Taylor & Francis
Taylor & Francis Group

Does Predictive Policing Lead to Biased Arrests? Results From a Randomized Controlled Trial

P. Jeffrey Brantingham[a,*], Matthew Valasik[b], and George O. Mohler[c,*]

[a]UCLA Department of Anthropology, Los Angeles, CA; [b]Department of Sociology, Louisiana State University, Baton Rouge, LA; [c]Department of Computer and Information Science, Indiana University-Purdue University, Indianapolis, IN

ABSTRACT

Racial bias in predictive policing algorithms has been the focus of a number of recent news articles, statements of concern by several national organizations (e.g., the ACLU and NAACP), and simulation-based research. There is reasonable concern that predictive algorithms encourage directed police patrols to target minority communities with discriminatory consequences for minority individuals. However, to date there have been no empirical studies on the bias of predictive algorithms used for police patrol. Here, we test for such biases using arrest data from the Los Angeles predictive policing experiments. We find that there were no significant differences in the proportion of arrests by racial-ethnic group between control and treatment conditions. We find that the total numbers of arrests at the division level declined or remained unchanged during predictive policing deployments. Arrests were numerically higher at the algorithmically predicted locations. When adjusted for the higher overall crime rate at algorithmically predicted locations, however, arrests were lower or unchanged.

1. Introduction

Place-based predictive policing is based on two core ideas: (1) mathematical forecasting methods can be used to anticipate future crime risk in narrowly prescribed geographic areas; and (2) the delivery of police resources to those prediction locations disrupts the opportunity for crime (Bowers, Johnson, and Pease 2004; Mohler et al. 2011). Randomized controlled experiments of predictive policing conducted in Los Angeles provided evidence that algorithmic methods not only predict two-times as much crime as existing best practice, but also double the amount of crime prevented (Mohler et al. 2015). While this treatment effect can be measured in the field, the specific mechanisms by which predictive policing delivers greater crime reduction are not immediately obvious.

The prevailing view, derived from experiments in hot spot policing (Sherman and Eck 2002; Braga and Bond 2008), is that the presence of police in a given place removes opportunities for crime even without any direct contact with potential offenders (Sherman and Weisburd 1995; Weisburd 2008; Loughran et al. 2011). This general deterrent effect persists for some time after police have departed (Koper 1995; Sherman and Weisburd 1995) and appears to diffuse into nearby areas, where the police were not concentrating their efforts (Clarke and Weisburd 1994; Weisburd et al. 2006; Telep et al. 2014). General deterrence is not the only mechanism by which crime might be prevented by police patrol, however. Direct interference via stops, searches, detentions short of arrest, and arrest, may prevent crime by physically incapacitating potential offenders (Sherman and

Eck 2002; Weisburd and Eck 2004). This use of selective incapacitation may have immediate effects on crime (Wyant et al. 2012), especially if prolific offenders are the ones being arrested. Incapacitation may have longer term effects if those prolific offenders are subsequently removed from the community.

Considerable evidence suggests that explicit and implicit bias can have a major impact on *who* gets stopped, searched, and detained. Reasonable concern therefore exists that predictive policing can exacerbate such biases and reinforce any tendency for police to target minority individuals and communities (Ferguson in press). Such concern exists even if the forecasting methods used to drive predictive policing refrain from incorporating data that would be an explicit source of bias. If predictive policing indirectly exacerbates bias, any crime control benefits would need to be weighed in terms of their discriminatory costs. In the worst case, documented benefits might be derived solely from bias induced by predictions. In other words, predictions absent such bias would yield no crime control benefits at all. Here, we seek to evaluate whether predictive policing leads to patterns of arrest biased against minority individuals.

2. Predictive Policing and Racial Bias

Research has demonstrated that a racial bias exists in the business of policing including the racial profiling of vehicles (Harris 1999; Smith and Petrocelli 2001; Meehan and Ponder 2002; Novak 2004; Farrell and McDevitt 2006; Ridgeway 2006; Warren et al. 2006; Epp, Maynard-Moody, and Haider-Markel 2014;

CONTACT P. Jeffrey Brantingham ✉ branting@ucla.edu ⬤ UCLA Department of Anthropology, 341 Haines Hall, Los Angeles, CA 90095.
*Geroge O. Mohler and P. Jeffrey Brantingham are co-founders and directors of PredPol, a company offering predictive policing services to law enforcement agencies.

Baumgartner et al. 2016; Horrace and Rohlin 2016), pedestrian stops (Harris 1994; Gelman, Fagan, and Kiss 2007; Fagan et al. 2010; Rios 2011; Fagan et al. 2015; Goel, Rao, and Shroff 2016; Stuart 2016), traffic tickets (Dunn 2009), drug enforcement and arrests (Black and Reiss Jr. 1970; Smith, Visher, and Davidson 1984; Beckett, Nyrop, and Pfingst 2006; Kochel, Wilson, and Mastrofski 2011; Lynch et al. 2013), use of force (Schuck 2004; Legewie 2016; Buehler 2017; Nix et al. 2017), and even in the decision to shoot white or black criminal suspects while in a training simulator (Geller and Toch 1995; Plant and Peruche 2005). While the mechanisms driving these observed patterns of racial disparity (i.e., racial profiling, stereotyping/cognitive bias, deployment, racial animus/prejudice) remain difficult to disentangle, as (Warren et al. 2006) attests, there is little doubt that racial disparities in policing outcomes do exist.

Racial bias of predictive policing algorithms has been the focus of a number of recent news articles (*https://www.sciencenews.org/blog/science-public/data-driven-crimeprediction-fails-erase-human-bias*; *https://www.pbs.org/newshour/nation/column-big-data-analysis-police-activity-inherently-biased*), and concerns have been raised by several national organizations (for example the ACLU and NAACP) (*https://www.aclu.org/other/statement-concern-about-predictive-policing-aclu-and-16-civil-rights-privacy-racial-justice*) and recent research articles (Brayne 2017; Jefferson 2017; Ferguson, in press). In regards to place-based predictive policing methods that forecast a time and location where a crime may occur, the concern is that racially biased police practices may be directed toward some areas rather than others (Ferguson, in press). In addition, knowing that they are in a prediction area may heighten the awareness of police officers in ways that amplify biases (Ferguson 2012). That is, a minority individual observed in a prediction area may be more likely to be subject to biased police actions than the same individual observed outside of a prediction area.

Lum and Isaac (2016) conducted a simulation study of predictive policing focused on drug arrests in Oakland, CA. Their goal is to ascertain if racially biased outcomes are possible, or even amplified, with place-based predictive policing methods. The algorithm Lum and Isaac analyze is a space-time Hawkes process that was the method used in the Los Angeles Predictive Policing Experiment (Mohler et al. 2015). In particular, given a city grid indexed by n, the probabilistic rate $\lambda_n(t)$ of events in cell n at time t is determined by

$$\lambda_n(t) = \mu_n + \sum_{t_n^i < t} \theta \omega e^{-\omega(t - t_n^i)}, \tag{1}$$

where t_n^i are the times of events in cell n in the history of the process, μ_n is a baseline rate of events, and $\theta \omega e^{-\omega t}$ reflects the increase in risk following a recent crime. Lum and Isaac showed that if the events t_n^i correspond to racially biased drug arrests, then $\lambda_n(t)$ will increase after an arrest, leading to more police resources deployed to cell n in the future. Thus, a feed back loop may be possible where more arrests then occur in cell n, leading to a further increase in λ_n. A similar concern is raised by Ferguson (2017), who notes that arrests in a prediction area "memorializes" that location as "hot," which guarantees that that it will show up again as a prediction area producing further arrests. Selbst (2017a) also warns that racially biased outcomes

may become entrenched in place-based predictive policing given data collected via discriminatory policing practices. Constitutional policing tenets such as reasonable suspicion and probable cause may offer little protection against bias embedded in the data (Degeling and Berendt 2017; Selbst 2017b). Though all of these studies deal with hypothetical scenarios or thought experiments, they succeed in demonstrating that careful attention needs to be paid to whether predictive policing produces biased arrests (Moses and Chan 2016; Brayne 2017; Degeling and Berendt 2017; Jefferson 2017).

In practice, the majority of hotspot and place-based predictive policing algorithms focus not on arrests, but on crimes predominantly reported to the police by the public (e.g., robbery, burglary, assault) (Johnson n.d.; Black 1970; Mohler et al. 2015; Ferguson 2017). Thus, the goal is to send police resources to areas where crimes have been reported by victims, thus preventing future crimes in those areas. While a feedback loop for reported crime may be possible, in this case the self-reinforcement is toward places where citizens are placing calls for service. Therefore, we focus on the question of whether predictive policing produces arrests biased against minorities when the inputs to the system are reported crime incidents, rather than arrests. We run a set of hypotheses tests on empirical arrests recorded during the Los Angeles predictive policing experiment. We ask three related questions: (1) Did arrest of minority individuals differ between control and treatment conditions in test divisions? (2) Did arrest rates overall differ between control and treatment conditions in test divisions? and (3) Did the rate of arrests per crime differ across treatment and control conditions.

3. Predictive Policing Experiments in Los Angeles

A randomized controlled trial (of predictive policing was conducted in three divisions of the Los Angeles Police Department (LAPD) between November 2011 and January 2013. The three participating LAPD divisions were Foothill (FH), North Hollywood (NH), and Southwest (SW). Only a brief outline of the experiment is presented here. Additional details of the algorithmic procedures, experimental design, and main effects are presented in (Mohler et al. 2015).

Each day of the experiment police patrol officers were handed patrol maps with 20 target areas marked as 500×500 foot boxes. Officers were informed that the target areas were locations where the risk of crime was highest for their shift. They were encouraged to patrol target areas during any available discretionary time. What officers did not know was that the mission maps distributed to them each day were designed either by an algorithmic forecasting method (see Mohler et al. 2011, 2015), or by an analyst from within the division using all of the technological and intelligence assets at their disposal. Which mission map officers received on any given day was randomized creating a treatment condition (algorithmic forecast) and control condition (analyst forecast). In this repeated-measures experimental design, treatment days were considered exchangeable with control days (Mohler et al. 2015).

The outcome of interest was the difference in reported crime between control and treatment days. The crime types targeted were burglary, car theft, and burglary theft from vehicle (BTFV). Historically, these crime types account for as much as 60% of the serious crime in the City of Los Angeles. In addition to this

outcome measure, we collected information on the amount of time police officers spent in prediction areas under each of the experimental conditions (Mohler et al. 2015). Officers used their in-car computer terminals to register when they were entering and exiting prediction locations. This "dosage" was aggregated by day for a total amount of time (in minutes) spent in prediction locations.

Across the three test divisions, patrol officers using the algorithmic predictions produced an average 7.4% drop in crime at mean patrol dosage. By contrast, use of the best-practice predictions produced an average 3.5% drop in crime at mean patrol dosage. Individually, the decrease in crime associated with algorithmic predictions was statistically significant, while that with best-practice predictions was not. The evidence presented in Mohler et al. (2015) suggests that police patrol, when influenced by accurate predictions about the timing and location of crime, may provide some additional crime deterrence value. However, the standard errors of the estimates are relatively large making the slopes difficult to distinguish statistically and the precision on these rates is insufficient to conclude that the algorithmic forecasting crime rate drop is greater than the analyst crime rate drop.

3.1. Data

3.1.1. Defining Experimental Periods

Predictive policing experiments in each of the three LAPD divisions started at different times and ran for different total durations (Figure 1). The experiment in Foothill Division ran for 172 days from November 7, 2011, to April 27, 2012. In North Hollywood, the experiment ran for 167 days from March 31, 2012, to September 14, 2012 and in Southwest Division, the experiment ran for 239 days between May 16, 2012, to January 10, 2013. Total crime and arrests counts differ as a function of experimental duration in each division.

3.1.2. Defining Control and Treatment Days

Control and treatment missions were designed independently, but in parallel each day of the experiment. Recall that treatment missions were based on algorithmic forecasting, while the control missions were based on existing best practice of analysts. Once mission designs were finalized, a control or treatment mission was chosen randomly for deployment. This randomization was done independently each day for each division taking part in the experiment. On occasion, the analyst was not present on a randomly designated control day and therefore control missions were not available for those days. We exclude treatment days from these days to ensure fair comparison. In Foothill Division, there were a total of 124 test days with successful random assignment, after discarding days on which the analyst was not present to design control missions. The 124 test days were evenly divided with 62 control and 62 treatment days. There were 152

total test days in North Hollywood Division. These included 82 control and 70 treatment days. In Southwest Division, there were 234 total days, including 117 control and 117 treatment days.

3.1.3. Defining Arrests

An arrest is generally understood to mean the taking into custody of an individual by the police given probable cause that a violation of the law has occurred. An arrest, as recorded by the LAPD, should not be conflated with other down-stream processes of the criminal justice system. An arrest does not imply booking, continued detention, nor whether those individuals are ultimately prosecuted for a crime. Arrests also should not be conflated with contacts between the public and police that did not result in arrest, even if such contacts were contentious. In general, police can exercise many alternatives to arrest in seeking to enforce laws and ensure order including behavioral directives, warnings and brief detention without arrest. On average, the LAPD makes about 1.5 million public contacts per year, about 24,000 of these contacts (1.6%) are arrests (Beck 2016). Here, arrests are taken at face value, without considering anything beyond the official record that an individual was taken into custody.

We do not distinguish between arrests for different types of crimes. In 2012, the LAPD made arrests under 520 different criminal codes representing 25 broad classes of crimes such as aggravated assault, robbery, burglary and larceny. Our primary focus is on whether the practice of policing introduces new biases into arrest patterns, not whether bias might be differentially present in arrests for different types of crimes.

3.1.4. Defining Racial-Ethnic Groups

The LAPD collects demographic information as part of the arrest process including age, sex, and race-ethnicity of the individuals arrested. This information may be elicited from the individual or inferred by the arresting officer. The LAPD recognizes the categories Asian, black, Latino, white and other, which combined constitute 97.7% of all arrests on average. Occasionally, other categories such as Filipino, Korean, and Pacific Islander appear within the data. Given the sometimes fraught history between the LAPD and communities of color (see Herbert 1997; Muniz 2015; Martinez 2016) we focus on patterns in the arrest of black and Latino individuals and therefore report results for these two groups and for white individuals.

3.1.5. Defining Crimes

We define crimes as those incidents reported to the LAPD that are classified by the LAPD crime coding system into 226 recognized crime types. In a typical year, the LAPD collects reports on approximately 180,000 crimes. Again, we do not distinguish between different types of crimes.
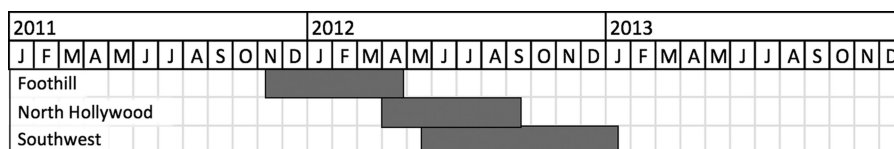

Figure 1. The exposure periods in Foothill, North Hollywood, and Southwest Divisions of the LAPD.

Table 1. Arrests vs. ethnicity on control and treatment days in three LAPD Divisions. † indicates counts adjusted for treatment days to a rate (per 82 days) to match control. Cochran–Mantel–Haenszel test $M^2 = 0.72581$, $df = 2$, p-value $= 0.6957$. Woolf test p-value $= 0.94337$.

Division	Race/Ethnicity	Control Arrests	Treatment Arrests
FH	Black	117	118
FH	Latino	700	756
FH	White	199	218
NH	Black	241	228†
NH	Latino	927	880†
NH	White	446	423†
SW	Black	1887	1818
SW	Latino	843	811
SW	White	110	73

4. Results

We now turn to a consideration of potential biases induced by predictive policing. We test three null hypotheses: (1) arrest of minority individuals did not differ between control and treatment conditions in test divisions; (2) arrest rates overall did not differ between control and treatment conditions in test divisions; (3) the rate of arrests per crime was unchanged across treatment and control conditions.

The LAPD experiment was designed to test for differences in predictive accuracy and impact on crime between control and treatment (predictive policing) conditions. Here we examine arrest patterns on control and treatment days (Table 1). Because in North Hollywood there was a lower number of treatment days ($n = 70$) compared to control days ($n = 82$) in the experiment, we adjust North Hollywood treatment counts to a rate per 82 days to be comparable to control. We conduct a Cochran–Mantel–Haenszel test (Agresti and Kateri 2011) to examine whether ethnicity is independent of the treatment condition. The CMH test is a generalization of a chi-square test, where the test is repeated across a strata, in this case several divisions. Here, we do not find evidence to reject the null hypothesis that treatment arrests are independent of ethnicity, the CMH test p-value is 0.6957 (Table 1). The CMH test assumes that the treatment effect is homogeneous across the three divisions, we test this assumption using a Woolf test and do not find evidence to reject the null hypothesis of homogeneity (p-value 0.94337).

In Table 2, we test whether the total number of arrests were higher or lower on treatment days at the division level. Here, we find that arrests were unchanged in Foothill and North Hollywood (using a chi square goodness-of-fit test for equal rates) and arrests were slightly lower in Southwest on treatment days (marginally significant at the 0.06 level).

We also examine arrest patterns in control and treatment boxes (Table 3). Recall that each participating division of the LAPD was allocated 20 prediction boxes, each 500×500 feet in

Table 2. Total arrests on control vs treatment days aggregated by division. p-value is given for a chi square goodness-of-fit test for equal rates. † indicates counts adjusted for treatment days to a rate (per 82 days) to match control.

Division	Control Arrests	Treatment Arrests	p-value
FH	1016	1092	0.1
NH	1614	1531†	0.14
SW	2840	2702	0.06

Table 3. Arrests in control and treatment boxes for all three LAPD Divisions. † indicates counts adjusted for treatment days to a rate (per 82 days) to match control. Cochran–Mantel–Haenszel test $M^2 = 1.8567$, $df = 2$, p-value $= 0.3952$. Woolf test p-value $= 0.720528$.

Division	Condition	Black	Latino	White
FH	Control	2	23	8
FH	Treatment	9	37	15
NH	Control	9	19	6
NH	Treatment	15†	31†	11†
SW	Control	71	25	6
SW	Treatment	158	44	6

Table 4. Arrests on control vs. treatment days in boxes aggregated by division. p-value is given for a chi-square goodness-of-fit test for equal rates. † indicates counts adjusted for treatment days to a rate (per 82 days) to match control.

Division	Control Arrests	Treatment Arrests	p-value
FH	33	61	0.004
NH	34	57†	0.015
SW	102	208	<0.001

size, each day. Control and treatment days were randomly determined. Again, we test the association between treatment and arrest ethnicity using a CMH test. Here, we do not find evidence to reject the null hypothesis that treatment arrests in boxes are independent of ethnicity, the CMH test p-value is 0.3952 (Woolf test p-value $= 0.720528$).

In Table 4, we test whether arrests were higher or lower on treatment days within predicted boxes. Here, we find a statistically significant increase in arrests (approximately double) in treatment boxes in all three divisions. To understand why arrests are higher in treatment boxes, in Table 5 we adjust the arrest rate to control for the higher overall rate of crime in treatment boxes. Our null hypothesis is that if predictive policing is "fair," then the percentage of crimes in an area leading to an arrest should be equal across treatment and control conditions. Here we find this to be the case in North Hollywood (arrests per crime 7.0% control, 5.6% treatment, p-value 0.34) and Southwest (arrests per crime 10.3% control, 10.7% treatment, p-value 0.75). The arrest rate per crime is slightly lower in treatment boxes in Foothill (arrests per crime 14.9% control, 8.5% treatment, p-value 0.009).

5. Discussion and Conclusions

The stated goal of the analyses presented above was to assess the degree to which arrest rates were impacted by the introduction of predictive policing in three divisions patrolled by the LAPD. Special attention was paid to arrest rates by the race-ethnicity of

Table 5. Predictive policing arrests in boxes as a proportion of reported crimes in treatment vs. control boxes. p-value is given for a 2 sample proportions test. † indicates counts adjusted for treatment days to a rate (per 82 days) to match control.

Division	Control	Treatment	p-value
FH arrests/crimes	33/221	61/714	0.0085
	14.9%	8.5%	
NH arrests/crimes	34/485	57/1021†	0.3421
	7.0%	5.6%	
SW arrests/crimes	102/993	208/1937	0.7452
	10.3%	10.7%	

the individuals detained. Our null hypotheses were: (1) arrest of minority individuals did not differ between control and treatment conditions in test divisions; (2) arrest rates overall did not differ between control and treatment conditions in test divisions; (3) the rate of arrests per crime was unchanged across treatment and control conditions.

The evidence presented does not allow us to reject null hypothesis (1). There is no significant difference in the arrest proportions of minority individuals between treatment and control conditions. We also cannot reject hypothesis (2) at the division level. Arrest rates overall are the same on control and treatment days within the test divisions as a whole. However, we do reject null hypothesis (2) at the box level. Arrests were higher overall in treatment prediction boxes. We therefore tested hypothesis (3) to see if the higher arrest rate in treatment boxes is explained by an overall higher crime rate in treatment boxes. We fail to reject the null hypothesis (3). Arrest rates per crime do not differ across treatment and control conditions.

Clearly, arrests are a common part of day-to-day police operations. The introduction of predictive policing did not increase arrests overall, though treatment prediction boxes did see significantly more arrests than control prediction boxes. The increase arrests in treatment prediction boxes are perhaps understandable given that algorithmic crime predictions are more accurate than those produced by existing best practice (Mohler et al. 2015).

The present study has several important limitations. Arrests are an imperfect proxy for other types of police contacts including stops, searches and detentions short of arrest. It is possible that predictive policing induced increases in these other categories of police contacts, without a concomitant impact on arrests. For this to hold true, it would have to be the case that the rate of arrest actually declined as these other precursor contacts increased, leaving overall arrest numbers unchanged. This hypothetical downward adjustment in arrests would have to hold not only for the experimental deployment period overall, but also for randomly assigned treatment days. We do not have sufficient data to exclude such dynamics, but they seem improbable on the face of it.

Second, the analyses do not provide any guidance on whether arrests are themselves systemically biased. Such could be the case, for example, if black and Latino individuals experienced arrest at a rate disproportionate to their share of offending (Rosenfeld and Fornango 2014). The current study is only able to ascertain that arrest rates for black and Latino individuals were not impacted, positively or negatively, by using predictive policing. Future research could seek to test whether the situational conditions surrounding arrests and final dispositions differ in the presence of predictive policing.

Finally, the results reported herein pertain for the narrowly-defined place-based predictive policing model used in the Los Angeles predictive policing experiment (Mohler et al. 2015). This model focused on reported crime data for a limited set of crimes including burglary, car theft and burglary from vehicle and used only information on crime location and time. Predictions were made for small 500 × 500 foot boxes and changed every day. Under those conditions, we can conclude that predictive policing did not result in biased arrests. Whether the same outcomes would hold given changes in implementation is

uncertain. If the exact same data types and methods are applied in a different location there may be reason to be optimistic. However, if the data types change, for example to focus on discretionary crimes or arrests, or if personal information is incorporated into predictions, then pessimism may be warranted. At the same time, we should ask whether police would be negligent if they had data or information that led to accurate forecasts of crime risk but failed to act on it for fear of potential bias (Ferguson 2017). Continued empirical scrutiny along with careful policy development will be needed to guard against bias in predictive policing and ensure fairness in outcomes.

Acknowledgments

Funding

References

Agresti, A., and Kateri, M. (2011), "Categorical Data Analysis," in *International Encyclopedia of Statistical Science*, Berlin: Springer, pp. 206–208. [4]

Baumgartner, F. R., Epp, D. A., Shoub, K., and Love, B. (2016), "Targeting Young Men of Color for Search and Arrest During Traffic Stops: Evidence From North Carolina, 2002–2013," *Politics, Groups, and Identities*, 1–25. [2]

Beck, C. (2016), *Use of Force Year-End Review Executive Summary 2015*, Report, Los Angeles Police Department. [3]

Beckett, K., Nyrop, K., and Pfingst, L. (2006), "Race, Drugs, and Policing: Understanding Disparities in Drug Delivery Arrests," *Criminology*, 44, 105–137. [2]

Black, D. J. (1970), "Production of Crime Rates," *American Sociological Review*, 35, 733–748. [2]

Black, D. J., and Reiss Jr, A. J. (1970), "Police Control of Juveniles," *American Sociological Review*, 35, 63–77. [2]

Bowers, K. J., Johnson, S. D., and Pease, K. (2004), "Prospective Hot-Spotting The Future of Crime Mapping?" *British Journal of Criminology*, 44, 641–658. [1]

Braga, A. A., and Bond, B. J. (2008), "Policing Crime and Disorder Hot Spots: A Randomized Controlled Trial" *Criminology*, 46, 577–607. [1]

Brayne, S. (2017), "Big Data Surveillance: The Case of Policing," *American Sociological Review*, 82, 977–1008. [2]

Buehler, J. W. (2017), "Racial/Ethnic Disparities in the Use of Lethal Force by US Police, 2010–2014," *American Journal of Public Health*, 107, 297–297. [2]

Clarke, R., and Weisburd, D. (1994), "Diffusion of Crime Control Benefits: Observations on the Reverse of Displacement," *Crime Prevention Studies*, 2, 165–184. [1]

Degeling, M., and Berendt, B. (2017), "What is Wrong About Robocops as Consultants? A Technology-Centric Critique of Predictive Policing," *AI & Society*. [2]

Dunn, R. A. (2009), "Measuring Racial Disparities in Traffic Ticketing Within Large Urban Jurisdictions," *Public Performance & Management Review*, 32, 537–561. [2]

Epp, C. R., Maynard-Moody, S., and Haider-Markel, D. P. (2014), *Pulled Over: How Police Stops Define Race and Citizenship*, Chicago, IL: University of Chicago Press. [1]

Fagan, J., Braga, A. A., Brunson, R. K., and Pattavina, A. (2015), "Stops and Stares: Street Stops, Surveillance and Race in the New Policing." *Fordham Urban Law Journal*, 43, 539–614. [2]

Fagan, J., Geller, A., Davies, G., and West, V. (2010), *Street Stops and Broken Windows Revisited: The Demography and Logic of Proactive Policing in a Safe and Changing City*, Race, Ethnicity and Policing, New York: New York University. [2]

Farrell, A., and McDevitt, D. J. (2006), "Rhode Island Traffic Stop Statistics Data Collection Study 2004–2005," Report, Institute on Race and Justice. [1]

Ferguson, A. G. (2012), "Predictive Policing and Reasonable Suspicion," *Emory LJ*, 62, 259–324. [2]

—— (2017). *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*, New York: New York University Press. [2,5]

—— (in press). *Policing Predictive Policing*, Washington University Law Review, 94. [1,2]

Geller, W. A., and Toch, H. (1995), *And Justice for All: Understanding and Controlling Police Abuse of Force*, Washington, DC: Police Executive Research Forum Washington. [2]

Gelman, A., Fagan, J., and Kiss, A. (2007), "An Analysis of the New York City Police Department's Stop-and-Frisk Policy in the Context of Claims of Racial Bias," *Journal of the American Statistical Association*, 102, 813–823. [2]

Goel, S., Rao, J. M., and Shroff, R. (2016), "Precinct or Prejudice? Understanding Racial Disparities in New York City's Stop-and-Frisk Policy," *The Annals of Applied Statistics*, 10, 365–394. [2]

Harris, D. A. (1994), "Frisking Every Suspect: The Withering of Terry," *UC Davis L. Rev.*, 28, 1. [2]

—— (1999), "The Stories, the Statistics, and the Law: Why Driving While Black Matters," *Minn. L. Rev.*, 84, 265. [1]

Herbert, S. K. (1997), *Policing Space: Territoriality and the Los Angeles Police Department*, University of Minnesota Press. [3]

Horrace, W. C., and Rohlin, S. M. (2016), "How Dark is Dark? Bright Lights, Big City, Racial Profiling," *Review of Economics and Statistics*, 98, 226–232. [2]

Jefferson, B. J. (2017), "Predictable Policing: Predictive Crime Mapping and Geographies of Policing and Race," *Annals of the American Association of Geographers*. [2]

Johnson, S. D. (n.d.) *Prospective Crime Mapping in Operational Context: Final Report*, London, UK: Home Office. [2]

Kochel, T. R., Wilson, D. B., and Mastrofski, S. D. (2011), "Effect of Suspect Race on Officers Arrest Decisions," *Criminology*, 49, 473–512. [2]

Koper, C. S. (1995), "Just Enough Police Presence: Reducing Crime and Disorderly Behavior by Optimizing Patrol Time in Crime Hot Spots," *Justice Quarterly*, 12, 649–672. [1]

Legewie, J. (2016), "Racial Profiling and Use of Force in Police Stops: How Local Events Trigger Periods of Increased Discrimination," *American Journal of Sociology*, 122, 379–424. [2]

Loughran, T. A., Paternoster, R., Piquero, A. R., and Pogarsky, G. (2011), "On Ambiguity in Perceptions of Risk: Implications for Criminal Decision Making and Deterrence," *Criminology*, 49, 1029–1061. [1]

Lum, K., and Isaac, W. (2016), "To Predict and Serve?" *Significance*, 13, 14–19. [2]

Lynch, M., Omori, M., Roussell, A., and Valasik, M. (2013), "Policing the Progressive City: The Racialized Geography of Drug Law Enforcement," *Theoretical Criminology*, 17, 335–357. [2]

Martinez, C. (2016), *The Neighborhood Has Its Own Rules: Latinos and African Americans in South Los Angeles*, New York: NYU Press. [3]

Meehan, A. J., and Ponder, M. C. (2002), "Race and Place: The Ecology of Racial Profiling African American Motorists," *Justice Quarterly*, 19, 399–430. [1]

Mohler, G., Short, M., Brantingham, P., Schoenberg, F., and Tita, G. (2011), "Self-Exciting Point Process Modeling of Crime," *Journal of the American Statistical Association*, 106, 100–108. [1,2]

Mohler, G., Short, M. B., Malinowski, S., Johnson, M., Tita, G. E., Bertozzi, A. L., and Brantingham, P. J. (2015), "Randomized Controlled Field Trials of Predictive Policing," *Journal of the American Statistical Association*, 110, 1399–1411. [1,2,3,5]

Moses, L. B., and Chan, J. (2016), "Algorithmic Prediction in Policing: Assumptions, Evaluation, and Accountability," *Policing and Society*, 1–17. Doi:10.1080/10439463.2016.1253695. [2]

Muniz, A. (2015), *Police, Power, and the Production of Racial Boundaries*, New Brunswick, NJ: Rutgers University Press. [3]

Nix, J., Campbell, B. A., Byers, E. H., and Alpert, G. P. (2017), "A Bird's Eye View of Civilians Killed by Police in 2015," *Criminology & Public Policy*, 16, 309–340. [2]

Novak, K. J. (2004), "Disparity and Racial Profiling in Traffic Enforcement," *Police Quarterly*, 7, 65–96. [1]

Plant, E. A., and Peruche, B. M. (2005), "The Consequences of Race for Police Officers' Responses to Criminal Suspects," *Psychological Science*, 16, 180–183. [2]

Ridgeway, G. (2006), "Assessing the Effect of Race Bias in Post-Traffic Stop Outcomes Using Propensity Scores," *Journal of Quantitative Criminology*, 22, 1–29. [1]

Rios, V. M. (2011), *Punished: Policing the Lives of Black and Latino Boys*, New York, NY: NYU Press. [2]

Rosenfeld, R., and Fornango, R. (2014), "The Impact of Police Stops on Precinct Robbery and Burglary Rates in New York City, 2003–2010," *Justice Quarterly*, 31, 96–122. [5]

Schuck, A. M. (2004), "The Masking of Racial and Ethnic Disparity in Police Use of Physical Force: The Effects of Gender and Custody Status," *Journal of Criminal Justice*, 32, 557–564. [2]

Selbst, A. D. (2017a), "Disparate Impact in Big Data Policing," *Georgia Law Review*, 49, 1–77. [2]

—— (2017b). "A Mild Defense of Our New Machine Overlords," *Vanderbilt Law Review En Banc*, 70, 87–104. [2]

Sherman, L. W., and Eck, J. E. (2002), "Policing for Crime Prevention," in *Evidence-Based Crime Prevention*, London: Routledge, pp. 295–397. [1]

Sherman, L. W., and Weisburd, D. (1995), "General Deterrent Effects of Police Patrol in Crime Hot Spots: A Randomized, Controlled Trial," *Justice Quarterly*, 12, 625–648. [1]

Smith, D. A., Visher, C. A., and Davidson, L. A. (1984), "Police Control of Juveniles," *Equity and Discretionary Justice: Influence of Race on Police Arrest Decisions*, 75, 234–249. [2]

Smith, M. R., and Petrocelli, M. (2001), "Racial Profiling? A Multivariate Analysis of Police Traffic Stop Data," *Police Quarterly*, 4, 4–27. [1]

Stuart, F. (2016), *Down, Out, and Under Arrest: Policing and Everyday Life in Skid Row*, Chicago: University of Chicago Press. [2]

Telep, C. W., Weisburd, D., Gill, C., Vitter, Z., and Teichman, D. (2014), "Displacement of Crime and Diffusion of Crime Control Benefits in Large-Scale Geographic Areas: A Systematic Review," *Journal of Experimental Criminology*, 10, 515–548. [1]

Warren, P., Tomaskovic-Devy, D., Smith, W., Zingraff, M., and Mason, M. (2006), "Driving While Black: Bias Processes and Racial Disparity in Police Stops," *Criminology*, 44, 709–738. [1]

Weisburd, D. (2008), "Place-Based Policing," *Ideas in American Policing*, 9, 1–15. [1]

Weisburd, D., and Eck, J. E. (2004), "What Can Police Do to Reduce Crime, Disorder, and Fear?" *The ANNALS of the American Academy of Political and Social Science*, 593, 42–65. [1]

Weisburd, D., Wyckoff, L. A., Ready, J., Eck, J. E., Hinkle, J. C., and Gajewski, F. (2006), "Does Crime Just Move Around the Corner? A Controlled Study of Spatial Displacement and Diffusion of Crime Control Benefits," *Criminology*, 44, 549–592. [1]

Wyant, B. R., Taylor, R. B., Ratcliffe, J. H., and Wood, J. (2012), "Deterrence, Firearm Arrests, and Subsequent Shootings: A Micro-Level Spatio-Temporal Analysis," *Justice Quarterly*, 29, 524–545. [1]