UC Irvine

UC Irvine Previously Published Works

Title

Randomized Controlled Field Trials of Predictive Policing

Permalink

<https://escholarship.org/uc/item/1br4975j>

Journal

Journal of the American Statistical Association, 110(512)

ISSN

0162-1459

Authors

Mohler, GO
Short, MB
Malinowski, S
et al.

Publication Date

2015-10-02

DOI

10.1080/01621459.2015.1077710

Peer reviewed

Randomized controlled field trials of predictive policing

G. O. Mohler^{*} M. B. Short[†] Sean Malinowski[‡] Mark Johnson[§]
G. E. Tita[¶] Andrea L. Bertozzi^{||} P. J. Brantingham^{**}

Abstract

The concentration of police resources in stable crime hotspots has proven effective in reducing crime, but the extent to which police can disrupt dynamically changing crime hotspots is unknown. Police must be able to anticipate the future location of dynamic hotspots to disrupt them. Here we report results of two randomized controlled trials of near real-time Epidemic Type Aftershock Sequence (ETAS) crime forecasting, one trial within three divisions of the Los Angeles Police Department and the other trial within two divisions of the Kent Police Department (UK). We investigate the extent to which i) ETAS models of short term crime risk outperform existing best practice of hotspot maps produced by dedicated crime analysts, ii) police officers in the field can dynamically patrol predicted hotspots given limited resources, and iii) crime can be reduced by predictive policing algorithms under realistic law enforcement resource constraints. While previous hotspot policing experiments fix treatment and control hotspots throughout the experimental period, we use a novel experimental design to

^{*}Department of Mathematics and Computer Science, Santa Clara University

[†]School of Mathematics, Georgia Institute of Technology

[‡]Los Angeles Police Department

[§]Kent Police Service

[¶]Department of Criminology, University of California, Irvine

^{||}Department of Mathematics, University of California, Los Angeles

^{**}Department of Anthropology, University of California, Los Angeles

allow treatment and control hotspots to change dynamically over the course of the experiment. Our results show that ETAS models predict 1.4-2.2 times as much crime compared to a dedicated crime analyst using existing criminal intelligence and hotspot mapping practice. Police patrols using ETAS forecasts led to a average 7.4% reduction in crime volume as a function of patrol time, whereas patrols based upon analyst predictions showed no significant effect. Dynamic police patrol in response to ETAS crime forecasts can disrupt opportunities for crime and lead to real crime reductions.

1 Introduction

Crime events arise out of interactions between local, place-based environmental conditions [5, 53] and the situational decision-making of offenders and victims [33, 24]. In theory, police patrol can prevent crime on a day-to-day basis by altering or disrupting the environmental conditions suitable for crime [42, 54]. However, experimental studies attempting to measure the effectiveness of different police patrol strategies show mixed results [12, 3]. Random patrol has a negligible impact [25] because the risk of crime is not uniformly distributed in space or time [43, 20]. Random patrol therefore allocates resources to locations that have little or no associated crime risk. Hotspot policing, by contrast, concentrates overwhelming resources in direct response to non-uniform crime patterns [44], leading to crime suppression not only at deployment locations, but also over a surrounding region through a diffusion of benefits [17]. Displacement of offenders appears to be incomplete in both theory [45] and practice [55, 40], indicating that hotspot policing produces a net reduction in crime.

However, the use of overwhelming resources obscures the relationship between patrol time and effect on crime. For example, in the Philadelphia foot patrol experiment [40], hotspots were patrolled for 16 hours a day, 5 days a week, by new recruits emerging from the police academy. Treatment hotspot areas exhibited 23% fewer crimes than control areas, but it is

not known whether the same result could have been achieved at more modest deployment levels. At a more realistic scale, Sherman and Weisburd [44] identified small clusters of addresses and sought targeted patrol in treatment clusters of 3 hours per day. While even this relatively low level of dosage was difficult to achieve, they did observe significant effects on both calls for service and crime and disorder in treatment hotspots. Given high call-to-service volumes and shrinking police budgets, 1-2 hours per day of patrol in determined hotspots is likely a more realistic number for many police agencies.

While studies such as the foot patrol experiment show that the rate at which patrol officers deter crime increases as the amount of patrol time increases, the goal of predictive policing is to increase the rate of crime deterrence under fixed resources by choosing those hotspots with the highest projected crime rates as the targets for patrol. In the Philadelphia foot patrol experiment, 120 hotspots were selected by ranking all potential hotspots according to the crime volume over the previous three years, weighting previous years 1, 2, and 3 by 1.0, 0.5, and 0.25, respectively. Then each pair of hotspots in the sequential list were randomly partitioned to either treatment or control, yielding 60 treatment hotspots that were fixed throughout the experiment. If a more accurate ranking method were used to dynamically select 60 different hotspots each day, then our hypothesis is that a greater net reduction in crime would be possible under the same patrol resource constraints.

A number of algorithmic methods have been proposed for estimating crime hotspot risk. Multivariate models of crime hotspots include fixed environmental variables such as demographics [51], income levels [31], and distance to crime attractors [51, 31, 27] to augment variables defined by crime incident volume. Risk factors that vary in space but not time help reduce variance and are effective for estimating long-term, chronic hotspots. Alternatively, short-term hotspots are estimated in dynamic hotspot mapping models of repeat victimization [2, 9, 23, 22, 13], where kernel density estimation is applied over a sliding window, typically on the order of several weeks, across the experimental period. In this paper we

use an Epidemic-type aftershock sequence model (ETAS) [36, 35] that estimates the risk associated with both long-term hotspots and short-term models of near-repeat risk. Analogous to ETAS models of seismic activity [32], a stationary point process representing fixed environmental risk is estimated via an Expectation-Maximization algorithm simultaneously with a short-term triggering kernel reflecting short-term, dynamic risk.

We conduct two experiments using the ETAS model, one in three divisions of the Los Angeles Police Department and one in two divisions of the Kent (UK) Police Department. In the experiments, the ETAS algorithm (treatment) was put head-to-head with hotspots maps produced each day and shift by dedicated crime analysts (control). We used silent tests where predictions were not deployed to the field in order to evaluate the predictive accuracy of the treatment method compared with control. In the three divisions within Los Angeles, we additionally conducted single-blind field trials where officers used either treatment or control predictions each day as the focus of their patrols. Because hotspot locations dynamically changed each day, we used an experimental design where days were randomly assigned to treatment or control. Thus each day all officers within an entire division received either ETAS predictions or analyst predictions that were identical in appearance save for the hotspot locations. Our results show that ETAS models predict 1.4-2.2 times as much crime compared to a dedicated crime analyst using existing criminal intelligence and hotspot mapping practice. In the three divisions in Los Angeles, police patrols using ETAS forecasts led to an average 7.4% reduction in crime volume as a function of patrol time, whereas patrols based upon analyst predictions showed less than half of the treatment effect at a level that was not statistically significant.

The outline of the paper is as follows. In Section 2 we provide details on our methodology. In Section 2.1 we outline our experimental design. In Section 2.2 we provide details on the ETAS algorithm and the cloud-based software architecture used in the experiment. In Section 2.3 we provide details on the methodology used by the crime analysts to produce

the control hotspot maps. In Section 2.4 we provide details on the methodology used by the patrol officers in the field. In Section 2.5 we describe the data utilized by the ETAS algorithm and the analyst to produce hotspot maps. In Section 3 we present our findings. Section 3.1 presents results of the silent tests to evaluate model accuracy. In Section 3.2 we analyze police patrol activity in the field tests. In Section 3.3 we evaluate the impact of treatment and control predictions on crime rates in Los Angeles. Section 4 concludes with a discussion of the study in the broader context of crime forecasting and predictive policing.

2 Methodology

2.1 Design of predictive policing field experiments

Controlled experiments in social policy settings face unique constraints [18, 19]. The ideal counterfactual for a policing experiment would be to compare how a target subject would have responded in the absence of policing (i.e., a pure placebo). However, a placebo controlled design that would remove policing from the equation is not feasible for practical and ethical reasons. Rather, policing experiments are structured similarly to drug controlled trials [28], where a target policing strategy or tactic, designated as treatment, is compared with existing policing practice, designated as control. Randomization of subjects among treatment and control groups is the primary mechanism whereby a valid counterfactual is constructed. Subjects randomly assigned to a treatment group are exchangeable with those assigned to a control group. Thus any effect of treatment would have been observed among control subjects had they been otherwise randomly assigned to the treatment condition. Conversely, no effect would have been observed among treatment subjects had they been randomly assigned to the control condition. Exchangeability is also essential in controlling for the vast number of unobserved factors. Indeed, it is both theoretically and practically impossible to observe

and restrict all aspects of subject behavior. Randomization helps to ensure that treatment and control subjects are exchangeable with respect to these random effects.

Hotspot policing experiments have typically followed a partially random block design [4, 40]. For example, Braga and Bond [4] run a randomized complete block hotspot policing study that is of similar scale to the study here (individual hotspots on the order of 0.0115 sq miles, only 30% larger than the predictive policing boxes). Thirty-four hotspots were matched in 17 homogenous pairs and then randomly assigned to control and treatment conditions. However, their study is focused on the effects of targeting ‘shallow’ problem solving at identified hotspots with routine patrol occurring in control hotspots. We are concerned with experimental comparison of different methods for identifying hotspots with equivalent policing tactics used in the different types of hotspots.

A different experimental design is required for evaluating predictive policing methods because the locations of hotspots are no longer fixed throughout the experiment. Instead, hotspot locations change based upon changes in estimated risk from the predictive model employed. Rather than dividing patrol areas into partially matched blocks, in our experimental design patrol “mission maps” were generated independently by a crime analyst (control) and ETAS algorithm (treatment) each day. Mission maps were identical in outward appearance, save for the exact placement of prediction boxes. Only after each set of predictions was generated was it randomly determined, using a Bernoulli random number generator, whether the control or treatment condition would be deployed to the field for the next 24 hour period (6AM-6AM). Importantly, command staff, supervisors, and patrol officers in each of the deployment areas were not aware of the distinction between treatment and control mission maps. Thus, days were randomly allocated to treatment or control and are considered exchangeable. Because analyst and ETAS prediction boxes are not constrained to be in the same place each day, but rather build up areas of exposure dynamically from day to day, it is feasible to measure the impact on crime at the aggregate scale of policing divisions. This

is a unique feature of our experiment as previous hotspot policing studies by design only measure impact within the small area comprising the treatment and control hotspots.

Of course, choice among different experimental designs often involves tradeoffs. A random block design that compares treatment and control regions [44, 4, 40] may ensure that treatment and control days are independent of one another, where there is no such guarantee for our design (but see further remarks regarding this in the Discussion). But this choice of a random block design comes at a cost whereby treatment and control regions cannot be guaranteed to be exchangeable. Indeed, the police officers assigned to and crime problems of one region are never strictly equivalent to those of another region no matter how much statistical matching one does. In keeping with the repeated measures type of experimental design used here, we are able to establish greater control over the type of confounds inherent to random block designs. Specifically, in our experiment the police officers involved are the same on both control and treatment days and they confront the same crime environment. We thus explore an experimental protocol with different strengths and weaknesses than those of random block design.

2.1.1 Silent tests

We conducted silent experiments across two divisions of the Kent (UK) Police Department, Maidstone and Sevenoaks, from Jan 14, 2013 to April 7, 2013, and the Southwest Division of the Los Angeles Police Department (LAPD) from May 16, 2012 to January 10, 2013. Silent tests of predictive accuracy were conducted to control for biases that would be introduced by directed patrol in response to predictions. Each day both the ETAS algorithm and a crime analyst (one in Kent and one in Southwest) selected twenty 150 x 150 m prediction boxes for each of two 12 hour shifts in each division. Crime analysts in both settings followed their respective existing best practice in selecting predictions (see Section 2.3 below). The task set before each analyst was to place predictions to maximize the number of crimes captured

over the subsequent twenty four hour period. In Kent, predictions were not deployed to the field. We therefore compare head-to-head accuracy on each day of the silent test in Kent. In Southwest, it was randomly determined each day of the experiment whether to give ETAS or analyst predictions to the officers for directed patrol. We therefore compare the predictive accuracy of ETAS on days that ETAS predictions were not deployed to the field with the predictive accuracy of analyst predictions on days that analyst predictions were not deployed to the field. We discarded days on which the analyst was absent, yielding a total of 58 silent test days in Kent and 234 silent test days (117 ETAS and 117 analyst) in Southwest. To maintain independence of the ETAS and analyst box selection processes, mutual exclusivity of predictions was not enforced (see Section 3.3 below). In Maidstone and Sevenoaks Divisions in Kent, 13.5% and 14.1% of control and treatment boxes overlapped, respectively. In Southwest Division in Los Angeles, 9.1% of boxes overlapped.

2.1.2 Field tests

To test whether predictive policing can impact crime, we conducted single-blind, randomized controlled field trials across three LAPD Divisions: Foothill from November 7, 2011 to April 27, 2012; North Hollywood from March 31, 2012 to September 14, 2012; and Southwest from May 16, 2012 to January 10, 2013. Following the finalization of daily mission maps it was randomly determined whether the ETAS or analyst missions would be deployed to the field. LAPD command staff, supervisors, and patrol officers were not aware of the distinction between control and treatment conditions, consistent with a single-blind experimental design [28]. Patrol officers were directed to use available time to “get in the box” and police what they saw. Thus, control and treatment conditions contrast directed patrol patterns in space and time, not differences in field tactics (see Section 2.4 below for further details on police activity). We tracked the amount of available time spent in prediction boxes using the existing in-car call logging system. As with silent accuracy testing, field-deployed missions

consisted of twenty 150 x 150 m prediction boxes for each twelve hour shift per division. Prediction boxes were therefore free to vary in their locations each shift. We discarded days the analyst was absent, yielding samples of 124, 152, and 234 test days in each division, respectively. Since each day was randomly assigned to control or treatment, variation in both policing tactics and patrol time was independent of experimental condition, providing an opportunity to measure the impact of different patrol levels [29, 46]. A total of 62 control days and 62 treatment days in Foothill, 82 control days and 70 treatment days in North Hollywood, and 117 control and 117 treatment days in Southwest were randomly assigned during the trial (Total = 510 days).

2.2 Epidemic-Type Aftershock Sequence Model for Crime Prediction

Building on a foundation of reaction-diffusion models of crime [45], we treat the dynamic occurrence of crime as a continuous time, discrete space epidemic-type aftershock sequence point process [36, 35, 32]. Policing areas were first discretized into 150 x 150 m square boxes. The conditional intensity, or probabilistic rate $\lambda_n(t)$ of events in box n at time t was determined by,

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

where t_n^i are the times of events in box n in the history of the process. The ETAS model has two components, one modeling place-based environmental conditions that are constant in time and the other modeling dynamic changes in risk. Rather than modeling fixed environmental characteristics of a hotspot explicitly using census data or locations of crime attractors, long term hotspots are estimated from the events themselves. In particular, the background rate μ is a nonparametric histogram estimate of a stationary Poisson process

[32]. If over the past 365 days a grid cell has a high crime volume, the estimate of μ will be large in that grid cell. The size of the grid cells on which μ is defined can be estimated by Maximum Likelihood and in general the optimum size of the grid cell will decrease with increasing data. However, for a fixed area flagged for patrol, a greater number of small hotspots are more difficult to patrol than a small number of large hotspots. The 150m x 150m hotspots were chosen in this study to be the size of a city block in Foothill and were then held constant across all of the experimental regions. The number of days of data used as input to the ETAS model, 365 days, was also chosen subjectively, though is consistent with other hotspot policing studies that use 1-2 years of data to select hotspots.

The second component of the ETAS model is the triggering kernel $\theta\omega e^{-\omega t}$ that models “near-repeat” or “contagion” effects in crime data. The exponential decay causes grid cells containing recent crime events to have a higher intensity than grid cells with fewer recent events and the same background rate. The main difference between the ETAS model and prospective hotspot maps [2] that model near-repeat effects is the introduction of the background rate μ . Whereas prospective hotspot maps only take into account short term hotspot dynamics, the ETAS model estimates both long term and short term hotspots and systematically estimates the relative contribution to risk of each via Expectation-Maximization [36, 35]. Given an initial guess for the parameters θ , μ , and ω , the EM algorithm is applied iteratively until convergence by alternating between the following two steps:

E-step

$$p_n^{ij} = \frac{\theta\omega e^{-\omega(t_n^j - t_n^i)}}{\lambda_n(t_n^j)}, \quad (2)$$

$$p_n^j = \frac{\mu_n}{\lambda_n(t_n^j)}, \quad (3)$$

M-step

$$\omega = \frac{\sum_n \sum_{i < j} p_n^{ij}}{\sum_n \sum_{i < j} p_n^{ij}(t_n^j - t_n^i)}, \quad (4)$$

$$\theta = \frac{\sum_n \sum_{i < j} p_n^{ij}}{\sum_n \sum_j 1}, \quad (5)$$

$$\mu = \frac{\sum_n \sum_j p_n^j}{T}, \quad (6)$$

where T is the length of the time window of observation.

The EM algorithm can be intuitively understood by viewing the ETAS model as a branching process [36]. First generation events occur according to a Poisson process with constant rate μ . Events (from all generations) each give birth to N direct offspring events, where N is a Poisson random variable with parameter θ . As events occur, the rate of crime increases locally in space, leading to a contagious sequence of “aftershock” crimes [36] that eventually dies out on its own, or is interrupted by police intervention; the former occurs naturally so long as $\theta < 1$, while the latter is unaccounted for by the model. In the E-step, the probability that event j is a direct offspring of event i is estimated, along with the probability that the event was generated by the Poisson process μ . Given the probabilistic estimate of the branching structure, the complete data log-likelihood is then maximized in the M-step, providing an estimate of the model parameters. For a detailed treatment of the EM algorithm in the context of ETAS see [49] or [30] where the EM algorithm is shown to be equivalent to projected gradient ascent optimization on the log-likelihood.

An ETAS crime forecast consists of the top 20 boxes in rank order displayed on a map and available for directed police patrol. The algorithm was implemented as a fully automated cloud based software system (see Figure 1). For each new crime report of the target crime types entered into the police records management system, the date, time, crime type, and address were automatically sent to secure, cloud-based data center servers. Events were geocoded in the cloud servers using Google API and stored for later use. Each day at 4AM local time the ETAS algorithm was run and each 150m x 150m cell in each division and shift was assigned a probability score by the algorithm. The top 20 cells were then displayed on a

Google map that could be accessed by a user via a secure internet connection. The software also allowed the analyst on control days to place 20 boxes (identical in size and appearance to ETAS boxes) on the Google map instead of ETAS generated boxes. Reports were printed out at roll call and distributed to the officers.

Figure 1: The ETAS model was implemented as a fully automated cloud based machine learning system. Java software was installed on the police RMS server that encrypted and sent crime report data (address, date, time, crime type) once every hour to cloud-based servers where the data was geocoded using Google API and stored for later use. Once a day at 4am the ETAS parameters were re-estimated using the previous 365 days of crime data (up to the latest records) and boxes were ranked by intensity λ_n . Patrol commanders logged on via a web interface prior to each patrol shift and printed out a report that was then distributed to officers. The report contained a map of hotspot locations as well as a list of nearest cross streets of the hotspot locations.

The ETAS model was automated to produce treatment missions using the most recent

crime data available in the live records management system. Predictions were therefore generated in near real-time given natural lags in recorded crime data. Missions were set daily at 5AM for two sequential shifts from 6AM-6PM and 6PM-6AM. In Foothill, mission predictions were made for the 24 hour period starting at 4PM. The number of prediction boxes and prediction time windows in each policing division were identical to those generated by analysts.

2.3 Analyst Control Predictions

We compare predictions derived from ETAS to the best practices of trained crime analysts. Analysts were tasked to place a fixed number of 150 x 150 m boxes within their operational environment in specified time windows. The goal of box placement was to identify a small set of locations where the analysts expected crime to be most likely to occur. In Kent, one analyst was responsible for generating 80 prediction boxes each day, evenly divided among two policing divisions (Maidstone and Sevenoaks) and two shifts (7AM-7PM and 7PM to 7AM). In Los Angeles, three crime analysts acting independently were responsible for placing prediction boxes for Foothill, North Hollywood, and Southwest Divisions. In Foothill Division, 20 predictions were generated for each 24 hour period beginning at 4PM each day. In both North Hollywood and Southwest Divisions, 40 prediction boxes were generated each day evenly divided between two shifts (6AM-6PM and 6PM-6AM).

Crime analysts were free to use any information and methods at their disposal to generate predictions. Not surprisingly, therefore, the approaches taken to crime prediction differed substantially between Kent Police and Los Angeles Police Department analysts. Kent Police adhere to an intelligence-led policing approach [39] as specified by the UK National Intelligence Model [14]. The goal of the intelligence-led approach is to fuse detailed individual-level information on chronic or prolific offenders, who are responsible for a large fraction of the

crime, with information on recent criminal incidents, known high-risk issues, and priority locations (i.e., hotspots), and use this information to guide resource allocation. The intelligence-led approach is explicit in its view that criminal intelligence helps police avoid excessive responses to random events. When tasked with placing prediction boxes, the Kent analyst relied primarily on information about the activity anchor points (e.g., home and work locations) of chronic offenders and their known associates. Information about recent crime events and their distribution in space and time was evaluated as relevant in relation to criminal intelligence. For example, a residential burglary that occurs in proximity to a place frequented by a known prolific offender is much more important in driving prediction placement than a burglary occurring in a location with no known connection to a prolific offender. Analysts in England and Wales consider crime mapping and analysis in support of intelligence-led policing a core function, though only about half of analysts from police forces report that such activities occupy most of their time [52].

The emphasis in Los Angeles was reversed. The Los Angeles Police Department follows a COMPSTAT [50] policing model focused on the analysis of seven-day crime maps supplemented with ad hoc street-level intelligence. The LAPD analysts emphasized small clusters of recent crimes as signaling emerging problems in need of short-term response to interrupt the formation of crime hotspots. Box placement therefore reflected the colloquial mantra about COMPSTAT as being “cops-on-the-dots”. Here the analyst approach was “box-on-the-dots”. Hot spot mapping is in widespread use within US policing [8]. A substantial majority of police forces serving populations of 100,000 people or more make use of hot spot mapping for both evaluation of past performance and tactical resource allocation [15, 38, 41].

In both Kent and Los Angeles an easy-to-use web interface was developed to allow analysts to choose prediction locations by clicking on a Google Map. While the physical task of placing predictions could be completed in a matter of minutes, in practice analysts spent considerable time engaged in analysis of crime and intelligence data before committing to

predictions. In Kent, the analyst spent between 3-4 hours each day researching chronic offender information and ongoing crime problems before placing their 80 total prediction boxes in the operational environment. In Los Angeles, the three analysts each spent 1.5-2 hours per day inspecting crime maps and recent crime reports before placing their 40 total prediction boxes (20 in Foothill).

2.4 Policing Activity Under Experimental Conditions

Police officers and patrol teams develop preferences that govern where they like to patrol and the types of policing tactics they like to use. For example, certain teams may prefer policing misdemeanor crimes, others may prefer to interdict chronic violent or property offenders, while still others may prefer non-crime community interactions. It is reasonable to suppose that some of these policing tactics will have more of an impact on suppressing crime than others, though there is continuous, vigorous debate surrounding policing tactical effectiveness. We are agnostic about the benefits of different tactical choices of officers in the field. In any case, LAPD command staff in each of the three test divisions did not mandate specific tactical actions. Rather, patrol officers were encouraged to use available time to “get in the box” and use their discretion to select appropriate field tactics upon entering prediction locations. In this approach we follow closely [44] in that we neither restrict how police deal with individual hotspots, nor do we restrict the amount of time they engage the problem. Nevertheless, we were concerned with potential bias arising from any non-random distribution of police tactical preferences across control and treatment groups.

The experiment was designed to help guard against such biases. Since each day is randomly assigned to control or treatment condition at the start of the day, over time each experimental condition is expected to receive equal exposure to the range of policing tactics present among the patrol teams. For example, an officer or patrol team that prefers so-called

“broken windows” policing – enforcing against misdemeanor infractions with the notion that it will disrupt felonies by the same individuals – will be provided equal opportunity to engage those tactics on control and treatment days (both inside and outside of prediction boxes). Moreover, because the outward appearance of control and treatment missions was identical, except for the exact placement of boxes, there were no cues that would encourage patrol teams to systematically adjust their tactics to be different on control or treatment days.

An ideal test of successful randomization of officer or patrol team tactical preferences would involve collection of detailed activity logs. This approach was not feasible given the scale and duration of the present experiment. Rather, we used officer time on mission in control and treatment prediction boxes as our primary measure of officer activity. Officers generated a time-stamped call event on their in-car mobile data terminal upon entering and leaving any mission box. We refer to the total amount of time spent within prediction boxes as officer time on mission and each incident of an officer spending time within a box as a mission stop.

2.5 Crime Types and Predictive Accuracy

The experiment focused on burglary, car theft, burglary-theft from vehicle, criminal damage, violence against the person (including sexual offences) and robbery as target crime types in Kent. These crimes together comprise 56% of the crime volume recorded by Kent. Burglary, car theft, and burglary-theft from vehicle were the target crime types in Los Angeles, which together comprise 55% of the reported crime volume. All of these crime types are known to be underreported [37] and therefore comparisons are made between control and treatment conditions subject to the same underreporting constraint. Reported crimes were tabulated according to whether they occurred strictly inside or outside active control and treatment boxes. A crime occurring strictly inside an active prediction box is considered a successful

prediction. We do not consider here “near misses” or crimes that occurred within some small distance of an active prediction box. An active prediction was defined as a box generated before the start of a patrol shift and held constant in that location for the duration of the shift. Active predictions were replaced at the end of each shift by new active boxes. In the case of silent accuracy testing, active predictions were not made available to patrol officers. For deployed testing, patrol officers had access to maps of active prediction boxes during their patrol hours.

3 Results

3.1 Accuracy of ETAS vs Analyst Predictions

We first analyze predictive accuracy. In Table 1, we count the number of crimes occurring in active prediction boxes under silent experimental conditions. The analyst in Kent predicted 6.8% (Maidstone) and 4.0% (Sevenoaks) of crimes successfully compared to 9.8% and 6.8% by the ETAS model, a factor of 1.4 and 1.7 improvement in accuracy significant at the $P = 0.03$ and $P = 0.04$ level, respectively. In Southwest Division, over the 117 days in which control conditions were not deployed to the field, the analyst successfully predicted 2.1% of crimes. Over the 117 days in which treatment conditions were not deployed to the field, the ETAS model successfully predicted 4.7% of crimes. Treatment conditions thus yielded a predictive accuracy 2.2 times greater than control in the absence of police patrol effects ($P = .0012$).

The difference in the accuracy between analyst and ETAS predictions reflects advantages that algorithmic approaches have in characterizing dynamic spatio-temporal patterns[36, 35, 51]. In LAPD Southwest Division, where the emphasis is on crime mapping, analyst predictions closely track recent crimes across the operational environment (Figure 2A,B). This

	ETAS				Analyst					
	Success	Total	Rate	PAI	Success	Total	Rate	PAI	Boost	P-value
Maidstone	60	615	9.8%	85.2	42	615	6.8%	59.6	1.4	0.0314
Sevenoaks	39	576	6.8%	55.7	23	576	4.0%	32.8	1.7	0.0409
Southwest	46	986	4.7%	3.5	20	933	2.1%	1.6	2.2	0.0184

Table 1: **Successfully predicted crimes under non-deployed conditions.** Successful predictions are defined as crimes occurring strictly within an active 150 x 150 m box. Total crime in Maidstone and Sevenoaks, Kent includes all burglary, car theft, theft from vehicle, and criminal damage occurring during 12 hour shifts where both analyst and ETAS predictions were generated. Total crime in Southwest, Los Angeles is all burglary, car theft, and theft from vehicle occurring on corresponding days where treatment (ETAS) or control (analyst) predictions were not deployed to officers. PAI is the predictive accuracy index, an area-standardized measure of accuracy. Boost is the increase in predictive accuracy associated with treatment. P-value corresponds to a one-tailed, two sample proportions test for the difference in mean predictive accuracy between treatment and analyst. Predictions are only calculated on shifts with exactly 20 prediction boxes. Random predictions correspond to an accuracy of 1.0% in Southwest and 0.1% in Maidstone and Sevenoaks.

indicates a tendency to judge overnight crime as more salient than other criteria in predicting future crime. In Kent Maidstone Division, the emphasis is on recent information about known prolific or chronic offenders, who can be responsible for a large proportion of crime [26]. The intelligence-led approach used in Kent produced more clustering in predictions over time (Figure 2E), reflecting an explicit rationale to not chase crimes that cannot be spatially associated with known chronic offenders [14] (Figure 2D).

Deployment of prediction boxes to the field and delivery of policing dosage to those boxes is expected to suppress some fraction of crime in those locations. Predictive accuracy should therefore decline in response to directed patrol. Table 2 shows that crimes did occur in treatment and control boxes deployed to the field. In LAPD Southwest Division, where we have a silent test of accuracy for comparison, predictive accuracy falls from 4.7% to 3.9%, a decline of -17% overall. Importantly, predictive accuracy for analyst control conditions remain effectively unchanged from the silent test (2.1%) to deployed settings (2.2%), con-

	ETAS				Analyst					
	Success	Total	Rate	PAI	Success	Total	Rate	PAI	Boost	P-value
Foothill	22	346	6.4%	16.9	11	347	3.2%	8.4	2.0	0.0244
N. Hollywood	21	611	3.4%	4.9	12	732	1.6%	2.4	2.1	0.0170
Southwest	38	981	3.9%	2.9	21	936	2.2%	1.7	1.7	0.0194
Total	81	1938	4.2%	6.8	44	2015	2.2%	3.5	1.9	0.0002

Table 2: **Successfully predicted crimes under deployed conditions.** Successful predictions defined as crime occurring strictly within an active 150 x 150 m box. Total crime is all burglary, car theft, and theft from vehicle occurring on corresponding treatment (ETAS) or control (analyst) days. PAI is the predictive accuracy index, an area-standardized measure of accuracy. Boost is the increase in predictive accuracy associated with treatment. P-value corresponds to a one-tailed, two sample proportions test for the difference in mean predictive accuracy between treatment and analyst. Predictions are only calculated on shifts with exactly 20 prediction boxes. This precludes some shifts in Southwest in which the analyst created fewer than 20 boxes, explaining the discrepancy between analyst and ETAS total crime numbers in deployed versus nondeployed conditions.

sistent with the observation that policing dosage has limited effect under control conditions (see below). Table 2 also presents accuracies for treatment and control boxes under deployed conditions in LAPD Foothill and North Hollywood Divisions. Silent tests of accuracy were not conducted in these policing areas and therefore similar comparisons are not possible at present.

It is easy to predict a large fraction of crime if one designates a large area as being at risk for crime [9]. In the experimental design used here, by contrast, we only place a tiny fraction of the total land area under prediction at any one time. LAPD Foothill, North Hollywood, and Southwest Divisions total 119.48, 64.87, and 33.96 km^2 in area, respectively. Only 0.38%, 0.69%, and 1.32% of the total areas, respectively, were under prediction. In Kent, the fraction of land area under prediction was even smaller. Sevenoaks is 370 km^2 in area and prediction boxes covered only 0.12% of the total at any one time. Maidstone is 393 km^2 in areas and prediction boxes covered only 0.11% of the total land area at any one time. The fraction of crime predicted relative to the fraction of land area under prediction

is very large. In Maidstone, for example, 9.8% of crime is accurately predicted in just 0.11% of the land area, representing area-standardized predictive accuracy of 85.2.

3.2 Crime analyst methodology

The goal of this study was to compare a fully automated statistical algorithm (ETAS) for determining patrol hotspots to existing best practice in the Los Angeles and Kent police departments. Rather than attempting to control for analyst methodology or documenting analyst methodology through surveys, we perform a post-experiment analysis of the hotspots selected by the analysts during the trial. In Table 3, we compare analyst success rate with hotspot predictions corresponding to N-day hotspot maps. In Table 4, we compare analyst and hotspot prediction locations by calculating the number of identical hotspots for each day of the experiment aggregated over the experimental period. The hotspot maps are determined by ranking grid cells by the number of crimes occurring in the past N days. The analyst in Maidstone significantly outperforms 3 day hotspot maps in predicting crime on a shift-by-shift basis. The success rate of analysts in Sevenoaks and Southwest Division is statistically indistinguishable from 3 day hotspot maps. All three analysts are statistically indistinguishable from 7 day hotspot maps in predicting crime. By comparison, ETAS substantially outperforms both 3 day and 7 day hotspot maps. Consistent with the difference between ETAS and analysts, ETAS doubles the amount of crime predicted relative to hotspot maps with short measurement time windows.

In spite of the statistical similarity in performance in predicting crime, analyst predictions do deviate from strict hotspot mapping a fraction of the time. We observe in Table 4 that 13-40% of analyst hotspots are identical to either 3, 7, 14, 21, or 28 day hotspot maps, depending on the analyst and division. In Table 5, we observe that analysts also place a significant number of hotspots to account for near-repeat victimization [11, 48]. The percentage of

	3 day						7 day					
	Success	Total	Rate	PAI	Boost	P-value	Success	Total	Rate	PAI	Boost	P-value
Maidstone	18	615	2.9%	24.1	2.3	0.0007	38	615	6.2%	50.8	1.1	0.3219
Sevenoaks	17	576	3.0%	25.8	1.4	0.1671	20	576	3.5%	30.3	1.2	0.3205
Southwest	17	933	1.8%	1.4	1.2	0.3092	27	933	2.9%	2.2	0.7	0.8495

Table 3: **Successfully predicted crimes for N-day hotspot maps.** Successful predictions defined as crime occurring strictly within an active 150 x 150 m box. Total crime is all burglary, car theft, and theft from vehicle occurring on corresponding control (analyst) days. PAI is the predictive accuracy index, an area-standardized measure of accuracy. Boost is the increase in predictive accuracy associated with control compared with 3 day and 7 day hotspot maps. P-value corresponds to a one-tailed, two sample proportions test for the difference in mean predictive accuracy between treatment and analyst. Predictions are only calculated on shifts with exactly 20 prediction boxes.

analyst hotspot maps in agreement or adjacent to N-day hotspot maps ranges from 47% to 67%. There are several possible explanations for the remaining analyst hotspots not identical or adjacent to hotspots determined through hotspot mapping. One explanation is that analysts use a more complex function (compared to simple crime counts) when determining hotspot rankings. Another is that the remaining hotspots are determined through street level intelligence, for example location of parolees or known crime attractors. A third explanation is that constraints on working memory influence the selection of 20 hotspots for each shift. In Table 6, we observe some evidence for the third explanation as 7-20% of analyst hotspots contained no crime throughout the course of the experiments.

3.3 Prediction Similarity Through Time

A distinctive feature of the present experiment is that crime prediction locations are free to vary on a shift-by-shift basis. In spite of this freedom, a fraction of prediction locations occur repeatedly over the course of the experiment. Here we investigate the degree to which treatment and control predictions fall in the same locations on subsequent days. The method we use is analogous to a Hamming distance between prediction sets. We first choose a focal

Analyst	3-day	7-day	14-day	21-day	28-day	Remainder
Maidstone	.105	.067	.049	.033	.014	.732
Sevenoaks	.064	.073	.073	.043	.022	.725
Southwest	.088	.077	.054	.038	.028	.716
Foothill	.073	.045	.002	.004	.006	.869
N. Hollywood	.202	.154	.023	.013	.008	.600

Table 4: **Analyst hotspot selection and hotspot mapping.** Fraction of analyst hotspots in agreement with N-day hotspot maps. Those analyst hotspots agreeing with a N-day hotspot map are not included in the fraction for M-day hotspot maps when $M > N$. Numbers are rounded and may not sum exactly to one.

Analyst	3-day	7-day	14-day	21-day	28-day	Remainder
Maidstone	.310	.133	.100	.059	.028	.369
Sevenoaks	.239	.189	.146	.065	.031	.330
Southwest	.182	.120	.078	.051	.040	.530
Foothill	.265	.241	.041	.010	.012	.430
N. Hollywood	.307	.207	.045	.019	.0165	.405

Table 5: **Analyst hotspot selection and near-repeat victimization.** Fraction of analyst hotspots in agreement with N-day hotspot maps or adjacent to boxes selected by a N-day hotspot map. Those analyst hotspots agreeing with or adjacent to a N-day hotspot map are not included in the fraction for M-day hotspot maps when $M > N$. Numbers are rounded and may not sum exactly to one.

Maidstone	Sevenoaks	Southwest	Foothill	N. Hollywood
.135	.117	.077	.194	.139

Table 6: **Analyst hotspot selection and zero crime rate hotspots.** Fraction of analyst hotspots containing no crime throughout the course of the experimental period.

patrol shift and record the locations of all prediction boxes during that shift. We then look at corresponding patrol shifts on each subsequent day N and score the fraction of predictions falling in the same locations as boxes on the focal day. For example, if all of the twenty treatment predictions for the daytime shift on Feb 1, 2013 were repeated during the daytime shift on Feb 2, 2013, then the fraction of boxes repeated on day $N = 1$ is 1. If ten of the prediction from Feb 1, 2013 were repeated on Feb 11, then the fraction of boxes repeated on day $N = 10$ is 0.5. Note that static hotspot policing experiments establish treatment and control hotspots at the outset and then hold those constant for every measurement event [4, 44, 55]. The corresponding fraction of repeated predictions in such static cases is 1 for each and every day of the experiment.

The fraction of overlapping predictions in Southwest Division are presented in Figure 3 for each combination of treatment and control conditions. On average, 97% of treatment boxes on a focal day are repeated on the immediately subsequent day. Fifteen days later, 25% of the boxes have been replaced with boxes in unique locations, leaving 75% of the boxes in the same places as on the focal day. Thirty days after the initial prediction set, 62% of the initial boxes remain on average. By contrast, only 27% of control boxes on a focal day are repeated on the immediately subsequent control day. The fraction of repeated control boxes declines rapidly, but then rebounds near day 7 suggesting a tendency for analysts to make predictions by day of the week. Thirty days following the initial set of control predictions only 3% of the same prediction boxes remain on average. Control boxes and treatment boxes follow one another at even lower frequencies. Only about 9% of treatment boxes follow control boxes on an immediately subsequent day, and vice versa. The fraction decays to 4% and 6%, respectively, at thirty days. In general, treatment and control prediction locations are decoupled over time.

3.4 Patrol Activity in Los Angeles

We next consider measures of patrol officer activity in Los Angeles. Like specific field tactics, we expect time on mission to vary across patrol teams in part as a matter of preference, but also that control and treatment conditions should be equivalently exposed to this variation. Figure 4A confirms this expectation. Patrol teams spent equivalent amounts of time in control and treatment boxes (Wilcoxon rank-sums test $W = 7983$, $p = 0.8192$). This pattern was consistent regardless of whether the patrol team logged large or small numbers of minutes in prediction boxes.

The number of minutes officers spent in a prediction box during a single stop is exponentially distributed up through stops of 120 minutes in duration ($s(t) = 4360.1e^{-0.049t}$, $R^2 = 0.9943$) (see Figure 4B). The mean duration of a directed mission stop was 20.27 minutes (sdev = 25.11 minutes). Aggregating all observed mission stops per day, the average time on mission increased in Foothill and Southwest Division over the course of the experiment, and remained stable in North Hollywood (see Figure 5A, 5B, 5C). Increases in total time on mission coincided with the concentration of activity in a subset of prediction boxes. Treatment boxes on average received 57 minutes more exposure to policing per day than control boxes, but the difference is marginally not significant (mean control = 552 minutes, sdev = 376 minutes; mean treatment = 617 minutes, sdev = 430 minutes; $F = 3.288$, $P = 0.07036$). The day-to-day variation in time on mission is high, reflecting the directive to use available time to police prediction boxes. We exploit this variability in patrol time to measure the impact of control and treatment conditions on changes in crime volume in the following section.

3.5 Impact on Crime Rates in Los Angeles

We consider changes in crime volume for all three LAPD divisions as a whole, rather than exclusively within prediction boxes. This is a significant difference with other studies [44, 40,

4], which rely on random assignment of matched hot spots into treatment and control blocks. A random block design has a disadvantage that one can only detect crime rate changes between matched hotspots, not overall. Because we randomly assign days to treatment and control and, more importantly, allow prediction locations to change twice daily, we are able to measure the overall effects of predictive policing.

Regression models of the form $Y_{ij} = \mu_j + \beta_k X_{ij} + \epsilon_{ij}$ are constructed to assess the relationship between daily crime volume and patrol time on mission. Here Y_{ij} is the daily crime volume on day i in division j , X_{ij} is the police patrol time in minutes on day i in division j across all active prediction boxes, and ϵ_{ij} is the uncorrelated error. The coefficient μ_j is an estimate of the mean crime volume per day in division j in the absence of directed patrol, while β_k is an estimate of the impact of increasing patrol dosage under experimental condition k . The variance in daily crime volume is high for the three focal crime types, ranging from a minimum of 1 to a maximum 18 crimes per day. Nevertheless, crime rates are measurably lower with increasing patrol time in prediction boxes under some conditions. There is a statistically significant negative relationship between crime volume and time on mission for combined control and treatment conditions ($\beta_{\text{combined}} = -7.54 \cdot 10^{-4}$, $P = 0.0208$) (Table 6). Treatment conditions alone are also statistically significant ($\beta_{\text{treatment}} = -9.78 \cdot 10^{-4}$, $P = 0.0221$). Control conditions alone display no significant relationship between daily crime volume and patrol time ($\beta_{\text{control}} = -4.66 \cdot 10^{-4}$, $P = 0.364$).

Treatment conditions are expected to yield one less crime per 1000 minutes of police patrol time in ETAS predicted locations. The impact translates to 4.3 fewer crimes per week at mean patrol levels, or a reduction of 7.4% on a mean 58.17 crimes per division per week in the absence of patrol. Control conditions would yield less than half the reduction at equivalent patrol levels, assuming statistical significance could be achieved with a larger sample size. The factor 2.1 difference in crime reduction between control and treatment conditions is in agreement with the factor of 1.4 to 2.2 increase in accuracy of ETAS over

	μ_{Foothill}	$\mu_{\text{N. Hollywood}}$	$\mu_{\text{Southwest}}$	β_k	SE
Combined conditions	6.59 ***	9.36 ***	8.98 ***	$-7.54 \cdot 10^{-4} *$	0.000326182
Treatment (ETAS)	6.79 ***	9.35 ***	9.12 ***	$-9.78 \cdot 10^{-4} *$	0.000427327
Control (analyst)	6.60 ***	9.2 ***	8.79 ***	$-4.66 \cdot 10^{-4}$	0.000513346

Table 7: **Parameter estimates for multiple regression.** μ_j is the estimated mean crime rate per day in each division, β_k is the estimated average change in crime volume per minute of patrol time under different experimental conditions k . Standard errors are for estimated β_k . * $P < 0.05$ and *** $P < 0.001$ levels.

analyst predictions observed in silent accuracy testing.

In both Kent and Los Angeles, analyst predictive accuracy is a dramatic improvement over a random baseline, with the intelligence-led approach used in Kent outperforming the more exclusive focus on recent crime events followed by LAPD analysts. However, patrol patterns derived from analyst predictions would need to be significantly increased to see a measurable impact on crime. By contrast, ETAS models systematically assess the relative importance of long- and short-term contributions to dynamic crime risk. Prediction boxes therefore track the probability of where and when crimes are most likely to occur. The distribution of treatment prediction boxes is distinct from both the density of crime and distribution of analyst predictions (Figure 2C, 2F). Predictive accuracy is approximately doubled and crime reductions are therefore detectable at lower patrol levels.

4 Discussion

In this paper we introduced an experimental design that allows for the comparison of a short-term crime forecasting algorithm with existing intelligence-led and hotspot policing practices. Hotspot locations change daily based upon updated forecasts and police patrols respond dynamically. Our results serve to validate the premise of predictive policing that crime rates can be reduced under fixed policing resources by increasing the crime rate in

hotspots flagged for directed patrol. This is accomplished by using an algorithmic approach to hotspot selection, systematically estimating the relative contributions to risk of short-term and long-term crime patterns.

While the focus in this study was to compare algorithmic approaches to those employed by analysts, we note that the methods also outperform standard hotspot mapping techniques in retrospective forecasts. The accuracy of statistical models of crime hotspots has been compared to that of hotspot maps in [51] where generalized additive models are constructed for robbery prediction and [35] where chronic homicide hotspot maps are compared to marked point process models in terms of accuracy. In both studies the statistical models are shown to significantly outperform standard hotspot mapping approaches, where in [35] and Equation (1) chronic hotspot mapping is a sub-model. Statistical estimation allows one to rigorously estimate the contributions of both chronic and short-term risk to improve model accuracy and at the same time make the analyst’s job easier through automation of the model. The partitioning of the probabilistic rate of crime $\lambda_n(t)$ into stationary and dynamic components captures implicitly the roles that environmental heterogeneity and contagion-like processes play in generating crime risk [21]. In this sense, our modeling approach parallels methods that focus on contagion [23] and risk facilities [1, 10]. However, it is critical to note that these parallels are only *implicit* mechanisms in the ETAS model. Future research should evaluate whether explicit treatment of environmental heterogeneity and contagion within an algorithmic framework would improve performance.

The present study suggests that ETAS crime forecasting outperforms crime forecasting based on intelligence-led and hotspot mapping models. The results find a balance between ecological and external validity. Unlike controlled tests that pit different algorithms against historical data in sterile conditions [9, 36], the present experiment was conducted head-to-head with practitioners under all the pressures that confront crime analysts and police on a day-to-day basis. Without such ecological validity, practitioners would never believe

that an algorithmic approach could match, let alone outperform the human expert using existing best practice. Such appears to be the case here. Yet this ecological validity in the specific settings of the experimental trials could compromise the ability to generalize the results. Much depends on how representative the practices of the Los Angeles and Kent Police Department are in the broader policing world. While there is considerable tactical and strategic variation in policing practice across departments, the use of crime mapping, hotspot maps, and simple criminal intelligence to identify chronic offenders is as close to universal as anything. In a survey of US police departments, for example, 100% of departments serving populations of 500,000 people or more relied regularly on these analytical methods [15, 41]. The proportion of police departments falls to 56% for those serving populations of 50,000-100,000. Police departments in England and Wales ranked crime mapping in support of intelligence-led policing at the top of the analytical tasks they perform as a proportion of time [52]. The comparison of algorithmic predictions to best practices in intelligence-led and hotspot mapping by two prominent departments therefore would seem to have reasonable external validity especially as department size increases.

More challenging is an assessment of whether similar effect sizes can be expected in other settings. Comparisons between regions based on existing studies are problematic [2]. Tables 1 and 2 show that area standardized accuracies (PAI) across regions vary widely. In the present study area-standardized accuracies range from 3.5 to 85.2 for ETAS and 1.2 to 50.8 for analysts in Southwest and Maidstone, respectively. Variation in area-standardized accuracies is driven by differences in the physical sizes of the regions in question [47], the mixtures of crimes included in predictions [9], and the time scales over which predictions are rendered [47]. Comparison of the relative performance of different forecasting methods within regions helps solve some of these problems. Since different forecasting methods tested within a region are subject to the same constraints of physical area, mix of crimes, and forecasting time scales, the gap between their performance is less subject to such confounding factors.

This is the approach followed by [9] where kernel density estimation yields an average boost over other hotspotting methods of 1.7 to 1.8. As shown here and in [36], ETAS provides a boost over kernel density estimation, providing a basis for positing an expected magnitude of results elsewhere. Ultimately, whether this hypothesis proves true is an empirical question that will require further experimental and comparative work.

Beyond the question of broader applicability, we might also ask whether the observed impact on crime is significant as a practical matter. At mean patrol levels (31 minutes per box per day), the ETAS algorithm corresponds to 4.3 fewer crimes per week per division. While a crime reduction of 7.4% may appear small, the potential savings to society are quite large when estimated using the methodology of McCollister, French and Fang [34]. In Table 8 we display the societal costs per crime calculated in [34], broken down by costs to the victim, police and court system, and offender (2013 dollars). We then use the relative frequencies of GTA, burglary theft from motor vehicle, and burglary (.292, .584, .184) along with the average crime rate per LAPD division at zero patrol levels to estimate the weekly cost of crime per division without predictive policing (Table 9). We then estimate the savings per week achieved by a 4.3 crime reduction per division on average and extrapolate the savings across LAPD. We project a \$17,258,801 annual savings to LAPD if ETAS were to be used for 31 minutes per day in each hotspot compared to no patrol. Patrols allocated under analyst conditions are projected to achieve less than half the half the savings (\$8,223,519).

From an ethnographic perspective, we believe predictive policing operates at a local level through short-term disruption of criminal opportunities. A representative scenario arises when an officer shows up at a location designated as high risk. An offender who lives or works in that area [6] sees the officer and decides to lay low or even run [16]. In that time they are laying low they are in no position to commit a crime. If the offender comes out a few hours later and again sees the officer in the same or a nearby hotspot, the deterrence effect may last well beyond those particular policing events. Short-term crime prevention of this type is

offense type	victim cost	police & court cost	offender cost	total per crime
GTA	\$6,615	\$4,184	\$598	\$11,398
Burglary	\$1,474	\$4,465	\$737	\$6,676
BTFV	\$519	\$3,115	\$176	\$3,812
Total Costs	\$8,608	\$11,765	\$1,512	\$21,886

Table 8: **Cost per crime.** Cost of crime for the crime types targeted in the LAPD experiment calculated from MCollister et al. converted into 2013 dollars.

offense type	divisional cost per week w/o ETAS	divisional savings per week w/ETAS	annual savings per division	annual savings LAPD (projected)
GTA	\$70,971	\$5,246	\$272,808	\$5,728,959
Burglary	\$47,845	\$3,537	\$183,913	\$3,862,182
Theft	\$94,988	\$7,022	\$365,127	\$7,667,660
Total Costs	\$213,805	\$15,805	\$821,848	\$17,258,801

Table 9: **Cost savings per LAPD division.** Savings per LAPD division calculated using the methodology of MCollister et al. and a mean reduction of 4.3 crimes per division per week under ETAS patrol.

different from long-term crime prevention methods aimed at fixing the root causes of crime such as building positive police-community relations, combating chronic drug dependencies or changing the built environment to reduce risk. We are not arguing that predictive policing is a singular approach, nor is it a replacement for other proven crime prevention strategies. Rather, accurate predictions about where and when crime is most likely to occur simply enhance the ability of police officers (compared to social workers, counselors, teachers, etc) to achieve short-term crime prevention more effectively.

In future work, greater crime reductions may be achieved by improving predictive algorithms. For this purpose, maximum likelihood estimation may not be the best method for optimization. Non-smooth loss functions such as precision are likely more relevant for predictive policing, where the goal is to maximize the percentage of crime in k flagged hotspots. Learning to rank methods from information retrieval [7] can potentially be used for this purpose in order to prevent risk estimates in the top ranked hotspots from being biased by the influence of medium and low-risk hotspots, as is the case with Maximum Likelihood.

The experimental protocol used here is similar to a repeated measures crossover design wherein the same experimental subjects are exposed to different treatment conditions. Here each day is randomly assigned to control and treatment conditions and thus control and treatment events can follow one another in different sequences. We have assumed independence of days. Nevertheless, it is possible that patrol in an ETAS hotspot on one day may influence crime in an analyst hotspot on the next day (or vice-versa). There is some evidence that patrol effects last on a shorter time scale (~ 2 hours) [29] and that displacement or diffusion of benefits in time may not extend out to a day or longer (or the effect may be small). The random allocation of experimental conditions by day helps to alleviate some of this concern. In general, control days are equally likely to be followed by another control day or a treatment day (and vice versa). Thus treatment and control days experience equal exposure to any possible contamination effects. Future experiments might deploy buffer days

or randomization by weeks to further separate treatment and control conditions in time. Longer experimental periods would be needed to achieve similar significance to this study, however. Alternatively, statistical corrections might be developed if contamination across conditions was clearly indicated.

More research is also needed to better understand the relationship between police patrol and crime. Whereas officer call logs were used in this study, in-car GPS can provide a greater level of precision and also provide information on officer activity when not on a predictive policing mission. Furthermore, given the deterrent effect observed in our field experiments, patrol times and locations may be important variables to include in predictive policing models to close the feedback loop between officers, criminals, and the algorithms linking them together.

Acknowledgements

This work was supported in part by NSF grant DMS-0968309, AFOSR MURI grant FA9550-10-1-0569, ONR grant N000141010221, ARO grant W911NF1010472, ARO MURI grant W911NF-11-1-0332, ONR grant N000141210838, ARO MURI grant W911NF-11-1-0332, and NSF Grant SES-1343123. Special thanks to Greg Ridgeway and Omar Qazi. G.M. and P.J.B. are co-founders of PredPol, a company offering predictive policing services to law enforcement agencies. G.M., P.J.B., M.B.S., A.L.B., and G.E.T. hold stock in PredPol.

References Cited

- [1] Kate Bowers. Risky facilities: Crime radiators or crime absorbers? a comparison of internal and external levels of theft. *Journal of Quantitative Criminology*, 30(3):389–

414, 2014.

- [2] Kate J Bowers, Shane D Johnson, and Ken Pease. Prospective hot-spotting the future of crime mapping? *British Journal of Criminology*, 44(5):641–658, 2004.
- [3] Anthony A Braga. Hot spots policing and crime prevention: A systematic review of randomized controlled trials. *Journal of Experimental Criminology*, 1(3):317–342, 2005.
- [4] Anthony A Braga and Brenda J Bond. Policing crime and disorder hot spots: A randomized controlled trial*. *Criminology*, 46(3):577–607, 2008.
- [5] Paul J Brantingham and Patricia L Brantingham. *Environmental criminology*. Sage Publications Beverly Hills, CA, 1981.
- [6] PL Brantingham and PJ Brantingham. Notes on the geometry of crime, environmental criminology, 1991.
- [7] Christopher JC Burges. From ranknet to lambdarank to lambdamart: An overview. *Learning*, 11:23–581, 2010.
- [8] Spencer Chainey and Jerry Ratcliffe. *GIS and Crime Mapping*. John Wiley & Sons, Chichester, 2005.
- [9] Spencer Chainey, Lisa Tompson, and Sebastian Uhlig. The utility of hotspot mapping for predicting spatial patterns of crime. *Security Journal*, 21(1):4–28, 2008.
- [10] John E Eck, Ronald V Clarke, and Rob T Guerette. Risky facilities: Crime concentration in homogeneous sets of establishments and facilities. *Crime prevention studies*, 21:225, 2007.
- [11] Graham Farrell and Ken Pease. *Once bitten, twice bitten: Repeat victimisation and its implications for crime prevention*. © Crown Copyright, 1993.

- [12] David P Farrington and Brandon C Welsh. Randomized experiments in criminology: What have we learned in the last two decades? *Journal of Experimental Criminology*, 1(1):9–38, 2005.
- [13] Matthew Fielding and Vincent Jones. 'disrupting the optimal forager': predictive risk mapping and domestic burglary reduction in trafford, greater manchester. *International Journal of Police Science & Management*, 14(1):30–41, 2012.
- [14] National Centre for Policing Excellence and United Kingdom. Guidance on the national intelligence model. 2005.
- [15] Robert R. Friedmann, Richard Rosenfeld, and Nadia Borissova. Improving crime data project. Technical report, 2010.
- [16] Alice Goffman. *On the Run: Fugitive Life in an American City*. University of Chicago Press, 2014.
- [17] Rob T Guerette and Kate J Bowers. Assessing the extent of crime displacement and diffusion of benefits: a review of situational crime prevention evaluations*. *Criminology*, 47(4):1331–1368, 2009.
- [18] Glenn W Harrison and John A List. Field experiments. *Journal of Economic Literature*, pages 1009–1055, 2004.
- [19] James J Heckman. Econometric causality. *International Statistical Review*, 76(1):1–27, 2008.
- [20] S.D. Johnson, W. Bernasco, K.J. Bowers, H. Elffers, J. Ratcliffe, G. Rengert, and M. Townsley. Space–time patterns of risk: A cross national assessment of residential burglary victimization. *Journal of Quantitative Criminology*, 23(3):201–219, 2007.

- [21] Shane Johnson. Repeat burglary victimisation: a tale of two theories. *Journal of Experimental Criminology*, 4(3):215–240, 2008. 10.1007/s11292-008-9055-3.
- [22] Shane D Johnson. *Prospective crime mapping in operational context: Final report*.
- [23] Shane D Johnson, Kate J Bowers, Dan J Birks, and Ken Pease. Predictive mapping of crime by promap: accuracy, units of analysis, and the environmental backcloth. In *Putting crime in its place*, pages 171–198. Springer, 2009.
- [24] Kees Keizer, Siegwart Lindenberg, and Linda Steg. The spreading of disorder. *Science*, 322(5908):1681–1685, 2008.
- [25] George L Kelling, Tony Pate, Duane Dieckman, and Charles E Brown. The kansas city preventive patrol experiment. 1974.
- [26] David M Kennedy. Pulling levers: Chronic offenders, high-crime settings, and a theory of prevention. *Val. UL Rev.*, 31:449, 1996.
- [27] Leslie W Kennedy, Joel M Caplan, and Eric Piza. Risk clusters, hotspots, and spatial intelligence: risk terrain modeling as an algorithm for police resource allocation strategies. *Journal of Quantitative Criminology*, 27(3):339–362, 2011.
- [28] Roger E Kirk. *Experimental design*. Wiley Online Library.
- [29] C. Koper. Just enough police presence: Reducing crime and disorderly behavior by optimizing patrol time in crime hot spots. *Justice Quarterly*, 12:649–672, 1995.
- [30] Erik Lewis and George Mohler. A nonparametric EM algorithm for multiscale Hawkes processes.
- [31] Hua Liu and Donald E Brown. Criminal incident prediction using a point-pattern-based density model. *International journal of forecasting*, 19(4):603–622, 2003.

- [32] D. Marsan and O. Lengline. Extending earthquakes’ reach through cascading. *Science*, 319(5866):1076–1079, 2008.
- [33] Ross L Matsueda, Derek A Kreager, and David Huizinga. Deterring delinquents: A rational choice model of theft and violence. *American Sociological Review*, 71(1):95–122, 2006.
- [34] Kathryn E McCollister, Michael T French, and Hai Fang. The cost of crime to society: New crime-specific estimates for policy and program evaluation. *Drug and alcohol dependence*, 108(1):98–109, 2010.
- [35] George Mohler. Marked point process hotspot maps for homicide and gun crime prediction in chicago. *International Journal of Forecasting*, 30(3):491–497, 2014.
- [36] GO Mohler, MB Short, P.J. Brantingham, FP Schoenberg, and GE Tita. Self-exciting point process modeling of crime. *Journal of the American Statistical Association*, 106(493):100–108, 2011.
- [37] Clayton J Mosher, Terance D Miethe, and Timothy C Hart. *The mismeasure of crime*. Sage Publications, 2010.
- [38] Timothy C O’Shea, J Keith Nicholls, Jonathan Archer, Elton Hughes, and Jana Tatum. *Crime analysis in America: Findings and recommendations*. US Department of Justice, Office of Community Oriented Policing Services, Washington, DC, 2003.
- [39] Jerry H Ratcliffe. *Intelligence-led policing*. Routledge, 2012.
- [40] Jerry H Ratcliffe, Travis Taniguchi, Elizabeth R Groff, and Jennifer D Wood. The philadelphia foot patrol experiment: a randomized controlled trial of police patrol effectiveness in violent crime hotspots*. *Criminology*, 49(3):795–831, 2011.

- [41] Brian A. Reaves. *Local Police Departments, 2007*. Bureau of Justice Statistics Report NCJ 231174. 2010.
- [42] Robert J Sampson, Stephen W Raudenbush, and Felton Earls. Neighborhoods and violent crime: A multilevel study of collective efficacy. *Science*, 277(5328):918–924, 1997.
- [43] Lawrence W Sherman, Patrick R Gartin, and Michael E Buerger. Hot spots of predatory crime: Routine activities and the criminology of place*. *Criminology*, 27(1):27–56, 1989.
- [44] L.W. Sherman and D. Weisburd. General deterrent effects of police patrol in crime ?hot spots?: A randomized, controlled trial. *Justice Quarterly*, 12:625–648, 1995.
- [45] Martin B Short, P Jeffrey Brantingham, Andrea L Bertozzi, and George E Tita. Dissipation and displacement of hotspots in reaction-diffusion models of crime. *Proceedings of the National Academy of Sciences*, 107(9):3961–3965, 2010.
- [46] Cody W Telep, Renée J Mitchell, and David Weisburd. How much time should the police spend at crime hot spots? answers from a police agency directed randomized field trial in sacramento, california. *Justice Quarterly*, (ahead-of-print):1–29, 2012.
- [47] Lisa Tompson and Michael Townsley. (looking) back to the future: using spacetime patterns to better predict the location of street crime. *International Journal of Police Science & Management*, 12(1):23–40, 2010.
- [48] M. Townsley, R. Homel, and J. Chaseling. Infectious burglaries: A test of the near repeat hypothesis. *British Journal Of Criminology*, 43(3):615–633, 2003.
- [49] Alejandro Veen and Frederic P Schoenberg. Estimation of space-time branching process models in seismology using an EM-type algorithm. *Journal of the American Statistical Association*, 103(482):614–624, 2008.

- [50] William F Walsh. Compstat: an analysis of an emerging police managerial paradigm. *Policing: An International Journal of Police Strategies & Management*, 24(3):347–362, 2001.
- [51] X. Wang and D.E. Brown. The spatio-temporal modeling for criminal incidents. *Security Informatics*, 1(1):1–17, 2012.
- [52] Ruth Weir and Mark Bangs. *The use of Geographic Information Systems by crime analysts in England and Wales*. Citeseer, 2007.
- [53] David Weisburd. Place-based policing. *Ideas in American Policing*, 9:1–15, 2008.
- [54] David Weisburd and John E Eck. What can police do to reduce crime, disorder, and fear? *The Annals of the American Academy of Political and Social Science*, 593(1):42–65, 2004.
- [55] David Weisburd, Laura A Wyckoff, Justin Ready, John E Eck, Joshua C Hinkle, and Frank Gajewski. Does crime just move around the corner? a controlled study of spatial displacement and diffusion of crime control benefits*. *Criminology*, 44(3):549–592, 2006.

Figure 2: Crime is widely distributed in both (A) LAPD Southwest Division and (D) Kent Police Maidstone District. Intense crime hotspots are more prominent in Maidstone. (B) Control prediction boxes placed by the Southwest Division crime analyst tracked day-to-day point locations of crime leading to wide dispersion of boxes with numerous intermediate intensity prediction hotspots. (C) Treatment boxes placed by the ETAS model in Southwest are densely clustered and distinctive from both crime and analyst prediction box distributions. (E) Control prediction boxes placed by the Maidstone analyst are clustered compared with the Maidstone crime distribution and contrast with the dispersed predictions of the Southwest Division analyst, reflecting the greater focus on the activity locations of known offenders. (F) ETAS predictions in Maidstone are even more tightly clustered relative to crime and analyst prediction boxes. Shown in (A) is a random sample of 1000 crime locations in LAPD Southwest Division and in (D) all 485 recorded events in Maidstone, Kent. Shown in (B-C, E-F) are random samples of approximately 1000 prediction boxes representing $\approx 15\%$ of control and treatment predictions in Southwest and $\approx 10.3\%$ of control and treatment predictions in Maidstone. Maps show kernel density estimates with constant bandwidth across all cases.

Figure 3: Fraction of prediction boxes repeated on day N following a focal treatment or control day. Comparisons are between predictions delivered for the same shift.

Figure 4: (A) Individual patrol units in LAPD Southwest Division varied in the total amount of time dedicated to prediction missions over the course of the experiment. For example, some patrol units spent nearly 10,000 minutes in total in prediction boxes, while others spent less than 1000 minutes in total. The amount of time spent in control and treatment boxes per patrol unit, however, is approximately equivalent. Randomization by day to control or treatment successfully apportioned variation in patrol time engagement among units across experimental conditions. Plotted are cumulative times on treatment and control missions for each patrol unit that was active in Southwest Division between May 20, 2012 and January 10, 2013. The regression line shows that for a patrol unit spending on average 1000 minutes in control predictions boxes they spent on average 1017 minutes in treatment prediction boxes. (B) The frequency distribution of minutes per stop in unique 150 x 150 m prediction boxes. The exponential form of stop durations through 120 minutes suggests that officer available time is controlled by Poisson arrivals of new calls for service.

Figure 5: Total minutes on mission per day in (A) Foothill, (B) North Hollywood, and (C) Southwest Divisions. Color represents the distribution of total minutes of patrol in boxes receiving greater (blue) or lower (yellow) amounts of patrol per day. Variation in total minutes on mission reflects day-to-day variation in officer available time, but increases over the course of the experiment in Foothill and Southwest Divisions. The total number of boxes visited per day varies in phase with total minutes on mission. Peak days for total minutes on mission also show concentration of activity in a subset of all boxes. Control and treatment conditions are combined in all three figure panels.