

The geography, incidence, and underreporting of gun violence: new evidence using ShotSpotter data

Jillian B. Carr* and Jennifer L. Doleac†

April 2016

Abstract

This paper provides new evidence on the extent of underreporting of gun violence. Criminal activity is often selectively underreported – that is, underreported in a non-random manner. This can make it difficult to understand public safety problems and devise effective policy strategies to address them. However, new surveillance technologies are facilitating the collection of more accurate data on crime. In this paper, we describe data on gunfire incidents, recorded using a tool called ShotSpotter. We compare those data with previously-available data on gun violence (reported crime and 911 calls) to estimate baseline correlations between these measures as well as the causal effect of gunfire incidents on reporting. Using data from Washington, DC, and Oakland, CA, we find that only 12% of gunfire incidents result in a 911 call to report gunshots, and only 2-7% of incidents result in a reported assault with a dangerous weapon. These extremely low reporting rates have important implications for research on gun violence. The characteristics and research potential of ShotSpotter data are relevant to surveillance data more broadly; while such data have not yet been exploited by social scientists, they could be extremely valuable for crime research and policy.

We thank Bill Gale, John Pepper, and participants in the 2015 APPAM Big Data Pre-Conference Workshop for valuable feedback.

*Department of Economics, Purdue University, carr56@purdue.edu

†Corresponding author; Batten School of Leadership and Public Policy, University of Virginia, jdoleac@virginia.edu

1 Introduction

The selective underreporting of criminal activity is a primary concern for researchers and policymakers. No traditional source of crime data includes the full universe of criminal activity. This is concerning because the reporting of criminal activity probably varies across communities and crimes in a non-random manner. Of even greater concern, many policy interventions that researchers might like to evaluate (such as hiring additional police officers) likely affect both the amount of crime and the rate at which crime is reported. Any analyses using traditional crime data will therefore produce biased estimates, and often that bias will be unsigned (Pepper, Petrie, and Sullivan, 2010).

High-tech surveillance data could be a game-changer. This paper describes a new source of data on gun violence, from a widely-used surveillance tool called ShotSpotter. It compares the ShotSpotter (surveillance) data with the best available data on gun violence from traditional sources (data on reported crimes and 911 calls), and uses them to consider the degree of crime underreporting as well as the value of the ShotSpotter data for policy analyses.

Using data from Washington, DC, and Oakland, CA, we find evidence of severe underreporting of gun violence. An individual gunfire incident provides an exogenous shock to the likelihood of a 911 call or crime report, allowing us to estimate the effect of gunfire on reporting. In DC, only 12.4% of gunfire incidents result in a 911 call to report shots fired. Any time a gun is fired with the intent to injure someone (or to threaten someone with injury), an assault with a dangerous weapon (AWDW) has been committed. However, only 2.3% of gunfire incidents in DC and 6.4% of gunfire incidents in Oakland result in a reported AWDW. Unless the remaining 93-98% of gunfire in these cities is not intended to be threatening (which seems highly unlikely), this type of crime is reported at extremely low rates. The effect of gunfire incidents on reporting varies by Police District, implying that underreporting is a larger problem in some places than others. Together, these estimates suggest that analyses based on traditional crime data could be misleading.

2 Traditional data on criminal behavior and gun violence

For unbiased analyses of policy effects on crime, we need data that capture all criminal activity, or at least a random subset of all criminal incidents. This does not describe traditional sources of crime data.

The best-known datasets on criminal behavior are the Uniform Crime Reports (UCR) and National Incident-Based Reporting System (NIBRS), both maintained by the FBI. The UCR and NIBRS provide information on the number of reported crimes at the reporting-agency level (typically a city or county); NIBRS also includes richer detail on these offenses for the subsample of jurisdictions that choose to participate

in the program. These are, technically, administrative data – which typically implies they are complete and high-quality – but they are collected from individual jurisdictions across the country, and those jurisdictions do not always respond accurately or completely. Additionally, and by nature, they miss any criminal activity that is not reported to law enforcement and recorded as a crime. Both the UCR and NIBRS arguably improve upon large-scale surveys such as the National Crime Victimization Survey (NCVS), which asks respondents to recall crimes from previous months that they may or may not have reported to the police. In reality, these data sources are complementary, due to concerns about the selective underreporting of crime. All are, ultimately, imperfect proxies for true criminal activity.

An increasing number of academic papers rely instead on detailed administrative data from local agencies. Local administrative data on individuals arrested for or convicted of crimes provide more flexibility in terms of the issues researchers can address (for instance, tracking individuals over time to measure recidivism). However, arrests and convictions are, again, imperfect proxies for criminal behavior. For instance, racial disparities in how individuals are perceived and/or treated by law enforcement, victims, and witnesses, could affect the likelihood that they are included in these datasets, conditional on the same underlying behavior. Such sample selection could bias the apparent effects of crime-prevention policies.

Data on gun violence are particularly problematic. The ideal data on gun violence would be something along the lines of every incident where a gun threatened someone else’s safety. As a proxy for this, researchers can use administrative data on reported crime that include weapons used (such as NIBRS), but — as noted above — not all gun violence is reported to police. Gun violence research typically focuses on gun homicides to minimize concerns about underreporting, but obviously not all shots fired (or guns wielded as a threat) result in death, and in fact as medical technology improves fewer gunshot injuries will result in fatalities. Many jurisdictions can provide data on 911 calls reporting shots fired. However, gunshots that do not hit anyone are often not reported to police, and this selective underreporting may be particularly problematic in the most violent neighborhoods if residents don’t trust the police to be helpful. The NCVS asks whether respondents were victims of a crime committed with a firearm, but these data are subject to the usual concerns about the validity of survey responses and self-selection of respondents; they also do not include precise times or locations of crime incidents. The Centers for Disease Control and Prevention (CDC) maintains city-level data on fatal injuries (from death certificates) and nonfatal injuries (from hospital emergency rooms), but these will obviously not include information about gun violence that does not result in injury, or individuals who avoid hospitals for fear of being arrested. Research based on any of these traditional data can provide suggestive evidence, at best.

This situation is distressing, given the important, often life-and-death nature of questions related to criminal justice policy and violent crime. But there is good news: improvements in technology are changing

this status quo. As law enforcement and governments increase their use of surveillance tools, they collect a great deal of objective data on criminal activity. These data have not yet been exploited by social science researchers, but have the potential to transform the field. In this paper, we consider a new source of data, generated by ShotSpotter technology, from two large cities. These data measure the quantity, timing, and locations of gunfire incidents with greater accuracy than do reported crime or 911 call data. More importantly, the accuracy of the data is unlikely to be affected by most interventions that aim to reduce gun violence. For this reason, using ShotSpotter data as an outcome measure is more likely to produce meaningful, unbiased results in policy evaluations.

3 Description of ShotSpotter and evidence on accuracy

ShotSpotter consists of audio sensors implemented throughout a targeted area (on top of buildings and in similar discrete locations), which detect the sound of gunfire and triangulate its location. An algorithm analyzes the recorded sound and determines whether it was gunfire or another loud noise such as construction or fireworks. If it is confirmed to be gunfire, the relevant information (including time, location, and a recording of the incident) is sent to local police so that they can quickly go to the scene.

ShotSpotter has been adopted by over ninety jurisdictions across the United States. Not all of these jurisdictions are major cities, though places that implement ShotSpotter tend to have higher crime rates than average. The firm releases an annual "National Gunfire Index" summarizing the data its system generates in cities across the United States (ShotSpotter, 2015). However, they do not make the incident-level data publicly available. We have obtained incident-level data directly from several jurisdictions where the data are considered public record.

The appeal of ShotSpotter-generated data is that they likely provide a more accurate count of "true" gunfire incidents than data such as reported crime or 911 calls. In addition, they include timestamps and geocodes that are far more precise than those in reported crime and 911 call data. However, ShotSpotter data are not perfect. In all data on gunfire – indeed, in all data on any criminal activity – there are two potential problems: false positives (detected incidents that were not actually gunshots) and false negatives (actual gunshots that were not detected). The technology on which ShotSpotter data are based has improved over time, but there is limited independent, published evidence of its current accuracy. Redwood City conducted a field trial in 1997; it found that ShotSpotter detected "nearly 80 percent of the test shots" and "was able to triangulate (locate) 84 percent of the test shots (N = 26 of 31 shooting events) within a median margin of error of 25 feet" (Mazerolle, Frank, Rogan, and Watkins, 2000). In a 2006 study financed by the National Institute of Justice, ShotSpotter "detected 99.6 percent of 234 gunshots at 23 firing locations", and

"located 90.9 percent of the shots within 40 feet" (Goode, 2012), though the report noted that quality of implementation will be key to the sensors' success (Litch and Orrison, 2011).

To address the number of false positives, ShotSpotter funded an independent study of the technology's use and effectiveness. The researchers interviewed employees at seven police departments that use ShotSpotter. The interviewees were asked what share of ShotSpotter alerts were actual gunfire. Their responses ranged from 50% to 97% (averaging 67%). However, these estimates were based on individuals' perceptions, not an analysis of actual data. Also, it is important to note that it is typically impossible to distinguish false positives from gunshots that cannot be corroborated by other evidence (e.g., a 911 call or evidence found at the scene). That said, large spikes in detected gunfire incidents on New Year's Eve and July 4th suggest that the algorithm sometimes confuses fireworks and firecrackers with gunfire (though local residents would probably not be able to distinguish between these sounds either, out of context). At this point, there is no reliable evidence about the rate of false positives in actual ShotSpotter data, and this is an area where future research would be helpful.

False negatives and false positives in any dataset are a concern for researchers if they are not random. If they are random noise, they introduce measurement error, which could increase the standard errors on empirical estimates when the data are used as an outcome measure. However, if they are non-random – for example, if gunfire in particular areas or by particular people is systematically less likely to be detected – this will introduce selection bias. False negatives and false positives are a greater concern for research on gun violence if they are affected by the policy interventions being studied (Pepper, Petrie, and Sullivan, 2010). The typical concern is that many policies that could reduce gun violence – like increasing policing in dangerous neighborhoods – probably also affect the reporting of gunfire; this makes it difficult to determine the true effect on criminal activity when using reported crime as the outcome measure. The promise of ShotSpotter data is that (1) both measurement error and selection bias should be much lower than when using reported crime data or 911 call data, and (2) the detection of gunfire by ShotSpotter will be unaffected by policy interventions that aim to reduce gun violence.

3.1 Description and availability of ShotSpotter data

ShotSpotter data typically include the following information on each incident: date, time, location (latitude and longitude), and whether the incident consisted of a single gunshot or multiple gunshots.

These data on gunfire incidents exist for many cities in the United States, but not all. The cost of the technology limits its coverage: not all cities choose to implement it, and those that do target the most violent neighborhoods.

At this time, data are freely-available to researchers for only a small subset of ShotSpotter cities. Most local contracts give ShotSpotter ownership of the data, so the data are not considered public record in most places. (Note that ShotSpotter is open to selling the data to researchers.) We have obtained data on gunfire incidents from the following cities: Washington, DC; Oakland, CA; Beloit, WI; Redwood City, CA; Youngstown, OH; Canton, OH; Peoria, IL; and Nassau County, NY. Animated maps of the incidents over time in each jurisdiction are available at <http://jenniferdoleac.com/maps/>.

Their limited availability makes ShotSpotter data most useful for city-specific research questions. They are currently less useful for projects where cross-city comparisons are necessary. However, as coverage expands, and as ShotSpotter becomes more open to collaborating with researchers, this could change.

4 Estimating the underreporting of gun violence

We use the timing of gunfire incidents to test the effect of gunfire on crime reports and 911 calls. This allows us to consider the extent of underreporting.

4.1 Data

To compare ShotSpotter data with traditional sources of crime data, we focus on Washington, DC, and Oakland, CA – the two largest cities in our sample. We were able to obtain incident-level, timestamped, geo-coded data on ShotSpotter incidents and traditional crime measures (reported crime and/or 911 calls) for the following time periods: January 2011 through June 2013 in DC, and January 2008 through October 2013 (excluding January through July of 2011) for Oakland.¹ Figures 1 and 2 show heatmaps of the gunfire incidents in DC and Oakland, respectively.

Because we are interested in whether gunfire incidents are reported, we focus on the categories of reported crimes and 911 calls that are most likely to be associated with such an incident. The reported crimes of interest are homicide and assault with a dangerous weapon. The 911 call categories of interest are all calls, calls for police assistance, and calls to report gunfire.

We obtained incident-level data on reported crimes and 911 calls from Washington, DC. These data include the date, time, and location of each event, but these are generally less precise than in the ShotSpotter data. The time and date of the reported crime will often be an estimate based on when the crime was discovered and/or when it was reported to police. The time and date of the 911 call is when the call was received at dispatch. In all cases, the location is an address rather than a latitude and longitude. In the

¹In Oakland, ShotSpotter data are unavailable for January through July of 2011, apparently due to a technical problem with the sensors. We thus exclude those months from the analysis.

case of 911 calls, the address is typically that of the caller, not necessarily where the crime occurred; it is therefore a particularly noisy measure for calls reporting gunshots. Because ShotSpotter covers a relatively large area in Washington, DC, we aggregate to the police district (PD) level.² Each observation is the number of incidents occurring during a particular hour in a particular PD.

We obtained incident-level data on reported crimes (but not 911 calls) from Oakland, CA. The data characteristics are similar to those in DC. We aggregate incidents to the city level, so each observation is the number of incidents occurring during a particular hour, city-wide.

As mentioned above, there are large spikes in ShotSpotter-detected incidents on New Year’s Eve and near the July 4th holiday. This is likely due to real increases in celebratory gunfire as well as large numbers of false positives due to fireworks or firecrackers that make it through the algorithm’s filter because they sound so similar to gunshots. To limit the effect of such outlier events on our results, we exclude January 1st, the week of July 4th, and December 31st.

Summary statistics for DC are in Table 1. Summary statistics for Oakland are in Table 2.

4.2 Empirical Strategy

Our goal is to estimate how gunfire incidents affect the number of 911 calls (total calls, calls for police, and calls to report gunshots), and the number of reported violent crimes (homicides and AWDWs).³ The intuition is that a gunshot incident is the initial event, and crimes are reported and/or 911 calls made as a result of that incident. Conditional on a variety of time and location fixed effects, the occurrence of a gunfire incident can be thought of as an exogenous shock to the likelihood of a 911 call or reported crime during that hour. The estimated effect of gunfire on reports tells us about the extent of underreporting of gun violence.

Descriptively, we see suggestive evidence that underreporting is a problem: In DC, there are 4,483 hours during our sample period with at least one gunshot incident. Only 982 of those hours also had at least one 911 call to report shots fired, and only 453 of the 4,483 hours with gunfire saw at least one AWDW reported. Only 110 of the 4,483 hours with gunfire incidents included both a 911 call to report shots fired *and* a reported AWDW.

To more rigorously test for the effects of gunfire incidents on reporting, we construct a balanced panel of data on the number of ShotSpotter incidents, reported crimes, and 911 calls, by hour, by location (PD or

²ShotSpotter has been implemented in PDs 3, 5, 6, and 7, so we restrict our attention to these areas.

³Note that an incident where someone was shot but not killed, or shot at but not hit, would be considered an AWDW.

city). We consider the effect of the number of gunfire incidents using the following specification:

$$\begin{aligned}
 Outcome_{i,h,p} = & \alpha + \beta_1 Gunshots_{h,p} + \beta_2 LaggedOutcome_{i,h,p} + \\
 & \lambda_{Location} + \gamma_{HourOfDay} + \delta_{MonthOfYear} + \theta_{Year} + e_{i,h,p},
 \end{aligned}
 \tag{1}$$

where i is the outcome of interest (type of reported crime or 911 call), h is the hour of the day, and p is the location (PD in DC, or city-wide in Oakland). A number of fixed effects (location, hour of day, month of year, and year) control for omitted variables that might affect both the number of gunfire incidents and reported crime, to better isolate the effect of a specific gunfire incident. The fixed effects also absorb variation in the outcome measure to help us detect gunfire’s effects. The outcomes of interest are: number of homicides, number of AWDWs, number of 911 calls, number of 911 calls for police assistance, and number of 911 calls to report gunfire. (Recall that we have 911 call data for DC only.) *Gunshots* is the number of gunshot incidents detected by ShotSpotter during the given hour in that location. *LaggedOutcome* is the number of reports or calls in the same location during the previous hour or during the same hour the previous day (i.e., 24 hours earlier); this controls for the recent level of reported crime or 911 calls in that location (which might not be fully absorbed by the fixed effects). To deal with serial correlation (observations in one hour are not independent from observations in the next hour), robust standard errors are clustered by date. The coefficient of interest is β_1 .

4.3 Results

The results of these regressions are presented in Tables 3 through 5.

In Washington, DC, there are strong correlations between gunshot incidents and reported crime, and between gunshot incidents and 911 calls. The first column of Table 3 shows the unconditional effects of a gunfire incident on the number of reported crimes and 911 calls. These suggest that in DC, on average, there are 0.0055 homicides for every one gunfire incident, and 0.0339 AWDWs for each gunshot incident. Turning to 911 calls, there are 0.6442 calls for service for every one gunshot incident detected by ShotSpotter in a given hour; 0.5690 are calls for police assistance, and 0.1427 are calls to report gunshots. These numbers provide baseline correlations between these different measures of gun violence, though they don’t necessarily tell us whether the reports or calls were the result of the gunfire incidents. It could be that gunfire and other reported crimes simply tend to happen at the same time (for instance, in a dangerous neighborhood late at night there might be an AWDW on one block and an unrelated gunfire incident nearby). We’ll want to control for typical levels of crime reports and calls to isolate the effect of the gunfire.

Columns 2–5 each add an additional set of fixed effects, to estimate the causal effect of gunfire on

reporting. Controlling for the hour of the day has the biggest effect on the estimated coefficient. Most of the controls have little to no impact, which suggests that the timing of gunshot incidents is indeed random with respect to other factors that might affect crime reports or 911 calls. Columns 6 and 7 add the lagged outcome measure (lagged by 1 hour or 24 hours, respectively), which have little effect on the estimates. Column 7 provides our preferred estimates for the causal effect of gunfire on crime reports and 911 calls, and shows that the number of gunfire incidents has a statistically-significant effect on all of the outcomes of interest. These estimates imply that 0.5% of gunfire incidents result in a homicide – a very small share. Studies that rely on homicide as the outcome variable therefore ignore a great deal of gunfire. 2.3% of gunfire incidents result in a reported AWDW. Note that while not all gunfire kills someone, most gunfire is probably intended to be threatening (and so constitutes an AWDW). 22.0% of gunfire incidents result in a 911 call (which would include calls for an ambulance), 20.7% result in a 911 call for police assistance (which would include calls to report a crime), and only 12.4% result in a 911 call to report gunfire. Even if someone calls 911 from the scene to report an injury or crime, we might expect that other residents nearby would call to report hearing gunshots. This is apparently uncommon.

Table 4 shows Column 7 for each PD separately. In Police District 3, 0.8% of gunfire incidents results in a homicide – the highest percentage among all four PDs (suggesting that gunfire is more likely to be deadly here than elsewhere) – but only 9.3% of incidents result in a 911 call to report gunshots – the lowest percentage among all four PDs. The share of gunfire incidents that results in 911 calls ranges from 13% in to 48%. In the PD with the highest 911 call rate – Police District 5 – gunfire incidents were the most likely to result in a reported AWDW, suggesting that residents’ willingness to call the police matters for whether a crime is recorded. These estimates imply that underreporting varies from place to place, which could result in selection bias in empirical estimates based on traditional data sources.

Table 5 presents results for reported crime in Oakland. Column 1 shows that, on average, there are 0.0162 homicides and 0.1034 reported AWDWs for each gunfire incident. As above, these should be interpreted as baseline correlations.

Columns 2–4 each add an additional set of fixed effects to estimate the causal effect of gunfire incidents on reports and calls. As in DC, controlling for hour of the day makes the most difference. Columns 5 and 6 add lagged outcome measures (lagged by 1 hour or 24 hours, respectively); these have very small effects on the coefficients. The estimates in Column 6 reveal that the number of gunfire incidents has a statistically-significant effect on the number of homicides and AWDWs. The magnitudes imply that 1.0% of gunfire incidents result in a homicide, and 6.4% of gunfire incidents results in a reported AWDW.

These results present strong evidence that gunfire incidents are underreported in these cities. Note that the estimated effect of gunfire on both homicide and AWDWs varies across these jurisdictions. On average,

gunfire in Oakland is about twice as deadly as gunfire in DC (though only slightly more deadly than gunfire in DC's Police District 3). Homicide is reported with near-accuracy, so this difference is probably not due to a difference in reporting. There might be differences across cities in terms of the nature of victims' injuries, or the speed with which they get to a hospital, that could affect the likelihood of surviving.

5 Discussion

This paper describes a new source of data on gun violence that is unaffected by selective underreporting and therefore could be more useful for researchers than traditional crime data. These data are representative of the highly-scalable "big data" generated by surveillance technology.

We show that underreporting of gun violence is a real concern in two major cities: Washington, DC, and Oakland, CA. In DC, only 12.4% of gunfire incidents result in a 911 call to report shots fired, and only 2.3% of gunfire incidents result in a reported AWDW (the crime that is committed when someone fires a gun in a threatening manner). In Oakland, 6.4% of gunfire incidents result in a reported AWDW, still very low. These results are consistent with a model of violent crime where neither the victim or offender is interested in involving the police (e.g. gang or drug-related violence). We also find evidence that the extent of underreporting varies across areas within the city. In DC, the probability that a gunfire incident results in a 911 call to report shots fired ranges from 9.3% in Police District 3 to 18.0% in Police District 5. Given these low reporting rates, and the variation across communities, it is very likely that policy interventions that might affect crime also affect reporting rates. This has important implications for research that uses reported crime data as an outcome measure.

Indeed, there are many contexts where changes in reporting could make it difficult to measure policy effects on crime: police hiring or police slowdowns (including testing for a so-called "Ferguson effect"), interventions aimed at improving trust between residents and police, greater use of surveillance cameras to deter criminal activity, neighborhood watches, and so on. Carr and Doleac (2015) analyze the effect of juvenile curfews on gun violence, using ShotSpotter data as the main outcome measure, and find that curfews increase the number of gunfire incidents. That study would not have been possible with traditional crime measures, because curfews also affect reporting rates (more police scrutiny could increase reporting, while fewer witnesses out on the streets could decrease reporting). Indeed, when they conduct the same analysis using reported crime and 911 call data, the estimates suggest the opposite (and incorrect) conclusion: crime reports and 911 calls fall when curfews are in effect.

ShotSpotter data provide an important alternative to homicide as an outcome measure that is not affected by underreporting. While we would all like to see fewer homicides, current murder rates are too low to detect

policy effects in many contexts. In the study of juvenile curfews, for instance, there were only three homicides during the sample period, but hundreds of gunfire incidents. ShotSpotter data provided sufficient statistical power to detect an effect, while homicides did not.

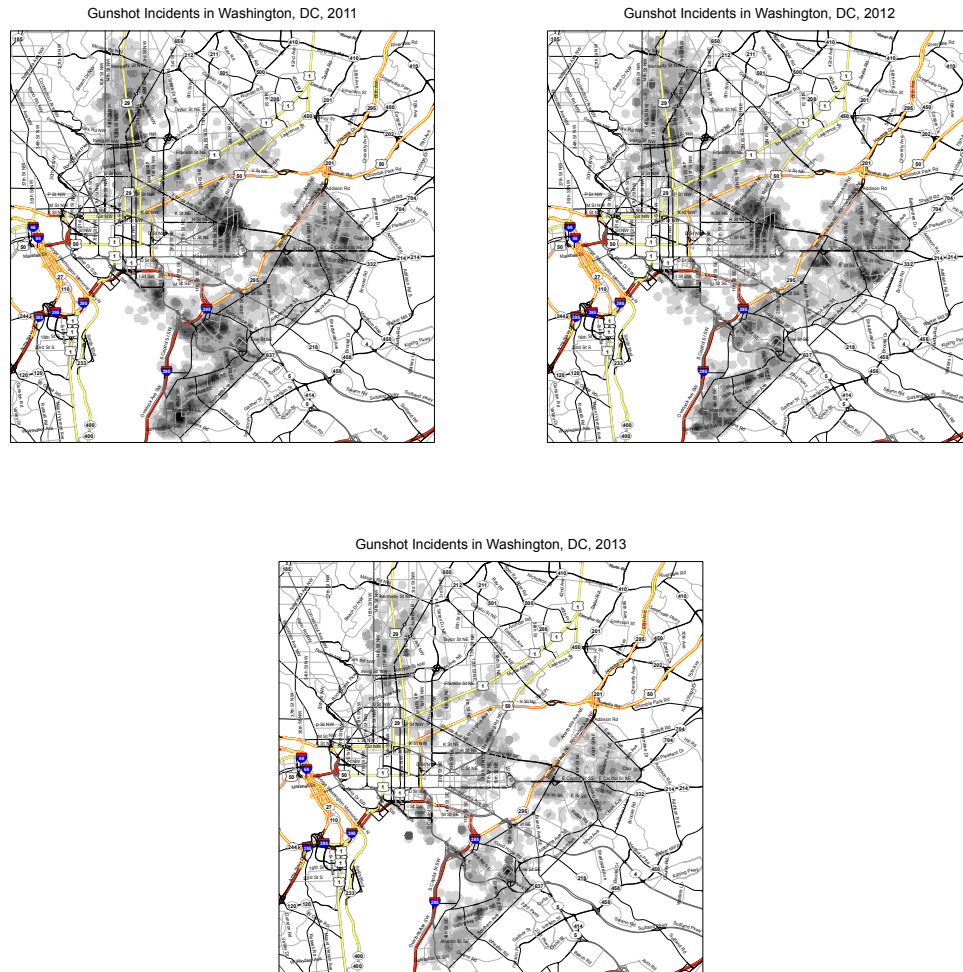
As technology improves, a wide variety of surveillance data should be extremely helpful to crime researchers and policymakers, because – like the data from ShotSpotter’s audio sensors – they do not rely on reporting by victims, witnesses, or the police. Increasing our use of such data will lead to a better understanding of crime patterns, as well as more accurate empirical estimates of policy effects.

References

- Carr, J. B., and J. L. Doleac (2015): "Keep the Kids Inside? Juvenile Curfews and Urban Gun Violence," Working paper, available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2486903.
- Goode, E. (2012): "Shots Fired, Pinpointed and Argued Over," *New York Times*.
- Litch, M., and G. A. Orrison (2011): "Draft Technical Report for SECURES Demonstration in Hampton and Newport News, Virginia," *US DOJ report*.
- Mazerolle, L. G., J. Frank, D. Rogan, and C. Watkins (2000): "Field Evaluation of the ShotSpotter Gunshot Location System: Final Report on the Redwood City Field Trial," *US DOJ report*.
- Pepper, J., C. Petrie, and S. Sullivan (2010): "Measurement Error in Criminal Justice Data," in *Handbook of Quantitative Criminology*, ed. by A. Piquero, and D. Weisburd. Springer.
- ShotSpotter (2015): "National Gunfire Index," Available at: <http://shotspotter.com/download-2015-ebook>.

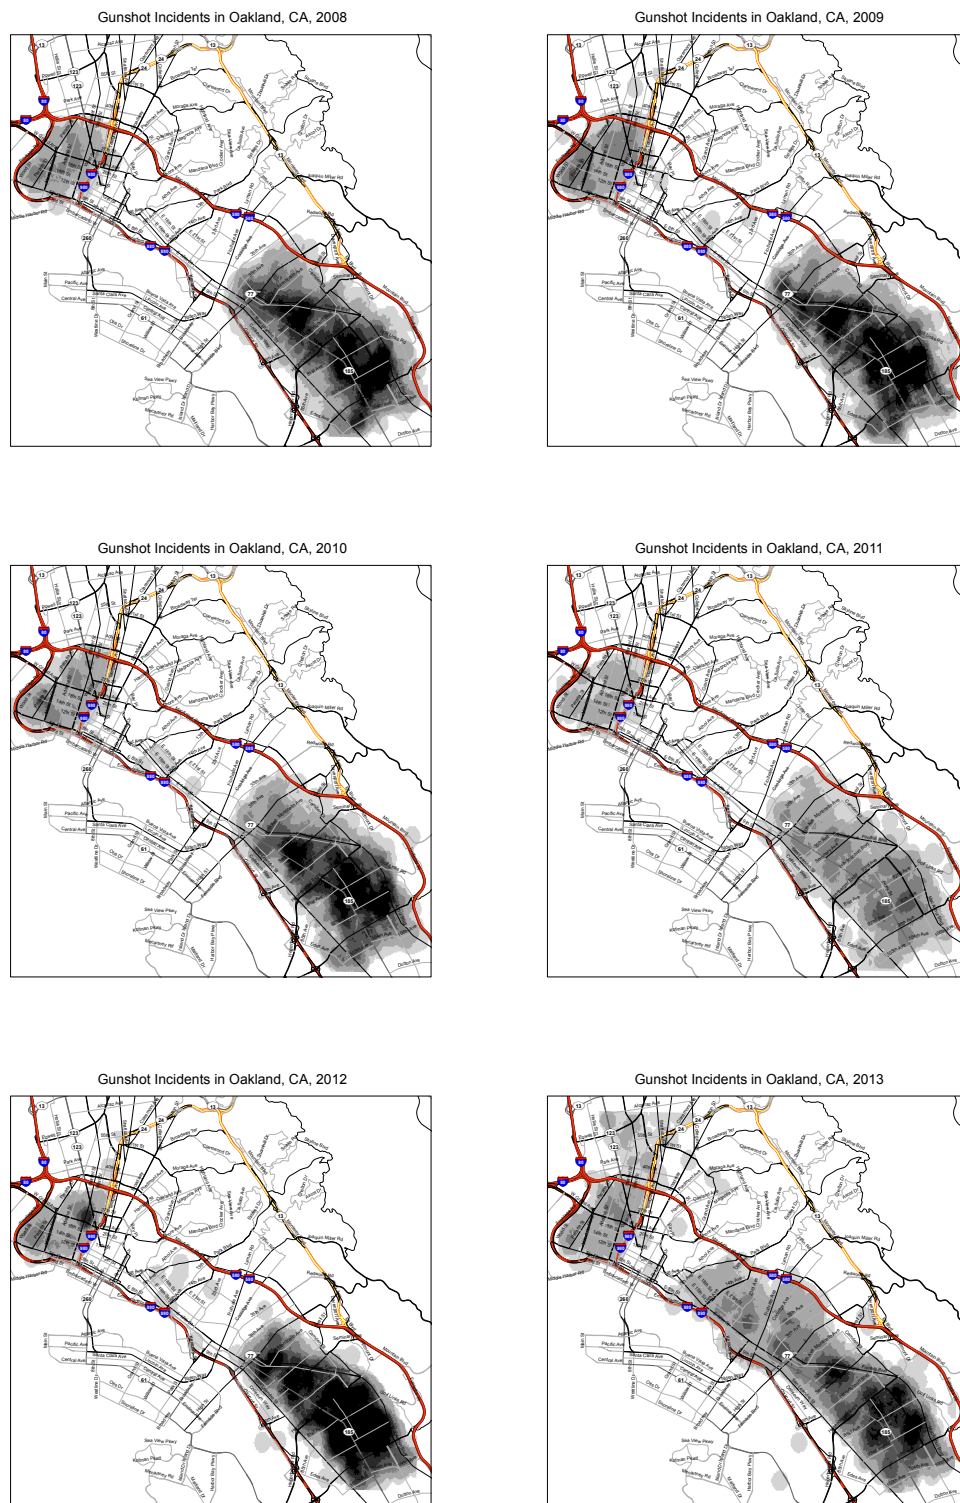
6 Figures and Tables

Figure 1: Washington, DC: Heatmaps of ShotSpotter-detected gunshot incidents, by year



Notes: Shaded regions show the location of detected gunfire in Washington, DC, in each year (January 2011 through June 2013). Darker regions signify more gunfire. Note that ShotSpotter sensors target Districts, 3, 5, 6, and 7.

Figure 2: Oakland, CA: Heatmaps of ShotSpotter-detected gunshot incidents, by year



Notes: Shaded regions show the location of detected gunfire in Oakland, CA, in each year (January 2008 through October 2013, excluding January through July of 2011). Darker regions signify more gunfire.

Table 1: Crime by Year in Washington, DC

	2011	2012	2013*
All Days			
SST-detected incidents	5196	3920	1514
Reported homicide	105	88	38
Reported AWDW	2179	2294	1073
911 calls	183955	193821	141963
911 calls for police	142675	149868	110205
911 calls reporting gunshots	1712	1685	1112
Excluding Outlier Days (NYE, week of July 4th)			
SST-detected incidents	2388	2234	898
Reported homicide	94	81	33
Reported AWDW	2015	2130	991
911 calls	170491	179334	131192
911 calls for police	132025	138704	101637
911 calls reporting gunshots	1522	1520	995

* 2013 data include January through June only.

Data source: DC Metropolitan Police Department.

Table 2: Crime by Year in Oakland, CA

	2008	2009	2010	2011*	2012	2013**
All Days						
SST-detected incidents	3260	3642	2852	1062	3622	2738
Reported homicide	201	202	192	91	218	124
Reported AWDW	1518	1007	1208	695	1497	1092
Excluding Outlier Days (NYE, week of July 4th)						
SST-detected incidents	2984	3110	2266	984	2968	2349
Reported homicide	180	190	182	90	203	113
Reported AWDW	1388	935	1153	689	1412	988

* 2011 data include August through December only

** 2013 data include January through September only

Data source: Open Oakland.

Table 3: Effects of gunfire incidents on traditional crime measures in Washington, DC

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Homicide							
Gunshots	0.0055*** (0.0011)	0.0054*** (0.0011)	0.0048*** (0.0011)	0.0048*** (0.0011)	0.0048*** (0.0011)	0.0048*** (0.0011)	0.0048*** (0.0011)
Observations	85056						
AWDW							
Gunshots	0.0339*** (0.0033)	0.0303*** (0.0033)	0.0236*** (0.0032)	0.0237*** (0.0032)	0.0233*** (0.0032)	0.0233*** (0.0032)	0.0233*** (0.0032)
Observations	85056						
911 Call							
Gunshots	0.6442*** (0.0410)	0.5497*** (0.0402)	0.2999*** (0.0344)	0.2901*** (0.0338)	0.2552*** (0.0314)	0.2191*** (0.0297)	0.2204*** (0.0305)
Observations	85056						
911 Call for Police							
Gunshots	0.5690*** (0.0352)	0.5219*** (0.0349)	0.2741*** (0.0296)	0.2663*** (0.0291)	0.2360*** (0.0268)	0.2025*** (0.0255)	0.2072*** (0.0261)
Observations	85056						
911 Call to Report Gunshots							
Gunshots	0.1427*** (0.0074)	0.1400*** (0.0074)	0.1236*** (0.0073)	0.1234*** (0.0073)	0.1229*** (0.0073)	0.1221*** (0.0072)	0.1242*** (0.0074)
Observations	85056						
Controls:							
Police District FE		X	X	X	X	X	X
Hour of Day FE			X	X	X	X	X
Year FE				X	X	X	X
Month of Year FE					X	X	X
Lagged Outcome (t - 1 hour)						X	
Lagged Outcome (t - 24 hours)							X

* $p < .10$, ** $p < .05$, *** $p < .01$. Data source: DC Metropolitan Police Department. Standard errors clustered by date.

Table 4: Effects of gunfire incidents on traditional crime measures in Washington, DC, by Police District

	District 3	District 5	District 6	District 7
Homicide				
Gunshots	0.0080** (0.0041)	0.0066* (0.0034)	0.0046*** (0.0017)	0.0035** (0.0016)
Observations	21264			
AWDW				
Gunshots	0.0271*** (0.0089)	0.0389*** (0.0088)	0.0221*** (0.0052)	0.0181*** (0.0051)
Observations	21264			
911 Call				
Gunshots	0.1343 (0.0937)	0.4826*** (0.0917)	0.1326*** (0.0425)	0.2683*** (0.0470)
Observations	21264			
911 Call for Police				
Gunshots	0.1458* (0.0857)	0.4034*** (0.0771)	0.1356*** (0.0367)	0.2564*** (0.0419)
Observations	21264			
911 Call to Report Gunshots				
Gunshots	0.0933*** (0.0202)	0.1799*** (0.0199)	0.1060*** (0.0092)	0.1249*** (0.0135)
Observations	21264			
Controls:				
Police District FE	X	X	X	X
Hour of Day FE	X	X	X	X
Year FE	X	X	X	X
Month of Year FE	X	X	X	X
Lagged Outcome (t - 24 hours)	X	X	X	X

* $p < .10$, ** $p < .05$, *** $p < .01$. Data source: DC Metropolitan Police Department. Standard errors clustered by date.

Table 5: Effects of gunfire incidents on traditional crime measures in Oakland, CA

	(1)	(2)	(3)	(4)	(5)	(6)
Homicide						
Gunshots	0.0162*** (0.0022)	0.0099*** (0.0022)	0.0100*** (0.0022)	0.0101*** (0.0022)	0.0100*** (0.0022)	0.0100*** (0.0022)
Observations	44376					
AWDW						
Gunshots	0.1034*** (0.0057)	0.0642*** (0.0059)	0.0648*** (0.0059)	0.0646*** (0.0059)	0.0642*** (0.0058)	0.0637*** (0.0059)
Observations	44376	44376	44376	44376	44375	44352
Controls:						
Hour of Day FE		X	X	X	X	X
Year FE			X	X	X	X
Month FE				X	X	X
Lagged Outcome (t - 1 hour)					X	
Lagged Outcome (t - 24 hours)						X

* $p < .10$, ** $p < .05$, *** $p < .01$. Data source: Open Oakland. Standard errors clustered by date.