TEMPLE UNIVERSITY
Center for Security and Crime Science

THE PHILADELPHIA PREDICTIVE POLICING EXPERIMENT

SUMMARY OF EXPERIMENT AND FINDINGS

Jerry H. Ratcliffe, Ralph B. Taylor, Amber P. Askey, Kevin Thomas,
John Grasso, Kevin J. Bethel, Ryan Fisher, Josh Koehnlein

Project background

The Philadelphia Predictive Policing Experiment was a two-year collaboration between Temple University's Center for Security and Crime Science, housed in the Department of Criminal Justice, and the Philadelphia Police Department. This NIJ-funded research project was the first place-based, randomized experiment to study the impact of different patrol strategies on violent and property crime in predicted crime areas. The experiment's goal was to learn whether different operationally-realistic police responses to crime forecasts, estimated by a predictive policing software program, would reduce crime. The experiment was directed by Philadelphia Police Deputy Commissioner Kevin Bethel (retd.) and coordinated by the Director of Research and Planning at the Philadelphia Police Department, Kevin Thomas. The evaluation team was led by Dr. Jerry Ratcliffe and Dr. Ralph Taylor.

The experimental design

We tested two operational questions about police patrol. If police are able to dedicate a car to predicted crime areas, would it be better to use a marked car or an unmarked car? The visible car would emphasize deterrence and prevention. The unmarked (plain-clothes) car would allow officers to conduct surveillance and approach crime undetected in a more intelligence-led and apprehension mode. We also examined if it was sufficient to just tell officers on roll call where the predicted grids were each day without having a car dedicated to the task.

Sherman (2013) suggested that "the bare minimum, rock-bottom standard for EBP (evidence based policing) is this: a comparison group is essential to every test to be included as 'evidence' that a police decision causes a certain result". We identified the treatment and comparison areas (control districts) by block randomizing the police districts. Randomization was based on socio-economic status and a crime harm index so districts were comparable across different experimental conditions.

A property crime phase ran for 90 days from June 1, 2015 through August 25, 2015. A break was scheduled to allow the police department to prepare for, and police, the visit of Pope Francis to the City of Philadelphia. Subsequently, the violent crime phase ran for 92 days from November, 1, 2015 through January 31, 2016.

Designed to target the time of the day with most crime problems, in each district three property crime grid predictions (500 feet by 500 feet) were active from 8am to 4pm, and three violent crime grid predictions were active in each district for an eight hour period from 6pm to 2am. After removing the airport and the lowest crime district, the remaining 20 districts were randomly assigned to one of four experimental conditions:

- **Awareness** districts informed officers of the predicted target areas for that shift. Officers were asked to focus on those areas when they were able, but no cars were dedicated to the grids.
- **Marked** car districts built upon the awareness model by dedicating a single marked vehicle to patrol of the predicted crime areas for the entirety of the shift.
- **Unmarked** car districts were similar to the implementation for the marked districts except for the use of plain-clothes officers and unmarked police vehicles instead of uniformed resources.
- **Control** districts were districts where police personnel did not have access to the crime prediction software, so they maintained a standard patrol strategy.

In the car districts, officers were instructed to remain in the grids as much as possible, but were not instructed about how long, how frequently, using what activities, or in what order they should be patrolled.

Results

We examined both predicted mission grid cells and the cells immediately surrounding them. This latter overall *treatment area* is a more realistic test because officers had to drive around nearby streets to get to the mission grids. We found that these treatment areas, over a week, would typically experience one or more property crimes about 60 percent of the time. The marked car treatment showed substantial benefits for property crime (31% reduction in expected crime count), as well as temporal diffusion of benefits to the subsequent 8-hour period (40% reduction in expected crime count). No other intervention demonstrated meaningful crime reduction. Some violent crime results ran contrary to expectations, but this happened in a context of extremely low crime counts in predicted areas. The reductions identified were probably not substantial enough to affect city or district-wide property crime. To extrapolate, if each of the 21 geographic districts dedicated a marked car to three grids for an 8-hour shift each day, we estimate a reduction in 333 Part I property crimes per year.

There were signs of a temporal diffusion of benefits. In the eight hours following the property crime patrols, the marked car districts were associated with a reduction in crime compared to the control areas. The chance of a 4pm to midnight shift (after the predictive patrols) suffering any property crime was 40% lower after a marked car had been on patrol (but again, actual numbers were low). Unfortunately, the relative rarity of crime on an hour-by-hour basis in such limited geographic areas hindered the ability to make confident inferences about any crime reductions regarding the experimental conditions. In other words, while the percentages were substantial, the results were

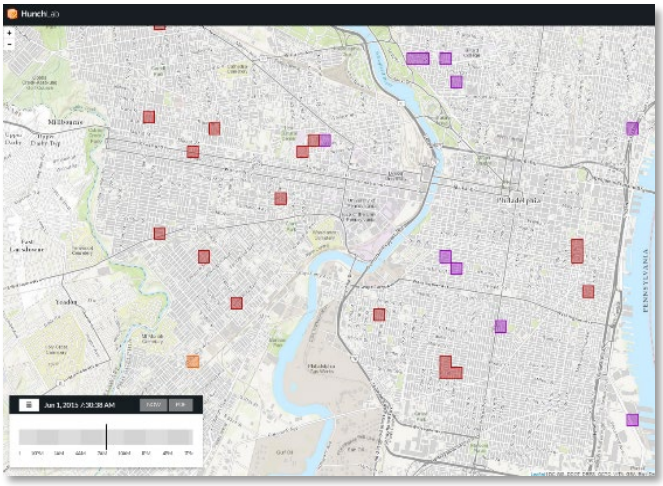

not statistically significant due to floor effects. As a result, there were no discernable crime reduction benefits associated with the violent phase of the experiment. And again, we also did not find any benefits with the property crime awareness or unmarked car interventions.

The experiment used the HunchLab predictive policing software (now called "ShotSpotter Missions"). The software predicted twice as much crime as would be expected if crime were uniformly distributed, even though the software was constrained from making its best predictions by the experimental design. The experiment did reveal the technical challenges in predicting crime in a limited number of small 500' x 500' grids.

Implementation

We conducted just over 100 ride-alongs with officers in marked and unmarked vehicles. We gathered 101 individual field note entries completed through both phases of the experiment, with 79 conducted with the property crime phase, and the remaining 22 with the violent crime phase. Officers patrolled the treatment areas to varying levels throughout the shift, with officers getting to the treatment areas earlier for the 8am to 4pm property crime phases, but officers patrolling treatment areas more extensively later in the shift for the violent crime phase (6pm to 2am). At least 50 percent saturation of treatment areas was achieved for 3.5 and 3.75 hours, respectively, for property and violent crime phases.

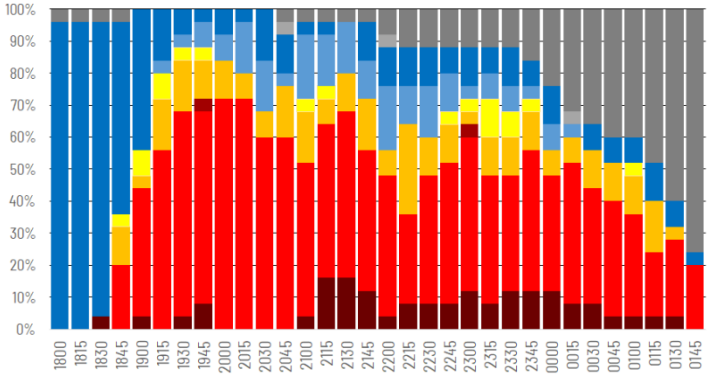

Further information

For additional and current information, visit the project website at **jratcliffe.net/philadelphia-predictive-policing-ex.**

Most of these results are reported in greater detail in Ratcliffe, JH, Taylor, RB, Askey, AP, Thomas, K, Grasso, J, Bethel, KJ, Fisher, R, Koehnlein, J (in press) **The Philadelphia Predictive Policing Experiment**. *Journal of Experimental Criminology*.

You can also find qualitative findings from the ride-alongs in Ratcliffe, JH, Taylor, RB, & Fisher, R (in press) **Conflicts and congruencies between predictive policing and the patrol officer's craft**. *Policing and Society*.

Suggested citation for this pdf: Ratcliffe, JH, Taylor, RB, Askey, AP, Thomas, K, Grasso, J, Bethel, KJ, Fisher, R, Koehnlein, J (2019) *The Philadelphia Predictive Policing Experiment: Summary of experiment and findings.* November 2019. Online at jratcliffe.net/philadelphia-predictive-policing-ex.