Algorithmic Policing*

Ranae Jabri†

November 2023

Abstract

Predictive policing is pervasive yet understudied. This paper examines the impacts of algorithmic policing and how it affects racial profiling. Using a novel dataset on predictive policing and natural experiment research design, I estimate that algorithmic policing decreases serious violent and property crimes, but exacerbates racial disparities in arrests in traffic incidents and serious violent crimes. The evidence suggests a threefold increase in arrests of Black motorists when the neighborhood is targeted in comparison to when it is not. Algorithmic policing can prevent crime at the cost of increasing racial disparities in arrests, underscoring racial equity implications of algorithmic targeting.

JEL codes: H76, J15, K42

Keywords: algorithm, law enforcement, artificial intelligence, police, inequality, race

*I am indebted to Pat Bayer, Rob Garlick, Seth Sanders, and Chris Timmins for their advice and support. I benefited from thoughtful comments and questions from Desmond Ang, Bocar Ba, Jason Baron, Phil Cook, Will Dobbie, Dhammika Dharmapala, Jennifer Doleac, Paul Dudenhefer, Erica Field, Amanda Geller, Felipe Goncalves, William Hubbard, Matt Masten, Giovanni Mastrobuoni, Arnaud Maurel, Emily Owens, John Rappaport, Cailin Slattery, Juan Carlos Suarez Serrato, Megan Stevenson, Bryan Stuart, Catherine Tucker, Pengpeng Xiao, Crystal Yang, and numerous seminar participants. I thank the jurisdiction for providing access to the data that made this paper possible and PredPol for providing access to their site. I gratefully acknowledge fellowship support from the Graduate School of Duke University and the Alfred P. Sloan Foundation (G-2021-16980). Any errors are my own.

†Mailing address: A02 Social Sciences Building Camperdown NSW 2050 Australia; Telephone: +14158878281; Email: ranae.jabri@sydney.edu.au

1 Introduction

Artificial intelligence (AI) technologies have permeated many aspects of our lives, from firms to hospitals to the criminal justice system. Predictive policing tools — algorithms that predict crime risk at hyperlocal levels and high frequency — are increasingly used in the hopes of improving efficiency.[1] Such tools function as follows: every shift, law enforcement receives electronic maps delineating geographic areas where they are instructed to patrol. The hope is that allocating patrols to these places with the highest crime risk will proactively deter crime (Becker, 1968). At the same time, law enforcement may use algorithmic policing to justify more discretion or profiling (Ferguson, 2012). Nonetheless, little is known about the effects of algorithmic policing or the extent to which it affects racial profiling. Understanding the racial equity implications of algorithmic policing is of particular concern given the evidence of racial disparities and bias in policing[2], and subsequent criminal justice processes.[3]

Even beyond algorithmic policing, the effects of local police presence remain open for investigation. Estimating causal effects of localized police presence is complicated by an endogeneity problem: the locations where law enforcement patrols within jurisdictions are likely related both to crime and arrest rates and to the racial composition of neighborhoods (Chen et al., 2021). An established literature studies the effects of large-scale, long-term police deployments.[4] However, we know little about the causal effects of local police presence (Blanes i Vidal and Mastrobuoni, 2018; Weisburd, 2021), and have even less evidence about whether local police presence has disproportionate racial impacts or any potential trade-offs.

In law enforcement agencies that have adopted predicted policing, patrols are allocated in a

[1] A 2012 survey of more than 500 U.S. police agencies conducted by the Police Executive Research Forum found that 38% of police agencies were already using predictive policing, and 70% were planning to use predictive policing by 2017 (Police Executive Research Forum, 2014).

[2] (Knowles et al., 2001; Grogger and Ridgeway, 2006; Antonovics and Knight, 2009; Anwar and Fang, 2006; Horrace and Rohlin, 2016; West, 2018; Fryer Jr, 2019; MacDonald and Fagan, 2019; Goncalves and Mello, 2021; Ba et al., 2021b; Feigenberg and Miller, 2022; Hoekstra and Sloan, 2022; Grosjean et al., 2022)

[3] See Agan (2022) and Doleac (2022) for a recent review.

[4] Papers have investigated the impacts of increases in police hiring (Levitt, 1997; McCrary, 2002; Levitt, 2002; Lin, 2009; Chalfin and McCrary, 2018; Weisburst, 2019; Chalfin et al., 2022), large-scale city-wide deployments (Di Tella and Schargrodsky, 2004; Klick and Tabarrok, 2005; Draca et al., 2011), and long-term police deployments such as to more traditional crime hotspots (Weisburd and Telep, 2014; MacDonald et al., 2016; Blattman et al., 2021).

systematic way based on a predicted crime risk score. For every shift in a jurisdiction, predictive policing tools generate high-frequency, hyperlocal data on where law enforcement should patrol, and I can isolate areas with similar crime risk where law enforcement is not induced to patrol. That is, I can compare areas with similar crime risk but different levels of police presence.

In this paper, I investigate the impacts of police presence induced by predictive policing algorithms on crime incidence and racial disparities in arrests. The answers to these open questions about algorthmic policing inform broader policy questions surrounding the trade-offs involved in neighborhood targeting of police presence: namely, what are costs and benefits of targeting police presence to an area? Are people of color disproportionately affected? To date, no studies have considered both the efficacy and equity implications of the causal impacts of local police presence.

To examine these questions, I collect a unique dataset describing predictive policing box locations and algorithm input data, crime incidents, and arrests from a major urban jurisdiction in the United States. The jurisdiction uses PredPol, a leading predictive policing technology (from a software company of the same name) that was one of the first to be deployed in the United States. PredPol, claims that its tool is "used to help protect one out of every 33 people in the United States."[5] The PredPol software predicts crime risk for 500 ft–by–500 ft area geographic units in a jurisdiction for every patrol shift. Of the thousands of boxes in the jurisdiction, for every shift, the top 24 boxes with the highest crime risk in each district are designated as predictive policing boxes or "PredPol boxes." Patrol officers receive the list of PredPol box locations for their district to patrol for their assigned shift. I assemble a dataset of 2.3 million box–shifts, where I can observe for each box for a given shift whether it has been designated as a PredPol box, the input data used by PredPol to predict the crime risk, the crime incidents and arrests that occur, and demographic information on any arrestees.

To estimate the causal effects of algorithmic policing, I use key features of the predictive policing institutional setting to isolate quasi-experimental variation in police presence induced by predictive policing algorithms. In 2019, there was an arbitrary change in the PredPol system that

[5] `https://blog.predpol.com/predpol-named-to-govtech100-list-for-5th-straight-year`

slightly altered the predictive policing boxes delivered to law enforcement.[6] Law enforcement did not know there was a change and only saw the new sets of predictive policing boxes every shift after it. For each shift, I observe the boxes that would have been PredPol boxes in the absence of the change. This set of predictive policing boxes that are not delivered after the change have similar underlying crime risk and make up a control group for predictive policing boxes that are delivered to police before the change. My research design compares the outcomes across PredPol boxes with the outcomes across boxes with similar crime risk that would have been PredPol boxes had the arbitrary change in the algorithm not occurred (conditional on box fixed effects). Additionally, using institutional knowledge about how PredPol predicts crime risk along data on input variables used to generate predictive policing boxes, I account for how PredPol boxes are designated by controlling for a proxy of PredPol's crime risk measures.

My first hypothesis is that algorithm-induced police presence decreases crime incidence. A neighborhood being designated a PredPol box increases the perceived police presence and probability of arrest, which should increase the expected cost of committing a crime and decrease the likelihood that a crime in committed in the box (Becker, 1968). At the same time, police have discretion to stop and arrest civilians.[7] PredPol boxes with higher levels of crime risk can motivate officers to exert higher effort as these areas may have higher marginal returns to their efforts. As a result, there may be high-discretion incidents that are more likely to be discovered that may have been previously undetected.

My second hypothesis is that algorithm-induced police presence has disproportionate racial impacts on arrests on racial minorities. Officers know that a neighborhood being designated as a PredPol box indicates the highest level of crime risk for a given shift, which may in turn warrant more officer discretion.[8] Discretion can enable greater police discrimination in targeted areas. This can affect both serious and high discretion incidents like traffic stops, which are the most common

[6]Before this change, PredPol generated its boxes using the violation codes for aggravated assault, auto burglary, motor vehicle theft, robbery, and shots fired. Under the change, two more violation codes (residential and commercial burglary) were added to the original set.

[7]*Whren v. United States*, 517 U. S. 806, 813 (1996). *Atwater v. Lago Vista*, 532 U. S. 318, 323–324 (2001).

[8]Wardlow

reasons for contact with the police.[9]

First, I find that algorithm-induced police presence decreases serious property and violent crime on the order of 3 crimes per 1000 PredPol boxes — or 30%— over a 12-hour shift, and do not find evidence of displacement to surrounding boxes. This result is robust to my controlling for district-level time trends or including district–time fixed effects. I also find suggestive evidence that reported traffic incident increase, which offers auxiliary evidence that law enforcement spends more time in PredPol boxes.

The result is quantitatively similar when I use an alternative research design. I use a regression discontinuity design (RDD) to estimate the effect of algorithm-induced police presence. PredPol predicts a continuous measure of underlying crime risk, and there are algorithmic cutoffs at which a box counts as a PredPol box. While I do not observe the continuous crime risk measure, I use institutional details from PredPol's marketing materials and a publication by authors affiliated with PredPol (Mohler et al., 2015) along with the algorithmic input data used in generating predictive policing boxes to predict an estimate of the continuous crime risk measure using a machine learning model. Using this estimate of the continuous crime risk measure as a running score in an RDD (Boehnke and Bonaldi, 2019), I compare the outcomes of predictive policing boxes that marginally pass the PredPol box threshold with boxes that marginally fall short of this threshold. This alternative framework yields the qualitatively similar result that algorithm-induced police presence decreases serious property and violent crime.

Next, I examine whether algorithm-induced police presence has disproportionate impacts by race on arrests. I use a nested model building on the main empirical strategy to estimate the effects of algorithm-induced police presence for Black, Hispanic and white individuals. Using the nested model, I calculate the counterfactual mean arrests for predictive policing boxes had the same box not been targeted. Then, to test whether algorithm-induced police presence has racially disparate impacts on arrests, I compare the causal effects *by race and ethnicity* weighted by the counterfactual of how many arrests would have occured for each group had the same boxes not

[9]Bureau of Justice Statistics. https://www.bjs.gov/index.cfm?tid=702&ty=tp

been predictive policing boxes.

My findings reveal that algorithm-induced police presence has uneven impacts on arrests of Black motorists for traffic incidents relative to the same types of arrests of white motorists. The evidence suggests a threefold increase in arrests of Black motorists when the neighborhood is targeted in comparison to when it is not. For arrests for serious violent crimes, I find disproportionately larger effects for Black than from white individuals. There is a decrease in arrests for white individuals when the neighborhood is targeted of the magnitude of 0.68 times the number of counterfactual arrests under under no targeting, with no proportional decrease in arrests for Black individuals. Overall, these findings are consistent with algorithmic policing exacerbating racially disproportionate law enforcement behavior.

I document striking heterogeneity across PredPol boxes in predominantly Black vs non–predominantly Black communities. I find that algorithmic policing increases racial disparities in arrests in traffic incidents in predominantly Black communities designated as PredPol boxes but do not find evidence of an effect on racial disparities in non–predominantly Black communities. These results reveal a disproportionate burden of algorithmic surveillance falls on Black individuals in predominantly Black communities. These results are particularly concerning given the evidence that PredPol is more likely to class minority areas as having higher crime risk (Lum and Isaac, 2016).

While there is evidence that algorithm-induced police presence decreases crime serious violent and property crime in the predictive boxes, there is also evidence that racial disparities in arrests increase, raising serious equity concerns about programs targeting police presence to specific neighborhoods by designating them predictive policing boxes. Moreover, the Black individuals who are more likely to be arrested when a box is designated as a predictive policing box are not those who are less likely to commit crimes.

This article makes three main contributions. First, to my knowledge, the paper provides the first evidence on the effects of algorithmic policing — police presence induced by predictive policing. Mastrobuoni (2020) finds that using predictive policing algorithms to investigate robberies can increase case clearance rates, which can decrease crime through incapacitation. Two experiments (Mohler et al., 2015; Ratcliffe et al., 2020) study the efficacy of predictive policing to allocate patrol in comparison to status quo approaches at the district–level.[10] Compared to these papers, my paper is a much more granular study of what actually happens at predictive policing boxes, including who is affected. While these experiments study the comparative advantage of predictive policing in prediction and allocation, I examine effects of actual police presence induced by predictive policing algorithms.

My paper contributes the first evidence on the disproportionate impacts of algorithmic policing on people and communities of color. It complements a closely related paper by Brantingham et al. (2018) on racial bias in algorithmic vs human prediction. Brantingham et al. (2018) examine whether algorithmic prediction flags areas with more racially disproportionate arrests as high risk than does human prediction. To compare the predictive capability of algorithms with that of humans, patrol officers in the experiment are blind as to whether the hotspots they are instructed to patrol are human- or algorithm-predicted high risk areas. The authors fail to find statistically significant evidence of more racial bias in algorithmic prediction than in human prediction. In my study, the officers know that PredPol boxes are algorithmically targeted which allows me to examine the effect of algorithmic policing on enforcement behavior and arrests. I find that the causal impact of a neighborhood's being targeted with algorithmic policing is more racially disproportionate arrests in traffic incidents and serious violent crimes.

Second, this paper contributes to the literature on the effects of policing on crime as predictive policing targets patrol to specific areas.[11] By estimating the effects of local police presence at a granular level with respect to time and space, I shed light on *how* police presence decreases crime and the trade-offs involved. My findings suggest law enforcement deters crime through

[10]Mohler et al. (2015) finds that using algorithms to predict and target patrol decreases district-level crime as compared to human prediction. Ratcliffe et al. (2020) finds that using algorithms to predict and target marked patrol cars decreases district-level property crime compared to the status quo where staff have no access to predictive policing software.

[11]See Durlauf and Nagin (2011) and Chalfin and McCrary (2017) for a review as well as papers referenced earlier; there is evidence that increasing the probability of apprehension through increasing police hiring and police deployments deters crime. Moreover, the probability of apprehension may affect the extent to which punishment severity deters crime (Gonzalez et al., 2022).

their presence rather than arrests. My paper is closely related to papers on the effects of traditional hotspot policing (Weisburd and Telep, 2014; Braga et al., 2019; Blattman et al., 2021).[12] Compared to predictive policing boxes, which update every shift and have uncertain locations, traditional hotspots are fixed over time, with certain patrol locations, and may simply displace rather than reduce crime.[13] To my knowledge, this literature does not study the effects of hotspot policing on racial disparities in arrests.

Moreover, I contribute to the open question of whether local police presence has disproportionate racial impacts. A closely related paper by Chalfin et al. (2022) also provides evidence on the efficacy and equity implications of increasing policing. While Chalfin et al. (2022) leverage quasi-experimental increases in police hiring to investigate the race-specific effect of a larger police force in a city over time, I focus on the effects of hyperlocal targeting of police presence induced by predictive policing, speaking to a different, important policy question: what happens when police enter a block and whether people of color are disproportionately affected.

Finally, this paper contributes to the literature on feedback in algorithms. Predictive models such as PredPol are considered particularly contentious because of concerns that they may amplify pre-existing racial inequities in policing patterns or data (O'Neil, 2017). PredPol claims to not explicitly use race in its models; however, past data may reflect historical patterns and biases in policing. Models may replicate and further amplify these disparities if police discover more crime in predictive policing boxes, which can create a negative feedback loop (Lum and Isaac, 2016). My findings show that algorithmic feedback can happen. The sign of the estimated causal effects of algorithm-induced police presence differs by type of crime, revealing that the direction of the feedback depends on the crime types used to predict the predictive policing boxes. These results reveal a racial equity–conscious policy must carefully carefully consider the types of crimes included in prediction.

[12] See Braga et al. (2019) for a review.

[13] When hotspots are fixed over time, potential offenders may become aware of the location of police presence and may leave areas with high levels of crime risk control. If it is costly for potential offenders to divert criminal activity when they observe police presence, the estimated effects of predictive policing boxes may be less susceptible to displacement effects than those of traditional hotspots.

The paper proceeds as follows: Section 2 provides an overview of the context that I study. Section 3 describes the novel data that I collect for my analysis. Section 4 outlines the quasi-experimental research design that I use to identify the effects of predictive policing on crime incidents. Section 4.2 finds quantitatively similar results using an alternative research design. Section 5 tests for disproportionate racial impacts of algorithm-induced police presence on arrests. Finally, Section 6 discusses the findings and policy implications and concludes.

2 Institutional Setting

2.1 Predictive Policing

Predictive policing algorithms uses past crime data to predict — at high frequency — high-crime-risk areas where law enforcement is instructed to patrol. The hope is that proactively patrolling the places with highest crime risk in a city will prevent crimes from occurring. Founded in 2012, PredPol offers one of of the first predictive policing algorithmic tools to go on the market in the United States.[14]

The PredPol software systematically produces a list of predictive policing boxes for every shift based on a crime risk score that it predicts for every geographic unit in a jurisdiction. According to PredPol marketing and a publication by coauthors affiliated with the company (Mohler et al., 2015), the PredPol software splits a jurisdiction into 500 ft–by–500 ft geographic units or boxes and then predicts a continuous measure of crime risk for a set of crime types for all boxes and for every shift. PredPol uses only crime times, crime types, and crime GPS coordinates to predict this continuous crime risk measure as an exponential decay function of the crime lags in each box and a crime time-invariant box effect. The functional form of the crime risk probabilistic rate (λ_{it}) of

[14] https://blog.predpol.com/predpol-named-to-govtech100-list-for-5th-straight-year

events in box i at time t (Mohler et al., 2015) is as follows:

$$\lambda_{it} = \mu_i + \sum_{t_i^n < t} \theta\omega e^{-\omega(t-t_i^n)} \tag{1}$$

where t_i^n are times of events in box i in the history of the process in the window being used for prediction T (T is suggested to be 365 days), μ_i is a baseline Poisson process rate (constant long-term background rate) or time-invariant box parameter, and $\theta\omega e^{\omega(t-t_i^n)}$ is an exponential decay "contagion" effect in crime data to capture short-term dynamics. The top 24 boxes per district per shift predicted by PredPol to have the highest crime risk are designated as predictive policing boxes or PredPol boxes. Therefore, a box i in district d at time t is a PredPol box if its crime risk for the box at time t (λ_{it}) is greater than or equal to the 24th-highest-risk box in the district ($max_{dt}^{24}\{\lambda_{1t}, \ldots, \lambda_{It}\}$).[15] PredPol boxes are delivered to law enforcement through an online interface and patrol reports. Figure C1 shows an example of a patrol report from a PredPol guide.

2.2 Jurisdiction

I study a major urban jurisdiction with a population of over 1 million people in a metropolitan statistical area among the fifty largest in the United States. I promised the jurisdiction that I would not reveal its name. The race/ethnicity breakdown of its population is approximately 15–20% non-Hispanic Black, 40–45% non-Hispanic white, and 25–30% Hispanic. The violent and property crime rates for the largest city in the jurisdiction are above the state and US national median rates. A large law enforcement agency using predictive policing, with over 2000 sworn officers and civilian employees, serves the jurisdiction. The jurisdiction has a uniform patrol division, assigned at the district and shift level, whose officers patrol in law enforcement uniforms. PredPol instructs officers to go to PredPol boxes for "about 6 minutes per hour."[16] In the jurisdiction that I study, patrols are instructed to go to PredPol boxes in their down time between calls for service. Patrols are also instructed to patrol in the PredPol boxes as they normally would.

[15] $PredPolBox_{idt} = 1(\lambda_{idt} \geq max_{dt}^{24}\{\lambda_{1t}, \ldots, \lambda_{It}\})$
[16] https://www.predpol.com/law-enforcement/#predPolicing

My research design exploits an exogenous change in the PredPol system in the set of PredPol boxes delivered to law enforcement in the jurisdiction. PredPol boxes for “All Crimes” are delivered to law enforcement patrol units during the day shift in all districts of the jurisdiction. Prior to 11/20/2019, auto burglary offenses, vehicle theft offenses, robbery offenses, assault offenses and shots fired calls for service were the crime types used to generate the “All Crimes” PredPol boxes. After 11/20/2019, two more crime types—residential burglary and commercial burglary—were added to the set of crime types used to generate the PredPol boxes for “All Crimes,” and the PredPol boxes for the new, expanded set of crime types began to be delivered to law enforcement. I refer to the PredPol boxes predicted using the original set of crimes types as All Crimes PredPol boxes and the PredPol boxes predicted using the new set with the two additional crime types as All Crimes-Plus PredPol boxes. Within the window around the exogenous change on 11/20/2019[17], the daily numbers of All Crimes and All Crimes-Plus crime types during the day shift are highly correlated, with a correlation of 0.92.

I interviewed the key law enforcement decision-maker in charge of PredPol in the jurisdiction. According to the law enforcement decision-maker, the change came about randomly. The change was implemented in all districts of the jurisdiction at once and is unlikely to be correlated with underlying time-varying unobservables at the district level that could be driving crime or arrest outcomes. Moreover, this change was not salient to law enforcement officers, who saw only that the PredPol boxes covered “All Crimes” in the patrol reports. Without the change, law enforcement would have continued to receive the All Crimes PredPol boxes for the original crime types; after the change, law enforcement received the All Crimes-Plus PredPol boxes for the original plus two more crime types. A notable feature of this institutional setting is that the PredPol software continued to predict the All Crimes PredPol boxes in the background even after they were no longer delivered to law enforcement.

[17] The date range 5/20/2019–3/1/2020, which I use for the quasi-experimental empirical strategy estimation window

3 Data

Estimating the effect of algorithm-induced police presence on crime incidents and racial disparities in arrests requires detailed data on predictive policing box locations, crime outcomes, arrests, and demographic information on any arrestees. To conduct my analysis, I assemble a unique data set that makes this analysis possible, combining novel data on (1) predictive policing box locations, (2) crime incident/call for service data for the crime types used in PredPol crime prediction (the input data to predict PredPol boxes), and (3) incident and arrest data from the jurisdiction. I collect the PredPol box data, which include the location of PredPol boxes for every shift in every district and the crime types for which the crime risk and PredPol boxes are predicted.

Obtaining data on actual police presence requires access to automatic vehicle locating (AVL) systems, which more police departments are starting to use for tracking. Unfortunately, the jurisdiction that I study does not record this kind of data.[18] While I do not have data on police presence itself, I use the predictive policing policy instrumentally to identify the effect of police presence in an intent-to-treat analysis. I describe each novel source of data that I collect. Then, I detail the box–shift-level panel dataset that I construct to conduct my analysis. Appendix A contains more information about the data collection.

1. **PredPol box location data ("algorithm output data"):** I observe the locations of 500 ft–by–500 ft PredPol boxes (GPS coordinates) and the shift/date for which they were generated. I observe the location of the All Crimes PredPol boxes that PredPol predicted in the background even after they were no longer delivered. Unfortunately, I only observe the All Crimes-Plus PredPol boxes for a lead-up of just a few days before the boxes actually began to be delivered. I create a list of the locations of the All Crimes and All Crimes-Plus PredPol boxes over a three year period (3/1/2018–3/1/2020). I call the boxes on this list the "ever-PredPol boxes."

[18] The use of such data for research is rare, though Weisburd (2021) and Blanes i Vidal and Mastrobuoni (2018) use AVL data to study the effect of police presence on overall crime outcomes, identifying the effects using plausibly exogenous shifts in police presence.

2. **Crime data used in PredPol box prediction ("algorithm input data"):** I observe the underlying crime and call for service data used to predict PredPol boxes — the PredPol algorithm input data. For each offense among the types of crimes used in prediction, I observe the date, start time, end time, offense type, address and GPS location. I also see whether each offense was excluded from prediction.[19] I map offenses to the box–shifts in which they occur using the GPS coordinates and incident start time.

3. **Incident/arrest data from jurisdiction:** The jurisdiction provides incident- and arrest-level data for my analysis. For each incident that results in an incident report, I observe the incident nature, incident report date and time, address (which I geocode using the *Google Maps Application Programming Interface*), suspect race, and victim race. For each incident, I observe the arrests that occurred, the race, age, gender of any arrestee(s), and whether the arrestee is Hispanic (ethnicity).

Panel dataset construction: Using the list of ever-PredPol boxes, I create a panel data set of box–time (date/shift) observations over a three-year period (3/1/2018–3/1/2020). A box i is included if it is ever a PredPol box over the three-year period. For every box–time observation, I observe the PredPol box treatment status of box i at time t. The outcomes of interest are crime incidence in box i at time t, arrests of Black individuals in box i at time t, arrests of Hispanic individuals in box i at time t, and arrests of white individuals in box i at time t. In Appendix A, I describe in detail how I construct the outcome variables by mapping crime incidents and arrests to the box–shifts in which they occur.

3.1 Descriptive Statistics

I follow the boxes that will be active PredPol boxes at least once over the three-year period from 3/1/2018 to 3/1/2020. There are 8,224 ever-PredPol boxes in the three-year panel dataset of ever-PredPol boxes that I assemble. Table 1 summarizes the average outcomes for all of these boxes

[19] Offenses can be excluded because they have a long duration (start to end time), because they are not properly geocoded, or because they are a duplicates.

over the 10-month study period in the "Total" column. Of these 8,224 boxes active at least once over the three-year period from 3/1/2018 to 3/1/2020, 1,937 boxes are active at least once over the 10-month period from 5/20/2019 to 3/1/2020; Table 1 shows summary statistics for these boxes that are ever active in the "Active" column. There are 6,287 boxes that are active sometime over the three-year period but never active in the 10-month window; Table 1 shows summary statistics for these boxes in the "Never Active" column. Note that this table summarizes averages for the boxes-shifts over the 10-month period, not only the outcomes from when the box is active.

Overall, the active boxes (that will be active at least once during the 10-month period) are active for an average of 21.262 shifts and have more crime incidents and more overall arrests, for Black, white and Hispanic individuals. The pattern emerges for both violent and property crimes (which together make up index crimes), and traffic incidents. This pattern is particularly striking for traffic incidents, as PredPol does not predict traffic incidents, underlining the importance of addressing the endogeneity problem that arises from Active boxes having higher underlying crime risk. These boxes that are active over the study period have more shots fired called in and are also slightly more likely to have Black individuals—both as victims and suspects— involved in the incident than the boxes that are never active.

Figure 1 plots the distribution of the number of day shifts in which a box is an active PredPol box over this 10-month period, using the sample of boxes ever active over this ten-month period. Among the boxes active at least once over the 10-month period, 49.61% are active for only five or fewer shifts. A total of 84.05% of these boxes are active for thirty or fewer day shifts, and 92.93% of them are PredPol boxes for 100 or fewer day shifts. No boxes are active every day over this period; the box that is most frequently a PredPol box is active for 267 shifts. This paper uses this high-frequency variation in boxes switching in and out of the PredPol box designation in every shift. Only 24 boxes per district are treated every shift. The vast majority of boxes that become active over the 10-month period are active very few times, highlighting the sparse nature of the treatment and the unlikelihood of a violation of the stable unit treatment values assumption (SUTVA) that the potential outcomes for a box vary with the treatment being assigned to other

Table 1: Summary statistics

	Never Active Boxes (1)	Active (2)	Total (3)
Violent Crimes	0.047	0.230	0.090
Arrests	0.028	0.128	0.052
Black	0.016	0.085	0.032
White	0.007	0.023	0.011
Hispanic	0.005	0.020	0.009
Property Crimes	0.152	1.169	0.391
Arrests	0.025	0.261	0.080
Black	0.012	0.116	0.037
White	0.008	0.082	0.025
Hispanic	0.005	0.062	0.018
Traffic Incidents	0.044	0.092	0.055
Arrests	0.041	0.086	0.052
Black	0.011	0.031	0.016
White	0.009	0.020	0.012
Hispanic	0.020	0.034	0.023
Shots fired called in	0.022	0.103	0.041
Black persons involved (fraction)	0.221	0.270	0.237
Black suspect (fraction)	0.209	0.252	0.224
Black victims (fraction)	0.243	0.293	0.260
Number of shifts box is active	0.000	21.262	5.008
N	6,287 (76.4%)	1,937 (23.6%)	8,224 (100.0%)

Notes: Total sample of all boxes ever designated as an active PredPol box over the three-year period from 3/1/2018 to 3/1/2020. Average statistics over the 10-month window from 5/20/2019 to 3/1/2020 for all boxes are in Column (3); statistics for boxes active at least once over the 10-month window are in Column (2); and statistics for boxes active sometime at least once over the three-year period but never active in the 10-month window are in Column (1). Crime and arrest outcomes are averages over box–shift outcomes. The outcomes "Black person[/suspect/victim]" are the fractions of total people involved in index crime incidents in the box who appear as Black in incident reports. The "Number of shifts box is active" outcome is the average number of shifts during which a box is designated a PredPol box over the 10-month window.

boxes.[20]

Figure 2 plots (1) the average number of index crimes and (2) the fraction of victims who are Black, by bin of the number of shifts in which a box is a PredPol box over the 10-month study window for all boxes designated as PredPol boxes within the 10-month period. Boxes more frequently designated as PredPol boxes have more index crimes on average. This is not surprising because PredPol predicts index crimes. In general, the percent of victims who are Black also increases with the number of shifts in which a box is active. If we take these details together, the vast majority of boxes are PredPol boxes highly infrequently (less than 5–10% of the time). Among these boxes, a lower proportion of victims are Black (only 25–30% of victims are Black). Very few boxes are active very often, and these boxes have more index crime on average.

Appendix Table B2 and Appendix Figure C3 show the distribution of boxes by the boxes' victim racial composition or fraction of victims who are Black. The mode of the distribution is 0% of victims who are Black: over 60% of the boxes do not have any index crime incidents with Black victims over a three-year period. There are several other points of the distribution with masses: nearly 9% of the boxes have 40-50% Black victims, and approximately 15% of the boxes have over 90% Black victims over a three-year period.

4 Research Design and Results

There are empirical complications to estimating the effects of algorithm-induced police presence: predictive policing box locations are located in the areas with the highest predicted crime risk within a district. Directly estimating the effect of algorithm-induced police presence, without accounting for the underlying crime risk associated with PredPol box designation, would lead to omitted variable bias. To circumvent these endogeneity concerns, I use a natural experiment research design to isolate quasi-experimental variation in police presence induced by predictive policing algorithms.

[20] Appendix Table B1 shows the percent of the time that a box is a PredPol box in the 10-month period using the sample of all boxes. Over 90% of boxes are a PredPol box at most 5% of the time.

Figure 1: Distribution of number of day shifts a box is designated as a predictive policing box from 5/20/2019 to 3/1/2020

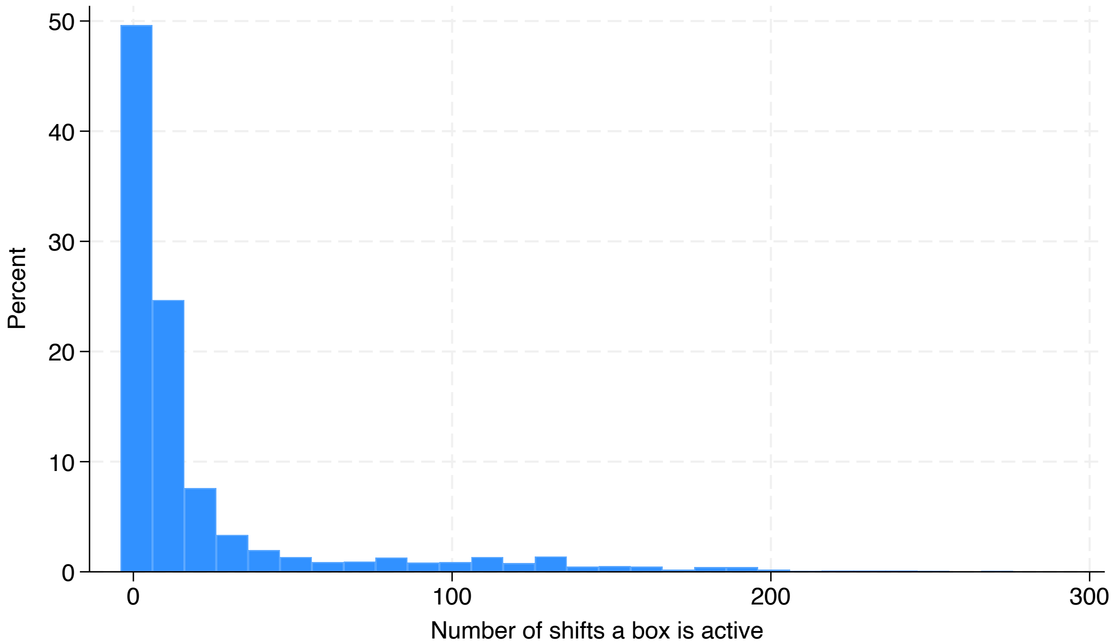

Figure 2: Average across boxes of number of index crimes and percent of victims who are Black by the number of day shifts a box is designated as a predictive policing box from 5/20/2019 to 3/1/2020

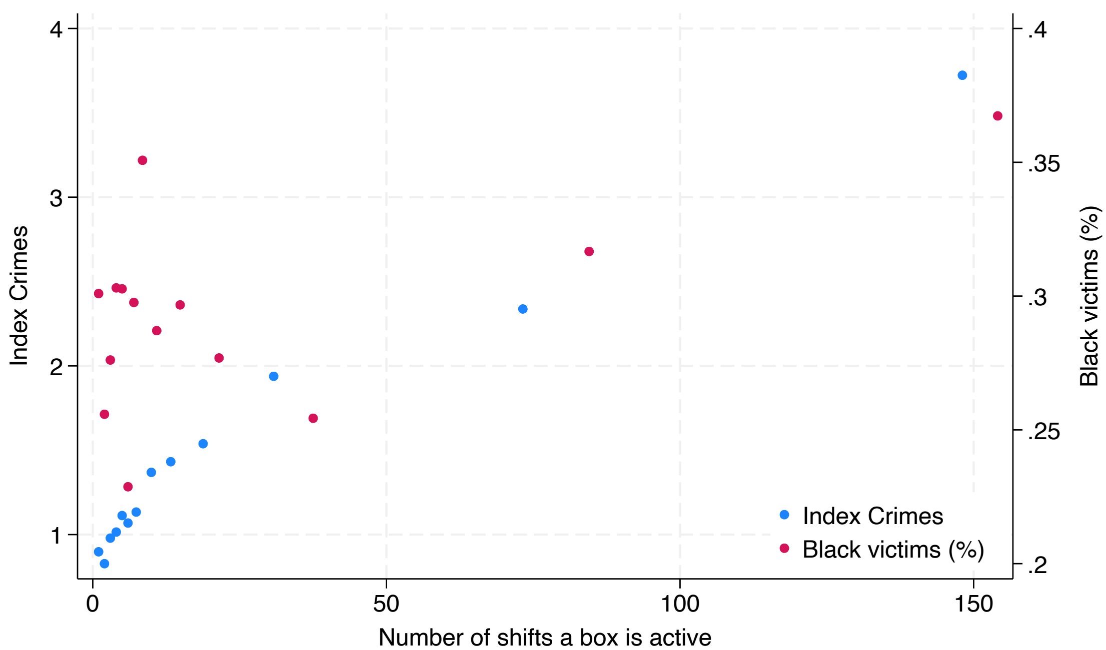

Notes: Figures 1 and 2 use the sample of the 1,924 boxes that are ever designated as a predictive policing box from 5/20/2019 to 3/1/2020.

My research design exploits an arbitrary change in the algorithm that slightly shifted the generation of the PredPol boxes delivered to police during the day shift. Before the change, PredPol generated its boxes using aggravated assault, auto burglary, motor vehicle theft, robbery, and shots fired violation codes. After the change, two more violation codes (residential and commercial burglary) were added to the original set.

Several features of the institutional setting and data make it possible to isolate quasi-experimental variation in algorithm-induced police presence from this change. First, after the change, law enforcement only sees the new sets of PredPol boxes every shift, and the change is not salient to law enforcement. Second, the change happens to all districts the jurisdiction at once, and is unlikely to be correlated with underlying district–level time–varying unobservables that could be driving crime or arrest outcomes. Third, after the change, for every shift, I observe where predictive policing boxes for the original crimes would have been if the change had not happened. This daily set of PredPol boxes that are not delivered after the change have similar underlying crime risk and make up a control group for predictive policing boxes that are delivered to law enforcement before the change. (Appendix Figure C4 illustrates this arbitrary change.)

My research design compares the outcomes of the active All Crimes PredPol boxes to those of the inactive All Crimes PredPol boxes with similar crime risk (which would have been active if not for the change), accounting for box fixed effects. The identifying variation comes from box i's switching in and of active PredPol box designation before the change together with box i's *not* switching in and out of PredPol box designation after this arbitrary change.

Using a window around this change, the following model estimates the effect of algorithm-induced police presence in predictive policing boxes, β:

$$Y_{it} = \beta Active\_AC_{it} + \delta AC_{it} + \xi Active\_ACPlus_{it} + \mu_i + \varepsilon_{it} \tag{2}$$

where Y_{it} are crime incidents in box i at time t, $Active\_AC_{it}$ is an indicator for whether box i is an active All Crimes PredPol box at time t, AC_{it} is an indicator for whether box i is an All Crimes

PredPol boxes at time t (active or inactive), and $Active\_ACPlus_{it}$ is an indicator for whether box i is an active All Crimes-Plus PredPol box at time t. $Active\_ACPlus_{it}$ is included to account for which boxes are actually active after the switch and any potential effect on crime, as well as for the overlap between All Crimes and All Crimes-Plus PredPol boxes. I can include box fixed effects μ_i to account for any time-invariant box characteristics, such that I can use the variation over time as boxes switch in and out of PredPol box designation and the variation in boxes that would have been included in the set of boxes switching in and out of PredPol box designation if not for the change to the algorithm. When a box is an All Crimes PredPol box ($AC_{it} = 1$) before the change ($Active\_AC_{it} = 1$), the box is a PredPol box that was actually delivered to law enforcement. When a box is an All Crimes PredPol box ($AC_{it} = 1$) after the change ($Active\_AC_{it} = 0$), the box was not delivered and is not an actual PredPol box.

To address any endogeneity concerns regarding whether active and inactive All Crimes PredPol boxes have different time-varying underlying crime risk, I account for how PredPol predicts crime risk using algorithm input data. I augment the model in equation 2 employing this selection-on-observables strategy. PredPol predicts a crime risk measure λ_{it} for box i at time t using the history of crime y_{it} for the box i and a background time-invariant box i parameter. I account for PredPol's underlying crime risk measure λ_{it} by controlling for crime lags for box i and including box fixed effects:

$$
\begin{aligned}
Y_{it} = & \beta Active\_AC_{it} + \delta AC_{it} + \xi Active\_ACPlus_{it} \\
& + \sum_{j=1}^{T} \gamma_j y_{it-j} + \mu_i + \phi_{dt} + \varepsilon_{it}
\end{aligned}
\tag{3}
$$

y_{it-j} are summed crime lags for both the crimes included in PredPol box prediction for All Crimes PredPol boxes and the crimes included for All Crimes-Plus PredPol boxes. β is the effect of algorithm-induced police presence in predictive policing boxes. By accounting for a proxy of PredPol's underlying crime risk measure, I address concerns that time-varying underlying crime risk (an omitted variable) could be correlated with ε_{it}. As a robustness check, I also include district–time fixed effects/district time trends ϕ_{dt} to control for any unobservable time-varying district-level trends that could be driving changes or outcomes. Standard errors are clustered at the box level. To estimate effects, I use a window around the arbitrary change from 5/20/2019 to 3/1/2020, cutting off of the sample around the start of the COVID-19 pandemic.

The identifying assumption is that, conditional on a box i being an All Crimes PredPol box at time t and an active All Crimes-Plus PredPol box at time t and on the inclusion of box fixed effects, the crime lags for both All Crimes and All Crimes-Plus PredPol boxes (PredPol's underlying crime risk proxy), and district–time fixed effects/district time trends, box i's designation at time t as an active All Crimes PredPol boxes must be orthogonal to omitted variables that could also affect crime in box i:

$$E[Active\_AC_{it} \cdot \varepsilon_{it} | AC_{it}, Active\_ACPlus_{it}, \mu_i, y_{it-j}, \phi_{dt} \ldots] = 0 \tag{4}$$

In this respect, a first concern could be that active and inactive All Crimes PredPol boxes and All Crimes-Plus PredPol boxes are different boxes. I address this concern by using box fixed effects to account for any time-invariant box unobservables that could be driving the results. A second concern could be that the box-level underlying crime risk that varies over time is an omitted variable correlated with ε_{it}. The arbitrary character of the change begins to address this concern in that I compare active with inactive All Crimes PredPol boxes. Moreover, I proxy for PredPol's underlying crime risk λ_{it} to account for the possibility that a box i happens to be inactive and not an All Crimes-Plus box when it has low crime risk. Finally, another concern is that district-level unobservables are correlated with ε_{it}. An example of this could be that there are district-level crime trends driving the results over time. However, the algorithm change happens to all districts at once and is unlikely to have been correlated with time-varying unobservables that could be driving the outcomes. I control for district time trends, or include district–time fixed effects to address this concern.

In a typical event study framework, balancing on past variables is expected. Such a balance check is not exactly applicable to the PredPol setting, as boxes switch in and out of PredPol box designation or treatment status for every shift. For example, a box that is a PredPol box in shift 5 may have been a PredPol box in shift 4. However, I do find, analogously to a balance check, that accounting for past crime does not change the results. If active and inactive All Crimes PredPol boxes have different past crime levels or histories, we would expect, in contrast, that the estimates change when past crime levels are accounted for.

Next, I briefly discuss mechanisms that could underlie the effect of predictive policing boxes and presence on reported crime incidents. First, police presence in predictive policing boxes may deter crime (Becker, 1968). Police presence increases the probability of apprehension and increases criminals' expected cost of committing a crime, which should decrease the likelihood that a crime is committed. Second, there may be more police–civilian interactions in PredPol boxes as a result of patrols being targeted there. "Crime" incidents that would have been previously undetected may be discovered with increased police presence, which would increase reported crimes. This mechanism is more likely to affect crimes where police have to be present to discover a crime or when they have discretion. On the other hand, this mechanism is also less likely to affect property crimes, where people report stolen items themselves.

4.1 Results

I examine the effect of algorithm-induced police presence on serious property and violent crime incidents[21], which corresponds to β from equation 3. Table 2 shows the results across several specifications. Overall, I find that algorithm-induced police presence statistically significantly deters around 2.909 to 3.077 crimes per thousand boxes, a nearly 30% reduction over the PredPol box mean of 9.348 crimes per thousand boxes. This finding is robust across specifications. Column (1) shows the baseline specification estimate, which additionally includes box fixed effects and controls for crime lags to address any additional endogeneity concerns about the quasi-experiment. In

[21] Serious property crimes include auto, commercial, and residential burglary and vehicle theft.Serious violent crimes include aggravated assault and robbery.

Table 2: Effect of algorithmic policing on serious property and violent crime incidents

	Box			Expanded box
	(1)	(2)	(3)	(4)
Active All Crimes PredPol Box	-3.077**	-2.943**	-2.909**	-4.844**
	(1.474)	(1.478)	(1.478)	(2.244)
Outcome mean	1.685	1.685	1.685	8.847
PredPol Box outcome mean	9.348	9.348	9.348	25.325
Box ID fixed effects	Yes	Yes	Yes	Yes
Underlying crime risk	Lags	Lags	Lags	Lags
District–time	No	Trends	Fixed effects	Fixed effects
Clusters	8224	8224	8224	8224
Observations	2352064	2352064	2352064	2352064

Notes: This table presents the estimates of β from equation 3. Serious property and violent crime incidents include aggravated assault, burglary, robbery, and motor vehicle theft. Sample of all box–shifts of boxes ever designated as All Crimes and All Crimes-Plus PredPol boxes over a three-year period. Regressions control for crime lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Lags include 7-day-shift lags and 12-month lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Standard errors are in parentheses and are clustered at the box level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Estimates and outcome means are multiplied by 1000.

particular, I use the box fixed effects and crime lags to proxy for PredPol's prediction of underlying crime risk. Column (2) shows the estimate for the baseline model additionally controlling for district time trends. Column (3) shows the estimate for the baseline model additionally including for district–time fixed effects. The estimates are robust across specifications; subsequent analysis of other outcomes utilize the augmented specification in Column (3). The estimates are consistent with results from prior papers finding that police patrols deter crime; (Weisburd, 2021) finds that a 10% decrease in police presence causes a 7% increase in crime.

I test whether there is crime displacement to areas around the corner from predictive policing boxes. Column (4) shows the effects of algorithm-induced police presence on serious property and violent crime incidents in predictive policing boxes and the 8 adjacent boxes, defined as the expanded box; the estimate is from the baseline specification augmented with district–time fixed effects. I expand the outcome to include the surrounding boxes to study whether crime is displaced to surrounding boxes. The effect of algorithm-induced police presence on the expanded box area is a decrease of 4.844 crimes per 1000 boxes, a 19% reduction over the PredPol expanded box outcome mean of 25.325 crimes per 1000 expanded boxes—evidence that crime is not being displaced to the areas around PredPol boxes. Overall, increased algorithm-induced police presence deters serious property and violent crime incidents in the expanded box.

My empirical approach focuses on algorithmically targeted areas opens the black box of predictive policing, and complements the work of (Mohler et al., 2015) and (Ratcliffe et al., 2020) on the effects of on overall crime of predictive policing as a method of allocating patrols vis-à-vis status quo approaches. I find decreases in serious violent and property crime in targeted areas and no evidence of displacement of crime to surrounding boxes. Mohler et al. (2015) find that predictive policing patrols reduced crime by 7.4% as a function of patrol time, while for status quo approaches, they find no statistically significant effect; Ratcliffe et al. (2020) find that using predictive policing–allocated patrols with marked cars results in a 31% decrease in property crime from the expected crime count.Taking my results on decreases in crime in predictive policing–targeted areas together with the results of (Mohler et al., 2015) and (Ratcliffe et al., 2020), it is plausible that that crime is not displaced further than adjacent areas (for which I am able to rule out the hypothesis of displacement). However, as my paper estimates treatment effects at the box and expanded box levels, I cannot speak to whether crime is displaced to areas outside the units of analysis, just as district- or city-level approaches cannot determine whether crime is displaced to areas outside the focal district or city.

Finally, in Table 3, I estimate the effect of algorithm-induced police presence on auxiliary measures capturing whether patrols actually go to PredPol boxes—traffic incidents and shots fired being called in. Columns (1) and (2) use the baseline specification additionally including district–time fixed effects. First, I find suggestive evidence that algorithm-induced police presence increases the number of traffic incidents by 0.779 traffic incidents per 1000 boxes over the expected average 1.675 traffic incidents per 1000 PredPol boxes over a 12 hour shift. While this effect is only statistically significant at the 10 % level, the finding is intuitive, as traffic incidents involve lower-

Table 3: Auxiliary evidence that law enforcement spend more time in PredPol boxes: Effect of algorithmic policing on traffic incidents and shots fired called in

	Traffic Incidents (1)	Shots fired called in (2)
Active All Crimes PredPol Box	0.779*	0.578**
	(0.428)	(0.261)
Outcome mean	0.393	0.143
PredPol Box outcome mean	1.675	0.486
Box ID fixed effects	Yes	Yes
Underlying crime risk	Lags	Lags
District–time	Fixed effects	Fixed effects
Clusters	8224	8224
Observations	2352064	2352064

Notes: This table presents the estimates of β from equation 3. Sample of all box–shifts of boxes ever designated as All Crimes and All Crimes-Plus PredPol boxes over a three-year period. Regressions control for crime lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Lags include 7-day-shift lags and 12-month lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Standard errors are in parentheses and are clustered at the box level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Estimates and outcome means are multiplied by 1000.

level offenses that are unlikely to be detected without the physical presence of law enforcement. When a patrol enters a predictive policing box, there may be a rise in police–civilian interactions, including traffic stops, which can result in traffic incidents being documented. Second, I examine the effects of algorithm-induced police presence on shots fired being called in. In the data, these calls for service are highly unlikely to become crime incidents. I find a statistically significant increase of 0.578 shots fired called in per 1000 boxes — a large increase relative to the expected average 0.486 shots fired called in per 1000 PredPol boxes over a 12 hour shift. While I do not directly observe patrol locations and presence, these estimates validate that patrols do indeed spend time in predictive policing box locations.

Overall, these findings show that predictive policing deters serious property and violent crime, but there is also suggestive evidence that it increases traffic incidents, an example of a lower-level offense unlikely to be detected without police presence. Taken together, these findings suggest that the effects of algorithm-induced police presence differ by crime type.

4.2 Alternative Research Design

As an alternative research design, I use the discontinuity in PredPol box treatment status at the algorithmic thresholds to estimate treatment effects. If I observed the underlying crime risk score λ, I could use it as a running score in a sharp regression discontinuity design (RDD). However, I do not observe this underlying crime risk score λ_{it}, only the data inputs used to predict it and the top 24 boxes in each shift designated as PredPol boxes. Moreover, while Mohler et al. (2015) have published the functional form of the model to predict λ_{it}, I cannot directly estimate this model, either, because I do not observe λ_{it}. I observe whether a box i at time t is a PredPol box (that is, among the top 24 boxes with the highest crime risk in its district at time t), which is a highly nonlinear function of λ_{it}. Therefore, I use machine-learning methods to predict λ, which I use in an RDD framework.

I apply the synthetic regression discontinuity design (SRDD) design framework of Boehnke and Bonaldi (2019). Using institutional details about how PredPol predicts underlying crime risk

(λ_{it}) along with the algorithm input data and PredPol box treatment status, I predict an estimate of λ_{it}. This estimate of λ_{it} is a synthetic running score $\hat{\lambda}_{it}$, which I use as a running variable in a sharp RDD.

Boehnke and Bonaldi (2019) propose a two-stage framework to identify and estimate the local average treatment effect using RDD when the running variable is unobservable but the treatment status is known. In the first stage, a synthetic score is predicted based on treatment status, and the second stage uses the synthetic score as a running variable in the RDD conditional on treatment status. The framework does not require that there be a continuous running score explicitly calculated by the decision-maker and underlying treatment assignment; it simply requires that the treatment assignment "can be described as if it were implicitly based on such a score" (Boehnke and Bonaldi, 2019). The PredPol context is a relevant setting in which to apply this framework, as I observe the PredPol box treatment status, which is explicitly based on the (unobserved) underlying crime risk prediction λ_{it} (Mohler et al., 2015).

The Boehnke and Bonaldi (2019) framework drops misclassified boxes to guarantee a discontinuity in the probability of PredPol box treatment AC_{it} at τ, the threshold of $\hat{\lambda}_{it}$:

$$\beta = \lim_{q \downarrow \tau} E[Y|\hat{\lambda} = q, AC_{it} = 1] - \lim_{q \uparrow \tau} E[Y|\hat{\lambda} = q, AC_{it} = 0] \tag{5}$$

The identifying assumptions are:

1. Continuity and smoothness of the unobserved running variable λ_{it} and synthetic score $\hat{\lambda}_{it}$

2. Perfect prediction of treatment status by the synthetic score in the first stage

In theory, the identifying assumption that the synthetic score perfectly predicts treatment status in the first stage is fulfilled in the PredPol context based on institutional knowledge. PredPol maintains that only three data points—crime type, crime time/date, and crime GPS coordinates—are used to predict crime risk and PredPol boxes. Therefore, these three data points should perfectly predict crime risk and PredPol box treatment status. Next, I apply the two-stage framework of

Boehnke and Bonaldi (2019) to examine the effects of algorithm-induced police presence in the PredPol setting:

4.2.1 First stage: Estimating the underlying crime risk synthetic running score

In the first stage, I estimate the continuous score $\hat{\lambda}_{it}$ underlying PredPol box treatment status:

$$AC_{it} = h(i, d, y_{it-1}, \ldots) \tag{6}$$

where AC_{it} is an indicator variable for whether box i is an All Crimes PredPol box. I use a multilayer perceptron (MLP) neural network to predict whether a box i at time t is a PredPol box using input vector $x = (i, d, y_{it-1}, \ldots)$, where i is the box ID, d is the district ID, and $\{y_{it-1}, \ldots\}$ is 1 year of crime lags for all the crime types used to predict All Crimes PredPol boxes. The MLP neural network is a kind of so-called deep neural network, which is a universal function approximator that thrives in large-scale data settings. I implement and train the neural network model using the open-source Keras/TensorFlow Python libraries.[22]

The neural network has two fully connected hidden layers, followed by rectified linear and sigmoidal activation functions, respectively.[23] For variables that take discrete values (discrete variables), I use an embedding function that maps the discrete variables to continuous features. For instance, PredPol includes a time-invariant box parameter in its crime risk model; I model this using an embedding space for both the district and box ID discrete variables; e.g., I map the discrete d and box i index to learned vector representations. There are more non-PredPol boxes than PredPol boxes; to address this imbalance, I use the class count weight to weight the loss function (Keras/TensorFlow feature).

The neural network outputs a PredPol box designation probability $\hat{\lambda}_{it}$ between 0 and 1 for every

[22] http://tensorflow.org. I used version 2.4.1 with GPU support.

[23] I explored the number and size of the MLP hidden layers and the stochastic gradient descent and adaptive moment estimation (Adam) learning algorithms.

observation x:

$$AC^{pred} = \begin{cases} 0 & \text{if } \hat{\lambda}_{it}(x) < 0.5 = \tau \\ 1 & \text{if } \hat{\lambda}_{it}(x) \geq 0.5 = \tau \end{cases} \tag{7}$$

I use a two-year sample of ever-PredPol boxes from 3/1/2018–3/1/2020, with 1 year of crime lags for the crime types used to predict All Crimes PredPol boxes. I randomly split the sample into a training set to train the neural network (60% of the sample) and a test set to test model performance out of sample (40% of the sample). Table 4 shows the performance of the neural network in prediction accuracy in the test set, defined as the percent of boxes for which PredPol box treatment status is correctly predicted. The best model is defined as the model with the best overall prediction accuracy (on the test set) that also equalizes prediction accuracy for PredPol boxes and non-PredPol boxes. The best performance in the test set achieves overall prediction accuracy of approximately 92.14%. Boehnke and Bonaldi (2019) also use machine learning for their first-stage prediction, achieving high accuracy of 97.9% in their validation set. The size of the data and the memory required to train the model limited how extensive the training and investigation of the MLP predictor could be. In the future, it may be possible to improve on the prediction accuracy.

4.2.2 Second stage: RDD using the synthetic running variable

In the second stage, I use the synthetic running variable $\hat{\lambda}_{it}$ as the running variable in a sharp RDD conditional on treatment status, estimated in the test set data:

$$Y_{it} = \alpha + \beta AC_{it} + P(\hat{\lambda}_{it}) + \varepsilon_{it} \tag{8}$$

where Y_{it} is the crime incidence in box i at time t, AC_{it} is the indicator for whether box i is an All Crimes PredPol box at time t, and $P(\hat{\lambda}_{it})$ is a polynomial function of the estimate of the continuous risk score, or synthetic running score, $\hat{\lambda}_{it}$. β is the local average treatment effect at the

Table 4: Out-of-sample (test set) predictive accuracy of multilayer perceptron (MLP) neural network

	Non-PredPol boxes (1)	PredPol boxes (2)	Overall (3)
MLP predictive accuracy	92.19%	92.11%	92.14%

Notes: I use an MLP neural network to predict whether a box i at time t is a PredPol box using input vector $x = (i, d, y_{it-1}, \dots)$, where i is the box ID, d is the district ID, and $\{y_{it-1}, \dots\}$ is 1 year of crime lags for all crime types used to predict All Crimes PredPol boxes. The neural network has two fully connected hidden layers, followed by rectified linear and sigmoidal activation functions, respectively.

margin of All Crimes PredPol box treatment conditioning on treatment status. Following Boehnke and Bonaldi (2019), I drop misclassified boxes to guarantee the discontinuity in probability of PredPol box treatment at τ, the threshold of $\hat{\lambda}_{it}$. I further restrict the sample to the period before 11/20/2019 when All Crimes PredPol boxes are active during the day.[24] I account for within-box correlation of errors over time with clustered standard errors. Following Boehnke and Bonaldi (2019), I use the bias-corrected RD estimator of Calonico et al. (2014) to "perform inference that is robust to the choice of bandwidth for the estimation of the local polynomials near the threshold" (Boehnke and Bonaldi, 2019).[25]

4.2.3 Results

Table 5 shows the effects of algorithm-induced police presence (Active All Crimes PredPol Box coefficient) on serious property and violent crime incidents (aggravated assault, burglary, robbery and vehicle theft). Column (1) shows the results from my main empirical strategy presented in Section 4, β from equation 3. Column (2) applies the SRDD framework of Boehnke and Bonaldi (2019) to further examine the effect of predictive policing box presence on crime as a robustness check. Column (2) finds a reduction of 4.954 serious property and violent crime incidents per

[24] For this period, $AC_{it} = Active\_AC_{it}$.

[25] This draft does not yet account for first-stage variation in the second-stage standard errors.

Table 5: Effect of algorithmic policing on serious property and violent crime incidents estimated using synthetic RDD

	Synthetic RDD (2)
Active All Crimes PredPol Box	-4.954***
	(1.337)
Conventional p-value	0.000
Robust p-value	0.000
Outcome mean	1.769
PredPol Box outcome mean	7.773
Observations	1344148

Notes: Serious index crime incidents include aggravated assault, burglary, robbery, and vehicle theft. Standard errors are in parentheses and are clustered at the box level. Estimates and outcome means are multiplied by 1000. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$.

1000 boxes, which is a 63.7% reduction relative to the average number of crimes in PredPol boxes during the day shift, 7.773 crimes per 1000 boxes. The two frameworks isolate different sources of quasi-experimental variation to estimate the treatment effects. Moreover, the SRDD estimates the local average treatment effect of algorithm-induced police presence in predictive policing boxes around the threshold of treatment. I find that the estimates from both frameworks have the same sign, implying robustness of the main research design and providing compelling support for my conclusions on the effects of algorithm-induced police presence.

5 Testing for Disproportionate Racial Impacts

In this section, I examine whether algorithm-induced police presence in predictive policing boxes has disproportionate racial impacts. Officers receive the locations of PredPol boxes, and can be differentially impacted by a box being designated a PredPol box.

To test whether algorithmic policing has differential effects by race, I examine the effect of algorithmic policing on law enforcement and arrests, in particular for traffic incidents and violent crimes.[26] Law enforcement has discretion to stop civilians.[27] Lower-level offenses that may otherwise have gone undetected may be more likely to be reported as crimes when they occur in a PredPol box. Police also have discretion to make arrests for lower-level offenses.[28] Moreover, traffic stops are the most common reasons for contact with the police.[29]

To examine whether algorithm-induced police presence has disproportionate racial impacts, I test whether arrests caused by algorithmic policing are disproportionately composed of minority arrests relative to white arrests. To determine whether the number of arrests is "disproportionate", I construct a counterfactual for how many arrests would have happened from each group if the box were not a PredPol box. A box not designated a PredPol box will still receive police presence through calls for service and patrol spending time there. To test whether the marginal arrests are disproportionately Black arrestees, I compare the number of Black marginal arrests with that of white marginal arrests, weighted by how many arrests would have happened in the non-PredPol box (counterfactual or inframarginal arrests). It is important to note that I am testing a relative measure of discrimination—how law enforcement affects Black compared to white motorists—and I do not explicitly account for whether an arrest was made in error.

To estimate the marginal arrests due to algorithmic policing by race, I use a nested model

[26]Unfortunately, the jurisdiction does not observe stops; such data could be used with arrests to perform a hit rate test to test for discrimination in the spirit of Knowles et al. (2001). The jurisdiction does observe traffic incidents, but not all traffic stops are written up in an incident report. All arrests in traffic stops have an incident report, however, and are necessarily be a traffic incident.

[27]Justice Sotomayor wrote in *Utah v. Strieff* in 2015, "This Court has allowed an officer to stop you for whatever reason he wants——so long as he can point to a pretextual justification after the fact. Whren v. United States, 517 U. S. 806, 813 (1996)."

[28]Justice Sotomayor wrote in *Utah v. Strieff* in 2015: "The officer's control over you does not end with the stop. If the officer chooses, he may handcuff you and take you to jail for doing nothing more than speeding, jaywalking, or 'driving [your] pickup truck . . . with [your] 3-year-old son and 5-year-old daughter . . . without [your] seatbelt fastened.' Atwater v. Lago Vista, 532 U. S. 318, 323–324 (2001)."

[29]Bureau of Justice Statistics. https://www.bjs.gov/index.cfm?tid=702&ty=tp

building on the main empirical strategy:

$$\begin{aligned} y_{it}^{R} = \sum_{r=b,h,w} 1\{R = r\}(\beta_r Active\_AC_{it} + \delta_r AC_{it} + \xi_r Active\_ACPlus_{it} \\ + \sum_{j=1}^{T} \gamma_{j,r} y_{it-j} + \mu_{i,r} + \phi_{dt,r}) + \varepsilon_{it} \end{aligned} \tag{9}$$

where y_{it}^{r} measures the number of arrests in box i at time t of individuals of race or Ethnicity R and r takes the values b (Black), h (Hispanic) and w (white). $Active\_AC_{it}$ is an indicator for box i being an active All Crimes PredPol box at time t, AC_{it} is an indicator that box i is All Crimes PredPol box at time t, and $Active\_ACPlus_{it}$ is an indicator that box i is an active All Crimes-Plus PredPol box at time t. β_r is the effect of police presence on arrests for individuals of race r.

The observed mean of arrests of group r in active All Crimes PredPol boxes is:

$$y_{obs,r} = \beta_r + \delta_r + \sum_{j=1}^{T} \gamma_{j,r} y_{it-j} + \mu_{i,r} + \phi_{dt,r} \tag{10}$$

The counterfactual mean of arrests of group r in active All Crimes PredPol boxes if treatment had not occurred is:

$$\begin{aligned} y_{cf,r} &= \delta_r + \sum_{j=1}^{T} \gamma_{j,r} y_{it-j} + \mu_{i,r} + \phi_{dt,r} \\ &= y_{obs,r} - \beta_r \end{aligned} \tag{11}$$

The arrests of group r that are caused by algorithm-induced police presence in active All Crimes PredPol boxes is the difference between $y_{obs,r}$ and $y_{cf,r}$ = β_r. Therefore, the counterfactual arrest mean in the absence of treatment is the number of inframarginal arrests.

To test whether marginal arrests caused by algorithmic policing disproportionately correspond to Black arrestees, I compare the Black with the white marginal effects, weighted by how many arrests would have happened without the PredPol box (inframarginal arrests):

$$\frac{\textit{Black marginal arrests}}{\textit{Black inframarginal arrests}} > \frac{\textit{White marginal arrests}}{\textit{White inframarginal arrests}}$$
$$\frac{\beta_{black}}{y_{cf,black}} > \frac{\beta_{white}}{y_{cf,white}} \tag{12}$$

Arrests are a function of potential offender and law enforcement behavior. Potential offenders do not know where PredPol boxes are, and should not be differentially impacted by algorithm-induced police presence compared to police presence not induced by algorithms. We would expect the effect on offenders to be the same whether they see an officer in a PredPol box or not, because offenders do not have this knowledge. In order for potential offender behavior to drive results, algorithm-induced police presence would need to affect Black vs white potential offender behavior differentially compared to how police presence not induced by algorithms (the counterfactual) does, which is unlikely.

First, I test the hypothesis that algorithmic policing has a disproportionately larger impact on high-discretion arrests of Black than of white motorists for traffic incidents. Law enforcement has greater discretion in traffic incidents than in other types of crimes. Moreover, law enforcement have a lot of discretion in who to stop and whether arrest. Second, I test the hypothesis for serious violent and property crime outcomes. Even though algorithm-induced police presence may deter crimes, law enforcement behavior may still change. Law enforcement may affect apprehension outcomes for serious property and violent crimes if law enforcement are ready to apprehend offenders.

5.1 Results

In Section 4.1, I found suggestive evidence of an increase in reported traffic incidents. Table 6 and Figure 3 show the results on whether there are racial disparities in the impacts of algorithm-induced police presence on arrests in traffic incidents. I find that arrests for traffic incidents due to algorithm-induced police presence in predictive policing boxes are disproportionately likely to be of Black arrestees. Algorithm-induced police presence has a marginally statistically significant positive effect on the number of Black arrests; I find insignificant effects for Hispanic and white arrests. The number of inframarginal arrests by race and ethnicity roughly mirrors the jurisdiction's population breakdown. Using the estimates for marginal and inframarginal arrests from equations 9 and 11, I test equation 12. Ultimately, after I account for how many arrests in traffic incidents would have happened without the predictive policing presence for each group, there is evidence of racial disparities in the effects of algorithm-induced police presence, with Black arrestees disproportionately likely to be arrested for traffic incidents in predictive policing boxes. I report p-values for a more conservative two-sided test, while a one-sided test is likely warranted based on the literature. The evidence suggests a three-fold increase in arrests for Black motorists when a box is designated a PredPol box compared to when it is not.

I present the results of the tests of disproportionate racial impacts of algorithm-induced police presence on arrests for serious violent crime (aggravated assault and robbery) in Figure 4 and Table 7. In Figure 4, the left panel shows the estimates for predictive policing boxes, and the right panel shows the estimates for the expanded box outcome (predictive policing boxes and surrounding boxes). In Section 4.1, I found a decrease in serious violent and property crime. I find that arrests for serious violent crime due to algorithm-induced police presence in predictive policing boxes and surrounding boxes are disproportionately likely to correspond to Black arrestees.

First, I find that algorithm-induced police presence has a statistically significant negative effect on the number of white arrests for both predictive policing boxes and expanded boxes. The estimates for Black marginal arrests and Hispanic marginal arrest estimates are statistically insignificant for both the box and the expanded box outcome. Second, in the absence of the box being designated a PredPol box, the counterfactual number of Black arrests is nearly double that of white arrests for both predictive policing boxes and expanded boxes. Third, I examine the marginal arrests weighted by the inframarginal arrests by race to test whether arrests due to algorithm-induced police presence are disproportionately made up of arrests of people of color, relative to the counterfactual outcome had there been no predictive policing box there.

Table 6: Test of disproportionate racial impacts of algorithm-induced police presence on arrests in traffic incidents

	(1)
White effect β_w	-0.162
	(0.241)
Black effect β_b	0.344*
	(0.181)
Hispanic effect β_h	0.131
	(0.136)
White inframarginal arrests $y_{cf,w}$	0.388
Black inframarginal arrests $y_{cf,b}$	0.109
Hispanic inframarginal arrests $y_{cf,h}$	0.210
White effect/infra.	-0.417
Black effect/infra.	3.156
Hispanic effect/infra.	0.624
P-value: Black effect/infra.= White effect/infra.	0.047
P-value: Hispanic effect/infra. = White effect/infra.	0.354
Box ID fixed effects	Yes
Underlying crime risk	Lags
District–time	Fixed effects
Clusters	8224
Observations	7056192

Notes: This table presents estimates of β_w, β_b, β_h, $y_{cf,w}$, $y_{cf,b}$, and $y_{cf,h}$ from equations 9 and 11. P-values are from two-sided tests. Sample of all box–shifts of boxes ever designated as All Crimes and All Crimes-Plus PredPol boxes over a three-year period. Regressions control for crime lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Lags include 7-day-shift lags and 12-month lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Standard errors are in parentheses and are clustered at the box level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Estimates and outcome means are multiplied by 1000.

Figure 3: Test of disproportionate racial impacts of algorithm-induced police presence on arrests in traffic incidents

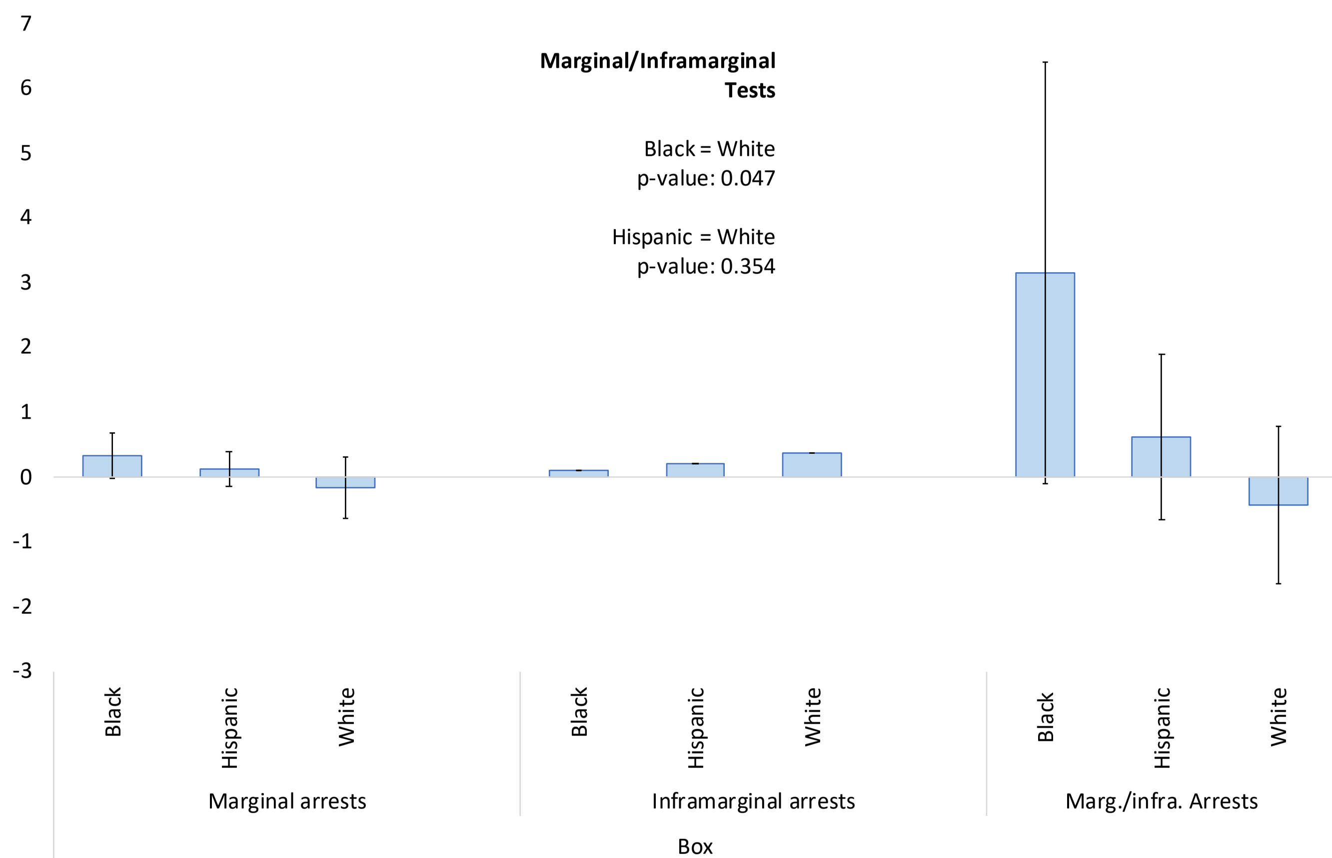

Notes: The left set of bars plots the estimates of the effect of algorithm-induced police presence by race (β_w, β_b, β_h from equation 9); the center set of bars plots the number of inframarginal arrests by race ($y_{cf,w}$, $y_{cf,b}$, and $y_{cf,h}$ from equation 11); the right set of bars plots the effect weighted by the number of inframarginal arrests by race. P-values are from two-sided tests testing equation 12.

Table 7: Test of disproportionate racial impacts of algorithm-induced police presence on arrests for serious violent crime

	Box (1)	Expanded box (2)
White effect β_w	-0.337**	-0.638**
	(0.166)	(0.277)
Black effect β_b	-0.084	0.470
	(0.307)	(0.635)
Hispanic effect β_h	0.018	0.098
	(0.127)	(0.169)
White inframarginal arrests $y_{cf,w}$	0.412	0.940
Black inframarginal arrests $y_{cf,b}$	0.990	1.945
Hispanic inframarginal arrests $y_{cf,h}$	0.095	0.204
White effect/infra.	-0.818	-0.679
Black effect/infra.	-0.084	0.242
Hispanic effect/infra.	0.189	0.480
P-value: Black effect/infra.= White effect/infra.	0.085	0.019
P-value: Hispanic effect/infra. = White effect/infra.	0.471	0.169
Box ID fixed effects	Yes	Yes
Underlying crime risk	Lags	Lags
District–time	Fixed effects	Fixed effects
Clusters	8224	8224
Observations	7056192	7056192

Notes: This table presents estimates of β_w, β_b, β_h, $y_{cf,w}$, $y_{cf,b}$, and $y_{cf,h}$ from equations 9 and 11. P-values are from two-sided tests. Serious violent crime includes aggravated assault and robbery. Sample of all box–shifts of boxes ever designated as All Crimes and All Crimes-Plus PredPol boxes over a three-year period. Regressions control for crime lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Lags include 7-day-shift lags and 12-month lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Standard errors are in parentheses and are clustered at the box level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Estimates and outcome means are multiplied by 1000.

Figure 4: Test of disproportionate racial impacts of algorithm-induced police presence on arrests for serious violent crime

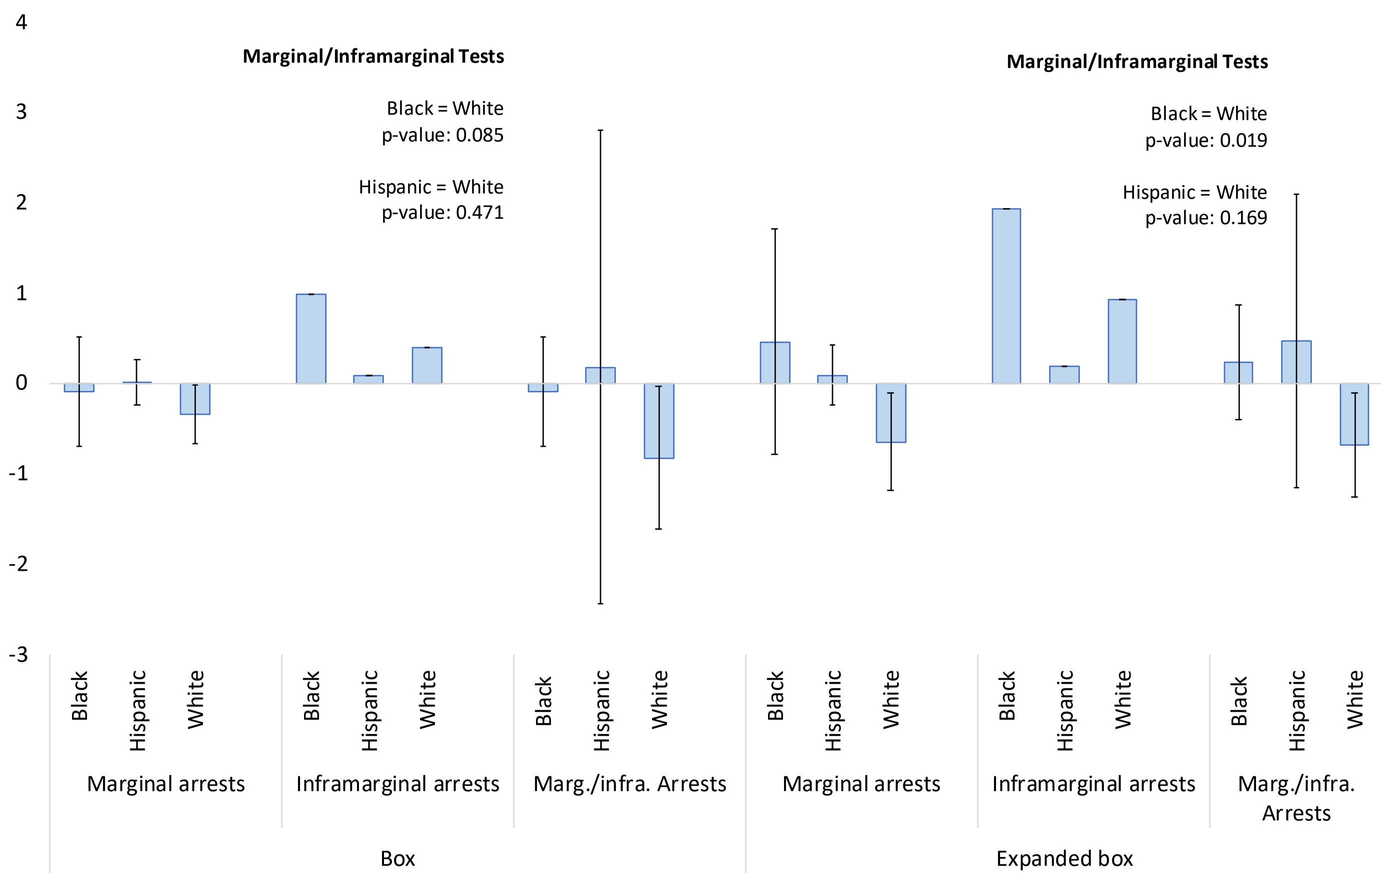

Notes: The left panel shows the estimates for predictive policing boxes and the right panel the estimates for predictive policing and surrounding boxes (the expanded box outcome). The left set of bars per panel plots the estimates of the effect of algorithm-induced police presence by race (β_w, β_b, β_h from equation 9); the center set of bars per panel plots the number of inframarginal arrests by race ($y_{cf,w}$, $y_{cf,b}$, and $y_{cf,h}$ from equation 11); the right set of bars per panel plots the effect weighted by the number of inframarginal arrests by race. P-values are from two-sided tests testing equation 12.

For the box outcome, I find suggestive evidence that is marginally statistically significant at the 10% level that there are racial disparities in arrests for serious violent crime. The estimate of marginal arrests weighted by the inframarginal arrests is negative for white arrestees. However, there is no proportional decrease for Black arrestees, and the estimate of marginal arrests weighted by the inframarginal arrests is negative and close to zero for Black arrestees. For the expanded box outcome in the right panel, there is a similar fall in the marginal arrests weighted by the inframarginal arrests for white arrestees and an increase in the number of arrests for Black arrestees. Moreover, I find statistically significant racial disparities in the impacts of algorithm-induced police presence, with a p-value of 0.019. There is a small number of Hispanic inframarginal arrests, and I cannot conclude whether there are racial disparities in the impacts of algorithm-induced police presence on Hispanic compared to white individuals. Arrests due to algorithm-induced police presence in predictive policing boxes and surrounding boxes are disproportionately of Black arrestees, providing evidence of disproportionate racial impacts of algorithm-induced police presence on arrests for serious violent crime. There is a reduction in arrests for serious violent crimes for white individuals around 0.67 times when a box is a PredPol box compared to not; there is no proportional decrease in arrests for Black individuals.

In Appendix Table B3, I fail to reject that there are racial disparities in the impacts of algorithm-induced police presence on arrests for serious property crime (burglary, vehicle theft). Examining the estimates of marginal arrests by race, I find suggestive evidence that is statistically significant at the 10% level that algorithm-induced police presence decreases the number of Black and white arrests for serious property crime. It is worth noting that the number of inframarginal Black arrests is over double the number of white arrests and over three times the number of Hispanic arrests.

5.2 Heterogeneity by Box Racial Composition

How are racial disparities in arrests exacerbated by algorithmic policing? I explore heterogeneity in the effects by neighborhood racial composition to better understand which boxes are driving the results. I found that Black motorists are disproportionately more likely than white motorists to be arrested in traffic incidents when a box is a PredPol box. Are officers more likely to arrest Black motorists in high-discretion traffic incidents when a box in predominantly Black neighborhood is targeted as a PredPol box? In other words, are results consistent with racially disproportionate law enforcement behavior being exacerbated in PredPol boxes in minority communities?

To test this hypothesis, I examine the effect of algorithmic policing on racial disparities across predominantly Black minority communities and non-Black communities. I focus on high-discretion arrests in traffic incidents. Traffic incidents are not predicted by PredPol, and there is no reason to believe that, in the absence of a change in officer enforcement, a box would have more traffic incidents arrests or more racial disparaties in traffic incidents when it is designated a PredPol box.

To explore heterogeneity by box racial composition, I use a split sample analysis. Focusing on minority communities that are predominantly Black, I create an indicator variable for whether a box is likely in a predominantly Black minority community; the indicator variable for a given box i is equal to 1 if more than 35% of the victims of index crimes over a three-year period are identified as Black in incident reports.[30] I categorize 3297 (40% of total) boxes as part of minority communities that are predominantly Black; the remaining 4927 (60% of the total) boxes are not (see Appendix Figure C3 and Appendix Table B2 for descriptive statistics on box racial composition).

Dividing the sample into boxes that are predominantly Black and boxes that are not, I estimate the nested model in equation 9 on both samples. This is equivalent to nesting the model again by interacting equation 9 with the indicator variable for whether a box is in a predominantly Black community. All the estimated parameters (β_{black}, β_{white} , . . .) can differ by box racial composition, which is a more flexible approach than one that forces estimated effects to be the same even with boxes of a different box racial composition. As before, I include box fixed effects, using the variation as boxes switch in and out of PredPol box designation, to address any potential concerns that the estimates are picking up differences in underlying time-invariant characteristics (e.g., in racial composition).

[30] To my knowledge, it is impossible to access racial demographic information at the level of granularity at which the 500 ft–by–500 ft boxes are defined.

5.2.1 Results

In Section 5.1, Black motorists are disproportionately more likely than white motorists to be arrested for traffic incidents when a box is a PredPol box; the evidence suggests that Black motorists are three times more likely to be arrested when a box is a PredPol box than when it is not.

Table 8 (and Appendix Figure C5) report the race-specific effects of algorithm-induced police presence—Black marginal arrests and white marginal arrests—for predominantly Black boxes and non-Black boxes. Black motorists are disproportionately more likely than white motorists to be arrested for traffic incidents when a box is a PredPol box, and even more so when a box is in a predominantly Black community. There is a marginally statistically significant effect of police presence in a PredPol box on arrests of Black motorists overall. The split-sample analysis reveals that boxes in predominantly Black areas drive this positive increase and the widening of the racial disparities. For predominantly Black areas, there is an increase of 0.733 arrests in traffic incidents, double the 0.344 marginal arrests that are due to algorithmic policing in traffic incidents across all boxes. The marginal effect of police presence in a PredPol box on racial disparities widens even further in boxes in predominantly Black compared to boxes in non-Black neighborhoods.

On the other hand, the race-specific effects of boxes in non-Black areas are very small, though the individuals estimates are not statistically significant. For boxes in non-Black areas, the effect on Black marginal arrests is 0.067, and that on white marginal arrests is a decrease of 0.025—that is, very small and close to zero. There is no evidence that PredPol boxes have an effect on racial disparities in arrests in non-Black areas. These results are consistent with predominantly Black neighborhoods driving the results.

PredPol boxes signal to officers that an area is higher crime risk, which could potentially bring them closer to the probable cause threshold required for stops and searches governed by the Fourth Amendment. However, if officers are using PredPol boxes as a signal that an area is riskier, they should have more interactions in all PredPol boxes, regardless of the racial composition of the community in which the box is located. These uneven results are consistent with PredPol boxes enabling officers to exercise more discretion, which could allow them to discriminate in these tar-

Table 8: Race-specific effects of algorithm-induced police presence on arrests in traffic incidents: Heterogeneity by neighborhood racial composition

	Predom. Black Neighborhoods Sample (2)	Non predom. Black Neighborhoods Sample (3)
Black marginal arrests β_b	0.733**	0.067
	(0.300)	(0.219)
Hispanic marginal arrests β_h	0.169	0.136
	(0.129)	(0.217)
White marginal arrests β_w	-0.381	-0.025
	(0.472)	(0.261)
Box ID fixed effects	Yes	Yes
Underlying crime risk	Lags	Lags
District–time	Fixed effects	Fixed effects
Clusters	4794	3430
Observations	4113252	2942940

Notes: Column (1) and (2) reports the estimates of race-specific effects of algorithm-induced police presence—Black marginal arrests (β_b) and white marginal arrests (β_w)—on arrests in traffic incidents, from 9, in the subset of predominantly Black and non-Black boxes, respectively. Boxes are classified as predominantly Black if more than 35% of the victims of index crimes over a three-year period are identified as Black in incident reports. Boxes are non-Black if 35% or fewer of the victims of index crimes over a three-year period are identified as Black in incident reports. Standard errors are in parentheses and are clustered at the box level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Estimates and outcome means are multiplied by 1000.

getted areas. Ultimately, these findings are consistent with officers treating individuals differently by race and additionally by the racial composition of their neighborhood. We do not see evidence that greater discretion enabled by PredPol is applied evenly across all boxes designated as PredPol boxes. The results – of disproportionately more arrests of Black motorists in predominantly Black communities – suggest that a disproportionate burden of algorithmic surveillance falls on Black individuals in predominantly Black communities.

6 Discussion and Conclusion

As algorithmic targeting is increasingly used, it is important to understand societal tradeoffs involved. This paper investigates the effects of algorithm-induced police presence on crime incidents and racial disparities in arrests using two natural experiment research designs. My findings indicate that algorithm-induced police presence deters serious violent and property crime, with no evidence of displacement of crime to the immediate surrounding areas. I validate that law enforcement outcomes are shaped by officer deployments to predictive policing boxes, finding an increase in shots fired called in and suggestive evidence that algorithm-induced police presence increases the number of traffic incidents. There is evidence that algorithm-induced police presence has disproportionate racial impacts on arrests for serious violent crimes in PredPol boxes and surrounding boxes and on arrests in traffic incidents in PredPol boxes. Arrests in predominantly Black neighborhoods drive these racial disparities. Overall, the results indicate that algorithmic policing matters, and the disproportionate burden of algorithmic surveillance falls on Black individuals in predominantly Black neighborhoods.

My focus on algorithmically targeted areas sheds light on the open questions of how algorithmic policing shapes outcomes and whether it has disproportionate impacts by race. Using the predictive policing institutional setting, I examine the local impacts of algorithm-induced police presence on crime and racial disparities in arrests. From this setting, we learn more about the effects of targeting of police patrols on racial disparities in arrests, which we know little about.

I provide new evidence on the effects of algorithmic policing, and local police presence more generally, on racial disparities in arrests. Chalfin et al. (2022) find that the effects of larger police forces differ by race. A larger police force size causally increases the number of low-level "quality-of-life" offenses such as drug possession and disorderly conduct, in particular for Black individuals. Similarly to Chalfin et al. (2022), I find suggestive evidence that algorithm-induced police presence increases the number of traffic incidents and find that the impact of this presence is disproportionately larger for Black arrestees than for white arrestees in traffic incidents—where police have more discretion. Chalfin et al. (2022) also find that larger police forces decrease serious property and violent crimes, in particular for Black suspects. While I also find that serious property and violent crime decrease, I find that white individuals are arrested less when a place is designated as a predictive policing box, with no proportional decrease for Black individuals. Overall, my findings suggest that increased local police presence has disproportionate racial impacts on Black relative to white individuals. While Chalfin et al. (2022) study the effect of a larger police force at the city-level, I study the local effects of algorithmic policing in targeted areas, and am able to study effects by neighborhood racial composition. These different institutional settings and levels of analysis answer different policy questions, which could drive different findings, highlighting the need for future research.

Ultimately, my finding that the impacts of algorithm-induced police presence on crime differ by crime type has implications for concerns about feedback in algorithms. I find that, among the violation code types used to generate predictive policing boxes, algorithm-induced police presence increases shots fired offenses called in. In the jurisdiction that I study, traffic incidents are not used to generate predictive policing boxes. However, I find suggestive evidence that algorithm-induced police presence increases reported traffic incidents, underlining that lower-level offenses where police have discretion are of particular concern for algorithmic feedback.

While there is evidence that algorithm-induced police presence deters certain crime types, there are important equity implications of agencies' using predictive policing to target areas, as there is also evidence that Black arrestees are disproportionately arrested for certain crimes as a result of algorithm-induced police presence. The empirical strategies that I develop and use can be applied to data from other cities to further study impacts of algorithmic policing, to speak more to the external validity of this analysis. It is also possible to use this framework to study how victims and other stakeholders are affected by algorithmic policing to further understand how algorithms affect our lives. It will be important in the future to holistically consider the costs and benefits of using predictive policing technologies, especially considering the nuanced effects by race that I find. Future work can consider the general equilibrium effects of algorithmic targeting on racial disparities in arrests in cities that use these technologies.[31]

[31] Fu and Wolpin (2018), Galiani et al. (2018) and Ba et al. (2021a) use equilibrium models to examine where and how patrols should be in cities and across cities.

References

Agan, A. Y. (2022). Racial disparities in the criminal legal system: Shadows of doubt and beyond. *Journal of Economic Literature*.

Antonovics, K. and Knight, B. G. (2009). A new look at racial profiling: Evidence from the boston police department. *The Review of Economics and Statistics*, 91(1):163–177.

Anwar, S. and Fang, H. (2006). An alternative test of racial prejudice in motor vehicle searches: Theory and evidence. *American Economic Review*, 96(1):127–151.

Ba, B., Bayer, P., Rim, N., Rivera, R., and Sidibé, M. (2021a). Police officer assignment and neighborhood crime. Technical report, National Bureau of Economic Research.

Ba, B. A., Knox, D., Mummolo, J., and Rivera, R. (2021b). The role of officer race and gender in police-civilian interactions in chicago. *Science*, 371(6530):696–702.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217.

Blanes i Vidal, J. and Mastrobuoni, G. (2018). Police patrols and crime.

Blattman, C., Green, D. P., Ortega, D., and Tobón, S. (2021). Place-based interventions at scale: The direct and spillover effects of policing and city services on crime. *Journal of the European Economic Association*, 19(4):2022–2051.

Boehnke, J. and Bonaldi, P. (2019). Synthetic regression discontinuity: Estimating treatment effects using machine learning.

Braga, A. A., Turchan, B. S., Papachristos, A. V., and Hureau, D. M. (2019). Hot spots policing and crime reduction: an update of an ongoing systematic review and meta-analysis. *Journal of experimental criminology*, 15(3):289–311.

Brantingham, P. J., Valasik, M., and Mohler, G. O. (2018). Does predictive policing lead to biased arrests? results from a randomized controlled trial. *Statistics and Public Policy*, 5(1):1–6.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Cameron, D. (2022). Justice department admits: We don't even know how many predictive policing tools we've funded. *Gizmodo*.

Chalfin, A., Hansen, B., Weisburst, E. K., and Williams Jr, M. C. (2022). Police force size and civilian race. *American Economic Review: Insights*, 4(2):139–58.

Chalfin, A. and McCrary, J. (2017). Criminal deterrence: A review of the literature. *Journal of Economic Literature*, 55(1):5–48.

Chalfin, A. and McCrary, J. (2018). Are us cities underpoliced? theory and evidence. *Review of Economics and Statistics*, 100(1):167–186.

Chen, M. K., Christensen, K. L., John, E., Owens, E., and Zhuo, Y. (2021). Smartphone data reveal neighborhood-level racial disparities in police presence. *arXiv preprint arXiv:2109.12491*.

Di Tella, R. and Schargrodsky, E. (2004). Do police reduce crime? estimates using the allocation of police forces after a terrorist attack. *American Economic Review*, 94(1):115–133.

Doleac, J. L. (2022). Racial bias in the criminal justice system. In *A Modern Guide to the Economics of Crime*, pages 286–304. Edward Elgar Publishing.

Draca, M., Machin, S., and Witt, R. (2011). Panic on the streets of london: Police, crime, and the july 2005 terror attacks. *American Economic Review*, 101(5):2157–81.

Durlauf, S. N. and Nagin, D. S. (2011). Imprisonment and crime: Can both be reduced? *Criminology & Public Policy*, 10(1):13–54.

Feigenberg, B. and Miller, C. (2022). Would eliminating racial disparities in motor vehicle searches have efficiency costs? *The Quarterly Journal of Economics*, 137(1):49–113.

Ferguson, A. G. (2012). Predictive policing and reasonable suspicion. *Emory LJ*, 62:259.

Fryer Jr, R. G. (2019). An empirical analysis of racial differences in police use of force. *Journal of Political Economy*, 127(3):1210–1261.

Fu, C. and Wolpin, K. I. (2018). Structural estimation of a becker-ehrlich equilibrium model of crime: Allocating police across cities to reduce crime. *The Review of Economic Studies*, 85(4):2097–2138.

Galiani, S., Cruz, I. L., and Torrens, G. (2018). Stirring up a hornets' nest: Geographic distribution of crime. *Journal of Economic Behavior & Organization*, 152:17–35.

Goncalves, F. and Mello, S. (2021). A few bad apples? racial bias in policing. *American Economic Review*, 111(5):1406–41.

Gonzalez, R. M., Komisarow, S., and Jabri, R. (2022). When does crime respond to punishment?: Evidence from drug-free school zones.

Grogger, J. and Ridgeway, G. (2006). Testing for racial profiling in traffic stops from behind a veil of darkness. *Journal of the American Statistical Association*, 101(475):878–887.

Grosjean, P., Masera, F., and Yousaf, H. (2022). Inflammatory Political Campaigns and Racial Bias in Policing*. *The Quarterly Journal of Economics*, 138(1):413–463.

Hoekstra, M. and Sloan, C. (2022). Does race matter for police use of force? evidence from 911 calls. *American Economic Review*, 112(3):827–60.

Horrace, W. C. and Rohlin, S. M. (2016). How dark is dark? bright lights, big city, racial profiling. *Review of Economics and Statistics*, 98(2):226–232.

Klick, J. and Tabarrok, A. (2005). Using terror alert levels to estimate the effect of police on crime. *The Journal of Law and Economics*, 48(1):267–279.

Knowles, J., Persico, N., and Todd, P. (2001). Racial bias in motor vehicle searches: Theory and evidence. *Journal of Political Economy*, 109(1):203–229.

Levitt, S. D. (1997). Using electoral cycles in police hiring to estimate the effect of police on crime. *The American Economic Review*, 87(3):270–290.

Levitt, S. D. (2002). Using electoral cycles in police hiring to estimate the effects of police on crime: Reply. *American Economic Review*, 92(4):1244–1250.

Lin, M.-J. (2009). More police, less crime: Evidence from us state data. *International Review of Law and Economics*, 29(2):73–80.

Lum, K. and Isaac, W. (2016). To predict and serve? *Significance*, 13(5):14–19.

MacDonald, J. M. and Fagan, J. (2019). Using shifts in deployment and operations to test for racial bias in police stops. 109:148–51.

MacDonald, J. M., Klick, J., and Grunwald, B. (2016). The effect of private police on crime: evidence from a geographic regression discontinuity design. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 179(3):831–846.

Mastrobuoni, G. (2020). Crime is terribly revealing: Information technology and police productivity. *The Review of Economic Studies*, 87(6):2727–2753.

McCrary, J. (2002). Using electoral cycles in police hiring to estimate the effect of police on crime: Comment. *American Economic Review*, 92(4):1236–1243.

Mohler, G. O., 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(512):1399–1411.

O'Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.

Police Executive Research Forum (2014). Future trends in policing. Technical report, Office of Community Oriented Policing Services.

Ratcliffe, J. H., Taylor, R. B., Askey, A. P., Thomas, K., Grasso, J., Bethel, K. J., Fisher, R., and Koehnlein, J. (2020). The philadelphia predictive policing experiment. *Journal of Experimental Criminology*, pages 1–27.

Weisburd, D. and Telep, C. W. (2014). Hot spots policing: What we know and what we need to know. *Journal of Contemporary Criminal Justice*, 30(2):200–220.

Weisburd, S. (2021). Police presence, rapid response rates, and crime prevention. *Review of Economics and Statistics*, 103(2):280–293.

Weisburst, E. K. (2019). Safety in police numbers: Evidence of police effectiveness from federal cops grant applications. *American Law and Economics Review*, 21(1):81–109.

West, J. (2018). Racial bias in police investigations. *Unpublished manuscript*.

A Appendix: Data

A.1 Collection

This section describes how the PredPol data was collected. To my knowledge, PredPol has never shared their list of customers. Even beyond PredPol, it can be difficult to ascertain which law enforcements use predictive policing. Even the U.S. Department of Justice (DOJ), a leading source of funding to purchase predictive policing tools, "[has] kept no 'specific records', a senior official has said, with regard to which agencies have tapped a leading source of DOJ funding to purchase predictive policing tools" (Cameron, 2022). I created a list of around 60 law enforcement agencies that could be using PredPol from (1) PredPol's website, marketing materials, and papers papers, (2) media reporting law enforcement agencies that had signed, trialed or began using PredPol, and (3) records of public records requests regarding PredPol and responses from law enforcement agencies from Muckrock – a website that helps users to track public records requests, "adding correspondence and responsive records to the public domain".[32] I validated this list by web searching for public documents and further evidence of PredPol use, then I mailed three rounds of letters, starting in 2018, to law enforcement agencies to inquire whether they would be interested to share data. The agencies asked PredPol to grant me access to their PredPol sites and Application Programming Interface (API). I used the PredPol sites and API to gather the PredPol box location data (February 2017 to August 2020), and algorithm input data described in Section 3. The agencies also shared incident-level and arrest-level data described in Section 3. I detail how I construct the panel dataset at the box-shift level and how I merge in the incident-level and arrest-level data.

A.2 Panel Construction

I use the list of boxes that are ever day-shift PredPol boxes over a three-year period (3/1/2018-3/1/2020) ("ever-PredPol boxes") to create a box-shift-level panel dataset. The sample is 3 years worth of day shifts for all 8,224 ever-PredPol boxes. I exclude one day shift (February 14, 2019)

[32] Muckrock.com

with data irregularities.

A.3 PredPol Box data

For every box-shift, I observe whether a box is a PredPol box that is active. That is, I observe whether a box is an All Crimes PredPol box before 11/20/2019, and whether a box is an All Crimes Plus PredPol box after 11/20/2019. I also observe whether the box has an inactive All Crimes PredPol box after 11/20/2019.

A.4 Incident-level data

There are two source of incident-level data: PredPol algorithm input data and data from the jurisdiction. I use the incident address to map incidents to the boxes in which they occur. For crime types used for prediction (auto burglary offenses, vehicle theft offenses, robbery offenses, residential burglary offenses, commercial burglary offenses, assault offenses and shots fired calls for service), I observe the incident start time, which I use to map incidents to the shift in which they started. For incidents for crime types outside of this set of crime types (traffic incidents), I only observe the incident report time, which I use to map these incidents to the shift in which they occurred.

For the subset of crime offenses used in prediction (contained in both the second and third sources of data above), I examine the difference between the incident start time and the incident report time. The difference is large for burglary and motor vehicle theft offenses. Burglary and motor vehicle theft are non-violent property crimes that do not involve a personal threat of violence. Incidents might be reported at a later time after discovery. There is a smaller difference between incident and report time for assault and robbery offenses, which are violent crimes where there is a personal threat of violence. Based on this analysis, I exclude crime types from my analysis where there may be gaps between incident time and incident report time. I focus on traffic incidents where law enforcement are likely initiating traffic stops, and the incident start time. There is unlikely to be much time lapsed to the incident report time. In the future, I hope to get access to better incident start time data to perform analysis for more types of crime incidents.

A.5 Arrest-level data

Every arrest has a corresponding incident report, therefore I merge arrest-level information into the incident-level data. This means arrests are being mapped into the box-shift level data using the incident address and incident time. For arrests for incidents among the crime types used for prediction, I use the incident start time, and for arrests in traffic incidents, I use the incident report time.

I do observe the arrest date/time and address if an arrest is able to be physically made at the time of the incident report. However, there are also arrests that happen after the incident report time for which I observe the arrest date/time and address. Because of these unresolved data complications, I rely on the incident address and incident time to map arrests to box-shifts rather than the arrest date/time or arrest address.

B Appendix: Tables

Table B1: Distribution of percent of shifts in which a box is designated as a predictive policing box over the quasi-random experiment window (5/20/2019-3/1/2020)

Percent of time box is a PredPol Box	Percent of boxes	Percent of victims who are Black
0	76.47	24.31
(0-5]	17.21	28.94
(5-10]	2.44	27.22
(10-15]	0.75	25.82
(15-20]	0.49	28.71
(20-40]	1.40	32.43
(40-60]	0.89	35.09
(60-80]	0.32	32.57
(80-100]	0.02	53.30

Notes: This table summarizes the percent of shifts in which a box is designated an All Crimes predictive policing box over the quasi-random experiment window (5/20/2019–3/1/2020) for all boxes.

Table B2: Distribution of box victim racial composition

Black victims (Percent for box)	Percent of boxes
0	60.20
(0-10]	0.32
(10-20]	1.85
(20-30]	4.04
(30-40]	3.53
(40-50]	8.75
(50-60]	0.65
(60-70]	3.21
(70-80]	1.56
(80-90]	1.16
(90-100]	14.74

Notes: This table summarizes the percent of boxes with the specifiedpercentages of victims identified as Black in crime incident reports for index crime incidents that occur over three years.

Table B3: Test of disproportionate racial impacts of algorithm-induced police presence on arrests for serious property crime

Dep. var.: 1000 x (Arrests for serious property crime)	(1)
White effect β_w	-0.550*
	(0.284)
Black effect β_b	-0.888*
	(0.506)
Hispanic effect β_h	-0.208
	(0.992)
White inframarginal arrests $y_{cf,w}$	0.550
Black inframarginal arrests $y_{cf,b}$	1.265
Hispanic inframarginal arrests $y_{cf,h}$	0.359
White effect/infra.	-1.000
Black effect/infra.	-0.702
Hispanic effect/infra.	-0.579
P-value: Black effect/infra.= White effect/infra.	0.639
P-value: Hispanic effect/infra. = White effect/infra.	0.863
Box ID fixed effects	Yes
Underlying crime risk	Lags
District–time	Fixed effects
Clusters	8224
Observations	7056192

Notes: This table presents estimates of β_w, β_b, β_h, $y_{cf,w}$, $y_{cf,b}$, and $y_{cf,h}$ from equations 9 and 11. P-values are from two-sided tests testing 12. Serious property crimes include residential burglary, commercial burglary, auto burglary, and motor vehicle theft. Sample of all box–shifts of boxes ever designated as All Crimes and All Crimes-Plus PredPol boxes over a three-year period. Regressions control for crime lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Lags include 7-day-shift lags and 12-month lags summing the crimes included in prediction for All Crimes and All Crimes-Plus PredPol boxes. Standard errors are in parentheses and are clustered at the box level. * p < 0.10 ** p < 0.05 *** p < 0.01. Estimates and outcome means are multiplied by 1000.

Table B4: Race-specific effects of algorithm-induced police presence on arrests for violent crime incidents: Heterogeneity by neighborhood racial composition

	A. Box			B. Expanded box		
	All	Predom. Black	Non predom. Black	All	Predom. Black	Non predom. Black
	(1)	(2)	(3)	(4)	(5)	(6)
Black marginal arrests β_b	-0.055	-0.554	0.284	0.463	0.402	0.557
	(0.306)	(0.671)	(0.241)	(0.635)	(1.413)	(0.466)
Hispanic marginal arrests β_h	0.023	-0.069	0.078	0.101	0.084	0.127
	(0.127)	(0.139)	(0.194)	(0.169)	(0.207)	(0.251)
White marginal arrests β_w	-0.332**	-0.270	-0.388*	-0.632**	-0.597	-0.645*
	(0.165)	(0.273)	(0.210)	(0.276)	(0.400)	(0.366)
Box ID fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Underlying crime risk	Lags	Lags	Lags	Lags	Lags	Lags
District–time	FE	FE	FE	FE	FE	FE
Clusters	8224	4794	3430	8224	4794	3430
Observations	7056192	4113252	2942940	7056192	4113252	2942940

Notes: This figure presents estimates of race-specific effects of algorithm-induced police presence—Black marginal arrests (β_b) and white marginal arrests (β_w)—on violent crime incident arrests, from equation 9. Panel A reports results for the box, and Panel B reports results for the extended box. Column (1) and Column (4) report race-specific effects estimated in the sample of all boxes. Columns (2) and (5)report race-specific effects estimated in the subset of predominantly Black boxes. Columns (3) and (6) report race-specific effects estimated in the subset of non-Black boxes. Boxes are classified as predominantly Black if more than 35% of the victims of index crimes over a three-year period are identified as Black in incident reports. Boxes are non-Black if 35% or fewer of the victims of index crimes over a three-year period are identified as Black in incident reports. Standard errors are in parentheses and are clustered at the box level. * $p < 0.10$ ** $p < 0.05$ *** $p < 0.01$. Estimates and outcome means are multiplied by 1000.

C Appendix: Figures

Figure C1: Sample of a PredPol daily prediction report

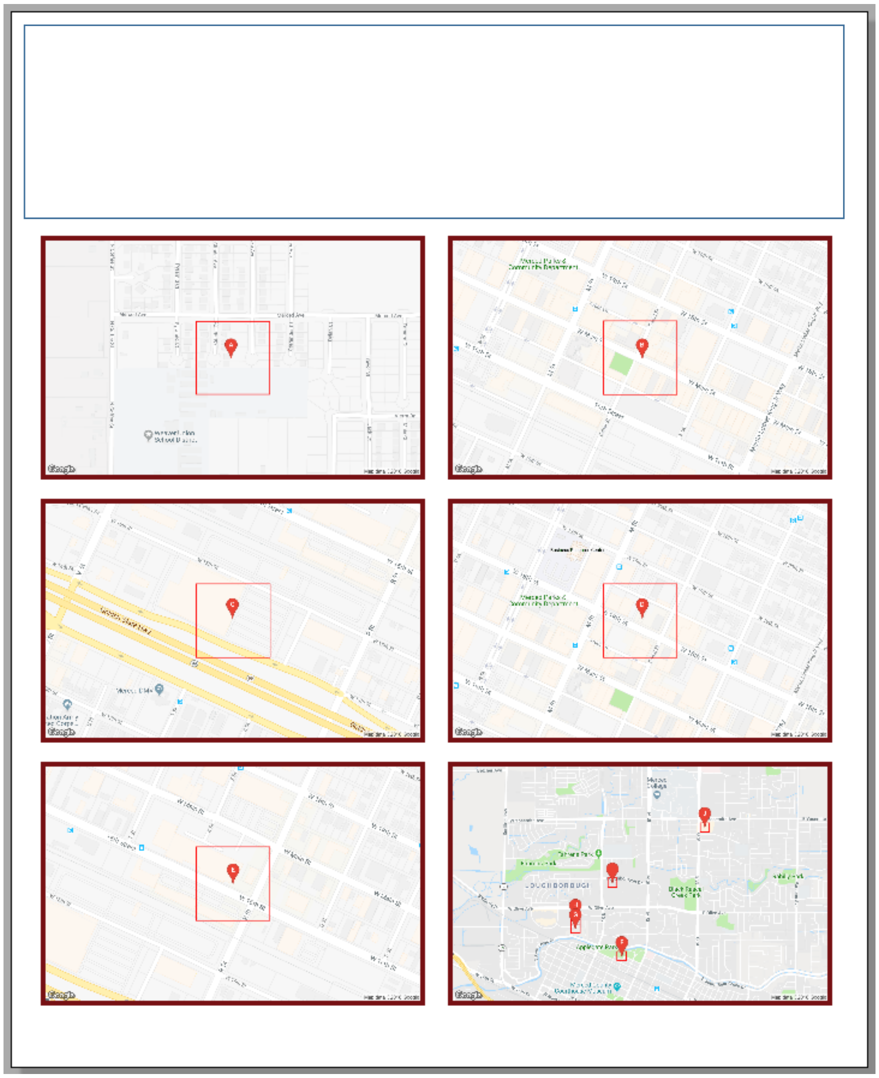

Notes: This is a sample to illustrate what law enforcement sees; it does not necessarily come from the jurisdiction that I study in this paper. PredPol boxes are shown on maps and also identified by the approximate intersection (blocked out here). The date, shift, and crimes types predicted are blocked out.

Figure C2: Map of PredPol predictive policing boxes from the PredPol website

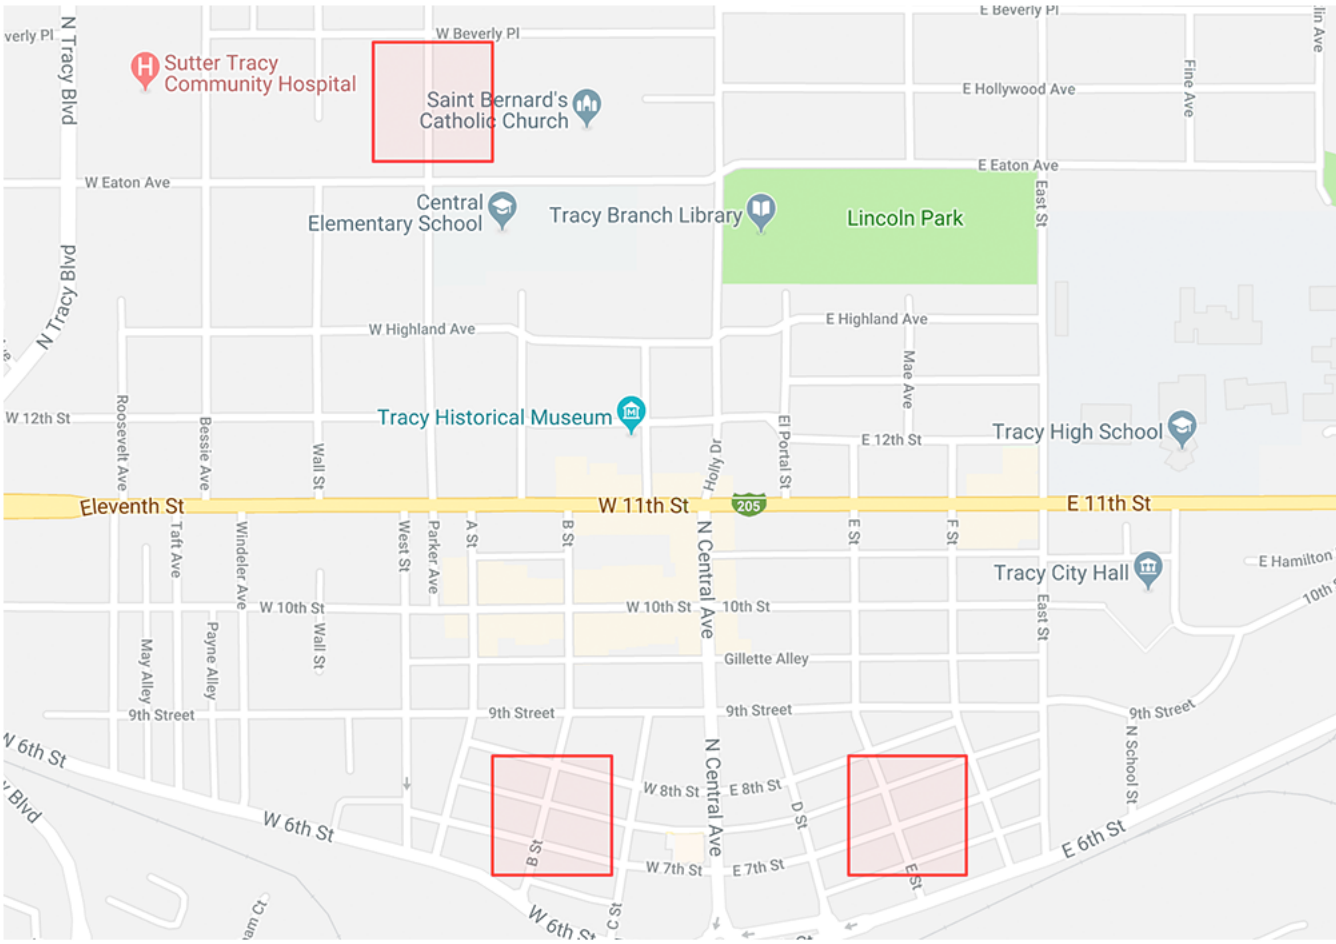

Notes: This is an example of a map of PredPol boxes from the PredPol website. PredPol boxes are the red boxes. This figure illustrates the size of PredPol boxes, which are 500 ft–by–500 ft and span 1–3 blocks. The map does not necessarily correspond to the jurisdiction that I study in this paper

Figure C3: Distribution of box victim racial composition

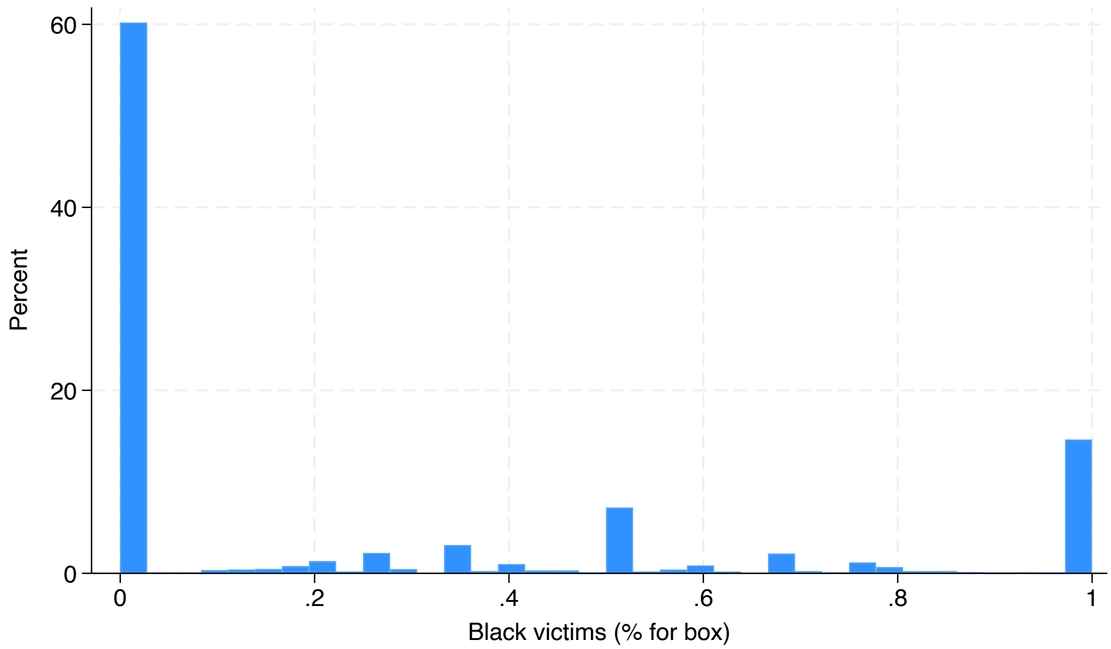

Notes: This figure plots the percent of boxes with the specified percentages of victims identified as Black in crime incident reports for index crime incidents that occur over three years.

Figure C4: Illustration of active predictive policing box quasi-experimental research design

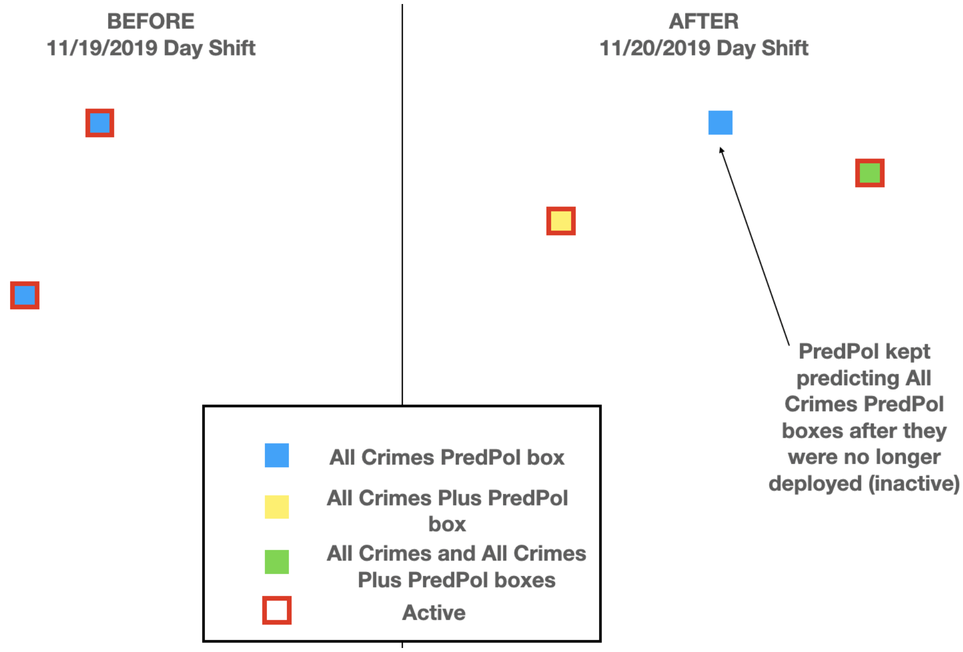

Notes: Prior to 11/20/2019, law enforcement received the All Crimes PredPol boxes in blue and was instructed to patrol in these boxes. After 11/20/2019, PredPol continued to generate All Crimes PredPol boxes even when they were no longer delivered to law enforcement; these serve as the control group in my research design. After 11/20/2019, law enforcement began to receive the All Crimes-Plus PredPol boxes in yellow and was instructed to patrol in these boxes. There is also overlap between the All Crimes and All Crimes-Plus PredPol boxes after the change since there is overlap in the crime types used for prediction. My research design compares the outcomes in the All Crimes PredPol boxes before the change with the outcomes in the All Crimes PredPol boxes after the change (which were not delivered and therefore law enforcement was not instructed to patrol there), accounting for the All Crimes-Plus PredPol boxes that *were* delivered to law enforcement after the change (where officers were instructed to patrol) and box fixed effects.

Figure C5: Race-specific effects of algorithm-induced police presence on arrests in traffic incidents: Heterogeneity by neighborhood racial composition

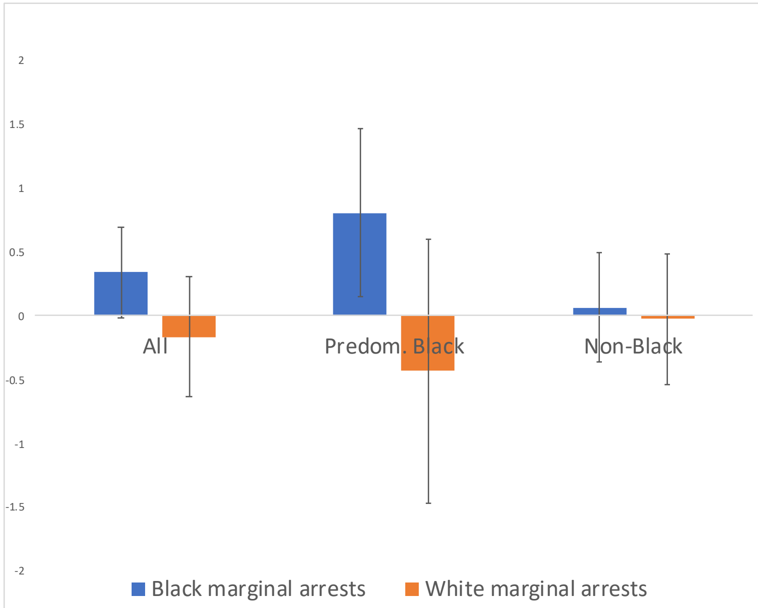

Notes: This figure plots estimates of race-specific effects of algorithm-induced police presence—Black marginal arrests (β_b) and white marginal arrests (β_w)—on arrests in traffic incidents. The left set of estimates are estimated with equation 9 and the sample of all boxes. The middle and right sets of estimates are estimated with equation 9 in the sample of predominantly Black boxes and non-Black boxes, respectively. Boxes are classified as predominantly Black if more than 35% of the victims of index crimes over a three-year period are identified as Black in incident reports. Boxes are non-Black boxes if 35% or fewer of the victims of index crimes over a three-year period are identified as Black in incident reports.

Figure C6: Race-specific effects of algorithm-induced police presence on arrests for violent crime incidents: Heterogeneity by neighborhood racial composition

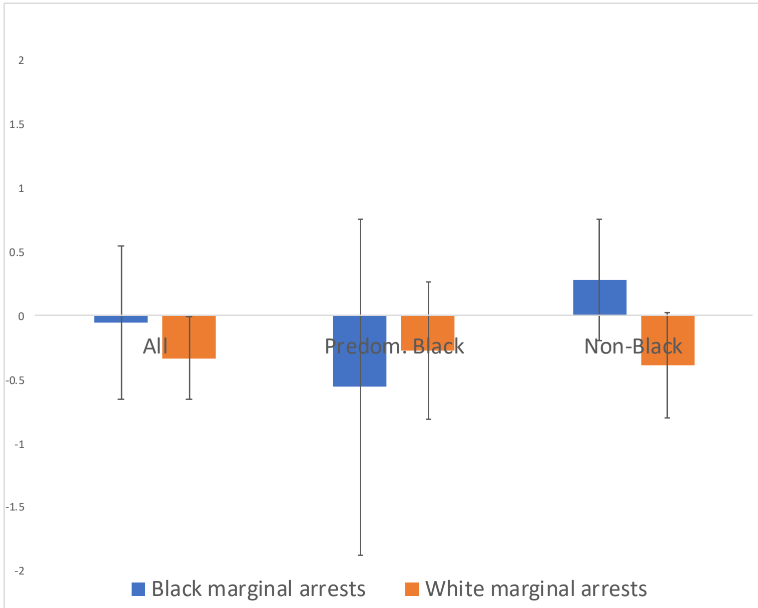

Notes: This figure plots estimates of race-specific effects of algorithm-induced police presence—Black marginal arrests (β_b) and white marginal arrests (β_w)—on arrests for violent crime incidents. The left set of estimates are estimated with equation 9 and the sample of all boxes. The middle and right sets of estimates are estimated with equation 9 in the sample of predominantly Black boxes and non-Black boxes, respectively. Boxes are classified as predominantly Black if more than 35% of the victims of index crimes over a three-year period are identified as Black in incident reports. Boxes are non-Black if 35% or fewer of the victims of index crimes over a three-year period are identified as Black in incident reports.

Figure C7: Race-specific effects of algorithm-induced police presence on arrests for violent crime incidents (Extended box): Heterogeneity by neighborhood racial composition

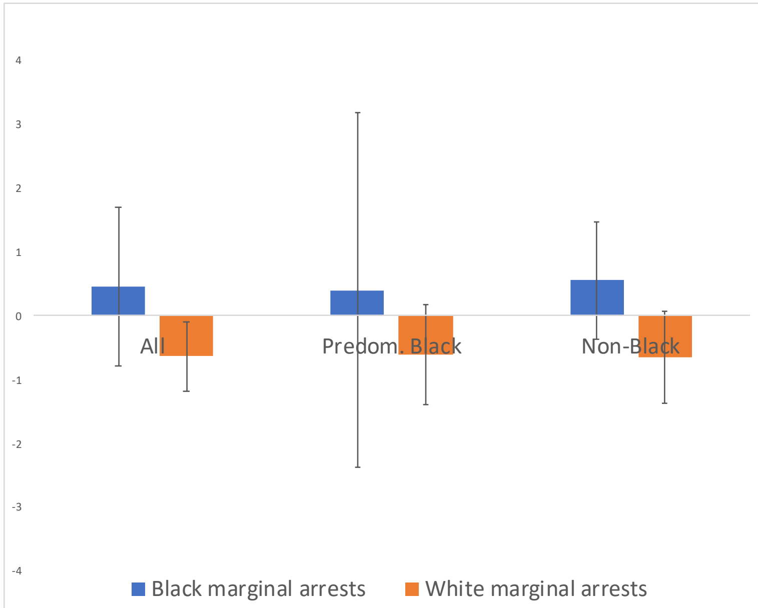

Notes: This figure plots estimates of race-specific effects of algorithm-induced police presence—Black marginal arrests (β_b) and white marginal arrests (β_w)—on arrests for violent crime incidents. The left set of estimates are estimated with equation 9 and the sample of all boxes. The middle and right sets of estimates are estimated with equation 9 in the sample of predominantly Black boxes and non-Black boxes, respectively. Boxes are classified as predominantly Black if more than 35% of the victims of index crimes over a three-year period are identified as Black in incident reports. Boxes are non-Black if 35% or fewer of the victims of index crimes over a three-year period are identified as Black in incident reports.