J. Benefit Cost Anal. 2017; 8(3):360–368
doi:10.1017/bca.2017.25

Jillian B. Carr

Estimating the Effects of Police Technology Using Quasi-Experimental Methods

Abstract: Law enforcement agencies are adopting a variety of new surveillance technologies at a fast pace. These technologies could have substantial benefits in terms of public safety, but, for many of them, their ability to reduce crime is unknown. Although a small experimental literature addresses some of these technologies, many of the implementations have been too small to provide an accurate measurement of their potential. In this paper, I explore the advantages and make general suggestions about the use of quasi-experimental methodologies in estimating the public safety benefits of police technology. I also consider the specific case of license plate readers and provide some examples of difference-in-differences approaches that could be used to study their efficacy.

Keywords: difference-in-differences; law and regulation; license plate readers; police; quasi-experiment; science and technology; surveillance.

JEL classifications: K14; C23.

1 Introduction

The use of technology in policing is increasing dramatically – jurisdictions are investing in a host of new technologies with the intention of making officers more efficient, fair and productive. Surveillance technologies in particular enable agencies to do more police work without increasing manpower. Some technologies have been adopted because they essentially automate a task that police previously did manually, increasing the speed and volume of the activity, (e.g., searchable databases containing things like DNA, outstanding warrants and sex offender addresses). Others allow law enforcement to "see" and "hear," as well as record, events that they would otherwise lack the manpower to observe using audio sensors and cameras.

Quantifying the ability of such technologies to reduce crime using randomized control trials can be challenging for cost, feasibility and political reasons, and

Jillian B. Carr: Assistant Professor, Department of Economics, Krannert School of Management, Purdue University, 403 W. State St, West Lafayette, IN 47907-2056, USA, e-mail: carr56@purdue.edu

the effectiveness of many of these technologies remains unknown. In this paper, I intend to suggest some strategies for identifying quasi-experimental approaches that could be used to analyze a host of policing strategies, some of which have already been used to successfully study DNA databases (Doleac, 2017) and sex offender registries (Agan, 2011). Determining whether and how much such strategies reduce crime is a first order objective for any researcher with the intention of performing a benefit-cost analysis of adopting any of these new policing technologies.

I apply these strategies to a particular surveillance tool that has become widely used – license plate readers. Although this technology is important in its own right, it also serves as a representative of a class of technology-based interventions that could alter a would-be offender's perceptions of likely police responses, including gunshot detection sensors (often referred to by their proprietary name "ShotSpotter") and body-worn cameras.

2 Background on license plate readers

License plate readers (or LPRs) allow jurisdictions to scan a massive number of vehicles' license places without having to expend additional labor resources. The scans are checked against law enforcement databases to determine whether the vehicle or plates have been reported stolen and whether the driver to whom the car is registered has an outstanding warrant or unpaid fines. The readers are often mounted on patrol cars and set to notify the officers when a "hit" is found, indicating that it turned up a match in one of the databases. The scans can also be stored for use in other types of investigations.

One of the objectives of installing a license plate reading system is to assist in solving auto theft crimes. Solving these existing crimes may have the effect of preventing additional crimes through a number of channels.

First, as cases are cleared and adjudicated, the offenders responsible for them are likely to face incarceration, which prevents them from committing crimes at least in the short run. These effects are often described as "incapacitating" would-be recidivist offenders by removing the opportunity to commit crime in communities. Estimates from Owens (2009), suggest that incapacitation effects are substantial, particularly for prime-age offenders in their early twenties.

Second, if potential criminals perceive an increase in the probability of detection, they may be deterred from committing crimes. In seminal work in the economics of crime, Becker (1968) suggests that in addition to the costs and benefits considered by all economic actors, potential criminals also consider the likelihood of detection and the punishment that will be conferred if they are found guilty. Any perceived increase in likelihood of detection would alter this basic calculus directly (for empirical evidence that additional police reduce crime, see Di Tella & Schargrodsky, 2004). Importantly, this deterrence depends on the salience of the policy. If the policing does not appear to be more substantial, potential offenders may never consider this additional likelihood that they are caught. Jurisdictions may even choose to implement the system with a large public information campaign in an attempt to amplify effects. Some empirical evidence of these detection effects can be found in Carr and Doleac (2017) who show that gunshots recorded by ShotSpotter sensors actually increase after curfews remove potential witnesses from the street.

Some jurisdictions have reported using license plate readers to solve other crimes, increasing the potential benefits of the technology in terms of crime reduction. Although it is not necessarily the intention of adopting the technology, this added investigative benefit can have the same effects (in terms of incapacitating individual criminals and deterring crimes in general) on other types of crime. Even if individuals connected to criminal networks are not cognizant of the technology causing the increase in detection, they may attribute a perceived increase in clearance rates to an overall increase in police efficiency.

Determining whether any of these mechanisms have an effect on crime is an empirical challenge, and attempting to disentangle them from one another is even yet more challenging. In this piece, I intend to offer some suggestions for methodological approaches. In the next section, I detail the current experimental research on license plate readers, giving special attention to some of the drawbacks to using a randomized control trial. In the following section, I make some specific suggestions for using quasi-experimental methodology to overcome the limitations of randomized control trials in this setting.

3 Experimental literature

Researchers have studied the effects on auto theft clearance rates and deterrent effects in a number of jurisdictions using randomized control trials. Lum, Hibdon, Cave, Koper and Merola (2011) study two adjacent jurisdictions in suburban Washington D.C.: Alexandria City and Fairfax County, Virginia. Taylor, Koper and Woods (2012) and Koper, Taylor and Woods (2013) focus on Mesa, Arizona.

All of these studies follow a similar empirical process. The researchers select "hotspots" or "hot routes" which represent geographic concentrations of auto theft or routes traveled after auto theft. Then they design randomized control trials in which some of the selected geographies are assigned to LPR use (both mobile and stationary applications) and others are assigned to standard practices. The Taylor et al. (2012) and Koper et al. (2013) studies also tested a "manual" license plate lookup treatment in which officers were instructed to look up plates without the use of the LPR technology. Importantly, in all of these studies, no announcements were made to the community about the new technology.

All three studies find no evidence in support of any deterrent effects on auto-related crimes. Taylor et al. (2012) do provide evidence that the readers lead to more "hits" than officers looking up plates manually, which lends support for the efficiency of the technology and suggests that there is some scope for incapacitation effects if more offenders are successfully incarcerated as a result.[1] Unless other would-be offenders perceive this increase in efficiency (or hear about the implementation of the new technology), deterrent effects are unlikely to follow as an immediate result.

The results in all of these papers exhibit strong internal validity and thorough experimental design. Randomized control trials are the gold standard in policy analysis – they allow for the strictest interpretation of results in a causal light, and allow the experimenter to design an intervention that can isolate the mechanisms behind any observed effects. In these studies, though, the results may not fully represent the potential effects of such technology because the scope of the treatment may have been too small to have any effect on the public knowledge of its use. For example, in the Lum et al. (2011) study, there are only 19 "hits" in Fairfax and 14 in Alexandria in which the scanning technology identified a car in need of additional attention. In a follow-up survey detailed in Lum, Merola, Willis and Cave (2010), 90% of responding residents claimed that they did not know whether LPRs are used in their jurisdiction.[2]

Randomized control trials are not always feasible, and as in the studies described above, their scale can be limited by resources and stakeholder buy-in. Small-scale randomized control trials such as these may not mirror "real-world" applications in their scope. In the context of license plate readers, the interventions seem too short to make a difference and also too low profile. The presence of potential deterrent effects due to license plate readers depends on the community's knowledge of the intervention, so the limited scale of these experiments and the resulting

1 They also find that there is less auto theft two weeks after the intervention in the areas that received the manual lookup treatment, but not the license plate reader treatment. To the extent that manual lookups appear to be more active, the manual treatment may have been more visible to residents, explaining the (potentially) surprising effect of manual lookups.

2 Similarly, in the Taylor et al. (2012) study, each hotspot received only 8 hours of treatment over the course of the study.

lack of public awareness could cause observers to underestimate the technology's potential.

4 Quasi-experimental methodologies

An alternative to using randomized control trials to study these technologies is to use quasi-experimental methods that exploit natural experiments, which take advantage of real-world scenarios that mimic the scientific method. They create inadvertent "treatment" and "control" groups that can be used to attribute causal effects to the programs of interest, and any criteria that randomly selects these groups, such as a lottery or eligibility rule, creates an ideal natural experiment. When intentional and unintentional randomization schemes are not options, quasi-experimental methods can still provide a solution. A difference-in-differences approach is less demanding of the experimental set-up and ideal difference-in-differences scenarios are often created by reasonable real-world implementation strategies. I will focus on this estimation strategy for the broad applicability of the method, acknowledging that it requires greater defense of the model identification than other alternatives.

This strategy compares the difference in *trends* in outcomes of treated and untreated units (likely geographic in this application) over time. Importantly, the validity of such an approach does not depend on the assumption that the treated and untreated units are ideal counterfactuals in absolute terms, but only that the untreated units provide an appropriate counterfactual for what would have happened to the trends in crime outcomes in the treated units, absent treatment.

For example, a jurisdiction that adopts license plate readers likely has a higher rate of auto theft than a jurisdiction that does not, but both are subject to the same overarching social trends and any other large-scale shocks that could affect crime. The use of a comparison group in this setting allows the experimenter to exploit the time series variation in the treatment (comparing before and after implementation), while controlling for the fact that external factors could cause changes in crime that could (otherwise) wrongly be attributed to the policy. Similarly, within a jurisdiction, a savvy police department would choose to use the technology in areas with high rates of auto theft using a strategy like hot spot analysis. These areas will be substantially different from those not selected, but they would be subject to all of the same local phenomena like local elections, periods of good or bad weather, or economic influences like manufacturing plant closures. For this reason, the crime trends in the untreated places will likely provide an appropriate counterfactual for the trends in treated places absent treatment.

Empirically, the estimation of a difference-in-differences model rests on the use of fixed effects in a panel regression (see, for reference, Angrist & Pischke, 2008, Chapter 5). Unit fixed effects are used in order to control for time-invariant characteristics of the unit (e.g., Mesa, Arizona has high auto theft rates) and time fixed effects are used to control for time-specific shocks that affect all units (e.g., national-level economic events). After controlling for these effects, the remaining variation is the difference in trends between treated and untreated groups.

An important component to any difference-in-differences analysis is the careful consideration of whether trends in comparison units truly do provide an appropriate counterfactual for trends in treated units. The first line of defense in supporting this claim is to compare the trends in the outcome variables before the time of implementation. This is often done using an event study framework, where the researcher examines whether the differences in trends coincide with implementation of the treatment or they commence before. Any deviation before implementation suggests that the comparison group is not a good counterfactual.

In the following two subsections, I explain the application of difference-in-differences methodologies to across-jurisdiction and within-jurisdiction settings, respectively.

4.1 Across-jurisdiction methodologies

A potentially obvious option is to use jurisdictions that never adopted such policies as a comparison group for those that did. Although it is reasonable to think that adopting agencies are different from non-adopting agencies in a number of ways, the important assumption is not that the treated and untreated groups are the same, but that the comparison group provides a reasonable counterfactual for the *trends* that would have occurred in the treated jurisdictions absent treatment.

In the case of license plate readers, preliminary information on the use of the technology across law enforcement agencies has already been collected as a part of the Lum et al. (2010) study – the authors performed a survey of jurisdictions about their use of this technology. In order to get a more recent measure of use, a follow-up survey could be performed, or the manufacturers of such products could be contacted to obtain a list of the jurisdictions using the technology currently and when they began to use it. Using detailed information on the timing of adoption to construct a measure of when the technology was in use would allow for the identification of the causal effects of license plate readers in the set of adopting jurisdictions, relative to the effects of general societal influences measured in those that did not adopt.

In order to test the hypothesis that deterrent effects are driven by the knowledge of license plate reader use, it would also be useful to find the first mention of the technology in local media outlets or police department press releases. In fact, I would consider the public announcement and the date of first use separately. The original survey included questions about the ways that departments use the technology (e.g., fixed vs. mobile use), and these measures would be useful in the study I outline here to explore heterogeneous effects by implementation style.

After details on the use of the technology by jurisdictions are obtained, the data contained in the National Incident-Based Reporting System (NIBRS) could provide an appropriate jurisdiction-level dataset for analysis. The NIBRS data contain detailed information on the offenses and the offenders that could be used to examine competing hypotheses about the underlying mechanisms.

4.2 Within-jurisdiction methodologies

Alternatively, geocoded crime data could be obtained from a jurisdiction or a set of jurisdictions in order to look at more localized effects. These two different levels of analysis could even be presented in the same study to strengthen claims.

A scenario that could lend itself to within-jurisdiction analysis is one where the rollout of the license plate readers is staggered temporally across subgeographies, likely due to resource or staffing constraints. For example, if a police department only has the ability to train half of the officers needed for an auto theft prevention task force, they may do so by training only the officers from some districts. This would allow those districts to start the program earlier than others. The areas where the officers are yet to be trained would serve as a comparison group. Any scenario in which parts of a city are treated at different times could provide such an opportunity.

Again, citizen awareness of the use of this new technology (or at least the awareness of rising clearance rates) is a necessary condition for deterrence to occur. Absent official announcements and media coverage, this awareness is tied to the observation of the technology's use or the knowledge that it has increased auto theft clearance rates. The intensity of treatment across communities within a jurisdiction is therefore related to the outcomes measuring the technology's successes. Rather than just studying the presence or absence of the intervention, an alternative measure would be to use the number of plates scanned, the number of "hits" or the number of cases cleared using the technology. If a jurisdiction announced where the technology was debuting over time, then that would be an even more direct way to measure citizen awareness. It is more likely that a jurisdiction would make a general announcement about the LPRs and then quietly roll them out over time.

In that scenario, it might make sense to employ an interrupted time series analysis in which the researcher tests for breaks in trends in the outcomes of interest at the time of the intervention in addition to using the difference-in-differences methods that I have described. This could allow for the separation of effects due to formal announcement and the observation of local LPR use and effectiveness.[3]

5 Conclusion

Natural experiments, while sometimes difficult to identify, are useful in scenarios in which the cost of a large-scale randomized control trial is prohibitively expensive or logistically impossible. In this piece, I have described methods that use naturally occurring variation in the policy environment to study the effects of license plate reader use on crime. Specifically, I suggest two ways to use a difference-in-differences methodology to exploit the variation in the timing and intensity of the use of the readers to obtain plausibly causal estimates of the deterrent effects.

These difference-in-differences strategies can be used in a way that explores an important part of the puzzle that I believe is understudied in the crime technology literature broadly – the role of public knowledge of the use of policing technology and the resulting deterrent effects. This caveat could explain the lack of detected effects in much of the existing experimental literature on LPRs. License plate readers can only deter crime if the populous is aware of their use and the effect that they have on the clearance rate of crimes.

This particular nuance is common to other surveillance technologies as well (such as gunshot audio sensors and criminal justice databases), making disentangling these effects of utmost importance. The difference-in-differences methodology can be applied to these other policing innovations as well, as long as the common trends assumption can be supported. Quasi-experimental methodologies may be especially useful when the scale of a randomized control trial would limit public knowledge of the innovation, and they may even allow researchers to explore these issues of citizen awareness directly.

Like many new technologies in policing, further study into the efficacy of license plate readers at deterring crime and their other benefits and costs is critical. As police agencies make decisions about how to use their scarce resources,

3 A full discussion of the interrupted time series methodology is beyond the scope of this piece, but in such a scenario it could prove useful with the caveat that the identifying assumption is much harder to support than that of a difference-in-differences model. It requires that no other determinants of the outcome variable are changing at the same time as the treatment so that any observed changes can be ascribed to the treatment.

having a full picture of the benefits and costs of competing uses will enable them to use those resources in the most beneficial ways.

References

Agan, Amanda (2011). Sex Offender Registries: Fear without Function? *The Journal of Law and Economics*, *54*(1), 207–239.

Angrist, Joshua D. & Pischke, Jörn-Steffen (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Becker, Gary S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy*, *76*(2), 169–217.

Carr, Jillian B. & Doleac, Jennifer L. (2017). Keeps the Kids Inside? Juvenile Curfews and Urban Gun Violence. *Review of Economics and Statistics*, forthcoming.

Di Tella, Rafael & Schargrodsky, Ernesto (2004). Do Police Reduce Crime? Estimates Using the Allocation of Police Forces After a Terrorist Attack. *American Economic Review*, *94*(1), 115–133.

Doleac, Jennifer L. (2017). The Effect of DNA Databases on Crime. *American Economic Journal: Applied Economics*, *9*(1), 165–201.

Koper, Christopher S., Taylor, Bruce G. & Woods, Daniel J. (2013). A Randomized Test of Initial and Residual Deterrence from Directed Patrols and Use of License Plate Readers at Crime Hot Spots. *Journal of Experimental Criminology*, *9*(2), 213–244.

Lum, Cynthia, Hibdon, Julie, Cave, Breanne, Koper, Christopher S. & Merola, Linda (2011). License Plate Reader (LPR) Police Patrols in Crime Hot Spots: an Experimental Evaluation in Two Adjacent Jurisdictions. *Journal of Experimental Criminology*, *7*(4), 321–345.

Lum, Cynthia, Merola, Linda, Willis, Julie & Cave, Breanne (2010). License Plate Recognition Technology (LPR): Impact Evaluation and Community Assessment. Center for Evidence-Based Crime Policy, George Mason University.

Owens, Emily G. (2009). More Time, Less Crime? Estimating the Incapacitative Effect of Sentence Enhancements. *The Journal of Law and Economics*, *52*(3), 551–579.

Taylor, Bruce G., Koper, Christopher S. & Woods, Daniel J. (2012). Combating Vehicle Theft in Arizona: A Randomized Experiment with License Plate Recognition Technology. *Criminal Justice Review*, *37*(1), 24–50.