the variables' observed variations (standard deviations) and correlations with each other, and motivate a discussion of hypothesis testing and statistical significance. It is not our intent to address the overall topic of how to do analytical modeling or assess tests of significance–these can be complicated issues and are best addressed in the context of social science estimation or prediction techniques.

We begin our examples of data templates by considering ways to address stop data. The most general way to expand the traditional comparison of stop data by race is to create a table structure that supplies information about additional factors, such as gender, time of day, or type of stop. Figure 3 displays a template for disaggregating data by type of stop, controlling for minority status. In this example, stop data are evaluated separately for driving-related stops and vehicle equipment stops, although other types of stops could be considered as well. This type of analysis has the advantage of being very easy to display and understand, but has the disadvantage that only a limited number of factors can be displayed in a single table.

Figure 3. Disaggregate Data by Type of Stop

		Car stops		
Race	**Population**	**All**	**Driving related**	**Vehicle equipment**
Minority	XX.X%	XX.X%	XX.X%	XX.X%
Nonminority	XX.X%	XX.X%	XX.X%	XX.X%
Total	100.0%	100.0%	100.0%	100.0%

- Ways to disaggregate: by gender, time of stop, location, etc.
- This list of variables informs what data to collect.

When looking at aggregate stop data, operational experts and researchers alike begin to think about other ways to consider what the data might be telling them. For example, it may be that most officers are deployed in minority neighborhoods, or that minorities are more likely to live in urban neighborhoods with more traffic. These and a host of similar considerations can be used to motivate the collection of additional data elements that can be used to control for other influences and illuminate the data findings.

One of the most common considerations is the influence of the stop location, which can be mapped to a distinct neighborhood or census tract. Figure 4 offers a sample template of one way to examine data by stop location. Narrowing down to census tract generally yields a better idea of the characteristics of the neighborhood and, theoretically, gets closer to the real candidate population that is available to be stopped. However, it is still possible that many nonresidents may be stopped, depending on the nature of the neighborhood and the type of highways or streets that run through it.

Figure 4. Analyze Data by Stop Location–Neighborhood or Census Tract

Tract number	Number of stops	Minority share		Percent difference (%)
		Stops	Population	
XXXXX	XX	XX%	XX%	XX%
XXXXX	XX	XX%	XX%	XX%
XXXXX	XX	XX%	XX%	XX%
XXXXX	XX	XX%	XX%	XX%

- Is the difference between the stops' minority share and the population's minority share *statistically significant?*

This type of tabular result allows a researcher to compare the minority share of stops in a specific census tract with the minority share of the population in that census tract. If there are differences, two questions to ask are how large are the differences and would the community consider them to be meaningful? Again, most evaluators would also explore other census tract characteristics, such as the median income for the tract, major through streets, presence or absence of shopping centers, and so on.

The effort to control for many factors simultaneously will often lead to using regression analysis to investigate the issues. Multiple regressions have the advantage of being able to estimate the separate influences of several factors at once. This allows the independent effect of one factor to be considered while holding constant, or taking into consideration, the effect of a range of other factors. Figure 5 is a sample template of the results of a multiple regression to evaluate stop rates by census tract, holding constant several important factors like the minority share of the tract's population. (The variables or factors controlled for in figure 5 are illustrative of the estimation technique, but should not be considered to exhaust the possible list of variables.)

Figure 5. Explaining tract-specific stop rates as a function of tract characteristics

Observation unit = Census tract; Dependent variable = Stop rate			
Explanatory variable	**Coefficient**	**Relative risk ratio**	**Standard error**
Minority population share	XX	XX	XX
Unemployment rate	XX	XX	XX
No. of officers deployed	XX	XX	XX
Calls for service	XX	XX	XX
Average income	XX	XX	XX
Male population share	XX	XX	XX
Urban/rural	XX	XX	XX
Business/residential	XX	XX	XX

- Controlling for other tract characteristics, is the minority population share a significant determinant of stop activity?

This regression approach can offer compelling information because it is often argued that differences between the minority share of stops and the minority population share at the aggregate level are due to unequal distributions of officers across areas/neighborhoods in the city. If this is the case, evaluating at the census tract level should yield useful information regarding how stops differ across these areas and why. In particular, it is also often argued that the correlation between police presence in a neighborhood and the race mix of the neighborhood is really a result of the correlation between race and income/poverty status

in U.S. cities. If this is true, then it is important to include both income and race in the analysis (in our terminology, hold constant the separate effects of income and race).

Another way to use the information shown in figure 5 is to use the estimated model to generate predictions of the amount of stop activity that would be expected in each census tract, given the levels of the variables or factors observed in each tract. Figure 6 shows comparisons between predicted stop activity and actual stop activity by census tract. In this example, the expected rates are used as benchmarks for actual rates.

When the actual rates are substantially different from the predicted rates (i.e., when there are outliers), it may be useful to consider what supplementary information might explain the differences. For example, if actual stop rates are substantially lower than the predicted rates, it might be that red light cameras are being used to monitor intersections and that, consequently, officers are turning their efforts in other directions.

Figure 6. Use Predicted Stop Rates for Benchmark Comparisons

Tract number	Stop rate		Number of stops	
	Actual	Predicted	Actual	Predicted
XXXXX	XX%	XX%	XX	XX
XXXXX	XX%	XX%	XX	XX
XXXXX	XX%	XX%	XX	XX
XXXXX	XX%	XX%	XX	XX

- Are actual numbers of stops and stop rates *significantly* different from predicted?
- Can supplemental data explain differences?

Another way to disentangle stop data is to look for patterns in stop outcomes, such as conducting a search or issuing a citation. For example, to analyze search outcomes, the first step is to explore whether, given that the stop has occurred, minorities are more likely than Whites to be searched. In addition, it is interesting to compare the results of searches by race (and possibly by other factors as well).

Looking at the search results can indicate how effective, or productive, searches are across racial subcategories, and can also indicate what type of illegal product was most frequently discovered (e.g., weapons, drugs). A sample template for this type of analysis can be seen in figure 7.

Figure 7. Use Data on Search to Analyze Profiling as Well as Search Yields

	Stops		Search results		
Race	**All**	**Search**	**Nothing**	**Weapon**	**Drugs**
Minority	XX%	XX%	XX%	XX%	XX%
Nonminority	XX%	XX%	XX%	XX%	XX%
Total	100%	100%	100%	100%	XX%

- Conditional on being stopped, are minorities *significantly* more likely to be searched?
- Conditional on being searched, are minorities *significantly* more likely to have contraband?
- Do searches have high overall yield rates?

Some authors have focused especially on stop outcomes because data on stop outcomes are considered to be cleaner than data on stops alone. Specifically, since it is usually not possible to identify the observed violator population (see figure 1), analysts don't have any information about drivers who could have been stopped, but weren't. In contrast, since stop outcomes are conditional on the occurrence of a stop, it is possible to identify who experienced each outcome, and equally important, who did not. In this case, therefore, there is no ambiguity about having identified the denominator correctly.

Knowing which drivers experienced a specific outcome and which did not also enables researchers to use regression analysis to estimate the likelihood that a specific event will occur as a function of the drivers' characteristics and characteristics of the stop. In other words, regression techniques allow us to look at a number of different variables that may be influencing an outcome, and figure out which of the variables actually have an important and substantial impact.

Figure 8 shows a template for reporting the results from a multiple regression technique that estimates the separate impacts of various factors that influence whether a person who has been stopped is subsequently searched.[4] The citizen race variable in the list of independent (or influencing) variables is the one of interest with respect to racial profiling questions. In this case, the model estimates the effect of the driver's race on the likelihood of being searched, holding constant, or separately taking into account, the influence of other factors, such as citizen gender, car age, reason for stop, time of day, and various officer characteristics.

Figure 8. Control for Multiple Factors Using Regression Analysis of Search

Observation unit = Stop; Dependent variable = Search/no search			
Explanatory variable	**Coefficient**	**Relative risk ratio**	**Standard error**
Citizen race	XX	XX	XX
Citizen gender	XX	XX	XX
Car age	XX	XX	XX
Officer race	XX	XX	XX
Officer's tenure	XX	XX	XX
Officer's unit	XX	XX	XX
Reason for stop	XX	XX	XX
Time of stop (day/night)	XX	XX	XX

- Is citizen race significant even when controlling for other aspects of the stop?

[4] The statistical model used here is a logit model. A logit is a non-linear regression model that is frequently used in analyses when the outcome variable being modeled is a probability (such as the probability that a stopped cited driver was searched or not), rather than a continuous variable (such as the number of citations issued in a tract). Logit models are better than linear models for predicting probabilities because the predictions generated by a logit always fall between zero and one. Logit models are often used for dependent variables with two possible outcomes (denoted as 0 or 1 outcomes), but can also be used for more complex dependent variables. For technical notes on the logit model, see Peter Kennedy, A Guide to Econometrics, MIT Press, 5th edition, 2003; William H. Greene, Econometric Analysis, Prentice Hall, 5th edition, 2003; and David Knoke and Peter J. Burke, "Log Linear Models (Quantitative Applications in the Social Sciences), Sage Publications, 1980.

In figure 8, officer characteristics were included in the list of explanatory variables as a suggestion rather than as an example of typical control variables. Although most departments seem certain that the stopping officer's race does not influence the race of the stopped (or cited or searched) citizen, very few departments have actually put this to the test. However, it is important to consider using variables that control for officer characteristics.

For example, minority officers might do racial profiling of minorities, or they might do "reverse" profiling. The key is that different results on officer race and experience level may tell different stories. In popular terms, this notion is framed by asking whether profiling is a question of White vs. Black or a question of Blue vs. others. The primary reasons that officer data have not been included in analyses are concerns about privacy and liability. Therefore, the suggested method for including officer information is to use identification codes so that data can be matched to specific officers without identifying them by name.

Another way officer identification can be used to analyze racial profiling is to examine officers' stops and compare them with some benchmark. A benchmark can be externally provided, by taking a standard set by other departments, for example. Another approach is to create an internal, department-specific benchmark by calculating an average behavior among a department's own officers, and then comparing individual officer data with the average for the department, unit, branch, precinct, or neighborhood, depending on what exact behavior is being evaluated.

Using hypothetical data, Figure 9 shows an example of the type of analysis that can be performed. This example allows an evaluator to consider whether there are officers who stop significantly more (fewer) minorities than the unit's overall average. It is important to stress that, although this approach identifies officers who are performing above or below an expected average range, this does not necessarily imply whether they are doing anything wrong. There may be many factors that would easily explain why a particular officer had an unusually high or low minority share of stops over a given period of time, such as a month. However, looking at such a metric allows department managers to consider these factors, and determine whether any unusual patterns might merit additional investigation or consideration in the future.

Figure 9. Analyze Data at the Officer Level and Compare with Benchmark

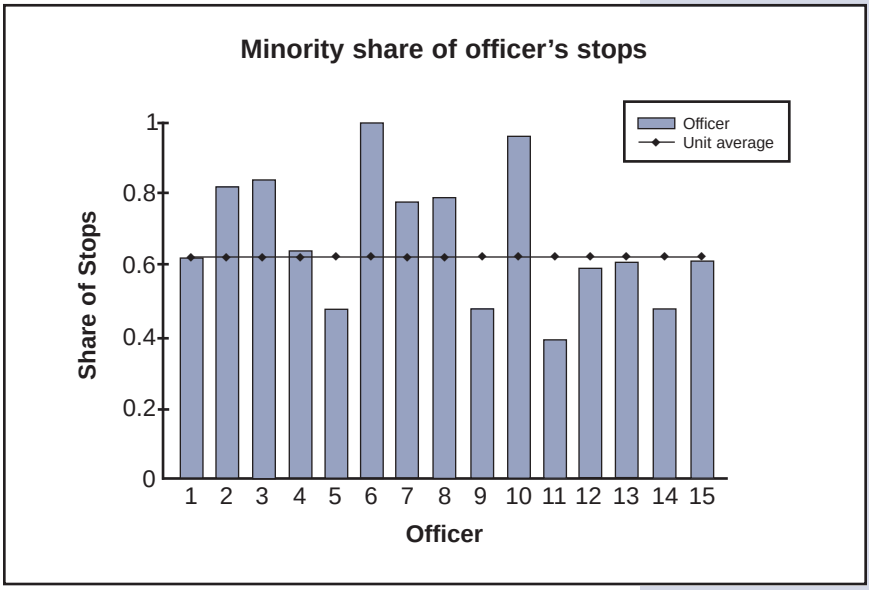

- Which officers stop significantly more minorities than the unit's average?
- Can supplemental data explain differences?

The key to this type of evaluation is picking the right benchmark to use as a comparison, and then identifying outliers and using other data to examine and explain them. For example, the race mix of stops for each officer is likely to depend on the area he or she patrols, and, within each area, on the time of day during which he or she is on duty.[5] Therefore, the benchmark should incorporate as many relevant factors as possible.

It is also possible to use officer identification codes to examine issues other than racial profiling. Figure 10 provides a very generalized example of the potential use of stop data to explore officer efficiency. For a group of hypothetical officers, the figure compares the rates at which each officer issued citations after making stops to the group's average citation rate. Thus, in this case the group average is the designated internal benchmark.

[5] Some areas or neighborhoods have very different nighttime and daytime populations.

Figure 10. Using Stop Data to Analyze Officer Efficiency

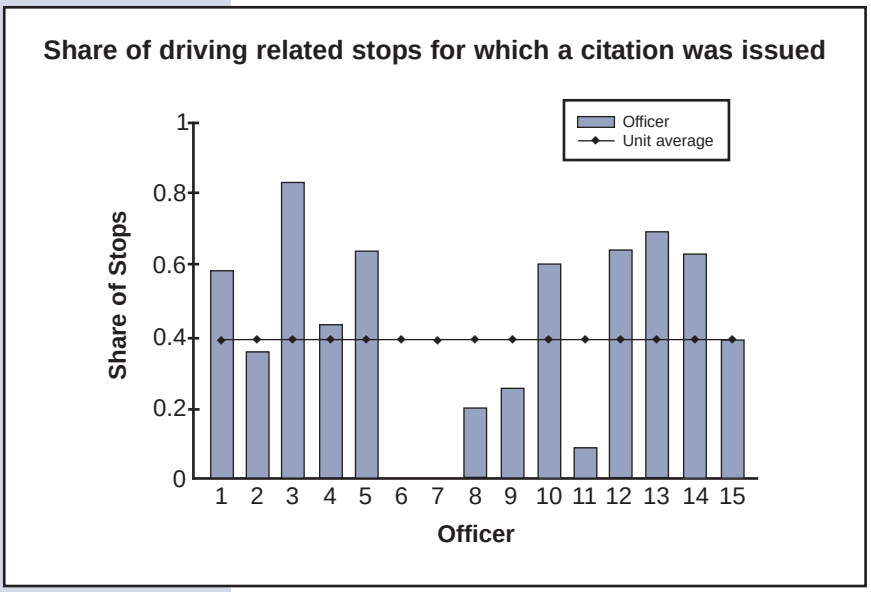

- Which officers give fewer citations than the unit average?
- What is the relevant benchmark? It should reflect department and unit goals.

Again, the key is to select the correct benchmark and then identify outliers and use other data to explain them. Officers who have unusual citation rates may or may not have excellent reasons that explain why their rates are unusually high or low. For example, different units may have different operational mandates. The traffic unit is goaled to monitor traffic and ensure traffic safety. So, officers in this unit are more likely to be looking for specific types of violations that are associated with danger and may be more likely to issue tickets. In contrast, officers in other units will make car stops and traffic stops more opportunistically, and potentially for different reasons. These and many other factors (e.g., time of day officer is on duty) must be taken into consideration when comparing across officers.

Summary

The methods of analysis shown in this report are only a sampling of the types of analysis that can be performed in the context of racial profiling or officer evaluations. As indicated before, the type of analysis that should be done depends on the nature of the department, the characteristics of the surrounding population, and the specific questions that the department and citizen groups may be posing. The questions being asked will determine what data should be collected and how it should be analyzed–both with respect to modeling technique and the level of detail that is maintained (e.g., stops, searches, officer identification). Generally speaking, multivariate analysis will require careful attention to properly control for the various factors that influence stop rates, citations, and searches. A productive way to approach this type of analysis is through a partnership between local police operational experts and social science researchers versed in applying analytical methods. This type of partnership can greatly enhance the processes of data collection, analysis, and interpretation.

Additional Resources

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

Fridell, Lorie, Robert Lunney, Drew Diamond and Bruce Kubu with Michael Scott and Colleen Laing. *Racially Biased Policing: A Principled Response*. Washington, DC: Police Executive Research Forum, 2001.

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

FOR MORE INFORMATION

U.S. Department of Justice
Office of Community Oriented Policing Services
1100 Vermont Avenue, N.W.
Washington D.C. 20530

To obtain details on COPS programs, call the
COPS Office Response Center at 800.421.6770
Visit the COPS internet address listed below.

Created March 2005 e03180002