

Duke Law Journal

VOLUME 54

MARCH 2005

NUMBER 5

ASSESSING THE COUNTERFACTUAL: THE EFFICACY OF DRUG INTERDICTION ABSENT RACIAL PROFILING

KATHERINE Y. BARNES[†]

ABSTRACT

This Article investigates the costs and benefits of racial profiling in the context of drug interdiction. I begin by reviewing the empirical economic and civil rights literature regarding the existence and rationality of racial profiling and then build an explicit model of a trooper's decision to search a stopped vehicle. Estimating the model using stop and search data from a portion of Interstate 95 in Maryland, I find that the Maryland State Police use the motorist's race as a factor in deciding which stopped vehicles to search. This result persists even after controlling for many other descriptive variables that impact the trooper's decision to search. I then introduce an additional model that controls for race's role in the search decision and estimates the counterfactual: the change in the amount of drugs

Copyright © 2005 by Katherine Y. Barnes.

[†] Associate Professor, Washington University School of Law; J.D., University of Michigan Law School, 2000; Ph.D., University of Minnesota, 2003. I would like to thank Sam Bagenstos, Lee Epstein, Samuel Gross, Pauline Kim, Richard Lempert, Andrew Martin, Robert Pollack, Margo Schlanger, Nancy Staudt, and seminar participants at Washington University School of Law and the Work, Families, and Public Policy Workshop at Washington University in St. Louis for helpful comments on earlier drafts of this Article. In addition, Deborah Jeon, Elliott Wolf, and Amy Cruice at the ACLU of Maryland and Mark Bowen, assistant attorney general of Maryland, were very helpful in allowing me access to the data. Finally, I would like to thank Cary Hawkins, Jeffrey Wax, and Ken Simpson for their able research assistance. The views expressed in this Article are my own and are in no way endorsed or approved by the ACLU of Maryland or the attorney general of Maryland.

the police would find if they ignored race as a factor in the search decision. Applying that model, I find that race is the strongest predictor of identifying drug traffickers, but that racial profiling comes at significant cost, as black motorists who are subject to search are also more likely to be innocent than their white counterparts.

TABLE OF CONTENTS

Introduction.....	1090
I. The Racial Profiling Debate	1095
A. Civil Rights Empiricists.....	1097
B. Economic Models of Racial Profiling.....	1100
II. The Decision to Search	1107
A. Model Specification	1110
1. Characteristics of the Driver	1110
2. Characteristics of the Vehicle.....	1110
3. Characteristics of the Encounter.....	1112
B. Results: The MSP's Search Profile.....	1113
III. Quantifying the Costs and Benefits of Profiles	1119
A. An Initial Look at the Data (and Why It Is Misleading) ..	1120
B. Benefits to Profiling: Contraband Seized.....	1124
C. Benefits to Profiling: Couriers Arrested	1130
D. Costs to Profiling: Innocent Motorists Searched.....	1133
E. The Final Analysis: Costs and Benefits Compared.....	1137
Conclusion	1140

INTRODUCTION

Although the term is less than twenty years old,¹ “racial profiling”—the act of investigating a particular racial group because of a belief that members of the group are more likely to commit certain crimes—has a long history in the United States. The practice began before the country did, in colonial times, when blacks were subject to greater policing because of a belief that they were more likely to commit crimes.² Perhaps the largest single act of racial

1. Jerome H. Skolnick & Abigail Caplovitz, *Guns, Drugs, and Profiling: Ways to Target Guns and Minimize Racial Profiling*, 43 ARIZ. L. REV. 413, 419 n.36 (2001).

2. RICHARD MIDDLETON, *COLONIAL AMERICA: A HISTORY, 1607–1760*, at 283–91 (1992).

profiling in this country, the Japanese American internment,³ occurred much more recently. During World War II (WWII), 120,000 individuals of Japanese ancestry—two-thirds of whom were American citizens—were forced from their homes because of the belief that they were more likely to commit espionage or sabotage against the United States.⁴ German and Italian Americans were also the object of racial profiling during WWII, but of a less severe variety: they had to register with the police, abide by curfews, and submit to loyalty interviews.⁵ As with Japanese Americans, people in those ethnic groups were considered more likely to commit acts of sabotage and espionage.⁶

The Japanese American internment serves as a cautionary tale about racial profiling after a direct attack on our nation. One hundred twenty thousand people were interned, at least in part, because of racist, degrading stereotypes. But it may be the case that Japanese Americans *were* more likely to engage in espionage and sabotage than other ethnic groups.⁷ Even if that were so, internment cannot be considered an appropriate policy response. Fruitful discussion of what *would* be an appropriate policy, however, requires the unpleasant acknowledgment that race or ethnicity may be associated, as a

3. See Exec. Order No. 9066, Feb. 19, 1942, 7 Fed. Reg. 1407 (Feb. 25, 1942) (authorizing the secretary of war to designate “military areas” from which “any and all persons may be excluded”); General DeWitt’s Public Proclamation No. 1, March 2, 1942, 7 Fed. Reg. 2320 (March 26, 1942) (designating military areas on the West Coast); *see also* Civilian Exclusion Order No. 34, May 3, 1942, 7 Fed. Reg. 3967 (May 28, 1942) (excluding persons of Japanese ancestry from Alameda County, California).

4. This belief has little support in historical fact. See Eric L. Muller, *Betrayal on Trial: Japanese-American “Treason” in World War II*, 82 N.C. L. REV. 1759, 1760–62 (2004) (discussing the extremely limited data on disloyalty among the Japanese and Japanese Americans interned).

5. See, e.g., Proclamation No. 2526, Dec. 8, 1941, 6 Fed. Reg. 6323 (Dec. 10, 1941) (referring to Germans); Proclamation No. 2527, Dec. 8, 1941, 6 Fed. Reg. 6324 (Dec. 10, 1941) (referring to Italians). These measures were taken under the auspices of Executive Order 9066, the same order that granted authority for the Japanese American internment. Indeed, on its face, Executive Order 9066 applied to all people, as it did not mention particular ethnicities and only referenced “enemy aliens” in passing. Exec. Order No. 9066, Feb. 19, 1942, 7 Fed. Reg. 1407 (Feb. 25, 1942).

6. A small number of German and Italian Americans, primarily recent immigrants, were also interned during the war. See *generally* Proclamation No. 2526 (German enemy aliens); Proclamation No. 2527 (Italian enemy aliens).

7. One might suggest, for example, that people of Chinese ancestry would be less likely to engage in espionage and sabotage for the Japanese. This is, of course, speculation, as was the assumption that the Japanese Americans interned in World War II were more likely to be saboteurs, but both of these examples simply suggest that we do not know what the empirical data would demonstrate.

descriptive matter, with some criminal activities.⁸ Whether the Japanese Americans were more likely to be spies or saboteurs, and, if so, how much more likely, bears on the best strategy for preventing such activities and need not reinforce stereotypes or cause society to succumb to irrational hysteria. If the internment incapacitated only a small number of potential saboteurs, there may have been more efficient ways of dealing with the threat. Without quantifying the costs and benefits of using ethnicity as a factor in deciding who the likely saboteurs were, one cannot separate empirical facts from stereotypes. One may, of course, argue that in such circumstances, constitutional rights trump the empirical situation and render it irrelevant. But when the United States has faced a grave threat, society has not treated constitutional rights as sacrosanct. Policymakers curtail rights in an effort to combat the perceived threat, and, as with the Japanese American internment, courts may well sanction the sacrifice of individual rights if they are convinced that the situation is grave enough.⁹ Understanding the empirical component is, therefore, critical to discussions of both policy and the law.

Few people would confuse the war on drugs with WWII, but both have been occasions for racial profiling. One would think that the disentangling of fact from stereotype would be easier in the context of drug interdiction than it was in the earlier era. Although the United States knew little about the relative prevalence of saboteurs among Japanese immigrants and Japanese Americans compared to other categories of Americans, there is a growing body of data on highway stops and searches that, theoretically, could illuminate racially differential rates of drug use and drug trafficking in the United States. Thus, empirical fact could be separated from baseless stereotype and racist speculation in deciding whether race is a reliable marker of the likelihood that a given driver is carrying drugs.

But even these data are not free of assumptions and stereotypes. Not only does society shape views—including stereotypes and

8. Recent scholarship acknowledges that some response may have been prudent, but that the extreme measures used were too drastic. Eric L. Muller, *Inference or Impact? Racial Profiling and the Internment's True Legacy*, 1 OHIO ST. J. CRIM. L. 103, 116–17 (2003).

9. See *Korematsu v. United States*, 323 U.S. 214, 220 (1944) (“[W]hen under conditions of modern warfare our shores are threatened by hostile forces, the power to protect must be commensurate with the threatened danger.”).

assumptions—but views also shape society and how it is described empirically. Look for drug couriers only among blacks and some will be found; however, no *white* drug couriers will be found within this group. Then, look at the people who have been arrested and jailed, and it is easy to see how the initial perceptions of criminality get reinforced. Over decades, a stereotype that blacks are more likely to engage in criminal activity can transform itself into large and statistically significant differences in arrest and incarceration rates for blacks and whites.

There is some data suggesting that American society has experienced this form of self-fulfilling racial bias in drug prosecutions. For example, black youth are seen as more likely than white youth to engage in drug use and violent behavior; arrest and incarceration rates reinforce this perception.¹⁰ Yet self-reported data¹¹ consistently reveals that except for homicide, black and white youth have similar propensities to commit violent crimes, and blacks are, if anything, less likely than whites to use illegal drugs.¹²

Because biases can be reinforced through action, an initial focus on blacks as drug couriers will be validated and reinforced by the data: police will believe that their profile is accurate because they arrest more blacks for drug possession. Put simply, the data on stop and search rates and the data on arrest and incarceration rates alone cannot demonstrate that blacks are more likely to be guilty of drug-related offenses. Without a determination of this empirical issue, one cannot separate fact from stereotype in the debate surrounding racial profiling as a tool of drug interdiction.

10. In 2001, blacks accounted for 34.5 percent of the drug violation arrests and 44.5 percent of the prison and jail population, but only 13.2 percent of the general population. FED. BUREAU OF INVESTIGATION, CRIME IN THE UNITED STATES 252 t.43 (2001), *available at* <http://www.fbi.gov/ucr/01cius.htm>; ALLEN J. BECK ET AL., PRISON AND JAIL INMATES AT MIDYEAR 2001, U.S. DEP'T OF JUSTICE, BUREAU OF JUSTICE STATISTICS BULLETIN 12 t.14 (April 2002), *available at* <http://www.ojp.usdoj.gov/bjs/pub/pdf/pjim01.pdf>; U.S. CENSUS BUREAU, POPULATION ESTIMATES 1, *available at* <http://www.census.gov/popest/national/asrh/NC-EST2003/NC-EST2003-03.pdf>.

11. Self-reported data is based upon a random sample of individuals, rather than a biased selection method, and therefore is not systematically biased to overrepresent black and Hispanic drug use.

12. In 2001, 6.8 percent of whites and 6.9 percent of blacks surveyed reported using illegal drugs in the last month. U.S. DEP'T OF HEALTH & HUMAN SERVS., OFFICE OF APPLIED STUDIES, 2001 NATIONAL SURVEY ON DRUG USE & HEALTH: DETAILED TABLES fig. 2.12 (on file with the *Duke Law Journal*).

This Article seeks to separate fact from stereotype by estimating empirically the extent to which the stereotypes used to justify racial profiling are based on actual propensities to commit drug crimes. In particular, the Article investigates the counterfactual: what would police find if they did not use race in deciding whom to search during highway traffic stops? In doing so, I seek to quantify the potential benefits and costs of using race as a factor in the decision to search a stopped car. What would society lose, in terms of the amount of drugs seized, if the police did not engage in racial profiling? What would society gain, in terms of the number of innocent motorists who would not be searched, if racial profiling was not used?

Before answering these questions, however, it will be useful to articulate a more precise definition of racial profiling. As I use the phrase, racial profiling occurs when race or ethnicity is used as a factor in deciding whom to subject to more intrusive police contact (generally a stop, search, or arrest), excluding those cases in which the police are looking for a specific suspect whose race is known. Racial profiling involves the belief—be it racist, rational, or a combination of the two—that certain crimes are committed disproportionately by a particular race. Thus, stopping and searching a car speeding on the highway because the driver is black and the officer believes (or has been told) that blacks are more likely to be drug couriers is racial profiling. Stopping and searching a car speeding on the highway because the driver is black and the officer has been told that a black man sped away from a drug bust is not racial profiling, although it may still be inappropriate, depending on the level of individualized suspicion.¹³

In Part I, I briefly review the racial profiling debate. In particular, I explore how the empirical literature has failed to model

13. One extreme example of an inappropriate police response to a witness description involving race is found in *Brown v. City of Oneonta*, 221 F.3d 329 (2d Cir. 1999), in which the Second Circuit upheld the dismissal of Equal Protection claims against the city. *Id.* at 334. In *Brown*, the victim, robbed at knifepoint, described her attacker as a young black man and said that his right hand had been cut during the robbery. *Id.* The police responded by stopping over two hundred black individuals and asking to see the individual's right hand. *Id.* The City of Oneonta had fewer than three hundred black residents (not all young men), and approximately one hundred fifty black college students (also not all men) residing within its borders. *Id.* While deeply troubling, the police department's use of race in this case was not racial profiling as I define it. For an argument that removing witness identification from the definition of racial profiling masks the racial justice issues at stake, see generally R. Richard Banks, *Race-Based Suspect Selection and Colorblind Equal Protection Doctrine and Discourse*, 48 UCLA L. REV. 1075 (2001).

the decision to search explicitly or to quantify the costs and benefits of racial profiling, although these models are critically relevant to the racial profiling debate. Then, in Part II, I introduce a model that determines whether the Maryland State Police (MSP) engaged in racial profiling on Interstate 95 (I-95) between 1997 and 2003. Specifically, I build a model of a state trooper's decision to search a stopped vehicle and determine whether race is a salient factor in that model, after controlling for other factors that may influence that decision. Based upon the results of this model, I conclude that the MSP did engage in racial profiling on I-95. In Part III, I seek to quantify the effects of this racial profiling. I provide estimates of different drug interdiction outcomes, based upon different profiling strategies. Because the data on searches represent what the troopers found when they searched those vehicles they believed were most likely to contain drugs, rather than a random sample of stopped vehicles, I control for this selection criteria in the model I estimate. Based upon these models, I suggest several possible profiles for the MSP to use in its drug interdiction efforts and compare their costs and benefits.

I. THE RACIAL PROFILING DEBATE

The term “racial profiling” entered the mainstream media in 1987.¹⁴ At that time, the term referred to the police's use of pretextual traffic stops, focused on blacks and Hispanics, as a method of drug interdiction. Colloquially, this was known as the offense of “driving while black.” The racial profiling debate focused on the facts—whether racial profiling occurred—rather than the normative questions behind the facts—whether racial profiling was constitutional, and, if so, whether it should be allowed. By early 2001, most commentators, including then-U.S. Attorney General John Ashcroft, had denounced racial profiling.¹⁵ Civil rights advocates argued that the available data suggested that racial profiling was rampant and was, at best, lazy policing that was counterproductive to

14. Skolnick & Caplovitz, *supra* note 1, at 419 n.36.

15. *Confirmation Hearing on the Nomination of John Ashcroft to Be Attorney General of the United States: Hearing Before the Senate Comm. on the Judiciary*, 107th Cong. 492 (2001) (response of Senator Ashcroft, stating that racial profiling is “wrong and unconstitutional no matter what the context”).

drug interdiction goals;¹⁶ conservative commentators countered that the data showed that police used variables correlated with both race and drug possession in deciding whom to stop and search, and that no invidious discrimination was evident.¹⁷ Finally, economists had a different perspective on the racial profiling issue: to them, the question was not whether racial profiling occurred, but whether it was a rational use of limited police resources.¹⁸

The debate shifted significantly after the terrorist attacks of September 11, 2001. One could not help but notice that all nineteen hijackers were young Arab men; it seemed foolish to ignore this information.¹⁹ Those commentators who previously argued about whether racial profiling was happening were now arguing over the conditions under which racial profiling was acceptable, either

16. See, e.g., DAVID A. HARRIS, *PROFILES IN INJUSTICE: WHY RACIAL PROFILING CANNOT WORK* 78–90 (2002) (citing hit rate statistics to demonstrate that “[r]acial profiling is neither an efficient nor an effective tool for fighting crime”); William H. Buckman & John Lamberth, *Challenging Racial Profiles: Attacking Jim Crow on the Interstate*, 10 TEMP. POL. & CIV. RTS. L. REV. 387, 388 (2001) (noting that “as few as one in thirty stops net contraband even as small as a single joint” and that “[t]his figure could be achieved in random stops of all travelers”); Deborah A. Ramirez et al., *Defining Racial Profiling in a Post-September 11 World*, 40 AM. CRIM. L. REV. 1195, 1197–1201 (2003) (reviewing empirical data).

17. See, e.g., Heather Mac Donald, *The Racial Profiling Myth Debunked*, CITY J., Spring 2002, at 63–64 (“It turns out that the police stop blacks more for speeding because they speed more. Race has nothing to do with it.”), available at http://www.city-journal.org/html/eon_3_27_02hm.html; Heather Mac Donald, *The Myth of Racial Profiling*, CITY J., Spring 2001, at 26 (“Arlington has a 10 percent black population, but robbery victims identify nearly 70 percent of their assailants as black [and a]s long as those numbers remain unchanged, police statistics will also look disproportionate.”), available at http://www.city-journal.org/html/11_2_the_myth.html; John Derbyshire, *In Defense of Racial Profiling: Where Is Our Common Sense?*, NAT’L REV. ONLINE, Feb. 19, 2001, at <http://www.nationalreview.com/19feb01/derbyshire021901.shtml> (“So long as race is only one factor in a generalized approach to the questioning of suspects, it may be considered.”).

18. See, e.g., John Knowles, Nicola Persico, & Petra Todd, *Racial Bias in Motor Vehicle Searches: Theory and Evidence*, 109 J. POL. ECON. 203, 205 (2001) [hereinafter KPT] (proposing a test to distinguish between racially prejudiced searches and searches that use race as a factor only because it helps to “maximize successful searches”). For a more detailed discussion of the economics literature regarding racial profiling, see *infra* Part I.B.

19. There is, of course, a huge leap from the descriptive statement “the terrorists were young Arab men” and the speculative statement “young Arab men are more likely to be terrorists.” Two points should be made. First, not all terrorists are young Arab men. As has been noted elsewhere, the second deadliest terrorist attack on U.S. soil in recent history was committed by Timothy McVeigh, a European American. Samuel R. Gross & Katherine Y. Barnes, *Road Work: Racial Profiling and Drug Interdiction on the Highway*, 101 MICH. L. REV. 651, 749 (2002). Second, the percentage of terrorists who are Arab men is likely quite different from the percentage of Arab men who are terrorists. The former may be reasonably large—again, a supposition—and may well be a larger percentage than 0.25 percent, the percentage of the U.S. population constituting Arab men, but the latter is almost certainly quite small.

constitutionally or as a policy matter.²⁰ The shift in the debate was certainly constructive because discussion of the normative question was missing from the initial debate. The evolving debate should not, however, lose sight of the factual questions. The first empirical question is whether the police engage in racial profiling. The second and more important question is whether the empirical assumption underlying racial profiling—that a particular minority group is more likely to engage in a given criminal activity—is accurate and, if so, what policy implications follow.

The empirical literature on racial profiling has grown dramatically in the past several years as many jurisdictions have begun collecting data. As Professor Bernard Harcourt notes in his recent article, the literature is split into two separate strains.²¹ Legal scholars focus on the racial disparity between blacks and whites in stop, search, and find rates on the highway in an attempt to demonstrate that the police engage in racial profiling. Economists, in contrast, focus on the efficiency of police practices, using variations on a theoretical model of incentives for both the police and the individuals carrying drugs on the highway. Neither group of scholars, however, quantifies the benefits of racial profiling in terms of the additional drugs seized or drug traffickers arrested, nor does either group quantify the cost of racial profiling in terms of the number of innocent motorists searched. The goal of this Article is to fill this gap in the debate.

Before moving to my empirical analysis, however, it is useful to describe the current empirical literature more fully.

A. *Civil Rights Empiricists*

Civil rights empiricists focus on the threshold racial profiling question: do the police engage in racial profiling? Civil rights

20. See Samuel R. Gross & Debra Livingston, *Racial Profiling Under Attack*, 102 COLUM. L. REV. 1413, 1414 & n.3 (2002) (gathering examples of the post-September 11 discussion of the benefits of racial profiling).

21. Bernard E. Harcourt, *Rethinking Racial Profiling: A Critique of the Economics, Civil Liberties, and Constitutional Literature, and of Criminal Profiling More Generally*, 71 U. CHI. L. REV. 1275, 1276–77 (2004). While a generalization, this dichotomy provides a useful outline of the recent empirical literature on racial profiling. Perhaps the exception to my dichotomy is Professor Harcourt's recent article itself, which provides a theoretical economic model of racial profiling. See *id.* at 1291–95, 1315–22 (dividing empirical racial profiling scholarship into “economics” and “civil liberties” strains). Although I do not use the same terminology as Professor Harcourt, the categories are essentially the same.

empiricists do not explicitly model the differing incentives of police and drivers. Nor do they quantify any potential benefit from racial profiling; unlike economists, they assume, for the most part, that any benefit from racial profiling is negligible at best.

Professor David Harris epitomizes the approach of civil rights scholars.²² In his book, *Profiles in Injustice*, Harris details aggregate hit rates²³ in several jurisdictions, arguing that equal hit rates in the face of disproportionate search rates demonstrates that the police are using race in their decisionmaking.²⁴ Equal hit rates, according to Harris, demonstrate that blacks and whites offend at equal rates; a disparity in search rates, therefore, must be the result of racial animus.²⁵ Several commentators, myself included, have criticized Harris' reliance on hit rates as the sole evidence of racial profiling;²⁶ hit rates alone cannot prove or disprove whether the police engage in racial profiling.²⁷ Because the police do not choose whom to search randomly, their selection criteria create a bias in hit rates; it is impossible to tell what the underlying rates of actual guilt are without information about the individuals who are not searched.²⁸

In addition, from the policymaking perspective, a focus on simple hit rates masks the importance of other potential drug interdiction goals, such as, for example, focusing on drug traffickers rather than drug possessors, or maximizing the amount of drugs seized. More importantly, however, the analyses of Harris and others do not control for nonracial explanations of the observed racial disparity.²⁹ If,

22. John Lamberth's multiple studies of racial profiling also fall within this category and have the same methodological flaws. For a list of his many studies, most commissioned by police departments themselves, see <http://www.lamberthconsulting.com/about-racial-profiling/research-articles.asp>.

23. The term "hit rates" is shorthand for the proportion of searched vehicles in which the police found contraband to the total number of searched vehicles.

24. HARRIS, *supra* note 16, at 78–82.

25. *Id.* at 80. Professor Harris also argues that racial profiling makes for bad policing—that, as economists would say, it is irrational and inefficient—although he does not engage the theoretical possibility that racial profiling could be rational. *Id.* at 79.

26. Gross & Barnes, *supra* note 19, at 690–91 (describing a scenario in which equal hit rates would not indicate racial profiling); Harcourt, *supra* note 21, at 1316–17 (criticizing Harris' focus on hit rates).

27. For an extended discussion of this proposition, see Gross & Barnes, *supra* note 19, at 690–91.

28. See *infra* Part III.A and Figure 2 for a graphical example of how selection bias works.

29. See Buckman & Lamberth, *supra* note 16, at 394 (providing a blueprint for challenging racial profiling in court that relies solely on racial comparisons); Gross & Barnes, *supra* note 19,

for example, the police search young men in luxury cars at a higher rate, this may result in a statistically significant racial disparity in search rates if blacks are overrepresented in the population of young men in luxury cars driving on the highway, because the analyses do not control for alternative explanations of the data, such as searching young men in luxury cars.

In a joint project with Professor Samuel Gross, I have also investigated whether the MSP engage in racial profiling.³⁰ While recognizing that hit rates alone cannot prove discrimination, we argue that hit rates, combined with some information about underlying base rates of criminal activity, demonstrate that the MSP engage in racial profiling.³¹ While we investigate separately several possible explanations for the racial disparity in search rates, we do not model explicitly the decision to search a vehicle. Because of this modeling decision, we cannot differentiate between the portion of the disparity explained by racial profiling and the portion of the disparity explained by profiling on nonracial characteristics that are correlated with race. Thus, we do not quantify the extent of racial profiling. Rather, we conclude that the disparity is large enough that race must be a factor, because other variables could not explain why 60 percent of the motorists searched on I-95 were black, when blacks constituted only 17 percent of the drivers.³² Without a systematic model of the decision to search, however, we cannot simultaneously control for different possible factors that contribute to the large disparity and therefore cannot quantify the portion of the racial disparity explained by racial profiling.

In sum, although the civil rights empiricists focus on one important aspect of racial profiling—whether it occurs—they neglect to model the complex process by which police make search decisions. Because of this, civil rights empiricists provide too simple a portrayal of the process in general, and racial profiling in particular.

at 694 (asserting that “[t]he most likely explanation for this [racial] disparity is the obvious one: Maryland State troopers took race into account in deciding who to stop”).

30. Gross & Barnes, *supra* note 19, at 692.

31. *Id.* at 694. We assume that similar national rates of drug use across races imply similar base rates of drug possession on I-95. *Id.* at 691.

32. *Id.* at 692.

B. *Economic Models of Racial Profiling*

In contrast to the civil rights empiricists who focus on the preliminary existence question, the economics literature on racial profiling focuses on whether racial profiling is rational (efficient) or irrational (inefficient).³³ Professors John Knowles, Nicola Persico, and Petra Todd (KPT) provide the leading example of this type of empirical analysis.³⁴ KPT create a theoretical economic model of “hit” rates—the percentage of searched individuals found with contraband—given the cost of searching. Drawing on the economic literature of discrimination in labor markets, KPT model racial animus as a “taste” for discrimination; specifically, racist police officers have a lower search cost for black motorists than for white ones.³⁵ KPT demonstrate that, under the assumptions of their model, equal hit rates are consistent with rational discrimination rather than racial animus—or, to put it more forcefully, equal hit rates imply that any difference in search rates is justified by rational factors, whether based on race or not.³⁶ They estimate the model using individual-level data on the searches conducted by the MSP from January 1995

33. One might think of economists’ focus on rationality as a determination of whether racial profiling has the limited benefit of being efficient, rather than based solely on racial animus.

34. KPT, *supra* note 18, at 227–28.

35. *Id.* at 205.

36. *Id.* at 228. KPT’s model is flexible enough to allow for different metrics of police success. KPT focus on hit rates—the percentage of searches that find contraband—but also estimate the role of racial animus in police behavior assuming that the police want to maximize the percentage of searches that find not just any quantity of drugs, but a quantity sufficient to constitute a felony. *Id.* at 224–26. In contrast to their hit rates model, KPT find that felony hit rates for whites and Hispanics are significantly lower than the black felony hit rate, suggesting that rational police behavior would be to search blacks more often. *Id.* at 226. However, KPT discount this conclusion based on felony hit rates, because it assumes police gain nothing from finding a small amount of drugs. *Id.*

KPT also investigate the profiling of several other groups, including men and individuals driving luxury cars, and find no evidence of discrimination for these groups. *Id.* at 222–23. Economist David Bjerck criticizes KPT’s approach, noting that the model requires troopers to oversample black motorists even though, in equilibrium, black and white motorists have the same hit rates. David Bjerck, *Racial Profiling, Statistical Discrimination, and the Effect of a Colorblind Policy on the Crime Rate* 14–15 (Mar. 25, 2005) (working paper), available at <http://socserv.socsci.mcmaster.ca/bjerck/papers.html>. According to Bjerck, KPT’s model is directly contradictory to a straightforward definition of racial unbiasedness, where troopers search blacks and whites in proportion to their relative likelihood to carry drugs. *Id.* at 16–17. This criticism, however, ignores both the dynamic equilibrium KPT’s model presents, and the difficulty of accurately assessing the likelihood of carrying drugs for blacks and whites.

through January 1999³⁷ and find that blacks and whites have statistically indistinguishable hit rates.³⁸ KPT conclude that the police do not engage in discrimination based upon racial animus—that is, the police behave as if they incur the same cost to searching blacks and searching whites.³⁹ The advantage to KPT’s approach to identifying racial prejudice in police search behavior is that it does not require KPT to control for other potential explanatory variables: the model suggests that an aggregate difference in hit rates for blacks and whites determines whether racial animus is present. This is also a disadvantage, however, because it translates into unrealistic predictions. Specifically, the model implies that all groups of individuals, regardless of other characteristics, are equally likely to carry drugs on the highway.⁴⁰ If, for example, black motorists chose to carry drugs more often than whites, the police would respond by increasing their level of policing blacks, and black motorists would respond to the increase by lowering their frequency of carrying drugs. In essence, the model predicts that any criminal profile is useless in increasing the hit rate of searches.⁴¹ Additionally, in differentiating between racial animus and rational discrimination, KPT’s model does not answer the threshold legal question: do the police engage in racial profiling? Instead, KPT seek only to differentiate between types of racial profiling—the “good” kind, which is rational and based on blacks having an a priori higher payoff to carrying drugs, and the “bad” kind, which is irrational, inefficient, and solely based on racial animus.⁴²

37. KPT, *supra* note 18, at 216. This is an earlier version of the search data I use in Parts II and III, although KPT do not use the stop data at all.

38. *Id.* at 228.

39. *Id.* at 219–22. KPT’s analysis does find that Hispanics have a lower hit rate than whites, which they conclude is evidence of racial animus against Hispanics. *Id.* at 222.

40. *Id.* This prediction is a consequence of KPT’s model, in which each driver faces the same cost/benefit tradeoff of being caught with drugs versus choosing not to carry drugs. *Id.* at 211.

41. Police use of profiles, however, is still helpful in deterring criminal behavior.

42. *Id.* at 205. Economists Kate Antonovics and Brian Knight suggest that this is the legally salient question. Kate L. Antonovics & Brian G. Knight, *A New Look at Racial Profiling: Evidence from the Boston Police Department* 2–3 (Nat’l Bureau of Econ. Research, Working Paper No. 10,634, 2004), available at <http://papers.nber.org/papers/W10634>. On a practical level, Antonovics and Knight may simply be suggesting that causation is difficult to prove in racial profiling cases; a rational use of race might, in fact, be a rational use of some unobserved characteristic of the motorists that is simply correlated with race. Nonetheless, whether the police engage in racial profiling is an important threshold question, as it determines what level of scrutiny is applied in an Equal Protection challenge. See *Grutter v. Bollinger*, 539 U.S. 306,

The KPT model and its extensions⁴³ assume that all troopers act in the same way.⁴⁴ Professors Kate Antonovics and Brian Knight and Professors Shamena Anwar and Hanming Fang,⁴⁵ in contrast, build models that explicitly rely on heterogeneous troopers. These authors provide a clever strategy, based on the race of the troopers, to identify whether racial animus motivates the troopers. They base their statistical tests on the simple premise that if discrimination is purely rational, all troopers, regardless of the *trooper's* race, would search black motorists at the same rate.⁴⁶ Different search rates⁴⁷ for black motorists depending on whether they are stopped by white or black troopers are taken as evidence of racial animus. Applying their model to data on stops and searches by the Boston police department, Professors Antonovics and Knight find evidence of racial animus

326 (2003) (reaffirming that the use of race in government decisionmaking triggers strict scrutiny under the Equal Protection Clause). At best, the KPT model is agnostic with respect to whether race is an explicit factor in the trooper's decision to search.

43. Several economic studies build on KPT's work. See, e.g., Vani K. Borooah, *Racial Bias in Police Stops and Searches: An Economic Analysis*, 17 EUR. J. OF POL. ECON. 17, 32–33 (2001) (estimating a model similar to KPT's using data on stops and searches in ten separate areas in England); Dhammika Dharmapala & Stephen L. Ross, *Racial Bias in Motor Vehicles Searches: Additional Theory and Evidence*, 3 CONTRIBUTIONS TO ECON. ANALYSIS & POL'Y, No. 1, art. 12 (2004), at 4–7, available at <http://www.bepress.com/bejeap/contributions/vol3/iss1/art12/> (extending KPT to account for the fact that potential offenders can bypass the highway altogether, thus avoiding detection); Harcourt, *supra* note 21, at 1354–71 (mapping out whether racial profiling is rational, and which population should be profiled, using a series of potential elasticities and offending rate parameters); Rubin Hernandez-Murillo & John Knowles, *Racial Profiling or Racist Policing? Bounds Tests in Aggregate Data*, 45 INT'L ECON. REV. 959, 960 (2004) (extending the KPT model to situations in which aggregate, rather than individual-level data, is available); Nicola Persico & Petra Todd, *Using Hit Rates to Test for Racial Bias in Law Enforcement: Vehicle Searches in Wichita 1* (Nat'l Bureau of Econ. Research, Working Paper No. 10,947, 2004), available at <http://papers.nber.org/papers/W10947> (extending the KPT model to allow for heterogeneous police and motorists).

44. See, e.g., KPT, *supra* note 18, at 205–06 (assuming that “the police maximize the number of successful searches, [and] net the cost of searching motorists”).

45. Antonovics & Knight, *supra* note 42, at 3; Shamena Anwar & Hanming Fang, *An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence 3–4* (Nat'l Bureau of Econ. Research, Working Paper No. 11,264, 2004), available at <http://papers.nber.org/papers/W11264>.

46. Antonovics & Knight, *supra* note 42, at 4; Anwar & Fang, *supra* note 45, at 6. Because identification of their full model is impossible, Antonovics and Knight assume that black and white troopers are equally prejudiced. That is, they assume that the difference in search cost of black and white motorists is equal, and they compare the white trooper/black motorist search rate to the black trooper/white motorist search rate. Antonovics & Knight, *supra* note 42, at 12.

47. Technically, both Professors Anwar and Fang and Professors Antonovics and Knight use the search rate for stopped cars, that is, the ratio of the number of vehicles searched to the number of vehicles stopped. Antonovics & Knight, *supra* note 42, at 15–16, 32; Anwar & Fang, *supra* note 45, at 25.

against black motorists.⁴⁸ Professors Anwar and Fang make the same basic finding using data from the Florida state police.⁴⁹

Taken as a group, these economic models have several flaws. First, they assume that the police stop and search vehicles differentially on the basis of the race of the driver—that is, that they engage in racial profiling. The economic models focus on whether doing so is rational. Thus, these economic models cannot answer the threshold legal question of whether the police use race in their decisionmaking. Second, while aggregation of data allows the researchers to estimate the models without controlling for every possible legitimate reason to search, the models cannot control for different levels of specificity in profiles for blacks and whites.⁵⁰ If, for example, the police have a more accurate profile of black drivers because of more information, this would result in a higher hit rate for blacks, which could be masking a taste for discrimination.⁵¹ One possible scenario would be that police obtain information about drug couriers by gathering demographic information on drug couriers who are in prison. As these drug couriers are disproportionately black, the police may have more accurate information regarding the profile of black drug couriers as compared with their white counterparts. If, in fact, the police do engage in significant racial profiling (as the authors of these studies assume), it is logical that they would gain more information about black drug possessors than their white counterparts, as these are the people they investigate more fully.

In addition, all of the models discussed to this point focus on hit rates as the primary measure of efficiency.⁵² This translates to assuming that the police's goal is to deter drug use generally, rather than drug trafficking more specifically. While multiple goals, including the goal of deterring drug use, may be optimal from a policy

48. Antonovics & Knight, *supra* note 42, at 19–20.

49. Anwar & Fang, *supra* note 45, at 26.

50. Antonovics and Knight, in particular, identify racial animus on the basis of different search behavior in cross-race interactions between troopers and motorists, Antonovics & Knight, *supra* note 42, at 12–13, and therefore their results are quite sensitive to this assumption. They address this criticism by providing evidence that the troopers do not have a different amount of information about cross-race motorists. *Id.* at 23–24.

51. A taste for discrimination, under these models, reduces hit rates because the troopers have a lower threshold of suspicion to search black motorists, and therefore will search more innocent black motorists.

52. KPT do investigate the alternative goal of increasing felony hit rates, but it is not the focus of their article. *See supra* note 36.

perspective, the goals may conflict, making them difficult to achieve concurrently.⁵³ Finally, the models do not quantify the costs and benefits of profiling.

While not strictly an empirical piece, Professor Harcourt offers a theoretical model of racial profiling in his recent article “Rethinking Racial Profiling.”⁵⁴ Based on the assumptions that blacks are somewhat more likely to engage in criminal behavior on the highway than white motorists, and that their utility of engaging in criminal activity is more inelastic—that is, blacks will alter their criminal behavior less than whites will when the police change their level of searching—Harcourt concludes that racial profiling is likely counterproductive in that it leads to a higher crime rate (because the ignored group will offend more often) and to more costly searching.⁵⁵ Harcourt also discusses what he calls the “ratchet effect,” which is the idea that racial profiling has a broader impact on a profiled population due to increased fines, incarceration rates, felony disenfranchisement, and other similar costs.⁵⁶ Harcourt does not, however, explicitly model the ratchet effect; it is an additional cost that is not taken into account in his model, which focuses on individual-level costs and benefits.

Professor Harcourt also argues that the empirical literature’s focus on hit rates is misplaced.⁵⁷ Instead, his model minimizes the total social cost of the criminal activity.⁵⁸ Here, Harcourt is particularly concerned with the economists’ rather narrow definition of success, based solely on increasing hit rates, rather than actually lowering crime rates.⁵⁹ As his piece is theoretical, however, Harcourt provides no insight into how to measure social costs. A broad term like “total social cost” suggests that almost any measure is incomplete. While I certainly agree that the focus on finding drugs is too narrow a definition of success, I doubt that the police are

53. See *infra* Part III.E.

54. Harcourt, *supra* note 21, at 1303–06.

55. *Id.* at 1279, 1301–02.

56. *Id.* at 1279–83, 1329–35.

57. *Id.* at 1303–05.

58. *Id.* at 1303. While theoretically possible to incorporate the costs of the ratchet effect into the model via the rather amorphous “social cost” of crime, it is clear from Harcourt’s exposition that this is not his intent. Similarly, Harcourt does not question whether there are more efficient ways to discourage drug trafficking than by searching vehicles on the highway, which deters only drug trafficking on the highway.

59. *Id.* at 1276–77.

farsighted enough to have considered minimizing total social cost as a goal, rather than something more short term and easily measurable, like maximizing the amount of drugs seized.⁶⁰

While Professor Harcourt's economic model is an advance, he does not question the initial assumption that the police engage in racial profiling. Thus, like the rest of the economic models, his cannot answer the initial question whether the police use race as a factor in deciding whom to search. Nor can his model, or the other economic models described above, quantify the harm or benefit from racial profiling; they simply determine whether it is economically efficient⁶¹ or the product of racial animus. This provides an incomplete picture of the costs and benefits of racial profiling.

None of the empirical investigations discussed above explicitly determines the factors that influence the decision to search a stopped vehicle or simultaneously controls for nonracial factors that may influence the decision to search. Without such a model, the investigators must make assumptions about the underlying crime rate that may not be borne out by the data. Most importantly, the studies assume that the base drug-possession rates for blacks and whites are similar. The authors, myself included, provide empirical data to suggest that these assumptions are reasonable, but we do not prove the assumptions correct.

Why has this prior work largely failed to model the decision to search explicitly? One simple reason: lack of data.⁶² In order to estimate a model of the decision to search, one needs detailed individual-level data on the motorists searched, and, more importantly, on the motorists who are not searched.

60. For example, commendations and promotions may well be tied to high-profile drug busts. In addition, individual police departments will not necessarily take total social cost into account; they may simply work to keep the criminal activity out of their neighborhood. The neighborhood is, after all, their constituency.

61. As narrowly defined, economic efficiency simply balances the direct cost to searching, in terms of the trooper's time, against the direct benefit of searching, in terms of the amount of drugs seized. *But cf.* Harcourt, *supra* note 21, at 1300 (arguing that economists define efficiency too narrowly in the context of the criminal justice system).

62. Data collection has become a focus of antiracial profiling efforts over the past several years. *See, e.g.*, DEBORAH RAMIREZ ET AL., A RESOURCE GUIDE ON RACIAL PROFILING DATA COLLECTION SYSTEMS: PROMISING PRACTICES AND LESSONS LEARNED (2000), available at <http://www.ncjrs.org/pdffiles1/bja/184768.pdf>. For an up-to-date list of jurisdictions that track at least some information related to racial profiling, see the Racial Profiling Data Collection Resource Center at the Northeastern University website, <http://www.racialprofilinganalysis.neu.edu>.

One group of investigators does have individual-level data of stops, including information on who was searched, and whether contraband was found.⁶³ Professor Nicholas Lovrich and his coauthors use data on 677,514 traffic stops in Washington State from March 2002 to October 2002.⁶⁴ They estimate several decision models, including the decision to search.⁶⁵ While cautioning that their study is not completed and results are somewhat preliminary, they find that race plays a role in the decision to search, even after controlling for such other factors as seriousness of the violation, sex of the driver, the officer's race and sex, the officer's experience, time of day, and location.⁶⁶

The empirical analysis conducted by Professor Lovrich and his colleagues is a large step forward. They provide the first systematic model of the decision to search a stopped vehicle and find that race is an important factor in the decision. They do not, however, quantify the gains from that decision, nor can they with the data set they describe. Because they do not have information on what was found in the search, they are limited to modeling whether a search occurred, rather than more broadly modeling the gains from searching. In order to discuss the policy implications of racial profiling, however, it is critical not only to determine whether the police use racial profiling, but also to quantify the costs and benefits. Both the civil rights empiricists, in focusing on the first question, and the economists, in focusing on a limited version of the second question, have failed to

63. NICHOLAS LOVRICH ET AL., WSP TRAFFIC STOP DATA ANALYSIS PROJECT 39–111 (June 1, 2003), available at <http://www.wsp.wa.gov/reports/wsptraff.pdf>. Professors Antonovics and Knight also explicitly model the decision to search using individual-level data. See Antonovics & Knight, *supra* note 41, at 15–16. Their model, however, does not control for other important variables beyond the race of the trooper and the race of the motorist because they do not have the detailed data necessary to do so. Their model is therefore significantly underspecified.

64. LOVRICH ET AL., *supra* note 63, at 93.

65. *Id.* at 107–08.

66. *Id.* at 109–10. They also control for the level of discretion in the search and find that race has the same effect in low-discretion searches (searches incident to arrest, impound searches, and warrant searches) as in high-discretion searches (consent searches, Terry searches, and K-9 searches). They find this inconsistent with the premise of racial profiling, which would suggest that more discretion allows race to play a greater role in the decisionmaking process. As the authors also believe this to be inconsistent with prior literature on discretionary searches, they posit that they have not controlled for all the factors that are driving the decision to search. *Id.* It is worth noting that this prior literature is based upon the intuition of many commentators, rather than empirical fact.

integrate these two fundamental questions and therefore have failed to provide a complete response to these questions.

In sum, while all of the prior literature on racial profiling has contributed to society's empirical knowledge of racial profiling in the context of drug interdiction, the authors have not quantified the costs and benefits of racial profiling. The primary difficulty with quantifying the costs and benefits is lack of data: what data is available contains information on the people that the police thought were worthy of investigation, rather than all motorists on the highway. Chances are, the nice little old lady next door is not going to arouse suspicion. But, she may be the person who carries a large quantity of drugs on her trip down to Florida. Or, to return to the example of the Japanese American internment during WWII, how many English American spies did the police overlook because they concentrated only on Japanese Americans? Looking for the wrong people—or more generally, for only a select subset of people—biases the data. Empiricists call this particular type of bias “selection bias” because it is created when data are collected on a subpopulation that was selected nonrandomly.

Unlike prior empirical work, I am able to control for selection bias and to quantify the costs and benefits of racial profiling. To do this, I use an empirical model called the “Heckman selection model” with a combined data set of searches and stops on the highway.⁶⁷ The first step in implementing the Heckman selection model is to determine how drivers are selected for searching; that is, to build a mathematical model that determines the police decisions to search stopped vehicles. In this Article, I do this with data from the MSP.

II. THE DECISION TO SEARCH

I model the average trooper's decisionmaking strategy by estimating the probability that a trooper will search a stopped vehicle with a given set of characteristics.⁶⁸ The model I build relies on data collected on all vehicles stopped by the MSP on a section of I-95. Thus, the universe of individuals included in my dataset is all stopped

67. For further discussion of the Heckman selection model, see *infra* notes 109 & 111 and accompanying text.

68. I use a probit model to estimate the probability p_i that a car with characteristics X_i will be searched. The probit model estimates a model of the form: $\Upsilon^{-1}(p_i) = X_i b$ where b is the vector of estimated coefficients that determine the effect of each X_i on p_i and $\Upsilon^{-1}(\cdot)$ is the inverse cumulative distribution function of the Normal distribution.

vehicles, rather than all motorists driving down the highway. Of course, that set of data is not completely unbiased, because the decision to stop a vehicle driving down the highway is also not completely random; it is, however, likely a more random process with respect to race and the likelihood of drug possession than the decision to search, as the police have very limited information about a vehicle when they decide to stop it. The police might know the race of the driver—indeed, stereotypical racial profiling involves troopers who angle their cars to see into vehicle interiors before pulling behind a vehicle and stopping it—but if the police are more likely to stop black motorists, this would tend to bias the estimated search probabilities for blacks downward.⁶⁹

Before describing the specific modeling choices, let me first provide a brief description of the datasets that I use to estimate the model.⁷⁰ I use data on stops and searches on a portion of I-95 in Maryland. The search dataset and stop dataset are separate, generated in response to two related lawsuits charging the MSP with using racial profiling in traffic stops and searches on the highway.⁷¹ The search dataset is the product of a consent decree in a federal civil rights class action against the MSP.⁷² The search dataset contains all searches of vehicles in Maryland by the MSP from January 1995 to December 2003.⁷³ The information available for each search includes the grounds the MSP used to justify the search (probable cause or consent); the reasons the MSP asked for consent or believed they had probable cause; characteristics of the car (such as make, model, state of registration, and model year); characteristics of the driver (such as race and sex); date and time of stop; as well as what was found during the search (quantity and type of drugs, guns, and money). There are approximately 15,000 searches in the dataset.

69. That is, if one assumes that stops are random events, and further assumes that the relative search rates of stopped vehicles across race are equal, one would *underestimate* the overall differential search rate for blacks if blacks were profiled at both stages of the process. For a discussion of the salience of the empirical difference between stopped vehicles and all vehicles on the highway, see *infra* note 110 and accompanying text.

70. The dataset is available at <http://law.wustl.edu/Academics/Faculty/Barnes/index.html>.

71. *Wilkins v. Md. State Police*, No. CCB-93-468 (D. Md. filed 1993) (settled by consent decree in 2003). Its companion case was filed in 1998. *NAACP v. Md. Dep't of State Police*, No. CCB-98-1098 (D. Md. filed Apr. 10, 1998) (settled by same consent decree in 2003).

72. *Wilkins*, No. CCB-93-468.

73. Although the search dataset contains information on all searches throughout Maryland, I use only the searches on a portion of I-95, in order to match the search dataset to the geographically limited stops dataset.

The dataset on stops arose out of concerns that the MSP were not creating a complete dataset of searches. At that point, the litigation had focused primarily on one group of MSP: those assigned to the JFK Barracks.⁷⁴ This dataset records every stop on the 49.5-mile stretch of I-95 north of Baltimore to the Delaware state line that defines the jurisdiction of the JFK Barracks. The stops data begin in May 1997 and record the location, date, and time of the stop; the make of the car; any traffic code violation; whether a ticket, warning, or safety violation was issued; and the race and sex of the driver. In a subset of the cases, the data also record the posted speed limit and the speed of the vehicle. There are 199,409 stops in the dataset from May 1, 1997, to December 31, 2003. Narrowing the search dataset to the time period and geographic area covered by the stop dataset yields 2,583 traffic stops for which search data are also available.⁷⁵

Although stops are far more likely than searches to be conducted randomly with respect to race and the likelihood of drug possession, as I mention above, there is still a possibility of some bias. If black

74. See Memorandum in Support of Plaintiffs' Motion for Enforcement of Settlement Agreement and for Further Relief, *Wilkins v. Md. State Police* (D. Md. 1997) (Civil Action No. CCB-93-468), at Parts III-IV (arguing that the I-95 corridor and MSP troop numbers 62 (the JFK Barracks) and 55 conduct the vast majority of searches) (on file with the *Duke Law Journal*); Order, *Wilkins v. Md. State Police* (D. Md. Apr. 22, 1997) (Civil Action No. CCB-468), ¶ 1 (finding a pattern or practice of racial discrimination in vehicle searches on the I-95 corridor) (on file with the *Duke Law Journal*); see also Elliott Wolf, ACLU, *The Use of Racial Profiling by the Maryland State Police in Drug Interdiction from 1995-2002*, at 2-3 (discussing stop and search patterns of the JFK Barracks) (on file with the *Duke Law Journal*).

75. For the purposes of this study, the primary problem with the stops and searches databases maintained by the MSP is that the two datasets are not linked. Thus, there is no variable in the stops data that indicates whether a search was conducted, or what that search found. One large data project has been the matching of searches to stops based on the time, location, direction of travel, trooper, vehicle make, and race and sex of the driver. In order to match a search to a particular stop, I require that the search occur within two hours of the stop, that the direction of travel match, if available in both datasets, and that the location of the stop be within five miles of the location of the search (to account for variations in data collection). In addition, I match searches based on four additional factors: make of vehicle, name of trooper, race of driver, and sex of driver. I use the order of this list as a tiebreaker. Thus, if two different stops match three criteria for the same search, the stop that matches all but the sex of the driver is preferred to the stop that matches all but the race of the driver, etc. With this rather stringent algorithm, I am able to match 1,248 of 2,583 total searches, or 48.3 percent of the searches on I-95 during the appropriate time period. Stop data from July 1, 2002, through December 31, 2003, did not contain the time of the stop; for this subset of data, stops were matched by date and the above criteria. For the stops and searches before July 1, 2002, when time data was available, 931 of 1486 searches were matched, for a match rate of 62.6 percent. For all of the data, under the assumption that the search data is more accurate, if the data sets diverge on control variables, I use the search data for any value that does not match the stop data.

drivers are more likely to be stopped for a given traffic violation than white drivers, evidence of bias at the search level will be attenuated. If white drivers are more likely to be stopped, the situation would be reversed, but a white bias in stops is unlikely.⁷⁶

A. *Model Specification*

The primary question I address is whether race is a part of the profile the MSP use in determining whether to search a stopped vehicle. To analyze this question, I use a regression-based model that incorporates a variety of independent variables, including the driver's race. Other independent variables both control for alternate reasons for the search and describe more accurately the decision to search overall. I discuss the variables of interest below.

1. *Characteristics of the Driver.* The aggregate I-95 data reveal that men were driving a majority of the cars stopped by the MSP. This disparity grows with respect to searches: 73.6 percent of vehicles searched had male drivers. Received wisdom suggests that men—in particular young men—are more likely than women to be carrying drugs. Thus, driver gender may be a part of the trooper's profile, and I include the driver's sex as an independent variable to control for this possibility.⁷⁷ In addition, race and sex may interact in important ways. For example, black men may be more likely than white men to carry drugs, but black women may not be more likely to carry drugs than white women. Because of this possibility, I include variables that capture the interaction of sex and race.

2. *Characteristics of the Vehicle.* I also include the make and age of the vehicle as independent variables. To capture the effects of the vehicle's make, I include binary variables representing the two categories of vehicles considered to have an innately higher

76. For example, evidence suggests that blacks are slightly more likely to speed than whites, making a white bias less likely. See Gross & Barnes, *supra* note 19, at 664, 687–88.

77. I do not have data for two particular characteristics of the driver, age and income, and therefore cannot control directly for either. I control for the driver's income indirectly through the make and age of the vehicle driven. I do not have a proxy for age in the data, however, so I cannot control for that variable. As blacks and Hispanics are, on average, younger populations than are whites, if MSP select younger drivers for search more often, this could be masked as a selection on the basis of race.

likelihood of search: luxury cars⁷⁸ and large commercial trucks.⁷⁹ These variables capture the typical example of racial profiling— young black men driving expensive cars—and the possibility that professional truck drivers have a higher tendency to be searched. In controlling for the vehicle’s age, I include a variable to capture whether it is an older model (ten or more years old).⁸⁰ I do not include a variable for the color of the car because of insufficient data, although trooper lore suggests that black or red cars are searched more often.⁸¹

Finally, I use a series of regional variables to control for the state in which the vehicle is registered. I break the states up into several groups representing different regions along I-95, the major north-south highway on the East Coast, and the rest of the country. The specific groups are Maryland, Eastern states north of Maryland (excluding New York), New York, Eastern states south of the District of Columbia, the District of Columbia, and the rest of the country.⁸² It is important to control for region because the Federal Drug Enforcement Administration (DEA) created a profile in the early 1990s suggesting that drugs traveled south from New York City to the Baltimore/D.C. metropolitan area, while cash traveled north from Baltimore to New York.⁸³ Whether or not the DEA profile is accurate, the state of origin will influence a trooper’s decision to search if that decision is based upon the DEA profile. More

78. Luxury models include Acura, Audi, Bentley, BMW, Cadillac, Hummer, Infiniti, Jaguar, Land Rover, Lexus, Lincoln, Mercedes, Porsche, Range Rover, Rolls Royce, Saab, Sterling, and Volvo.

79. Large trucks include International, Mack, Peterbilt, and Kenworth trucks.

80. I experimented with different definitions of “older car”; none made a substantive change in the results of the model.

81. There is, for example, a widely held belief that red cars cost more to insure because their drivers are, on average, more reckless. See CarInsurance.com, FAQs, <http://www.carinsurance.com/kb/content10059.aspx> (last visited Oct. 17, 2005) (dispelling this myth).

82. In the tables that follow, I denote these categories as Maryland Plates, Northeast Plates, New York Plates, Southeast Plates, DC Plates, and Non-East Coast Plates, respectively.

83. David Kocieniewski, *New Jersey Argues That the U.S. Wrote the Book on Race Profiling*, N.Y. TIMES, Nov. 29, 2000, at A1; see also *United States v. Wilson*, 853 F.2d 869, 875 (11th Cir. 1988) (discussing a DEA course that taught Georgia officers that “drug couriers are frequently Hispanics”); *United States v. Layman*, 730 F. Supp. 332, 334, 337 (D. Colo. 1990) (noting that officers were trained by DEA agents to use drug courier profiles that included race as a factor); TASK FORCE ON GOV’T OVERSIGHT, OPERATION PIPELINE REPORT (1999), available at <http://www.aclunc.org/discrimination/webb-report.html> (noting the California Highway Patrol’s use of racial profiling as part of the DEA’s “Operation Pipeline” protocol).

generally, troopers may believe that individuals traveling from specific regions may be more or less likely to carry drugs, and therefore they may search vehicles from those regions more or less often.

3. *Characteristics of the Encounter.* A typical explanation for the higher percentage of blacks being stopped and searched is that blacks drive differently,⁸⁴ the implication being that individuals who drive less safely are more likely to possess drugs.⁸⁵ I control for this potential explanation by including several variables capturing the violation upon which the stop was based. The two primary violations were speeding (65.8 percent of the stops) and seatbelt violations (16.6 percent of the stops). Two other violations were registration infractions such as missing or expired registration or tags (3.8 percent of the stops) and broken or missing lights (1.6 percent of the stops).⁸⁶ Finally, while an infrequent violation, I control for driving while intoxicated (DWI) (about 0.8 percent of the stops). The dataset also provides information regarding whether a ticket, warning, or safety violation was issued, so I use variables to capture those outcomes as well. Because of the small number of safety violations, I include them with warnings and look only at whether a ticket was issued. In some specifications, I also include a variable that controls for driving at night, defined as between 9:00 p.m. and 5:00 a.m., because troopers may be more likely to search late at night, both for logistical reasons and because traveling late may be correlated with criminal activity. Data on time of day are missing for all stops after June 30, 2002; for this reason, all models that use nighttime as an independent variable must drop all the data from July 1, 2002, through December 31, 2003, a total of almost 50,000 stops.⁸⁷

I include a variable to capture the direction of travel because of the DEA profile, mentioned above, that suggests that drugs travel south on I-95 from New York City to Baltimore. Again, whether or

84. See Gross & Barnes, *supra* note 19, at 687–88 (discussing this possibility).

85. Blacks may be more likely to be stopped because of unsafe driving, but unless unsafe driving is positively correlated with drug possession, there is no reason to search unsafe drivers disproportionately more often.

86. Many of the stops include more than one violation.

87. The models I report in Tables 1–2 and 4–8 *infra* are estimated using the full dataset and drop nighttime as an independent variable. I indicate in footnotes when the alternative model—including nighttime as a variable but excluding stops and searches made after June 30, 2002—provides different results.

not the DEA profile is accurate, if the troopers rely on the profile, the direction of travel should influence a police officer's decision to search. Finally, I also include indicator variables for all police officers performing more than one hundred stops during the time period of the study, May 1, 1997, to December 31, 2003, because troopers may set significantly different thresholds for the probability of finding drugs above which they will search a vehicle.⁸⁸

B. Results: The MSP's Search Profile

The results from a model with these independent variables demonstrate that the driver's race is the most salient factor in a trooper's decision to search a stopped vehicle. I focus on consent searches because they arguably allow more discretion on the part of the police officer.⁸⁹ Nonconsent searches are based on a wide range of circumstances—including a vehicle search incident, an arrest, and an officer noticing the odor of marijuana when approaching a vehicle. While the individual MSP officers still have discretion to ignore the marijuana odor, nonconsent searches do not provide information about racial profiling as I define it. Because the troopers have probable cause to believe a crime has been committed, the issue is one of selective enforcement of the laws, rather than an a priori belief that blacks and Hispanics are more likely to commit a crime. Table 1 provides the details of the results from this model. Black drivers who are stopped face a 0.80 percent search rate—2.6 times higher than the 0.31 percent probability of search faced by otherwise-similar white drivers. Hispanics fair worse. They are searched 3.5 times more often than whites. This is the effect of race after controlling for other relevant variables and is statistically significant with p-values for the black and Hispanic variables of less than 0.0005.⁹⁰ The most likely

88. Out of the 240 troopers who stopped more than 100 vehicles, 20 of the troopers never conducted a search. Thus, variables for these 20 troopers perfectly predict no search (or a probability of searching equal to zero). I therefore drop the stops by these troopers from the model, as they provide no further explanatory power. A total of 14,781 stops were dropped from the full dataset for this reason.

89. *But see* LOVRICH ET AL., *supra* note 63, at 107–10 (describing potential means by which “non-discretionary” searches may still allow room for discretionary decisionmaking).

90. A p-value represents the probability that a result as extreme or more extreme would occur when there was, in fact, no relationship between the independent variable and the decision to search. For example, in this case, the p-value of less than 0.0005 represents the probability that after controlling for other relevant variables, the estimated difference in search rates between blacks and whites would be as great or greater if there were no relationship

explanation for this is that the MSP engaged in large-scale racial profiling.⁹¹

In addition, female drivers who have been stopped are far less likely to have their vehicles searched than male drivers. The interaction between race and sex is also significant, which means that the difference in search rates between men and women changes depending on the race of the men and women compared. Overall, Hispanic male drivers are the most likely to be searched and are 7.4 times more likely to be searched than white women, even after controlling for vehicle characteristics and the particular characteristics of the encounter. Black men are 5.5 times more likely to be searched than white women. The MSP's target in searching is clearly black and Hispanic men.

Beyond the primary findings that race is the most salient factor in the MSP's profile and that black and Hispanic men are the primary target of the MSP, there are several other interesting findings. Vehicles from Washington, D.C., and non-East Coast states are slightly more than 1.5 times as likely to be searched as vehicles from Maryland; this difference is statistically significant at the 0.0005 level. Vehicles registered in New York are slightly more likely to be searched, with a relative risk⁹² of 1.21. As the data does not provide more specific information than state of registration, it may be that those searched are disproportionately from New York City in particular, but I cannot determine this. Those searched are, however, disproportionately from Washington, D.C., non-East Coast states, and, to a lesser extent, from New York.

between race and the decision to search. A coefficient with a p-value of 0.05 or less is generally described as statistically significant.

91. If a very powerful independent variable that is highly correlated with race were omitted from the model, the conclusion of racial profiling would be invalid. But this is very unlikely, especially when race is a stronger predictor of searches than variables that are part of the DEA's drug interdiction profile.

92. The relative risk of a search is the additional risk faced by a driver with a specific characteristic, expressed as a ratio, compared to the baseline of a white male motorist driving a newer car with Maryland plates who did not get a ticket. For each variable, the relative risk is given by:

$$\frac{\text{Pr (Search given variable (e.g., black))}}{\text{Pr (Search given baseline)}}$$

TABLE 1. DETERMINANTS OF THE DECISION TO SEARCH

Variable	Relative Risk	P-Value
Baseline	1.00	N/A
Race/Sex of Driver ⁹³ (Control: White Male)		
Black Male	2.59	0.000 ⁹⁴
Hispanic Male	3.48	0.000
White Female	0.47	0.004
Black Female	0.68	0.601
Hispanic Female	1.26	0.001
Luxury Vehicle/White Driver	0.51	0.005
Luxury Vehicle/Black Driver	2.73	0.057
Luxury Vehicle/Hispanic Driver	5.11	0.177
State of Origin (Control: Maryland Plates)		
Northeast Plates	1.15	0.171
New York Plates	1.21	0.000
Southeast Plates	1.70	0.167
DC Plates	1.53	0.000
Non-East Coast Plates	1.71	0.000
Seatbelt Violation	0.40	0.112
DWI Violation	0.82	0.489
Registration Violation	0.85	0.268
Lights Violation	0.82	0.659
Speeding Violation	0.89	0.117
Traveling South	1.12	0.632
Ticket Issued	0.96	0.001
Older Vehicle	1.32	0.000

93. In the standard statistical practice, these variables are not interaction terms, but instead are dummy variables for each demographic group, comparing the group (Hispanic women, for example) against white men. This is equivalent to the standard statistical practice of using “interaction” terms, in which a variable labeled “Hispanic X female” would estimate whether the difference in search rates between Hispanic men and women was the same as the difference in search rates between white men and women. I report dummy variables in the tables to make the demographic differences more transparent. Other relative risks may be found by division:

$$\text{Relative Risk of Variable A to Variable B} = \frac{\text{Relative Risk of Variable A}}{\text{Relative Risk of Variable B}}$$

Thus, for example, the relative risk of a Hispanic man being searched as compared to a white woman is: $3.48 / 0.47 = 7.4$.

94. A p-value of 0.000 simply means that the true p-value is less than 0.0005; rounded to three significant digits, this is 0.000.

Turning to the effect of different predicate violations on the decision to search, only one variable predicts whether a consent search occurs: if a ticket is issued, a driver is slightly less likely to be searched. While small differences are often statistically significant in such large datasets, the relative risk of 0.96 is sufficiently close to 1 to be of little practical significance. Other violations are not statistically significant at the 5 percent level.⁹⁵

The type of vehicle stopped also impacts the likelihood of a search. Older vehicles are 1.32 times more likely to be searched than other vehicles. Luxury vehicles driven by white motorists are significantly less likely to be searched; white motorists stopped in luxury vehicles are about half as likely to be searched as similar motorists in nonluxury cars. The benefit of driving luxury vehicles disappears, however, for black motorists. Overall, black motorists driving luxury cars are 2.7 times as likely to be stopped as white motorists driving nonluxury cars. This effect is marginally statistically significant, with a p-value of 0.057.⁹⁶ As compared to black motorists in nonluxury vehicles, black motorists stopped in luxury cars are equally likely to be searched (relative risk of 1.03). As compared to white motorists in luxury cars, black motorists in luxury cars are 5.4 times more likely to be searched if stopped. This difference is significant at the 5 percent level. No other vehicle characteristic is statistically significant, suggesting that the MSP do not rely on these traits in their decisions to search stopped vehicles.

Although not listed in Table 1, several individual troopers have search rates significantly different from the average, ranging from a search rate of 0 percent (zero searches out of 3,319 stops) to a search rate of 10.3 percent (111 searches out of 1,079 stops); the average search rate is 0.6 percent. Despite the DEA profile, southbound

95. In the alternative model that includes nonconsent searches and controls for these probable cause searches, whether a vehicle was speeding is a statistically significant predictor of a search. Specifically, speeders are less likely to be searched, with a relative risk of 0.40. This implies that nonconsent searches are disproportionately nonspeeders (or, conversely, consent searches are disproportionately performed on speeders) and therefore may suggest that speeding is being used as a pretext for stops.

96. Very few Hispanic motorists were stopped driving luxury vehicles; only 275 Hispanic drivers driving luxury vehicles were stopped, as compared with over 14,000 whites and 7,000 blacks. With such a small amount of data, the coefficient for Hispanic motorists driving luxury vehicles is not statistically significant. Based upon the magnitude and direction of the estimated effect—Hispanic motorists in luxury cars are over ten times as likely to be searched after being stopped than their white counterparts—the same pattern may hold, but without more data, it is impossible to determine whether this is the case.

motorists are not more likely to be searched, at least after controlling for other relevant variables. Looking more closely, however, it appears that the trooper fixed effects may be masking a southbound effect; several troopers *only* search southbound vehicles. Because of the trooper fixed effects, the southbound searches for these troopers provide no additional information with which the model can identify the southbound variable. Essentially, to the model, these searches do not matter. In a model without trooper fixed effects, southbound drivers are 1.25 times as likely to be searched.⁹⁷

In sum, the MSP determine which stopped vehicles to search on the basis of their personal threshold for searching and the race and sex of the driver, with black and Hispanic men being searched most often. The other factors in the troopers' profile include traveling southbound and driving an older vehicle registered in the southeast or in non-East Coast states.

To make the multivariate nature of the predictions more transparent, Table 2 provides a list of potential drivers on the highway and their individual probabilities of being searched, based on a given set of characteristics.

97. This is statistically significant, with a p-value of 0.003.

TABLE 2. SEARCH PROBABILITIES FOR
DIFFERENT DEMOGRAPHIC GROUPS

Driver Description ⁹⁸	Pr(Search)	Relative Risk
White male from Maryland in newer car, traveling north	.31%	1
Black male from D.C. in newer luxury car, traveling south	.93%	3.0
Hispanic male from New York in older luxury car, traveling south	2.30%	7.5
White female from New York in older car, traveling south	.24%	0.8
Black female from southeastern U.S. in newer car, traveling north	.36%	1.2
Hispanic female from northeastern U.S. in older car, traveling north	.58%	1.9
White male from non-East Coast state in newer luxury car, traveling north	.28%	0.9
Black male from southeastern U.S. in older luxury car, traveling south	1.88%	6.1
White male from non-East Coast state in older car, traveling south	.53%	1.7
Hispanic male from non-East Coast state in older car, traveling south	3.28%	10.7

Table 2 demonstrates that the MSP focus the bulk of their attention on black and Hispanic motorists. As compared to white men stopped in newer, nonluxury vehicles from Maryland traveling northbound during the day, black and Hispanic motorists, depending on the other characteristics, are 1.2 to 10.7 times more likely to be searched. The category of driver facing the highest overall probability of being searched is a black male motorist from a non-East Coast state, stopped in an older luxury car, traveling southbound. This is the strongest version of the MSP's profile and confirms that the MSP engage in racial profiling. The driver's race is clearly one element of a trooper's decision to search a stopped vehicle.

98. The estimates provided in this table hold all other variables constant at their medians (the baseline value).

III. QUANTIFYING THE COSTS AND BENEFITS OF PROFILES

Having determined which motorists, once stopped, face the highest likelihoods of being searched by the MSP, I turn to the second empirical project of this Article: quantifying the costs and benefits of the MSP's use of racial profiling. What does engaging in racial profiling gain the MSP? One justification is that the MSP find more drugs because they focus on black and Hispanic motorists.⁹⁹ This Part estimates a model that determines what the gains to racial profiling are in terms of drug users and couriers arrested and total drugs seized.¹⁰⁰ Along the way, the model also quantifies one direct cost of racial profiling: the number of innocent motorists searched.¹⁰¹

Why quantify the benefits of racial profiling? Critics of racial profiling argue that such quantification serves little good; even if blacks carry drugs more often, or are more likely to be drug couriers than whites, the gains from using race in a search profile are minimal.¹⁰² After all, drug interdiction on the highway involves picking the drivers most likely to be carrying drugs out of hundreds of thousands of motorists each day. It is an (almost) futile exercise. How could profiling based upon a broad measure that cuts across many other salient demographics be significantly beneficial? On the flip side, if drug interdiction on the highway is so difficult, why should the police not be allowed to use race if, in fact, it works? A first look at the data suggests that racial profiling may be advantageous, but it also suggests a high cost, in terms of greater numbers of innocent motorists searched.

99. See, e.g., Mac Donald, *Myth of Racial Profiling*, *supra* note 17, at 20 (arguing that racial profiling increases the amount of drugs seized).

100. The cut points determining the boundaries between "user" and "courier" are based on government estimates of the average prices and the amounts spent by users of various contraband drugs in 1998. See WILLIAM RHODES ET AL., OFFICE OF NATIONAL DRUG CONTROL POLICY, WHAT AMERICA'S USERS SPEND ON ILLEGAL DRUGS 1988-1998, at 12, 16, 22-23 (2000). The cut points were chosen to balance the cost versus weight for different drugs. For marijuana, the cut point is 455g = \$5,120 or five years' supply. For crack and powder cocaine, the cut point is 50g = \$7,450 or about nine months' supply. For heroin, the cut point is 10g = \$10,290 or about eleven months' supply. Finally, for other drugs, the amounts were recoded, very roughly, into dosage units; a cut point of 150 dosage units defines a drug courier. See Gross & Barnes, *supra* note 19, at 696 n.152, for further details.

101. I recognize that this is not a complete measure of the costs of racial profiling, even on an individual level. Other costs, not easily quantified, include the loss of dignity, the stigma of being searched, and the deterioration of police-citizen relationships in minority communities.

102. See HARRIS, *supra* note 16, at 84-86 (noting that in "stops and searches," police rarely find evidence of crime); Gross & Barnes, *supra* note 19, at 750-52.

A. *An Initial Look at the Data (and Why It Is Misleading)*

Table 3 presents the frequency distribution of the amount of crack and powder cocaine seized by the MSP, by race. Eighty-three percent of the drivers found with larger amounts (over 50 grams) of these drugs are black or Hispanic (25 out of 29). This table makes clear why police have a greater tendency to search cars driven by blacks and Hispanics: that is where the drugs are.¹⁰³

TABLE 3. NUMBER OF SEARCHES FINDING
CRACK AND POWDER COCAINE, BY QUANTITY

Quantity	# White Drivers	# Minority Drivers	Total
None	523	654	1177
Trace–50 grams	17	26	43
50–100 grams	0	3	3
100–500 grams	2	10	12
500–1,000 grams	0	5	5
Over 1,000 grams	2	7	9
<i>Total</i>	544	705	1,249

Or is it? Recall that the search data I have are not a random sample of drivers on the highway; nor are they a random sample of drivers who were stopped by the MSP.¹⁰⁴ Instead, the data are a sample of those individuals, already stopped for some reason, whom the MSP believed were most likely to be carrying drugs. From Part II.B, we know that these individuals are not a random sample of the population. Indeed, those searched are disproportionately black and Hispanic men who are traveling south in older cars. Without controlling for this selection bias, there is no way of confirming that the police are using the right profile. It may be that white drug couriers are driving on I-95 as well, but because they do not fit the MSP's profile, they do not get searched. The MSP would have no idea that these white drug couriers were driving on the highway or being stopped and not searched. The data would seem to confirm that the MSP were doing their job: arresting the drug couriers. To the MSP, it


103. Or, at least, where crack and powder cocaine are. The pattern is similar if one looks at all drugs. Thirty-nine out of forty-nine drivers found with large amounts of drugs are black or Hispanic.

104. See *supra* p. 1107 for the discussion of selection bias inherent in the dataset.

would just be an unfortunate fact of life that drug couriers are almost all black or Hispanic.¹⁰⁵ The models I introduce below in Parts III.B, C, and D control for this selection bias and allow me to discover if the MSP's perceptions of reality based on the searches they conduct are accurate, or if they are missing an entire subgroup of drug traffickers because of a poor search profile.

Figure 1 provides a different perspective on the search data, this time using data on all drugs, rather than just crack and powder cocaine, by presenting the number of searches finding no drugs (innocent drivers), some drugs (drivers who are drug users), and large amounts of drugs (drivers who are drug traffickers).¹⁰⁶ Figure 1 makes clear that, while minority drivers comprise a larger percentage of the drug dealers, there is a high cost to searching more minorities: because innocent drivers are more likely to be black or Hispanic, more innocent drivers will be searched when the MSP engage in racial profiling.

FIGURE 1. TYPES OF DRIVERS BY MINORITY STATUS



105. Cf. Kathy Barrett Carter & Ron Marsico, *Whitman Fires Chief of State Police*, STAR LEDGER (New Jersey), Mar. 1, 1999, at 1A (providing an example of comments by the chief of the New Jersey State Police on the racial stratification of drug trafficking); C.J. Chivers, *Ex-Police Leader's Claim of Bias Attacked*, N.Y. TIMES, Oct. 4, 1999, at B4 (noting that other documents also describe drug trafficking along racial lines).

106. I define the threshold amount of drugs to be considered a drug courier by reference to government estimates. See *supra* note 100. These thresholds are based upon a combination of street price and number of dosage units. Gross & Barnes, *supra* note 19, at 696 n.152.

Or does it? What if the MSP are just more discerning when they choose to search a vehicle driven by a white person? If they require a higher threshold of suspicion before subjecting whites to a search, they would search fewer innocent whites, and greater numbers of innocent blacks and Hispanics.¹⁰⁷


Again, because of selection bias, this data does not show that minorities are more often innocent of drug crimes; instead, it shows that in the population of drivers searched, minorities accounted for two-thirds of the innocent drivers searched. It may be that blacks are less likely to carry drugs on the highway, but these data do not show this directly because the data have not been corrected for selection bias.

Figure 2 illustrates the problem of selection bias graphically. For ease of discussion, the three groups in Figure 1—innocents, users, and couriers—are collapsed into two categories: innocent motorists and motorists carrying drugs. Panel A provides a graph of the search data actually observed. Note that there are slightly more guilty minority motorists in the data, suggesting that racial profiling has some benefits. Panel B adds hypothetical data of motorists who were stopped but not searched. The hypothetical data are represented by hatch marks; they are reasonable if the troopers do engage in racial profiling, meaning that they search vehicles driven by minorities stopped on the highway at much higher rates.¹⁰⁸ Panel C simply shades in the hypothetical data, illustrating that the underlying guilt rate is the same for both groups—minorities and whites in stopped vehicles are carrying drugs about 25 percent of the time, and the relative percentages of white and minority motorists stopped on the highway are 75 percent and 25 percent, respectively. The white motorists, however, are searched at a much lower rate; they drive down the highway carrying drugs with little chance of being caught, even if a trooper stops them. Minority motorists, in contrast, are very likely to get caught once stopped by a trooper.

107. This searching strategy is a form of racial profiling that biases the data. Racial profiling is just one selection mechanism that would bias the data; the model I use controls for whatever selection criteria the MSP actually use.

108. The hypothetical data are simply for illustrative purposes. I created the hypothetical data with two goals: first, in order to make the final graph—Panel C—demonstrate the same offense level across minority status and second, to create a reasonably accurate racial mix of highway drivers.

FIGURE 2. SELECTION BIAS AT WORK



These graphs simply illustrate that the data describing every search the MSP have conducted are necessarily incomplete; without knowing what the hypothetical data on the stopped vehicles look like, one cannot make inferences about racial profiling using the search data. Luckily, there are statistical methods that provide a “best guess” of what the hypothetical data look like, and thereby control for the bias inherent in the decision to search selected vehicles.

So what would these figures look like if the MSP did not engage in racial profiling? How many innocent drivers and drug couriers would be searched? And how much contraband would the MSP find? These are the questions answered by the next three models.

B. Benefits to Profiling: Contraband Seized

The first model I estimate investigates which variables in a trooper’s search profile are most predictive of the amount of contraband likely to be found. If the MSP’s goal is to maximize the amount of drugs seized, this model determines what the MSP’s profile should be—factors that truly correlate with carrying large amounts of drugs. In order to control for the bias introduced because of the MSP’s nonrandom selection of vehicles to search, I use a Heckman selection model that estimates the quantity of drugs carried by a random sample of stopped vehicles on I-95, even though the data are not a random sample of stopped vehicles.¹⁰⁹ The Heckman selection model first estimates a selection submodel that determines which stopped vehicles get searched; this submodel is essentially the same as the model described earlier in Part II. Then, the Heckman model uses the results of that submodel described in Part II.B to control for the selection bias in a second submodel that estimates the outcome of interest (in this case, quantity of drugs seized).¹¹⁰

109. See generally James J. Heckman, *Sample Selection as a Specification Error*, 47 *ECONOMETRICA* 153 (1979).

110. The selection model does not control for the potential selection bias inherent in data of all stopped vehicles, rather than all vehicles—i.e., the bias due to the initial selection of which vehicles to stop. This selection bias is almost certainly less salient than the bias introduced by the decision to search; the troopers have much less information when deciding to stop a vehicle than when deciding to search a stopped vehicle. To test this assumption, I re-ran the Heckman models using only stops and searches that occurred at night, when the troopers have even less information and have more difficulty discerning the race of a motorist before stopping the vehicle. The substantive results do not change significantly. Because only slightly less than half the data is used, however, the statistical significance for several variables is lower.

While it may seem that the Heckman selection model gives one something (information about the drugs carried by individuals who are not searched) for nothing (no data is actually available), this is not the case. The Heckman selection model is not magic;¹¹¹ rather, the idea is simply to extrapolate the information from the fact that the individuals were not selected to be searched. Essentially, the selection model exploits the observable differences in the two populations—those searched, and those only stopped—to correct for the bias in looking only at the quantity of drugs actually possessed by the searched subgroup. Through this process, the model provides a best guess of what (and how much) drugs the unsearched motorists would carry, using this to create estimates based upon the entire population of stopped vehicles, rather than just the subset of those that were searched.¹¹² In order to do this, as with any statistical model, the data must have some variability—that is, for any given characteristic, some motorists with that characteristic must be searched, and some must not be. This variability is what identifies the parameters, allowing for an accurate estimation of the parameters.

In addition, to work well, the two Heckman submodels should include different independent variables, which provide the traction necessary to identify, or estimate accurately, the model.¹¹³ Without

111. This is not to belittle the idea in any way; indeed, James Heckman won the 2000 Nobel Prize in Economics for his insights in this area. See Nobelprize.org, *The Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel: 2000* (noting that Heckman won “for his development of theory and methods for analyzing selective samples”), at <http://www.nobel.se/economics/laureates/2000> (last modified Apr. 14, 2005). But Heckman, too, recognizes the limits of his model. See James Heckman & Bo Honoré, *The Empirical Content of the Roy Model*, 58 *ECONOMETRICA* 1121, 1122 (1990) (noting that the Heckman model can be identified by model assumptions alone, and that care must be taken to make sure that the data are driving the results).

112. As an example, suppose that the selection equation—just like the one estimated in Part II.B—estimates that Troopers *A* and *B* search 5 percent of the vehicles they stop, and find, respectively, an average of two grams and four grams of cocaine per stop, after controlling for other factors. Suppose further that the model estimates that Trooper *C* searches 10 percent of the vehicles stopped, and finds seven grams of cocaine on average, after controlling for other factors. The Heckman selection model then assumes that a motorist’s chance of being searched is correlated with the amount of drugs the motorist carries. The Heckman selection model estimates that motorists who have a 10 percent chance of being searched carry about seven grams of cocaine, and motorists who have a 5 percent chance of being searched carry about three grams of cocaine. The model extrapolates to motorists who have even lower (or higher) chances of being searched, and thereby provides a best estimate of what the unsearched motorists stopped on the highway are carrying.

113. Identification, in statistical terms, is an important concept concerning what information is driving the results: information from the data or from the model assumptions. Ideally, one

different independent variables in the two submodels, there is nothing in the data itself that allows an investigator to distinguish between different hypotheses; any estimation of coefficients results primarily from modeling assumptions, rather than from the data.

My identification strategy relies upon the individual trooper fixed effects. Essentially, I assume that individual troopers differ significantly in the threshold of suspicion they require in order to conduct a search, and, therefore, individual troopers have different base search probabilities. Once the decision to conduct a search is made, however, I assume that every trooper would find the same contraband—that is, that the troopers would be equally thorough once the initial decision to conduct a full search is made.

The model estimates the probability of drug possession in the population of stopped cars by extrapolating across different officers. Table 4 provides a simplified version of how my identification strategy works. For clarity, I merge individual troopers into three groups based on their individual probability of conducting a search. I create three groups: troopers whose search probability is low (less than 0.2 percent), medium (0.2 percent to 0.5 percent), or high (more than 0.5 percent). Focus first on the second column of data in Table 4, which provides the percentage of drug couriers found among the searches conducted. As the probability of searching increases, so does the probability of finding a drug courier—suggesting that MSP who search more vehicles are selecting the right vehicles. Extrapolating to the concept of searching all stopped cars suggests that the underlying percentage of drug couriers among stopped vehicles is greater than 4.4 percent. This upward trend presages the formal results from this model, which estimate an overall rate of 7.2 percent for all stopped cars.

In contrast, the first and third columns show no monotonic relationship. This suggests that there is no clear direction in which the population of stopped cars differs from the set of searched cars. And again, the formal results agree: the models estimate no significant selection bias—that is, they find no statistically significant difference between the population of stopped cars and the set of searched cars, in terms of either the hit rate for drug possession or the value of drugs possessed.

wants to identify the results from the data rather than the model assumptions, which are somewhat arbitrary.

TABLE 4. TROOPER-SPECIFIC PROBABILITY OF SEARCH AND CONTRABAND FOUND (NO CONTROL FOR SELECTION BIAS)

Search Probability	Value of Drugs		
	Seized per Search	% Drug Courier	% Hit Rate
Low (less than 0.2%)	\$1.45	1.2	36.6
Medium (0.2-0.5%)	\$1.33	2.6	33.7
High (more than 0.5%)	\$1.63	4.4	34.9

Having discussed the Heckman model in general and my identification strategy in particular, I turn to the actual model and its results. There are several independent variables that I incorporate into the model to determine whether they are correlated with the quantity of drugs actually carried by a particular motorist. The first of these variables is race: if race is statistically significant in this model, it implies that the amount of drugs traveling up and down the highway in Maryland varies by the race of the driver carrying them—and therefore that racial profiling has some benefit in terms of the amount of drugs seized by the MSP.¹¹⁴ Whether the driver's race is a salient factor in determining the quantity of drugs seized is the primary question this Section seeks to answer. It is also important, however, to control for other variables that possibly correlate with the amount of drugs carried in order to isolate the relative importance of race. These other variables, if statistically significant, are also interesting in their own right; they constitute actual indicators of the likelihood that a vehicle is carrying drugs and should, therefore, define the MSP's search profile. For this reason, I include them.¹¹⁵ In order to estimate this model, I combine the total amounts seized for crack and powder cocaine, heroin, and marijuana, based upon the street value of the drugs seized.¹¹⁶ I limit my analysis to these because the "other" category is too varied to convert to a simple cash value. Dispensing

114. This does not mean that racial profiling needs to target blacks or other minorities. It may be that whites carry larger quantities of drugs than blacks or Hispanics do, implying that profiling whites would increase the quantity of drugs seized.

115. Thus, I control for sex of driver, direction of travel, the vehicle's region of origin, whether the vehicle is a luxury model, and the interaction of luxury vehicle with race, age of vehicle, and whether the driver received a ticket. *See supra* Part II.A.

116. Because crack and powder cocaine represent the bulk of the drugs seized, I ran several alternate models using the total amount of drugs seized, with different composite measures of total drugs seized (equalizing either estimated number of dosages or estimated value of drugs seized). The substantive results were unchanged. For ease of discussion, I concentrate on powder drugs in this portion of the analysis.

with this category should not affect the results significantly, as these drugs are only found in 1.4 percent of the searches.

Table 5 provides the model's estimate of the effect of each variable on the value of the drugs seized.¹¹⁷ The average stopped vehicle driven by a white male from Maryland traveling northbound carries \$6 worth of drugs. The model estimates an additional \$2.7 million of drugs among those stopped vehicles that are not searched.¹¹⁸ Turning to the specific factors that correlate with greater drug possession, the model indicates that the driver's race plays only a marginal role in predicting the value of drugs carried. There is no statistically significant difference in the estimates of the value of drugs carried by black, Hispanic, and white men after controlling for other factors. Hispanic women, however, carry \$2 worth of drugs on average, or one-third the amount that white men do. This difference is statistically significant at the 5 percent level.¹¹⁹ A somewhat surprising result is that white motorists driving luxury vehicles carry almost \$20 worth of drugs on average, although this finding is only marginally significant. In contrast, black motorists driving luxury vehicles carry only \$4.69 on average—about three-quarters of the amount carried by white men in nonluxury vehicles, and less than one-quarter as much as white motorists in luxury vehicles. This is directly contradictory to the MSP's chosen profile, which samples black men, particularly those in luxury vehicles, at higher rates. These results are also contrary to what most law enforcement personnel believe, and even to what many civil rights activists believe to be the case.¹²⁰ Nonetheless, the model provides a best estimate of the mean value of drugs carried by the general population of drivers stopped on the highway, rather than the significantly skewed quantity of drugs carried by the subgroup of drivers searched. This conclusion suggests that the MSP profile's focus on black and Hispanic men¹²¹ does a poor job of maximizing the value of drugs seized on I-95.

117. To estimate the model, I transform the cash value of the drugs seized to the logarithmic scale in order to stabilize the variance. Specifically, I use the transformation $y = \log(\text{cash value} + 1)$ to avoid taking the logarithm of zero. In the tables that follow, however, I transform the logged value back to the standard scale for ease of discussion.

118. To make this estimate, I simply calculate the estimated average value of drugs carried by each type of motorist and sum across all nonsearched vehicles.

119. Given the number of nonsignificant relationships found and the magnitude of the p-value, even this relationship may represent random variation.

120. *See supra* note 83.

121. *See supra* Part II.B.

TABLE 5. ESTIMATES OF THE VALUE OF
DRUGS SEIZED IN STOPPED VEHICLES

Variable ¹²²	Value of Drugs Seized ¹²³	P-Value
Baseline	\$5.98 ¹²⁴	N/A
Race/Sex of Driver (Control: White Male)		
Black Male	\$8.93	0.231
Hispanic Male	\$6.47	0.828
White Female	\$12.69	0.237
Black Female	\$14.86	0.906
Hispanic Female	\$1.87	0.016
Luxury Vehicle/White Driver	\$19.30	0.094
Luxury Vehicle/ Black Driver	\$4.69	0.015
Luxury Vehicle/ Hispanic Driver	\$16.33	0.954
State of Origin (Control: Maryland Plates)		
Northeast Plates	\$3.98	0.270
New York Plates	\$2.93	0.168
Southeast Plates	\$2.37	0.016
D.C. Plates	\$6.21	0.979
Non-East Coast Plates	\$2.17	0.006
K-9 Alert	\$4.94	0.527
Grounds "Nervous" ¹²⁵	\$4.01	0.308
Owner Not Present	\$6.54	0.696
Ticket	\$9.26	0.020
Traveling South	\$11.30	0.009
Older Vehicle	\$4.76	0.312

122. Statistically significant variables and their corresponding values are shown in bold.

123. This column represents the average additional (above the baseline) value of drugs that a driver of this description will have, holding all other variables fixed at their medians, which is the baseline value.

124. This is the average value of drugs seized for "baseline" motorists: white men speeding north in their newer, nonluxury cars with Maryland plates.

125. Being nervous when stopped seems ubiquitous; in this context, an individual stopped on the highway is "nervous" if the officer conducting the search lists nervousness as one of the grounds for the search. The results of the model confirm the intuition that after controlling for other factors, nervousness of the driver, as reported by an MSP trooper after the fact, is not correlated with the quantity of drugs carried.

In the course of controlling for factors other than race that might affect the value of drugs carried, the model predicts what the MSP's profile should be. Few factors are statistically significant; the value of drugs carried does not, apparently, vary across many factors. Three variables do have some predictive power. Motorists in stopped vehicles from off the East Coast carry an average of \$2.17 worth of drugs or about two-thirds less, for white men, than motorists from Maryland. Again, this contradicts the MSP's profile, as the MSP searched vehicles with non-East Coast plates 1.7 times as often as vehicles with Maryland plates.¹²⁶ Similarly, vehicles from the southeastern United States carry only \$2.37 worth of drugs. While these vehicles were not oversampled as compared to Maryland state vehicles, from an efficiency point of view, the MSP should search southeastern vehicles less often, rather than equally often. The MSP did get one factor right, however: drugs do travel south. A vehicle stopped southbound carries, on average, \$11.30 worth of drugs, or almost double the value of drugs that would be traveling north in a similar vehicle.

C. Benefits to Profiling: Couriers Arrested

Thus far, I have focused on the MSP's goal of maximizing the quantity of drugs seized by measuring the value of the drugs seized. But it may be that the MSP sets as a goal arresting the maximum number of drug couriers, instead of, or in addition to, maximizing the quantity of drugs seized. Seizing drugs removes them from the street, but arresting people keeps drug couriers incapacitated for years. It may be that the additional benefit of seizing about \$300,000 (2,000 grams) versus only \$150,000 (1,000 grams) of cocaine is marginal at best. The motorist will be incarcerated for years in either case,¹²⁷ and while the money involved is a considerable sum for the average person, seizing an extra kilogram of cocaine is not even a drop in the

126. As would be true with any model, the data can only determine the profile in use *during the time period of data collection*. Thus, the data show that the profile used from May 1997 to December 2003 did not mirror the profile that would have yielded the maximum quantity of drugs seized. The data cannot determine whether the profile was appropriate at some prior time, but then became obsolete as drug couriers adapted to the policing strategy. The issue of such dynamic behavior is beyond the scope of this paper.

127. In both cases, the individual would be eligible to be tried as a drug kingpin. MD. CODE ANN., CRIM. LAW §§ 5-612, 5-613 (2002) (setting a threshold of 448 grams, or 16 ounces, of cocaine to be tried and sentenced as a "drug kingpin").

bucket of the drug traffic in Baltimore alone.¹²⁸ The second potential benefit of racial profiling I investigate, therefore, is the number of drug couriers arrested. I return to investigating all drugs because the definition of drug courier requires only one comparison across drug types—the cut point for defining the amount of drugs carried by a user versus a courier—and therefore using all drugs does not pose analytical problems as significant as those that exist when trying to model the value of drugs seized across drug types.¹²⁹ To do this, I model the probability that the driver of a stopped vehicle is a drug courier. Again, because the search data I use does not contain data on what the MSP would have found had they searched stopped vehicles randomly, I need to control for the bias that the MSP's selection criteria create by using a Heckman selection model.¹³⁰

The Heckman selection model's estimates of the prevalence of drug couriers among the general population of stopped vehicles are provided in Table 6 and suggest a different profile from the one discussed in Section B. Again, with the exception of Hispanic women, race has no direct effect on the outcome—in this case, the probability of being a drug courier. However, the interaction between race and driving a luxury vehicle is highly salient; white and black motorists driving luxury cars have a small probability of carrying sufficient drugs to be considered drug couriers. Black men in luxury vehicles are 0.58 times as likely to be drug couriers as white men in nonluxury cars. In contrast, Hispanic men driving luxury vehicles are 7.55 times more likely to be drug couriers than the baseline of white men in nonluxury vehicles. Thus, if the MSP's goal were to maximize drug couriers arrested, racial profiling against Hispanic motorists driving luxury vehicles would further this goal, but racial profiling targeted at black motorists in luxury vehicles would be counterproductive.

128. Gross & Barnes, *supra* note 19, at 751–52 (providing a rough estimate that the 33 kilograms seized per year represented about 0.5 percent of the cocaine consumed in the Baltimore/D.C. metropolitan area per year).

129. See *supra* note 100 (describing the cut points used to discern a drug user from a drug courier). A small subset of the drivers carry more than one drug. I define these individuals to be drug couriers if any one of the drugs they carry meets the appropriate threshold. If, instead, I were to combine the drug amounts by standardizing the amount of each separate drug seized to a percentage of the amount necessary to be deemed a drug courier, no additional drivers would be labeled drug couriers.

130. In this case, I use a variant of the standard Heckman selection model that allows me to estimate a model with a binary dependent variable. See G.S. MADDALA, LIMITED-DEPENDENT VARIABLES AND QUALITATIVE VARIABLES IN ECONOMETRICS 221–23 (1986) (describing the Heckman selection model).

TABLE 6. ESTIMATES OF DRUG COURIERS
ARRESTED IN STOPPED VEHICLES

Variable ¹³¹	Relative Risk	P-Value
Baseline: White Male, Non-East Coast Plates, Consent Search ¹³²	1	N/A
Race/Sex of Driver (Control: White Male)		
Black Male	2.09	0.184
Hispanic Male	2.43	0.271
White Female	1.46	0.757
Black Female	5.33	0.617
Hispanic Female	0.00 ¹³³	0.002
Luxury Vehicle/White Driver	0.00 ¹³⁴	0.016
Luxury Vehicle/Black Driver	0.58	0.023
Luxury Vehicle/Hispanic Driver	7.55	0.012
State of Origin (Control: Maryland Plates)		
Northeast Plates	1.41	0.646
New York Plates	1.51	0.631
Southeast Plates	1.30	0.726
D.C. Plates	3.19	0.424
Non-East Coast Plates	0.81	0.779
K-9 Alert	2.07	0.126
Grounds “Nervous”	2.00	0.379
Owner Not Present	1.14	0.764
Speeding	0.77	0.610
Ticket	1.00	0.992
Traveling South	11.15	0.001
Older Vehicle	0.62	0.317

131. Statistically significant variables and their corresponding values are shown in bold.

132. The baseline probability of a trooper arresting a drug courier after stopping the vehicle is 1.0 percent.

133. Again, in the dataset, no Hispanic women—out of 3 searched—carried a large quantity of drugs. Thus, the model’s best estimate for the relative risk of being a drug courier for Hispanic women is zero, even though the true relative risk is very likely nonzero, although still small.

134. Similarly, in the dataset, no white men driving luxury vehicles—out of 29 searched—carried a large quantity of drugs. They did, however, carry just under this amount of drugs. Thus, the model’s best estimate for the relative risk of being a drug courier for white men driving luxury vehicles is zero, even though the true relative risk is very likely nonzero, although still small.

Almost no other variables, in fact, are statistically significant in predicting whether a stopped driver is a drug courier. Drivers traveling southbound are 11.15 times more likely to be drug couriers than their northbound counterparts when compared at the baseline value. Once again, Hispanic women are significantly less likely than others to be drug couriers; in fact, the estimated probability is so low as to be almost zero. State of registration, anxiety when stopped, and a K-9 unit alerting to drugs are all insignificant. The fact is, very few of the variables are predictive at all—except for southbound travel, and race when combined with driving a luxury car.

These results are not directly contradictory with the profile described in Section B; the MSP can maximize the value of drugs seized, as well as the number of drug couriers arrested, by profiling southbound vehicles, while avoiding Hispanic women. The question is whether they should also profile Hispanic motorists that drive luxury vehicles. The model predicting drug couriers suggests that the benefits are high; the model predicting the total value of drugs suggests that there is no benefit, but also no cost, in terms of the value of drugs seized. This may seem contradictory—after all, the same motorists who are drug couriers must be carrying a large amount of drugs. It is consistent, however, with insufficient data in the first model, which attempts to provide much more specific information (an exact value of drugs seized), compared to the simpler question of whether the amount exceeds some threshold. Simply put, stopping only 275 Hispanic motorists in luxury vehicles does not provide sufficient variability in the value of drugs seized to identify the “Luxury Vehicle/Hispanic Driver” parameter with enough precision to discern an effect. Collection of more data may help in this respect, but given that the MSP only stopped 275 such motorists over a period of six and one-half years, it may be a long wait to gather sufficient data.

D. Costs to Profiling: Innocent Motorists Searched

In order to assess the impact of various profiling strategies on innocent motorists, the final model I investigate determines what factors predict whether a driver is carrying any amount of drugs,¹³⁵ that is, the ubiquitous “hit rates” on which many commentators

135. Once again, I use any type of drug found in this analysis, rather than limiting the analysis to powder drugs found.

focus.¹³⁶ I do this for two reasons. First, arresting drug users is another potential goal that the MSP may have in their drug interdiction program. As being arrested is, at best, a huge inconvenience, arresting drug users makes drug use more expensive and less attractive on average. Second, investigating hit rates necessarily means investigating their opposite: fail rates. And fail rates are important because they describe which subgroup of innocent drivers (those not carrying any drugs) is most burdened by racial profiling. If, for example, blacks are less likely to carry drugs than whites, but more likely to be drug couriers, one would conclude two things: profiling blacks would increase the number of drug couriers arrested, but it would also increase the number of innocent drivers searched.

It turns out that a variant of this hypothetical is true, when confined to Hispanic motorists driving luxury vehicles. Table 7 provides the details.¹³⁷ Compared to whites, Hispanic motorists in luxury vehicles are 1.01 times more likely to be innocent. Of course, 1.01 is very close to 1, so this statistically significant finding does not have much substantive weight. If, however, the MSP decided to focus on Hispanic motorists in luxury vehicles more than they do now, the absolute number of additional innocent Hispanics may become quite large. Black motorists driving luxury vehicles are *less* likely to be

136. See, e.g., HARRIS, *supra* note 16, at 78–82 (citing hit rate statistics to argue that “[r]acial profiling is neither an efficient nor an effective tool for fighting crime”); Antonovics & Knight, *supra* note 42, at 6–7 (same); Gross & Barnes, *supra* note 19, at 689–93 (same); KPT, *supra* note 18, at 219–24 (modeling whether a motorist carries drugs).

137. Before discussing what the three models I estimate in this Part together imply about the MSP’s profiling strategy, let me address one further point. The model predicts that the baseline rate of innocence, for white men from Maryland, is 73 percent, or, equivalently, that 27 percent of stopped vehicles have some contraband. This is a very high percentage. The model also estimates that the MSP’s selection criteria do not bias the results; the MSP would not get different results if they randomly searched stopped vehicles. These two facts are consistent: the underlying base rate of offense for stopped vehicles is approximately equal to the offense rate of those vehicles that were selected to be searched. This high offense rate suggests that some selection is happening at earlier stages—in particular, in the trooper’s decision to stop a speeding car, or in the motorist’s decision to carry drugs while driving on I-95. Most likely, the decision to stop is not random and is correlated with having drugs. For example, reckless drivers may be overrepresented in both stopped vehicles and vehicles that contain drugs. For this to be a problem with my analysis of the use of *race* in profiles, however, the overrepresentation of reckless drivers in stops also needs to correlate with race. While I do not have individual-level data on the motorists who were not stopped, and therefore cannot control for selection using a Heckman-style model, I did re-analyze the results using only stops and searches that occurred at night and obtained very similar results. Assuming that it is much more difficult to ascertain the race of motorists at night before stopping them, this suggests that my analysis of the racial element of the MSP’s profiling is robust to the selection bias inherent in choosing which vehicles to stop.

either drug couriers or users; they are 1.1 times more likely to be innocent than the baseline driver. In sum, even if the MSP did not engage in racial profiling, blacks and Hispanics would be overrepresented among innocent motorists. Profiling blacks and Hispanics exacerbates this problem.¹³⁸

What other factors influence innocence rates beyond race? Sex of the driver is unimportant; apparently drug use is equal opportunity when it comes to gender. Drivers from the southeast and D.C. are innocent 0.89 and 0.58 times more often than Maryland drivers, respectively. Southbound motorists are about two-thirds as likely to be innocent as their northbound counterparts, when compared at the baseline value. Whether a K-9 unit alerted to drugs predicts a higher likelihood of drug possession. Being particularly nervous or driving a vehicle without the owner present are also salient characteristics, making searching an innocent driver somewhat less likely, although both these factors are only marginally statistically significant. Finally, receiving a ticket, rather than a warning or a safety violation, predicts a significantly lower innocence rate—these drivers are 0.40 times as likely to be innocent as those who do not receive tickets. No other factor is statistically significant.¹³⁹

138. Profiling black and Hispanic motorists may also have created this problem in the first place, in that black and Hispanic motorists may choose to carry drugs less often because they are searched more often. This is the type of game theoretic argument that economists use in their models. Even if the MSP changed their policy, however, this would not necessarily change the behavior of black and Hispanic motorists; after all, they drive through many jurisdictions quite quickly. Only a wholesale change in the practice of most jurisdictions would lead to a significant change in drug possession.

139. In the alternate model that uses both consent and nonconsent searches, probable cause is the largest determinant of innocent rates. Vehicles searched due to probable cause have innocent drivers 1.5 times more often than vehicles subject to consent searches.

TABLE 7. DETERMINANTS OF THE INNOCENCE RATE

Variable ¹⁴⁰	Relative Risk of Innocence	P-Value
Baseline: ¹⁴¹ White Male with Non-East Coast Plates	1	N/A
Race/Sex of Driver (Control: White Male)		
Black Male	0.96	0.289
Hispanic Male	1.10	0.459
White Female	0.94	0.373
Black Female	0.86	0.562
Hispanic Female	1.38	0.857
Luxury Vehicle/White Driver	0.67	0.000
Luxury Vehicle/Black Driver	1.10	0.033
Luxury Vehicle/Hispanic Driver	1.01	0.009
State of Origin (Control: Maryland Plates)		
Northeast Plates	0.74	0.542
New York Plates	0.97	0.204
Southeast Plates	0.89	0.002
D.C. Plates	0.56	0.003
Non-East Coast Plates	0.92	0.726
K-9 Alert	0.78	0.002
Grounds “Nervous”	0.83	0.056
Owner Not Present	0.61	0.078
Speeding	0.78	0.631
Ticket	0.40	0.027
Traveling South	0.66	0.000
Older Vehicle	0.66	0.741

140. Statistically significant variables and their corresponding values are shown in bold.

141. As in Table 1, this value represents the baseline probability that a white male motorist stopped on the northbound highway in a newer inexpensive car with Maryland plates is carrying drugs.

E. The Final Analysis: Costs and Benefits Compared

Table 8 summarizes the three models that I have discussed in this Part, and what each model suggests about the optimal profile for the MSP to use. Such a profile would include those variables that predicted more drugs seized and drug couriers arrested, while simultaneously minimizing the effect on innocent motorists. Only two variables unambiguously belong to this category: southbound travel and luxury vehicle. Together, they provide a small benefit from profiling that does not impose a greater burden on innocent motorists than random searches. One cannot say the same about racial profiling, however. Race is the best predictor of drug couriers available in the data; but using race as a factor in deciding which vehicles to search places a huge cost on innocent motorists.

TABLE 8. COMPARISON OF RACIAL
PROFILING'S BENEFITS AND COSTS

Profile	D Value of Drugs Seized	Arrested Drug Couriers ¹⁴²	Effect of Innocents ¹⁴³
Baseline: White Male, Non-East Coast Plates	\$5.98	1	1
Race/Sex of Driver (Control: White Male)			
Black Male	— ¹⁴⁴	—	—
Hispanic Male	—	—	—
White Female	—	—	—
Black Female	—	—	—
Hispanic Female	-\$4.12	0.00	1.38
Luxury Vehicle / White Driver	\$13.32	0.00	0.67
Luxury Vehicle / Black Driver	-\$1.29	0.58	1.10
Luxury Vehicle / Hispanic Driver	—	7.55	—
State of Origin (Control: Maryland Plates)			
Northeast Plates	—	—	—
New York Plates	—	—	0.97
Southeast Plates	-\$3.61	—	0.89
D.C. Plates	—	—	—
Non-East Coast Plates	-\$3.81	—	0.92
Grounds "Nervous"	—	—	—
K-9 Alert	—	—	—
Owner Not Present	—	—	—
Ticket	\$3.28	—	0.40
Traveling South	\$5.32	11.15	—
Older Vehicle	—	—	—

142. This column represents the relative risk of finding and arresting a drug courier, as compared to the baseline value, which holds all variables at their medians.

143. This column represents the relative risk of searching an *innocent* motorist with the given characteristics compared to the baseline value, which holds all variables at their medians.

144. All cells marked "—" have no statistically significant effect.

An optimal profile would have a positive impact on the quantity of drugs seized and drug couriers arrested and a negative effect on innocents. Second best would be to have a nonnegative (positive or zero) relationship on the quantity of drugs seized and drug couriers arrested, and a nonpositive (negative or zero) effect on innocents. Two factors satisfy the first criteria, but only in the negative: according to Table 8, Hispanic women and black motorists driving luxury vehicles should not be searched. Doing so would decrease drugs seized, find fewer drug couriers, and increase the number of innocents searched. While the models provide some traction on whom the MSP should avoid, none of the factors satisfy the optimal criteria to determine on whom the MSP should focus. However, several factors satisfy the criteria for a second-best profile. In order to maximize these three goals, the MSP should profile (1) southbound travelers (2) Hispanic men driving luxury vehicles; and (3) motorists who receive tickets.¹⁴⁵ The gain in quantity of drugs seized would be somewhat small, but the gain in identified drug couriers could be large. Table 9 compares the results of using a few sample profiles against the option of random searching.

TABLE 9: PROFILING COSTS AND
BENEFITS FOR DIFFERENT DEMOGRAPHIC GROUPS

Driver Description	Value of Drugs Seized	Pr(Drug Courier)	Innocence Rate
No Profile; Random Searches	\$13.39	7.2%	70.5%
Black Men in Luxury Vehicles	\$6.74	5.3%	79.6%
Vehicles Traveling South	\$15.82	12.1%	71.1%
White Men	\$11.15	5.0%	70.7%
Luxury Vehicles Traveling South	\$28.13	3.4%	61.0%
Hispanic Men in Luxury Vehicles	\$18.32	21.7%	80.2%
White Women	\$21.48	6.7%	64.1%
Hispanic Women	\$1.50	0%	99%

The first thing that stands out about Table 9 is that Hispanic women are truly the optimal subgroup to avoid. They carry very little

145. It is possible that the decision to write a ticket is made after a search is performed, in which case this factor would not be salient.

drugs and are estimated by the model to be innocent almost all the time. While this is in part due to little variation in the data, the fact remains that only one Hispanic woman searched had any contraband, and it was worth only \$1.49.

Surprisingly, white women also appear to be good candidates for profiling. They carry large average amounts of drugs, are reasonably likely to be drug couriers, and have a low innocence rate. These effects, however, are not significantly different from profiling white men, who do not have such advantages. Interestingly, profiling vehicles driving south provides close to the same benefits as profiling Hispanic men in luxury vehicles, with fewer innocent motorists searched. In addition, Table 9 clearly presents that profiling black men in luxury vehicles is counterproductive, as compared to random searches. This finding is contrary to the MSP's actual profile, and much of the literature on racial profiling as well.

CONCLUSION

As the past twenty years have suggested, there is no magic bullet in the war against drugs. Any policy requires trade-offs; using racial profiling as a strategy in drug interdiction is no exception. While I do not directly weigh the costs and benefits in this Article, it is clear that there is scant payoff to using profiles that turn largely on race. Not only do such profiles have the potential for exacerbating racial tensions by increasing the propensity to search innocent minorities, but using profiles can also be directly counterproductive, leading the police to ignore signs that suggest a search is in order, particularly when they have stopped white men or women.

My empirical analysis demonstrates that the MSP engaged in racial profiling between May 1, 1997, and December 31, 2003. They searched black and Hispanic motorists stopped on the highway more than 1.5 times as often as white motorists, and these searches cannot be explained by nonracial cues. This profile provides some benefits to the MSP, particularly in terms of drug couriers arrested. But it also comes with specific costs: profiling black motorists in luxury vehicles yields fewer drugs and fewer drug courier arrests, and it increases the number of innocent motorists subjected to a search. The same is true for the MSP's profiling of Hispanic women. This empirical study suggests that there may be some nonracial factors that would be effective in a profile, particularly driving southbound, but random searches also appear to work about as well as other profiles, with the

exception of avoiding black motorists driving luxury vehicles and Hispanic women. While this strategy would provide gains in drugs seized and drug couriers arrested, without increasing the cost imposed on innocent motorists, it is equivalent to profiling all other groups, including white motorists. This, of course, may not be politically feasible or normatively desirable, but it is important to note that it is just as effective a law enforcement tool as profiling Hispanic motorists driving luxury vehicles.

Returning to the example of the Japanese American internment, what insight might this empirical study on racial profiling in the context of drug interdiction provide? First, a reality that is counterintuitive to many—that black motorists are not more likely to carry large amounts of drugs, and, in fact, are more likely to be completely innocent of drug crimes—may in fact be true. Taking the analogy further, it may be true that the Japanese American internment was counterproductive, in that interning a different group of people would have provided significantly more benefits. This is not to say, of course, that any internment is necessarily appropriate. I only wish to emphasize that the reality that people believe to be true at a given time—that the Japanese Americans were more likely to be dangerous traitors, or that black and Hispanic motorists are more likely to carry drugs (and large amounts at that)—may be false, and relying upon these assumptions may be directly counterproductive.

I come, finally, to the question that this Article speaks to but does not answer: is racial profiling worth the cost? By quantifying the benefits and some of the costs of racial profiling, this Article provides a first step in answering this question. There it stops, however, allowing the reader to answer this normative question.