RESEARCH ARTICLE

Journal of Sociotechnical Inquiry
Volume 1, Issue 1, 2026, pp. 42–50

What Follows a ShotSpotter Alert? Gunshot Detection, Recorded Gun Crime and Neighborhood Exposure in Chicago from 2017 to 2024

Mariam El-Ghazaly[1,*]

[1]Department of Statistics & Insurance, Menoufia University, Egypt

ABSTRACT

Acoustic gunshot detection is judged by how many shots it detects and how fast it sends police. Residents meet it as police cars arriving on their block. This study compares the two views for Chicago's ShotSpotter system. It links all 222,108 alerts sent between January 2017 and September 2024 to gun-related crime records, shooting victimizations and arrest flags within 500 metres and 60 minutes. A placebo that moves each alert one week forward or back estimates how often such records occur at the same place by chance. About 18.3% of alerts had a nearby gun-related crime record, against 2.0% under the placebo, so the excess linked to alerts was 16.4 percentage points. Only 6.4% had a nearby shooting and 2.5% a gun-crime record marked with an arrest. Alerts the system labelled as a single gunshot matched a record 7.8% of the time, against 25.7% for multiple gunshots. Alerts were placed and concentrated by design: 99.2% fell in 12 of 22 police districts, home to 49.8% of residents but 78.4% of Black residents and 15.4% of white residents. The average Black resident lived in a beat with 11.8 times as many alerts per square kilometre as the average white resident, while the gap in recorded shootings was 4.4 times. The design is correlational and cannot verify each alert. The results show that a system counted in detections produced, on the ground, mostly police responses that left no gun-crime record, and that these responses fell mainly on Black and Latino neighborhoods.

Keywords: gunshot detection, ShotSpotter, policing technology, Chicago, algorithmic alerts, racial disparities

1. Introduction

Acoustic gunshot detection systems listen for gunfire through microphones mounted on buildings and poles. When several sensors pick up a loud impulsive sound, software estimates its location and classifies it as possible gunfire. In Chicago, a ShotSpotter analyst then reviewed the alert and either dismissed it or passed it to the Chicago Police Department (CPD), which sent officers to the location (City of Chicago, 2024). Chicago used the system from 2017 until 22 September 2024, when the city ended its use. In those years, the system sent more than 220,000 alerts to police.

The vendor and many police agencies describe the system in terms of what it detects. Its key measures are the number of gunfire incidents found, the share not reported by 911 callers, and the minutes saved between a shot and police arrival (Mazerolle et al., 1998; Choi et al., 2014). On these measures the system looks productive, because each alert is a detection and each detection is a dispatch. The people who live near the sensors meet the system in a different form. For them, an alert is a police car that arrives quickly, often at night, looking for a shooter or shell casings. Whether that visit finds anything is not part of the detection count.

In 2021, Chicago's Office of Inspector General (OIG) examined that second view. Using police dispatch records for January 2020 to May 2021, it found that 4,556 of 50,176 CPD responses to ShotSpotter alerts (9.1%) ended with a disposition indicating a gun-related crime (Chicago Office of Inspector General, 2021). The OIG report covered 17 months and used internal dispatch data that the public cannot access. It did not study where the alerts fell or whom they affected.

This article extends that question to the full life of the Chicago program and to its geography, using only public data. It treats the alert as a machine classification that turns a sound into a police deployment. The study asks three research questions. RQ1: How many alerts did the system send, and how were they distributed across police beats compared with recorded shootings and residents? RQ2: How often was an alert followed by an official record of a gun-related crime, a shooting or an arrest near the same place and time, beyond what the place alone would predict, and does this differ by the system's own alert labels? RQ3: Who lived where the alerts were concentrated, and was the racial gap in alert exposure larger than the racial gap in recorded shootings?

The answers show a gap between the two ways of seeing the system. Measured by detections, the system was busy and growing: alerts rose from 15,110 in 2018 to 43,503 in 2023. Measured by what police later recorded nearby, about four in five alerts left no trace of a gun-related crime within 500 metres and an hour. The deployments were concentrated by design in majority Black and Latino districts. The article argues that this gap between detection counts and recorded outcomes is a property of the whole sociotechnical system rather than of the sensors alone, and that it should shape how such systems are evaluated.

2. Related work and framing

2.1 Technology in policing and how it is evaluated

Research on data-driven policing shows that new tools change what police notice and where they go. Brayne (2017) studied the Los Angeles Police Department and found that big data tools widened the range of people under routine surveillance and moved discretion into software and databases. Her later book traced how such tools shaped daily police routines (Brayne, 2020). Ferguson (2017) argued that data-driven policing can carry old patterns of suspicion into new technical forms, with race shaping where the data come from. Lum and Isaac (2016) showed that a predictive policing algorithm trained on Oakland drug arrests would send officers mainly to Black neighborhoods, even though survey data suggested drug use was spread across the city.

Evaluations of predictive policing have mostly asked whether the tools reduce crime. A field trial in Los Angeles and Kent found that algorithmic forecasts predicted crime better than analysts did (Mohler et al., 2015). A later trial found no clear evidence that predictive deployments raised arrest rates for minority groups (Brantingham et al., 2018). In Chicago, an earlier person-based prediction pilot did not reduce gun violence among those on the list, though listed people were more likely to be arrested for a shooting (Saunders et al., 2016). C. Lum et al. (2017) found that police tend to adopt technology through existing reactive routines, which limits what it can change. These studies share a point: the measures used to judge a tool are chosen by those who build or buy it, and they can miss what the tool does to the people it is used on.

2.2 Acoustic gunshot detection

Studies of gunshot detection have followed two lines. The first asks whether the technology speeds up police response. An early Dallas trial found response times fell by about one minute but officer workload rose (Mazerolle et al., 1998). A study in a Massachusetts city found faster dispatch and response but no clear gain in case resolution (Choi et al., 2014). A recent study of Denver, Milwaukee and Richmond, California, found that agencies were alerted to about twice as many shooting incidents after adoption, and that staff time spent on these calls rose two- to threefold (Lawrence & Novak, 2025).

The second line asks whether the technology reduces gun violence or improves case outcomes. In St. Louis, a time-series study found fewer citizen calls about shots fired after installation but no drop in reported gun crime (Mares & Blackburn, 2012). A later panel study of the same city found little evidence of a lasting effect on serious violence (Mares & Blackburn, 2021). In Philadelphia, a partially randomized trial found that acoustic sensors raised recorded gunshot incidents by 259% but did not raise confirmed shootings, so officers spent more time at scenes with no evidence of a shooting (Ratcliffe et al., 2019). In Kansas City, shots-fired calls in the sensor area were more likely to be coded as unfounded, and the system did not improve evidence collection or clearance (Piza et al., 2025). A study of large United States counties found no significant effect of adoption on firearm homicides or arrests (Doucette et al., 2021).

Two studies speak directly to Chicago. Topper and Ferrazares (2026) used the staggered rollout across CPD districts and found that, once the system arrived, officers answered unrelated 911 calls more slowly and those calls were less likely to end in an arrest. Their work shows that alerts consume police time that would otherwise go elsewhere. The OIG (Chicago Office of Inspector General, 2021) found that most responses ended without evidence of a gun crime and that officers sometimes cited the frequency of alerts in an area as a reason for stopping people. Economists have also used ShotSpotter data as a better measure of gunfire than 911 calls, because many shots are never reported (Carr & Doleac, 2018). That use treats each alert as true gunfire, which is the assumption the OIG questioned.

2.3 Classification, measures and place

Three ideas frame this study. First, a classification system makes some things visible and countable and leaves others out (Bowker & Star, 1999). An alert labelled 'multiple gunshots' is a category produced by software and an analyst, and the label then travels into police records as a reason to act. Second, public measures change the work they describe. Espeland and Sauder (2007) called this reactivity: once a number such as detections is used to judge a program, people organize around producing that number. Third, algorithmic systems must be studied as part of the social settings in which they act, not as stand-alone models (Selbst et al., 2019; Kitchin, 2017).

Place matters as well. Gun violence clusters in a small share of street segments and neighborhoods (Braga et al., 2010; Weisburd, 2015). Placing sensors in high-violence districts follows that pattern. The same choice also places the system's false or unconfirmed alerts in those districts. Research on police stops shows that the burden of suspicion-based contact falls unevenly by race, even after accounting for local crime (Gelman et al., 2007; Goel et al., 2016). Being stopped repeatedly shapes how residents see their standing as citizens (Epp et al., 2014). Distrust of police can in turn reduce calls for help in Black neighborhoods (Desmond et al., 2016). A detection system whose deployments land mainly in those neighborhoods enters this history.

This study builds on the OIG report in three ways. It covers the whole program from 2017 to 2024 rather than 17 months. It uses only public data, so others can repeat it. It adds the geography of alerts and compares it with the geography of recorded shootings and residents.

3. Data and methods

3.1 Data

The main source is the City of Chicago's 'Violence Reduction - Shotspotter Alerts - Historical' dataset (City of Chicago, 2024). It lists every alert sent to CPD from the first deployment in January 2017 to the end of the contract. The file downloaded in October 2026 holds 222,108 alerts, from 13 January 2017 to 22 September 2024. Each record gives the time, block, beat, district and community area of the alert, a shifted latitude and longitude, the number of rounds, and the system's alert type: 'single gunshot', 'multiple gunshots' or 'gunshot or firecracker'. The dataset has no field for what police found. The city notes that sensors were placed in 12 police districts (2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 15 and 25) and that alert counts may rise over time because detection improved, even when gunfire did not rise.

Outcome records come from two other public CPD datasets on the same portal. The first is 'Crimes - 2001 to Present', which

holds one record per reported crime incident with an arrest flag (City of Chicago, n.d.). I kept records from December 2016 to October 2024 whose offense involves a gun. The selection follows the list of gun-related dispositions used by the OIG: homicide; armed robbery, aggravated battery and aggravated assault with a handgun or firearm; reckless firearm discharge; and unlawful use or possession of a firearm or ammunition. Gun offender registration offences and air rifles were excluded. This gave 143,331 records with coordinates. The second is 'Violence Reduction - Victims of Homicides and Non-Fatal Shootings'. I kept victims with a gunshot injury and merged victims of the same case into one shooting incident, which gave 21,636 incidents.

Neighborhood composition comes from the American Community Survey 2019–2023 five-year estimates at the census tract level (U.S. Census Bureau, 2024). I used counts of non-Hispanic white, non-Hispanic Black and Hispanic or Latino residents, people below the poverty line, and median household income. Tract counts were moved to the 275 police beats in the city's December 2012 beat map by area-weighted interpolation. This assumes that residents are spread evenly within each tract, which is a simplification. Beat median income is the population-weighted mean of tract medians.

The alert itself is the object of study, so it helps to set out how it was made. According to the city's documentation, an alert began when at least three sensors registered a sound that the vendor's software judged to be possible gunfire (City of Chicago, 2024). The software located the sound by triangulating between sensors, the general method of acoustic gunshot location (Mares & Blackburn, 2012). A human analyst at the vendor then listened to the recording and either dismissed it or released it to CPD with a type and a round count. Once released, the alert became an event in the city's dispatch system and a car was sent. The public dataset therefore contains only alerts that passed both the software and the analyst, and each one corresponds to a police deployment. It contains no record of the sounds that were dismissed, and no record of what the officers did or found.

3.2 Linking alerts to recorded outcomes

The public data contain no common key between alerts and crime records, unlike the dispatch event numbers the OIG used. I therefore linked records by place and time. Under the main rule, an alert counts as matched if at least one record lies within 500 metres and 60 minutes before or after it. The window runs both ways because crime records give the estimated time of the offense, which can fall before the alert is logged. The 500 metre radius allows for the shifting of coordinates that the city applies to protect privacy, which places records at block level.

Four outcomes were linked. The first is any gun-related crime record, which is the closest public analogue to the OIG measure. The second is a gunfire record: a shooting incident or a reckless firearm discharge, both of which mean a gun was fired. The third is a shooting incident alone. The fourth is a gun-related crime record whose arrest flag is true.

Some matches will be coincidence. Sensors sit in places with many shootings, so a gun crime may occur near an alert by chance. To estimate this, I repeated the match for each alert at the same place and clock time but seven days earlier and seven days later, and averaged the two. This placebo keeps the place, weekday and hour, so it captures the base rate of gun-crime records at that place and time of week. The excess is the observed match rate minus the placebo rate. Its 95% confidence interval uses standard errors clustered by beat.

The excess is an upper bound on what an alert itself led to, for two reasons. Many matched records would exist without the alert, because a shooting with a victim is usually reported by 911 callers or hospitals. A match also does not show that the responding officers found the evidence. It is also a lower bound on gunfire, because many shots without a victim or a report leave no crime record at all. The link therefore measures whether an alert coincided with something the police recorded as a gun crime, which is how the deployment was experienced on the official record. It does not measure whether a gun was fired.

I tested the rule with nine combinations of radius (250, 500 and 1,000 metres) and window (30, 60 and 120 minutes). I also repeated the main analysis after removing alerts that followed another alert within 10 minutes and 250 metres. These 33,607 alerts (15.1%) are likely to belong to the same event.

3.3 Concentration and exposure

For RQ1 and RQ3, I counted alerts, shooting incidents and gun-related crime records in each beat from 1 January 2019 to 22 September 2024. This 5.73-year period starts after all 12 sensor districts were active. I describe concentration with Lorenz curves, Gini coefficients and the share of each count held by the top 10% and 20% of beats. To measure exposure by group, I computed the population-weighted average of alerts per square kilometre per year in the beat where each group lives. The same measure was computed for recorded shootings, so the two gaps can be compared.

Within the 156 beats of the sensor districts, I fitted Poisson regressions of beat alert counts. The predictors were the log of shooting incidents plus one, the log of beat area, the percentage of Black and of Hispanic residents (per 10 points) and the log of median income. One model adds district fixed effects, which compare beats within the same district. A further model uses the number of matched alerts with the log of all alerts as an offset, so its coefficients describe the share of alerts that matched a gun-related crime record. All models report heteroskedasticity-consistent standard errors. These models are descriptive. Sensor locations are not public, so they cannot separate where sensors were placed from how often they fired.

All data were downloaded from public portals in October 2026. The analysis used R with the data.table, sf, fixest, ineq and ggplot2 packages. Every number reported below comes from that code.

4. Findings

4.1 How many alerts and where

The system grew fast. Alerts rose from 4,694 in 2017, when only some districts were covered, to 19,035 in 2019 and 40,399 in 2021. They peaked at 43,503 in 2023. In 2024 the system sent 30,284 alerts before it was switched off on 22 September. Figure 1 shows the monthly series. The tallest spikes fall in January 2022, January 2023 and January 2024, driven by New Year's Day. On 1 January and on 4 and 5 July, the system sent 11,146 alerts

in total, or 484.6 alerts per day, against 76.2 per day on all other days.

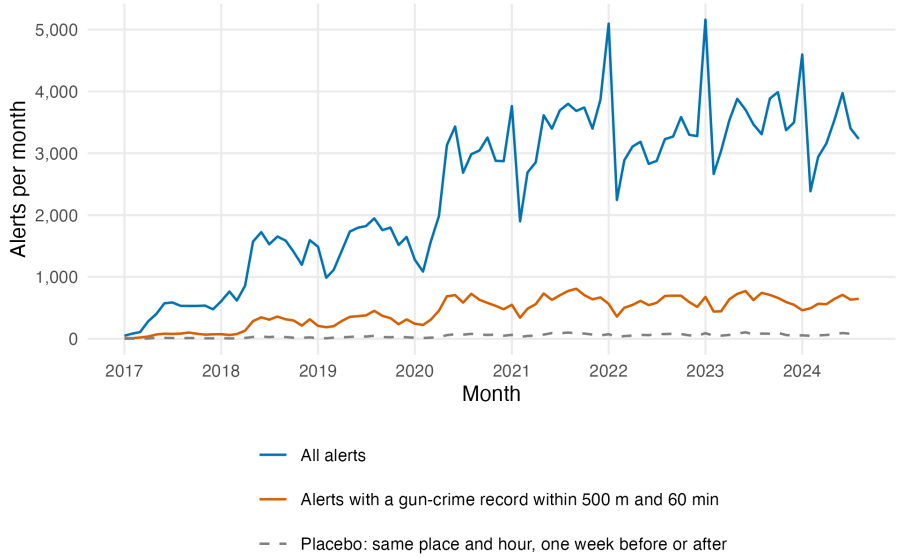

Figure 1. Monthly ShotSpotter alerts in Chicago, January 2017 to August 2024, with the number of alerts that had at least one gun-related crime record within 500 m and 60 minutes, and the same count under a placebo that shifts each alert by one week. September 2024 is omitted because the program ended on 22 September. Sources: City of Chicago Shotspotter Alerts and Crimes datasets.

The system's own labels show how it classified sounds. Of all alerts, 58.5% were labelled 'multiple gunshots', 32.9% 'single gunshot' and 8.5% 'gunshot or firecracker'. The median alert reported two rounds, with an interquartile range of one to five. The third label is notable: about one in twelve alerts carried, in its own label, the doubt that it was gunfire at all, and each still sent officers.

Alerts were concentrated by design. Of the 222,108 alerts, 99.2% fell in the 12 sensor districts; the rest fell just across district lines. Between 2019 and 2024, 89 of the city's 275 beats recorded no alert at all. Figure 2A maps annual alerts per square kilometre by beat. Of the 28 beats with the highest alert density, seven were in district 7, which covers Englewood, and three each were in districts 6, 9, 10 and 11.

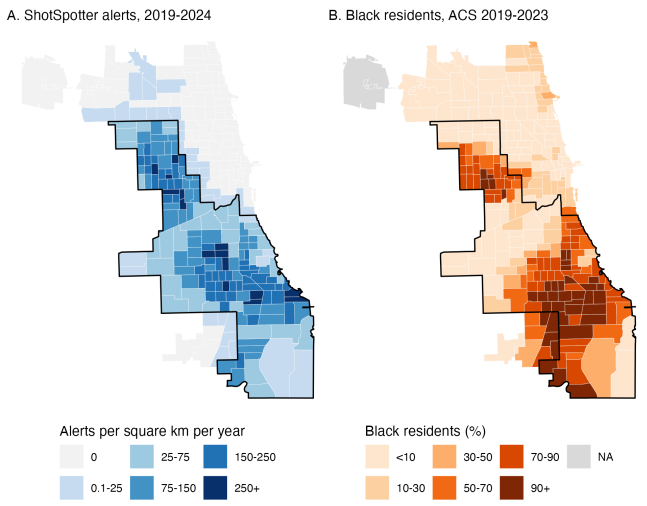

Figure 2. (A) ShotSpotter alerts per square kilometre per year by police beat, January 2019 to September 2024. (B) Percentage of residents who are non-Hispanic Black, by beat, from area-weighted ACS 2019–2023 tract estimates. The black outline marks the 12 police districts with sensors. Grey beats in panel B have no resident population. Sources: City of Chicago; U.S. Census Bureau.

Figure 3 compares the distribution of alerts across all 275 beats with the distribution of recorded shootings, gun-related crime records and residents. Alerts were the most concentrated of the four. The Gini coefficient was 0.58 for alerts, 0.46 for shooting incidents, 0.37 for gun-related crime records and 0.31 for residents. The top 10% of beats held 31.8% of alerts and 27.1% of shooting incidents. The top 20% held 53.9% of alerts and 46.1% of shootings.

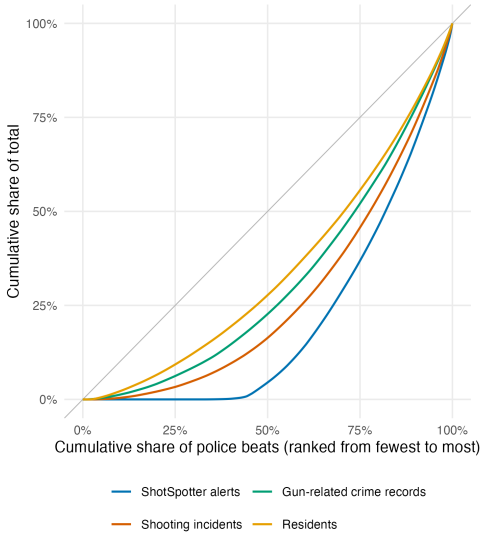

Figure 3. Lorenz curves of ShotSpotter alerts, shooting incidents, gun-related crime records and residents across Chicago's 275 police beats, January 2019 to September 2024. The diagonal shows an equal distribution. Sources: City of Chicago; U.S. Census Bureau.

The extra concentration of alerts came from the choice of where to place sensors, not from how alerts spread within the covered area. The sensor districts held 49.8% of residents and 56.8% of the city's area, but 83.6% of shooting incidents and 79.6% of gun-related crime records. Within the 156 sensor-district beats, alerts were spread about as evenly as shootings: the Gini coefficient was 0.27 for alerts and 0.29 for shootings. Beat alert counts in these districts were closely related to shooting counts (Pearson r = 0.64; Spearman ρ = 0.70). On average, the sensor beats produced 15.0 alerts for each recorded shooting incident.

4.2 What alerts were followed by

Table 1 reports the main linkage results. Of all alerts, 18.3% had at least one gun-related crime record within 500 metres and 60 minutes. Under the placebo, the same place and hour one week earlier or later, the rate was 2.0%. The excess linked to alert times was therefore 16.4 percentage points (95% CI 15.9 to 16.9). Put the other way, 81.7% of alerts had no gun-related crime record anywhere within 500 metres in the hour before or after.

Table 1. Share of ShotSpotter alerts with an official record nearby (within 500 m and 60 minutes), observed and under a one-week placebo, in Chicago from 2017 to 2024

Alert group	Alerts (n)	Any gun-related crime record, observed %	Placebo %	Excess, points (95% CI)	Gunfire record %	Shooting incident %	Gun-crime record with arrest %
All alerts	222,108	18.3	2.0	16.4 (15.9 to 16.9)	12.0	6.4	2.5
First alert in a 10-minute, 250 m cluster	188,501	16.1	1.9	14.2 (13.8 to 14.6)	10.3	5.3	2.2
OIG window, January 2020 to May 2021	45,017	19.6	1.9	17.7 (17.0 to 18.4)	12.6	6.9	3.2
Labelled multiple gunshots	130,009	25.7	2.1	23.7 (23.1 to 24.3)	17.5	9.5	3.0
Labelled single gunshot	73,180	7.8	1.8	5.9 (5.6 to 6.2)	3.9	1.8	1.8
Labelled gunshot or firecracker	18,917	8.1	1.6	6.5 (6.0 to 7.0)	4.6	2.8	1.7
1 to 7 July	4,828	14.7	1.9	12.8 (11.2 to 14.4)	10.0	5.6	2.6

Narrower outcomes were rarer. A gunfire record, meaning a shooting incident or a reckless firearm discharge report, lay near 12.0% of alerts (placebo 0.6%). A shooting incident with a victim lay near 6.4% (placebo 0.4%). A gun-related crime record flagged with an arrest lay near 2.5% of alerts (placebo 0.6%), so the excess for arrests was 1.9 points (95% CI 1.8 to 2.0). Removing likely duplicate alerts lowered the gun-crime match rate to 16.1%. Within the OIG's 17-month window, the public data hold 45,017 alerts, fewer than the 50,176 dispatches the OIG counted, and 19.6% of them matched a gun-related crime record.

The result did not depend much on the matching rule. Across the nine combinations of radius and window, the observed rate ranged from 15.0% (250 metres, 30 minutes) to 28.5% (1,000 metres, 120 minutes). The placebo rate rose with the size of the window, from 0.3% to 12.0%, so the excess stayed between 14.7 and 16.8 points. In other words, widening the net mostly added coincidental records. Even under the widest rule, 71.5% of alerts had no gun-related crime record nearby.

The rate also changed little over time, while alert volume nearly tripled. The excess was 17.7 points in 2019, 18.4 in 2020, 15.9 in 2022 and 15.4 in both 2023 and 2024. The share of alerts near a shooting incident fell from 7.8% in 2019 to 5.0% in 2024. More alerts therefore did not mean more alerts that matched a recorded shooting. Figure 1 shows the gap. From January 2020 to August 2024, matched alerts ranged from 224 to 809 a month, while total alerts ranged from 1,086 to 5,161.

The system's own labels sorted alerts by how often they matched, as Figure 4 shows. Alerts labelled 'multiple gunshots' matched a gun-related crime record 25.7% of the time and a shooting 9.5% of the time. Alerts labelled 'single gunshot' matched 7.8% and 1.8%. Alerts labelled 'gunshot or firecracker' matched 8.1% and 2.8%. The two weaker labels made up 41.4% of alerts, and under the main rule more than nine in ten of them had no nearby gun-crime record. Holiday alerts tell the same story. On New Year's Day and 4 and 5 July, only 7.6% of alerts matched a gun-related crime record, against 1.3% under the placebo.

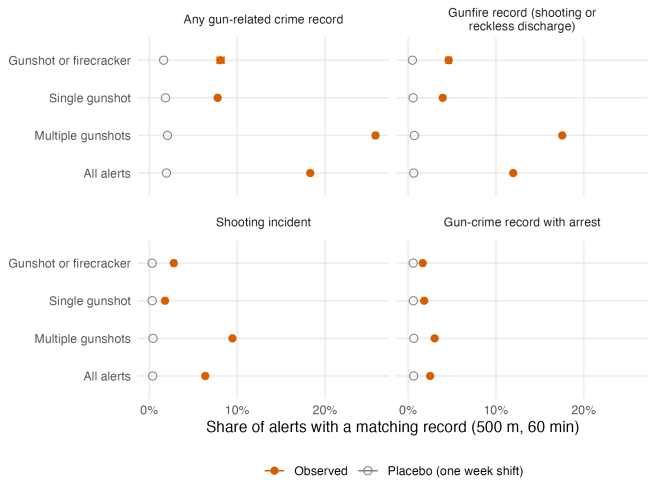

Figure 4. Share of ShotSpotter alerts with a matching record within 500 m and 60 minutes, by the system's alert label, for four outcomes. Filled points are observed shares with 95% Wilson confidence intervals (most are narrower than the point); open points are the one-week placebo. Sources: City of Chicago.

The matched records also show what kind of police record followed. The most common was a reckless firearm discharge report, which lay near 12,660 alerts (5.7% of all alerts). Aggravated battery with a gun, which is how CPD records most non-fatal shootings, lay near 10,868 alerts (4.9%). A homicide record lay near 2,396 alerts (1.1%). Across the 29,177 distinct gun-related crime records matched to any alert, 13.5% carried an arrest flag. Most of those arrests came from possession offences, which are usually recorded because an arrest was made: 89.4% of matched possession records carried an arrest flag. Only 2.7% of matched reckless discharge records and 4.6% of matched aggravated batteries with a gun did.

Seen from the other side, the system did detect many shootings. Of the 17,638 shooting incidents recorded in the sensor districts while the program ran, 51.6% had an alert within 500 metres and 60 minutes. From 2019, when coverage was complete, the share was 58.9%. Detection of real shootings and the low yield of alerts are two sides of the same pattern: the system flagged many shootings but also sent many more alerts that matched nothing in the police record.

4.3 Who lived where the alerts were

The sensor districts covered very different shares of each racial group. Of Chicago's 736,225 non-Hispanic Black residents, 78.4% lived in a sensor district. So did 65.9% of 794,991 Hispanic or Latino residents and 15.4% of 829,771 non-Hispanic white residents. Figure 2B shows how closely the sensor boundary follows the city's majority Black areas on the South and West Sides.

Exposure to alerts followed the same lines. The average Black resident lived in a beat with 122.9 alerts per square kilometre per year. The figure was 63.0 for the average Hispanic resident and 10.4 for the average white resident. The Black-to-white ratio was therefore 11.8. For recorded shooting incidents the same measure was 10.0 per square kilometre per year for Black residents and 2.2 for white residents, a ratio of 4.4. The racial gap in alert exposure was thus close to three times the racial gap in recorded shootings.

Table 2 groups beats into fifths by alert density. The fifth with the most alerts held 43.2% of all alerts. Its residents were 76.5% Black and 2.5% white, 30.2% lived below the poverty line, and the mean of tract median household incomes was $41,949. The fifth with no alerts was 59.9% white, with 9.2% in poverty and a mean tract median income of $119,633. Even that fifth recorded 845 shooting incidents over the period, none of which could produce an alert.

Table 2. Police beats grouped into fifths by ShotSpotter alert density, with population composition, recorded shootings and alert yield, January 2019 to September 2024

Fifth of beats by alert density	Alerts	Shooting incidents	Residents	Black %	Hispanic %	White %	Below poverty %	Mean tract median income ($)	Alerts per shooting	Alerts matched to a gun-crime record %
1 (lowest)	0	845	548,132	6.5	18.0	59.9	9.2	119,633	0.0	n/a
2	335	1,157	694,524	12.5	21.6	52.0	12.3	94,351	0.3	18.8
3	48,002	3,168	610,661	27.3	45.9	17.5	17.3	69,498	15.2	17.1
4	66,315	4,455	457,236	45.8	45.3	5.6	23.1	52,754	14.9	18.6
5 (highest)	87,181	6,353	311,980	76.5	18.7	2.5	30.2	41,949	13.7	19.0

The yield of alerts did not differ much across the covered fifths. In the third to fifth groups, between 17.1% and 19.0% of alerts matched a gun-related crime record, and there were 13.7 to 15.2 alerts per recorded shooting. Residents of the most exposed beats therefore did not receive more productive alerts. They received more alerts at about the same low yield.

The beat models give the same picture within the sensor districts. Without demographic terms, a 1% higher count of shootings went with about 0.68% more alerts (coefficient 0.68, SE 0.05). With district fixed effects and demographic terms, the coefficient on shootings was 0.60 (SE 0.06). In the same model, each 10-point higher share of Black residents went with 11% more alerts (incidence rate ratio 1.11, 95% CI 1.06 to 1.17), and each 10-point higher share of Hispanic residents with 13% more (1.13, 1.08 to 1.18). Income had no clear association (1.10 per unit of log income, 0.93 to 1.31). In the model for the share of alerts that matched a gun-crime record, racial composition had no association (rate ratio for Black share 1.00, 0.97 to 1.04), while beats with more shootings had higher yields (1.23 per unit of log shootings, 1.14 to 1.34). These are associations across 156 beats. Without sensor locations, they cannot say whether the extra alerts in Black and Latino beats reflect denser sensors, more unrecorded gunfire, or more non-gunfire sounds classified as gunfire.

5. Discussion

The findings show two accounts of the same system that do not meet. In the first, the system is judged by its output. It sent 222,108 alerts, more each year until 2023, and alerts came near about half of the shootings recorded in the covered districts. In the second, the system is judged by what its deployments left in the police record. About four in five alerts had no gun-related crime record within 500 metres and an hour, about 94% had no nearby shooting, and about 97.5% had no nearby gun-crime record marked with an arrest. This is the kind of gap the journal's theme describes: the system was evaluated on detections sent and experienced as police visits that mostly found nothing reportable.

The public-data estimates agree with the OIG's internal analysis

while extending it. The OIG found that 9.1% of responses produced a gun-related disposition in the dispatch record. In the same 17 months, 19.6% of public alerts had any gun-related crime record nearby, with an excess over the placebo of 17.7 points. The public figure is higher because it counts any nearby record, including crimes reported by 911 callers that the alert response did not produce. The two measures bracket the outcome: a gun-crime record that police tied to the alert in about one in ten responses, and some nearby gun-crime record in fewer than one in five. Over the full eight years, neither share rose as volume rose.

The alert label is where the technology is most visible in these data. The system sorted sounds into three classes, and those classes differed sharply in what followed. Multiple-gunshot alerts matched a gun crime record more than three times as often as single-gunshot alerts. Yet all three labels produced the same action: a dispatch. A label that says 'gunshot or firecracker' records the classifier's own doubt, but in practice that doubt did not change the response. In the language of Bowker and Star (1999), the category made the sound actionable while hiding the uncertainty inside it. The holiday spikes show the same thing at scale. On three days a year the system sent more than six times its usual daily volume, mostly to blocks where no gun crime was later recorded.

These alerts may have been true gunfire. Celebratory shooting on New Year's Eve is real, and many shots without victims are never reported. Ratcliffe et al. (2019) and Carr and Doleac (2018) both show that acoustic sensors find gunfire that 911 calls miss. The present data cannot settle whether a given unmatched alert was a gunshot. But that question is not the only one that matters for residents. Each alert sent officers to a block in an expectant state, looking for a shooter. Whether the sound was a gunshot nobody reported or a firework, the street-level result was the same: a rapid police presence that, in most cases, produced no crime record. The OIG found that officers sometimes used the general frequency of alerts in an area as part of their reason for stopping people (Chicago Office of Inspector General, 2021). The count of alerts thus became a feature of the place as well as a series of events.

The geography of those visits is the second part of the gap. Placing sensors where shootings are concentrated is a defensible choice, and the sensor districts did hold 83.6% of recorded shootings. But the placement turned a gap of 4.4 to 1 in recorded shooting exposure between Black and white residents into a gap of 11.8 to 1 in alert exposure. Part of this comes from area: the system covered all beats in its districts, including beats with few shootings, and none of the beats outside them, including the 845 shootings in the lowest fifth. Part comes from the yield of alerts. Each recorded shooting in the covered areas was accompanied by about 15 alerts, most of which matched no gun-crime record. Because yield did not differ by racial composition, the larger volume of unproductive deployments fell where the volume of alerts was larger, which was in Black and Latino beats.

This pattern fits the wider literature on data-driven policing. Lum and Isaac (2016) showed that a tool trained on police data sends officers back to the places already policed. Gunshot detection works differently, since it listens to the street rather than to police records. Its placement, however, was set by past shooting data, and its outputs then became new police data on those same places. Brayne (2017) described how data tools extend routine surveillance; the alert record here is one such extension, logged block by block. Topper and Ferrazares (2026) show that this volume of dispatches costs time on other 911 calls in the same districts. The combined result is a deployment pattern that intensifies police presence where residents already have the most contact with police, and where distrust can reduce reporting (Desmond et al., 2016).

The measures also shape how the system appears over time. Under a detection-centred view, the rise from 19,035 alerts in 2019 to 43,503 in 2023 reads as better coverage or more gunfire. The city's own documentation warns that improved sensors can raise counts without any change in gunfire (City of Chicago, 2024). Under an outcome-centred view, the same rise reads as more deployments at a steady or falling yield, since the share of alerts near a shooting fell from 7.8% to 5.1% over those years. Espeland and Sauder (2007) describe how a public number can come to stand in for the activity it was meant to summarize. Here, a count that records what the sensors heard came to stand in for what the police response achieved. Neither number is false. The problem is that only one of them was routinely published.

The findings suggest three practical points for evaluating such systems. First, a public evaluation should report outcomes per alert alongside alerts and response times. The city's own dataset contains no outcome field, which makes the detection count the default measure. Second, alerts with weaker labels could be handled differently from multiple-gunshot alerts, for example by review before dispatch, and evaluated separately. Third, any future deployment should publish where sensors are placed and the share of each group's residents covered, so that the burden of unconfirmed deployments can be weighed against the benefit of detected shootings. These points concern measurement. They do not by themselves say whether the system was worth keeping, which also depends on what its detected shootings led to for victims and cases.

6. Limitations

The study has clear limits. It cannot verify whether any alert was gunfire. A match measures coincidence with a police record, not the truth of the alert or the work of the responding officers. Some matched records would have existed without any alert, so the excess overstates what alerts produced; some real gunfire leaves no record, so the match rate understates gunfire.

The link between datasets is spatial and temporal, not a shared event number. Coordinates in both datasets are shifted to protect privacy, and offense times in crime records are sometimes estimates. The sensitivity analysis shows that the excess changes little across nine rules, but a different rule could still shift the levels. The public alert dataset also holds fewer alerts than the OIG counted for the same months, possibly because of differences in how alerts and dispatch events were defined, which the public documentation does not explain.

The demographic analysis uses area-weighted interpolation from tracts to beats, which assumes residents are spread evenly within

each tract. It describes the composition of places, not of the people who were present when officers arrived. Sensor locations are not public, so the beat models cannot separate sensor density from alert frequency. The study does not observe stops, searches or uses of force that followed alerts, which the OIG found are poorly recorded. All results are correlational.

7. Conclusion

Chicago's ShotSpotter program can be described as a detection system that sent 222,108 alerts in under eight years. It can also be described as a dispatch system whose deployments left a gun-related crime record nearby in fewer than one in five cases and an arrest-flagged gun-crime record in about one in forty. Both descriptions use the same alerts. The difference lies in what is counted. The first view counts the system's outputs; the second counts what followed them on the street.

The second view also shows where the costs fell. Alerts were placed in districts that held half the city's residents but nearly four in five of its Black residents. Within those districts, alerts tracked shootings, and the yield per alert was similar across neighborhoods. Because alerts were many and yields were low, the mostly unproductive deployments fell overwhelmingly on Black and Latino neighborhoods. Evaluations of gunshot detection, and of similar alert-based technologies, should report what each alert leads to and for whom, alongside how many alerts the system sends.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship and/or publication of this article.

Funding

This research received no external funding.

Ethics statement

This study is a secondary analysis of publicly available, de-identified administrative data. No human participants were contacted, and no ethics approval was required.

Data availability statement

All data are public and were accessed in October 2026. City of Chicago Data Portal: Violence Reduction - Shotspotter Alerts - Historical (https://data.cityofchicago.org/d/3h7q-7mdb); Crimes - 2001 to Present (https://data.cityofchicago.org/d/ijzp-q8t2); Violence Reduction - Victims of Homicides and Non-Fatal Shootings (https://data.cityofchicago.org/d/gumc-mgzr); PoliceBeatDec2012 boundaries (https://data.cityofchicago.org/d/n9it-hstw); PoliceDistrictDec2012 boundaries (https://data.cityofchicago.org/d/24zt-jpfn); Boundaries - Census Tracts - 2020 (https://data.cityofchicago.org/d/7qjk-677h). U.S. Census Bureau, American Community Survey 2019–2023 5-year table-based summary files, tables B03002, B17001 and B19013 (https://www2.census.gov/programs-surveys/acs/summary_file/2023/table-based-SF/). Chicago Office of Inspector General report (https://igchicago.org/wp-content/uploads/2023/08/Chicago-Police-Departments-Use-of-ShotSpotter-Technology.pdf). The analysis code will be deposited in a public repository on acceptance; it is withheld during review only to keep the review anonymous.

References

Bowker, G. C., & Star, S. L. (1999). Sorting things out: Classification and its consequences. MIT Press. https://doi.org/10.7551/mitpress/6352.001.0001

Braga, A. A., Papachristos, A. V., & Hureau, D. M. (2010). The concentration and stability of gun violence at micro places in Boston, 1980–2008. Journal of Quantitative Criminology, 26(1), 33–53. https://doi.org/10.1007/s10940-009-9082-x

Brantingham, P. J., Valasik, M., & Mohler, G. O. (2018). Does predictive policing lead to biased arrests? Results from a randomized controlled trial. Statistics and Public Policy, 5(1), 1–6. https://doi.org/10.1080/2330443X.2018.1438940

Brayne, S. (2017). Big data surveillance: The case of policing. American Sociological Review, 82(5), 977–1008. https://doi.org/10.1177/0003122417725865

Brayne, S. (2020). Predict and surveil: Data, discretion, and the future of policing. Oxford University Press. https://doi.org/10.1093/oso/9780190684099.001.0001

Carr, J. B., & Doleac, J. L. (2018). Keep the kids inside? Juvenile curfews and urban gun violence. Review of Economics and Statistics, 100(4), 609–618. https://doi.org/10.1162/rest_a_00720

Chicago Office of Inspector General. (2021). The Chicago Police Department's use of ShotSpotter technology (OIG File #21-0707). City of Chicago. https://igchicago.org/wp-content/uploads/2023/08/Chicago-Police-Departments-Use-of-ShotSpotter-Technology.pdf

Choi, K.-S., Librett, M., & Collins, T. J. (2014). An empirical evaluation: Gunshot detection system and its effectiveness on police practices. Police Practice and Research, 15(1), 48–61. https://doi.org/10.1080/15614263.2013.800671

City of Chicago. (2024). Violence reduction - Shotspotter alerts - Historical [Data set]. Chicago Data Portal. Retrieved October 2026, from https://data.cityofchicago.org/d/3h7q-7mdb

City of Chicago. (n.d.). Crimes - 2001 to present [Data set]. Chicago Data Portal. Retrieved October 2026, from https://data.cityofchicago.org/d/ijzp-q8t2

Desmond, M., Papachristos, A. V., & Kirk, D. S. (2016). Police violence and citizen crime reporting in the Black community. American Sociological Review, 81(5), 857–876. https://doi.org/10.1177/0003122416663494

Doucette, M. L., Green, C., Necci Dineen, J., Shapiro, D., & Raissian, K. M. (2021). Impact of ShotSpotter technology on firearm homicides and arrests among large metropolitan counties: A longitudinal analysis, 1999–2016. Journal of Urban Health, 98(5), 609–621. https://doi.org/10.1007/s11524-021-00515-4

Epp, C. R., Maynard-Moody, S., & Haider-Markel, D. P. (2014). Pulled over: How police stops define race and citizenship. University of Chicago Press. https://doi.org/10.7208/chicago/9780226114040.001.0001

Espeland, W. N., & Sauder, M. (2007). Rankings and reactivity: How public measures recreate social worlds. American Journal of Sociology, 113(1), 1–40. https://doi.org/10.1086/517897

Ferguson, A. G. (2017). The rise of big data policing: Surveillance, race, and the future of law enforcement. New York University Press. https://doi.org/10.18574/nyu/9781479854608.001.0001

Gelman, A., Fagan, J., & Kiss, A. (2007). An analysis of the New York City Police Department's "stop-and-frisk" policy in the context of claims of racial bias. Journal of the American Statistical Association, 102(479), 813–823. https://doi.org/10.1198/016214506000001040

Goel, S., Rao, J. M., & Shroff, R. (2016). Precinct or prejudice? Understanding racial disparities in New York City's stop-and-frisk policy. The Annals of Applied Statistics, 10(1), 365–394. https://doi.org/10.1214/15-AOAS897

Kitchin, R. (2017). Thinking critically about and researching algorithms. Information, Communication & Society, 20(1), 14–29. https://doi.org/10.1080/1369118X.2016.1154087

Lawrence, D. S., & Novak, K. J. (2025). Firearm shootings and the police response: Examining the impact of gunshot detection technology. Police Practice and Research, 26(6), 726–745. https://doi.org/10.1080/15614263.2024.2384879

Lum, C., Koper, C. S., & Willis, J. (2017). Understanding the limits of technology's impact on police effectiveness. Police Quarterly, 20(2), 135–163. https://doi.org/10.1177/1098611116667279

Lum, K., & Isaac, W. (2016). To predict and serve? Significance, 13(5), 14–19. https://doi.org/10.1111/j.1740-9713.2016.00960.x

Mares, D., & Blackburn, E. (2012). Evaluating the effectiveness of an acoustic gunshot location system in St. Louis, MO. Policing: A Journal of Policy and Practice, 6(1), 26–42. https://doi.org/10.1093/police/par056

Mares, D., & Blackburn, E. (2021). Acoustic gunshot detection systems: A quasi-experimental evaluation in St. Louis, MO. Journal of Experimental Criminology, 17(2), 193–215. https://doi.org/10.1007/s11292-019-09405-x

Mazerolle, L. G., Watkins, C., Rogan, D., & Frank, J. (1998). Using gunshot detection systems in police departments: The impact on police response times and officer workloads. Police Quarterly, 1(2), 21–49. https://doi.org/10.1177/109861119800100202

Mohler, G. O., Short, M. B., Malinowski, S., Johnson, M., Tita, G. E., Bertozzi, A. L., & Brantingham, P. J. (2015). Randomized controlled field trials of predictive policing. Journal of the American Statistical Association, 110(512), 1399–1411. https://doi.org/10.1080/01621459.2015.1077710

Piza, E. L., Arietti, R. A., Carter, J. G., & Mohler, G. O. (2025). The effect of gunshot detection technology on evidence collection and case clearance in Kansas City, Missouri. Journal of Experimental Criminology, 21(1), 127–152. https://doi.org/10.1007/s11292-023-09594-6

Ratcliffe, J. H., Lattanzio, M., Kikuchi, G., & Thomas, K. (2019). A partially randomized field experiment on the effect of an acoustic gunshot detection system on police incident reports. Journal of Experimental Criminology, 15(1), 67–76. https://doi.org/10.1007/s11292-018-9339-1

Saunders, J., Hunt, P., & Hollywood, J. S. (2016). Predictions put into practice: A quasi-experimental evaluation of Chicago's predictive policing pilot. Journal of Experimental Criminology, 12(3), 347–371. https://doi.org/10.1007/s11292-016-9272-0

Selbst, A. D., boyd, d., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In Proceedings of the Conference on Fairness, Accountability, and Transparency. Association for Computing Machinery. https://doi.org/10.1145/3287560.3287598

Topper, M., & Ferrazares, T. (2026). The unintended consequences of policing technology: Evidence from ShotSpotter. Journal of Human Resources. Advance online publication. https://doi.org/10.3368/jhr.1024-13929R2

U.S. Census Bureau. (2024). American Community Survey 2019–2023 5-year estimates, table-based summary file [Data set]. Retrieved October 2026, from https://www2.census.gov/programs-surveys/acs/summary_file/2023/table-based-SF/

Weisburd, D. (2015). The law of crime concentration and the criminology of place. Criminology, 53(2), 133–157. https://doi.org/10.1111/1745-9125.12070

Corresponding author: Mariam El-Ghazaly, Menoufia University. Email: mariam.elghazaly@gmail.com

Received: 2 August 2026. Accepted: 31 August 2026. Published: 7 September 2026.

© 2026 The Author(s). This is an open access article distributed under the terms of the CC BY 4.0 licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.