

## **Is Racial Profiling a Nondiscriminatory Policing Strategy?**

*Anderson v. Cornejo (2004) and the Economics of Police Search Procedures*

### **Facts**

**S**haron Anderson and eighty-nine other plaintiffs contended that they were discriminated against at Chicago's O'Hare Airport between March 1996 and August 1999. All ninety plaintiffs were black women who were chosen for nonroutine searches after returning to the United States from foreign travel. They claimed that Customs officials subjected them to pat-downs, strip searches, X-ray inspections, or body-cavity searches because of racial and sexual bias and not due to any reasonable suspicion that they were violating the law. None of the searches yielded illegal contraband.

To support their claim of discrimination, the plaintiffs referred to data collected by the General Accounting Office (GAO). Using a national sample of airline passengers, the report presented evidence that black women were far more likely to be subjected to X-ray searches than were other groups. Black women had a 6.4 percent chance of being X-rayed, compared to .73 percent for white women, .53 percent for white men, and 4.6 percent for black men. Thus, black women were eight times more likely than white women, twelve times more likely than white men, and 40 percent more likely than black men to be X-rayed. According to the plaintiffs, these numbers appear to leave little doubt that black women were being discriminated against.

## Court's Decision

The main issue the court had to deal with is whether the managers at the airport should be liable for the actions of their inspectors. The court ultimately reached the following conclusion: "There may well have been race or sex discrimination at O'Hare Airport, but the managerial defendants are not liable on account of discrimination practiced by the line inspectors." While this decision is of great importance to the litigants, it is not very interesting from an economic perspective. What is interesting, however, is that the court, while not relying on the GAO data for its decision, nevertheless pointed to another group of statistics that the women thought supported their case but that the court viewed otherwise.

In addition to the X-ray search rates listed above, the court's opinion referred to the GAO report's strip-search *hit rates*, meaning the percentage of searches that successfully yielded contraband. The hit rate for black women was 27.6 percent, compared to 19.5 percent for white women, 25.1 percent for white men, and 61.6 percent for black men. Thus, concluded the plaintiffs, significantly higher search rates were not coupled with significantly higher success rates in terms of finding contraband, further supporting the claim of discrimination against black women.

The court, however, viewed the hit-rate data in a different light:

Data from the GAO's report do not imply that Customs officials are searching black women (or any other group) but not similarly-situated passengers in other groups. The report's outcome-by-group tables—we gave one example above, concerning the success rate of strip searches—show that Customs officials search black women with (on average) the same degree of suspicion that leads them to search white women or white men. A 27.6% success rate for a particular kind of border search is not to be sneezed at. It may imply that the Customs officials are conducting too few searches, not too many.

In interpreting the hit-rate data, the court explicitly referred to (at that time) state-of-the-art economic research which concluded that different search rates for different groups does not necessarily imply discrimination. Instead, it is different hit rates that suggest discriminatory practices on the part of the authorities, and the discrimination is found against the group with the *lowest* hit rate. Thus, using the strip-search hit-rate numbers from the GAO report, it is white women who are most discriminated

against, black women and white men are treated similarly, and black men who are least discriminated against. These seemingly unusual results require far more explanation.

### **Economic Analysis**

The court in *Anderson* argued that a distinction must be made between search rates and hit rates. To understand this distinction in a model of police search, let's begin with a fairly simple conceptual example that for the time being does not involve the controversial issue of race. Consider two groups of people—the old and the young. The first important assumption is that these two groups are fairly easy to distinguish just by looking at them. Perhaps by “old” we mean over the age of fifty, and by “young” we mean under the age of twenty-five (but not a teenager). The second assumption is that these two groups have different propensities to commit crime. Let's say the young are more likely to sell illicit drugs than the old. Assume 35 percent of the young sell drugs but only 15 percent of the old do.

Next, consider the motives of the police officers who are conducting the searches. These searches will involve motor vehicle stops, a very common form of police search. The police position themselves in a way that allows them to easily distinguish the age of the driver or passengers. Assume for the time being that these police officers are *completely unbiased*. In other words, they have no interest in searching either group due to any biases against that group. The goal of the police is to maximize the probability that a search will find contraband. The police want to fight crime, and that is all they want to do. By putting all of these assumptions together, how should the police proceed with motor vehicle stops?

Let's begin with the police treating each group equally by stopping one in every ten vehicles with a young driver, and one in every ten vehicles with an old driver. This leads to an equal search rate of 10 percent for each group. While this has the appearance of being “fair,” it is not efficient from the perspective of maximizing the probability of finding contraband. To demonstrate this point, let's think about the next (or marginal) search only.

If the police know that the young are more likely to be carrying drugs than the old, the next search will yield a 35 percent chance of finding contraband if the police search the young but only a 15 percent chance if they search the old (these are the hit rates). To maximize the probability of finding contraband, then, the police need to focus their next search on the

young. In a model of rational criminal behavior, as the police change the intensity of search, the criminals are predicted to change their propensity to commit crimes. As the police increase the search rate of the young, the young will now be less likely to carry contraband because they will be more likely to be caught. Likewise, as the police reduce the search rate of the old, the old will now be more likely to carry contraband because they will be less likely to be caught. Numerically, let's say the hit rate among the young drops from 35 percent to 30 percent but rises among the old from 15 percent to 20 percent. While these rates are now closer, they are not yet equal. It is still more likely for the marginal search to find contraband on the young than on the old.

Along the same lines, the police should further continue to increase their search rate of the young and reduce their search rate of the old. Only when the propensity to carry contraband is the same across both groups should the police stop adjusting their search rates. Notice, however, that because the police began with a 10 percent search rate of each group, and then increased the rate for the young and reduced the rate for the old, the final (or efficient) search rates can no longer be identical: *the young will be searched at a higher rate than the old*. What will be identical are the hit rates, that is, the probability of finding contraband in either group. Perhaps, in the end, the search rate will be 15 percent for the young and 5 percent for the old, and the hit rate will be 25 percent for each group.

One of the key implications of this model is that efficient police searches may require different groups to be searched at different rates. This result may be thought of as condoning police discriminatory practices, but that would be a *gross* misunderstanding. Instead, this model is designed to *identify* discrimination, not encourage it. The model demonstrates that different search rates do not mean that the police are biased. Remember, in our example above, the police are explicitly assumed to be unbiased. Discrimination is identified by different hit rates. More specifically, it is the group with the *lower* hit rate that is being discriminated against. Why is this so?

As the police intensify their search of one group, that group will have a lower propensity to commit crime. If the police are biased against the young, for example, they may want to stop young drivers specifically because they don't like the young and want to impose costs upon them. So, the hit rate for the young will fall below 25 percent because the decision to search is not based solely on the objective of finding contraband. The more intensely a group is searched, the less likely contraband will be found on them, and the more they are biased against. Thus, the model leads to a simple test to identify the existence of police bias: if discrimination exists, different groups will

have different hit rates. Thus, the existence of discrimination cannot *solely* be inferred from differences in search rates.

This model of optimal police search is applicable in a wide range of settings, but much of the economic research focuses on identifying one particular type of police bias—racial bias. There is a large, long-standing, and widely varied body of research across many academic disciplines on racism in the criminal justice system. Numerous studies have identified racial bias from all agents (police officers, prosecutors, juries, and judges) of the criminal justice system. The argument that the criminal justice system discriminates against blacks often begins with one striking fact—while blacks make up only 12 percent of the total population in the United States, they make up 40 percent of the prison population. That is, the black population in prison grossly overrepresents the black population in the whole country.

But what are we to make of this fact? Consider another striking fact—while men make up only 50 percent of the total population in the United States, they make up nearly 90 percent of the prison population. That is, the male population in prison grossly overrepresents the male population in the whole country. Does this suggest that men are discriminated against in the criminal justice system, or is there a simpler explanation, such as men are more likely than women to commit crimes? In terms of race, the succinct question becomes: does race predict criminal behavior? If so, race can be used as an easily observable characteristic that aids the police in performing efficient searches.

This thinking still leads to a difficult question: if the crime rate is higher among blacks than it is among whites, why is this so? The most likely explanation is not that race in and of itself predicts criminal behavior, but that other characteristics that are correlated with race do, such as income, education, employment, and so on. These other factors may be difficult for police officers to observe when considering whom to search, so race acts as an observable proxy.

And there is no shortage of theories (well beyond the scope of this book) among criminologists and other social scientists that try to account for racial differences in crime rates, with the common objective of recommending social reforms to help combat crime at its root causes. Nevertheless, while it may be extremely important to understand why different groups have different propensities to commit crime, an efficient police search strategy needs to consider these differences *regardless* of why they exist.

The seminal economic study (Knowles, Persico, and Todd 2001) on police searches and racial bias begins with a striking fact of its own: between January 1995 and January 1999, blacks made up 63 percent of motorists

searched for illegal drugs or other contraband on Interstate 95 by Maryland state police, but blacks made up only 18 percent of motorists on the road. But without examining the hit rates of these searches, little can be said at this point about whether the police in Maryland are using racial profiling because they are biased against blacks. The study makes it clear that it is not *differences* in the way blacks and whites are treated by the police that identifies racism, but in determining the *reasons* for the differences. In this case, one reason may be that race is a predictor of crime. Another reason may be that the police are racially biased. The study tries to distinguish between these two competing explanations (which the authors state more formally as distinguishing between *statistical* discrimination and *racial* discrimination).

The study uses data on all motor vehicle searches (a total of 1,590) on a stretch of Interstate 95 for the time period stated above. First, the study assumes that the objective of the police is to maximize the number of successful searches, or hits. While the search rate is 63 percent for blacks and only 29 percent for whites, the study finds that the hit rates are almost identical (34 percent for blacks and 32 percent for whites). This finding does not support the belief that the differential search rates for blacks and whites are discriminatory. However, the study finds slight evidence of discrimination (that is, relatively low hit rates) against Hispanics and white women, but because those groups each make up a small percentage of the total searches, the results are weak.

The study also considers a slightly different objective for the police—maximizing the number of hits that involve a “large” amount of drugs, or “hard” drugs. For example, when searching for drugs that constitute a felony charge under Maryland’s laws, the study finds the hit rate for blacks to be 13 percent and for whites to be 3 percent. This result implies that it is whites who face discrimination, as their lower hit rate suggests that they are being searched too intensely relative to the nondiscriminatory rate of search.

In all, the study does not find evidence that the differential search rates indicate a racial bias against blacks. A follow-up study (Sanga 2009), however, expands on the former study in two ways: it extends the sample period from 1995 to 2006; and it includes all police searches of motor vehicles in Maryland (a total of 19,000), not just those along Interstate 95. With this larger data set, the new study finds that hit rates are 38 percent for whites, 28 percent for blacks, and 8 percent for Hispanics. This suggests racial bias against blacks, and especially against Hispanics, as the lower hit

rates for these groups (relative to whites) can be explained by the police searching them too intensely.

Another way to empirically investigate the existence of bias associated with racial profiling is to distinguish *police officers* by race, as opposed to just distinguishing the victims of search by race. If the police are not racially biased when they search individuals, black and white police officers should approach racial profiling in similar ways. If there is an efficient unequal search rate for two groups of individuals, any police officer with a nonbiased objective, regardless of race, should be applying the same technique. But if the race of a police officer affects the rate of search, this is an indication of possible racial bias.

One study (Antonovics and Knight 2009), using data from the Boston Police Department from 2001 to 2003, finds that the race of the police officer does have an effect on search rates. If the race of the officer is different from the race of the motorist, there is a higher probability that the motorist will be searched. This holds for white officers searching black motorists, and for black officers searching white motorists.

Another study (Sanga 2014) takes a different approach to testing how police behavior is affected by the race of the officer. Instead of using search rates and hit rates, this study looks at differences in police officers' *stop rates*. Stopping a suspect is only the first step in a series of potential actions a police officer can undertake. Once stopped, will a suspect quickly be released? Will the suspect be cited for one or more violations, and if so, which ones? Will the suspect and the vehicle be searched, and if so, how intensely? Will the suspect be arrested? Thus, officer discretion plays a large role in the decision on how to proceed once a stop occurs. Furthermore, how a researcher interprets data from the complete stop process depends largely on how the police officer reports what is done, and this report is unlikely to accurately reflect the officer's motives for all actions taken and, perhaps even more important, all actions *not taken*. For these reasons, empirical studies that use differential hit rates to identify police bias are unlikely to be capturing all the potential ways for police bias to surface: "Such discretion should make us wary to base inferences of racial discrimination on stop outcomes like citation or arrest codes or even the results of a search, especially since there are legal ways for officers to manipulate these both for and against the suspect" (Sanga 2014, 408).

On the other hand, when focusing *solely* on the initial decision to stop a suspect, there is far less police discretion to be concerned about. Of course there will be some discretion, as the officer will have to decide when to stop

a suspect. This decision can depend on many factors, some that depend on the behavior of the suspect and some possibly relating to officer bias. In addition, the officer is required by law not only to report the stop but to report it accurately. Thus, whatever information a stop yields to the researcher should be fairly precise, not only because the law requires accuracy but because there isn't much information for an officer to manipulate. And this is one of the downsides of using stop-rate data—they really don't yield much information. The study overcomes this deficiency, however, by finding a clever way to use stop rates to determine if the race of the police officer has an impact on stop behavior.

To determine if officer race matters, it would be ideal to examine the stop behavior of two different police officers—one white and one black—who are placed in *identical* policing situations. Let's say this scenario can be realized, and it is observed that the white police officer stops more black suspects and fewer white suspects than does the black officer. We can conclude from this fact that the white officer discriminates against black suspects *relative* to the black officer. The stop-rate data can be used to identify relative discrimination in this scenario because we are explicitly assuming that the policing situations are identical. Therefore, any differences in officer behavior must be due to some type of bias. What the stop rates don't tell us, however, is precisely what *type* of bias.

For example, assume that both officers are completely racially unbiased. What can account for their different stop behavior? Perhaps the white officer has a sincere, but mistaken belief that blacks have a higher propensity to commit crime relative to whites. In his mind, then, efficient policing would require him to devote more resources to stopping blacks than whites. If the black officer doesn't have this mistaken belief, the white officer's behavior will look like discrimination relative to the black officer's behavior. In this particular example (which is only one of many possibilities), the white officer uses different information than the black officer does but, by assumption, is not racially biased. This doesn't mean that another example can't be discussed in which it is racial bias that leads to different stop rates; it just means that there is no obvious explanation as to why the rates are different.

To replicate the ideal of identical policing situations as best as possible using real-world data, the study attempts to control for the area an officer patrols. Using data from Oakland, California, from 2005 to 2010, the study can track the neighborhood, day, and time an officer is on patrol. For example, a white police officer is assigned to patrol a specific neighborhood on Thursday at 6 p.m., and then a black police officer is



assigned to patrol the same neighborhood the next Thursday at 6 p.m. This allows, at least somewhat, for stop-rate data to be compared across similar policing situations to determine if one officer discriminates relative to the other.

When examining officer behavior specifically *within* neighborhoods, the stop-rate test suggests that relative discrimination does occur and that the race of the officer matters. But there is a peculiarity in the results. In minority neighborhoods, the study finds that black officers discriminate in favor of minorities and against whites relative to white officers. In white neighborhoods, however, black officers discriminate in favor of whites and against minorities relative to white officers. In other words, the *same officers* discriminate in favor of their own race in some neighborhoods but against their own race in others. This leads the study to conclude that *where* a suspect is stopped may be more important than by whom a suspect is stopped. Whatever is causing officer race to matter in explaining differential stop rates, the lack of a consistent bias directed toward a specific race makes it difficult for the study to accept that these differences are motivated by racial bias.

Rather than examining differential search rates, hit rates, or stop rates, a completely different approach to detecting racial bias in the criminal justice system is the focus of another study (Tomic and Hakes 2008). When detaining a suspect to perform a search, police officers are attempting to gather information to determine if an arrest should be made. Therefore, a low level of evidentiary standard, such as an officer's intuition or suspicion, must be met to justify the search. When actually arresting a suspect, however, a greater degree of evidentiary standard must be met since there must be a sufficient amount of information to justify the arrest. Thus, if racial bias is factoring into an officer's decision to arrest a suspect, compared to arrests that are not motivated by bias, there may be a greater chance of future dismissal of the charges due to the higher evidentiary standard that must ultimately be met.

The study examines different categories of crimes that involve differences in the discretion an arresting officer can use at the time of arrest. For example, a fully developed murder investigation that eventually leads to the arrest of a suspect is likely to involve a very low level of officer discretion, if any, in making the arrest. In contrast, when an officer feels compelled to make an "on-scene" arrest, such as with a drug or weapons offense, a high level of officer discretion is used. If racial bias is present, it is predicted to be observed more often in the arrests that involve a high level of officer discretion. Indeed, this is precisely what the study finds.

Using data from approximately 58,000 felony cases in the United States from 1990 to 1998, the study finds that blacks face a higher dismissal rate compared to whites for arrests that involve a high level of officer discretion, but the dismissal rates are not different for arrests that do not involve a high level of discretion. Thus, police officers' arrest discretion is found to be racially motivated. At first blush, it may seem comforting that whatever bias is present at the time of arrest, it is corrected further down the line in the criminal justice system. But unlike a racially motivated vehicle stop for a traffic violation, which causes short-term annoyance and minor costs, an actual arrest involves a fair amount of administrative costs, in addition to whatever social stigma costs are incurred by the defendant.

While these various studies try to disentangle the reasons for differential search rates across racial groups, what is by far the most likely case is that racial profiling *simultaneously* involves efficient policing and racial bias. Policy makers, then, face the difficult task of determining if the benefits of racial profiling as an effective policing technique outweigh the costs associated with racially biased police behavior. Complicating this task, policy makers have to confront a strong public outcry against racial profiling. This has led to several jurisdictions enacting reforms in an attempt to reduce or eliminate the practice. And this leads to an interesting question: what is the impact on crime rates when the police are no longer allowed to use racial profiling?

In an attempt to answer this question, one study (Heaton 2010) takes advantage of a high-profile, controversial policing incident in New Jersey in the late 1990s that led to immediate reforms. The study examines whether these reforms had an unintended impact on criminal behavior. More specifically, did these reforms reduce the arrest rates of minorities, thus leading to an increase in minority crime rates?

In April 1998, two white New Jersey state troopers shot at four men, three blacks and one Hispanic, who were driving in a van on the New Jersey Turnpike. The troopers pulled the van over for speeding, and as one of them approached the van on foot, it began to slowly roll backward. When the van did not come to a complete stop, the troopers fired eleven rounds into it, wounding three of the occupants. The driver of the van claimed that the van rolled backward because he did not properly place it into park. Furthermore, he claimed that he was not exceeding the speed limit in the first place. The troopers claimed that the van was speeding, although they did not have a radar detector with them to accurately measure the extent of the speeding. None of the occupants of the van were armed, and no drugs were found. The troopers claimed that they were

threatened when the van continued to roll backward and that their actions were justified to protect themselves and to prevent the van from fleeing.

This particular incident led to public outrage over the practice of racial profiling by the New Jersey state police, as well as among police officers in general. Not only was a civil suit filed against the state of New Jersey, the two troopers involved in the shooting were indicted on attempted murder and assault charges. Although the criminal charges were ultimately dismissed, the state paid nearly \$13 million to settle the lawsuit. Most significant, however, from a public policy perspective, the incident directly led to an explicit change in state police policy concerning racial profiling and vehicle stops.

In response to the incident, the governor of New Jersey ordered the state Attorney General's Office to write a detailed report documenting the state police's racial profiling practices and to offer a number of reforms to greatly reduce the incidence of such practices. Among the many reforms offered, these two provide a good representation of the report's objectives:

[The report] recommends that the State Police enhance and modify their training programs to make certain that the policies regarding racial profiling and the disparate treatment of minorities proposed in this Interim report are understood by all State Troopers who are assigned to patrol, their supervisors, and dispatchers.

[The report] recommends the development of a legislative initiative to create new official misconduct offenses to deal specifically with the use of police authority to knowingly or purposely violate a citizen's civil rights.

In short, the shooting incident led to concrete policy changes governing the behavior of the New Jersey state police.

Racial profiling became, and continued to be, an extremely well-publicized issue in New Jersey in the aftermath of the turnpike shooting. Other incidents relating to the shooting (controversial public statements by public officials, charges brought against the troopers, officials being fired) kept the local, regional, and national media busy. Thus, the public was made well aware of the reforms that would now guide police behavior. In addition to the heightened public scrutiny, the reforms themselves could directly affect the behavior of state troopers since they would now face disciplinary measures if racial bias was believed to be present in the performance of their duties.

Examining this anticipated change in police behavior, the study finds that the reforms did lead to a fairly substantial reduction in the number of arrests (between 16 percent and 33 percent) of blacks relative to whites for motor vehicle thefts, a crime that is often detected through vehicle stops. The study also finds that minorities respond to changes in police procedures in a similar fashion to the general population—higher arrest rates lead to less crime, and lower arrest rates lead to more crime—and so the lower arrest rates led to higher rates of motor vehicle theft in minority areas. Because the New Jersey policing scandal led to substantial policy changes that were enacted reasonably quickly, the study attributes the change in arrest and crime rates largely to the sudden changes in police procedures concerning racial profiling. In short, the reduction in racial profiling as a policing technique led to an increase in the minority crime rate (in this case, motor vehicle thefts).

In all, there is nearly universal acceptance of the fact that racial profiling exists. Individuals of different races (as well as other demographic characteristics) face very different search rates by the authorities. The key contribution of economic reasoning, however, is to emphasize that different search rates, even substantially different search rates, are not *necessarily* an indication of racial bias. Some studies conclude that racial profiling is an efficient policing technique involving statistical discrimination that takes advantage of the fact that different races have different propensities to commit crime. Other studies conclude that racial profiling involves racial discrimination: either police officers (in general) are biased against specific races, or the race of a particular police officer determines the direction of the bias for that officer. As with other empirical debates we have discussed, there is no easy way to identify and evaluate the advantages and disadvantages of such a controversial policing technique. One thing is certain: even if racial profiling involves efficient statistical discrimination, any policing procedure that relies on a demographic such as race is bound to lead to a substantial public outcry, making the public policy implications of using such a procedure all the more difficult to assess.

## Discussion Questions

1. Social welfare analysis begins with the identification of the social costs and social benefits of whatever issue is at hand. What are the social costs of racial discrimination in the criminal justice system? Can you identify any social benefits?

2. If it is the case that racial profiling simultaneously involves efficient policing and racial bias, how would you design public policy to deal with the problem? Would you apply a consistent policy across other demographic characteristics such as gender, age, ethnicity, and citizenship?
3. Many states have laws prohibiting employers from using criminal background checks when considering hiring prospective employees. The main benefit of these laws is that they help released offenders find legitimate employment opportunities to discourage future criminal behavior. But one of the main costs of these laws that economists have identified is that they may encourage racial discrimination. Why might that be? How would you evaluate the costs and benefits of these laws?

## References

- Antonovics, K., and B. G. Knight. 2009. "A New Look at Racial Profiling: Evidence from the Boston Police Department." *Review of Economics and Statistics* 91:163–77.
- Heaton, P. 2010. "Understanding the Effects of Antiprofiling Policies." *Journal of Law and Economics* 53:29–64.
- Knowles, J., N. Persico, and P. Todd. 2001. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy* 109:203–29.
- Sanga, S. 2009. "Reconsidering Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy* 117:1155–59.
- . 2014. "Does Officer Race Matter?" *American Law and Economics Review* 16:403–32.
- Tomic, A., and J. K. Hakes. 2008. "Case Dismissed: Police Discretion and Racial Differences in Dismissals of Felony Charges." *American Law and Economics Review* 10:110–41.