NBER WORKING PAPER SERIES

RACIAL DISPARITIES IN MOTOR VEHICLE SEARCHES CANNOT BE JUSTIFIED BY EFFICIENCY

Benjamin Feigenberg
Conrad Miller

Working Paper 27761
http://www.nber.org/papers/w27761

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2020

We thank Ian Ayres, Felipe Goncalves, Peter Hull, Jonathan Leonard, David Levine, Dan O'Flaherty, Steven Rivkin, Evan Rose, Yotam Shem-Tov, and seminar participants at USC, University of Illinois, Chicago, the Online Economics of Crime seminar, the Online Economics of Racism seminar, and NBER Summer Institute for helpful comments. We thank researchers at the Stanford Open Policing Project for providing data on Texas Highway Patrol stops. We thank the Fisher Center for Real Estate and Urban Economics at Haas for providing computing resources. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Racial Disparities in Motor Vehicle Searches Cannot Be Justified by Efficiency
Benjamin Feigenberg and Conrad Miller
NBER Working Paper No. 27761
August 2020
JEL No. J15,K42

ABSTRACT

During traffic stops, police search black and Hispanic motorists more often than white motorists, yet those searches are equally or less likely to yield contraband. We ask whether equalizing search rates by motorist race would reduce contraband yield. We use unique administrative data from Texas to isolate variation in search behavior across highway patrol troopers and find that, across troopers, search rates are unrelated to the proportion of searches that yield contraband. Our results imply that, in partial equilibrium, troopers can equalize search rates across racial groups, maintain the status quo search rate, and increase contraband yield.

Benjamin Feigenberg
Department of Economics
University of Illinois at Chicago
University Hall room 706
601 South Morgan Street
Chicago, IL 60607
bfeigenb@uic.edu

Conrad Miller
Haas School of Business
University of California, Berkeley
Berkeley, CA 94720
and NBER
ccmiller@berkeley.edu

I Introduction

During routine traffic stops, black and Hispanic motorists are more likely to be searched for contraband by police than white motorists (Pierson et al., 2019). These disparities invite allegations that police engage in *racial profiling*, using race as one factor when deciding whether to search someone. Many regard the practice as unjust: perceived profiling undermines trust in police (Epp et al., 2014) and racial disparities in search rates likely contribute to racial differences in arrests and exposure to police use of force (Fryer, 2019). However, equalizing search rates by race may reduce the effectiveness of policing if race is an informative predictor of criminal behavior (Persico, 2002).

In this paper we evaluate whether racial profiling in fact poses an equity-efficiency trade-off. The answer has important practical implications; recent legal scholarship argues that profiling is not legally permissible in the absence of "legitimate law-enforcement-related necessity" (Tiwara, 2019). Researchers have found that the percentage of searches that yield contraband—known as the *hit rate*—among black and Hispanic motorists is typically equal to or lower than the hit rate for white motorists (e.g., Pierson et al., 2019). Some argue that this pattern indicates equal or lower rates of offending among black and Hispanic motorists (Harcourt, 2004). This reasoning suggests that equalizing search rates across motorist racial groups would not decrease overall contraband yield.

This argument implicitly assumes that equalizing search rates across motorist racial groups would not change group-specific hit rates. However, this assumption fails to hold if troopers face diminishing returns to search. If racial profiling is (accurate) statistical discrimination applied by police to maximize the proportion of their searches that yield contraband, then the hit rate for the marginal black or Hispanic motorist—the last black or Hispanic motorist deemed suspicious enough to be searched—will be equal to that for the marginal white motorist. This logic motivates the Becker (1957, 1993) *outcome test*: to evaluate whether police are on the efficient frontier, test whether *marginal* hit rates are equal across motorist racial groups. If troopers face diminishing returns to search, the hit rate for the average and marginal search may differ significantly, and comparisons of average hit rates between motorist racial groups may be uninformative—the well-known *inframarginality problem* (Ayres, 2002).

We test for an equity-efficiency trade-off using data on traffic stops for speeding violations conducted by Texas Highway Patrol troopers. To examine whether equalizing search rates across motorist racial groups would reduce contraband yield, we exploit variation in search behavior across troopers. In our setting, the identity of the trooper conducting a speeding stop is plausibly exogenous conditional on the location and time of the stop. We measure variation across troopers in the rate at which they search motorists—their *search rate*. Across troopers, we estimate the relationship between search rates and the percentage of stops that yield contraband (the *unconditional* hit rate), where we calculate these rates separately by motorist racial group. Strikingly, this relationship is approximately linear within each motorist group, implying approximately constant returns to search across troopers. In other words, troopers who search motorists twice as often find contraband twice as often. We show that, under conditions consistent with our setting, this result implies that there is no inframarginality problem because average and marginal hit rates are similar. Among motorists searched with positive probability, troopers appear unable to distinguish between those who are more or less likely to carry contraband. Our findings imply that it is feasible for troopers to (1) search all motorist racial groups at the same rate, (2) maintain the status quo overall search rate, and (3) increase contraband yield.

Our analysis proceeds in four steps. First, we summarize racial disparities in search rates and hit rates. A unique feature of the data is that they contain identifying information on stopped motorists, which allows us to (a) track motorists across multiple stops and to (b) merge in additional data on motorist characteristics, including criminal history and neighborhood income. Conditional on stop location and time, we find that black and Hispanic motorists are about 170% and 70% more likely to be searched than white motorists. However, searches of these motorists are about 15% and 30% less likely to yield contraband. Controlling for stop history, criminal history and neighborhood income reduces the black-white and Hispanic-white disparities in search rates by about 50% and 35%. The residual black-white difference in search rates is similar in magnitude to the increase in search likelihood associated with multiple previous non-drug arrests and half of the increase associated with a prior drug arrest. Among stopped motorists with no arrest record at the end of our sample period, black and Hispanic motorists are about 200% and 80% more likely to be searched than white motorists. We investigate whether these stark racial disparities in search rates, even among motorists who do not engage in crime, are a necessary feature of contraband yield maximization.

Second, we introduce a simple model of trooper search behavior to frame our analysis. We build on Anwar and Fang (2006), where troopers decide whether to search a stopped motorist using a noisy signal for the motorist's guilt. Prior research assumes that troopers can strictly rank motorists by contraband risk and hence face strictly diminishing returns to search. This implies that troopers equalizing marginal hit rates across groups leads to an equity-efficiency trade-off. By contrast, we allow signals to be coarse in the sense that, among those motorists searched with positive probability, troopers are unable to distinguish between those who are more or less likely to carry contraband. In this case, the returns to search are constant, and equalizing hit rates at the margin no longer implies an equity-efficiency trade-off.

Third, we use between-trooper variation in search behavior to trace out the relationship between trooper search rates and unconditional hit rates—the between-trooper *search productivity curve (SPC)*—separately by motorist group. Our identifying assumption is that, conditional on the location and time of a stop, the identity of the trooper conducting the stop is as good as random. For a speeding violation at a given time of the week and on a given stretch of highway, the identity of the trooper making the stop will vary due to week-to-week variation in shift schedules and to within-shift variation in exact trooper location. We show that troopers vary substantially in the rate at which they search motorists following speeding stops. While motorists are searched in about 1% of stops, search rates range across troopers from 0% at the 10th percentile to 3% at the 90th percentile. Conditional on the location and time of the stop, motorists with different characteristics are stopped by troopers with similar search rates. We find that the between-trooper SPC is approximately linear for each motorist group. The slopes of groups-specific SPCs—which give the between-trooper returns to search—are similar for black and white motorists, while the returns to search are lower for Hispanic motorists.

The key threat to our approach is that, conditional on our measures of stop location and time, troopers vary in the composition of motorists they stop. This variation may exist because our location and time measures are not sufficiently granular or because, in the same environment, troopers vary in the motorists they decide to stop.

We address this concern using several approaches. We show that baseline SPC estimates do not change if we use more granular location and time measures and control directly for observable motorist characteristics. Baseline SPC estimates are also invariant to removing troopers that stop motorists with unusual observable characteristics from the analysis. We also corroborate our baseline SPC estimates using two alternative research designs. First, we find similar patterns when we rely only on within-motorist variation in outcomes among motorists stopped multiple times. Second, we compare stop outcomes on opposite sides of trooper patrol area borders in a spatial regression discontinuity (RD) design. Along the same highway route, the composition of troopers making traffic stops changes sharply across patrol area borders. We use this feature to validate trooper search rates as measures of the causal effect of trooper assignment on search likelihood and to confirm that average and marginal hit rates are approximately equal.

Fourth, we examine whether troopers themselves face constant marginal returns to search. The linear between-trooper SPCs we identify suggest that there is no equity-efficiency trade-off in our setting. However, search rate equalization across motorist racial groups requires that individual troopers change their search behavior. Determining what would happen in this scenario requires knowledge of the *within-trooper* SPC: the within-trooper relationship between search rates and contraband yield. The between-trooper SPC and average within-trooper SPC may differ if, for example, troopers that are better at screening also search motorists at higher rates. The two SPCs are identical if trooper screening ability and search propensity are independent.

We show suggestive evidence that the between-trooper and average within-trooper SPC are similar. Estimating the within-trooper SPC is complicated by the fact that within-trooper variation in search rates in part derives from changes in trooper screening ability and motorist composition. We instrument for trooper search rates using variation over time in coworker search rates, which we show is largely unrelated to observable variation in motorist characteristics. The within-trooper SPC slope we estimate is statistically indistinguishable from the between-trooper slope, though the within-trooper estimate is imprecise. We conclude that, in partial equilibrium, racial profiling does not present an equity-efficiency trade-off. Moreover, we present suggestive evidence that motorist racial group-specific *deterrence* effects are negligible at the margin, implying that predicted changes in search productivity are unlikely to be offset by changes in contraband carrying behavior.

Previous researchers have argued that disparate policing behavior is driven by racial bias (e.g., Pierson et al., 2019), and this channel offers one potential explanation for our findings. We examine how search patterns vary with several proxies for trooper preferences and beliefs: trooper race (Anwar and Fang, 2006; Close and Mason, 2007; Antonovics and Knight, 2009; West, 2018), local political preferences (Cohen and Yang, 2019), and citation behavior (Goncalves and Mello, 2018). We find that all trooper racial groups are more likely to search black and Hispanic motorists than white motorists, but the black-white disparity is smaller for black troopers. The black-white search disparity is also substantially larger in counties with higher Republican vote shares in the 2016 presidential election. While we find no clear link between black-white disparities in search and citation rates, in counties where troopers cite Hispanic motorists more often, troopers are also more likely to search Hispanic motorists.

Related Literature.—Our paper relates closely to a series of papers testing whether racial differences in search rates reflect racial bias or are driven by some combination of statistical discrimination and omitted variables. Research in this literature applies the reasoning of the Becker (1957, 1993) outcome test. Implementing this test in practice is made difficult by the inframarginality problem.

Two seminal papers develop tests of racial bias that attempt to circumvent the inframarginality problem. Both papers document large racial disparities in search rates and similar or lower hit rates for black and Hispanic motorists, but conclude there is no evidence of racial bias against black motorists. Knowles et al. (2001) develop an equilibrium model where troopers decide whether or not to search motorists and motorists decide whether or not to carry contraband. They show that if troopers are not racially biased, all motorists must carry contraband with equal probability in equilibrium. In their model there is no inframarginality problem because there is no difference between hit rates for the marginal and average searched motorists. While we find that marginal and average hit rates are similar empirically, there are at least two features of our setting that are inconsistent with the Knowles et al. (2001) framework. First, troopers vary systematically in their hit rates, implying variation in screening ability. Second, we find little evidence that motorists respond to variation in search risk by adjusting contraband carrying rates, at least in the range of search rates we observe.

Anwar and Fang (2006) argue that if troopers are not racially biased, the ranking of search and hit rates by white and black troopers should be unaffected by the motorist's race. In their framework, if rankings differ by the motorist's race, then either white or black troopers are biased (or both), though the approach cannot be used to detect absolute racial bias. Applying the Anwar and Fang (2006) test to our data, we do not find evidence of relative bias among black, white, and Hispanic troopers.[1]

More recently, Arnold et al. (2018) and Marx (2020) use variation in behavior among decision-makers to address inframarginality. Arnold et al. (2018) use the quasi-random assignment of defendants to bail judges with varying release tendencies to test whether the rates of pretrial misconduct differ for marginal black and white defendants. A key assumption for their approach is *strict mono-*

[1] Simoiu et al. (2017) address the inframarginality problem by jointly estimating decision thresholds and risk distributions parametrically.

tonicity—that bail judges share the same ranking of defendants by misconduct risk. Under strict monotonicity—an assumption that does not hold in our setting—our results would imply that Texas state troopers exhibit racial bias against Hispanic, but not black, motorists in the Becker (1957, 1993) sense, yet troopers search black motorists more often than white motorists with no associated efficiency gains. This highlights a limitation of the Becker (1957, 1993) outcome test: when the returns to search are constant, equalized marginal hit rates do not imply an equity-efficiency trade-off.[2] Marx (2020) applies a logic similar to that of Arnold et al. (2018) to the policing context. He uses search and hit rates for each combination of motorist race and trooper race to bound marginal hit rates for each motorist racial group. He finds suggestive evidence of bias against Hispanic motorists. We require a weaker condition than strict monotonicity to construct our counterfactuals: that the between-trooper SPC and average within-trooper SPC are equal. This condition is implied by the skill-independence condition of Chan et al. (2020) and the average monotonicity condition of Frandsen et al. (2020).

Methodologically, we build on Arnold et al. (2020) and Chan et al. (2020), who do not assume monotonicity and identify variation in both preferences and screening ability across decision-makers facing similar cases. Arnold et al. (2020) use the quasi-random assignment of defendants to judges to study racial disparities and discrimination in pretrial release rates in New York City courts. They first show how to non-parametrically identify racial discrimination, which they define as racial differences in pretrial release rates conditional on pretrial misconduct risk. The authors then use the relationship between judge release rates and their share of released defendants who commit pretrial misconduct to estimate a structural model of judge behavior. They conclude that discrimination is driven by both racial bias and statistical discrimination and judges vary in both degree of bias and screening ability. Chan et al. (2020) exploit the quasi-random assignment of patients to physicians to study the relationship between a physician's share of patients diagnosed with pneumonia and their share of patients who leave with undiagnosed pneumonia. They estimate a structural model to recover the distribution of diagnostic skill and preferences across physicians. A common feature of these two papers and ours is that they use *between-agent* variation in behavior to make inferences about policy counterfactuals that require *within-agent* changes in behavior. A key distinction is that we do not aim to identify the distribution of trooper preferences and screening ability. Instead, we are focused on whether racial profiling poses an equity-efficiency trade-off in the aggregate.

A second branch of the economics literature on racial profiling considers whether profiling is justified, either on efficiency or ethical grounds. Several papers make the point that, even if troopers are profiling as a form of unbiased statistical discrimination to maximize the rate at which their searches yield contraband, profiling may not be efficient if it affects whether motorists decide to carry contraband in the first place (Persico, 2002; Dominitz and Knowles, 2006). If the social goal of policing is to reduce crime (rather than maximize the rate at which contraband is found), then

[2]Our findings are consistent with MacDonald and Fagan (2019) who study a New York Police Department policy that increased search and frisk rates during civilian stops in specific locations, particularly of black and Hispanic civilians. They find that this increase in search and frisk rates did not correspond with a change in the hit rate.

efficient trooper behavior requires that they search motorist groups such that marginal deterrence effects are equalized across groups. We provide suggestive evidence that deterrence effects are limited in our setting, and so efficiency can be assessed on the basis of contraband yield.

More broadly, our paper contributes to a large literature examining racial disparities in the criminal justice system. Blacks and Hispanics are charged with more severe offenses (Rehavi and Starr, 2014; Goncalves and Mello, 2018), are more likely to be denied bail (Arnold et al., 2018) and convicted of a crime (Anwar et al., 2012), and are issued longer incarceration sentences (Rehavi and Starr, 2014). Racial disparities in search rates are an order of magnitude larger than those documented in other criminal justice contexts. Traffic stops are also the most common interaction between police and the public, accounting for 41% of police contacts in 2015 (Davis et al., 2018).

II Context and Data

We study the search behavior of highway patrol troopers. In Texas, the primary responsibility of highway patrol troopers is to enforce state traffic laws on highways and state roads, but they have authority to enforce state criminal law throughout the state. During a typical shift, troopers conduct an average of eight traffic stops. When conducting each traffic stop, a trooper will give either a warning or citation for the original traffic violation. Troopers may also decide to further investigate if they suspect that a motorist may be carrying contraband, such as illicit drugs or weapons. As part of their investigation, troopers may search the motorist or vehicle for contraband. Below, we describe the legal standard pertaining to the use of race as a factor in trooper decisions to search a motorist or vehicle. We then describe the combination of three datasets that we use to characterize patterns in search rates and outcomes by motorist racial group: (1) administrative data on traffic stops conducted by the Texas Highway Patrol, (2) administrative data on individual criminal histories in Texas, and (3) commercial address history data.

II.A Legal Framework

Whether police officers can legally use race as a factor in deciding to engage in routine activities, such as vehicle and motorist searches, remains controversial. In an early review of the relevant case law, Knowles et al. (2001) concludes that the legality of racial profiling is complex and context-specific. Legal scholars have also noted that constitutional challenges to racial profiling have largely been unsuccessful, often requiring plaintiffs to show evidence of "discriminatory purpose" (i.e., racial animus) underpinning the profiling behavior being challenged. An alternative avenue for redress is offered by Legal Code 34 U.S.C. § 12601 (Section 12601), which authorizes the Department of Justice to pursue cases against police departments engaged in unconstitutional practices. Indeed, the Department of Justice has historically taken action against a number of police departments for racially-targeted stops of pedestrians and motorists on this basis (Anderson, 2020), although documented transgressions in these departments were particularly egregious, including false arrests, illegal searches, and excessive use of force. In any case, Section 12601 cannot be used by private individuals seeking legal remedy for mistreatment, and its use by the Department of Justice is discretionary; since 2017, no new Section 12601 investigations have been initiated.

Recent legal scholarship, including Tiwara (2019), has challenged the legality of racial profiling on the basis of disparate impact liability arising under the Omnibus Crime Control and Safe Streets Act of 1968. Under this framework, a practice that has a disparate impact on minorities "may be permissible only if the police can demonstrate that it has a legitimate law-enforcement-related necessity for the use of the practice at issue" (Tiwara, 2019). In our context, evidence that state troopers search black and Hispanic motorists more frequently than white motorists without any associated efficiency gain would likely constitute a discriminatory practice on this basis.

II.B Traffic Stop Data

The primary data source we use is a comprehensive dataset of 16 million traffic stops of motor vehicles conducted by the Texas Highway Patrol between 2009 and 2015. For each stop, the data include the date, time, location (including GPS coordinates), motorist's race and ethnicity, motorist's gender, information on the motor vehicle (including make, model, and year), the associated violation(s)[3], whether a search was conducted, the rationale for each search, whether contraband was found, and the ID number of the trooper who conducted the stop. These data are similar to those used in earlier studies of racial profiling (see Anwar and Fang, 2006). A unique feature of the Texas data is that they include the motorist's full name and address. This identifying information allows us to augment the data in three ways: (1) we match multiple traffic stops to the same motorist, (2) we merge in criminal histories for each motorist using data described below, and (3) we use each motorist's address to identify their neighborhood (Census block group) median income.

The data only cover motorists who are stopped and not all motorists who could potentially be stopped. This constraint will be particularly relevant when we study variation across troopers in their search behavior because troopers may also vary in whom they decide to stop. If the composition of stopped motorists varies across troopers in ways that we cannot observe in the data, this will complicate our interpretation of between-trooper differences in search behavior. Due to this concern, we focus our analysis on what Epp et al. (2014) and Baumgartner et al. (2018) classify as *safety stops*, which they distinguish from *investigatory stops*. The goal of their classification is to distinguish stops by the trooper's motivation for the stop. In safety stops the trooper's motivation for conducting the stop is the traffic violation itself and not the characteristics of the motorist or vehicle. By contrast, in investigatory stops, troopers use minor traffic offenses as a pretext for pulling motorists over and potentially searching them or their vehicles. Troopers use more discretion in deciding whether to conduct an investigatory stop, and hence there is more potential for motorist characteristics to vary across troopers for these stops.

Our data do not identify the trooper's reason for the stop directly, but we infer this from the traffic or criminal violation(s) associated with the stop. To proxy for safety stops, we limit the

[3]The data include both stops that result in warnings and citations.

sample to stops that include a speeding violation.[4] This includes 61% of all stops. Consistent with our interpretation of speeding stops as predominantly safety stops, we show in Appendix B that variation across troopers in cited speeds is limited. In addition, in Section IV we measure and account for variation in the composition of stopped motorists across troopers.

We also limit our analysis to stops where the motorist has a Texas address and where the motorist is black, Hispanic, or white. A prior investigation found that Texas state troopers incorrectly recorded many Hispanic drivers as white (Collister, 2015). Following Pierson et al. (2019), we recategorize motorists as Hispanic if they have a surname such that at least 75% of people with that surname identify as Hispanic in the 2010 Census.[5]

Finally, to reduce variation in stop location across troopers, we limit our analysis to stops made on state and interstate highways. This restriction excludes stops made on farm-to-market roads, ranch-to-market roads, county roads, and city streets, which account for about 26% of stops but have far fewer stops per road than state and interstate highways. Appendix Table A1 summarizes the number of observations we drop with each sample restriction. After applying these restrictions, our sample includes 4,931,332 stops.

We divide motorists into four categories based on their history of previous traffic stops.[6] We assign all motorists who have not had a previous stop to the first category. Motorists with a prior stop but no prior search are assigned to the second category. Motorists with a prior search but no prior search that yielded contraband are assigned to the third category. Motorists with a prior search that yielded contraband are assigned to the fourth category.

In our setting, there are four types of searches: consent, probable cause, incident to arrest, and inventory. Inventory searches are searches that occur after a vehicle is ordered impounded. In these instances, police are free to search the inventoried vehicle subject to departmental search policy. Incident to arrest searches are searches that occur *following* an arrest. After an arrest, troopers can search the arrested individual for contraband and, under broad conditions, search the vehicle. Alternatively, troopers have the right to conduct a search if they have probable cause to believe a law has been broken. Last, in a consent search, a trooper conducts a search only after receiving permission from the motorist to do so. In our data, 84% of searches are consent and probable cause searches. When contraband is discovered following a search, the motorist may be arrested on charges related to the contraband discovered.

[4] To proxy for safety stops, Baumgartner et al. (2018) use stops where the trooper's stated reason for the stop is a speed limit violation, stop light or sign violation, other moving violation, or driving while intoxicated. They classify stops as investigatory if the stated reason is an equipment violation (e.g., broken taillight), a regulatory or paperwork violation (e.g., expired registration or license), seat belt violation, or if the stop results from a trooper looking for a particular individual (e.g., as part of a criminal investigation).

[5] For the subsample of motorists with arrest records, the correlation between this constructed measure of Hispanic ethnicity and the measure included in Texas administrative criminal history data is 0.74 (0.75 for males and 0.70 for females).

[6] Stop histories are constructed using all stops, not just those meeting our sample criteria.

II.C Criminal History Data

We construct motorist criminal histories using data from the Texas Computerized Criminal History System. These data are maintained by the Texas Department of Public Safety. State troopers have access to these same data when conducting stops. The data track state felony and misdemeanor criminal charges from arrest through sentencing up to 2015.[7] Agencies are required to report data for all offenses that are Class B misdemeanors or greater, including all offenses that would potentially lead to a confinement sentence. The data include information on each criminal charge, including the original arrest charge, date of arrest, final court charge, charge disposition, and, if the charge results in conviction, the final sentence. The data include arrest charges that are ultimately dropped. We use these data to create summary measures of each motorist's criminal history at the time they are stopped. The data also include an individual's full name, address, race and ethnicity, gender, and a unique individual ID.

We construct two criminal history indices, one based on all drug offense arrests and another for non-drug offense arrests. For the drug offense index, we divide motorists into three categories. The first category is motorists with no prior drug-related arrests. Among motorists with any prior drug-related arrest, the median number of prior drug-related arrests is one. We assign remaining motorists to the second and third criminal history categories if their number of prior drug-related arrests is one and greater than one, respectively. We construct an analogous index for non-drug offense arrests. Among motorists with any prior non-drug offense arrest, the median number of prior non-drug offense arrests is two.

II.D Address History Data

One shortcoming of the traffic stop data is that it does not include a unique motorist ID. The problem this presents is that for two traffic stops with the same associated motorist name but different addresses, we do not know whether these stops correspond to the same person. The criminal history data includes an individual identifier and allows us to construct a partial address history for a given person. But the addresses we observe in those data only correspond to the points in time when that person is arrested, if they have any criminal history at all.

To facilitate matching traffic stops and criminal history to a given motorist, we use commercial data on address history from Infogroup. These data are similar to address history data used in prior research, including Diamond et al. (2019) and Phillips (forthcoming). For each individual, the data include their full name and street addresses at which the individual lived with estimated dates of residence. The data extract we use includes the addresses histories for all individuals in the database with a Texas residence between 2005 and 2016.

We merge traffic stops and criminal history to individuals using full name and address, incorporating address history data to account for address changes. The data merge is described in more detail in Appendix A.

[7]The share of arrest records with an available sentencing record falls after 2013.

II.E Descriptive Statistics

We present descriptive statistics for our merged dataset in Table I. We report descriptive statistics for all stops and subset the data by motorist race. We do the same for all stops that lead to searches. The motorist is female in 36% of stops, white in 55% of stops, Hispanic in 35.9% of stops, and black in 9.1% of stops.[8] For about 43% of stops, the motorist has been stopped previously. For about 1.2% of stops, the motorist has been stopped and searched previously, and for about 30% of those stops, the motorist has also been found with contraband in a previous search. For about 9% of stops, the motorist has a previous non-drug arrest, and in about 3% of stops the motorist has a previous drug arrest. Troopers search motorists in 1.06% of stops and find contraband in 0.34% of stops.

[Table 1 about here.]

Black motorists are slightly less likely than white motorists to have been stopped in the past, but they are more likely to have been searched in the past. They are also more likely to have an arrest history. Consistent with past research on racial profiling (see Pierson et al., 2019), black motorists are nearly three times more likely to be searched than white motorists (corresponding search rates are 2.202% and 0.755%). For Hispanic motorists, stop history, criminal history, and search rates generally fall between white and black motorists. Hispanic and black motorists reside in neighborhoods with similar median incomes, while median neighborhood incomes for white motorists are higher.

Compared to all stopped motorists, searched motorists are about 18 percentage points more likely to be male and come from neighborhoods with median incomes that are 13 log points lower. Searched motorists are more than seven times more likely to have been searched previously, three times more likely to have a previous arrest unrelated to drugs, and six times more likely to have a previous drug-related arrest.

In Appendix B we summarize the joint determinants of search in a series of logistic regressions. Conditional on stop location and time, black and Hispanic motorists are about 170% and 68% more likely to be searched than white motorists. Conditioning further on motorist neighborhood income, expected neighborhood income given vehicle type, criminal history, and stop history reduces black-white and Hispanic-white odds ratios to 1.86 and 1.44. This rich set of controls can statistically account for about half of the black-white and 35% of the Hispanic-white disparities in search rates we estimate by conditioning on only stop time and location. The residual black-white difference in search rates is similar in magnitude to the increase in search likelihood associated with multiple previous non-drug arrests and half of the increase associated with a prior drug arrest.

The percentage of searches that yield contraband (the hit rate) is 31.9%. The hit rate for white motorists (37.4%) exceeds the hit rate for black motorists (34.0%), which exceeds the hit rate for Hispanic motorists (25.9%). This ranking is consistent with past research on racial profiling

[8] For comparison, in 2010 the age 15 and above Texas population was 51% female, 49% non-Hispanic white, 34% Hispanic, and 12% non-Hispanic black.

(see Pierson et al., 2019). In Appendix B we summarize the joint determinants of contraband yield among searches. Conditional on stop location and time, searches of black and Hispanic motorists are about 15% and 30% less likely to yield contraband than searches of white motorists. Conditioning further on both motorist income proxies, criminal history, and stop history attenuates these differences to about 10% and 25%.[9]

Appendix Table B1 describes the distributions search types, contraband types, and arrest outcomes. Black motorists are more likely to be subject to probable cause searches, and less likely to be subject to consent and inventory searches. Drugs make up 51.8% of contraband found, weapons make up 3.8%, and currency makes up 0.6%. In the remaining 44% of cases, contraband is characterized as "Other", a category that includes drug paraphernalia and open containers of alcohol. Across motorist racial groups, the most salient difference is that black motorists are about four and two percentage points more likely to be found with drugs and weapons than the pooled average, and are less likely to be found with "Other" contraband. We find that only 24.5% of stops that yield contraband lead to an arrest, similar to the rate documented in North Carolina by Baumgartner et al. (2018). This percentage is similar across motorist racial groups. The severity of arrest charges, as proxied by the average incarceration sentence associated with conviction, is also similar across groups.

III A Model of Trooper Search Behavior

We model trooper search behavior using a modified version of the model developed in Anwar and Fang (2006). Troopers decide whether to search a stopped motorist using a noisy signal for the motorist's guilt. The modification we make is to allow this signal to be coarse over some range so that troopers are unable to distinguish between more or less suspicious motorists in this range.

We begin with a continuum of motorists. We first consider the behavior of a single trooper and later consider heterogeneity in trooper preferences and screening ability. Suppose fraction π of motorists carry contraband. For each stopped motorist i, the trooper observes a noisy signal for the motorist's guilt, $\theta_i \in [0, 1]$. If the motorist is carrying contraband, the index θ is randomly drawn from a distribution with continuous probability density function (PDF) $f_g(\cdot)$ and cumulative distribution function (CDF) $F_g(\cdot)$; if the motorist is not carrying contraband, θ is randomly drawn from a continuous PDF $f_n(\cdot)$ and CDF $F_n(\cdot)$. (The subscripts g and n stand for "guilty" and "not guilty," respectively.)

We assume that $f_g(\cdot)$ and $f_n(\cdot)$ satisfy a modified version of the standard monotone likelihood ratio property (MLRP): $f_g(\theta)/f_n(\theta)$ is strictly increasing in θ for $\theta < \bar{\theta}$ and is *constant* for $\theta \geq \bar{\theta}$. The MLRP assumption on the signal distributions provides that a higher index θ signals that a motorist is more likely to be guilty. However, in our formulation, for sufficiently "suspicious" signals, there is a point at which signals are no longer informative at the margin about a motorist's likelihood of carrying contraband. In other words, signals are coarse in the sense that troopers can

[9] Interestingly, we also find that, conditional on motorist race, low-income motorists are more likely to be searched, but searches of low-income motorists are less likely to yield contraband.

identify the riskiest motorists but, within this group, are unable to distinguish between those who are more or less likely to carry contraband.[10]

Let G denote the event that a motorist is found with contraband if searched. When a trooper observes a motorist with signal θ, the posterior probability that the motorist is guilty of carrying contraband, $Pr(G|\theta)$, is given by Bayes's rule:

$$P(G|\theta) = \frac{\pi f_g(\theta)}{\pi f_g(\theta) + (1-\pi) f_n(\theta)}.$$

From the MLRP, we have that $P(G|\theta)$ is strictly increasing in θ for $\theta < \bar{\theta}$. For $\theta \geq \bar{\theta}$, this probability is constant and is given by

$$P(G|\theta \geq \bar{\theta}) = \frac{\pi f_g\left(\bar{\theta}\right)}{\pi f_g\left(\bar{\theta}\right) + (1-\pi) f_n\left(\bar{\theta}\right)}.$$

Following the literature, we assume that the trooper's objective is to maximize the rate that traffic stops yield contraband, net of search costs. We further assume that search costs are a convex function, $C(\cdot)$, of the overall proportion of stops that result in searches, σ.

Given this cost structure, troopers will choose some threshold θ^* where troopers will search any motorist with $\theta_i \geq \theta^*$. Given this search threshold, the overall search rate is

$$\sigma(\theta^*) = \pi(1 - F_g(\theta^*)) + (1-\pi)(1 - F_n(\theta^*)).$$

The trooper's problem is to choose θ^* that maximizes their objective function

$$\int_{\theta^*}^{1} P(G|\theta^*) f(\theta) d\theta - C(\sigma(\theta^*)),$$

where $f(\theta) = \pi f_g(\theta) + (1-\pi) f_n(\theta)$. Hence, the trooper will set a threshold θ^* to equalize the marginal costs and benefits of search for the marginal searched motorist:

$$P(G|\theta^*) = C'(\sigma(\theta^*)).$$

Given search threshold θ^*, the trooper's unconditional hit rate is

$$\eta(\theta^*) = \pi(1 - F_g(\theta^*)).$$

We define the contraband *yield* rate (or *hit rate*) as $\frac{\eta(\theta^*)}{\sigma(\theta^*)}$, the share of searches that yield contraband.

We denote the relationship between $\eta(\theta^*)$ and $\sigma(\theta^*)$ as the trooper's SPC. Equivalently, we define SPC implicitly as $\tilde{\eta}(\sigma) = \eta(\sigma(\theta^*))$. We present a theoretical example of this SPC in Figure I. By the MLRP, this relationship is linear where $\theta^* \geq \bar{\theta}$, and hence $\sigma(\theta^*)$ is low. As θ^* declines

[10] We allow for one region for the most suspicious motorists where $f_g(\theta)/f_n(\theta)$ is constant because this fits the pattern we observe in the data. The framework can be readily extended to allow for alternative locations of this "flat" region or multiple such regions.

below $\bar{\theta}$, the relationship becomes concave, as the marginal searched motorist is less likely to have contraband than inframarginal searched motorists.

[Figure 1 about here.]

III.A Trooper Heterogeneity

In practice, we will estimate feasible combinations of search rates and unconditional hit rates using variation in outcomes across troopers. More formally, let p index troopers. We will identify the set of outcomes for all troopers, $\{(\sigma_p^*, \eta_p^*)\}_{p \in P}$. We use this set to calculate the *between-trooper* SPC, the conditional expectation function for η_p^* given σ_p^*, which can be expressed as

$$\tilde{\eta}^{\mathrm{BT}}(\sigma) \equiv E[\eta_p^* | p \text{ s.t. } \sigma_p^* = \sigma].$$

This may differ from the SPC that an individual trooper faces if trooper-specific SPCs—the set of feasible outcomes for a specific trooper—are heterogeneous.

In our setting, troopers may vary in their search rates and unconditional hit rates because they face different search costs, $C_p(\cdot)$, which would lead to varying search thresholds, θ_p^*. Troopers may also vary in their ability to infer the contraband risk for each motorist in the sense that the signal distributions $f_g(\cdot)$ and $f_n(\cdot)$ may vary across troopers. In this case, troopers may vary in the unconditional hit rates they can achieve for a given search rate, leading to variation in trooper-specific SPCs. A uniformly higher SPC—meaning a trooper can achieve a (weakly) higher hit rate for every given search rate—corresponds to greater screening ability.

If troopers vary only in search costs, SPCs will not vary across troopers, and the between-trooper SPC we identify will correspond to each trooper's own SPC. This condition follows from the strict monotonicity assumption of Arnold et al. (2018). But if troopers vary in screening ability, SPCs will vary across troopers, and the between-trooper SPC we identify may no longer correspond to any particular trooper's SPC.

If trooper-specific SPCs vary, we can still define the *average within-trooper* SPC. We define the average within-trooper SPC as the average of trooper-specific SPCs,

$$\tilde{\eta}^{\mathrm{WT}}(\sigma) \equiv E_p[\tilde{\eta}_p(\sigma)].$$

For any given search rate σ, the value for the average within-trooper SPC is the average unconditional hit rate troopers would achieve if all troopers search at that rate. The between-trooper SPC we identify will correspond to the average within-trooper SPC if variation in trooper screening ability is independent of trooper search rates, σ_p^*. More formally, suppose there exists a function that assigns a skill α_p to each trooper j such that $\tilde{\eta}_p(\sigma) = \tilde{\eta}_{p'}(\sigma)$ for all search rates σ where $\alpha_p = \alpha_{p'}$. Then the between-trooper SPC identifies the average within-trooper SPC if α_p is independent of σ_p^*. This condition corresponds to the skill-propensity independence condition in Chan et al. (2020). The condition is weaker than the strict monotonicity assumption of Arnold et al. (2018), which would require that any motorist searched by a given trooper would have also been searched by any trooper with a higher search propensity, and any motorist not searched by a given trooper would not have been searched by any trooper with a lower search propensity.

III.B Disparities between Motorist Groups

The focus of this paper is on racial disparities in search rates and whether equalizing search rates across motorist racial groups would reduce contraband yield. Accordingly, we extend the model to allow for multiple motorist groups (e.g., black, Hispanic, and white motorists). In particular, we index groups by $r \in \{b, h, w\}$ and allow for group-specific signal distributions ($f_g^r(\cdot)$ and $f_n^r(\cdot)$) and search thresholds (θ_r^*), which imply group-specific SPCs. We also allow for search costs to depend on the search rates for each group so that costs are defined as

$$C(\sigma^b(\theta_b^*), \sigma^h(\theta_h^*), \sigma^w(\theta_w^*)).$$

By characterizing group-specific SPCs and identifying where troopers locate along those group-specific SPCs, we can determine whether troopers face an equity-efficiency trade-off.

There are two scenarios where no trade-off exists. In the first scenario, search productivity at the margin is unequal across motorist groups, and marginal productivity is lower for the group with the higher search rate. This corresponds to Panel B of Figure II. In this case, the Becker (1957, 1993) outcome test would identify troopers as biased.

In the second scenario, search productivity at the margin is equalized across groups, but $\theta_r^* \geq \bar{\theta}_r$ for $r \in \{A, B\}$, and search rates are unequal across groups. This scenario is depicted in Panel C of Figure II. Note that, in this scenario, troopers are unbiased in the sense of Becker (1957, 1993).[11] For comparison, Panel A of Figure II depicts a scenario in which an equity-efficiency trade-off is present because equalizing marginal hit rates requires unequal search rates.

[Figure 2 about here.]

IV Estimating the between-Trooper Search Productivity Curve

We have shown that black and Hispanic motorists are more likely to be searched than white motorists, while searches of black and Hispanic troopers are equally or less likely to yield contraband. The central question of this paper is whether equalizing search rates across motorist racial groups would decrease contraband yield. To answer this question, we first estimate the relationship between trooper search rates and unconditional hit rates. We present evidence that different troopers search equivalent groups of motorists at varying rates and examine how troopers' search rates relate to their search productivity.

[11] An alternative notion of bias is based on whether search rates are equal *among motorists with* $\theta_r^* \geq \bar{\theta}_r$. We are unable to test for this form of bias, however, because we cannot measure the total number of motorists meeting this condition.

An essential requirement of our approach is that we measure how outcomes vary across troopers for equivalent stops. There is no random assignment of troopers to stops in our context. Instead, we will make a *selection on observables* argument—conditional on the time and location of the stop, the identity of the trooper who conducts the stop is unrelated to other determinants of search and search outcomes. In practice, cross-trooper variation arises from week-to-week variation in trooper shift schedules and within-shift variation in trooper locations. Our primary analysis relies on between-trooper variation within assigned patrol areas ("sergeant areas") to isolate variation in search rates conditional on location. We will bolster the argument that we are identifying how different troopers treat equivalent stops by showing that our SPC estimates are robust to varying the set of included controls and troopers and are corroborated by two alternative research designs that rely on different identifying assumptions.

IV.A Conceptualizing the between-Trooper Search Productivity Curve

For each stop, let i denote the motorist and t denote the specific time. The functions $\ell(i,t)$ and $\tau(t)$ map each stop to its associated location and time category. Let P_ℓ denote the set of troopers working in location ℓ. We limit the analysis to trooper-by-location combinations where the trooper has conducted stops in that location for each time category. For each stop, the associated trooper must decide whether to conduct a search. Let $\text{SEARCH}_{itp} \in \{0,1\}$ denote the potential (search) outcome of the stop, which indicates whether trooper $p \in P_{\ell(i,t)}$ would conduct a search if they were conducting stop (i,t). Let $G_{it} \in \{0,1\}$ indicate whether the motorist is carrying contraband at the time of the stop. Hence, trooper p would find contraband in stop (i,t) if $G_{it} \times \text{SEARCH}_{itp} \equiv \text{CONTRABAND}_{itp} = 1$.[12] Finally, the function $p(i,t)$ maps a stop (i,t) to the trooper who conducts the stop in practice.

For every trooper $p \in P_\ell$, we can define what their search rate and unconditional hit rate would be if they conducted all of the searches conducted in location ℓ:

$$\sigma_{p\ell} \equiv E[\text{SEARCH}_{itp} | \ell(i,t) = \ell], \tag{1}$$

$$\eta_{p\ell} \equiv E[G_{it}\text{SEARCH}_{itp} | \ell(i,t) = \ell]. \tag{2}$$

We refer to these objects as search *propensities* and unconditional hit *propensities*.

We define our between-trooper SPC as

$$\tilde{\eta}^{\text{BT}}(\sigma) \equiv E_\ell[E_p[\eta_{p\ell} | p \text{ s.t. } \sigma_{p\ell} = \sigma]]. \tag{3}$$

In words, the between-trooper SPC is the relationship between trooper search propensities and unconditional hit propensities across troopers within a location, averaging across locations.

We are also interested in the between-trooper SPC for specific racial groups of motorists (black, Hispanic, and white). Let $r(i)$ indicate the race of motorist i, where $r \in \{b,h,w\}$. We define $\sigma^r_{p\ell}$

[12] Note that we assume no variation across troopers in their ability to identify contraband during a search.

and $\eta^r_{p\ell}$ analogously as a trooper's search propensity and unconditional hit propensity for motorists from group r and $\tilde{\eta}^{\text{BT}}_r$ as the between-trooper SPC for motorists from group r.

In practice, we do not observe search and unconditional hit propensities. Instead, for trooper p, we only observe the stop outcomes for stops conducted by trooper p in practice. To recover propensities, we rely on the following conditional independence assumption:

Conditional Independence (CI) Assumption. *Conditional on location $\ell(i,t)$ and time category $\tau(t)$, the race $r(i)$, guilt G_{it}, and potential search decisions $\{SEARCH_{itp}\}_{p\in P_{\ell(i,t)}}$ are independent of the trooper associated with the stop $p(i,t)$.*

We assess the plausibility of this assumption in Section IV.C. Let $S^r_{p\ell\tau}$ denote the set of stops conducted by trooper p in location ℓ at time category t of motorists from group r. Under this assumption, we can construct estimates for $\sigma^r_{p\ell}$ and $\eta^r_{p\ell}$ using the following weighted averages of observed trooper search rates and unconditional hit rates, $s^r_{p\ell}$ and $h^r_{p\ell}$:

$$s^r_{p\ell} \equiv \sum_\tau \underbrace{\left(\frac{1}{|S^r_{p\ell\tau}|}\sum_{(i,t)\in S^r_{p\ell\tau}} \text{SEARCH}_{it}\right)}_{p-\ell-\tau-r \text{ search rate}} \times \underbrace{\left(\frac{|\{(i,t)|\ell(i,t)=\ell;\tau(t)=\tau;r(i)=r\}|}{|\{(i,t)|\ell(i,t)=\ell;r(i)=r\}|}\right)}_{\tau \text{ share for } \ell - r}, \tag{4}$$

$$h^r_{p\ell} \equiv \sum_\tau \underbrace{\left(\frac{1}{|S^r_{p\ell\tau}|}\sum_{(i,t)\in S^r_{p\ell\tau}} \text{CONTRABAND}_{it}\right)}_{p-\ell-\tau-r \text{ unconditional hit rate}} \times \underbrace{\left(\frac{|\{(i,t)|\ell(i,t)=\ell;\tau(t)=\tau;r(i)=r\}|}{|\{(i,t)|\ell(i,t)=\ell;r(i)=r\}|}\right)}_{\tau \text{ share for } \ell - r}. \tag{5}$$

IV.B Measuring between-Trooper Variation in Search Propensities

We begin by documenting substantial between-trooper variation in search rates among stops with similar locations and times. The notion of location we use at baseline is the sergeant area. The Texas Highway Patrol Division defines six primary regions, which encompass a total of 21 districts and 157 sergeant areas. Outside of the state's most populous areas, sergeant areas typically cover one to two counties in their entirety. In contrast, there are multiple sergeant areas associated with each of the state's most populous counties. Our data identify the exact location of the stop, and we geocode the corresponding sergeant area using boundary shape files received in response to a Texas Public Information Act request.

We infer the sergeant area to which each trooper is assigned based on the trooper-specific distribution of stop locations. In a given calendar year (month), troopers conduct 85% (90%) of all stops in the same sergeant area, on average. Assigning troopers to the modal sergeant area in which they conduct stops in each year, we observe that roughly two-thirds of troopers are assigned to the same sergeant area during the entirety of the sample period. Within a given year, however, temporary reassignments are quite common. When sergeant area assignments are determined on a monthly basis, we calculate that over 70% of troopers experience a change in assignment at least once during the sample period. Roughly half of these reassignments are within the trooper's home region, with the remaining reassignments disproportionately concentrated in the set of districts adjacent to the US-Mexico border.[13]

The notion of time we use at baseline is the combination of quarter of day and whether the stop was conducted on a weekday or the weekend.

We apply additional sample restrictions that limit the analysis to troopers who have made a sufficient number of stops in a given location. For our pooled analysis, which pools motorists from all racial groups, we limit the analysis to trooper-by-location and time cells with at least 5 stops. We further limit to trooper and location cells with at least 100 stops. Panel A of Figure III depicts the number troopers meeting these criteria in each sergeant area. Finally, we limit to *locations* with at least 10 troopers meeting these criteria, leaving us with 1,951 troopers in 133 locations with 2,657 combinations of trooper and location accounting for 71% of stops. There are an average of 1,235 stops per trooper and location combination.[14]

[Figure 3 about here.]

Figure IV plots the distribution of $s_{p\ell}$ where each trooper-by-location combination is weighted equally. While the median trooper-by-location search rate is only 0.7%, there is a long right tail, indicating that a small number of troopers search at particularly high rates. The 90th percentile search rate is 3.4%.

[Figure 4 about here.]

We next look at search rates separately by motorist race. We denote the race-specific search rates by s^r_{pl} where $r \in \{w, b, h\}$. When examining race-specific search behavior, we apply different sample criteria to ensure that the troopers we include have made a sufficient number of stops for a specific motorist group. We limit the analysis to trooper-by-location-by-time-by-motorist-race cells with at least 5 stops and trooper-by-location-by-motorist-race cells with least 100 stops. We then limit to locations where, for each motorist racial group, there are at least 10 troopers meeting the sample criteria, leaving us with 1,289 troopers in 79 locations accounting for 58% of black motorist stops, 34% of Hispanic motorist stops, 53% of white motorist stops, and 46% of stops overall. The sample includes 1,064 troopers for black motorists, 1,238 for Hispanic motorists, and 1,280 for white motorists. For the sergeant areas we include in the race-specific analysis, Panel B of Figure III depicts the number of troopers who satisfy the sample criteria, averaging across motorist racial groups. There are an average of 266, 500, and 1,003 stops per trooper for black, Hispanic, and white motorists.[15]

[13] In an effort to strengthen border enforcement by increasing trooper presence, the Texas Department of Public Safety launched a "border surge" in 2014 that involved rotating troopers to the border region (Benning and Chavez, 2016).

[14] Descriptive statistics for the stops included in the pooled analysis are presented in Appendix Table B3.

[15] Descriptive statistics for the stops included in the race-specific analysis are presented in Appendix Table B4.

Mirroring racial differences in overall search rates, between-trooper variation in search rates is larger for non-white motorists. For white, black, and Hispanic motorists, the difference in search rates between the 10th and 90th percentiles of the trooper distribution is 2.8, 6.2, and 4.2 percentage points.

Our goal is to identify variation in how different troopers treat equivalent stops. However, for a fixed sergeant area and time category, the composition of stopped motorists may still vary across troopers. Sergeant area and time category may not fully capture variation in stop context. Even in the same environment, troopers may vary in the composition of motorists they decide to stop. For example, troopers may vary in whether they racially profile when deciding whom to stop for speeding (Horrace and Rohlin, 2016). We examine how search rates change when we condition on a larger set of stop and motorist characteristics. The objective is to learn whether a significant portion of the variation in search rates is due to differences in motorist composition and to isolate variation due to differences in trooper search behavior rather than differences in motorist composition.

To calculate search rates that adjust for differences in stop and motorist characteristics, we estimate the following linear probability model, separately by location:

$$\text{SEARCH}_{it} = \phi_{p(i,t)\ell(i,t)\tau(t)} + X_{it}\gamma + \delta_{m(t)} + \rho_{r(i,t)} + \epsilon_{it}, \tag{6}$$

where $\phi_{p(i,t)\ell(i,t)\tau(t)}$ are fixed effects for trooper-by-location and time combinations, $\delta_{m(t)}$ are fixed effects for the month of the stop, and $\rho_{r(i,t)}$ are fixed effects for the (highway) road of the stop. X_{it} is a vector of motorist characteristics, including race, gender, log of neighborhood median income, vehicle-based expected log neighborhood income, stop history, non-drug arrest history, and drug arrest history. We use this model to calculate search rates for each trooper-by-location-by-time combination, adjusting for motorist characteristics and stop month. We use these predicted search rates to construct an overall search rate for a trooper-by-location combination using the same weights as above. We denote this adjusted trooper search rate as $\tilde{s}_{p\ell}$.[16]

Appendix Figure B1 compares adjusted trooper search rates ($\tilde{s}_{p\ell}$) to unadjusted trooper search rates ($s_{p\ell}$) across trooper-by-location combinations after partialling out location fixed effects. The slope of the fitted line is 0.97, and the correlation is 0.99. Observable motorist characteristics explain virtually none of the variation in search rates across troopers. Instead, the variation is attributable to differences in trooper search behavior for observably similar stops. This finding provides support for our interpretation of trooper search rates as characterizing causal trooper search propensities.

[16]Formally, $\tilde{s}_{p\ell}$ is given by

$$\tilde{s}_{p\ell} = \sum_{\tau} \left(\hat{\phi}_{p\ell\tau} + E[X_{it}\hat{\gamma} + \hat{\delta}_{m(t)} + \hat{\rho}_{r(i,t)} | \ell(i,t) = \ell; \tau(t) = \tau] \right) P(\tau(t) = \tau | \ell(i,t) = \ell).$$

IV.C Trooper Search Rates and Motorist Characteristics

To further probe the CI assumption, we investigate the degree to which troopers with high and low search rates stop different types of motorists. Specifically, we examine how motorist characteristics predict the search rate of the trooper conducting the stop. As a benchmark, we estimate the analogous relationship between the same motorist characteristics and whether the stop leads to a search. In particular, we estimate linear regression models of the form

$$Y_{it} = \lambda_{\ell(i,t)} + \omega_{\tau(t)} + \delta_{m(t)} + X_{it}\gamma + \epsilon_{it}, \tag{7}$$

where Y_{it} is either SEARCH_{it} or leave-out trooper-by-location search rates, unadjusted ($s_{p\ell}^{-it}$) or adjusted ($\tilde{s}_{p\ell}^{-it}$). $\lambda_{\ell(i,t)}$ and $\omega_{\tau(t)}$ are location and time category fixed effects.

Table II shows that motorist characteristics predict trooper-by-location search rates, but the magnitude of the relationship is small.[17] Column (1) uses a linear probability model to examine how motorist characteristics predict whether a stop leads to search. Columns (2) and (3) use identical specifications to assess the extent to which these same motorist characteristics predict the unadjusted ($s_{p\ell}^{-it}$) and adjusted leave-out search rate ($\tilde{s}_{p\ell}^{-it}$) of the trooper conducting the stop. Where the outcome is the trooper search rate, the coefficients on all motorist characteristics are one to two orders of magnitude smaller. The relative magnitudes are particularly small for stop and arrest history, which would be difficult for troopers to observe before conducting the stop. Given the large size of our sample, many of these coefficients are statistically significant. In Section IV.D, we show that our SPC estimates are quantitatively similar whether or not we control for these characteristics directly. We also corroborate our baseline SPC estimates using two alternative research designs that rely on different identifying assumptions.

[Table 2 about here.]

As a robustness check, we repeat our main analyses after excluding troopers with the most selected set of stopped motorists. For varying κ, we remove the κ% of troopers with compositions of stopped motorists who deviate most from their expected composition given the time and location of their stops. We discuss how we identify these troopers in more detail in Appendix B. Columns (4)–(6) replicate (1)–(3) after removing stops conducted by the 20% of troopers with the most selected set of stopped motorists. In columns (5) and (6), which relate trooper search rates to stopped motorist characteristics, removing these troopers from the analysis further reduces the magnitude of coefficients on motorist characteristics. All results presented in the remainder of the paper are insensitive to excluding troopers with the most selected set of stopped motorists.

[17] Appendix Table B5 presents analogous estimates for trooper-by-location unconditional hit rates. The findings are similar.

IV.D Baseline Search Productivity Curve Estimates

For each location, we observe multiple troopers with varying search rates. We next calculate and compare unconditional hit rates across troopers within a location. We calculate trooper-by-location-specific unconditional hit rates analogous to the search rates constructed above, replacing the outcome with CONTRABAND_{it}. We denote the unadjusted and adjusted trooper-by-location-specific unconditional hit rates as $h_{p\ell}$ and $\tilde{h}_{p\ell}$. Each trooper demonstrates a feasible combination of search rate and unconditional hit rate, which we use to construct a between-trooper SPC. We then pool these location-specific SPCs across locations to construct an aggregate SPC.

Note that trooper-by-location-specific search rates and unconditional hit rates are, in principle, only directly comparable across troopers within a location.[18] Hence, we aggregate location-specific SPCs without relying on between-location comparisons. We do so in two ways. First, we plot the relationship between search rates and unconditional hit rates while adjusting for location fixed effects using the method of Cattaneo et al. (2019). We refer to this as the fixed effects (FE) approach. One problem with this is approach is that because the distribution of trooper search rates varies across locations, different portions of the SPC are estimated using varying sets of locations. In the second approach, within locations we divide troopers into quantiles by search rate, group quantiles across locations, and then plot the relationship between search rates and unconditional hit rates across quantiles. We refer to this as the quantile (Q) approach. The advantage of this approach is that each location is proportionally represented in each quantile. The interpretation of the slope of this relationship is the change in the aggregate unconditional hit rate associated with a change in the search rate quantile for all locations. We use deciles because each location has at least 10 troopers by construction.

Figure V summarizes the relationship between adjusted trooper search rates ($\tilde{s}_{p\ell}$) and unconditional hit rates ($\tilde{h}_{p\ell}$) using both approaches. In both approaches, we weight trooper-by-location combinations by number of stops. From the FE approach, the figure includes 95% confidence bands for the local linear relationship between adjusted trooper search rates and unconditional hit rates, the best linear fit in the depicted range, and the slope of the best linear fit, labeled as β^{FE}. From the Q approach, the figure includes the mean values for each decile, the best linear fit, and the slope of the best linear fit, labeled as β^{Q}.

[Figure 5 about here.]

There are two main findings to note. First, for both approaches the SPC is approximately linear. For the FE approach, we conduct a formal test for linearity following Cattaneo et al. (2019) and fail to reject the null hypothesis that the relationship is linear. Second, the two approaches provide similar SPC estimates: β^{FE} is 0.331 and β^{Q} is 0.346. These slopes indicate that a 1 percentage point increase in search rate is associated with about a 0.34 percentage point increase in the unconditional hit rate. Note that the ratio of the unconditional hit rate to the search

[18] In Sections IV.E.1 and IV.E.2 we show that trooper-by-location-specific search rates predict *between-location* differences in search behavior as well.

rate gives us the percentage of searches that yield contraband, the hit rate. The fact that the relationship between the unconditional hit rate and the search rate is linear implies that hit rates are approximately *constant* across quantiles of trooper search rates and hence are unrelated to trooper search rates.

In Figure VI we repeat the analysis separately by motorist racial group. Estimates in Panel A are derived using only white motorists, Panel B using only black motorists, and Panel C using only Hispanic motorists. There are two main findings to note. First, as above, each SPC is approximately linear. Second, the slopes for white and black motorists are comparable, while the slope for Hispanic motorists is smallest in magnitude. Using either the FE or Q approach, we cannot reject the null hypothesis of equal slopes for white and black motorists, while the slope for Hispanic motorists is significantly smaller than the slope for white motorists at the 1% level. Using the FE approach, the estimated SPC slopes for white, black, and Hispanic motorists are 0.395, 0.346, and 0.294. These slope estimates are comparable to overall group-specific hit rates described in Appendix Table B4. Using the Q approach, the estimated SPC slopes for white, black, and Hispanic motorists are 0.372, 0.388, and 0.312.[19]

[Figure 6 about here.]

IV.E Robustness Checks

The CI assumption underlying our approach is violated if, conditional on the sergeant area and time category of a stop, troopers vary in the composition of motorists they stop. This variation may exist because our location and time measures are not sufficiently granular or because, in the same environment, troopers vary in the motorists they decide to stop. In this section we present several robustness checks for the baseline results presented in Section IV.D.

In Section IV.E.1 we reestimate pooled and race-specific SPCs using only within-motorist variation in stop outcomes among motorists involved in multiple speeding stops. In Section IV.E.2 we compare stop outcomes on opposite sides of trooper patrol area borders in a spatial RD design to validate trooper search rates as estimates of causal trooper search propensities. In both sections we conduct additional tests for whether between-trooper SPCs are linear.

In Appendix Figure B4, we show that the slope of the pooled between-trooper SPC is stable if we exclude a varying proportion of troopers with compositions of stopped motorists who deviate most from their expected composition given the time and location of their stops. In Appendix Figure B5 we conduct a similar exercise for race-specific SPC slopes and find that slope estimates and their ordering across groups are stable when we vary the set of included troopers.

Another concern with our approach is that $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$, as estimates of their population analogs, $\sigma_{p\ell}$ and $\eta_{p\ell}$, are subject to correlated sampling error. This sampling error may bias our estimate

[19] Under the strict monotonicity assumption of Arnold et al. (2018), troopers' individual SPCs are the same as the between trooper SPCs we estimate. Under this assumption, the SPC estimates imply that Texas state troopers exhibit racial bias against Hispanic, but not black motorists. However, in Appendix B we document that strict monotonicity does not hold in our setting. We show that, conditional on their search rate, troopers vary in their hit rate, implying variation in screening ability across troopers.

of β. We account for sampling error in two ways in Appendix B. First, we apply empirical Bayes adjustments to $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$ (Morris, 1983). Second, we take a split-sample IV approach to estimation. We randomly split stops into two samples and estimate $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$ separately in each sample. In each sample, we regress $\tilde{h}_{p\ell}$ on $\tilde{s}_{p\ell}$ and location fixed effects, instrumenting for $\tilde{s}_{p\ell}$ using its pair estimate from the other sample. Reassuringly, both approaches yield β estimates that are statistically indistinguishable from the OLS estimates.

Finally, we verify that the key features of the pooled and race-specific SPCs are not sensitive to the particular hit rate definition employed. In our main specifications, we measure hit rates using an indicator for whether the trooper finds any contraband as recorded in the traffic stop data. This measure may mask heterogeneity in the significance of the contraband discovered across stops. In Appendix Figures B6 through B7 we replicate Figure V and Figure VI using an alternative outcome: an indicator for whether the contraband found leads to an arrest. The patterns are qualitatively similar. The SPCs are approximately linear, the slopes for white and black motorists are comparable, and the slope for Hispanic motorists is smaller in magnitude.

IV.E.1 Exploiting within-Motorist Variation

High and low search rate troopers may stop motorists who differ on unobservables correlated with contraband risk. To address this concern, we take advantage of the fact that we can match multiple stops to the same motorist. We look at sequential pairs of stops for the same motorist and measure the relationship between differences in stop outcomes and differences in the search behavior of the troopers conducting those stops. By looking at differences in stop outcomes for the same motorist, we net out time-invariant motorist-level determinants of search and contraband risk.

Consider a group of motorists stopped by two sets of troopers, one with high search costs and the other with low search costs. Suppose that the distribution of screening skill across troopers and the probability that a given motorist is carrying contraband are equal across sets. With diminishing returns to search, we expect the hit rate to be lower for the low search cost (and hence high search rate) set. Moreover, we expect this difference in hit rates to be increasing in the difference in search rates between the trooper sets. By contrast, if troopers are searching on the linear portion of the SPC, we expect constant hit rates.

To implement this idea, we first group sequential pairs of stops of the same motorist into deciles by their difference in trooper-by-location search rates, $\Delta_{it}\tilde{s}_{p\ell} = \tilde{s}_{p(i,t)\ell(i,t)} - \tilde{s}_{p(i,t')\ell(i,t')}$, where $t > t'$ and stops at t' and t are sequential for motorist i. Descriptive statistics for the sequential pairs of stops are presented in Appendix Table B7 and Appendix Table B8.

Panel A of Figure VII summarizes the characteristics of motorists involved in sequential pairs of stops, grouped into deciles on the horizontal axis based on their value of $\Delta_{it}\tilde{s}_{p\ell}$. The vertical axis depicts the average value of $P(SEARCH|X_{it})$, the predicted search rate for a motorist given their characteristics at the time of the initial stop.[20] The figure includes the best linear fit and a bin scatter. The measured relationship is flat. Motorists stopped by different troopers, as characterized by their search rates, do not markedly vary in their characteristics.

[Figure 7 about here.]

In Panel B of Figure VII we plot the relationship between $\Delta_{it}\tilde{s}_{p\ell}$ and Δ_{it}SEARCH. The relationship is linear with a slope of 0.733. If we suppose that Δ_{it}SEARCH is an unbiased measure of differences in trooper search propensities for motorist i, then the fact that this slope is less than one indicates that $\Delta_{it}\tilde{s}_{p\ell}$ exhibits forecast bias. This could be explained by one or a combination of two factors: (1) $\tilde{s}_{p\ell}$ is a biased estimate of $\sigma_{p\ell}$ or (2) trooper search propensities vary with motorist characteristics, and motorists stopped multiple times differ from typical motorists stopped in either location.[21] The overall search rate in the pooled SPC sample is 1.11% while the search rate for the subset of stops analyzed here is 0.93%.

Panel C of Figure VII plots the relationship between $\Delta_{it}\tilde{s}_{p\ell}$ and Δ_{it}CONTRABAND by decile. Again, the relationship is strikingly linear. If marginal motorists—motorists more likely to be searched by high search rate troopers—are less likely to carry contraband, we would expect the slope to be declining in $\Delta_{it}\tilde{s}_{p\ell}$. Instead, linearity is consistent with marginal motorists who are as likely to carry contraband as inframarginal motorists.

The patterns in Panels B and C of Figure VII imply a particular relationship between a given increase in search rates and an increase in the unconditional hit rate if we frame $\Delta_{it}\tilde{s}_{p\ell}$ as an instrument for Δ_{it}SEARCH. The slope we estimate is essentially the slope of an SPC in first differences. More concretely, we estimate the following model via just-identified two-stage least squares (2SLS), separately by motorist race:

$$\Delta_{it}\text{CONTRABAND} = \beta\Delta_{it}\text{SEARCH} + \epsilon_{it}, \tag{8}$$

where the first stage is

$$\Delta_{it}\text{SEARCH} = \pi\Delta_{it}\tilde{s}_{p\ell} + \zeta_{it}. \tag{9}$$

We repeat this exercise pooling all motorists and separately by motorist racial group. When we estimate the model for a specific motorist racial group, we replace $\Delta_{it}\tilde{s}_{p\ell}$ with its race-specific analog, $\Delta_{it}\tilde{s}^{r}_{p\ell}$.[22]

[20] We construct $P(SEARCH|X_{it})$ using the logistic regression model

$$P(\text{SEARCH}_{it} = 1|X_{it}) = \frac{e^{(X_{it}\beta)}}{1 + e^{(X_{it}\beta)}}$$

where X_{it} is a vector of motorist characteristics including motorist race, gender, log of neighborhood income, expected log income given vehicle, stop history, non-drug arrest history, and drug arrest history.

[21] A third possibility is that $\tilde{s}_{p\ell}$ differs from $\sigma_{p\ell}$ due to sampling error, leading to attenuation bias. However, the fact that empirical Bayes and split sample adjustments of $\tilde{s}_{p\ell}$ explored in Appendix B do not make a material difference indicates that sample sizes are sufficiently large for sampling error to not be a significant issue.

[22] Note that the set of stops included in the race-specific analysis is a subset of the stops included in the pooled analysis because $\tilde{s}^{r}_{p\ell}$ is measured for a smaller set of trooper-by-location combinations, as discussed in Section IV.B.

We report β estimates in Table III. For all motorists pooled and each motorist group, the slopes are somewhat smaller than, but comparable to, the baseline SPC slopes. As above, the SPC slope for white motorists exceeds the slopes for black and Hispanic motorists in magnitude. Time-invariant motorist unobservable characteristics cannot explain the pattern of SPC slopes we identify in Section IV.D.

[Table 3 about here.]

IV.E.2 Marginal Returns to Search at Sergeant Area Borders

We derive an alternative estimate for the marginal returns to search by comparing outcomes of stops conducted on opposite sides of sergeant area borders. Along the same highway route, the composition of troopers making traffic stops changes sharply across sergeant area borders, which designate the areas that troopers are assigned to patrol. If troopers assigned to one sergeant area search motorists at higher rates than troopers in a neighboring sergeant area, then motorists crossing from one sergeant area to the other will face sharp changes in their chances of being searched. This spatial feature of search rates suggests a natural RD research design. By comparing search and unconditional hit rates for speeding stops on either side of sergeant area borders, we measure the causal effect of changing from one set of troopers to another with higher search rates. This comparison provides another test for whether trooper search rates characterize the causal effect of trooper assignment on search likelihood and another measure of the between-trooper SPC slope.

Our identifying assumption is that the composition of motorists evolves continuously through sergeant area borders. This assumption is reasonable because sergeant area borders are defined only for administrative purposes; there is little reason to think the composition of motorists traveling on a given stretch of highway would change discontinuously at these boundaries. One possible exception is that some motorists are aware that their chances of being subject to a search change at sergeant area borders and adjust their travel or contraband carrying behavior accordingly. In practice, we find little evidence of a deterrence effect at sergeant area borders, a point we discuss in more detail in Section V.B.

For this exercise, we apply sample restrictions that differ from the sample restrictions described in Section IV.B.[23] There are 202 state and interstate highways crossing 319 distinct sergeant area borders. Each highway by border pair is a potential RD. We limit the RD analysis to highway and border intersections with at least 100 stops made in each corresponding sergeant area between 2 and 7 miles from the intersection. We are left with 424 intersections. These intersections as well as the corresponding highways and sergeant areas are shown in Appendix Figure B8.

For each intersection, we limit the analysis to speeding stops made in each corresponding sergeant area within seven miles from the intersection. We use the distance between the location

[23]Descriptive statistics for the stops included in this analysis are presented in Appendix Table B9.

of a stop and the intersection as the running variable. For each intersection, we set the distance as negative for the sergeant areas with the lower average trooper search rate.[24]

While sergeant area borders generate a discrete change in trooper patrol areas, in practice these borders do not provide a discontinuity in where troopers search. Troopers conduct some stops outside of their patrol area, and they are particularly likely to do so just outside sergeant area borders. Figure VIII pools all intersections and plots the share of stops conducted by troopers assigned to each adjacent sergeant area by distance from the intersection. More than two miles from the border, the share of stops conducted by troopers assigned to that corresponding sergeant area generally exceeds 70%. Approaching the border, this share falls to about 40%.

[Figure 8 about here.]

To add substantial statistical power to our test, we take a "donut" approach and exclude stops that occur within a two-mile window around the intersection, denoted by the dashed vertical lines in Figure VIII (Barreca et al., 2016). The trade-off is that, by excluding stops in this range, we can no longer take a non-parametric approach to identification. Instead we assume that, in the absence of contamination near the intersection, potential search outcomes would continue to evolve as they do outside of this range.[25] We use a bandwidth of seven miles, leaving us with 1,480,372 stops conducted between two and seven miles from the border.

Figure IX plots (leave-out) trooper search rates in Panel A and motorist characteristics in Panel B as a function of distance from the intersection. For each RD plot, within each set of stops corresponding to an intersection, we demean the outcome and then stack observations across intersections. Panel A shows that the search rates for troopers conducting stops are approximately constant across distances within each sergeant area. Across the boundary, extrapolated trooper search rates jump by 0.309 percentage points using a constant extrapolation and 0.287 percentage points using a linear extrapolation.

By contrast, as indicated in Panel B, the characteristics of stopped motorists vary only slightly across the threshold. Note that the statistical significance of this discontinuity does not necessarily indicate a discontinuity in the composition of motorists on either side of the border. Instead, it may simply reflect differences in trooper stop behavior.

[Figure 9 about here.]

Figure IX plots search rates (Panel C) and unconditional hit rates (Panel D) by stop location. In Panel C the pattern is noisier, but there is again a clear jump in extrapolated search rates at the boundary. The magnitude of the jump is 0.306 percentage points using a constant extrapolation

[24] We measure trooper search rates using all speeding stops conducted by a trooper, not just those made in the RD window or in a specific location.

[25] In contrast with typical applications of the donut RD design, we are not concerned about manipulation or error in the running variable (Barreca et al., 2016). Instead, we apply this approach because the change in treatment—in our case, the search rates of the troopers conducting the stops—at the border is muted by the fact that troopers make some stops just outside their assigned patrol area.

and 0.490 percentage points using a linear extrapolation. In the constant case, the magnitude of the jump is comparable to the magnitude of the corresponding change in trooper search rates. This result can be interpreted as a validation of trooper search rates as measures of causal trooper search propensities. In the linear case, the jump in search rates is larger, though the estimate is relatively imprecise.

The ratio of the increase in the unconditional hit rate to the increase in the search rate is 0.352 (with a standard error of 0.030) in the constant case and 0.333 (0.061) in the linear case.[26] This ratio provides an alternative estimate of the marginal slope of the between-trooper SPC. Both estimates are consistent with our SPC slope estimates of 0.331 and 0.346 in Section IV.D. They are also statistically indistinguishable from the overall hit rate, consistent with a linear between-trooper SPC.

V Policy Counterfactuals

We next use our SPC estimates to investigate how contraband yield would change if troopers equalized search rates across motorist racial groups. In Section V.A we describe the distinction between the between-trooper SPC and the *within-trooper* SPC and why the latter is central to constructing counterfactuals. We then provide suggestive evidence that the between-trooper and average within-trooper SPCs are similar. In Section V.B, we present evidence that motorist group-specific deterrence effects are negligible at the margin, implying that such deterrence effects are unlikely to substantively influence counterfactual conclusions. In Section V.C, we proceed to simulate and discuss counterfactuals.

V.A The within-Trooper Search Productivity Curve

We have documented the between-trooper relationship between search rates and unconditional hit rates. We find that this relationship is linear with racial group-specific slopes that suggest it is feasible for troopers to (1) search all motorist racial groups at the same rate, (2) maintain the status quo overall search rate, and (3) increase overall contraband yield. This counterfactual requires that individual troopers change their search behavior, so determining what would happen in this scenario requires knowing the SPCs faced by individual troopers. Yet, the *between-trooper* SPC need not be the same as *within-trooper* SPCs. In particular, a linear between-trooper SPC may still be consistent with troopers facing diminishing returns to search if troopers with more screening skill search at higher rates.[27]

[26] We estimate the ratios and associated standard errors by instrumenting for SEARCH_{it} using threshold crossing in a linear regression model for CONTRABAND_{it}.

[27] As we discuss in Section III.A, the distinction between the between-trooper SPC and within-trooper SPCs is related to the monotonicity conditions of Arnold et al. (2018), Chan et al. (2020), and Frandsen et al. (2020). Under the strict monotonicity condition of Arnold et al. (2018), each trooper faces the same SPC, which also corresponds to the between-trooper SPC. Strict monotonicity does not hold in our setting; in particular, we find systematic variation in hit rates between troopers. The between-trooper SPC and average within-trooper SPC are equal under a substantively weaker condition: independence of trooper screening ability and search propensity. This condition

Two recent papers closely related to ours, Arnold et al. (2020) and Chan et al. (2020), face a similar issue. Those papers identify cross-sectional variation in behavior across bail judges and doctors, respectively, and then conduct counterfactual exercises where they consider what would happen if decision-makers were made to change their behavior in some way. To infer how judge and doctor outcomes would change under counterfactuals, Arnold et al. (2020) and Chan et al. (2020) make parametric assumptions about the form of heterogeneity across agents and then use features of the cross-sectional distribution of agent behavior to identify parameters that characterize this heterogeneity.

By contrast, we aim to identify plausibly exogenous *within-trooper* variation in search behavior. Our ideal instrument would predict changes in a trooper's effective search costs that are orthogonal to changes in the composition of motorists they stop and their screening ability. We propose using changes in the search behavior of a trooper's coworkers as an instrument for that trooper's search behavior. The rationale is that troopers share managers (sergeants) and peers who influence their effective search costs. Our instrument captures shocks to search rates common across coworkers (as defined by sergeant areas). Unfortunately, we do not have a direct measure of manager influence or peer effects. Instead, we use this "black box" variation and show that it does not appear to be driven by variation in motorist composition. In this context, the exclusion restriction requires that changes in coworker search rates affect search outcomes only through changes in a given trooper's own search rate. One potential violation is that coworker search rates affect a trooper's search outcomes by influencing whether motorists carry contraband. However, in Section V.B we show that deterrence effects are likely to be negligible on the margin in our setting.

Let $\hat{\ell}(p, t)$ denote the assigned location of trooper p in the year corresponding to t. Let $s^{-p}_{\hat{\ell}(i,t)}$ denote the search rate of all troopers assigned to location ℓ in the year corresponding to t *excluding trooper* p.

We estimate a 2SLS system with first stage

$$\text{SEARCH}_{it} = \pi s^{-p}_{\hat{\ell}(i,t)} + X_{it}\gamma + \phi_{p(i,t)\ell(i,t)} + \delta_{y(t)} + \epsilon_{it}, \tag{10}$$

where $\phi_{p(i,t)\ell(i,t)}$ are trooper-by-location fixed effects and $\delta_{y(t)}$ are year fixed effects. The second stage is

$$\text{CONTRABAND}_{it} = \beta \text{SEARCH}_{it} + X_{it}\gamma + \phi_{p(i,t)\ell(i,t)} + \delta_{y(t)} + \zeta_{it}. \tag{11}$$

Table IV correlates the instrument with observable motorist characteristics and presents first stage and 2SLS estimates. In column (1) we estimate (10) but with $P(\text{SEARCH}_{it}|X_{it})$ as the outcome. The estimated coefficient on $s^{-p}_{\hat{\ell}(i,t)}$ is 0.012, indicating that a 1 percentage point increase in coworker search rates is associated with a 0.012 percentage point increase in $P(\text{SEARCH}_{it}|X_{it})$. This coefficient estimate is statistically significant but is small in magnitude. Columns (2) and (3) present first stage estimates with and without controls for motorist characteristics. The coefficient

corresponds to the skill-propensity independence condition in Chan et al. (2020) and is implied by the average monotonicity condition of Frandsen et al. (2020).

on $s^{-p}_{\hat{\ell}(i,t)}$ is 0.130 in column (2), indicating that a 1 percentage point increase in coworker search rates is associated with a 0.130 percentage increase in a trooper's own search rate. The coefficient is statistically significant, with a t-statistic of about 4.

Consistent with the limited relationship between coworker search rates and motorist characteristics shown in column (1), adding controls for motorist characteristics in column (3) only slightly attenuates this coefficient to 0.116. Columns (4) and (5) present IV estimates for β. Without controlling directly for motorist characteristics, the 2SLS estimate for β of 0.273 with a standard error of 0.087, in line with our between-trooper estimates. Adding motorist controls X_{it} has no material effect on this estimate. In Appendix Figure B9 we show that both the first stage and reduced form relationships are approximately linear.

[Table 4 about here.]

Overall, these results suggest that the marginal slope of the within-trooper SPC, averaged across troopers, is similar to the between-trooper SPC slope.[28]

V.B Deterrence Effects

To assess how contraband yield would change under counterfactual search rates, it is important to first gauge the expected responsiveness of contraband carrying behavior to motorist racial group-specific search intensity. While our analysis is not structured to explicitly characterize motorist racial group-specific deterrence effects, a number of factors suggest that changes in contraband carrying rates are likely to be negligible. First, if drivers respond to the overall search rate rather than to racial group-specific search rates (e.g., if drivers are uninformed regarding racial group-specific changes in search), then equalizing search rates across racial groups while keeping the overall search rate constant will not influence contraband carrying rates. To the extent that drivers do respond to racial group-specific search rate changes, then the relative (racial group-specific) elasticities of contraband carrying with respect to the search rate will determine the net impact of search rate equalization (see Bjerk, 2007 for a related discussion). In practice, changes in aggregate contraband carrying rates will be quite limited unless the difference in racial group-specific elasticities is large. Low observed rates of search and the fact that searches can only occur if a motorist is first stopped indicate, however, that these elasticities are likely to be low in the neighborhood of status quo search rates.

To provide support for the assertion that deterrence effects are likely to be negligible on the margin in our setting, we undertake two complementary exercises. The logic behind both exercises is that motorists cannot influence the trooper who conducts a stop within a sergeant area, but they can choose to not drive through a particular sergeant area while carrying contraband. Hence, if higher search rates deter motorists from carrying contraband, we expect motorist behavior to

[28] Note that this does not imply that all troopers face the same SPC. In fact, we see persistent differences in hit rates between troopers, implying differences in screening ability. The similarity of the between- and within-trooper SPCs suggests that trooper skill and search rates are largely unrelated.

be responsive to between-sergeant area variation in search rates, but not within-sergeant area, between-trooper variation in search rates.

In the first exercise, we compare SPC slopes estimated using within-sergeant area variation (as in Section IV.D) to SPC slopes estimated using cross-sergeant area boundary variation (as in Section IV.E.2). With significant deterrence effects, we would expect the SPC slope to be smaller when derived from cross-sergeant area boundary variation because the estimate would reflect the fact that motorists are less likely to carry contraband in high search rate sergeant areas. Instead, we find that the two slope estimates are statistically indistinguishable, suggesting that motorist behavior is not responding to sergeant area search rates.

In the second exercise, we compare the SPC slopes estimated using within-motorist variation (as in Section IV.E.1) among two sets of sequential pairs of stops. In one set, both stops are in the same sergeant area. In the other set, the two stops are in different sergeant areas. With significant deterrence effects, we would expect SPC slope derived using the second set of stops to be smaller. Instead, we find the two slopes to be statistically indistinguishable. In other words, we find no evidence of deterrence.

Importantly, the lack of measured deterrence effects is applicable only for the range of sergeant-area-level search rates we observe. Deterrence effects may be non-linear so that motorists are substantially more responsive to larger changes in search risk. For example, while we find that searches of white motorists are weakly more productive than searches of black or Hispanic motorists at the margin, suggesting that troopers could increase contraband yield by only searching white motorists, deterrence effects may be more relevant in that counterfactual. In addition, it is possible that motorists are insensitive to between-sergeant area variation in search rates because they are not aware of that variation. In that case, deterrence effects may be more relevant in counterfactuals where changes in motorist racial group-specific search rates are more salient.

V.C Equity and Counterfactuals

Based on the group-specific SPCs we estimate, it is straightforward to assess whether search rates can be equalized across motorist racial groups while maintaining search efficiency. To provide a benchmark, column (1) of Table V summarizes status quo search and hit rates by motorist group and in the aggregate. In column (2), we apply group-specific SPC slope estimates to calculate the predicted hit rate associated with status quo search rates (as expected, predicted hit rates are similar to those in column (1)).

To test for an equity-efficiency trade-off, in column (3) we characterize predicted changes in hit rates if all motorist racial groups were searched at a rate equal to the status quo aggregate search rate. The estimates indicate that equalizing search rates would modestly *increase* search productivity. Under the assumption that marginal changes in search rates do not influence contraband carrying behavior, this summary finding implies that there is no equity-efficiency trade-off present. In other words, it is feasible for troopers to search all motorist racial groups at the same rate, maintain the status quo overall search rate, and increase overall contraband yield.

[Table 5 about here.]

VI Are Search Disparities Driven by Racial Bias?

Black and Hispanic motorists are more likely to be searched, but racial disparities in search rates are not justified on efficiency grounds. As such, an outstanding question is why troopers elect to search black and Hispanic motorists more frequently. Two candidate explanations are that (1) troopers hold inaccurate beliefs regarding differences in contraband carrying behavior by race or (2) troopers search black and Hispanic motorists more frequently due to racial bias. Understanding the precise reason why black and Hispanic motorists are searched more frequently in the absence of efficiency gains is not central to our analysis, since our counterfactual conclusions hold regardless of the source of search rate disparities. Moreover, the legality of search rate disparities in the absence of efficiency gains does not likely depend on the specific mechanism that explains these disparities. Nonetheless, we briefly investigate the potential role of racial bias by assessing whether trooper-level racial disparities in search rates are associated with three factors: trooper race, local political preferences, and local disparities in citation rates.

A common test for racial bias in the policing literature is to compare the behavior of officers from different racial groups (Anwar and Fang, 2006; Close and Mason, 2007; Antonovics and Knight, 2009; West, 2018; Goncalves and Mello, 2018). The typical approach is to test whether black-white search rate disparities are smaller or reversed for black troopers, or whether Hispanic-white search disparities are smaller or reversed for Hispanic troopers. The premise is that if search disparities are driven by racial bias, we should expect biased troopers to favor motorists from the same racial group. We identify trooper race using 2015 personnel records for 2,469 troopers accounting for 84% of stops. Appendix Table B10 documents search rates and hit rates by both motorist and trooper race. Appendix Table B11, discussed in more detail in Appendix B, summarizes differences in black-white and Hispanic-white search odds ratios by trooper race that account for other stop and motorist characteristics. We find that all trooper racial groups are more likely to search black and Hispanic motorists than white motorists, but the black-white disparity is smaller for black troopers. The Hispanic-white disparity is similar for white and Hispanic troopers and smaller for black troopers.

We next examine whether, at the sergeant area level, racial group search rate disparities are associated with two proxies for trooper preferences and beliefs: racial disparities in citation rates (Goncalves and Mello, 2018) and the local Republican vote share in the 2016 presidential election (Cohen and Yang, 2019). For the 79 sergeant areas that we include in our estimation of race-specific SPCs, we calculate black-white and Hispanic-white search odds ratios using the following logistic regression model:

$$P(\text{SEARCH}_{it} = 1|X_{it}) = \frac{e^{\left(X_{it}\beta+\omega_{\tau(t)}+\delta_{m(t)}\right)}}{1+e^{\left(X_{it}\beta+\omega_{\tau(t)}+\delta_{m(t)}\right)}} \tag{12}$$

where X_{it} includes indicators for whether the motorist is female, black, and Hispanic.

To measure racial disparities in citation rates, for each sergeant area we estimate logistic regression models analogous to equation (12) where the outcome is replaced with an indicator for whether the stop led to a speeding citation. Overall, white motorists are cited in 28.2% of stops, while black and Hispanic motorists are cited in 34.9% and 37.5% of stops.[29] We next calculate the Republican vote share in each sergeant area in the 2016 presidential election. For sergeant areas that cover multiple counties, we take a weighted average of the county-level Republican vote shares where weights reflect the share of sergeant area stops conducted in each county.

Appendix Table B12 summarizes the joint, sergeant area-level relationship between search disparities, citation disparities, and Republican vote share. The relationship between black-white search disparities and vote share is both economically and statistically significant. A 10 percentage point increase in the Republican vote share is associated with an increase in the black-white odds ratio of 0.32. By contrast, there is no detectable relationship between black-white search disparities and citation disparities. We find a marginal positive relationship between Hispanic-white search and citation odds ratios, but no detectable relationship between the Hispanic-white search disparity and Republican vote share.

VII Conclusion

In this paper we use unique administrative data on traffic stops conducted by the Texas Highway Patrol to evaluate whether racial profiling poses an equity-efficiency trade-off. As in previous analyses, compared to white motorists, we find that troopers are more likely to search black and Hispanic motorists following stops for speeding, yet these searches are equally or less likely to yield contraband. In general, this finding does not imply that troopers can equalize search rates without sacrificing contraband yield due to the inframarginality problem: average search productivity may differ from search productivity at the margin. However, we show that the inframarginality problem is not empirically relevant in our setting. To do so, we exploit variation across troopers in search behavior and find that the relationship between trooper search rates and the proportion of stops that yield contraband is approximately linear. This finding suggests that, among motorists searched with positive probability, troopers are unable to distinguish between those who are more or less likely to carry contraband.

Overall, our results imply that it is feasible for troopers to (1) search all motorist racial groups at the same rate, (2) maintain the status quo overall search rate, and (3) increase overall contraband yield. Our findings highlight a limitation of the Becker (1957, 1993) outcome test: when the returns to search are constant, equalized marginal hit rates do not imply an equity-efficiency trade-off. More generally, our findings demonstrate that even if racial disparities in treatment cannot be definitively attributed to racial bias, such disparities may not be justifiable on efficiency grounds.

[29] Goncalves and Mello (2018) test whether troopers are more likely to give speeding "discounts" to white motorists than black motorists. There is no apparent bunching in speeds in our data.

References

Aaronson, Daniel, Lisa Barrow, and William Sander, "Teacher and Student Achievement in the Chicago Public Schools," *Journal of Labor Economics*, 2007, *25* (1), 95–135.

Anderson, April J., "Racial Profiling: Constitutional and Statutory Considerations for Congress," Technical Report, Congressional Research Service July 2020.

Antonovics, Kate and Brian G. Knight, "A New Look at Racial Profiling: Evidence fom the Boston Police Department," *Review of Economics and Statistics*, February 2009, *91* (1), 163–177.

Anwar, Shamena and Hanming Fang, "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence," *American Economic Review*, March 2006, *96* (1), 127–151.

_ , Patrick Bayer, and Randi Hjalmarsson, "The Impact of Jury Race in Criminal Trials," *Quarterly Journal of Economics*, May 2012, *127* (2), 1017–1055.

Arnold, David, Will Dobbie, and Crystal S. Yang, "Racial Bias in Bail Decisions," *Quarterly Journal of Economics*, November 2018, *133* (4), 1885–1932.

_ , _ , and Peter Hull, "Measuring Racial Discrimination in Bail Decisions," April 2020. Unpublished manuscript.

Ayres, Ian, "Outcome Tests of Racial Disparities in Police Practices," *Justice Research and Policy*, 2002, *4*, 131–142.

Barreca, Alan, Jason Lindo, and Glen Waddell, "Heaping-Induced Bias in Regression Discontinuity Designs," *Economic Inquiry*, 2016, *54* (1), 268–293.

Baumgartner, Frank R., Derek A. Epp, and Kelsey Shoub, *Suspect Citizens: What 20 Million Traffic Stops Tell Us About Policing and Race*, Cambridge University Press, 2018.

Becker, Gary S., *The Economics of Discrimination*, University of Chicago Press, 1957.

_ , "Nobel Lecture: The Economic Way of Looking at Behavior," *Journal of Political Economy*, June 1993, *101* (3), 385–409.

Benning, Tom and Andrew Chavez, "After 2 years and $800M, Texas' border boost has one solid outcome: more traffic tickets," November 2016.

Bjerk, David, "Racial Profiling, Statistical Discrimination, and the Effect of a Colorblind Policy on the Crime Rate," *Journal of Public Economic Theory*, 2007, *9* (3), 521–546.

Cattaneo, Matias, Richard Crump, Max Farrell, and Yingjie Feng, "On Binscatter," 2019. Unpublished manuscript.

Chan, David C., Matthew Gentzkow, and Chuan Yu, "Selection with Variation in Diagnostic Skill: Evidence from Radiologists," April 2020. Unpublished manuscript.

Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson, "Health Care

Exceptionalism? Performance and Allocation in the US Health Care Sector," *American Economic Review*, August 2016, *106* (8), 2110–2144.

Close, Billy R. and Patrick Leon Mason, "Searching for Efficienct Enforcement: Officer Characteristics and Racially Biased Policing," *Review of Law and Economics*, 2007, *3* (2), 263–321.

Cohen, Alma and Crystal S. Yang, "Judicial Politics and Sentencing Decisions," *American Economic Journal: Economic Policy*, February 2019, *11* (1), 160–191.

Collister, Brian, "Texas troopers ticketing Hispanic drivers as white," November 2015.

Correia, Sergio, Paulo Guimarães, and Thomas Zylkin, "Verifying the existence of maximum likelihood estimates for generalized linear models," August 2019. Unpublished manuscript.

Davis, Elizabeth, Anthony Whyde, and Lynn Langton, "Contacts Between Police and the Public, 2015," Special Report, Bureau of Justice Statistics October 2018.

Diamond, Rebecca, Tim McQuade, and Franklin Qian, "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, September 2019, *109* (9), 3365–3394.

Dominitz, Jeff and John Knowles, "Crime Minimisation and Racial Bias: What Can We Learn From Police Search Data?," *Economic Journal*, November 2006, *116* (515), F368–F384.

Epp, Charles R., Steven Maynard-Moody, and Donald P. Haider-Markel, *Pulled Over: How Police Stops Define Race and Citizenship*, Chicago: University of Chicago Press, 2014.

Frandsen, Brigham R., Lars J. Lefgren, and Emily C. Leslie, "Judging Judge Fixed Effects," June 2020. Unpublished manuscript.

Fryer, Roland, "An Empirical Analysis of Racial Differences in Police Use of Force," *Journal of Political Economy*, June 2019, *127* (3), 1210–1261.

Goncalves, Felipe and Steve Mello, "A Few Bad Apples? Racial Bias in Policing," January 2018. Unpublished manuscript.

Harcourt, Bernard E., "Rethinking Racial Profiling: A Critique of the Economics, Civil Liberties, and Constitutional Literature, and of Criminal Profiling More Generally," *University of Chicago Law Review*, 2004, *71* (4).

Horrace, William C. and Shawn M. Rohlin, "How Dark Is Dark? Bright Lights, Big City, Racial Profiling," *Review of Economics and Statistics*, May 2016, *98* (2), 226–232.

Knowles, John, Nicola Persico, and Petra Todd, "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy*, February 2001, *109* (1), 203–229.

MacDonald, John M. and Jeffrey Fagan, "Using Shifts in Deployment and Operations to Test for Racial Bias in Police Stops," *AEA Papers and Proceedings*, May 2019, *109*, 148–51.

Marx, Philip, "An Absolute Test of Racial Prejudice," May 2020. Unpublished manuscript.

Morris, Carl N., "Parametric Empirical Bayes Inference: Theory and Applications," *Journal of the American Statistical Association*, 1983, *78* (381), 47–55.

Persico, Nicola, "Racial Profiling, Fairness, and Effectiveness of Policing," *American Economic Review*, December 2002, *92* (5), 1472–1497.

Phillips, David, "Measuring Housing Stability with Consumer Reference Data," *Demography*, forthcoming.

Pierson, Emma, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Daniel Jenson, Amy Shoemaker, Vignesh Ramachandran, Phoebe Barghouty, Cheryl Phillips, Ravi Shroff, and Sharad Goel, "A large-scale analysis of racial disparities in police stops across the United States," March 2019. Unpublished manuscript.

Rehavi, M. Marit and Sonja B. Starr, "Racial Disparity in Federal Criminal Sentences," *Journal of Political Economy*, 2014, *122* (6).

Simoiu, Camelia, Sam Corbett-Davies, and Sharad Goel, "The Problem of Infra-Marginality in Outcome Tests for Discrimination," *The Annals of Applied Statistics*, 2017, *11* (3), 1193–1216.

Tiwara, Alisa, "Disparate-Impact Liability for Policing," *The Yale Law Journal*, 2019, *129* (1).

West, Jeremy, "Racial Bias in Police Investigations," 2018. Unpublished manuscript.

Figure I
A Theoretical Search Productivity Curve

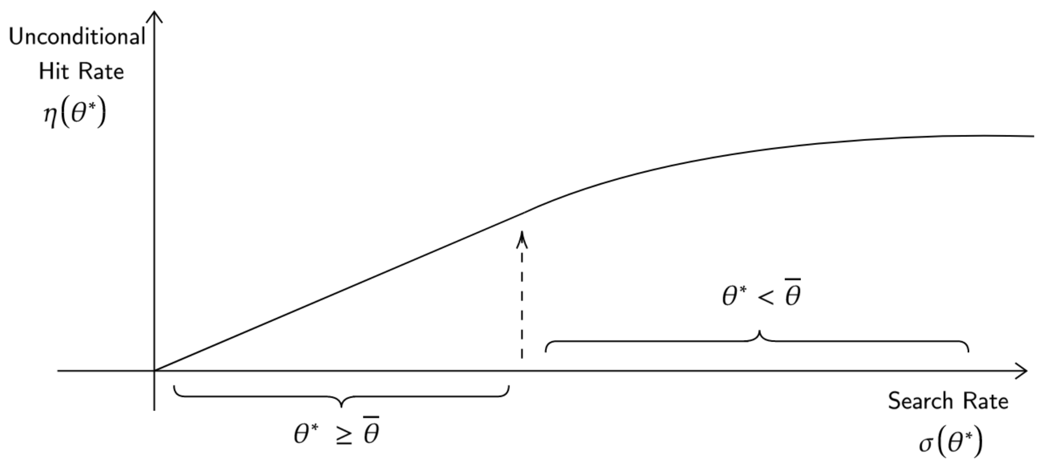

Note: This figure depicts a theoretical example of a trooper's search productivity curve (SPC), the relationship between the trooper's search rate, $\sigma(\theta^*)$, and unconditional hit rate, $\eta(\theta^*)$. For signal thresholds $\theta^* \geq \bar{\theta}$, the SPC is linear. For $\theta^* < \bar{\theta}$, the relationship is concave, as the marginal searched motorist is less likely to have contraband than inframarginal searched motorists.

Figure II
Theoretical Cases with and without an Equity-Efficiency Trade-Off

(a) Trade-Off

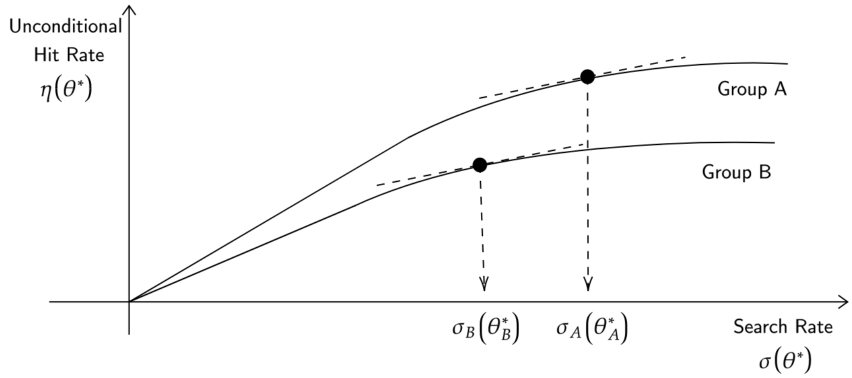

(b) No Trade-Off and $\theta_r^* < \bar{\theta}_r$

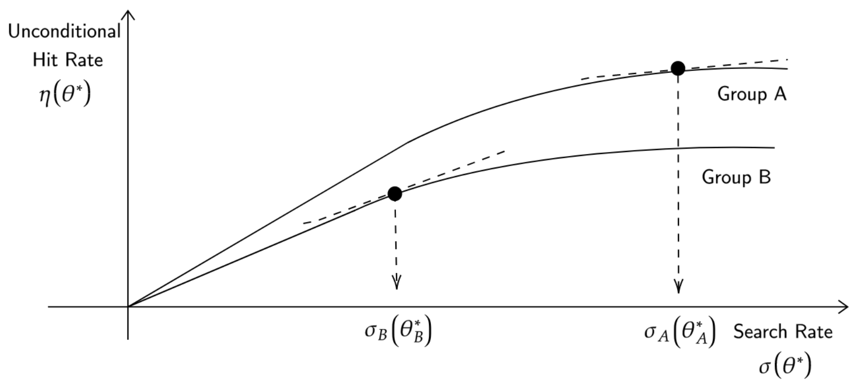

(c) No Trade-Off and $\theta_r^* \geq \bar{\theta}_r$

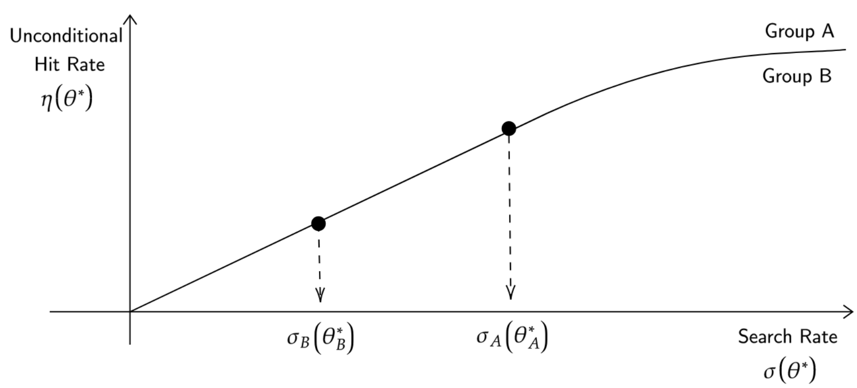

Note: These figures depict three scenarios for search productivity curves (SPCs) for two groups of motorists, Group A and Group B. In all scenarios, the search rate for Group A exceeds the search rate for Group B. In Panel A, the trooper faces diminishing returns to search within each group and equalizes marginal hit rates between groups. Equalizing search rates between groups while maintaining the overall search rate would decrease the hit rate. In Panel B, the trooper faces diminishing returns to search within each group but does not equalize marginal hit rates. In Panel C, the trooper equalizes marginal hit rates between groups but faces constant returns to search. In Panels B and C, the trooper can equalize search rates without reducing the overall hit rate.

Figure III
Number of Troopers Satisfying Sample Criteria by Sergeant Area

(a) Pooled SPC Criteria

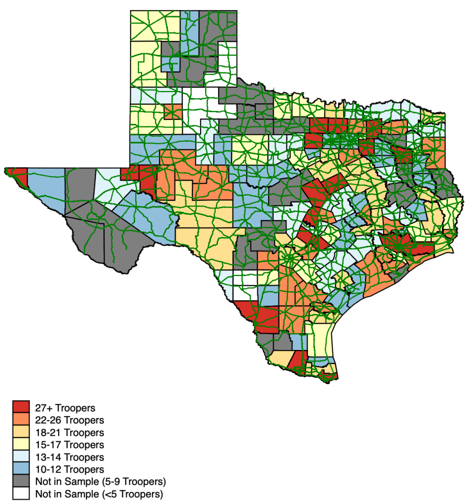

(b) Race-Specific SPC Criteria

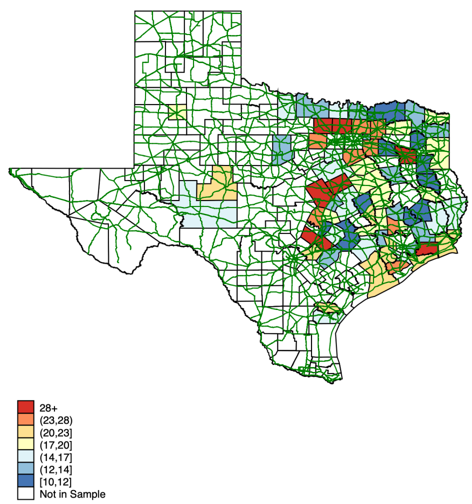

Note: These maps depict the number of troopers in each sergeant area that satisfy the sample criteria described in Section IV.B. State and interstate highways are superimposed in green. Panel A depicts the number of troopers who satisfy sample criteria for estimating the pooled search productivity curve (SPC). We include sergeant areas in the analysis if they have at least ten troopers meeting the sample criteria. For sergeant areas included in the estimation of race-specific SPCs, Panel B depicts the number of troopers who satisfy the sample criteria, averaging across motorist racial groups (white, black, Hispanic). Sergeant areas included in the estimation of race-specific SPCs have at least ten troopers meeting the sample criteria for each motorist racial group.

Figure IV
Distribution of Search Rates across Troopers and Locations

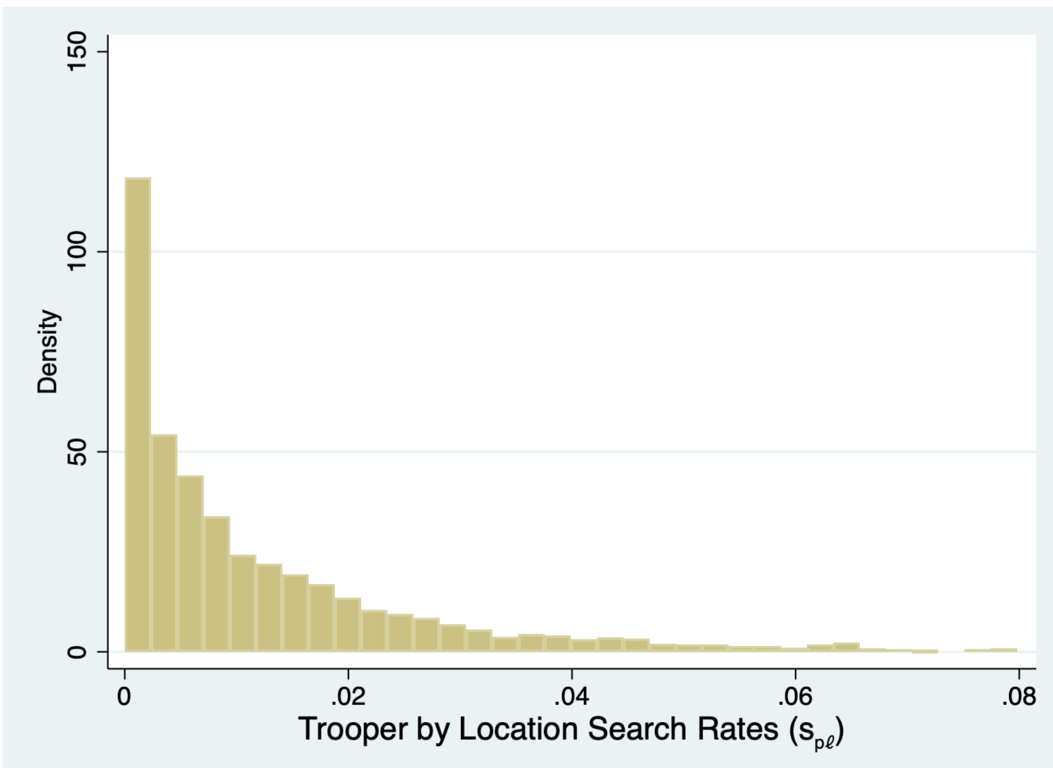

Note: This figure plots the distribution of search rates across trooper-by-location combinations ($s_{p\ell}$). Sample restrictions are described in Section IV.B.

Figure V
Between-Trooper Search Productivity Curve

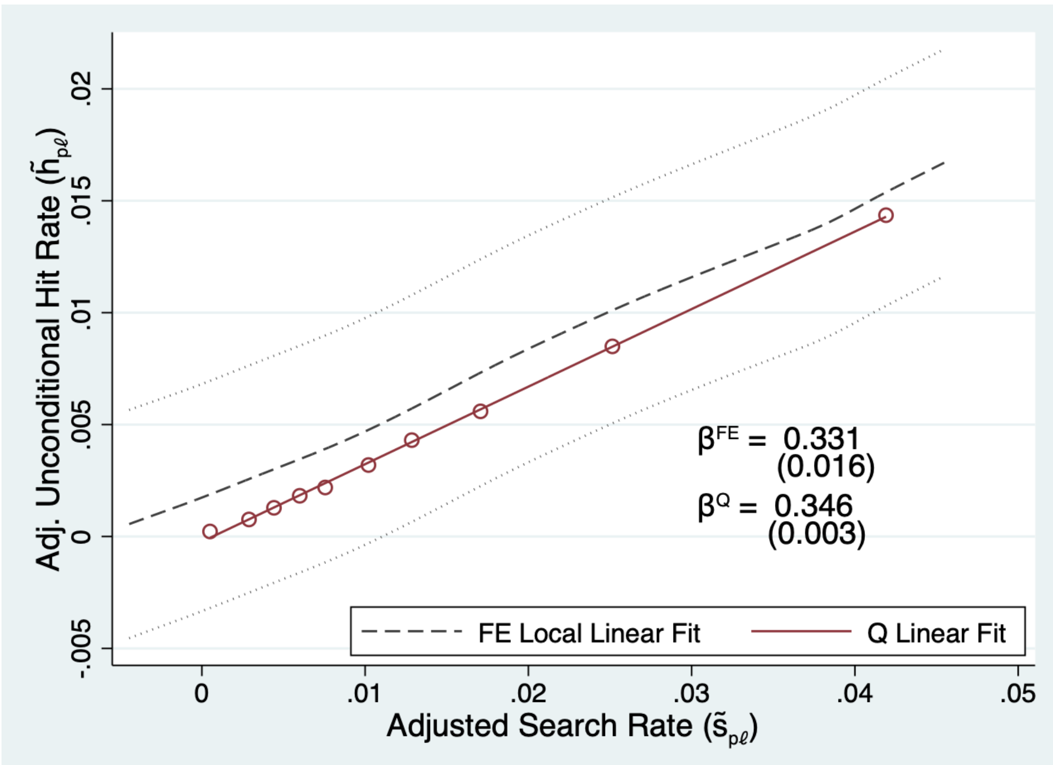

Note: In this figure we plot adjusted trooper unconditional hit rates ($\tilde{h}_{p\ell}$) against trooper search rates ($\tilde{s}_{p\ell}$) using two approaches described in Section IV.D: the fixed effects (FE) approach and the quantile (Q) approach. In the fixed effects approach we net out location fixed effects from both $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$ and plot the residuals. From the FE approach, the figure includes 95% confidence bands for the local linear relationship between adjusted trooper search rates and unconditional hit rates and the best linear fit and its slope. The local linear fit is derived using a Gaussian kernel with a rule-of-thumb bandwidth. In the quantile approach we divide troopers into quantiles by search rate within locations, group quantiles across locations, and then plot the relationship between search rates and unconditional hit rates across quantiles. From the Q approach, the figure includes the mean values for each decile and the best linear fit and its slope.

Figure VI
Between-Trooper Search Productivity Curve, by Motorist Race

(a) White Motorists

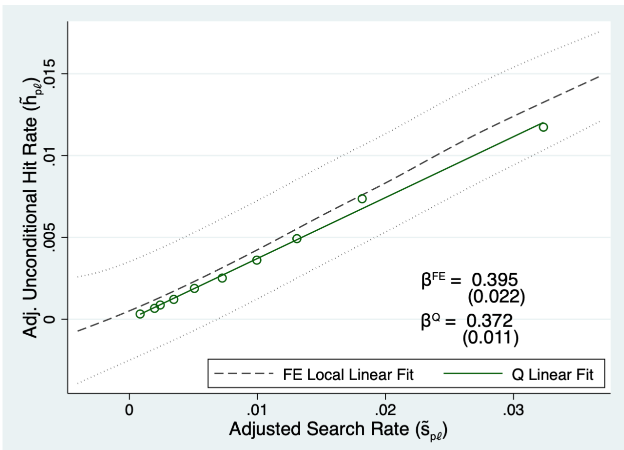

(b) Black Motorists

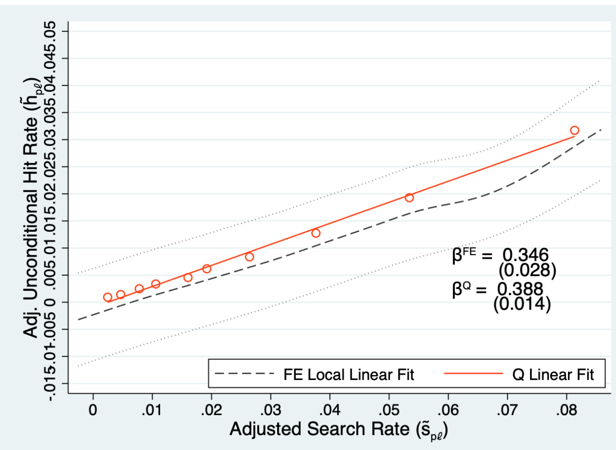

(c) Hispanic Motorists

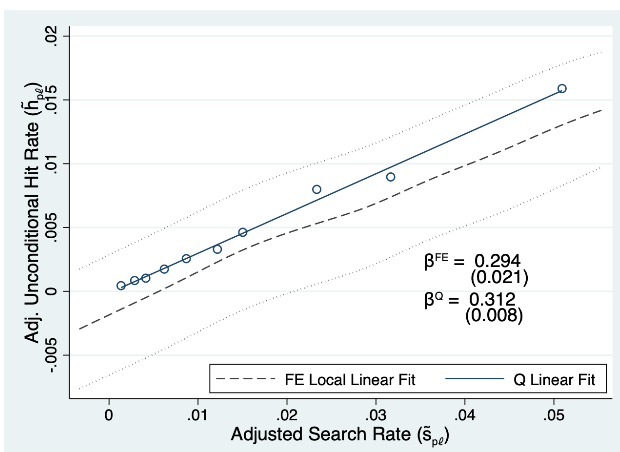

Note: In this figure we plot adjusted trooper unconditional hit rates ($\tilde{h}_{p\ell}$) against trooper search rates ($\tilde{s}_{p\ell}$) using two approaches described in Section IV.D: the fixed effects (FE) approach and the quantile (Q) approach. In the fixed effects approach we net out location fixed effects from both $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$ and plot the residuals. From the FE approach, the figure includes 95% confidence bands for the local linear relationship between adjusted trooper search rates and unconditional hit rates and the best linear fit and its slope. The local linear fit is derived using a Gaussian kernel with a rule-of-thumb bandwidth. In the quantile approach we divide troopers into quantiles by search rate within locations, group quantiles across locations, and then plot the relationship between search rates and unconditional hit rates across quantiles. From the Q approach, the figure includes the mean values for each decile and the best linear fit and its slope. Panel A, Panel B, and Panel C plot the search productivity curve (SPC) for white motorists, black motorists, and Hispanic motorists, respectively.

Figure VII
Selection and within-Motorist Differences in Trooper Search Rates

(a) Motorist Characteristics

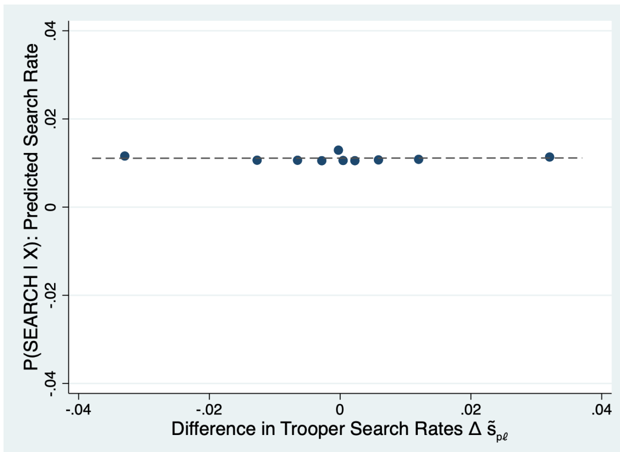

(b) ΔSearch Rate

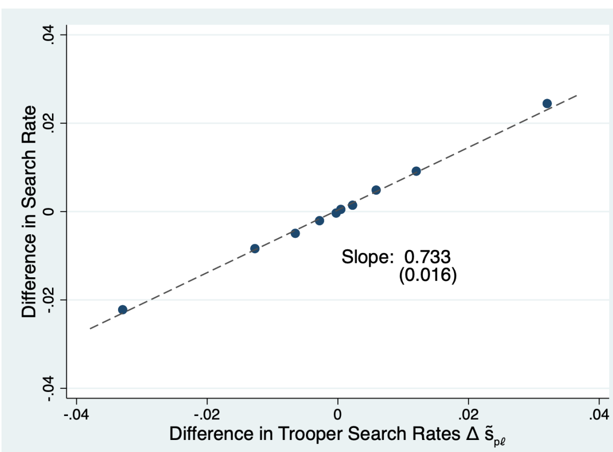

(c) ΔUnconditional Hit Rate

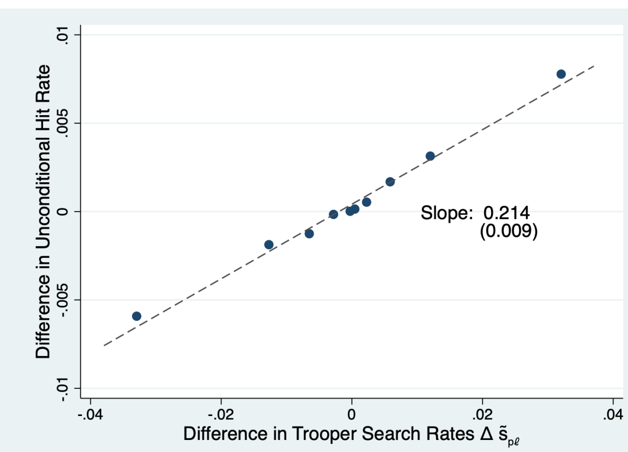

Note: This figure plots the relationship between the difference in trooper search rates associated with sequential pairs of stops for the same motorist, $\Delta_{it}\tilde{s}_{p\ell}$, and three variables: (1) motorist characteristics, (2) the difference in search rates between stops, and (3) the difference in unconditional hit rates between stops. Motorist characteristics are summarized by the probability of search given motorist characteristics at the time of the initial stop, $P(\text{SEARCH}_{it}|X_{it})$. Sequential pairs of stops are grouped by their decile value of $\Delta_{it}\tilde{s}_{p\ell}$.

Figure VIII
Trooper Assignments by Stop Location

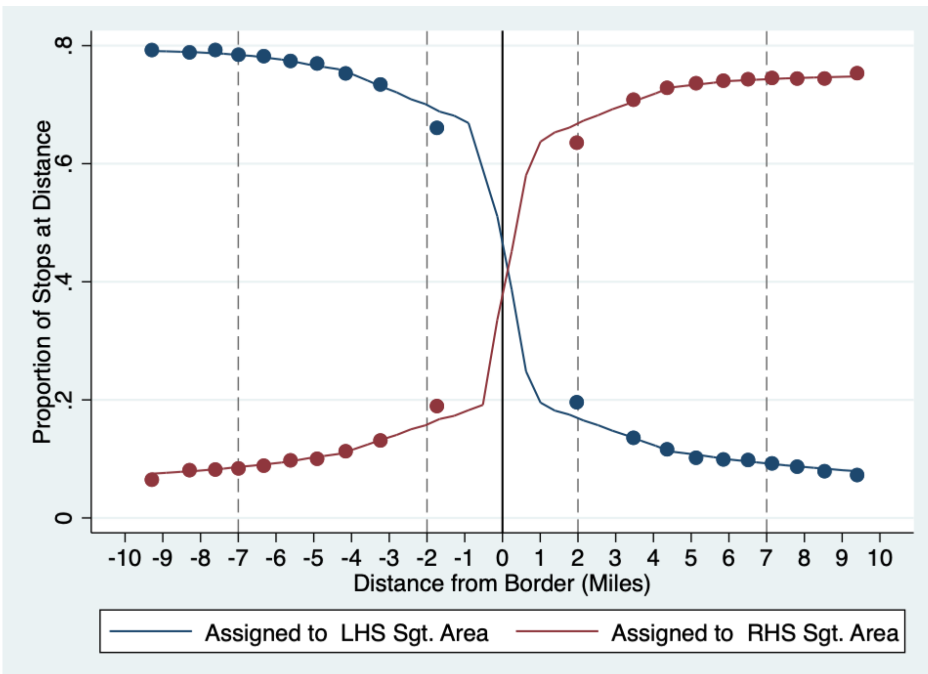

Note: This figure plots the share of stops conducted by troopers assigned to each adjacent sergeant area by travel distance from highway and sergeant area border intersections as described in Section IV.E.2. The data are limited to stops within 10 miles of the intersection. The figure includes a bin scatter, where stops are grouped by side of the border and into deciles by distance from the intersection. Stops conducted between 2 and 7 miles of the intersections are included in the regression discontinuity (RD) analysis.

Figure IX
Stop Characteristics and Outcomes by Stop Distance from Border

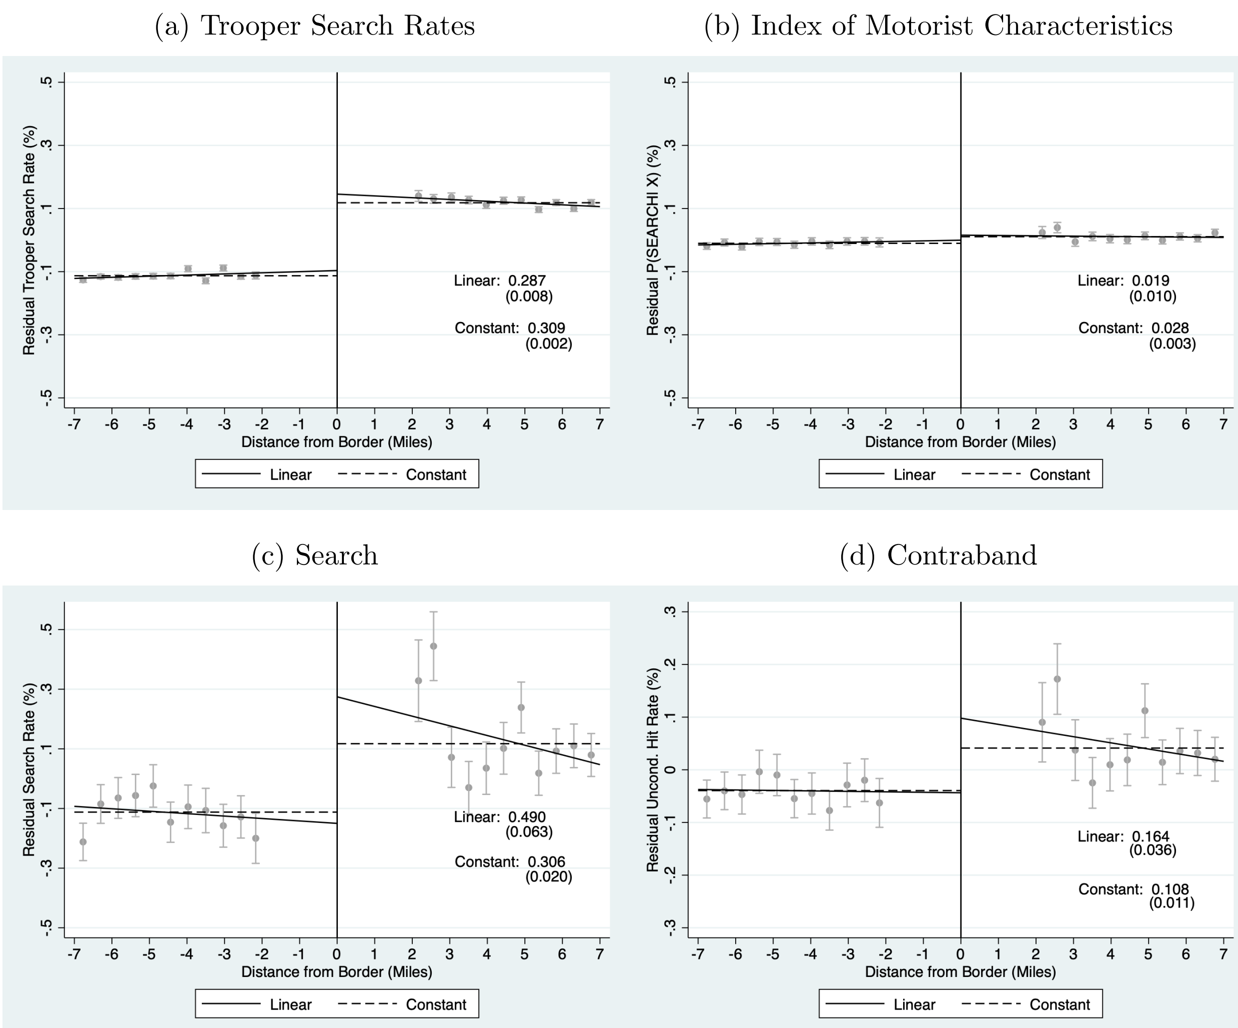

Note: These figures are stacked regression discontinuity plots for four outcomes: leave-out trooper search rates (Panel A); $P(\text{SEARCH}_{it}|X)$ (Panel B), an index of motorist characteristics; search (SEARCH_{it}) (Panel C); and contraband (CONTRABAND_{it}) (Panel D). The running variable is the travel distance from a stop to its corresponding highway by sergeant area border intersection. Sample selection criteria are described in Section IV.E.2. The figures include discontinuity estimates using linear and constant extrapolation.

Table I
Traffic Stop Descriptive Statistics

	All Stops				All Searches			
	Black	Hispanic	White	All	Black	Hispanic	White	All
Black	100	0	0	9.09	100	0	0	18.91
Hispanic	0	100	0	35.88	0	100	0	41.82
White	0	0	100	55.03	0	0	100	39.26
Female	38.95	32.59	37.67	35.96	16.32	15.58	20.80	17.77
Log Median Income	10.70	10.71	10.95	10.84	10.59	10.60	10.87	10.71
	(0.500)	(0.495)	(0.466)	(0.495)	(0.492)	(0.469)	(0.471)	(0.493)
Expected Log Income Given	-0.161	-0.083	0.081	0.000	-0.524	-0.456	-0.371	-0.436
Vehicle (Standardized)	(1.013)	(0.966)	(1.012)	(1.000)	(0.889)	(0.841)	(0.893)	(0.873)
Stop History:								
No Prior Stops	60.32	57.31	56.53	57.15	58.12	56.72	53.60	55.76
Prior Stop, No Search	37.92	41.13	42.61	41.66	32.01	35.12	37.50	35.47
Prior Search, No Contraband	1.130	1.205	0.534	0.829	4.669	4.429	3.796	4.226
Prior Search, Contraband	0.626	0.356	0.322	0.362	5.195	3.728	5.098	4.544
Non-Drug Arrest History:								
No Prior Non-Drug Arrests	87.02	89.41	93.48	91.43	64.46	72.42	71.80	70.67
1-2 Prior Non-Drug Arrests	7.158	6.830	4.296	5.465	14.88	14.46	14.24	14.45
3+ Prior Non-Drug Arrests	5.824	3.760	2.223	3.102	20.66	13.12	13.96	14.88
Drug Arrest History:								
No Prior Drug Arrests	93.92	96.29	97.55	96.77	73.48	82.39	79.42	79.54
1 Prior Drug Arrest	2.814	2.161	1.363	1.781	10.13	8.588	9.426	9.208
2+ Prior Drug Arrests	3.263	1.547	1.088	1.451	16.40	9.019	11.16	11.25
Search Rate	2.202	1.234	0.755	1.059	100	100	100	100
Unconditional Hit Rate	0.757	0.324	0.285	0.342	34.00	25.90	37.35	31.93
Observations	448,337	1,769,369	2,713,626	4,931,332	9,874	21,832	20,497	52,203

Sample restrictions are described in Section II. ‘Log Median Income’ refers to the median household income for the Census block group of the motorist’s residential address as measured in the 2009-2013 5-year American Community Survey. ‘Expected Log Income Given Vehicle’ is the average Log Median Income associated with a vehicle, where vehicles are classified as a combination of make, type (passenger car, pick-up truck, SUV), and age (above and below median given make and type), generating 204 total vehicle categories. We standardize Expected Log Income Given Vehicle to have mean zero and standard deviation one in our sample of stops.

Table II
Motorist Selection into Stops by Trooper Search Rate

				Excluding Most Selective Troopers		
	100 × SEARCH_{it}	100 × $s_{p\ell}^{-it}$	100 × $\tilde{s}_{p\ell}^{-it}$	100 × SEARCH_{it}	100 × $s_{p\ell}^{-it}$	100 × $\tilde{s}_{p\ell}^{-it}$
	(1)	(2)	(3)	(4)	(5)	(6)
Black	0.888	0.106	0.066	0.774	0.050	0.026
	(0.058)	(0.020)	(0.020)	(0.054)	(0.013)	(0.012)
Hispanic	0.299	0.028	0.009	0.274	0.007	-0.005
	(0.022)	(0.008)	(0.008)	(0.022)	(0.006)	(0.006)
Female	-0.513	-0.023	-0.018	-0.474	-0.010	-0.008
	(0.021)	(0.004)	(0.004)	(0.021)	(0.003)	(0.003)
Log Median Income	-0.312	-0.007	-0.006	-0.281	0.003	-0.001
	(0.019)	(0.005)	(0.005)	(0.019)	(0.004)	(0.004)
Expected Log Income Given	-0.295	-0.027	-0.025	-0.258	-0.011	-0.011
Vehicle (Standardized)	(0.012)	(0.003)	(0.003)	(0.011)	(0.002)	(0.002)
1-2 Prior Non-Drug Arrests	0.805	0.011	0.008	0.744	-0.003	-0.004
	(0.053)	(0.005)	(0.005)	(0.055)	(0.004)	(0.004)
3+ Prior Non-Drug Arrests	1.700	0.022	0.019	1.531	0.008	0.008
	(0.099)	(0.008)	(0.008)	(0.101)	(0.006)	(0.006)
1 Prior Drug Arrest	3.304	0.022	0.017	3.106	0.002	0.001
	(0.145)	(0.007)	(0.007)	(0.146)	(0.006)	(0.006)
2+ Prior Drug Arrests	5.324	0.029	0.020	4.823	-0.000	-0.004
	(0.213)	(0.010)	(0.010)	(0.218)	(0.008)	(0.008)
Prior Stop, No Search	-0.221	-0.015	-0.008	-0.192	-0.010	-0.006
	(0.017)	(0.004)	(0.004)	(0.017)	(0.004)	(0.004)
Prior Search, No Contraband	2.943	0.072	0.065	2.806	0.033	0.031
	(0.201)	(0.017)	(0.017)	(0.201)	(0.012)	(0.012)
Prior Search, Contraband	10.025	0.158	0.154	8.722	0.100	0.105
	(0.642)	(0.025)	(0.024)	(0.653)	(0.023)	(0.022)
Location by Time FEs	✓	✓	✓	✓	✓	✓
Joint F-Statistic	84.88	9.31	8.65	74.03	6.30	5.84
Observations	3,280,250	3,280,250	3,280,171	2,739,955	2,739,955	2,739,899

This table presents coefficients from estimates of equation (7), where in columns (2) and (3) we replace the outcome SEARCH_{it} with $s_{p(i,t)\ell(i,t)}^{-it}$ and $\tilde{s}_{p(i,t)\ell(i,t)}^{-it}$, leave-out trooper search rates corresponding to the trooper who conducted the stop. Columns (4)–(6) exclude stops conducted by the 20% of troopers with the most selected set of stopped motorists. Standard errors are clustered at the motorist level. 'Joint F-Statistic' refers to an F-test for the joint significance of all motorist characteristics.

Table III
Search Productivity Curve Estimates, within-Motorist Design

	Δ_{it}SEARCH	Δ_{it}CONTRABAND	
	OLS	OLS	2SLS
	(1)	(2)	(3)
Pooled			
$\Delta_{it}\tilde{s}_{p\ell}$	0.733	0.214	
	(0.016)	(0.009)	
Δ_{it}SEARCH			0.292
			(0.011)
Observations	694,246	694,246	694,246
White Motorists			
$\Delta_{it}\tilde{s}^{r}_{p\ell}$	0.750	0.272	
	(0.032)	(0.019)	
Δ_{it}SEARCH			0.362
			(0.021)
Observations	245,778	245,778	245,778
Black Motorists			
$\Delta_{it}\tilde{s}^{r}_{p\ell}$	0.778	0.247	
	(0.039)	(0.024)	
Δ_{it}SEARCH			0.317
			(0.025)
Observations	50,674	50,674	50,674
Hispanic Motorists			
$\Delta_{it}\tilde{s}^{r}_{p\ell}$	0.688	0.196	
	(0.036)	(0.020)	
Δ_{it}SEARCH			0.284
			(0.025)
Observations	87,320	87,320	87,320

This table presents coefficients for the two-stage least squares (2SLS) system described by equations (8) and (9). Each observation is a pair if sequential stops for a given motorist. Standard errors are clustered at the motorist level.

Table IV
Within-Trooper Returns to Search

	$P(\text{SEARCH}_{it} \mid X_{it})$	SEARCH_{it}		CONTRABAND_{it}	
	OLS	OLS	OLS	IV	IV
	(1)	(2)	(3)	(4)	(5)
SEARCH_{it}				0.273	0.261
				(0.087)	(0.095)
$s^{-p}_{\hat{\ell}(i,t)}$	0.012	0.130	0.116		
	(0.004)	(0.030)	(0.030)		
Motorist by Location FEs	✓	✓	✓	✓	✓
Year FEs	✓	✓	✓	✓	✓
Motorist Controls			✓		✓
Observations	3,206,583	3,206,583	3,206,583	3,206,583	3,206,583

This table presents estimates of equations (10) and (11). $s^{-p}_{\hat{\ell}(i,t)}$ denotes the search rate of all troopers assigned to location ℓ in the year corresponding to t *excluding trooper* p. Motorist characteristics include race, gender, log of neighborhood median income, vehicle-based expected log income, stop history, non-drug arrest history, and drug arrest history.

Table V
Counterfactuals

	Observed	Estimated SPC-Based	
	Status Quo	Status Quo	Equal Search Rates
All Motorists			
Search Rate	1.16	1.16	1.16
Hit Rate	34.20	34.74	35.89
		(0.57)	(0.56)
White Motorists			
Search Rate	0.82	0.82	1.16
Hit Rate	37.28	39.42	39.42
		(0.53)	(0.54)
Black Motorists			
Search Rate	2.24	2.24	1.16
Hit Rate	34.51	34.26	34.26
		(0.62)	(0.62)
Hispanic Motorists			
Search Rate	1.37	1.37	1.16
Hit Rate	30.14	29.36	29.36
		(0.57)	(0.57)

This table presents observed and simulated counterfactual search rates and hit rates by motorist racial group. Observed search rates and hit rates are calculated using the sample of stops used to estimate race-specific search productivity curves (SPCs). Sample criteria are described in Section IV.B. Simulated counterfactual search rates and hit rates are calculated using racial group-specific SPC estimates from Section IV.D.

A Appendix: Data Construction

We merge traffic stop data to commercial address history data from Infogroup using full name and address. We first use an address standardization algorithm, the Stata function `stnd_address`, to ensure that addresses are structured analogously across the two data sets, with separate fields for street address, unit number, etc. We also extract the address number. In addition, we manually standardize Texas city and town names in the traffic stop data. We standardize full names and extract suffixes. We then use the Stata command `reclink2` to perform a probabilistic linkage across the two data sources. We fuzzy match using the following fields: last name, first name, middle name, suffix, address number, street name, city, and zip code. We require that observations match exactly on the first letter of the first name and the first letter of the last name. For zip code, we define agreement discretely based on whether the fields match exactly. For all other fields, we utilize the bigram string comparator to assess the degree of agreement. The address history data includes an identifier that matches the same individual to multiple addresses. We use this identifier to match multiple stops to the same person. We are able to match 75% of stops to the address history data. For stops that we are unable to match, we create identifiers based on full name, street address, and zip code.

We then match the criminal history data to traffic stops using the full set of addresses associated with each person. We apply the same address and name standardization to the criminal history data, and apply the same fuzzy match.

Though Diamond et al. (2019) and Phillips (forthcoming) find that similar address history data from Infutor are of high quality, we are unable to match every stop to the address history data and these data may be incomplete. Hence, we may not correctly associate all stops and criminal history with the corresponding motorist.

B Appendix: Additional Analyses

B.1 Variation in Cited Speeds Across Troopers Is Limited

We limit our analysis to traffic stops associated with speeding violations because, as previous researchers have argued (see, for example, Baumgartner et al., 2018), we believe these stops are more likely to be motivated by the traffic violation itself, rather than some investigatory motive. For stops that are motivated by the traffic violation itself, we expect the composition of stopped motorists to be more similar across troopers, conditional on the location and time of the stop.

In this section we document the extent that cited speeds vary across troopers. Each traffic stop is associated with a speeding warning or a speeding citation. There is a citation in 34% of stops. Actual speeds are observed for citations but not warnings.

We rescale cited speeds by taking the difference between the log cited speed and log posted speed limit. We refer to this rescaled speed as the *log speed above limit*.

The average log speed above limit is 0.208, meaning the average cited speed is about 21% over the posted speed limit. The standard deviation of log speed above limit is 0.076. 99% of cited speeds are at least 10% above the speed limit. Cited speeds are similar for white and Hispanic motorists (about 20.6% above the posted speed limit), while cited speeds are slightly higher for black motorists (22.4% above the posted speed limit).

For each trooper-by-location combination, we calculate the citation rate and average log speed above limit. Within locations, the standard deviation of average log speed above limit across troopers is 0.027. The difference in average log speed above limit between the 10th and 90th percentile of troopers is only about 6% of the speed limit, which for the average speed limit is about 4 miles per hour.

Troopers that cite more often have lower average cited speeds, but the differences are minor. For every 10 percentage point increase in the citation rate, cited speeds decrease by 0.4%.

B.2 Detailed Search Outcomes by Motorist Race

Detailed outcomes of searches are summarized in Table B1.

B.3 Descriptive Analysis of What Predicts Search and Contraband Yield

We use the uniquely rich merged data set to answer two descriptive questions: (1) what motorist characteristics predict trooper search? And (2) among those searched, what motorist characteristics predict whether a search yields contraband? The answers to these questions clarify the degree to which race-based differences in search and hit rates can be explained by factors correlated with race but that have been unobservable to previous researchers.

B.3.1 Racial Disparities in Search Rates

One advantage of our setting relative to prior analyses is that we have a much richer set of motorist covariates. It is potentially the case that racial differences in search rates documented previously–and interpreted as evidence of racial profiling–could be explained, at least in a statistical sense, by other motorist characteristics that are observed by troopers but typically not observed by researchers. We investigate this possibility by examining whether conditioning on criminal history, stop history, and income affects measured race-based differences in search rates.

For each stop, let i denote the motorist and t denote the specific time. The functions $\ell(i,t)$, $\tau(t)$, and $m(t)$ map each stop to its associated location, time category, and month (e.g., June 2013), respectively. We categorize time by as the combination of quarter of day and whether the stop was conducted on a weekday or weekend. We estimate logistic regressions of the form

$$P(\text{SEARCH}_{it} = 1 | \ell(i,t), \tau(t), m(t), X_{it}) = \frac{e^{(\lambda_\ell + \omega_\tau + \delta_m + X_{it}\gamma)}}{1 + e^{(\lambda_\ell + \omega_\tau + \delta_m + X_{it}\gamma)}}, \tag{B.1}$$

where SEARCH_{it} is an indicator whether the stop of motorist i at time t led to a search; λ_ℓ, ω_τ, and δ_m are fixed effects for stop location, time category, and month of the stop; and X_{it} is a vector of motorist characteristics, including race, gender, log of neighborhood median income, stop history, non-drug arrest history, and drug arrest history. We also construct a second proxy for motorist income based on the vehicle involved in the stop. We classify vehicles by make, type (passenger car, pick-up truck, SUV), and age (above and below median given make and type), generating 204 total vehicle categories. We then calculate the mean of log of median neighborhood income among stopped motorists for each vehicle category. To our knowledge, this is the first paper to examine the relationship between trooper search behavior and motorist's criminal history, stop history, and neighborhood income.

Odds ratios for estimates of equation (B.1) are presented in Table B2. Across specifications, we vary the set of covariates included in the model, moving from more parsimonious specifications to more saturated models. In column (1) we include only a subset of motorist characteristics (X_{it}): motorist race and gender. The baseline search rate for white motorists is 0.76 percent. The coefficient for *black* of 3.00 indicates that, controlling only for gender, odds of search are 3 times higher for black motorists. Given the low probabilities in this context, odds and probabilities are similar, meaning search rates are also approximately 3 times higher for black motorists. For Hispanic motorists, search rates are about 58% higher. Conditional on motorist race, women are about 62% less likely to be searched. In column (2) we add separate fixed effects for stop location, time category, and month. Doing this reduces the *black* odds ratio slightly to 2.70, while the *Hispanic* odds ratio increases to 1.68. The coefficient for *female* is unaffected. In column (3) we add our income proxies as covariates. The coefficient for *log median income* is 0.68, indicating that a one standard deviation increase in neighborhood income (about 35 log points) is associated with about a 11% decrease in search rates. The association with *vehicle-based expected log income* is similar. Including the income proxies as controls reduces the *black* odds ratio to 2.20 and the *Hispanic* odds ratio to 1.45.

Column (4) adds motorist arrest history indices as explanatory variables. Previous arrests also predict searches, particularly drug arrests. The increase in search likelihood associated with black motorists relative to white motorists is similar in magnitude to the increase in search likelihood associated with multiple previous non-drug arrests and half of the increase associated with a prior drug arrest. Column (5) adds motorist stop history. Conditional on motorist demographics and arrest history, motorists who have been stopped previously but not searched previously are 18% less likely to be searched than motorists who have not been stopped previously, the omitted category. Motorists who have been searched previously but not found with contraband are 115% more likely to be searched, while motorists who have been previously found with contraband are 383% more likely to be searched.

Controlling for criminal and stop history reduces the black and Hispanic odds ratios to 1.86 and 1.44. Comparing columns (2) and (5), motorist income and criminal/stop history can statistically explain about 50% and 35% of the black-white and Hispanic-white disparities in search rates, respectively. Note that racial differences in stop history and likely criminal history already incorporate racial differences in police treatment. Hence, we think of these percentages as upper bounds on the share of black-white and Hispanic-white disparities that can be explained by these factors.

B.3.2 Racial Disparities in Hit Rates

Next, we estimate logistic models identical to (B.1) except that we replace the outcome with CONTRABAND_{it}, an indicator for whether a search yields contraband.[30] We limit estimation to stops that led to a search (i.e., where $\text{SEARCH}_{it} = 1$).

The results are presented in columns (6) through (10) of Table B2. The model specifications are analogous to those in columns (1) through (5).

There are four main findings. First, controlling for only motorist race and gender, searches of black and Hispanic motorists are about 14% and 41% less likely to yield contraband.

Second, hit rates are increasing in motorist income, and the magnitude of the relationship is economically significant. In columns (8) through (10) the coefficient for log neighborhood median income hovers around 1.27, indicating that a one standard deviation increase in neighborhood income is a associated with a nearly 10% increase in the hit rate. Interestingly, hit rates are unrelated to our vehicle-based proxy for motorist income.

Third, while previous drug arrests predict about a 40% increase in the hit rate, hit rates are weakly *lower* for motorists with previous non-drug arrests. For those with one or two previous non-drug arrests, the hit rate is the same as for those without any non-drug arrests; for those with more than two previous non-drug arrests, the hit rate is about 13% lower. This finding is particularly interesting given that previous non-drug arrests significantly *increase* a motorist's likelihood of being searched in the first place.

Fourth, for motorists who have been previously searched, the outcomes of those previous searches are highly predictive of contemporaneous outcomes. Relative to motorists with no stop history, searches of motorists who have been previously searched but not found with contraband are 37% less likely to yield contraband. Searches of motorists who have been previously found with contraband are 245% more likely to yield contraband.

B.4 Estimating Between-Trooper Search Productivity Curves

We estimate search productivity curve slopes using various specifications in Table B6.

One concern with our approach is that $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$, as estimates of their population analogs, $\sigma_{p\ell}$ and $\eta_{p\ell}$, are subject to sampling error, and those errors are correlated. This correlated sampling error may bias our estimate of β.

As one approach to accounting for this measurement error, we adjust estimates of trooper-location search propensities using an Empirical Bayes (EB) approach (Morris, 1983, Aaronson et al., 2007). We observe trooper-location search rates, which are estimates of search propensities.

[30] We show in Table B1 that the percent of stops yielding contraband that lead to an arrest and the severity of arrest charges, as proxied by the average incarceration sentence associated with conviction, are similar across motorist racial groups.

Some trooper-location estimates are derived from more observations and are thus more precise than others. The EB estimate for trooper-location $p\ell$ is a weighted average of the trooper-location search rate and overall search rate of the location, where the weight is a function of the reliability of the trooper-location $p\ell$ estimate. We follow the approach of Chandra et al. (2016) and use their Stata code to construct EB estimates for trooper-location search rates, $s_{p\ell}^{EB}$. We construct an analogous EB estimate of conditional search propensities, $h_{p\ell}^{EB}$, using the same weighting.

In Table B6 we show alternative estimates for the SPC slope from regressing $h_{p\ell}^{EB}$ on $s_{p\ell}^{EB}$ with location fixed effects. The slope we estimate is indistinguishable from the slope we get using unadjusted search and unconditional hit rates.[31]

To account for sampling error, we also take a split-sample IV approach to estimation. We randomly split stops into two samples and estimate $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$ separately in each sample. In each sample, we regress $\tilde{h}_{p\ell}$ on $\tilde{s}_{p\ell}$ and location fixed effects, instrumenting for $\tilde{s}_{p\ell}$ using its pair estimate from the other sample. Reassuringly, as shown in Table B6, this procedure yields β estimates that are statistically indistinguishable from the OLS estimates.

B.4.1 Excluding Selective Troopers

A key concern with our research design is that stopped motorists are not randomly assigned to troopers. We take a 'conditional on observables' approach and argue that, conditional on stop time and location, the identity of the trooper conducting the stop is as good as randomly assigned. However, even conditional on these stop contextual characteristics, we see motorist characteristics that predict search (e.g., race, income, criminal history) also predict the search propensity of the troopers that stop them. This relationship is quite weak (as discussed in Section IV.B), and controlling directly for observable motorist characteristics does not affect any of our conclusions. Nonetheless, this selection may introduce bias.

Here we take a complementary approach to assess whether our results are sensitive to this form of selection. We exclude troopers for whom we find the most evidence of motorist selection, and then repeat our analysis using this selected sample of troopers.

We first describe how we identify the troopers to exclude. The goal is to identify troopers that have a composition of stopped motorists that deviates most from what one would expect based on the time and location of their stops alone. We estimate the following logistic regression model:

$$P(\text{SEARCH}_{it} = 1 | X_{it}) = \frac{e^{(X_{it}\beta)}}{1 + e^{(X_{it}\beta)}} \tag{B.2}$$

where X_{it} is a vector of motorist characteristics including motorist race, gender, log of neighborhood income, expected log income given vehicle, stop history, non-drug arrest history, and drug arrest history. From this we calculate the search probability for each stop based on observable motorist

[31] Note that the SPC slopes presented here differ somewhat from the slopes presented in Figure V and Figure VI. The slopes in the main text are fit to local linear estimates for the relationship between search rates and unconditional hit rates over a more limited range of search rates. The slopes in Table B6 are derived from a linear regression of $h_{p\ell}^{EB}$ on $s_{p\ell}^{EB}$ with location fixed effects using all trooper-location combinations.

characteristics, P(SEARCH | X_{it}). Figure B2 depicts a histogram of P(SEARCH | X_{it}) across stops.

We then characterize troopers by their mean value of P(SEARCH | X_{it}) after conditioning on stop time and location. We estimate the following Poisson pseudo-likelihood regression model:

$$\log(E(P(\text{SEARCH}|X_{it})|\psi_{p(i,t)\ell(i,t)}, \ell(i,t), \tau(t), m(t)) = \psi_{p(i,t)\ell(i,t)} + \lambda_{\ell(i,t)} + \omega_{\tau(t)} + \delta_{m(t)} + \epsilon_{it} \tag{B.3}$$

where $\psi_{p(i,t)\ell(i,t)}$ are trooper by location fixed effects. We estimate the model using the pseudo-maximum likelihood estimator of Correia et al. (2019). Figure B3 depicts a histogram of $\psi_{p\ell}$ across troopers. If the assignment of motorists to troopers conditional on stop time and location were indeed as good as random, $\psi_{p\ell}$ would only vary across troopers due to chance. Troopers with large and positive (negative) values of $\psi_{p\ell}$ are stopping motorists with characteristics that predict high (low) search rates (e.g. non-white, low-income men with criminal histories) relative to other troopers making stops at the same times and in the same locations. We rank trooper by location combinations by $|\psi_{p\ell}|$, where combinations with the largest absolute values are 'most selective'.

In Figure B4, we show that the slope of the pooled between-trooper SPC is stable if we exclude a varying proportion of troopers with compositions of stopped motorists who deviate most from their expected composition given the time and location of their stops. In Figure B5 we conduct a similar exercise for race-specific SPC slopes and find that slope estimates and their ordering across groups are stable when we vary the set of included troopers.

B.4.2 Troopers Vary in Screening Ability

Our finding that average and marginal hit rates are similar is consistent with Knowles et al. (2001), who develop an equilibrium model where troopers decide whether or not to search motorists and motorists decide whether or not to carry contraband. They show that if troopers are not racially biased, all motorists must, in equilibrium, carry contraband with equal probability. In this model there is no inframarginality problem because there is no difference between hit rates for the marginal and average searched motorists.

However, there are at least two features of our setting that are inconsistent with the Knowles et al. (2001) framework. First, as we document in Section V.B, we find little evidence that motorists respond to variation in search risk by adjusting contraband carrying rates, at least in the range of search rates we observe. Second, as we show in this section, troopers vary systematically in their hit rates, implying variation in screening ability. This is inconsistent with Knowles et al. (2001), which assumes there is no screening.

Figure B10 documents that troopers vary in screening ability in two ways. Panel A plots adjusted search rates $\tilde{s}_{p\ell}$ against adjusted unconditional hit rates $\tilde{h}_{p\ell}$ for each trooper by location combination. Conditional on search rate, there is significant variation in unconditional hit rates.

This variation is not due to statistical noise alone. Panel B plots trooper by location hit rates in one randomly selected half of stops against the same trooper by location hit rate in the remaining half of stops. The estimated slope is 0.372, indicating that while some variation in hit rates is due to chance, there is systematic variation in hit rates across trooper by location combinations.

B.4.3 Racial Search Disparities by Trooper Race

In this section we examine differences in search behavior by trooper race. We identify trooper race using 2015 personnel records for 2,469 troopers accounting for 84% of stops. Table B10 documents search rates and hit rates by both motorist and trooper race.

We next measure differences in black-white and Hispanic-white search odds ratios by trooper race, accounting for stop and other motorist characteristics. We estimate logistic regression models analogous to equation (B.1) that include fixed effects for trooper race and interactions between motorist and trooper race. We limit the analysis to stops conducted by troopers that we identify as black, Hispanic, or white.

Table B11 presents coefficient estimates, where columns (1) through (5) are analogous to the same columns in Table B2. The black-white search disparity for black troopers is about 20% smaller than the same disparity for white troopers, and about 35% smaller than the same disparity for Hispanic troopers. The Hispanic-white disparity is similar for white and Hispanic troopers and smaller for black troopers.

Figure B1
Comparing Estimates of Trooper Search Rates

Note: This figure compares unadjusted trooper search rates ($s_{p\ell}$) to estimates of trooper search rates that adjust for additional stop and motorist characteristics ($\tilde{s}_{p\ell}$). Stops characteristics include the month and specific highway of the stop. Motorist characteristics include race, gender, log of neighborhood median income, vehicle-based expected log income, stop history, non-drug arrest history, and drug arrest history. The construction of trooper search rates is described in Section IV.B. The red dashed line is a 45° line. The slope of the best fit line is 0.99.

Figure B2

Distribution of $P(\text{SEARCH}|X_{it})$ Across Stops

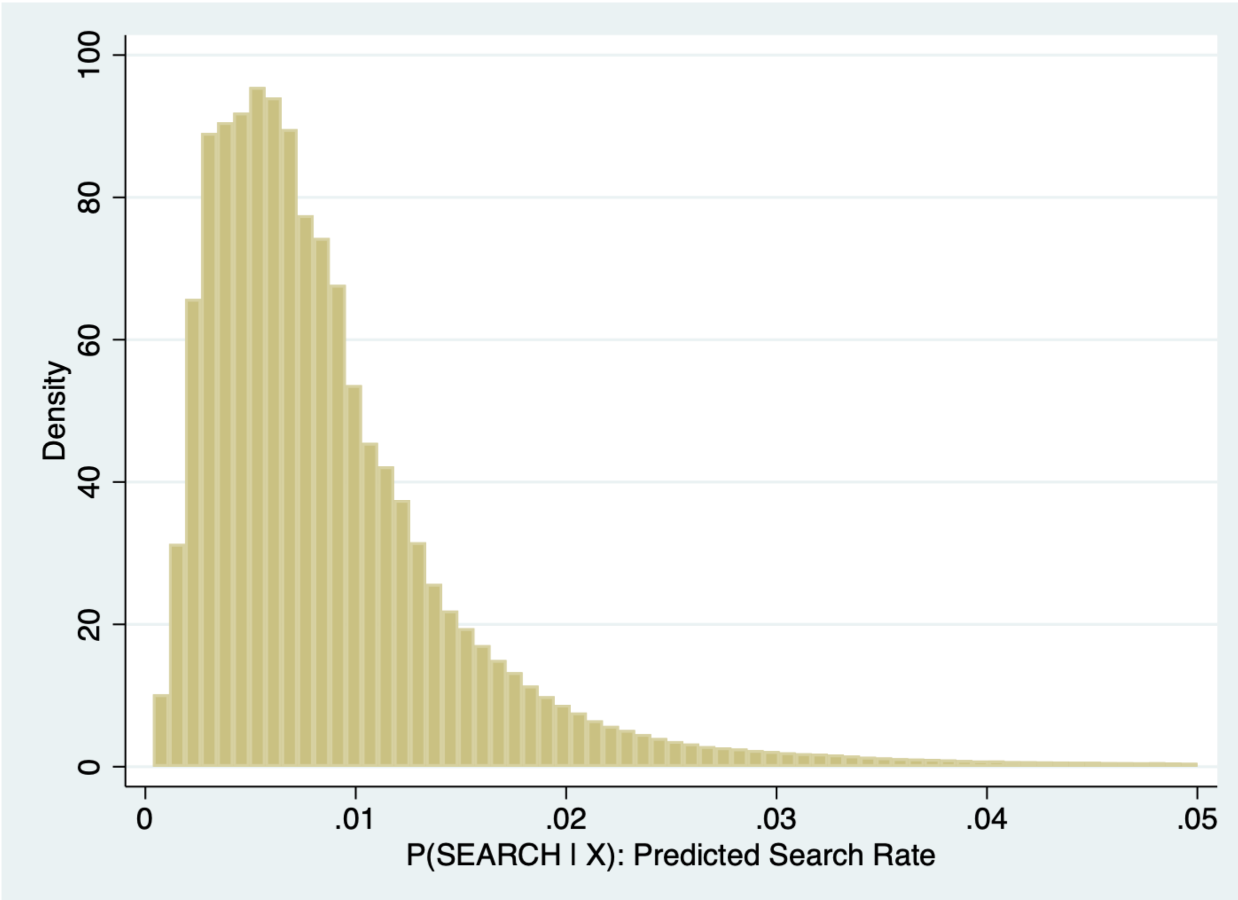

Note: This figure plots the distribution of P(SEARCH | X_{it}), search probability for each stop based on observable motorist characteristics. P(SEARCH | X_{it}) is estimated from equation (B.2) described in Section B.4.1.

Figure B3
Distribution of Motorist Selection Across Troopers

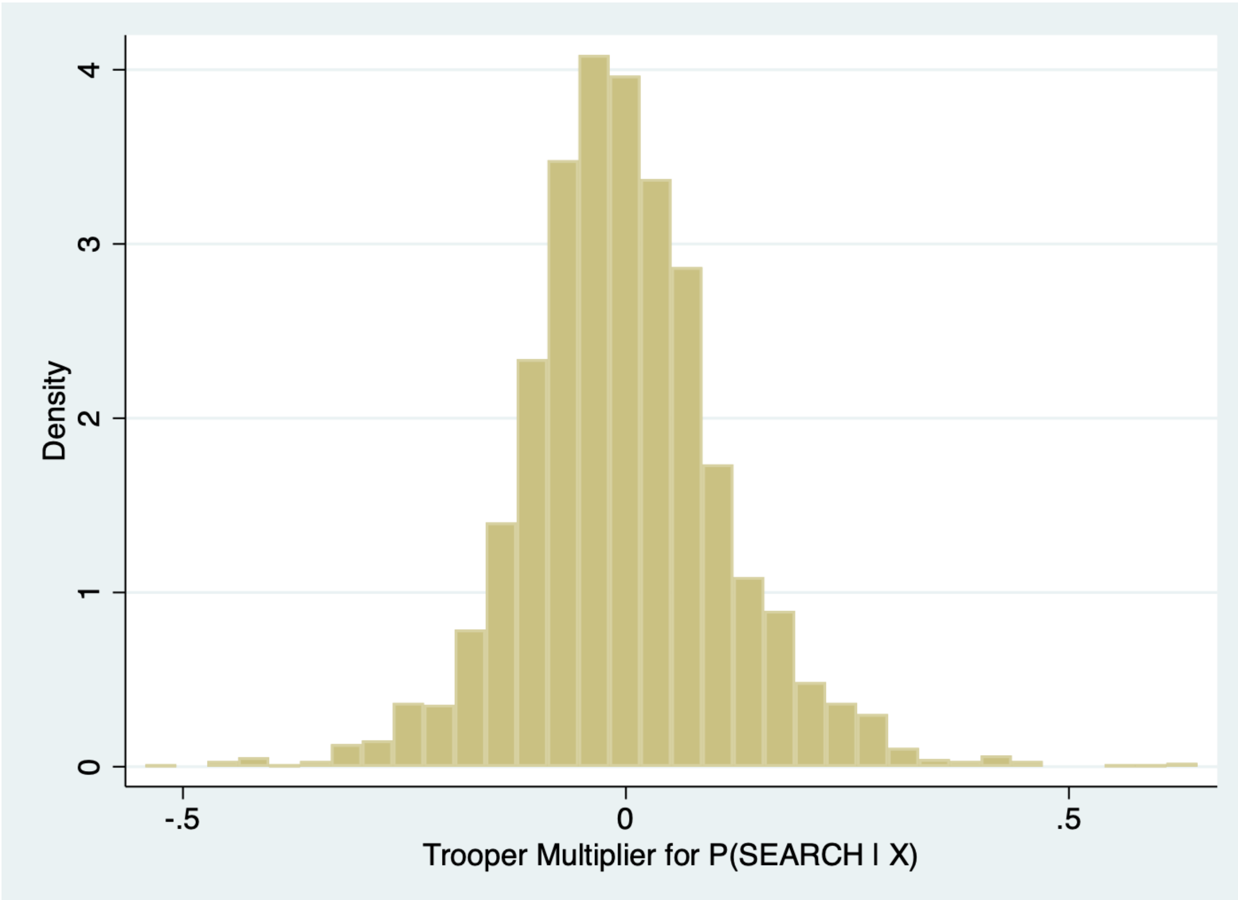

Note: This figure plots the distribution of $\psi_{p\ell}$ estimates derived from equation (B.3), described in more detail in Section B.4.1. The $\psi_{p\ell}$ terms are trooper by location fixed effects from a Poisson regression model for $P(\text{SEARCH}|X_{it})$, the search probability for each stop based on observable motorist characteristics. They summarize the degree to which motorist characteristics for those stopped by a given trooper in a given location deviate from what one would expect based on the time and location of their stops alone.

Figure B4

Stability of Pooled Search Productivity Curve Slope For Varying Trooper Exclusions

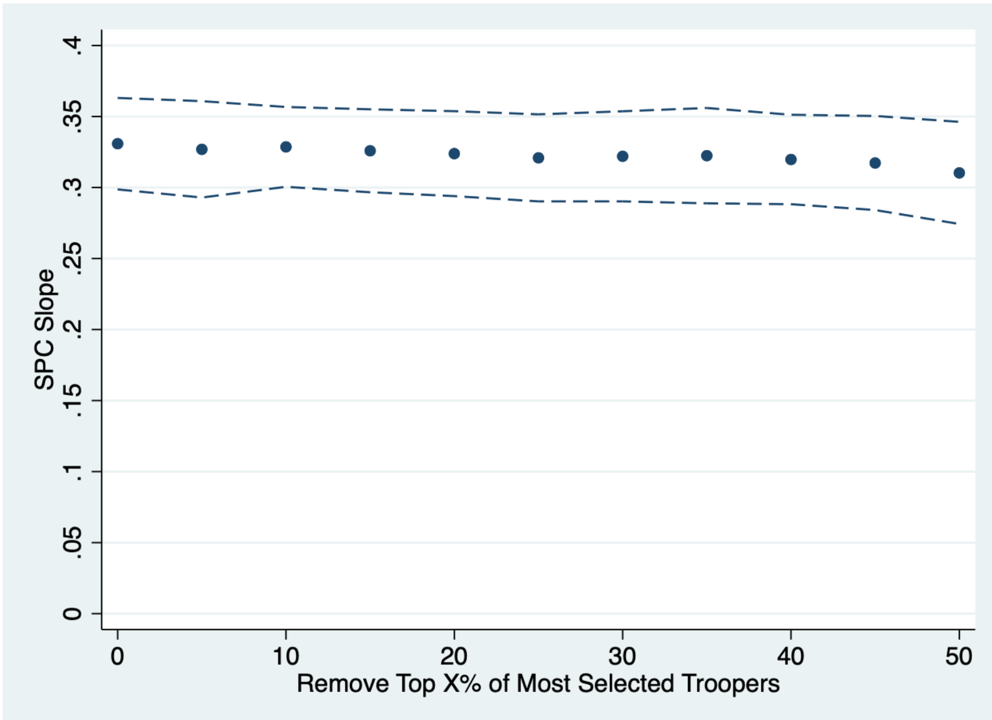

Note: This figure plots the slope of the relationship between trooper search rates ($\tilde{s}^r_{p\ell}$) and unconditional hit rates ($\tilde{h}^r_{p\ell}$) for varying samples of troopers using the fixed effects (FE) approach described in IV.D. For varying X, we remove the X% of troopers with compositions of stopped motorists that deviate most from their expected composition given the time and location of their stops. We discuss how we identify these troopers in more detail in Section B.4.1.

Figure B5
Stability of Race-Specific Search Productivity Curve Slopes For Varying Trooper Exclusions

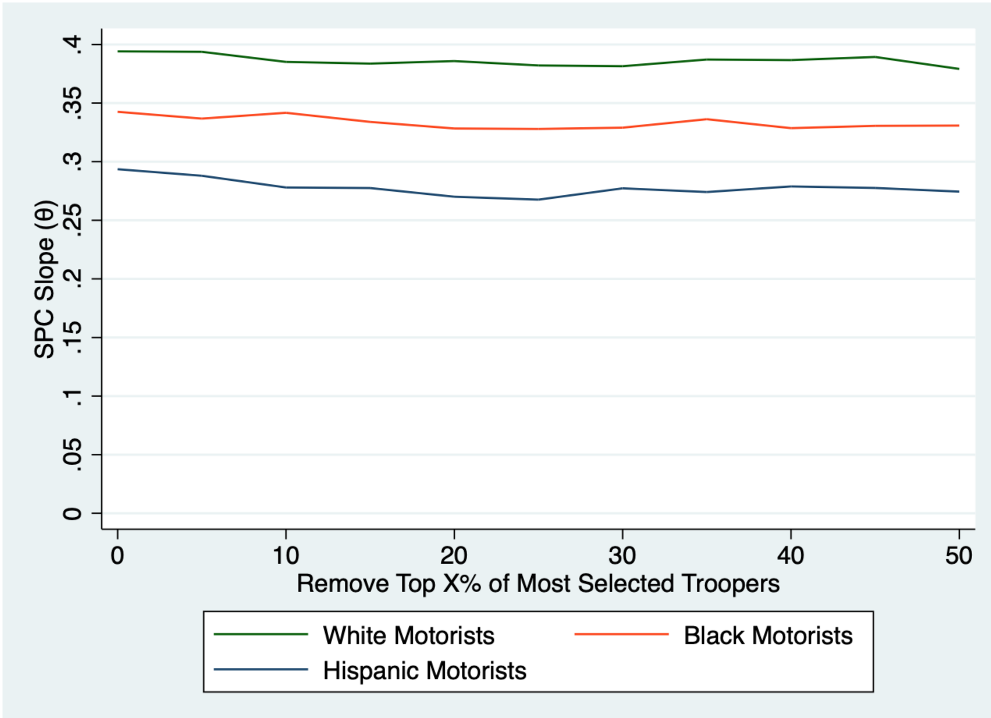

Note: This figure plots the slope of the relationship between trooper search rates ($\tilde{s}^r_{p\ell}$) and unconditional hit rates ($\tilde{h}^r_{p\ell}$) by motorist race and for varying samples of troopers using the fixed effects (FE) approach described in IV.D. For varying X, we remove the $X\%$ of troopers with compositions of stopped motorists that deviate most from their expected composition given the time and location of their stops. We discuss how we identify these troopers in more detail in Appendix B.4.1.

Figure B6
Between-Trooper Search Productivity Curve, Arrests

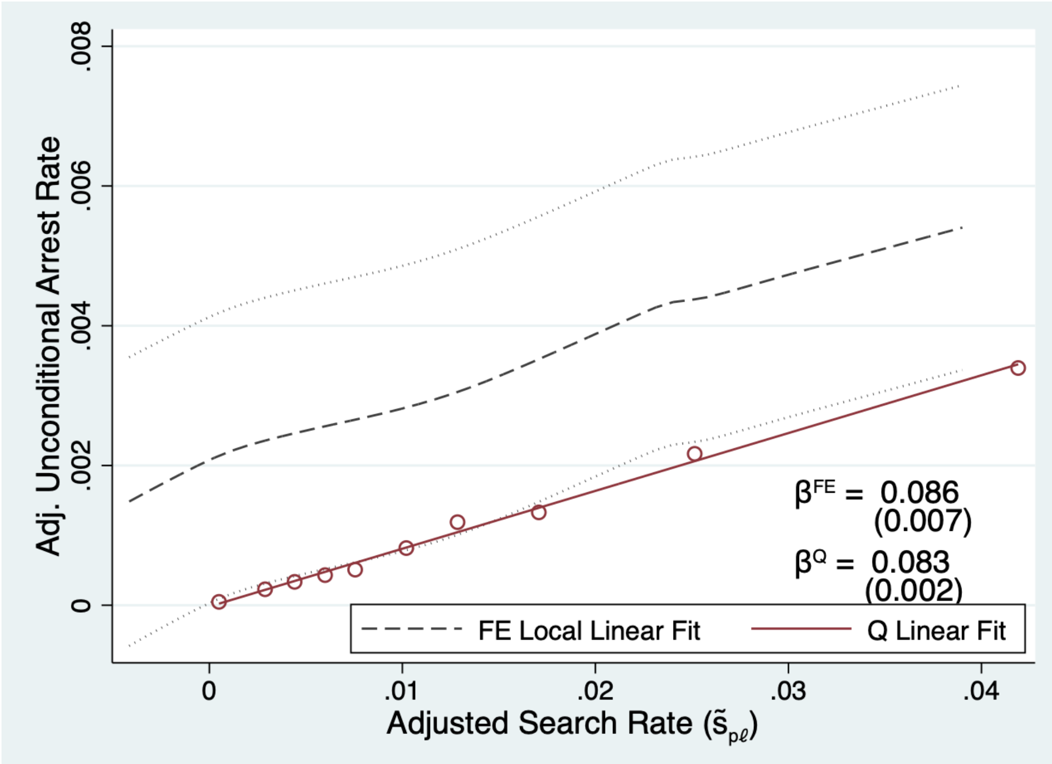

Note: In this figure we plot adjusted trooper unconditional arrest rates against trooper search rates ($\tilde{s}_{p\ell}$) using two approaches described in Section IV.D: the fixed effects (FE) approach and the quantile (Q) approach. In the fixed effects approach we net out location fixed effects from both $\tilde{s}_{p\ell}$ and adjusted trooper unconditional arrest rates and plot the residuals. From the FE approach, the figure includes 95% confidence bands for the local linear relationship between adjusted trooper search rates and unconditional arrest rates and the best linear fit and its slope. The local linear fit is derived using a Gaussian kernel with a rule-of-thumb bandwidth. In the Q approach within locations we divide troopers into quantiles by search rate, group quantiles across locations, and then plot the relationship between search rates and unconditional arrest rates across quantiles. From the Q approach, the figure includes the mean values for each decile and the best linear fit and its slope.

Figure B7

Between-Trooper Search Productivity Curve, by Motorist Race, Arrests

(a) White Motorists

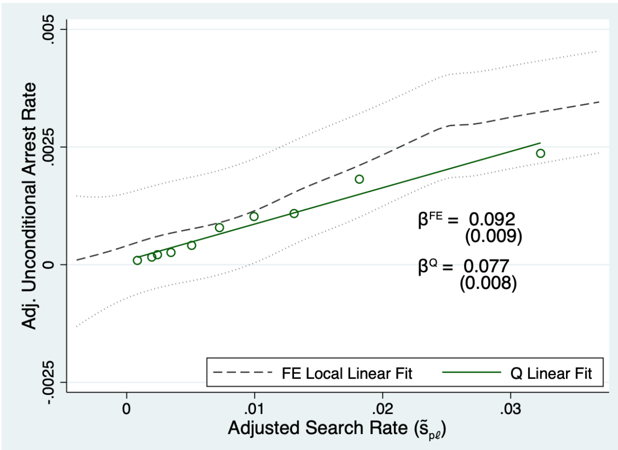

(b) Black Motorists

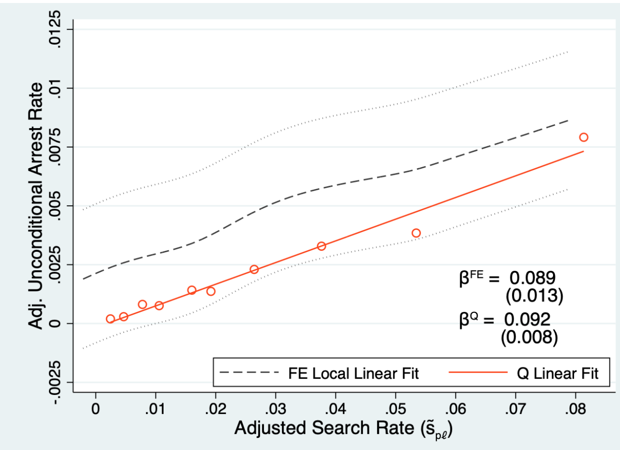

(c) Hispanic Motorists

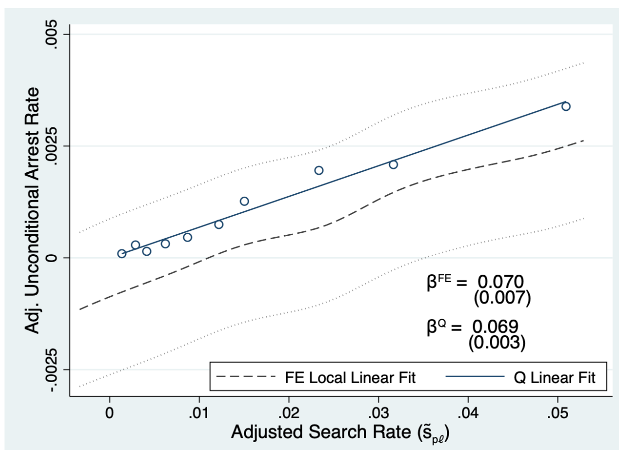

Note: In this figure we plot adjusted trooper unconditional arrest rates against trooper search rates ($\tilde{s}_{p\ell}$) using two approaches described in Section IV.D: the fixed effects (FE) approach and the quantile (Q) approach. In the fixed effects approach we net out location fixed effects from both $\tilde{s}_{p\ell}$ and adjusted trooper unconditional arrest rates and plot the residuals. From the FE approach, the figure includes 95% confidence bands for the local linear relationship between adjusted trooper search rates and unconditional arrest rates and the best linear fit and its slope. The local linear fit is derived using a Gaussian kernel with a rule-of-thumb bandwidth. In the Q approach within locations we divide troopers into quantiles by search rate, group quantiles across locations, and then plot the relationship between search rates and unconditional arrest rates across quantiles. From the Q approach, the figure includes the mean values for each decile and the best linear fit and its slope. Panel A, Panel B, and Panel C plot the search productivity curve (SPC) for white motorists, black motorists, and Hispanic motorists, respectively.

Figure B8
Map of Highway-Border Intersections Included in RD Analysis

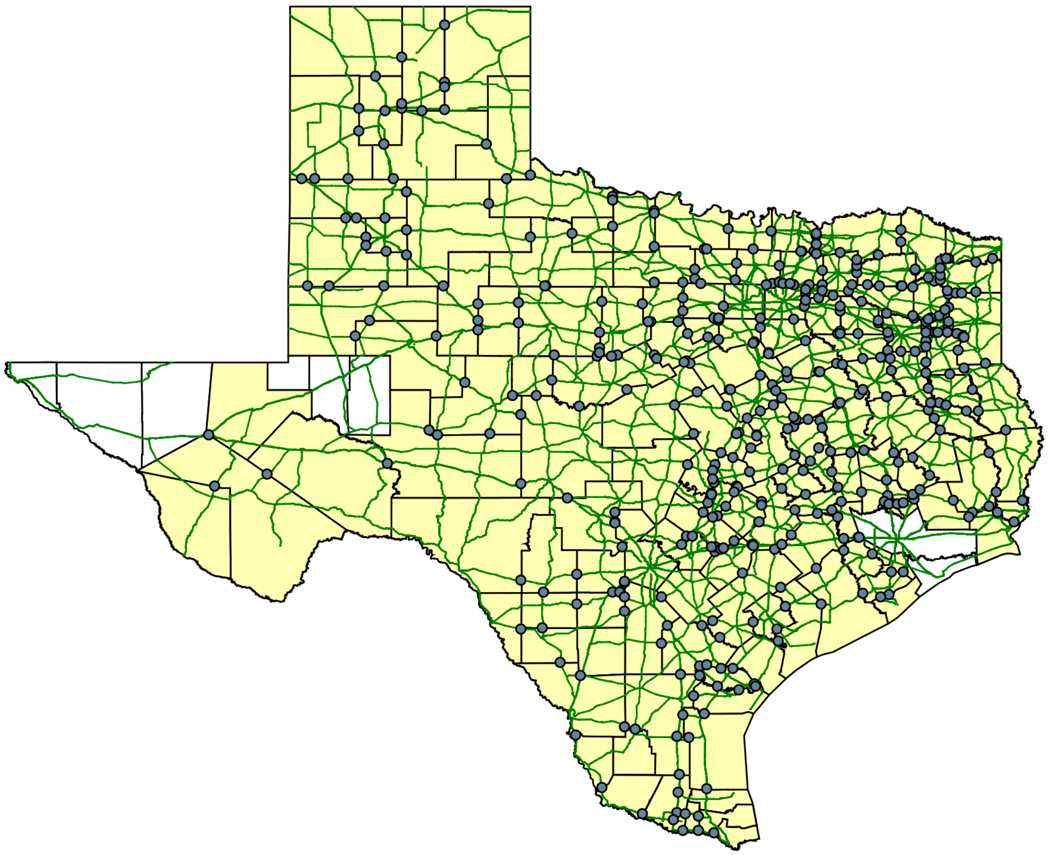

Note: This map depicts in blue the 424 intersections between state/interstate highways and sergeant area borders that define the spatial RD sample. The set of state and interstate highways associated with these intersections is superimposed in green. Sergeant areas included in the RD sample are shaded yellow.

Figure B9
Within-Trooper Variation in Search Rates

(a) First Stage

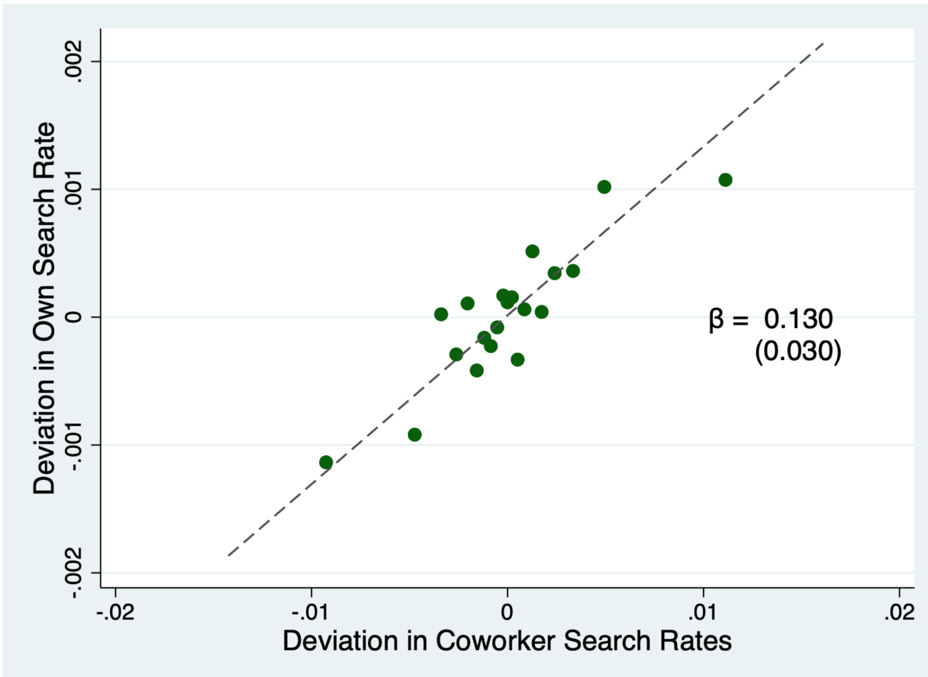

(b) Reduced Form

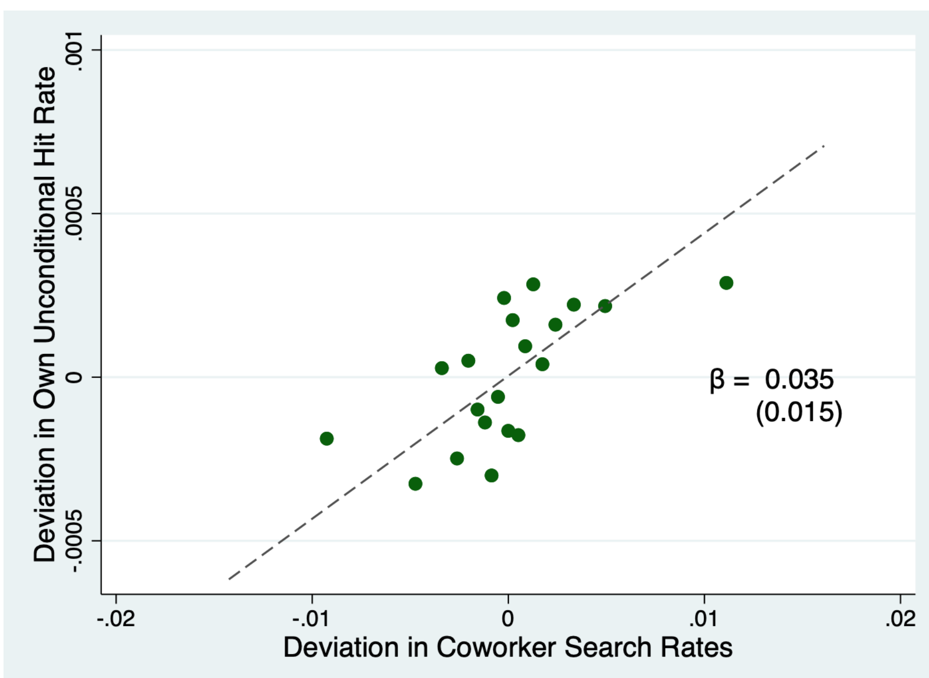

Note: These figures summarize the relationship between deviations in coworker search rates and deviations in troopers' own search rates (Panel A) and unconditional hit rates (Panel B). Observations are at the trooper-year level. Both plots include bin scatters where observations at grouped into ventiles based on the deviation in coworker search rates. The construction of these deviations is described in Section V.A.

Figure B10
Troopers Vary in Screening Ability

(a)

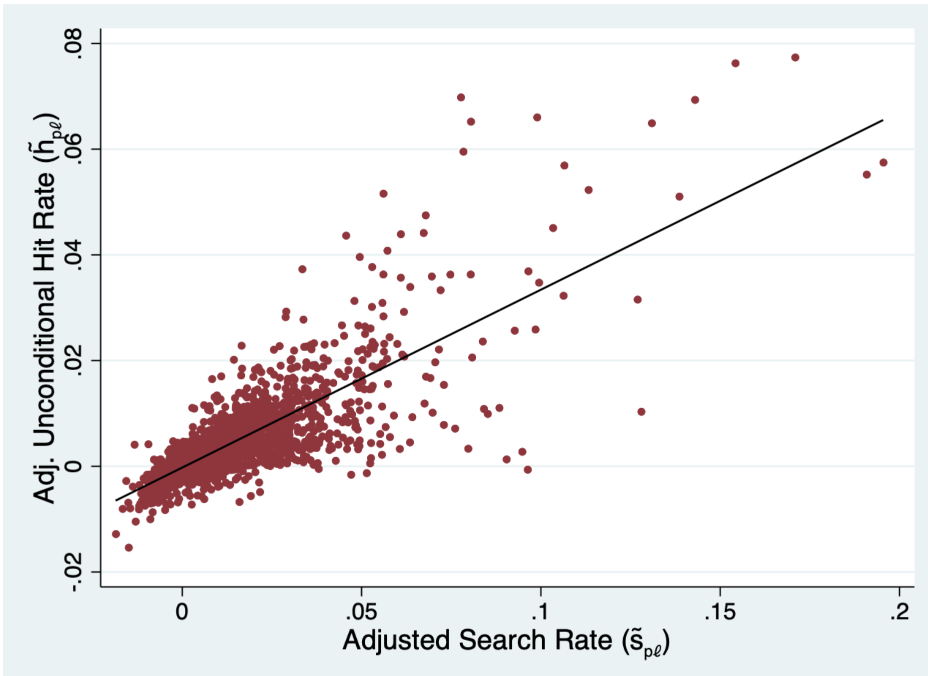

(b) Split Sample Hit Rates

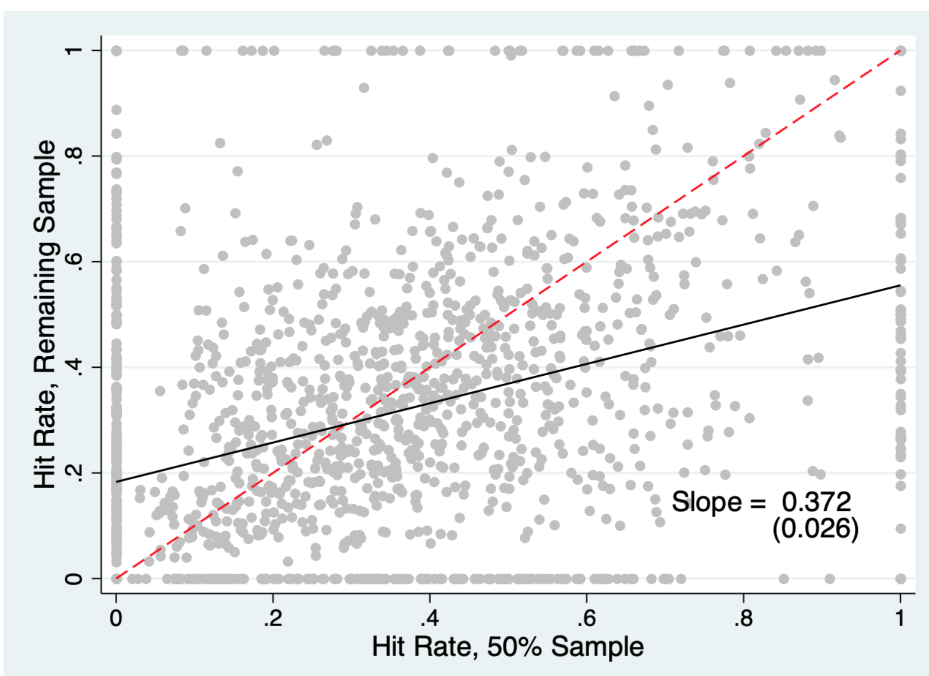

Note: These figures establish that hit rates vary systematically across troopers. Panel A plots adjusted search rates ($\tilde{s}_{p\ell}$) against adjusted unconditional hit rates ($\tilde{h}_{p\ell}$), where each marker represents a trooper by location pair. Panel B plots trooper by location hit rates in one randomly selected half of stops against the same trooper by location hit rates in the remaining half of stops.

Table A1
Sample Selection

Sample step	Observations Dropped	Observations Remaining
1. All stops conducted by Texas Highway Patrol between 2009 and 2015		15,956,460
2. Retains stops made on state and interstate highways	4,134,057	11,822,403
3. Drop stops with missing location information	532,933	11,289,470
4. Drop stops in the state capitol region	63	11,289,407
5. Retain stops of motorists with Texas addresses	1,624,350	9,665,057
6. Retain stops of passenger cars, pick-up trucks, and SUVs	779,841	8,885,216
7. Drop stops with missing motorist information	587,605	8,297,611
8. Retain stops of motorists that are white, black, or Hispanic	171,790	8,125,821
6. Retain stops with at least one associated speeding violation	3,187,592	4,938,229
3. Drop stops with missing trooper ID or stop outcomes	6,907	4,931,332

Table B1
Detailed Search Outcomes by Motorist Race

	Black	Hispanic	White	All
Consent	45.19	57.71	48.70	51.80
Incident to Arrest	4.223	4.993	5.079	4.881
Inventory	8.315	11.70	11.74	11.08
Probable Cause	42.27	25.60	34.48	32.24
Conditional on Contraband:				
Currency	0.501	1.135	0.168	0.564
Drugs	56.68	49.31	51.50	51.80
Weapon	5.217	3.004	3.751	3.792
Other	37.61	46.55	44.58	43.85
Arrest	24.31	23.83	25.04	24.48
Felony Arrest	9.254	8.122	7.165	7.911
Charge Severity (Days)				
Mean	98.35	88.07	90.70	91.35
90th Percentile	253.83	243.11	246.61	246.61

This table summarizes detailed search outcomes by motorist race. ‘Charge Severity’ refers to the average incarceration sentence associated with conviction for that arrest charge. Charge Severity is set to zero for searches that do not lead to an arrest.

Table B5
Motorist Selection into Stops by Trooper Unconditional Hit Rate

				Excluding Most Selective Troopers		
	100 × CONTRABAND_{it}	100 × $h_{p\ell}^{-it}$	100 × $\tilde{h}_{p\ell}^{-it}$	100 × CONTRABAND_{it}	100 × $h_{p\ell}^{-it}$	100 × $\tilde{h}_{p\ell}^{-it}$
	(1)	(2)	(3)	(4)	(5)	(6)
Black	0.273	0.042	0.029	0.221	0.017	0.010
	(0.028)	(0.010)	(0.010)	(0.024)	(0.005)	(0.005)
Hispanic	0.020	0.010	0.006	0.014	0.001	-0.001
	(0.009)	(0.004)	(0.003)	(0.009)	(0.003)	(0.003)
Female	-0.157	-0.009	-0.007	-0.140	-0.004	-0.004
	(0.010)	(0.002)	(0.002)	(0.009)	(0.001)	(0.001)
Log Median Income	-0.040	-0.002	-0.002	-0.031	0.002	0.000
	(0.009)	(0.002)	(0.002)	(0.009)	(0.002)	(0.002)
Expected Log Income Given	-0.092	-0.011	-0.011	-0.074	-0.004	-0.004
Vehicle (Standardized)	(0.005)	(0.002)	(0.002)	(0.005)	(0.001)	(0.001)
1-2 Prior Non-Drug Arrests	0.256	0.006	0.005	0.228	0.000	0.000
	(0.029)	(0.002)	(0.002)	(0.030)	(0.002)	(0.002)
3+ Prior Non-Drug Arrests	0.406	0.010	0.009	0.352	0.005	0.005
	(0.051)	(0.003)	(0.003)	(0.051)	(0.002)	(0.002)
1 Prior Drug Arrest	1.400	0.011	0.009	1.280	-0.000	-0.001
	(0.086)	(0.003)	(0.003)	(0.085)	(0.002)	(0.002)
2+ Prior Drug Arrests	2.144	0.015	0.012	1.806	0.002	0.000
	(0.123)	(0.005)	(0.005)	(0.118)	(0.003)	(0.003)
Prior Stop, No Search	-0.050	-0.006	-0.004	-0.034	-0.004	-0.002
	(0.009)	(0.002)	(0.002)	(0.009)	(0.002)	(0.002)
Prior Search, No Contraband	0.337	0.021	0.019	0.361	0.007	0.007
	(0.072)	(0.006)	(0.005)	(0.075)	(0.004)	(0.004)
Prior Search, Contraband	7.611	0.094	0.092	6.619	0.062	0.063
	(0.609)	(0.014)	(0.014)	(0.599)	(0.012)	(0.012)
Location by Time FEs	✓	✓	✓	✓	✓	✓
Joint F-Statistic	48.74	7.59	7.32	43.87	5.86	5.62
Observations	3,280,250	3,280,250	3,280,171	2,739,955	2,739,955	2,739,899

This table presents coefficients from estimates of equation (7), where in columns (2) and (3) we replace the outcome SEARCH_{it} with CONTRABAND_{it}, $h_{p(i,t)\ell(i,t)}^{-it}$ and $\tilde{h}_{p(i,t)\ell(i,t)}^{-it}$, leave-out trooper unconditional hit rates corresponding to the trooper who conducted the stop. Columns (4)–(6) exclude stops conducted by the 20% of troopers with the most selected set of stopped motorists. Standard errors are clustered at the motorist level. 'Joint F-Statistic' refers to an F-test for the joint significance of all motorist characteristics.

Table B2

What Predicts Odds of Searches and Contraband Yield?

Outcome:	Motorist/Vehicle Searched					Contraband Found | Search				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Black	3.00	2.70	2.20	1.87	1.86	0.86	0.85	0.91	0.90	0.91
	(0.04)	(0.04)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Hispanic	1.58	1.68	1.45	1.45	1.44	0.59	0.68	0.71	0.72	0.73
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)	(0.02)
Female	0.38	0.39	0.41	0.50	0.50	0.98	0.99	0.99	1.03	1.03
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)
Log Median Income			0.68	0.74	0.74			1.26	1.27	1.27
			(0.01)	(0.01)	(0.01)			(0.03)	(0.03)	(0.03)
Expected Log Income Given			0.65	0.68	0.69			1.00	1.01	1.01
Vehicle (Standardized)			(0.00)	(0.00)	(0.00)			(0.01)	(0.01)	(0.01)
1–2 Prior Non-Drug Arrests				1.79	1.76				1.00	1.00
				(0.03)	(0.03)				(0.03)	(0.03)
3+ Prior Non-Drug Arrests				2.06	2.00				0.87	0.87
				(0.04)	(0.04)				(0.03)	(0.03)
1 Prior Drug Arrest				3.05	2.74				1.48	1.41
				(0.06)	(0.05)				(0.06)	(0.05)
2+ Prior Drug Arrest				4.04	3.37				1.47	1.36
				(0.08)	(0.07)				(0.06)	(0.05)
Prior Stop, No Search					0.82					1.10
					(0.01)					(0.03)
Prior Search, No Contraband					2.15					0.63
					(0.06)					(0.04)
Prior Search, Contraband					4.83					3.45
					(0.13)					(0.18)
Time FEs		✓	✓	✓	✓		✓	✓	✓	✓
Location FEs		✓	✓	✓	✓		✓	✓	✓	✓
Month FEs		✓	✓	✓	✓		✓	✓	✓	✓
White Mean	0.757	0.757	0.757	0.757	0.757	37.24	37.24	37.24	37.24	37.24
Observations	4,931,332	4,438,416	4,438,416	4,438,416	4,438,416	52,203	46,485	46,485	46,485	46,485

This table presents odds ratio estimates for the logistic regression model (B.1). In columns (1) through (5) the outcome is SEARCH_{it}, an indicator of whether the stop of motorist i at time t led to a search. In columns (6) through (10) the outcome is CONTRABAND_{it}, an indicator for whether a search yields contraband. For these specifications, the sample is limited to stops that result in a search (i.e., where $\text{SEARCH}_{it} = 1$). Standard errors are clustered at the motorist level.

TABLE B3

TRAFFIC STOP DESCRIPTIVE STATISTICS, POOLED SPC SAMPLE

	All Stops				All Searches			
	Black	Hispanic	White	All	Black	Hispanic	White	All
Black	100	0	0	10.40	100	0	0	21.10
Hispanic	0	100	0	35.22	0	100	0	39.74
White	0	0	100	54.38	0	0	100	39.16
Female	39.25	32.79	38.16	36.38	17.12	15.75	20.87	18.04
Log Median Income	10.71	10.73	10.97	10.86	10.59	10.62	10.89	10.72
	(0.501)	(0.497)	(0.467)	(0.495)	(0.491)	(0.469)	(0.473)	(0.495)
Expected Log Income Given	-0.149	-0.063	0.087	0.009	-0.525	-0.456	-0.365	-0.435
Vehicle (Standardized)	(1.025)	(0.983)	(1.023)	(1.013)	(0.884)	(0.849)	(0.894)	(0.876)
Stop History:								
No Prior Stops	61.74	59.74	57.84	58.91	59.39	58.51	54.95	57.30
Prior Stop, No Search	36.60	38.91	41.31	39.97	31.22	33.81	36.66	34.38
Prior Search, No Contraband	1.078	1.027	0.528	0.761	4.560	3.961	3.625	3.956
Prior Search, Contraband	0.584	0.325	0.323	0.351	4.835	3.719	4.759	4.361
Non-Drug Arrest History:								
No Prior Non-Drug Arrests	87.04	89.42	93.34	91.30	65.28	72.12	71.63	70.49
1-2 Prior Non-Drug Arrests	7.306	6.940	4.422	5.609	15.18	14.71	14.61	14.77
3+ Prior Non-Drug Arrests	5.657	3.641	2.237	3.087	19.53	13.17	13.76	14.74
Drug Arrest History:								
No Prior Drug Arrests	93.95	96.30	97.45	96.68	73.55	81.75	79.19	79.02
1 Prior Drug Arrest	2.810	2.173	1.429	1.835	10.31	9.075	9.729	9.591
2+ Prior Drug Arrests	3.240	1.531	1.119	1.485	16.14	9.172	11.08	11.39
Search Rate	2.244	1.248	0.796	1.106	100	100	100	100
Unconditional Hit Rate	0.763	0.337	0.299	0.361	33.70	26.77	37.22	32.33
Observations	341,083	1,155,385	1,783,782	3,280,250	7,653	14,414	14,205	36,272

Sample restrictions are described in Section II and Section IV.B.

Table B4

Traffic Stop Descriptive Statistics, Race-Specific SPC Sample

	All Stops				All Searches			
	Black	Hispanic	White	All	Black	Hispanic	White	All
Black	100	0	0	12.94	100	0	0	25.02
Hispanic	0	100	0	28.32	0	100	0	33.39
White	0	0	100	58.73	0	0	100	41.59
Female	40.30	33.01	39.22	37.60	17.28	15.63	21.32	18.41
Log Median Income	10.71	10.81	10.97	10.89	10.59	10.68	10.91	10.75
	(0.503)	(0.494)	(0.468)	(0.490)	(0.490)	(0.470)	(0.470)	(0.494)
Expected Log Income Given	-0.155	-0.017	0.086	0.025	-0.531	-0.422	-0.366	-0.426
Vehicle (Standardized)	(1.030)	(1.018)	(1.034)	(1.032)	(0.880)	(0.869)	(0.906)	(0.890)
Stop History:								
No Prior Stops	61.76	63.97	58.73	60.60	59.57	62.04	55.84	58.84
Prior Stop, No Search	36.57	34.98	40.42	38.38	31.21	31.35	36.01	33.25
Prior Search, No Contraband	1.084	0.738	0.530	0.661	4.576	3.039	3.694	3.696
Prior Search, Contraband	0.584	0.319	0.326	0.357	4.639	3.571	4.453	4.205
Non-Drug Arrest History:								
No Prior Non-Drug Arrests	87.01	90.37	93.31	91.66	65.69	73.68	71.63	70.83
1-2 Prior Non-Drug Arrests	7.354	6.471	4.445	5.396	15.28	14.44	14.71	14.76
3+ Prior Non-Drug Arrests	5.631	3.157	2.243	2.940	19.03	11.88	13.66	14.41
Drug Arrest History:								
No Prior Drug Arrests	93.96	96.67	97.38	96.74	73.76	82.33	78.83	78.73
1 Prior Drug Arrest	2.807	1.949	1.464	1.775	10.30	8.431	9.875	9.500
2+ Prior Drug Arrests	3.232	1.384	1.154	1.488	15.94	9.235	11.30	11.77
Search Rate	2.239	1.365	0.820	1.158	100	100	100	100
Unconditional Hit Rate	0.766	0.413	0.308	0.397	34.04	30.00	37.32	34.06
Observations	283,012	619,360	1,284,325	2,226,319	6,337	8,457	10,532	25,326

Sample restrictions are described in Section II and Section IV.B.

Table B6
Robustness of Search Productivity Curve Slope Estimates

	Pooled	White Motorists	Black Motorists	Hispanic Motorists
Unadjusted Rates	0.347	0.375	0.400	0.312
	(0.016)	(0.022)	(0.028)	(0.022)
Covariate-Adjusted Rates	0.345	0.372	0.398	0.318
	(0.016)	(0.022)	(0.028)	(0.021)
EB-Adjusted Rates	0.348	0.381	0.404	0.309
	(0.016)	(0.024)	(0.029)	(0.023)
Split Sample (2SLS), First	0.348	0.366	0.443	0.335
	(0.017)	(0.031)	(0.037)	(0.026)
Split Sample (2SLS), Second	0.351	0.383	0.378	0.334
	(0.019)	(0.026)	(0.036)	(0.029)

This table presents the slope of the relationship between trooper search rates and unconditional hit rates conditional on location fixed effects for several specifications and for varying samples of motorists. Trooper-by-location combinations are weighted by number of stops. For the split sample models, we randomly split stops into two samples and estimate $\tilde{s}_{p\ell}$ and $\tilde{h}_{p\ell}$ separately in each sample. In each sample, we regress $\tilde{h}_{p\ell}$ on $\tilde{s}_{p\ell}$, instrumenting for $\tilde{s}_{p\ell}$ using its pair estimate from the other sample.

TABLE B7

TRAFFIC STOP DESCRIPTIVE STATISTICS, POOLED WITHIN-MOTORIST SAMPLE

	All Stops				All Searches			
	Black	Hispanic	White	All	Black	Hispanic	White	All
Black	100	0	0	9.67	100	0	0	19.91
Hispanic	0	100	0	32.88	0	100	0	37.23
White	0	0	100	57.45	0	0	100	42.86
Female	34.70	27.40	33.70	31.72	14.74	12.98	17.68	15.35
Log Median Income	10.68	10.71	10.92	10.83	10.57	10.62	10.89	10.73
	(0.494)	(0.483)	(0.447)	(0.476)	(0.491)	(0.474)	(0.455)	(0.490)
Expected Log Income Given	-0.196	-0.083	0.055	-0.015	-0.556	-0.418	-0.328	-0.407
Vehicle (Standardized)	(1.000)	(0.940)	(0.964)	(0.964)	(0.897)	(0.862)	(0.878)	(0.880)
Stop History:								
No Prior Stops	49.00	46.64	46.90	47.02	44.55	41.22	42.98	42.64
Prior Stop, No Search	48.67	51.39	52.04	51.50	38.30	42.63	42.65	41.78
Prior Search, No Contraband	1.565	1.527	0.667	1.037	7.131	7.284	5.582	6.524
Prior Search, Contraband	0.770	0.444	0.389	0.444	10.02	8.869	8.783	9.060
Non-Drug Arrest History:								
No Prior Non-Drug Arrests	83.97	85.94	91.51	88.95	58.01	66.54	66.62	64.87
1-2 Prior Non-Drug Arrests	9.079	9.123	5.654	7.126	18.83	17.27	18.65	18.17
3+ Prior Non-Drug Arrests	6.954	4.934	2.835	3.923	23.16	16.20	14.74	16.96
Drug Arrest History:								
No Prior Drug Arrests	92.92	95.18	96.98	96.00	68.19	75.84	76.40	74.56
1 Prior Drug Arrest	3.309	2.869	1.706	2.243	11.38	11.35	10.64	11.05
2+ Prior Drug Arrests	3.772	1.952	1.310	1.759	20.43	12.81	12.95	14.39
Search Rate	1.858	1.023	0.674	0.903	100	100	100	100
Unconditional Hit Rate	0.625	0.290	0.261	0.305	33.17	27.89	38.26	33.39
Observations	67,157	228,256	398,833	694,246	1,248	2,334	2,687	6,269

Sample restrictions are described in Section II, Section IV.B, and Section IV.E.1.

TABLE B8

TRAFFIC STOP DESCRIPTIVE STATISTICS, RACE-SPECIFIC WITHIN-MOTORIST SAMPLE

	All Stops				All Searches			
	Black	Hispanic	White	All	Black	Hispanic	White	All
Black	100	0	0	13.20	100	0	0	26.08
Hispanic	0	100	0	22.75	0	100	0	26.13
White	0	0	100	64.04	0	0	100	47.79
Female	36.36	29.50	35.98	34.55	15.23	15.09	19.81	17.38
Log Median Income	10.66	10.81	10.91	10.85	10.58	10.73	10.90	10.77
	(0.494)	(0.476)	(0.443)	(0.465)	(0.483)	(0.464)	(0.456)	(0.484)
Expected Log Income Given	-0.217	-0.049	0.039	-0.015	-0.575	-0.391	-0.329	-0.410
Vehicle (Standardized)	(0.999)	(0.989)	(0.975)	(0.985)	(0.875)	(0.902)	(0.898)	(0.899)
Stop History:								
No Prior Stops	48.52	50.72	47.43	48.32	46.01	45.48	42.53	44.21
Prior Stop, No Search	49.11	47.81	51.49	50.34	37.70	37.94	42.94	40.27
Prior Search, No Contraband	1.612	0.991	0.676	0.871	7.348	6.164	5.578	6.193
Prior Search, Contraband	0.754	0.479	0.411	0.471	8.946	10.41	8.948	9.331
Non-Drug Arrest History:								
No Prior Non-Drug Arrests	83.77	87.21	91.40	89.44	57.51	67.69	66.59	64.51
1-2 Prior Non-Drug Arrests	9.179	8.529	5.687	6.795	20.02	16.47	18.59	18.41
3+ Prior Non-Drug Arrests	7.055	4.261	2.912	3.766	22.47	15.83	14.82	17.08
Drug Arrest History:								
No Prior Drug Arrests	92.89	95.76	96.85	96.08	67.31	74.60	75.48	73.12
1 Prior Drug Arrest	3.308	2.433	1.780	2.131	12.03	10.31	11.45	11.30
2+ Prior Drug Arrests	3.799	1.806	1.373	1.792	20.66	15.09	13.07	15.58
Search Rate	1.853	1.078	0.700	0.938	100	100	100	100
Unconditional Hit Rate	0.643	0.400	0.264	0.345	34.08	36.77	37.42	36.38
Observations	50,672	87,323	245,777	383,772	939	941	1,721	3,601

Sample restrictions are described in Section II, Section IV.B, and Section IV.E.1.

Table B9

Traffic Stop Descriptive Statistics, Spatial RD Sample

	All Stops				All Searches			
	Black	Hispanic	White	All	Black	Hispanic	White	All
Black	100	0	0	11.16	100	0	0	22.96
Hispanic	0	100	0	33.27	0	100	0	38.91
White	0	0	100	55.57	0	0	100	38.13
Female	39.52	33.74	38.96	37.29	15.89	14.72	21.18	17.46
Log Median Income	10.73	10.76	10.98	10.88	10.59	10.64	10.92	10.73
	(0.502)	(0.508)	(0.472)	(0.502)	(0.491)	(0.477)	(0.468)	(0.498)
Expected Log Income Given	-0.133	-0.032	0.112	0.037	-0.517	-0.439	-0.354	-0.425
Vehicle (Standardized)	(1.036)	(1.016)	(1.046)	(1.039)	(0.877)	(0.869)	(0.928)	(0.896)
Stop History:								
No Prior Stops	62.03	61.16	58.46	59.76	61.45	60.46	55.88	58.94
Prior Stop, No Search	36.26	37.52	40.68	39.14	29.47	32.88	35.91	33.25
Prior Search, No Contraband	1.115	1.012	0.540	0.761	4.472	3.482	3.650	3.773
Prior Search, Contraband	0.594	0.311	0.317	0.346	4.607	3.180	4.558	4.033
Non-Drug Arrest History:								
No Prior Non-Drug Arrests	87.46	90.25	93.47	91.73	65.49	74.79	72.91	71.94
1-2 Prior Non-Drug Arrests	6.960	6.432	4.347	5.332	14.01	13.45	13.58	13.63
3+ Prior Non-Drug Arrests	5.579	3.321	2.180	2.939	20.50	11.76	13.51	14.44
	(22.95)	(17.92)	(14.60)	(16.89)	(40.38)	(32.22)	(34.19)	(35.15)
Drug Arrest History:								
No Prior Drug Arrests	94.27	96.61	97.48	96.83	75.54	83.21	78.65	79.71
1 Prior Drug Arrest	2.703	2.005	1.421	1.758	9.213	8.426	9.732	9.105
2+ Prior Drug Arrests	3.027	1.380	1.100	1.408	15.25	8.362	11.61	11.18
Search Rate	2.246	1.277	0.749	1.092	100	100	100	100
Unconditional Hit Rate	0.759	0.313	0.282	0.345	33.27	24.32	37.24	31.30
Observations	165,237	492,524	822,611	1,480,372	3,712	6,290	6,165	16,167

Sample restrictions are described in Section II and Section IV.E.2. These statistics refer to the first stop for each sequential pair of stops included in the analysis.

Table B10
Search and Hit Rates by Motorist and Trooper Race

	White Troopers	Black Troopers	Hispanic Troopers
All Motorists			
Search Rate	1.15	0.83	0.89
Hit Rate	34.4	27.9	26.8
White Motorists			
Search Rate	0.81	0.62	0.62
Hit Rate	38.4	30.5	36.9
Black Motorists			
Search Rate	2.38	1.49	2.22
Hit Rate	35.3	28.6	31.6
Hispanic Motorists			
Search Rate	1.41	0.85	0.96
Hit Rate	29.3	24.2	20.6
Number of Troopers	1,465	216	744

This table presents search and hit rates by motorist and trooper race. We identify trooper race from 2015 personnel records.

Table B11
Racial Disparities in Search Rates by Trooper Race

Outcome:	Motorist/Vehicle Searched				
	(1)	(2)	(3)	(4)	(5)
Black	3.08	2.73	2.23	1.92	1.90
	(0.05)	(0.04)	(0.04)	(0.03)	(0.03)
Hispanic	1.70	1.68	1.47	1.47	1.46
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Black Trooper	0.83	0.80	0.82	0.81	0.80
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Hispanic Trooper	0.75	0.80	0.80	0.81	0.81
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
Black $\times$ Black Trooper	0.79	0.81	0.80	0.79	0.80
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Black $\times$ Hispanic Trooper	1.21	1.17	1.16	1.14	1.15
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Hispanic $\times$ Black Trooper	0.75	0.71	0.71	0.72	0.72
	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)
Hispanic $\times$ Hispanic Trooper	0.92	1.02	0.99	0.96	0.96
	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Female	0.39	0.40	0.42	0.51	0.51
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Time FEs		✓	✓	✓	✓
Location FEs		✓	✓	✓	✓
Month FEs		✓	✓	✓	✓
Income			✓	✓	✓
Arrest History				✓	✓
Stop History					✓
Observations	3,790,428	3,790,428	3,790,428	3,790,428	3,790,428

This table presents odds ratio estimates for the logistic regression model (B.1) augmented with fixed effects for trooper race and interactions between motorist and trooper race. We identify trooper race from 2015 personnel records. We limit to stops conducted by black, Hispanic, and white troopers. Standard errors are clustered at the motorist level.

Table B12
Search Disparities, Citation Disparities, and Local Political Preferences

Black-White Gap

	Black-White Search Odds Ratio		
	(1)	(2)	(3)
Citation Odds Ratio	0.531		-0.152
	(0.487)		(0.544)
Republican Vote Share		3.245	3.367
		(0.678)	(0.978)
Observations	79	79	79
Adjusted R^2	0.023	0.222	0.223

Hispanic-White Gap

	Hispanic-White Search Odds Ratio		
	(4)	(5)	(6)
Citation Odds Ratio	0.573		0.468
	(0.220)		(0.245)
Republican Vote Share		0.674	0.416
		(0.381)	(0.400)
Observations	79	79	79
Adjusted R^2	0.051	0.032	0.061

This table presents estimates of linear regression models where the outcome is the sergeant area-specific black-white (Panel A) or Hispanic-white search odds ratio (Panel B) derived from equation (12). 'Citation Odds Ratio' refers to sergeant-area specific black-white (Panel A) or Hispanic-white citation odds ratio (Panel B) derived from equation (12) where the outcome is replaced with an indicator for whether the stop results in a citation rather than a warning. 'Republican Vote Share' refers to the Republican vote share in the 2016 presidential election. For sergeant areas that cover multiple counties, we take a weighted average of the county-level Republican vote shares where weights reflect the share of sergeant area stops conducted in each county. Robust standard errors are reported in parentheses.