The Unintended Consequences of Policing Technology: Evidence from ShotSpotter*

Michael Topper and Toshio Ferrazares**

Draft Date: October 29, 2024

Latest Version Link

Abstract

Technology is integral to police departments, automating officer tasks, but inherently changing their time allocation. We investigate this by studying ShotSpotter, a technology that automates gunfire detection. Following a detection, officers are dispatched to the scene, thereby reallocating their time. We leverage this shock to officers' time allocation using the rollout of ShotSpotter across Chicago police districts to study the effects on 911 call response. We find substantial consequences—officers are dispatched to calls slower (22%), arrive on-scene later (13%), and the probability of arrest is decreased 9%. Consequently, police departments must evaluate their resource capacities prior to implementing technologies.

JEL Codes: D04, H40, R41

* Previously circulated under the title "The Effect of ShotSpotter Technology on Police Response Times" **Department of Economics, University of California, Santa Barbara. (Corresponding author: michaeltopper@ucsb.edu). We would like to thank Heather Royer, Kevin Schnepel, Dick Startz, Bocar Ba, Peter Kuhn, Cody Tuttle, members of the 2023 Summer WEAI Conference, and attendees of the UCSB AMEL 2023 for their feedback on iterations of this paper. In addition, we would like to thank Max Kapustin and Ashna Arora for assisting with details regarding rollout dates in Chicago. Last, we would also like to thank Matt Chapman, Joseph Mahr, Pascal Sabino, and the firm Loevy & Loevy for their aid in collecting the data.

1 Introduction

In the contemporary police department, technology possesses the potential to serve as either a substitute or complement to human capital. In particular, police departments are utilizing technologies both as substitutes, effectively functioning as 'eyes-on-the-street' through facial recognition and traffic cameras, as well as collaborative complements in targeting high-crime areas. These technologies are seen as imperative for public safety moving forward, addressing the issues of both officer shortages and eroding public opinion of the police (Gallup, 2022). Nevertheless, the integration of officers and technology systems is fundamentally reshaping the nature of policing.

One quickly expanding and widely adopted police technology is ShotSpotter—an acoustic gunfire detection technology that is currently implemented in over 150 cities world-wide. ShotSpotter's primary intention is to rapidly dispatch police officers to violent crime scenes with the goal of reducing gun violence. The technology utilizes an array of microphones and sensors placed on streetlights and buildings that use machine learning algorithms to detect the sound of gunfire, triangulate its location, and alert police officers for rapid response. Because of its unique functionality, ShotSpotter bypasses the reliance on civilian reporting. In effect, previous studies have utilized this feature of ShotSpotter as a measure of underlying crime that is independent of reporting habits (Carr and Doleac, 2016, 2018; Ang et al., 2021). Consequently, it has been estimated that only 12% of gunfire is reported, leaving a significant portion of these occurrences unattended (Carr and Doleac, 2016). Therefore, ShotSpotter offers a solution wherein police officers are dispatched to additional instances of gunfire. In Chicago, the setting of this paper, this results in approximately 70 ShotSpotter-related dispatches each day, equating to 75 hours of officer investigation time.[1] This represents a two-fold increase in the number of gunfire reports that require officers to engage in rapid response.[2]

However, reallocating resources to gunfire detection changes an officer's time allocation. On one hand, this reallocation could be beneficial—ShotSpotter may place officers closer to locations that foster higher volumes of crime or provide valuable crime-reduction. In this situation, an officer's time of arrival may be reduced and crime may be mitigated. On the other hand, these investigations of previously unreported gunfire may incapacitate officers from attending to reports of other crimes in the form of 911 calls—a lifeline for citizens in distress. Given that there is little evidence suggesting that ShotSpotter is a productive crime reduction tool (Mares and Blackburn, 2012; Choi et al., 2014; Connealy et al., 2024), this may have far-reaching implications given the critical importance of 911 response times which have been shown to alter the probability of crime clearance (Blanes i Vidal and Kirchmaier, 2018) and victim injury (DeAngelo et al., 2023).[3] Furthermore, response times may affect timely medical treatment, as emergency medical personnel are required to delay their services until police arrive if their safety is

[1] A ShotSpotter investigation takes roughly 20 minutes to complete. While we cannot delineate between the number of officers dispatched to the scene for our entire sample period, we find, using another source of data from 2019-2023, that the average number of officers dispatched to a ShotSpotter detection is approximately 3.35. On the other hand, a lower bound, assuming only one officer dispatched to each ShotSpotter alert, would result in 23 total hours.

[2] This statistic is based on the average number of 911 dispatches relating to a 'Shots Fired' report and the average number of ShotSpotter dispatches post-implementation in all police districts.

[3] As discussed in Section 6, we independently test this claim and similarly find little evidence of the technology being an effective gun-crime reduction tool. Moreover, in Section 7, we find descriptive evidence that a 911 call may be more productive than responding to a ShotSpotter alert. However, given the data limitations, we cannot truly verify whether ShotSpotter dispatches are more or less productive than a 911 dispatch.

compromised.[4] Thus, while ShotSpotter is implemented with the intention of enhancing public safety, it may have unintended consequences that are socially costly.

In this paper, we utilize variation in timing from the staggered rollout of ShotSpotter across Chicago police districts from 2016-2022 to estimate the causal effect of ShotSpotter technology on the response times from 911 calls designated as Priority 1—the most frequent call classification in Chicago which pertains to life-threatening and time-sensitive events. Using 911 call dispatch data from the Chicago Police Department (CPD), we construct two measures of rapid response: the time from a 911 call to when a dispatcher finds an available police officer for dispatch (Call-to-Dispatch) and the time from a 911 call to when the officer arrives on-scene (Call-to-On-Scene). By applying a staggered difference-in-differences framework, we find that the implementation of ShotSpotter significantly increases both Call-to-Dispatch time and Call-to-On-Scene time by approximately one-minute (22%) and two-minutes (13%), respectively. These estimates are robust to a variety of sensitivity tests and estimators.

Moreover, we find that the delays in response times are driven by resource-constrained periods, consistent with the hypothesis that ShotSpotter is affecting police officers' time constraints. We test this using days when there are fewer officers on-duty and times of day with higher numbers of ShotSpotter detections. Each of these subsets show larger effect sizes during these resource-constrained periods, suggesting that ShotSpotter forces officers to make time trade-offs in favor of responding to ShotSpotter alerts. Consistent with this mechanism, response times from other time-sensitive calls (Priority 2) are also increased, and in addition, time-

[4]This is found from the Chicago EMS System Policies and Procedures: `https://chicagoems.org/wp-content/uploads/sites/2/2017/08/2017-PP_APPROVED.pdf`

insensitive calls (Priority 3) show suggestive evidence of longer delays, providing further evidence of heightened officer responsibilities.[5]

Consequently, these elevated response times come at a significant cost. In Section 5.3, we analyze the relationship between police response time and the likelihood of an arrest using both reduced-form estimates and an instrumental variables (IV) framework, leveraging the ShotSpotter implementation as an instrument for 911 response times. The reduced-form estimates show that Priority 1 calls are 9% less likely to have the perpetrator arrested, and this result holds in the IV approach, which directly tests the impact of longer response times for ShotSpotter-induced delays (compliers). These findings align with Blanes i Vidal and Kirchmaier (2018) who attribute faster response times to higher crime clearance rates.

In addition to these unintended consequences, we also find little evidence that ShotSpotter reduces gun violence or enhances police productivity beyond the 911 call system. In Section 6.1, we broaden the analysis to examine city-wide, aggregate counts of gun-related arrests, clearance, and victimization that are not restricted to 911 incidents. Importantly, none of these outcomes change significantly post-implementation despite a marked increase in policing resources to gun-incidents. These results align with other case studies in Chicago which do not find ShotSpotter to be effective as a violence-reduction tool (Manes, 2021; Ferguson and Witzburg, 2021; Connealy et al., 2024).

We contribute to a growing literature on the effect of technology on policing, the criminal justice system, and in a wider context, efficient workforce allocation and policies. While previous studies have found positive effects of crimi-

[5]A detailed analysis of Priority 2 and Priority 3 calls is available in Appendix E. In the main text, we focus on Priority 1 calls.

nal justice and police technology in the form of algorithmic bail decisions (Kleinberg et al., 2018), body-worn cameras (Zamoff et al., 2022; Ferrazares, 2023; Kim, 2019a), electronic monitoring (Williams and Weatherburn, 2022; Rivera, 2023), military-grade equipment (Harris et al., 2017; Bove and Gavrilova, 2017), predictive policing (Mastrobuoni, 2020; Jabri, 2021; Heller et al., 2022), and traffic cameras (Conover et al., 2023), we conversely find significant unintended consequences that are both fiscally and socially expensive.[6] As a consequence, our results provide further evidence that efficient allocation and effective policies are imperative for better policing outcomes (Getty et al., 2016; Ba et al., 2021; Kapustin et al., 2022a; Rivera and Ba, 2023; Adger et al., 2023), and on a larger scale, general workforce productivity (Hsieh and Klenow, 2009; Fenizia, 2022).

More broadly, this study adds to the claim that police departments are personnel-constrained, and potentially understaffed (Chalfin and McCrary, 2018). Similar studies have explored the elasticity of crime with respect to police presence and generally find that increased police presence lowers crime (Levitt, 1997; Chalfin and McCrary, 2018; Mello, 2019; Weisburst, 2019; Weisburd, 2021). Of these works, the most related is Weisburd (2021), which leverages changes in police locations, prompted by service calls, to explore a reduction in the availability of police officers that arises from increased demand for police officer time. However, in contrast to Weisburd (2021), this study unpacks a mechanism which determines response times, allowing us to explore how the time constraints of police officers affect their availability to respond to crime. We find that when police resources are stretched thin, the effectiveness of a police force to respond to crimes and arrest

[6]Chicago is estimated to spend approximately 8.9 million each year on ShotSpotter technology. For comparison, a 2016 estimate put body-worn cameras at 6.5 million annually.

perpetrators is diminished. As a result, our findings suggest that implementing a personnel-intensive policy should be paired with an increase in officer availability, achieved through hiring or redistributing responsibilities, in order to prevent under-policing in communities.

The paper proceeds as follows: Section 2 provides background information on dispatching procedures and implementation of ShotSpotter in Chicago, Section 3 discusses the data, Section 4 describes the empirical strategy, Section 5 presents the main results, mechanism, and effect on arrest probability, Section 6 discusses ShotSpotter's effectiveness as a gun-crime reduction tool, and Section 7 concludes.

2 Background

2.1 ShotSpotter Technology and Implementation in Chicago

ShotSpotter is an acoustic gunfire technology that uses a network of microphones and sensors on buildings and light-posts to detect gunfire sounds. These sounds are used to triangulate the location of gunfire, which is then relayed to police departments to rapidly deploy police officers to the crime scene. Over the past decade, this technology has seen significant expansion and is now operational in over 150 cities globally. The rationale for adopting ShotSpotter is to enable police departments to respond to gunfire faster and with more geographic precision. Moreover, the unique functionality of ShotSpotter allows police departments to bypass their reliance on civilian reporting, which only accounts for approximately 12% of gunfire occurrences (Carr and Doleac, 2016). While previous studies support some of these rationales in the form of geographic accuracy (Piza et al., 2023) and faster gun-related dispatch times (Choi et al., 2014), others have found little impact on gun violence (Mares and Blackburn, 2012; Connealy et al., 2024) and case resolution (Choi et al., 2014).

The technology relies on machine learning algorithms to classify sounds of gunfire.[7] When a gunshot is detected, the sensors triangulate the location of the noise and data/recordings of the incident are forwarded to ShotSpotter's Incident Review Center. At this center, a human reviewer assesses the data and flags for false-positives. Once a gunshot is confirmed, information regarding the location and number of shots fired are shared with the police department, where dispatchers then send officers to the scene. This entire process from gunshot noise to police dispatch is known as a *ShotSpotter dispatch*.

In Chicago, ShotSpotter technology has been implemented in 12 of the 22 police districts in order to respond to gun-related issues faster and with more geographic accuracy.[8] The staggered roll-out began in January 2017 in response to the large influx in gun violence in 2016.[9] ShotSpotter was first implemented in the districts with the highest rates of gun violence and later implemented in less

[7] According to ShotSpotter's website, from 2019 to 2021, the aggregate accuracy rate across all of their customers was 97% with a very small false-positive rate of approximately 0.5%, however this has not been independently tested.

[8] In Chicago, the average population of a police district is approximately 100k. However, there is substantial variance in size, ranging from approximately 60k-250k.

[9] This wide-scale adoption follows previous testing of select areas between 2003 and 2007, 2012, and again in 2016. However, to our knowledge, no district received district-wide coverage during this trial period and the extent of testing was small (https://www.cbsnews.com/chicago/news/chicago-police-testing-new-gunshot-detection-technology/). Moreover, there appears to be no ShotSpotter dispatches in the data prior to the official dates. In an abundance of caution, we conduct a leave-one-out analysis and find that the results are consistent.

violent districts.[10] The expansion ended in May 2018, with no further police districts receiving the technology. Appendix Figure F1 shows the locations of the 12 police districts in Chicago that received ShotSpotter technology. As mentioned, the areas where this technology is implemented (the South and West Chicago areas) experience higher rates of gun crime, on average.

2.2 Dispatching 911 Calls and ShotSpotter Alerts in Chicago

In Chicago, the coordination of emergency 911 calls involves two entities: the Office of Emergency Management and Communications (OEMC) and the Chicago Police Department (CPD). The OEMC oversees 911 calls and dispatches police officers from the CPD. Each 911 call is prioritized on a scale of imminent danger/threat ranging from Priority 1 (immediate dispatch) to Priority 3 (routine dispatch).[11]

When a 911 call is made, the call is received by an OEMC call-taker who records the caller's information, assigns a call type that characterizes the incident, and forwards this information to the dispatcher.[12] Next, the dispatcher assigns the event to an available CPD unit in the caller's police district. Once the scene has been cleared, officers notify the OEMC and mark themselves as available for future call assignments.

On the other hand, the coordination of ShotSpotter dispatches is a collaborative effort involving the OEMC, CPD, and the Strategic Decision Support Center

[10] Note that difference-in-differences relies on the assumption of common trends, not random assignment of the rollout.

[11] Technically, there are six priorities ranging from Priority 0-5. However, Priority 0, 4, and 5 are reserved for special cases such as police officers calling for emergency assistance, administrative meetings, or alternate responses that do not need a field unit, respectively.

[12] Later in Section 3.1, we define the beginning of a 911 call as the time when a call-taker assigns a call-type. This is done rapidly and allows us to more closely target delays due to police officers.

(SDSC). When gunfire is detected, ShotSpotter's headquarters sends vital information such as the location, time, estimated severity, and amount of shots being fired to the SDSC. The SDSC then synthesizes this information and notifies the OEMC to immediately dispatch a police officer to the location of the gunfire.[13]

Importantly, there is a clear distinction between 911 calls and ShotSpotter dispatches. A 911 call is the result of a civilian reporting a crime, while a ShotSpotter dispatch is a police dispatch to the location of a gunfire sound from ShotSpotter sensors. The focus of this paper is 911 calls, which we show to be impacted by the *presence* of ShotSpotter dispatches.

Both 911 calls and ShotSpotter dispatches share a variety of procedural similarities. For instance, each ShotSpotter dispatch is classified with the same distinction as a Priority 1 911 call, therefore necessitating immediate dispatch.[14] Hence, both Priority 1 911 calls and ShotSpotter dispatches share the same dispatch procedures and responding officers. Furthermore, the OEMC prioritizes both 911 calls and ShotSpotter dispatches to rapid response units and police officers within the police district of occurrence.[15] Only in rare circumstances are police officers assigned to these emergencies outside their district.[16]

Despite the similarities, officers must follow additional operating procedures when responding to a ShotSpotter alert. In particular, officers are instructed to

[13]Alternatively, officers can also receive ShotSpotter alerts directly on their phones or smartwatches, without them first being filtered through the SDSC and elect to respond to them this way if they are not already responding to a call.

[14]Priority 1 calls account for roughly 43% of all 911 calls during the sample period.

[15]Specifically, dispatchers prioritize dispatching police officers within the beat they are assigned to. Police beats are subsections within police districts.

[16]In particular, the dispatching order is in the following order of priority: rapid response unit or beat unit from the beat of occurrence, tactical unit, rapid response sergeant, sector sergeant, tactical sergeant, other field supervisor, and closest available unit.

canvass a 25-meter radius for victims, evidence, and witnesses, report any deficiencies in ShotSpotter data or alerts, and document if a case incident is ShotSpotter-related. Interestingly, the average ShotSpotter-related dispatch resolves faster than 911 calls (20 minutes to 65 minutes respectively), thus questioning the productivity of in comparison to a 911 call. In Section 7, we report descriptive evidence showing that approximately 2.2% of all ShotSpotter dispatches result in an arrest, compared to 3.5% of gun-related 911 calls prior to ShotSpotter implementation.

3 Data

3.1 Data Sources

The main sample contains several data sources from years 2016 to 2022 that are obtained through Freedom of Information Act requests to the Chicago Police Department (CPD). These data include dispatched 911 calls, sworn police officer shifts, incidents of crime, arrest reports, and district-level ShotSpotter activation dates.

The CPD 911 call dispatch data encompasses all 911 calls that led to the dispatch of a CPD officer. This administrative data is rich, containing information on the time of the 911 call, the time an officer is dispatched to the scene of the crime, and the time the officer arrives on-scene. Additionally, the data details the priority-level of the call, a brief description, a block-level location, the final disposition, and a case report number that links to arrests and incident reports.

Based on this information, we construct the two main outcome variables: the time from the start of a 911 call to an officer being dispatched (Call-to-Dispatch) and the time from the start of a 911 call to an officer's arrival (Call-to-On-Scene).

We define the start of the 911 call as the time that a 911 call-taker creates an event number for the associated incident—an action that typically occurs immediately following the call being received. Notably, while Call-to-Dispatch contains no missing data, approximately 45% of the Call-to-On-Scene information is missing. This is due to officers failing to report when they arrive at the scene (OIG, 2023). However, we present the main results for when we can observe both outcomes in Section 5, and also report that both outcomes are highly correlated in Appendix A. For the remainder of the paper, unless otherwise specified, the results are presented using the entirety of Call-to-Dispatch observations.

These two measures of rapid response capture separate degrees of police availability. First, if an officer is too busy, they will be delayed or unable to be dispatched. In particular, the officer will not be classified as available to take Priority 1 calls on the Computer Aided Dispatch (CAD) system, and a dispatcher will not assign them to a call. This increase in time would be observed as a higher Call-to-Dispatch time and is a function of the coordination between the dispatcher and an individual police officer. On the other hand, Call-to-On-Scene—which captures both the dispatch time and the time an officer takes to arrive on-scene— may increase independently of Call-to-Dispatch time if, for example, an officer is located farther away from their dispatch location.

The police shift data contains information on every start time, end time, and district/beat assignment worked by CPD staff. We restrict the shift data to include only police officers that are present for duty, and exclude administrative positions and higher level managerial roles such as police lieutenants and police chiefs. To assess officer availability, we construct the number of officer hours within a police district-day. By using on the number of officer hours rather than the number of shifts, we account for the possibility of overtime or early-leave.

The ShotSpotter activation dates indicate when each police district is equipped with ShotSpotter technology. However, since the records provide only the month of implementation, we rely on the raw data corresponding to ShotSpotter dispatches to determine the specific activation day for each police district. Nonetheless, we observe several small discrepancies in the activation dates when comparing to the number of ShotSpotter dispatches in District 6, 9, 10, and 15. In particular, these districts have no ShotSpotter dispatches until several months after their official activation date. Therefore, we adjust these four dates of activation to align with the onset of ShotSpotter alerts. This adjustment ensures that the effects observed are accurately attributed to police officers responding to ShotSpotter alerts. However, as a robustness check, we estimate the results using the official dates in Appendix Figure F2 and find that the results remain consistent.

Figure 1 plots the monthly trend of dispatches relating to both ShotSpotter and civilian reports of gunshots. In addition, the ShotSpotter activation dates are plotted with dashed red lines. In this figure, each police district exhibits an increase in ShotSpotter dispatches as time progresses, a pattern which is likely due to the increasing amounts of gun violence which began in the year 2020 (Grawert and Kim, 2023). Notably, this figure also depicts a substantial increase in police resources devoted to gunfire post-ShotSpotter implementation.

3.2 Sample Restrictions

The main sample is restricted to dispatched 911 calls classified as Priority 1—the highest priority level.[17] Priority 1 is defined as any situation that may involve an imminent threat to life, bodily injury, or major property damage/loss. By including only Priority 1 calls, the analysis focuses on calls that require the most time-sensitive response. However, Appendix E analyzes lower-priority calls of Priority 2 and Priority 3.

As an important distinction, dispatched 911 calls do not include dispatches for ShotSpotter gunshot detections. While ShotSpotter detections are classified as Priority 1 and responded with the same police units, these are not reported by civilians.[18] By implementing this restriction, we ensure that we are comparing similar distributions of civilian reports of crime before and after the ShotSpotter rollout.

Four additional restrictions are implemented to reduce potential noise and confounding factors in the response time data. First, observations with negative Call-to-Dispatch or Call-to-On-Scene times are removed (0.03%). Second, due to the sensitivity to outliers of the ordinary least squares estimator, Call-to-Dispatch and Call-to-On-Scene outliers that exceed three standard deviations from the mean are omitted (0.4% and 1.6%, respectively). Third, specific dates that coincide with celebratory gunfire and fireworks are removed (e.g., January 1, July 4, and December 31) to ensure the results are not driven by these particular days. Last, since 911 dispatchers can combine calls that coincide directly with a ShotSpotter alert,

[17] Priority 0 is actually the highest level of priority, but this is a special case reserved for situations where police or firefighters are calling for assistance in life-threatening situations. These are extremely rare, and make up only 0.01% of the top four priority dispatches.

[18] In some cases, ShotSpotter alerts can also be responded to by specialized tactical units which do not respond to Priority 1 calls.

we also remove any 911 call that is classified as a call for “Shots Fired.” Importantly, the results remain robust to including each of these restrictions, as shown in Appendix Figure F2.

3.3 Descriptive Statistics

Table 1 shows summary statistics of the main outcomes and other relevant variables. All statistics are based on dispatched Priority 1 911 calls unless otherwise noted. Panel A reports that the average Call-to-Dispatch time is roughly five minutes, while the average Call-to-On-Scene time is approximately 13 minutes. Additionally, the distribution of these outcomes are plotted in Figure 2 showing that response times can be particularly lengthy (1+ hours) in rare cases. Furthermore, the probability of making an arrest on a Priority 1 call is low, with an average of 3%.

In Panel B, Priority 2 and Priority 3 calls are reported to be less frequent than Priority 1. Priority 2 calls are defined as those in which timely police action has the potential to affect the outcome of an incident, while Priority 3 calls are those in which a reasonable delay in police action will not affect the outcome of the incident. Consistent with these definitions, Priority 2 and Priority 3 have slower response times for both Call-to-Dispatch and Call-to-On-Scene measures.

Furthermore, statistics on the number of Priority 1 911 dispatches, ShotSpotter dispatches, and number of officer hours, are reported in Panel C of Table 1—each measured at the district-day level. The average number of Priority 1 dispatches within each district-day is approximately 73, although these have considerable variability, with a maximum of 223. ShotSpotter dispatches are reported to be an average of approximately three per-district-day although this includes time periods and districts that do not necessary have ShotSpotter implemented. When restricting the sample to only post-ShotSpotter implementation dates, the average number of ShotSpotter dispatches in each treated district-day is six (∼70 city-wide). Finally, due to the high level of crime in the South and West locations of Chicago, the presence of officers varies considerably across districts, ranging from as little as 231 officer hours to as many as 6,558 officer hours. We later analyze this heterogeneity in Section 5.2 where we find longer response times when there are fewer officers.

4 Empirical Strategy

4.1 Baseline Specification

To estimate the causal effect of ShotSpotter technology on police response times, we estimate the following staggered difference-in-differences (DiD) equation using ordinary least squares (OLS):

$$ResponseTime_{cdt} = \beta ShotSpotter_{dt} + \eta_{\tilde{c}} + \delta_d + \gamma \mathbb{X}_{f(t)} + \varepsilon_{cdt} \tag{1}$$

where $ResponseTime_{cdt}$ represents the Priority 1 Call-to-Dispatch or Call-to-On-Scene time for call c, in police district d, at time t. The treatment variable, $ShotSpotter_{dt}$, is a binary indicator equal to one if police district d is equipped with ShotSpotter at time t. The model includes fixed effects for call types ($\eta_{\tilde{c}}$) and police districts (δ_d), as well as a vector of time-varying controls ($\mathbb{X}_{f(t)}$) which consist of day-by-month-by-year and hour-of-the-day fixed effects.[19] By including these terms, we control for fixed differences across different call distinctions, police districts, and time periods. The error term is ε_{cdt}. Moreover, standard errors are clustered by police district (N = 22) to allow for serial correlation within districts, although we also report wild cluster bootstrapped standard errors in the main results as recommended by Cameron et al. (2008) since the number of clusters is below 30. Intuitively, Equation 1 is comparing response times on days with ShotSpotter activated to days without ShotSpotter activated, while accounting for the expected differences in call types, police districts, and different times of the year and day.

4.2 Identification

The coefficient of interest in Equation 1 is β, which measures the average change in response times between days with and without ShotSpotter technology. To identify β as a causal effect, there are several assumptions that must be satisfied: response times in ShotSpotter districts would have continued on a similar trend to non-ShotSpotter districts in the absence of ShotSpotter, there is no change in 911 dispatching procedures post-ShotSpotter implementation, the distribution of 911 calls/dispatches did not change post-ShotSpotter, and there are no other changes that coincide with the timing of ShotSpotter that may affect response times.

The first identification assumption is that police districts that adopt ShotSpotter would have continued to have similar response times to non-ShotSpotter districts in the absence of adoption (i.e., *common trends*). Specifically, ShotSpotter adoption

[19] Each 911 call is given a final dispatch code. When controlling for the type of call, we use the final dispatch code as the distinction.

must not be correlated with a systematic rise or fall in response times. To address this concern, we estimate an event study framework given by the following model:

$$ResponseTime_{cdt} = \sum_{\substack{i=-12,\\ i\neq -1}}^{24} \beta^i \text{ShotSpotter}^i_{dt} + \eta_{\tilde{c}} + \delta_d + \gamma \mathbb{X}_{f(t)} + \varepsilon_{cdt} \qquad (2)$$

where $ShotSpotter^i_{dt}$ is a set of indicators that are set to 1 if ShotSpotter is adopted i months from time t in district d. Each period is relative to the month before ShotSpotter adoption. Twelve periods pre-ShotSpotter and 24 periods post-ShotSpotter are estimated to maintain a balanced panel, and the first and final periods are binned endpoints as described in Schmidheiny and Siegloch (2023).

Figures 3 and 4 show the event study estimations for Call-to-Dispatch and Call-to-On-Scene response times, and display little visual evidence of an upward or downward trend prior to the implementation. The error-bars represent 95% confidence intervals, while the coefficient estimates are reported in seconds. We report two sets of estimates in this visualization: the two-stage difference-in-differences imputation estimator (Gardner, 2021) and the OLS estimator. The two-stage difference-in-differences estimator is robust to the negative weights which arise in OLS estimates when there are heterogeneous treatment effects across groups and over time in staggered designs (de Chaisemartin and D'Haultfoeuille, 2020; Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Athey and Imbens, 2022). Unlike the estimators proposed in Sun and Abraham (2021) and Callaway and Sant'Anna (2021), this estimator allows us to maintain the preferred day-by-month-by-year fixed effects while simultaneously estimating monthly bins without aggregation. Moreover, this estimator allows for comparisons of treated units between *both* never-treated and not-yet treated units. In each set of estimations, there appears to be little evidence of a trend prior to ShotSpotter implementation. We later enhance this visual test in Section 5.1 (and more thoroughly in Appendix D) with the sensitivity test described in Rambachan and Roth (2023) where we allow for relaxations of the common trends assumption.

The second assumption states that there is no change in how police are dispatched to 911 calls post-implementation. We find no indication that the operating procedures for 911 calls change within the sample period (CPD, 2016).

Third, we assess whether the distribution of 911 calls is changing post-implementation. While the main findings are presented here, a detailed analysis is available in Appendix B. First, we estimate whether the number of 911 calls changes post-implementation using an event-study framework. Although the total number of calls remains stable, there is evidence that Priority 1 calls may slightly decrease (roughly two fewer calls). However, this reduction is likely to *reduce* officers' responsibilities, rather than strain them. Second, we test for changes in the likelihood of distinct call-types using a variation of Equation 1 with a binary variable for each call-type as the outcome. This involves 112 hypothesis tests, with p-values corrected for multiple hypothesis testing using the method by Benjamini and Hochberg (1995). After correction, we do not find substantial evidence of distribution changes, with only 2% of the call-types exhibiting significant changes in likelihood. Finally, as mentioned in Section 3.2, we remove any 911 call that is classified as a call for "Shots Fired," since dispatchers can combine calls that coincide with a ShotSpotter alert. Nevertheless, as shown in Appendix Figure F2, the results are robust when including these calls.

For the final assumption that there are no other changes that coincide with ShotSpotter implementation, we discuss two initiatives that are implemented at similar times as ShotSpotter: Strategic Decision Support Centers (SDSCs) and Body Worn Cameras (BWC). A thorough discussion is available in Appendix C, while the key takeaways are reported here. First, SDSCs have similar implementation dates to ShotSpotter with an average of 73 days *prior*, although only 57% of SDSCs are equipped with ShotSpotter as shown in Appendix Table C1. SDSCs are housed with software such as observation displays, geospatial predictive policing software, and social media monitoring. However, only one of these technologies coincides with the SDSC roll-out (geospatial predictive policing), and the others have been utilized in Chicago in years prior. While predictive policing software may change officer patrolling patterns, and therefore affect response times, Kapustin et al. (2022b) find that officers are responsive to the software recommendations in only two of Chicago's police districts. In Appendix C, we estimate the main results and the corresponding event studies while controlling for SDSC roll-out dates, and report consistent findings with the main results.[20] Second, BWCs are also implemented near ShotSpotter dates, although the district-timing differs by an average of 283 days. In Appendix Table C2, we control for the BWC implementation and find little differences from the main results. This aligns with intuition, as body worn cameras have been found to affect complaints (Kim, 2019b; Braga et al., 2022; Zamoff et al., 2022; Ferrazares, 2023) and stops (Braga et al., 2022; Zamoff

[20] In Section 5.2, we also present intensive margin estimates of ShotSpotter using the number of ShotSpotter dispatches as identifying variation. This variation is less correlated with the SDSC roll-out, and provides further evidence that ShotSpotter is causing the increase in response times.

et al., 2022), but are unlikely to affect an officer's ability to rapidly respond.

5 Results

In this section, we present the main estimates on the effect of ShotSpotter on Priority 1 response times using Equation 1. We show that the results are robust across various specifications, estimators, sample selections, and sensitivity tests. Moreover, we analyze dynamic effects and present evidence that ShotSpotter affects response times by constraining officer resources. Last, we show that increased response times lead to fewer perpetrators being arrested, thereby showing that ShotSpotter has costly implications.

Figure 5 serves as an intuitive preview of the main results, plotting only the raw data. We plot the average Call-to-Dispatch and Call-to-On-Scene times within each police district before/after ShotSpotter implementation. Consistent with the main results, districts that receive ShotSpotter show a substantial increase in the average Call-to-Dispatch and Call-to-On-Scene times. Notably, there does not appear to be significant visual evidence that average response times are different in districts that receive ShotSpotter in comparison to those that did not.

5.1 Main Results - Response Time Changes

Table 2 reports estimates from Equation 1 for Call-to-Dispatch (Panel A) and Call-to-On-Scene (Panel B) response times, where each coefficient estimate is reported in seconds. Recall that Call-to-Dispatch and Call-to-On-Scene are the length of time from when a 911 call is received to when a police is dispatched or subsequently arrives at the scene, respectively. First, in Column 1 of Table 2, we estimate Equation 1 with only the time and group fixed effects. We find a statistically significant increase in Call-to-Dispatch and Call-to-On-Scene times of 64 seconds and 101 seconds, respectively. Remarkably, the Call-to-On-Scene estimates are an additional 40 seconds on top of the delays in finding officers to dispatch. This suggests that ShotSpotter may not be placing officers in areas closer to the majority of other 911 emergencies, whereby travel time may be reduced.

Column 2 of Panel A and Panel B report estimates from the preferred specification outlined in Section 4.2 where we supplement the model in Column 1 with controls for time-of-day and the type of 911 call. When including these controls, the results for both response times are consistent, showing increases from the mean of approximately 22% and 13%, respectively. In Column 3, we restrict the observations to only those where we can observe both response times, as Call-to-On-Scene is reported less frequently. The results show that this restriction yields consistent conclusions, although we further explore this discrepancy in Appendix A.

Given the staggered difference-in-differences research design, Column 4 reports estimates that are robust to treatment heterogeneity across groups and over time using the two-stage difference-in-differences imputation estimator (Gardner, 2021). This estimator equally weights each district-date estimate, making it robust to the bias from negative weighting in the presence of treatment effect heterogeneity (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Athey and Imbens, 2022). We opt to use this estimator since it allows for comparisons of treated units between *both* never treated units and not-yet treated units and requires no aggregation, unlike similar approaches discussed in Callaway and Sant'Anna (2021). The estimates, albeit slightly larger than the preferred specification, remain consistent with the main findings.

Furthermore, we consider spillover effects in Column 4 by including an indicator variable (*Border Activated*) equal to one for any police district that is adjacent to a ShotSpotter-activated district. In effect, the coefficient on the indicator for a neighboring ShotSpotter district measures the spillover impacts of the implementation. As reported in both Panel A and Panel B, there does not appear to be evidence of spillover effects on response times. This result aligns with the standard dispatching procedures discussed in Section 2.2, where officers are only dispatched outside their beat/district of patrol in rare circumstances.

Next, to analyze the dynamic effects of ShotSpotter implementation over time, we estimate an event study using Equation 2. We estimate this model using both OLS and the Gardner (2021) robust estimator to account for potential treatment heterogeneity across groups and time periods. Figure 3 and Figure 4, for Call-to-Dispatch and Call-to-On-Scene respectively, show that the effect of ShotSpotter implementation takes several months post-implementation to significantly alter response times. In each figure, the red error-bars represent the 95% confidence intervals using OLS, while the blue error bars are derived from the Gardner (2021) estimator. We attribute the delayed effect in response times to the increasing number of ShotSpotter dispatches that occur over time, as previously shown in Figure 1.

Importantly, these results are robust to a variety of sample selections and sensitivity tests. First, Appendix Figure F2 shows estimations of Equation 1 relaxing the four additional restrictions outlined in Section 3.2, in addition to omitting year 2020 (Covid-19 pandemic), and using the official activation dates from a Freedom of Information Act request rather than the observed beginning of ShotSpotter dispatches. In each of these subsamples, the results remain consistent. Second, we perform a leave-one-out analysis in Appendix Figure F3 where Equation 1 is estimated 22 times, with each iteration excluding a unique police district. Given that the results remain consistent in each iteration, we rule out the possibility that these effects are driven by only one police district. Finally, in Appendix D, we conduct analysis following Rambachan and Roth (2023) to illustrate the sensitivity of the event study estimates to possible violations of parallel trends. Specifically, we evaluate the degree of nonlinearity we can impose on a linear extrapolation of the pre-treatment trend while maintaining a significant post-treatment average treatment effect. As explained further in Appendix C, we find that the average of all post-implementation periods maintain their statistical significance under both a linear extrapolation of the pre-period and increasing amounts of non-linearity for both response times.

5.2 Mechanism - Resource Constraints

In this subsection, we provide evidence that the longer response times associated with ShotSpotter are a result of the allocation of scarce police resources. Recall from Section 3.3 that post-implementation, there are approximately 70 ShotSpotter dispatches each day in Chicago—a two-fold increase in the number of gunfire-related incidents officers must respond to compared to pre-implementation. These dispatches are resource-intensive, taking an average of 20 minutes each, which collectively amounts to roughly 75 hours of officer time allocated to ShotSpotter.[21] To establish this link, we conduct three analyses to show that ShotSpotter creates longer response time delays on both the extensive margin (implementation) and the intensive margin (number of ShotSpotter dispatches).

First, on the extensive margin, we differentiate the effect of ShotSpotter by officer watch schedules, which represent times when officers begin and end their shift. This division allows us to examine periods with varying levels of ShotSpotter dispatches, wherein officers may be more or less constrained by attending to ShotSpotter investigations. Panel A of Figure 6 plots the distribution of ShotSpotter dispatches by the hour of the day and corresponding watch. As shown in the figure, the nighttime shifts of Watch 1 (11:00pm - 7:00am) and Watch 3 (3:00pm - 11:00pm) have significantly higher counts of ShotSpotter dispatches than Watch 2 (7:00am - 3:00pm).[22] In Panel B of Figure 6, we plot the coefficient estimates of Equation 1 by officer watch and show that watch times with higher levels of ShotSpotter dispatches have longer response time delays. On the x-axis, each coefficient estimate and 95% confidence interval is plotted for the corresponding watch number on the y-axis. For both outcomes, the effects correspond to the distribution of ShotSpotter dispatches in Panel A; Watch 1 and Watch 3 exhibit effects that are both statistically significant and larger in magnitude than Watch 2. Moreover, while the Call-to-On-Scene delays reach nearly 3 minutes in Watch 3, the Call-to-On-Scene estimates are near-zero for Watch 2, and are not statistically significant.

[21] As mentioned in the introduction, we calculate this using the average number of officers that are dispatched to ShotSpotter detections over a sample period of 2019-2023 (roughly three officers). Unfortunately, records retention schedules did not allow us to receive this data for our sample period.

[22] The typical police watches in Chicago last for 9 hours total with a 45-minute briefing to begin the shift. We use 8-hour intervals to account for these briefings.

Second, also on the extensive margin, we show that the longer response times are driven by district-days that have fewer officers on duty and, consequently, fewer police resources. In Columns 2 and 3 of Table 3, we divide the sample by the district-day median of officer availability, measured using the number of officer-working hours. Column 2 (above median, less constrained) shows percentage increases of 14% for Call-to-Dispatch and 8% for Call-to-On-Scene, while Column 3 (below median, more constrained) reports increases of 26% and 17% for Call-to-Dispatch and Call-to-On-Scene, respectively.

Finally, on the intensive margin, we exploit the number of daily ShotSpotter dispatches as an alternative source of variation to test whether more resources devoted to ShotSpotter results in longer response times. To do so, we modify Equation 1 by aggregating the data to the district-day level, restricting the sample to only ShotSpotter-implemented police districts and days of activation, and changing the variable of interest from a binary treatment to a continuous count of the number of ShotSpotter dispatches.[23] In effect, this model measures the marginal effect of an additional ShotSpotter dispatch on response times.

Column 4 of Table 3 shows that one additional ShotSpotter dispatch is associated with an increase in the average Call-to-Dispatch time of 6 seconds and an increase in the average Call-to-On-Scene time of 8 seconds. These results are statistically significant at the 1% level. However, we note that these results are under the assumption of a linear relationship between the number of ShotSpotter dispatches and response times. We show the plausibility of this assumption in Appendix Figure

[23]The corresponding model is $ResponseTime_{dt} = \zeta ShotSpotterDispatches_{dt} + \delta_d + \gamma_t + \varepsilon_{dt}$ where $ShotSpotterDispatches_{dt}$ is the number of dispatches attributed to ShotSpotter alerts in district d at time t, δ_d are police district fixed effects, and γ_t are day-by-month-by-year fixed effects.

F4 where we split the number of ShotSpotter dispatches into deciles and re-estimate the model. Interestingly, we find that each response time increases monotonically with ShotSpotter dispatches, which further suggests the incapacitation effect it has on officers.

Taken together, these findings underscore the significance of police resource allocation within a day. If ShotSpotter affects response times by overloading officer responsibilities, then it is imperative to reallocate the appropriate amount of staffing to times when ShotSpotter dispatches are more frequent.

5.3 Impact on Call Resolution

Although the findings demonstrate that ShotSpotter affects police officer response times, we acknowledge that this may not yield detrimental consequences if it does not reduce the likelihood of apprehending perpetrators. To address this concern, we investigate changes in the arrest probability of 911 calls using two methods: first, we estimate the reduced form where ShotSpotter implementation is the binary treatment variable, and the arrest probability is the outcome variable.[24] Second, we employ an instrumental variables (IV) design where the ShotSpotter implementation is the instrument, Call-to-Dispatch is the endogenous variable of interest, and Arrest Made is the binary outcome variable. By estimating an IV design, we estimate the marginal effect of an extra second in dispatch delay on arrest likelihood for only 911 calls that are induced to have a longer Call-to-Dispatch time by ShotSpot-

[24] We use two sets of arrest data which include arrests from the arrest database and case reports that end in arrests. We merge this to the 911 dispatch data using incident report number as the common identifier (RD). Based on conversations with the Chicago Police Department, this is the best way to map 911 calls to arrests.

ter (compliers). This more precisely isolates the direct implication of ShotSpotter response time delays on public safety, and moreover, is consistent with the reduced-form estimates. To corroborate these findings, we examine the final disposition of 911 calls using both methods and find evidence that officers are more likely to report dispositions that are consequential of a slower arrival.

Column 1 of Table 4 shows that the likelihood of an arrest made decreases because of the ShotSpotter implementation. Panel A shows the reduced-form estimates while Panel B shows second-stage of the two-stage least squares (2SLS) design—each multiplied by 100 for a percentage-point interpretation. In particular, the reduced-form estimates in Panel A show that the arrest likelihood decreases by 9% relative to the mean and is statistically significant at the 1% level. In Panel B, the results from the 2SLS estimation show that the marginal effect of an additional second of Call-to-Dispatch time results in a statistically significant 0.004 percentage point decrease in the likelihood of an arrest being made. Given the average increase in Call-to-Dispatch time is approximately 60 seconds (first-stage), this results in an average decrease of roughly 0.240 percentage points—consistent with the reduced form. To assess the validity of these findings, we present the reduced-form estimates using the Gardner (2021) estimator in Appendix Table F3 and additionally plot the corresponding event study in Appendix Figure F5 to assess the plausibility of the parallel trends assumption.[25] For 2SLS, we report the efficient F-statistic (Eff F-stat) from Olea and Pflueger (2013), and Anderson and Rubin (1949) confidence intervals with corresponding p-values for weak instruments. In each of these analyses, we find consistent results and little visual evidence of a pre-trend.

[25] In addition, we estimate the reduced form using logistic regressions rather than OLS. The results are shown in Appendix Table F2. The results remain consistent.

In Columns 2-4, we utilize the 911 call final dispositions and find evidence that officers are more likely to report dispositions that are consistent with increased response times. To begin, we analyze changes in the three most frequent final dispositions which account for approximately 70% of all 911 calls: Other Police Service (41%), No Person Found (21%), and Peace Restored (6%).[26] Consistent with response time delays, Column 3 in Panel A shows that officers are more likely to report that a person is not found by 0.751 percentage points (3%). This result is statistically significant at the 5% level. Moreover, Column 4 of Panel A demonstrates that there is a statistically significant 0.558 percentage point reduction (8%) in reports of peace being restored, thereby raising the notion that an officer has arrived too slowly and there is no longer a conflict to mitigate. In Panel B, each of these columns are consistent in magnitude when estimating 2SLS, although only Person Not Found has similar statistical inference.

Despite these corroborating results, we note that final dispositions are officer-reported, and there may be instances where an officer does not carefully delineate between codes (i.e., Person Not Found vs. Other Police Service). Hence, while these results are in-line with slower response times, we advise these final disposition results (Columns 2-4) be taken as suggestive.

[26] Each of these are dispositions in which an officer did not find what the initial service call was for. For instance, a civilian may call for a domestic disturbance only for an officer to arrive and not find any perpetrator. On the other hand, if an officer did find the perpetrator, they would report in their final disposition the corresponding crime code. Appendix Figure F6 plots the estimates for all other final disposition codes, which are particularly infrequent.

6 Discussion

6.1 Is ShotSpotter an Effective Gun-Violence Reduction Tool?

Although previous sections have shown that ShotSpotter reallocates police resources to the detriment of 911 calls, we now broaden the analysis to independently test whether ShotSpotter contributes to a reduction in gun violence and enhances police productivity beyond the 911 call system. As mentioned in Section 2, while ShotSpotter has been found to provide geographical accuracy (Piza et al., 2023) and faster gun-related dispatch times (Choi et al., 2014), there is also evidence that ShotSpotter has no impact on reducing gun violence (Connealy et al., 2024) and fails to result in police action upon arrival (Ferguson and Witzburg, 2021; Manes, 2021).

In this subsection, we analyze whether ShotSpotter affects *aggregate* counts of gun-related arrests, clearances, and victimization from across the city, regardless of the relation to a 911 call. In doing so, we test whether ShotSpotter provides gun regulation, results in valuable intelligence/evidence needed to solve crimes, or saves lives from gun violence. The results show little evidence of these potential benefits, thereby reconciling findings from other studies focusing on Chicago (Ferguson and Witzburg, 2021; Manes, 2021; Connealy et al., 2024).

To begin, we construct a new data set which includes daily-level counts of gun-related arrests, clearance, and victimization at the district-level. We define an outcome as gun-related if the charge description contains the word ‘gun’ or 'firearm'. Next, we estimate an event study framework similar to Equation 2:

$$Y_{dt} = \sum_{\substack{i=-12, \\ i \neq -1}}^{24} \beta^{i} \text{ShotSpotter}^{i}_{dt} + \delta_d + \gamma_t + \varepsilon_{dt} \tag{3}$$

where Y_{dt} is the outcome of interest in district d in time t. Day-by-month-by-year (γ_t) and district fixed effects (δ_d) are also included. Given the (non-negative) count nature of the outcomes, we estimate Equation 3 using Poisson estimation in lieu of OLS. Standard errors are clustered at the district level.

Figure 7 reveals no significant changes in the aggregate counts of gun arrests, gun crimes cleared, or gun victimization following the implementation of ShotSpotter. For each outcome, we find few statistically significant estimates and minimal evidence of a lasting effect. This contrasts with a marked increase in police allocation to gun-related incidents over this time period (Panel D), as reflected in the total number of 911 calls *combined* with ShotSpotter dispatches.[27] Taken together, despite the reallocation of resources to gun-related incidences, we do not find evidence that ShotSpotter enhances productivity or mitigates gun violence through arrests, clearances, or victimization. [28]

[27] As discussed more in Section 7, despite knowing the number of ShotSpotter dispatches across the entire sample period (2016-2022), we do not know more information about what happens on these dispatches aside from a small subsample (2019-2022) due to data restrictions. Chicago Police Department claims further information prior to 2019 does not exist.

[28] This lack of efficacy in Chicago mirrors concerns expressed in other cities, such as New York City, where the Comptroller Brad Lander has shared criticism that "the evidence shows that NYPD is wasting precious time and money on this technology and needs to do a better job managing its resources."

7 Conclusion

In this study, we analyze the adoption of a policing technology that crowds out police officer time and disrupts the availability of officers. We do so by exploring the effect of ShotSpotter technology on two measures of rapid response: Call-to-Dispatch and Call-to-On-Scene. Using a comprehensive dataset of all dispatched 911 calls over a seven-year period (2016-2022), we find that response times significantly increase following the implementation of ShotSpotter in Chicago. Specifically, we find that 911 dispatchers assign officers one-minute slower (Call-to-Dispatch) and officers arrive on-scene approximately two-minutes later (Call-to-On-Scene). These delays have significant implications, as they reduce the likelihood of resolving a 911 call with an arrest.

Furthermore, we find that ShotSpotter increases response times by re-allocating scarce police resources from 911 calls to ShotSpotter-detected gunfire alerts (ShotSpotter dispatches), resulting in a significant time trade-off. Given the substantial resources that ShotSpotter requires, police officers are forced to allocate a significant portion of their time to fulfill ShotSpotter requirements, thereby incapacitating them from attending to 911 calls. In particular, we show that the effects are driven by times of day when there are fewer police officers on-duty and when ShotSpotter dispatches are most frequent. On the intensive margin, we find that each additional ShotSpotter dispatch results in a six-second increase in Call-to-Dispatch time and an eight-second increase in Call-to-On-Scene time, showing further evidence that ShotSpotter is creating a costly time allotment.

Moreover, we find little evidence of that ShotSpotter is generating substantial benefits outside the 911 call system. City-wide counts of gun-related arrests, clearance, and victimization, each show little evidence of a significant change post-implementation, thus calling into question ShotSpotter's effectiveness as a crime-reduction tool. However, as a limitation of the analysis, we note that we cannot directly evaluate the productivity of a ShotSpotter dispatch in comparison to a 911 dispatch over the sample period.[29] Despite this shortcoming, we utilize a subset of the data (2019-2022) and find descriptive evidence that approximately 2.2% of all ShotSpotter dispatches result in an arrest.[30] As a benchmark, gun-related 911 calls in ShotSpotter districts prior to implementation result in an arrest approximately 3.5% of the time. While this is suggestive of a comparative decrease in productivity, we emphasize that an arrest is not the only productivity measure in a dispatch; police may gather valuable intelligence at the crime scene, or the presence of officers may produce a deterrence effect from subsequent crimes occurring in the area (Chalfin and McCrary, 2017). As a result, while we provide robust evidence to the cost of ShotSpotter implementation, further research is needed to understand the productivity of ShotSpotter dispatches to perform a rigorous cost-benefit analysis.

Hence, we cannot advocate for, nor against ShotSpotter, but aim to inform policymakers of the substantial unintended consequences that it, and similar technologies, create. Given the analysis, we find that ShotSpotter creates a resource constraint problem where officers have too many responsibilities. However, we cannot rule out the possibility that ShotSpotter may be an effective tool for police departments that are sufficiently staffed, yet recommend that police depart-

[29] Two reports from Chicago have raised concerns over ShotSpotter's productivity (Ferguson and Witzburg, 2021; Manes, 2021)

[30] Officers were not required to report whether an arrest was associated to ShotSpotter until after February 2019 according to a Freedom of Information Act request for such information. This number is found using the total number of distinct arrests that are associated with a ShotSpotter and dividing by the number of ShotSpotter dispatches post-February 2019.

ments carefully evaluate whether their operations can accommodate the intensive resources that this technology requires in order to mitigate the consequences. In Chicago, a back-of-the-envelope calculation shows that in order to eliminate the on-scene time delays, 36% more officers are needed.[31] This underscores the notion that police technology such as ShotSpotter, as of now, can possibly act as a valuable complement for police officers, but not as a perfect substitute.

[31] To calculate this, we estimate the specification outlined in Section 5.2, replacing the $NumberSSTDispatches_{dt}$ with the number of officers within district d at time t and the number of officers within district d at time t squared. The marginal effect of an additional officer on response times using this model is to 1.78 seconds increased in on-scene time. We then use the average increase in Call-to-On-Scene from Column 2 of Table 2 (103.7) and divide by the 1.78 to find the number of officers needed to negate this effect. Using the average number of officer hours (1277.86), and dividing by 8 (the average shift time), we find the average number of officers within a district (159.73). Finally, dividing the number of officers needed by the average number of officers within a district gives the percentage increase (36%).

References

Adger, C., Ross, M., and Sloan, C. (2023). The Effect of Field Training Officers on Police Use of Force.

Anderson, T. W. and Rubin, H. (1949). Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations. *The Annals of Mathematical Statistics*, 20(1):46–63. Publisher: Institute of Mathematical Statistics.

Ang, D., Bencsik, P., Bruhn, J., and Derenoncourt, E. (2021). Police Violence Reduces Civilian Cooperation and Engagement with Law Enforcement. *SSRN Journal*.

Athey, S. and Imbens, G. W. (2022). Design-based analysis in Difference-In-Differences settings with staggered adoption. *Journal of Econometrics*, 226(1):62–79.

Avdic, D. (2016). Improving efficiency or impairing access? Health care consolidation and quality of care: Evidence from emergency hospital closures in Sweden. *Journal of Health Economics*, 48:44–60.

Ba, B., Bayer, P., Rim, N., Rivera, R., and Sidibé, M. (2021). Police Officer Assignment and Neighborhood Crime.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300.

Blanes i Vidal, J. and Kirchmaier, T. (2018). The Effect of Police Response Time on Crime Clearance Rates. *The Review of Economic Studies*, 85(2):855–891.

Bove, V. and Gavrilova, E. (2017). Police Officer on the Frontline or a Soldier? The Effect of Police Militarization on Crime. *American Economic Journal: Economic Policy*, 9(3):1–18.

Braga, A. A., MacDonald, J. M., and McCabe, J. (2022). Body-worn cameras, lawful police stops, and nypd officer compliance: A cluster randomized controlled trial.

Butts, K. and Gardner, J. (2021). *did2s: Two-Stage Difference-in-Differences Following Gardner (2021)*.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-Based Improvements for Inference with Clustered Errors. *The Review of Economics and Statistics*, 90(3):414–427.

Carr, J. and Doleac, J. L. (2016). The Geography, Incidence, and Underreporting of Gun Violence: New Evidence Using Shotspotter Data.

Carr, J. B. and Doleac, J. L. (2018). Keep the Kids Inside? Juvenile Curfews and Urban Gun Violence. *The Review of Economics and Statistics*, 100(4):609–618.

Chalfin, A. and McCrary, J. (2017). Criminal Deterrence: A Review of the Literature. *Journal of Economic Literature*, 55(1):5–48.

Chalfin, A. and McCrary, J. (2018). Are U.S. Cities Underpoliced? Theory and Evidence. *The Review of Economics and Statistics*, 100(1):167–186.

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

Connealy, N. T., Piza, E. L., Arietti, R. A., Mohler, G. O., and Carter, J. G. (2024). Staggered deployment of gunshot detection technology in chicago, il: a matched quasi-experiment of gun violence outcomes. *Journal of Experimental Criminology*, pages 1–27.

Conover, E., Kraynak, D., and Singh, P. (2023). The effect of traffic cameras on police effort: Evidence from India. *Journal of Development Economics*, 160:102953.

CPD (2016). Cpd general order g03-01-01: Radio communications. *CPD Directives*.

de Chaisemartin, C. and D'Haultfoeuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review*, 110(9):2964–2996.

DeAngelo, G., Toger, M., and Weisburd, S. (2023). Police Response Time and Injury Outcomes. *The Economic Journal*, page uead035.

Fenizia, A. (2022). Managers and Productivity in the Public Sector. *Econometrica*, 90(3):1063–1084. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA19244.

Ferguson, J. and Witzburg, D. (2021). The chicago police department's use of shotspotter technology. *Chicago, IL: City of Chicago Office of Inspector General*.

Ferrazares, T. (2023). Monitoring Police with Body-Worn Cameras: Evidence from Chicago. *Journal of Urban Economics*, page 103539.

Fischer, A. and Roodman, D. (2021). fwildclusterboot: Fast wild cluster bootstrap inference for linear regression models (version 0.13.0).

Gallup, I. (2022). Confidence in U.S. Institutions Down; Average at New Low. Section: Politics.

Gardner, J. (2021). Two-stage differences in differences. arXiv:2207.05943 [econ].

Getty, R. M., Worrall, J. L., and Morris, R. G. (2016). How Far From the Tree Does the Apple Fall? Field Training Officers, Their Trainees, and Allegations of Misconduct. *Crime & Delinquency*, 62(6):821–839. Publisher: SAGE Publications Inc.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.

Grawert, A. and Kim, N. (2023). Myths and Realities: Understanding Recent Trends in Violent Crime | Brennan Center for Justice.

Harris, M. C., Park, J., Bruce, D. J., and Murray, M. N. (2017). Peacekeeping Force:

Effects of Providing Tactical Equipment to Local Law Enforcement. *American Economic Journal: Economic Policy*, 9(3):291–313.

Heller, S. B., Jakubowski, B., Jelveh, Z., and Kapustin, M. (2022). Machine Learning Can Predict Shooting Victimization Well Enough to Help Prevent It. Technical Report w30170, National Bureau of Economic Research.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and Manufacturing TFP in China and India*. *The Quarterly Journal of Economics*, 124(4):1403–1448.

Jabri, R. (2021). Algorithmic Policing. *SSRN Journal*.

Kapustin, M., Neumann, T., and Ludwig, J. (2022a). POLICING AND MANAGEMENT.

Kapustin, M., Neumann, T., and Ludwig, J. (2022b). Policing and Management.

Kim, T. (2019a). Facilitating Police Reform: Body Cameras, Use of Force, and Law Enforcement Outcomes.

Kim, T. (2019b). Facilitating police reform: Body cameras, use of force, and law enforcement outcomes. *Use of Force, and Law Enforcement Outcomes (October 23, 2019)*.

Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., and Mullainathan, S. (2018). Human Decisions and Machine Predictions*. *The Quarterly Journal of Economics*, 133(1):237–293.

Leslie, E. and Wilson, R. (2020). Sheltering in place and domestic violence: Evidence from calls for service during COVID-19. *Journal of Public Economics*, 189:104241.

Levitt, S. D. (1997). Using Electoral Cycles in Police Hiring to Estimate the Effect of Police on Crime. *The American Economic Review*, 87(3):270–290. Publisher: American Economic Association.

Manes, J. (2021). The state of illinois v. michael williams. *End Police Surveillance*.

Mares, D. and 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.

Mastrobuoni, G. (2020). Crime is Terribly Revealing: Information Technology and Police Productivity. *The Review of Economic Studies*, 87(6):2727–2753.

Mello, S. (2019). More COPS, less crime. *Journal of Public Economics*, 172:174–200.

OIG (2023). Chicago police department 911 response time data collection and reporting. *Chicago, IL: City of Chicago Office of Inspector General*.

Olea, J. L. M. and Pflueger, C. (2013). A Robust Test for Weak Instruments. *Journal of Business & Economic Statistics*, 31(3):358–369. Publisher: ASA Website _eprint: https://doi.org/10.1080/00401706.2013.806694.

Piza, E. L., Hatten, D. N., Carter, J. G., Baughman, J. H., and Mohler, G. O. (2023). Gunshot Detection Technology Time Savings and Spatial Precision: An

Exploratory Analysis in Kansas City. *Policing: A Journal of Policy and Practice*, 17:paac097.

Rambachan, A. and Roth, J. (2023). A More Credible Approach to Parallel Trends. *The Review of Economic Studies*, page rdad018.

Rivera, R. (2023). Release, detain or surveil? the effects of electronic monitoring on defendant outcomes. *Unpublished manuscript, Columbia University*.

Rivera, R. G. and Ba, B. A. (2023). The Effect of Police Oversight on Crime and Misconduct Allegations: Evidence from Chicago. *The Review of Economics and Statistics*, pages 1–45.

Schmidheiny, K. and Siegloch, S. (2023). On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization. *Journal of Applied Econometrics*, 38(5):695–713. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.2971.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Topper, M. (2023). *panelsummary: Create Publication-Ready Regression Tables with Panels*. https://github.com/michaeltopper1/panelsummary, https://michaeltopper1.github.io/panelsummary/.

Weisburd, S. (2021). Police Presence, Rapid Response Rates, and Crime Prevention. *The Review of Economics and Statistics*, 103(2):280–293.

Weisburst, E. K. (2019). Safety in Police Numbers: Evidence of Police Effectiveness from Federal COPS Grant Applications. *American Law and Economics Review*, 21(1):81–109.

Wilde, E. T. (2013). Do Emergency Medical System Response Times Matter for Health Outcomes? *Health Economics*, 22(7):790–806. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/hec.2851.

Williams, J. and Weatherburn, D. (2022). Can Electronic Monitoring Reduce Reoffending? *The Review of Economics and Statistics*, 104(2):232–245.

Zamoff, M. E., Greenwood, B. N., and Burtch, G. (2022). Who Watches the Watchmen: Evidence of the Effect of Body-Worn Cameras on New York City Policing. *The Journal of Law, Economics, and Organization*, 38(1):161–195.

8 Figures

[This page intentionally left blank]

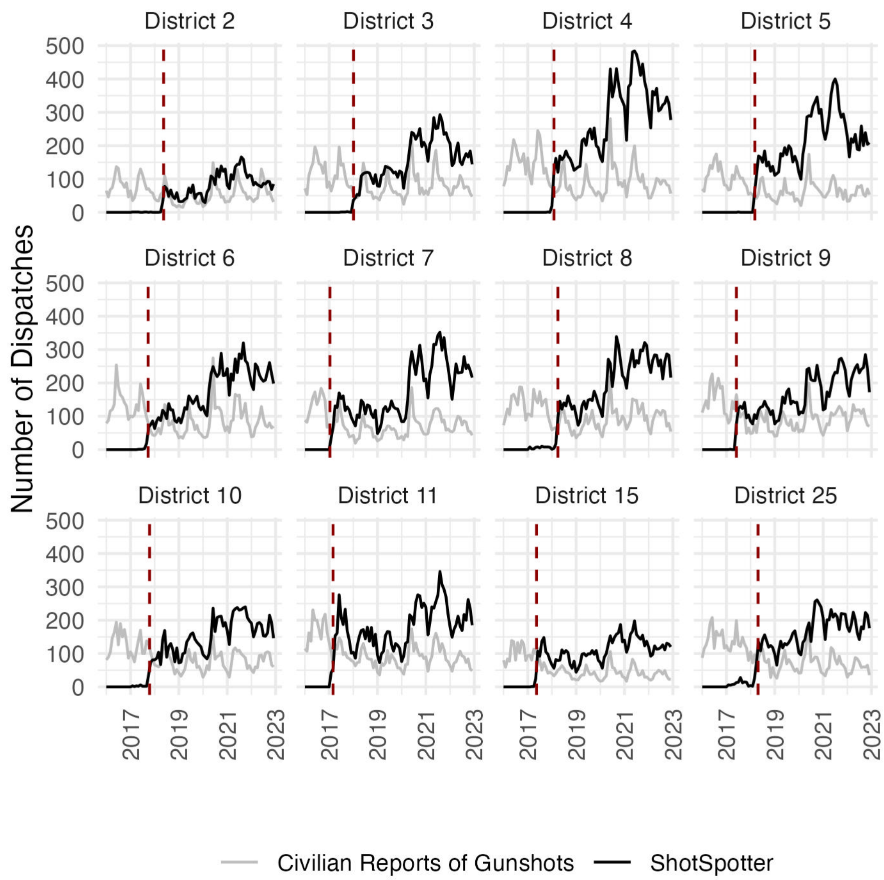

District	Enactment
2	2018-05-16
3	2018-01-04
4	2018-02-01
5	2018-03-07
6	2017-09-24
7	2017-01-13
8	2018-04-01
9	2017-06-01
10	2017-10-16
11	2017-03-01
15	2017-05-13
25	2018-04-24

Figure 1: ShotSpotter Alert Trends and Enactment Dates

Note: This figure depicts police districts that are implemented with ShotSpotter technology. Months are on the x-axis, while the y-axis is the number of ShotSpotter dispatches aggregated to the monthly level. The table on the right shows the corresponding implementation date for ShotSpotter technology. In Chicago, 12 of the 22 police districts have ShotSpotter technology. The dashed red line shows the implementation dates used in the main results. In some cases, the implementation date we use differs from the date given from the Chicago Police Department, since the ShotSpotter dispatches data does not align. Analysis using public records date is shown in Appendix Figure F2. Prior to implementation, some districts may observe some ShotSpotter dispatches if sensors in a neighboring district detect gunshots from afar. However, this is a rare occurrence.

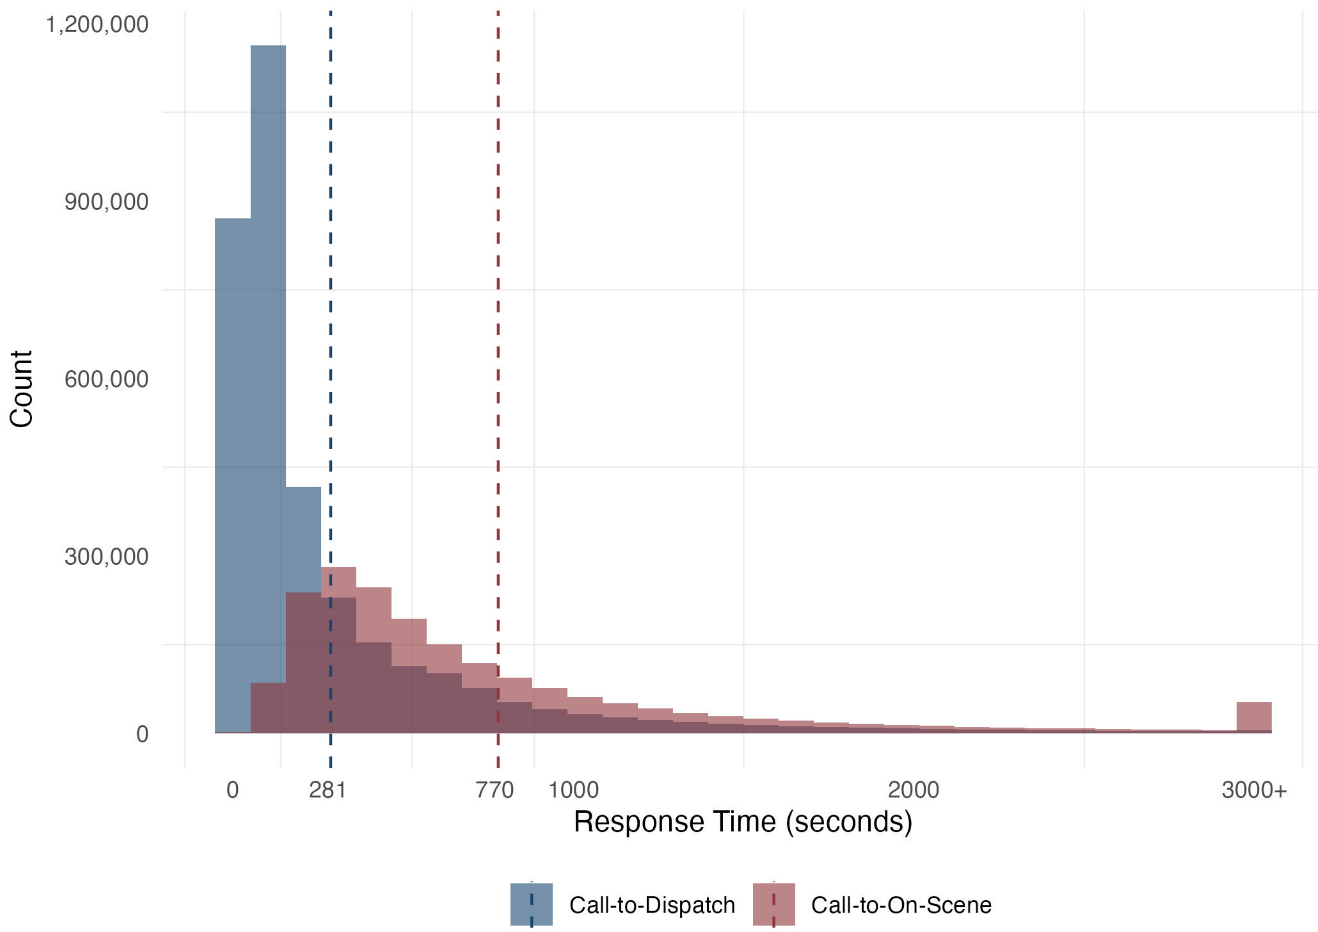

Figure 2: Distribution of Outcome Variables

Note: The two plotted variables are Call-to-Dispatch and Call-to-On-Scene. Call-to-Dispatch is the time from a 911 call to when a police officer is dispatched to the crime scene. Call-to-On-Scene is the time from a 911 call to the time a police officer arrives at the scene of the reported crime. This sample excludes outliers that are greater than three standard deviations from the mean for each outcome. Observations with response times higher than 3000 seconds are binned. However, the main results remain consistent when including these outliers, as shown in Appendix Figure F2. The dashed blue line represents the mean of Call-to-Dispatch time, while the dashed red line represents the mean of Call-to-On-Scene time.

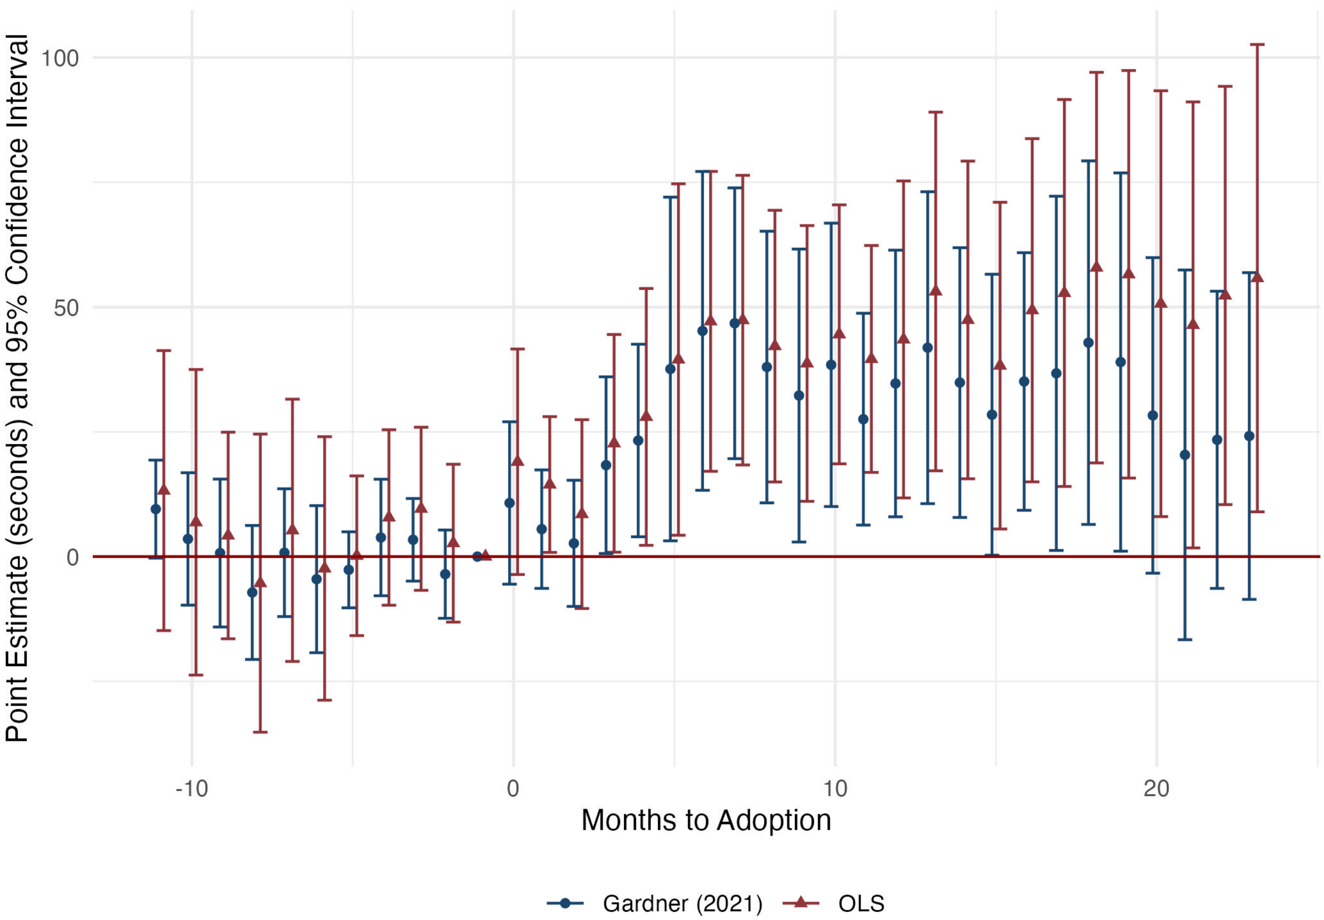

Figure 3: Event Study (Call-to-Dispatch)

Note: This figure shows the event study as specified in Equation 2 for Call-to-Dispatch times. Call-to-Dispatch is the amount of time from a 911 call to a police officer being dispatched to the crime scene. The x-axis denotes the number of months pre-/post-adoption of ShotSpotter technology. The y-axis denotes the 95% confidence intervals and point estimates (in seconds). The red error-bars/points represent confidence intervals/point estimates from OLS estimation while the blue are using the Gardner (2021) two-stage difference-in-difference estimator, which is robust to heterogeneous treatment effects in staggered adoptions. All pre-/post-periods are relative to the month before ShotSpotter adoption. Twelve pre-periods (24 post-periods) are estimated, but only 11 pre-periods (23 post-periods) are reported, as the -12 (+24) is a binned endpoint. Controls match the preferred specification. Standard errors are clustered at the district level.

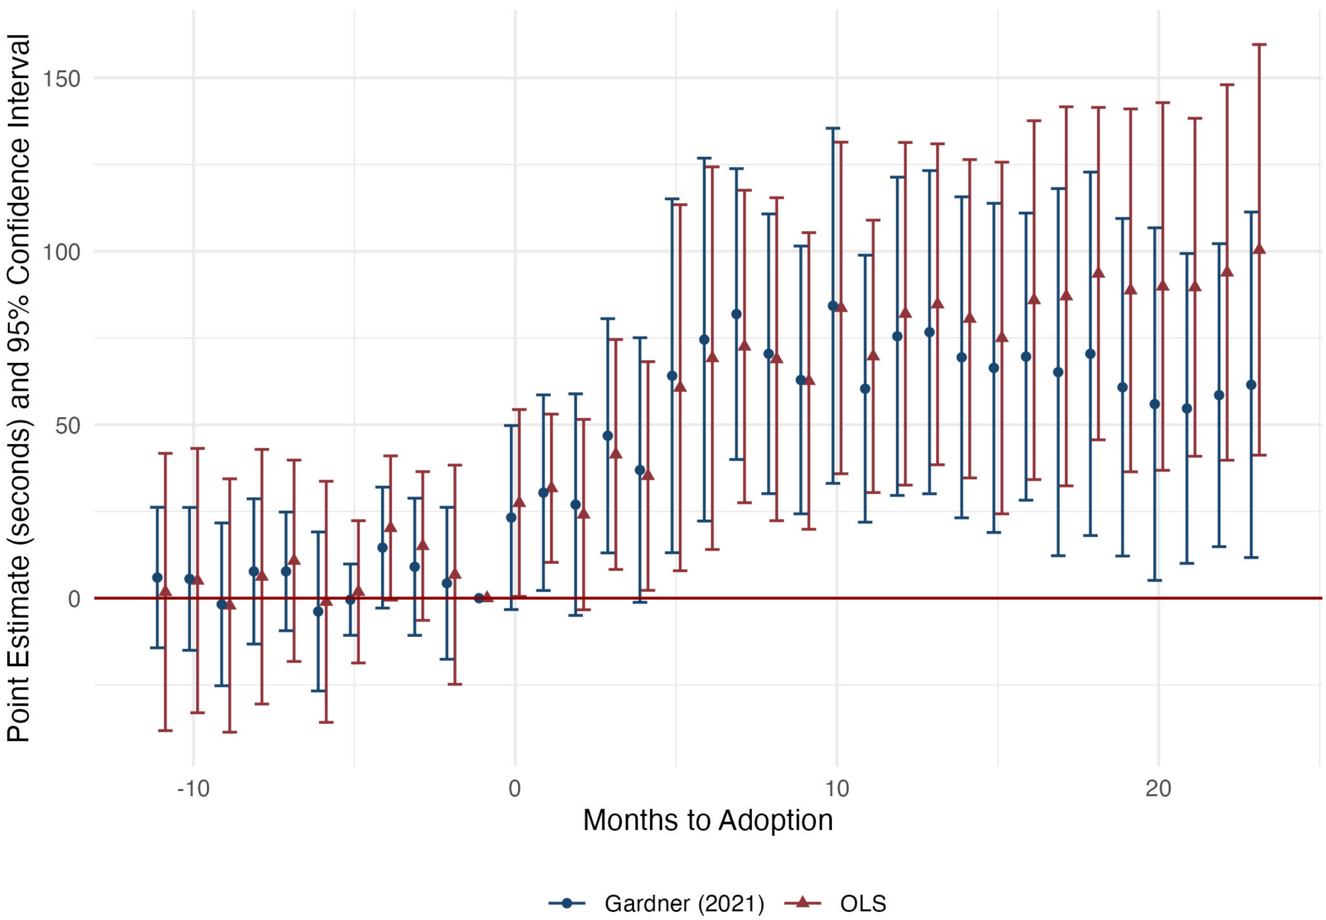

Figure 4: Event Study (Call-to-On-Scene)

Note: This figure shows the event study as specified in Equation 2 for Call-to-On-Scene times. Call-to-On-Scene is the amount of time from a 911 call to a police officer arriving to the crime scene. The x-axis denotes the number of months pre-/post-adoption of ShotSpotter technology. The y-axis denotes the 95% confidence intervals and point estimates (in seconds). The red error-bars/points represent confidence intervals/point estimates from OLS estimation while the blue are using the Gardner (2021) two-stage difference-in-difference estimator, which is robust to heterogeneous treatment effects in staggered adoptions. All pre-/ post-periods are normalized by the month before ShotSpotter adoption. Twelve pre-periods (24 post-periods) are estimated, but only 11 pre-periods (23 post-periods) are reported, as the -12 (+24) is a binned endpoint. Controls match the preferred specification. Standard errors are clustered at the district level.

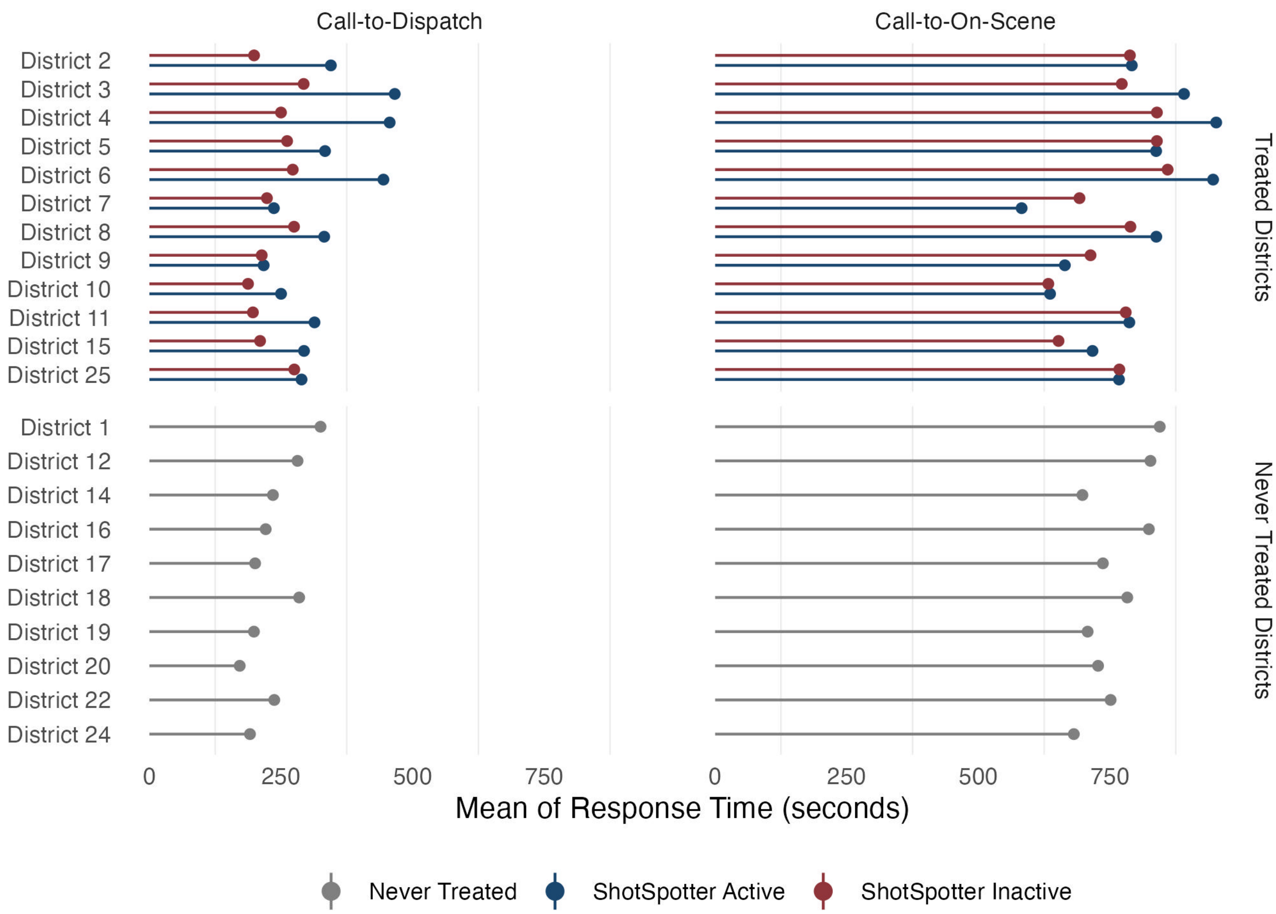

Figure 5: Average Outcomes in Police Districts

Note: Each police district is plotted on the y-axis, and the average of Call-to-Dispatch and Call-to-On-Scene (seconds) is on the x-axis. In the top panel, police districts that receive ShotSpotter technology are plotted. In the bottom panel, police districts that never receive ShotSpotter are plotted. All ShotSpotter-implemented districts have two distinctions: ShotSpotter Active and ShotSpotter Inactive. The red lines correspond to periods prior to ShotSpotter implementation, and the blue bars correspond to post-implementation. There are 12 of 22 police districts in Chicago that receive ShotSpotter technology.

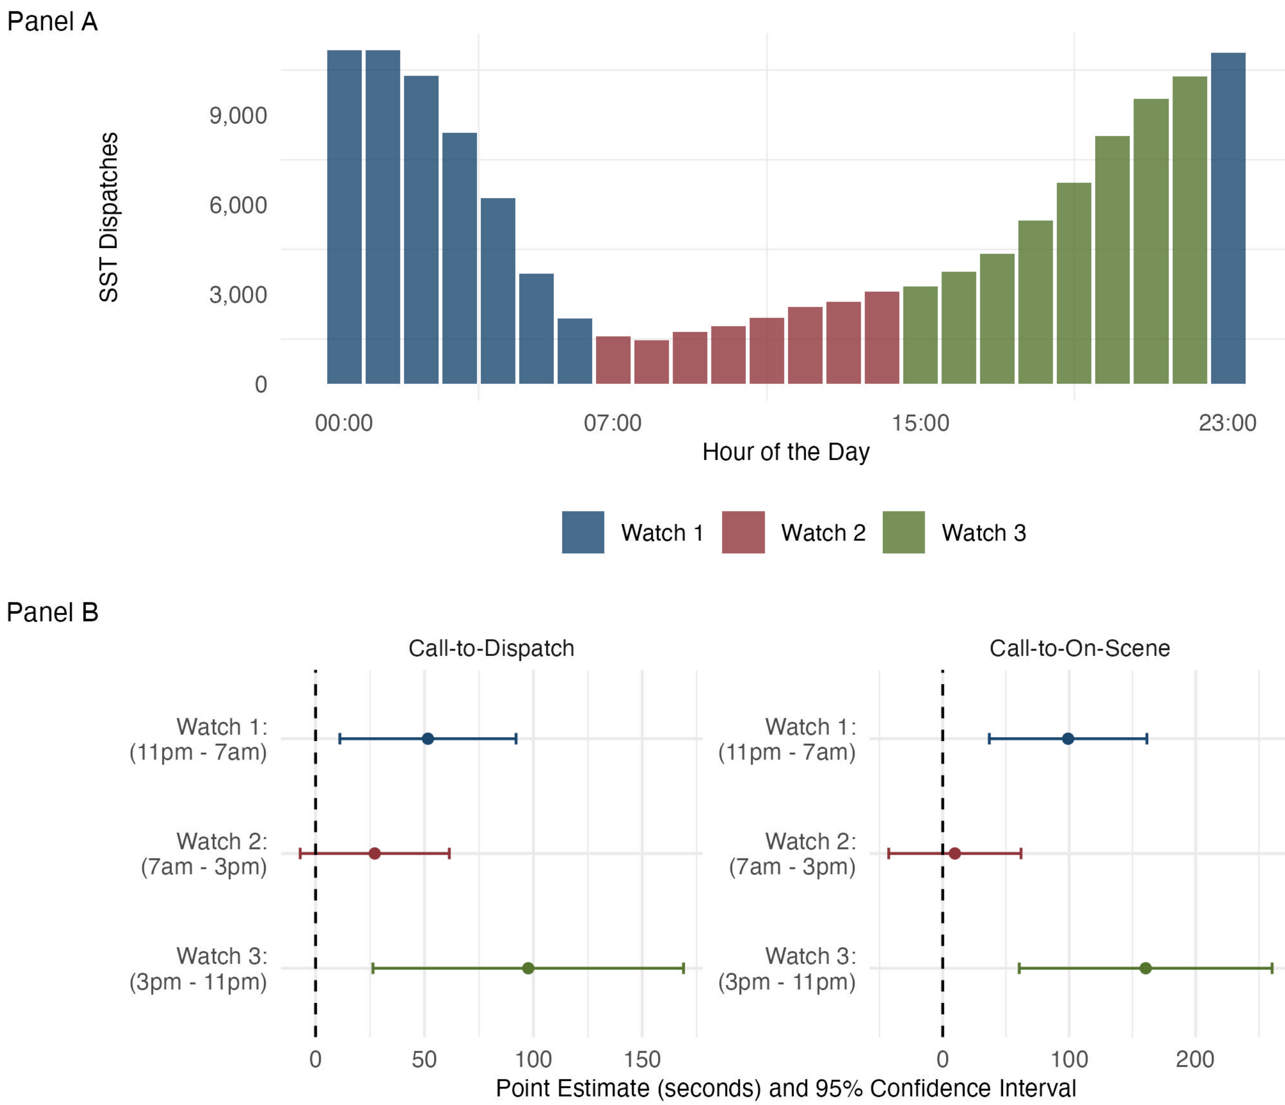

Figure 6: Effect of ShotSpotter by Officer Watch Times

Note: This figure shows that in times when officers are responding to more ShotSpotter (SST) detections, their response times are slower. In Panel A, the number of ShotSpotter dispatches are plotted by the hour of occurrence. The y-axis is the number of ShotSpotter dispatches, while the x-axis the hour of the day. In Panel B, Call-to-Dispatch and Call-to-On-Scene estimates using the specification in Equation 1 are shown along with the 95% confidence intervals, split by officer watch. There are three main watches in Chicago: Watch 1 (11:00pm-7:00-am), Watch 2 (7:00am-3:00pm), and Watch 3 (3:00pm-11:00pm).

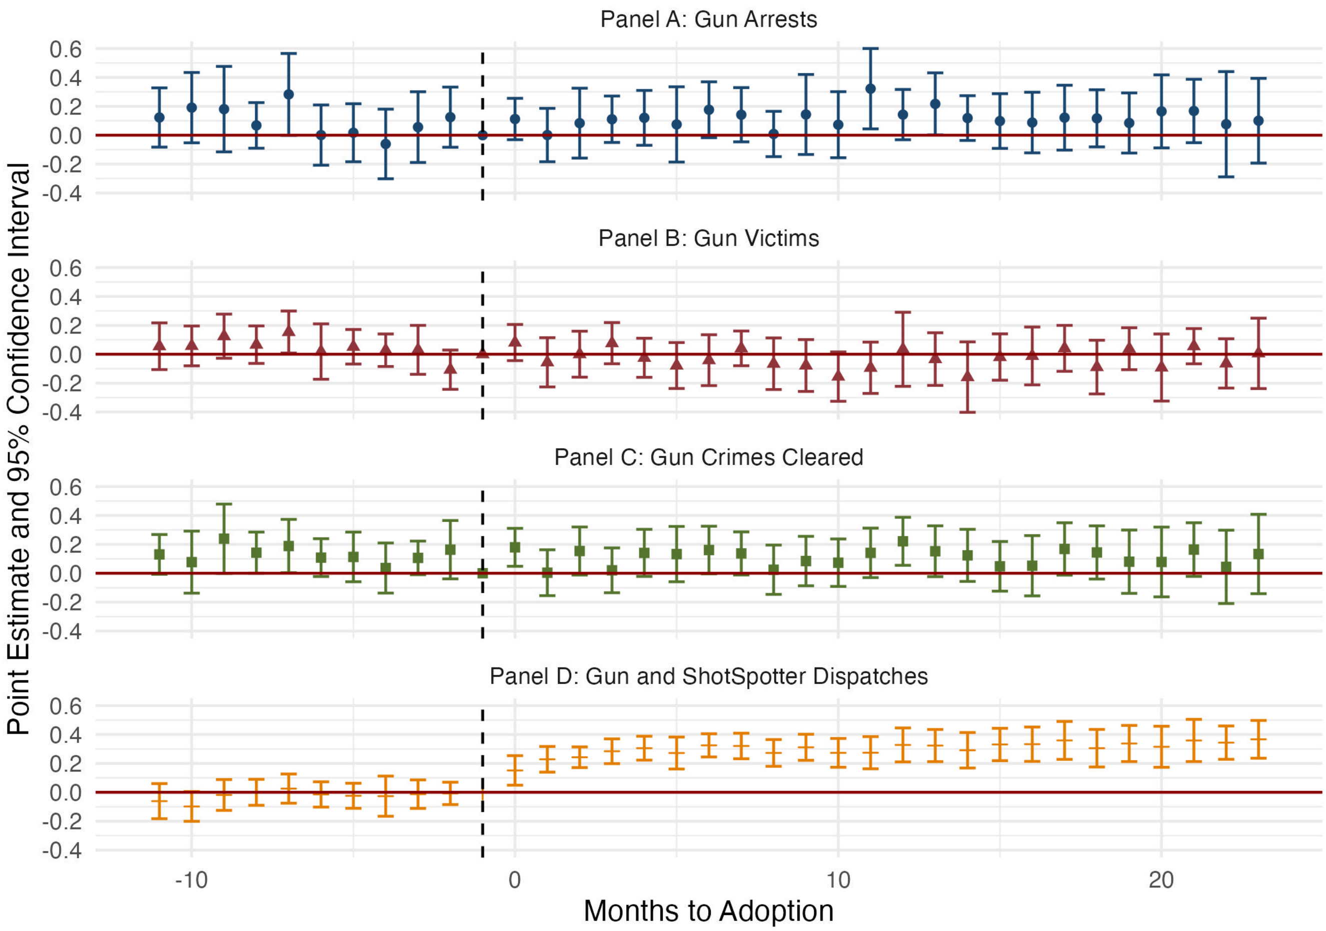

Figure 7: Event Study on Aggregate Outcomes (Poisson)

Note: This figure shows the event study as specified in Equation 3 for the district-level outcomes of gun-related arrests, victimization, and clearance. All outcomes are counts, and the estimation method is Poisson regression. Panel D shows the outcome of the number of gun-related 911 dispatches (i.e., 'person shot', 'person with gun', and 'shots fired') *combined* with ShotSpotter dispatches. The x-axis denotes the number of months pre-/post-adoption of ShotSpotter technology. The y-axis denotes the 95% confidence intervals and point estimates. All pre-/ post-periods are normalized by the month before ShotSpotter adoption. Twelve pre-periods (24 post-periods) are estimated, but only 11 pre-periods (23 post-periods) are reported, as the -12 (+24) is a binned endpoint. Controls include day-by-month-by-year and district fixed effects. Standard errors are clustered at the district level.

9 Tables

Table 1: Summary Statistics

	Mean	Std. Dev.	Min	Max	N
Panel A: Priority 1 Outcomes:					
Call-to-Dispatch	287.25	440.55	2.00	3,111.00	3,453,655
	(4.79 mins)	(7.34 mins)	(0.03 mins)	(51.85 mins)	
Call-to-On-Scene	779.21	786.87	11.00	7,671.00	1,935,142
	(12.99 mins)	(13.11 mins)	(0.18 mins)	(127.85 mins)	
Arrest Made	0.03	0.16	0.00	1.00	3,453,655
Panel B: Secondary Outcomes:					
Call-to-Dispatch (Priority 2)	362.04	524.78	2.00	3,577.00	1,604,670
	(6.03 mins)	(8.75 mins)	(0.03 mins)	(59.62 mins)	
Call-to-On-Scene (Priority 2)	964.45	901.10	14.00	6,615.00	776,284
	(16.07 mins)	(15.02 mins)	(0.23 mins)	(110.25 mins)	
Call-to-Dispatch (Priority 3)	1,013.00	1,258.18	2.00	6,550.00	3,283,969
	(16.88 mins)	(20.97 mins)	(0.03 mins)	(109.17 mins)	
Call-to-On-Scene (Priority 3)	1,915.38	1,820.18	10.00	11,702.00	1,226,065
	(31.92 mins)	(30.34 mins)	(0.17 mins)	(195.03 mins)	
Panel C: Other Variables:					
Priority 1 911 Dispatches	72.90	24.59	8.00	223.00	3,453,655
ShotSpotter Dispatches	2.95	4.17	0.00	57.00	3,453,655
Officer Hours	1,342.61	393.43	231.00	6,558.10	3,453,655

Note:
Units are in seconds unless otherwise noted. Data is at the call-level. Call-to-Dispatch represents the amount of time from the 911 call to an officer dispatching to the scene. Call-to-On-Scene is the time from a 911 call to when an officer arrives on-scene. Priority 1 Call-to-On-Scene is missing approximately 45 percent of on-scene times, although the results remain consistent when we can observe both response times. This is discussed further in Appendix A. Arrest Made is and indicator equal to one if the 911 call resulted in an arrest. Priority 1 refers to an immediate dispatch, Priority 2 a rapid dispatch, and Priority 3 a routine dispatch. Priority 2 and 3 analysis, while not in the main text, is available in Appendix E. Priority 1 911 Dispatches is the number of Priority 1 dispatches at the district-day level. ShotSpotter Dispatches is the number of dispatches due to ShotSpotter detections. Importantly, ShotSpotter Dispatches is also at the district-by-day level and includes days in which ShotSpotter is not implemented. The average number of ShotSpotter dispatches on post-implementation days is approximately 6. The average daily number of ShotSpotter dispatches across Chicago once all 12 districts have implemented ShotSpotter is approximately 70. Note that New Years Eve/New Years Day/Fourth of July are excluded from the sample as these days correspond with high amounts of celebratory gunfire. Officer Hours are the number of working hours sworn police officers work at the district-day level.

Table 2: Effect of ShotSpotter on Response Times (OLS)

	(1)	(2)	(3)	(4)	(5)
Panel A: Call-to-Dispatch					
ShotSpotter Activated	61.319**	64.074**	66.409**	71.696***	61.363**
	(21.840)	(22.648)	(23.566)	(22.620)	(21.892)
Border District Activated					21.535
					(16.663)
Mean of Dependent Variable	287.246	287.246	271.420	287.246	287.246
Observations	3,453,655	3,453,655	1,935,142	3,453,623	3,453,655
Wild Bootstrap P-Value	0.012	0.018	0.014		0.015
Panel B: Call-to-On-Scene					
ShotSpotter Activated	96.750***	102.394***	102.394***	121.020***	100.664***
	(26.077)	(28.649)	(28.649)	(27.646)	(28.035)
Border District Activated					24.366
					(17.710)
Mean of Dependent Variable	779.213	779.213	779.213	779.213	779.213
Observations	1,935,142	1,935,142	1,935,142	1,935,116	1,935,142
Wild Bootstrap P-Value	0.003	0.007	0.004		0.009
FE: Day-by-Month-by-Year	X	X	X	X	X
FE: District	X	X	X	X	X
FE: Call-Type		X	X	X	X
FE: Hour-of-Day		X	X	X	X
Gardner (2021) Robust				X	

Note:

* p < 0.1, ** p < 0.05, *** p < 0.01

Standard errors are clustered by district. All coefficient estimates are in seconds. Shotspotter is activated in 12 of the 22 police districts in Chicago. Panel A shows results for Call-to-Dispatch while Panel B shows results for Call-to-On-Scene. Column 1 reports only time and group fixed effects. Column 2 reports the preferred specification from Equation 1, which includes hour-of-day and call-type fixed effects. Column 3 includes only observations where both Call-to-Dispatch and Call-to-On-Scene can be observed. Column 4 reports estimates using the Gardner (2021) estimator which is robust to heterogeneous treatment effects across groups and time periods in staggered designs. Due to its two-stage method, some observations are dropped if unable to predict values in the first stage. Column 5 includes Border District Activated which is an indicator for when a district is adjacent to a ShotSpotter implemented district. Wild cluster bootstrap p-values using 999 iterations are also reported as the number of clusters (22) is below the threshold of 30 put forth in Cameron et al. (2008). The bootstrap cannot be performed using the Gardner (2021) estimator.

Table 3: Effect of ShotSpotter on Response Times Mechanisms (OLS)

	ShotSpotter Rollout			ShotSpotter Dispatches
		Officer Availability		
	Pooled	> Median	<= Median	Pooled
	(1)	(2)	(3)	(4)
Panel A: Call-to-Dispatch				
ShotSpotter Activated	64.131***	34.500**	85.180***	
	(22.379)	(13.630)	(27.959)	
Number SST Dispatches				6.094***
				(1.513)
Mean of Dependent Variable	281.890	239.951	323.077	269.365
Observations	3,582,560	1,775,086	1,807,474	47,933
Panel B: Call-to-On-Scene				
ShotSpotter Activated	102.682***	59.706***	138.102***	
	(28.724)	(21.061)	(37.671)	
Number SST Dispatches				8.023***
				(1.842)
Mean of Dependent Variable	770.863	711.409	827.843	770.462
Observations	1,997,102	977,332	1,019,770	47,932
FE: Day-by-Month-by-Year	X	X	X	X
FE: District	X	X	X	X
FE: Call-Type	X	X	X	
FE: Hour-of-Day	X	X	X	

Note:
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Standard errors are clustered by district. ShotSpotter Activated is a binary equal to one when a district has ShotSpotter technology (extensive margin). Number SST Dispatches refers to the number of ShotSpotter dispatches that occur within a district-day (intensive margin). All coefficient estimates are in seconds. Panel A reports results for Call-to-Dispatch while Panel B reports results for Call-to-On-Scene. Officer availability is measured by number of officer hours within a district-day. Column 2 corresponds to district-days that have officer hours above their district median (more officer availability), while Column 3 corresponds to district-days that have officer hours below their district median (less officer availability). Analyses for Columns 1-3 are on the extensive margin, and utilze call-level data. The coefficients for these analyses are interpreted as average effects. Analysis for Column 4 is on the intensive margin, and the data is aggregated to the district-day level. The coefficients of interest for Column 4 are interpreted as marginal effects. We aggregate to the district-day since the number of ShotSpotter dispatches is measured at the district-day. Because of this, we cannot use call-level data to correctly identify the marginal effects. Moreover, we restrict the sample to only post-implementation days for treated districts to ensure that only the intensive margin, rather than extensive margin, is identified. Further explanation of this model is given in Section 5.3.

Table 4: Effect of ShotSpotter on 911 Call Resolutions (OLS)

		Most Frequent Final 911 Dispositions		
	Arrest Made	Other Police Service	No Person Found	Peace Restored
	(1)	(2)	(3)	(4)
Panel A: Reduced Form				
ShotSpotter Activated	-0.238***	1.056*	0.751**	-0.558**
	(0.068)	(0.592)	(0.316)	(0.217)
Mean of Dependent Variable	2.553	41.047	21.367	6.283
Observations	3,453,655	3,453,655	3,453,655	3,453,655
Wild Bootstrap P-Value	0.002	0.039	0.002	0.001
Panel B: 2SLS (Second Stage)				
Call-to-Dispatch	-0.004**	0.016	0.012**	-0.009*
	(0.002)	(0.013)	(0.005)	(0.004)
Mean of Dependent Variable	2.553	41.047	21.367	6.283
Observations	3,453,655	3,453,655	3,453,655	3,453,655
Eff F-Stat	8.004	8.004	8.004	8.004
AR Conf.Int	[-0.012, -0.002]	[-0.001, 0.092]	[0.003, 0.029]	[-0.031, -0.002]
AR P-Value	0.000	0.074	0.018	0.010
FE: Day-by-Month-by-Year	X	X	X	X
FE: District	X	X	X	X
FE: Call-Type	X	X	X	X
FE: Hour-of-Day	X	X	X	X

Note:
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Standard errors are clustered by district. All coefficient estimates and means are in percentages. The dependent variable in Column 1 is an indicator equal to one if a 911 call resulted in an arrest. The dependent variable in Columns 2-4 is an indicator equal to one if a 911 call resulted in Other Police Service (Column 2), No Person Found (Column 3), or Peace Restored (Column 4). Columns 2-4 report the three most frequent 911 final dispositions: Other Police Service, No Person Found, and Peace Restored. The final disposition is the final result of what happened on the 911 call. Panel A shows the reduced form estimates of the ShotSpotter implementation, while Panel B shows the second stage of a 2SLS regression where Call-to-Dispatch time is the endogenous variable, and ShotSpotter implementation is the instrument. Hence, Panel B estimates show the marginal effect of an extra second of Call-to-Dispatch on arrest probability for calls that are induced to have longer Call-to-Dispatch times by ShotSpotter (compliers). First stage estimates are shown in the main results table where ShotSpotter implementation results in ~60 second increase for Call-to-Dispatch times on average. In Panel A, Wild cluster bootstrap p-values using 999 replications are also reported since the number of clusters (22) is below the threshold of 30 put forth in Cameron et al. (2008). In Panel B, we report the effective F-statistic (Eff F-Stat) from Olea and Pflueger (2013), and the confidence intervals and corresponding p-values from the Anderson and Rubin (1949) test which is robust to weak instruments.

Appendix A Missing Call-to-On-Scene Data

In this appendix, we conduct analyses regarding the notable amount of data missing for one of the key outcome variables, Call-to-On-Scene. Recall that Call-to-On-Scene denotes the time interval between a 911 call and an officer's arrival at the scene of the incident. While we find suggestive evidence that missing Call-to-On-Scene times are correlated with ShotSpotter implementation, this section outlines several reasons to maintain confidence in the main results despite this limitation.

A.1 Reasons for Missing Data

First, we note that the underlying reason behind a missing Call-to-On-Scene entry is an officer's failure to report to the dispatcher that they have arrived on-scene. This could be due to an officer forgetting to report, or more likely, an officer being immediately engaged on-scene. Importantly, we provide suggestive evidence that the latter is happening more frequently post-implementation of ShotSpotter due to officers being more time-constrained.

In Panel A of Appendix Table A1, we estimate the preferred specification from Equation 1 on an indicator for a missing Call-to-On-Scene time and find suggestive evidence of a correlation. Column 1 of Panel A reports a 3.5 percentage point increase in the likelihood of missing Call-to-On-Scene when ShotSpotter is implemented, which is statistically significant at the 10% level. However, Columns 2 and 3 show that this effect is driven by times in which there are fewer officers on duty, implying that ShotSpotter may be straining officers' time allotment. For instance, if an officer feels they have fallen behind, they may disregard relaying to the dispatcher that they have arrived to the scene. If this is the case, then the missing on-scene times may be larger than the non-missing times, thereby suggesting that the main results are biased downward.

A.2 Impact on Call-to-Dispatch Times

Second, we examine the impact of missing data on Call-to-Dispatch times—the time from a 911 call to when an officer is dispatched to the crime scene. Notably, Call-to-Dispatch times, a mechanism underlying Call-to-On-Scene times as discussed in Section 5, are 100% reported.

To begin, we supplement Equation 1 with an interaction between ShotSpotter implementation (ShotSpotter Activate) and an indicator for missing Call-to-On-Scene times (Missing On-Scene).[32] In doing so, we test whether there are differences in the effect of ShotSpotter on Call-to-Dispatch times between cases with missing and no missing data. Panel B of Appendix Table A1 reports no significant change in Call-to-Dispatch times when there is missing Call-to-On-Scene data. As shown across Columns 1-3, there is little evidence that Call-to-Dispatch times differ in a missing data case. Specifically, the coefficient on the interaction term is small and statistically insignificant. This result instills confidence that officers are likely still arriving on-scene at later times even in missing data cases, as there appears to be no change in Call-to-Dispatch times when on-scene times are missing.

A.3 Consistent Trends

Last, given that Call-to-Dispatch times are fully reported and there is no change when Call-to-On-Scene times are missing, we plot the event study coefficients from Figures 3 and 4 in Appendix Figure A1 which shows that there is a consistent time trend for each outcome variable. The convergence in trends reinforces the notion that even when Call-to-On-Scene data is absent, officers may still experience delays in reaching the scene due to slower dispatching procedures. This consistent pattern underscores the reliability of the

[32] The fixed effects are also interacted with Missing On-Scene.

Call-to-On-Scene findings.

Table A1: Analysis of Missing Call-to-On-Scene Data (OLS)

		Officer Availability	
	Pooled	> Median	<= Median
	(1)	(2)	(3)
Panel A: Missing Call-to-On-Scene			
ShotSpotter Activated	0.035*	0.029	0.039*
	(0.019)	(0.019)	(0.021)
Mean of Dependent Variable	0.440	0.447	0.432
Observations	3,453,655	1,725,816	1,727,839
Panel B: Call-to-Dispatch			
ShotSpotter Activated	66.409***	34.739**	92.687***
	(23.293)	(14.126)	(30.051)
ShotSpotter Activated x Missing	0.101	0.895	-2.956
	(33.317)	(20.676)	(42.251)
Mean of Dependent Variable	287.246	243.869	330.572
Observations	3,453,655	1,725,816	1,727,839

Note:
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Standard errors are clustered by district. All coefficient estimates are in seconds. In Panel A, the table shows regressions on a binary variable equal to one if Call-to-On-Scene is missing. Columns 2 and 3 are split by district-day medians of officer hours. In Panel B, Call-to-Dispatch time, which contains no missing data, is estimated with an additional interaction term which interacts Call-to-Dispatch time with the indicator for whether on-scene time is missing. The coefficient estimate on this term shows that that there is no difference in Call-to-Dispatch time when there is missing on-scene data. Note that in these specifications, the fixed effects are also interacted to get a similar interpretation as if there were two separate regressions estimated. All controls utilized in these regressions are consistent with the preferred specification and are estimated using OLS.

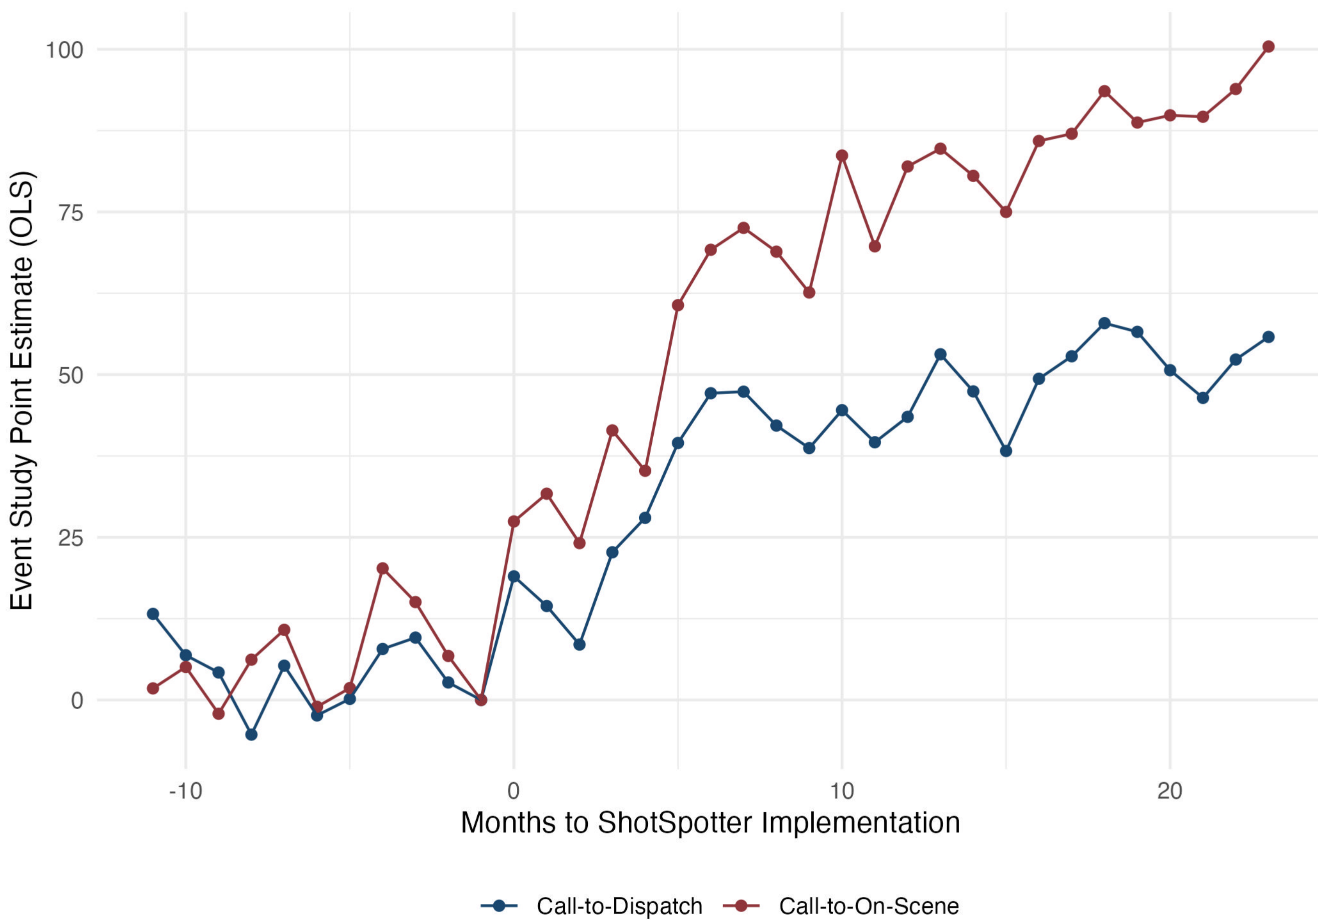

Figure A1: Event Study Point Estimates Trend
Note: This figure plots the point estimates of the event study specifications in Equation 2 for both Call-to-Dispatch (blue) and Call-to-On-Scene (red). In effect, this figure shows that the trends for each of these outcomes are similar. The y-axis denotes the point estimate in seconds, and the x-axis displays the number of months to ShotSpotter implementation. Recall that Call-to-Dispatch has no missing data, while Call-to-On-Scene is approximately 45 percent missing. This figure is intended to show that Call-to-Dispatch, a mechanism underlying slower on-scene times, has a similar trend to Call-to-On-Scene, suggesting that missing data may not be a substantial issue.

Appendix B Changes in Call Distribution

In this appendix, we discuss two sets of analyses to test the assumption that call composition is stable post-implementation of ShotSpotter. First, we test whether the number of 911 calls changes. Second, we test whether certain call classifications are more common. In each test, we find little evidence of a significant change in the number of calls or call composition.

B.1 Changes in the Number of 911 Calls

First, we test whether the number of 911 calls changes significant post-implementation. It is imperative that the number of calls is not increasing since this could result in the volume of calls causing response time delays rather than ShotSpotter. To test this, we estimate an event-study model using Poisson estimation, similar to Equation 3:

$$NumberCalls_{dt} = \sum_{\substack{i=-12, \\ i\neq-1}}^{24} \beta^i \text{ShotSpotter}^i_{dt} + \delta_d + \gamma_t + \varepsilon_{dt} \tag{B0}$$

where $NumberCalls_{dt}$ is the number of 911 calls in district d at time t. Day-by-month-by-year (γ_t) and district fixed effects (δ_d) are also included. Given the (non-negative) count nature of the outcomes, we estimate Equation B0 using Poisson estimation in lieu of OLS. Standard errors are clustered at the district level.

Figure B1 plots the event-study estimates for four outcomes: all 911 calls (Panel A), Priority 1 calls (Panel B), Priority 2 calls (Panel C), and Priority 3 calls (Panel D). Panel A shows that there is no significant change in the total number of 911 calls post-implementation which suggests that officers have a constant workload of call-volume. Priority 2 and 3 calls also exhibit similar results. However, there is evidence of a *decline* in the number of Priority 1 calls as shown in Panel B. We emphasize two points in light of this result: first, the decline is likely economically insignificant, as the DiD estimate translates to an average of two-fewer calls—a sum that is unlikely to greatly affect officer availability—and second, a decline in the number of calls results in *fewer* officer responsibilities, which would result in *more* officer availability. As shown in Table 3, when there are more officers available, the effects of ShotSpotter on response times are decreased.

B.2 Changes in Call Composition

Second, in order to assess whether call composition is changing, we estimate whether the likelihood of each 911 call-type changes post-implementation. We do this using a variation of the baseline OLS specification (Equation 1):

$$CallType_{cdt} = \beta ShotSpotter_{dt} + \delta_d + \gamma \mathbb{X}_{f(t)} + \varepsilon_{cdt} \tag{B0}$$

where $CallType_{cdt}$ is a binary indicator equal to one for each call classification for call c in district d at time t.[33] All other components of the model remain the same.

Since there are some infrequent call classifications, we combine any call-type that does not appear in every year into one category, resulting in 114 final classifications. Next, we estimate 114 iterations of Equation B0 with a different call-type as the binary outcome variable. This results in 114 estimates of parameter β—the average change in the probability of a call-type occurring post-ShotSpotter implementation. Given the large number of hypothesis tests, we correct the p-values using the method outlined in Benjamini and Hochberg (1995).

Figure B2 plots the corrected p-values from the 114 estimations of Equation B0 and shows that there is little evidence of significant changes in call composition. In particular, only 3% of the call classifications show statistically significant changes in likelihood

[33] Recall that we use the final dispatch code as the classification for each call-type.

which is in-line with expectations.

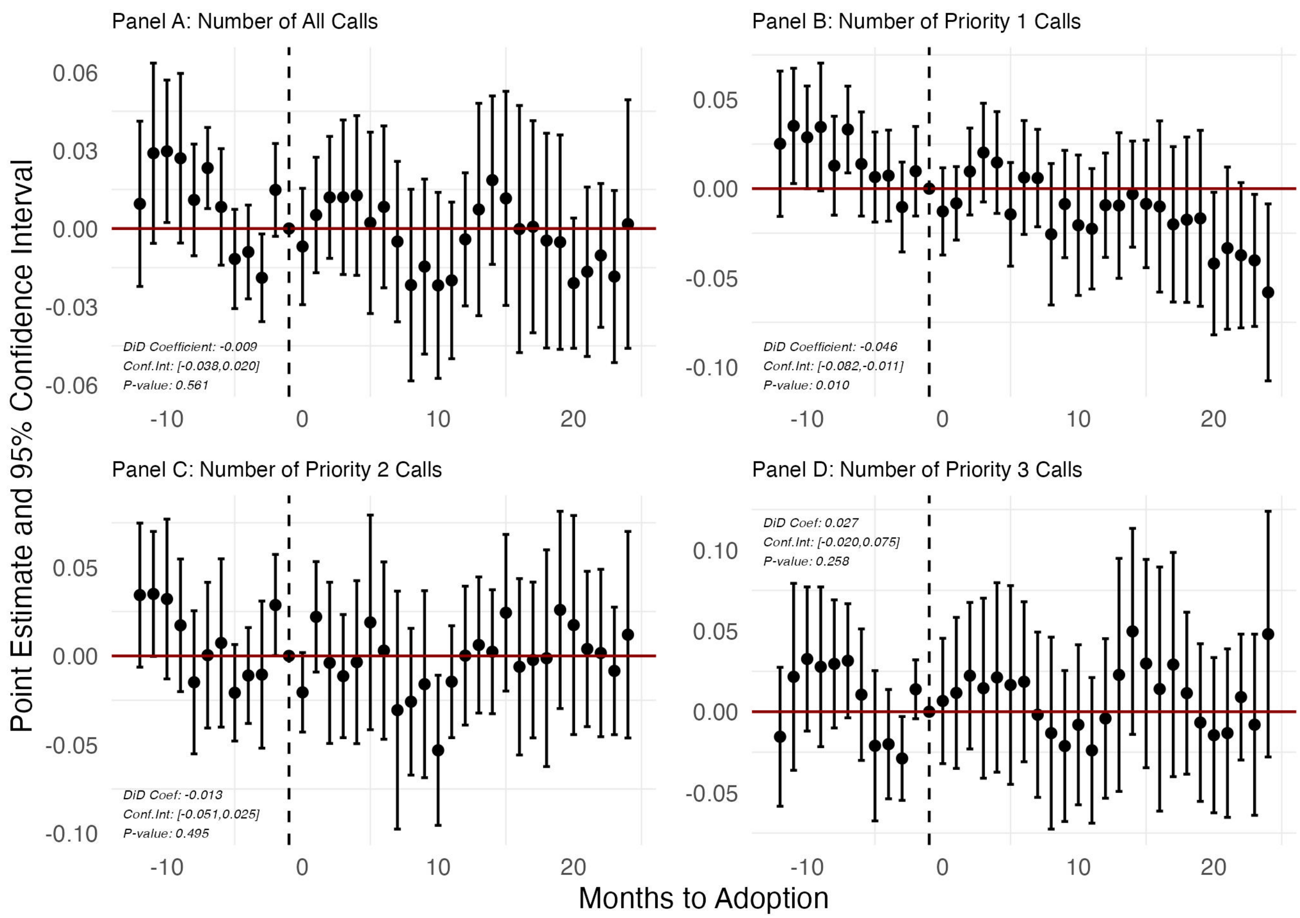

Figure B1: Event Study for Number of 911 Calls (Poisson)

Note: This figures shows the event study for the number of all 911 calls (Panel A), Priority 1 calls (Panel B), Priority 2 calls (Panel C), and Priority 3 calls (Panel D). Coefficients are estimated using Poisson regression due to the non-negative count nature of the data. Standard errors are clustered at the district level. The error bars represent the 95% confidence intervals and the point estimates are monthly bins. All pre-/post-periods are normalized by the month before ShotSpotter adoption. Twelve periods are estimated, but only 11 pre-periods and 23 post-periods are reported as the -12 and +24 are binned endpoints. All difference-in-differences estimates and corresponding confidence intervals and p-values are also reported in each panel.

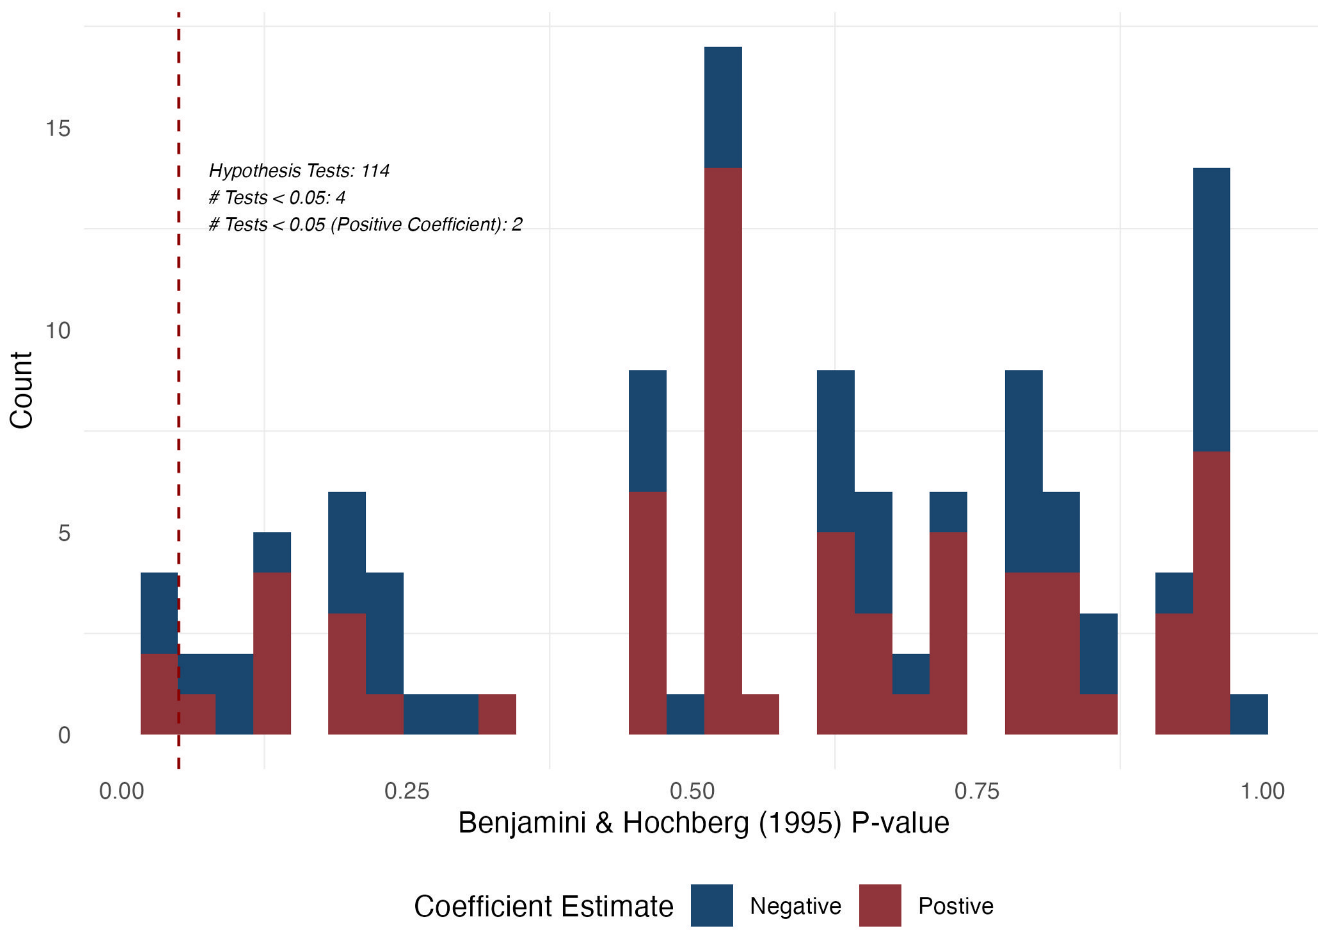

Figure B2: Distribution of P-values for Call-types (OLS)

Note: This figure plots the Benjamini and Hochberg (1995) p-values (i.e., corrected for multiple hypothesis testing) for 114 regressions using Appendix Equation B0. From the 114 regressions, only four tests are significant at the 5% level (3%). Red coloring denotes p-values which correspond to a positive coefficient estimate, meaning the call-type is more likely to occur post-implementation. On the other hand, blue coloring denotes p-values which correspond to a negative coefficient, meaning the call-type is less likely to occur pre-implementation. The dashed red line denotes the 0.05 value on the x-axis. The y-axis denotes the count of p-values.

Appendix C Coinciding Initiatives

In this appendix, we discuss two initiatives that were implemented in the Chicago Police Department (CPD) near the timing of ShotSpotter: Strategic Decision Support Centers and Body-worn Cameras. While neither of these exactly coincide with ShotSpotter implementation, we perform several sets of analyses to mitigate concerns that these are confounding the results.

C.1 Strategic Decision Support Centers

Strategic Decision Support Centers (SDSC) are command and control centers created to give police officers more awareness of what is occurring in their districts, and decide on responses. The main objective of SDSCs is to reduce crime, improve officer safety, and reduce service times. Each SDSC has staff members which include a dedicated supervisor (usually a sworn officer who is a lieutenant or sergeant) and a data analyst.

These support centers act as a hub for all of Chicago's policing technologies, whereby they can relay real-time information to police officers in the field. In particular, these centers are constantly analyzing data from automated license plate readers, social media monitoring, police observation cameras and devices, and geospatial predictive police software (Hunchlab).[34] While most of these technologies have already been in utilization by the CPD prior to SDSCs,[35] the Hunchlab software is implemented at the exact timing of an SDSC.

Importantly, as described in further detail in Kapustin et al. (2022b), the implementation of an SDSC did not include an infusion of officers in the form of new officers

[34]Hunchlab was bought by ShotSpotter in fall of 2018 and is now known as ShotSpotter Missions. We refrain from using this terminology, as it might be confusing to a reader.

[35]Automated license plate readers began as early as 2006, social media monitoring as early as 2014, and police observation cameras and devices as early as 2003.

being hired, existing officers being relocated, or officers working extra hours. Moreover, SDSCs were told not to implement new policing strategies, but to only assist department members with crime forecasting.

C.1.1 SDSC Technology Effect on Police Patrolling

There may be reason to suspect that Hunchlab, the geospatial predictive policing technology implemented with SDSCs, affects police response times. Hunchlab functions by creating location hot-spots in which police officers are supposed to visit more frequently in their patrols. These hot-spots are places where Hunchlab algorithms are predicting crime to occur. Hence, Hunchlab could affect response times by placing officers closer (or farther) to reported incidents of crime, or by placing them in areas where they are more likely to make arrests/stops and be unavailable for dispatch.

Despite this potential limitation, a thorough analysis of this exact technology is provided in Kapustin et al. (2022b). Specifically, they find that officers take Hunchlab recommendations in only two police districts (District 7 and District 9). The null results they report in the other police districts are attributed to commanders or officers disregarding the software's suggestions.

C.1.2 Main Results Controlling for SDSCs

In this subsection, we re-estimate the main specification and corresponding event studies on Call-to-Dispatch and Call-to-On-Scene times, as well whether an arrest was made on the call (Arrest Made) while controlling for the SDSC implementation. SDSCs are implemented in a district-by-district roll-out that is similar to ShotSpotter's implementation. Appendix Table C1 reports the districts and corresponding dates of their implementation. On average, SDSCs are implemented 76 days *prior* to ShotSpotter. Moreover, only 57% of districts with an SDSC has ShotSpotter.

Appendix Table C2, shows consistent findings of the effects of ShotSpotter on response times while controlling for the roll-out of SDSCs. In Columns 1, we use the OLS estimator while in Column 2, we use the Gardner (2021) estimator to account for possible treatment heterogeneity across groups and over time given the staggered design. In Panel A, Call-to-Dispatch times show increases of approximately one-minute, while in Panel B, Call-to-On-Scene times exhibit increases of two-minutes—each statistically significant at the 1% level. Moreover, Panel C shows that the likelihood of an arrest decreases by approximately 0.246 percentage points (9%). On the other hand, there appears to be a decrease in response times due to the SDSC roll-out on both Call-to-Dispatch and Call-to-On-Scene times, suggesting that the Hunchlab technology in the SDSCs is not incapacitating officers' availability, and that the SDSCs may provide some efficiency gains with the reorganization of intelligence software.

In Columns 3 and 4 of Appendix Table C2, we re-estimate the specifications from Columns 1 and 2, but exclude police districts 7 and 9 which have been found to have changes in police patrolling behavior following the SDSC rollout (Kapustin et al., 2022b). In doing so, we focus the analysis on districts officers do not take Hunchlab's recommendations and patrolling patterns are likely to be more stable. The results for both Call-to-Dispatch, Call-to-On-Scene, and arrests are consistent with the main findings.

Next, we estimate the event study specifications in Equation 2 while controlling for SDSC implementation. Appendix Figures C1, C2, and C3 plot the event studies for Call-to-Dispatch, Call-to-On-Scene, and Arrest Made using both the OLS estimator (red) and the Gardner (2021) estimator (blue). In each plot, the standard errors get significantly larger relative to the models without SDSC controls. This is likely due to the proximity of both ShotSpotter implementation and SDSCs. However, despite these larger standard errors, the pre-period shows no visual evidence of a violation of the common trends assumptions, and the post period results appear similar to the main event studies in Figures

3, 4, and F5.

C.2 Body-Worn Cameras

In this subsection, we show that controlling for the body-worn camera (BWC) implementation in Chicago has no effect on the response time results. As mentioned in the main text, the district implementation of BWCs differs by 283 days on average (see Appendix Table C1) from the ShotSpotter roll-out (see Appendix Table C1). Moreover, while body worn cameras have been found to affect complaints (Kim, 2019b; Braga et al., 2022; Zamoff et al., 2022; Ferrazares, 2023), arrests, and stops (Braga et al., 2022; Zamoff et al., 2022), there is little reason to suspect that they significantly affect an officer's ability to rapidly respond.

Columns 5 and 6 of Appendix Table C2 report the results for Call-to-Dispatch, Call-to-On-Scene, and Arrest Made (in percentage points) while controlling for BWC implementation. The results are consistent with the main findings, and the negative coefficient on BWC does not show any evidence of affecting response times.

Table C1: Implementation Dates of ShotSpotter/SDSC/BWC

District	ShotSpotter	SDSC	BWC	Difference SDSC	Difference BWC
2	2018-05-16	2018-03-01	2016-06-29	76 days	686 days
3	2018-01-04	2018-01-01	2017-11-06	3 days	59 days
4	2018-02-01	2018-01-01	2016-08-13	31 days	537 days
5	2018-03-07	2018-01-01	2017-11-20	65 days	107 days
6	2017-09-24	2017-03-15	2016-08-04	193 days	416 days
7	2017-01-13	2017-01-07	2017-05-01	6 days	108 days
8	2018-04-01	2018-03-01	2017-10-02	31 days	181 days
9	2017-06-01	2017-03-15	2016-08-18	78 days	287 days
10	2017-10-16	2017-03-15	2016-07-25	215 days	448 days
11	2017-03-01	2017-02-17	2017-06-05	12 days	96 days
15	2017-05-13	2017-03-15	2016-06-13	59 days	334 days
25	2018-04-24	2018-01-01	2017-12-04	113 days	141 days
1		2020-06-01	2017-03-10		
12		2018-03-01	2017-12-04		
14		2019-02-25	2016-06-01		
16			2017-11-20		
17		2019-02-25	2017-11-27		
18		2018-08-01	2017-03-31		
19		2019-02-01	2017-10-30		
20		2019-02-25	2017-10-23		
22		2019-02-25	2017-10-30		
24		2019-02-01	2017-10-16		

Note:
This table shows the implementation dates of ShotSpotter technology and Strategic Decision Support Centers (SDSC). SDSCs are implemented in similar, although not the same time period. The Difference column shows the number of days between the SDSC implementation and ShotSpotter activation. On average, this is approximately 73 days in districts that have both ShotSpotter and an SDSC. SDSCs contain many police prediction softwares, however, only Hunchlab, a location prediction software, is implemented in conjuction with these as the others had been previously used in Chicago prior to SDSCs. Hunchlab has been found to only change patrolling behaviors in districts 7 and 9 as discussed in Kapustin et al. (2022). Further robustness of the results including SDSC implementation dates as controls are shown in Appendix Table B2.

Table C2: Robustness of Estimates Controlling for Other Technologies (OLS)

	SDSC Controls				BWC Controls	
			Omitting Districts 7 and 9			
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Call-to-Dispatch						
ShotSpotter Activated	75.425***	71.454***	84.737***	90.120***	61.110***	71.461***
	(25.374)	(22.770)	(27.235)	(22.300)	(21.220)	(22.783)
SDSC Activated	-36.714**		-48.405**			
	(16.847)		(17.110)			
BWC Activated					-31.571	
					(20.928)	
Mean of Dependent Variable	287.246	287.246	294.523	294.523	287.246	287.246
Observations	3,453,655	3,453,623	3,084,625	3,084,600	3,453,655	3,453,623
Wild Bootstrap P-Value	0.002		0.007		0.003	
Panel B: Call-to-On-Scene						
ShotSpotter Activated	120.292***	120.721***	127.204***	146.018***	98.225***	121.014***
	(30.383)	(27.869)	(32.838)	(24.256)	(27.768)	(27.946)
SDSC Activated	-60.056***		-71.208***			
	(19.205)		(20.477)			
BWC Activated					-39.998	
					(26.159)	
Mean of Dependent Variable	779.213	779.213	799.595	799.595	779.213	779.213
Observations	1,935,142	1,935,116	1,708,425	1,708,405	1,935,142	1,935,116
Wild Bootstrap P-Value	0.001		0.002		0.003	
Panel C: Arrest Made						
ShotSpotter Activated	-0.246***	-0.256***	-0.257***	-0.266***	-0.236***	-0.252***
	(0.074)	(0.064)	(0.081)	(0.072)	(0.067)	(0.065)
SDSC Activated	0.026		0.011			
	(0.071)		(0.077)			
BWC Activated					0.019	
					(0.071)	
Mean of Dependent Variable	2.553	2.553	2.546	2.546	2.553	2.553
Observations	3,453,655	3,453,623	3,084,625	3,084,600	3,453,655	3,453,623
Wild Bootstrap P-Value	0.006		0.006		0.002	
Gardner (2021) Robust		X		X		X

Note:
* p < 0.1, ** p < 0.05, *** p < 0.01
Standard errors are clustered by district. Coefficient estimates are in seconds for Panel A and B, while estimates and means are in percent for Panel C. Columns 1 and 2 of Panel A show Call-to-Dispatch estimates when controlling for the implementation of Strategic Decision Support Centers (SDSC). In Columns 3 and 4, police districts 7 and 9 are omitted as Kapustin et al. (2022) shows that SDSCs affect police patrolling in these districts. Panel B is similar to Panel A, with the outcome of interest being Call-to-On-Scene times. In Panel C, the outcome of interest is a binary for whether a 911 call ended in an arrest. In Columns 5 and 6, we control for Body-Worn Camera (BWC) adoption. Note that in each specification, controls are consistent with the preferred specification. OLS estimates are reported in odd-numbered columns, while Gardner (2021) robust estimates are reported in even columns. The coefficient estimates of controls when using Gardner (2021) estimator are not reported as the two-stage method only returns the coefficient estimate of interest on the treated variable. In addition, the two-stage procedure may drop observations in the first stage if unable to predict values. This happens infrequently as shown in the observation counts, but is worth noting. Finally, wild cluster bootstrap p-values using 999 iterations are also reported as the number of clusters (22) is below the threshold of 30 put forth in Cameron et al. (2008). The bootstrap procedure cannot be performed using the Gardner (2021) estimator.

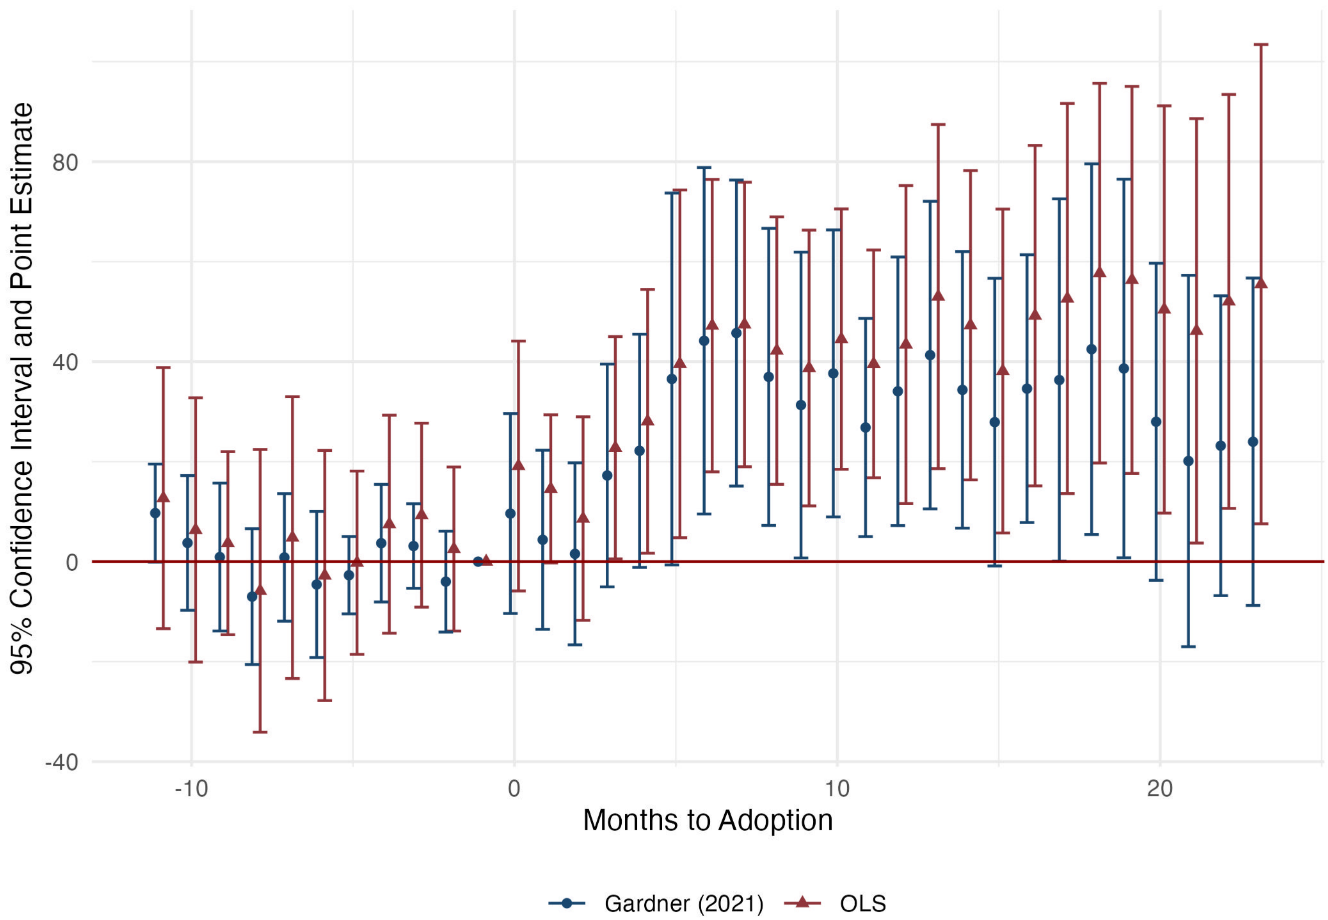

Figure C1: Event Study w/ SDSC Controls (Call-to-Dispatch)

Note: This figure shows the event study as specified in Equation 2 for Call-to-Dispatch times. Call-to-Dispatch is the amount of time from a 911 call to a police officer being dispatched to the crime scene. The x-axis denotes the number of months pre-/post-adoption of ShotSpotter technology. The y-axis denotes the 95% confidence intervals and point estimates (in seconds). The red error-bars/points represent confidence intervals/point estimates from OLS estimation, while the blue are from Gardner (2021) two-stage difference-in-difference estimators which are robust to heterogeneous treatment effects in staggered adoptions. All pre-/post-periods are normalized by the month before ShotSpotter adoption. Twelve periods are estimated, but only 11 pre-periods and 23 post-periods are reported as the -12 and +24 are binned endpoints. Controls match the preferred specification in addition to SDSC rollout. Standard errors are clustered at the district level.

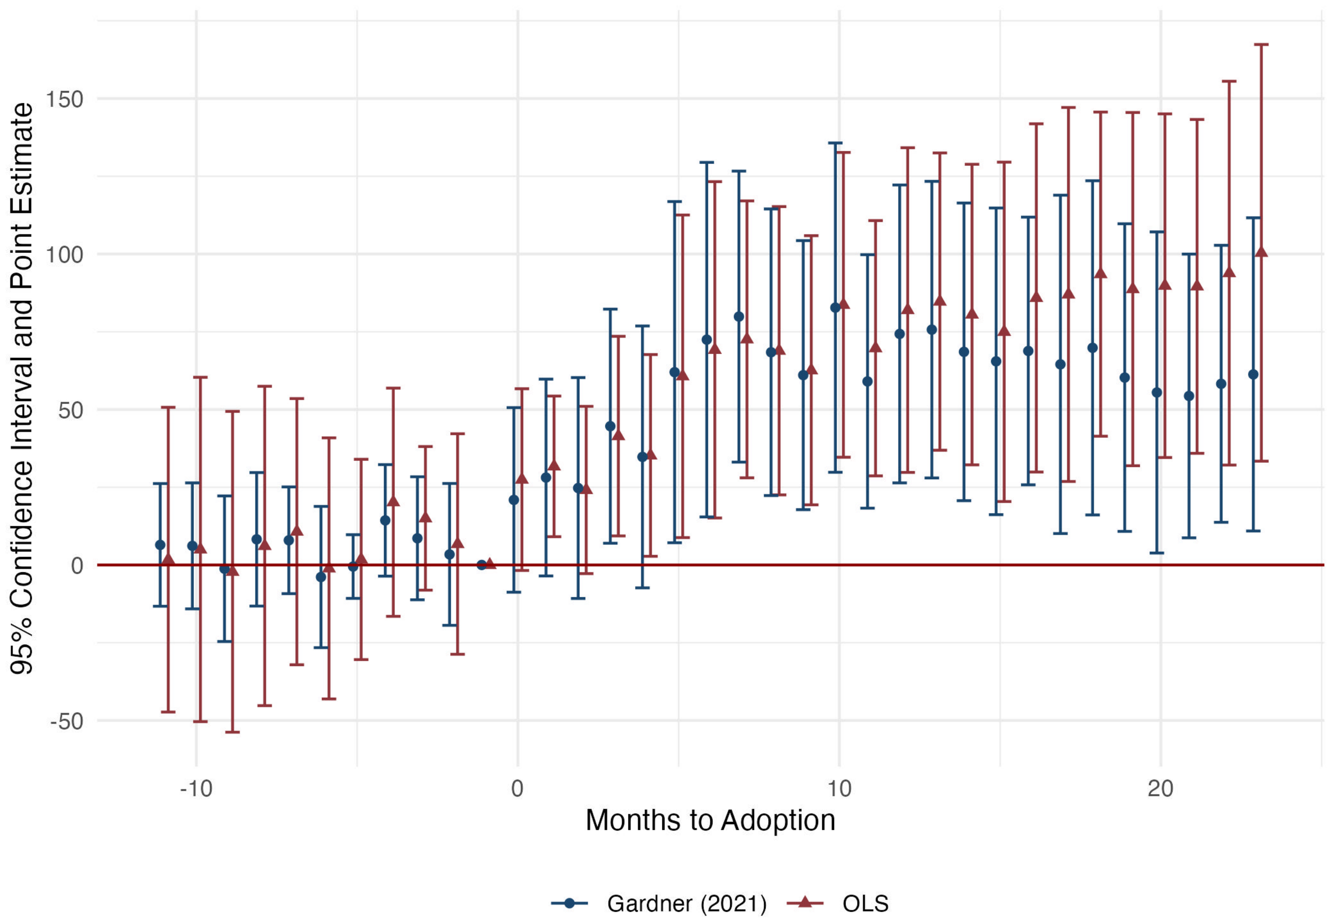

Figure C2: Event Study w/ SDSC Controls (Call-to-On-Scene)

Note: This figure shows the event study as specified in Equation 2 for Call-to-On-Scene times. Call-to-On-Scene is the amount of time from a 911 call to a police officer arriving to the crime scene. The x-axis denotes the number of months pre-/post-adoption of ShotSpotter technology. The y-axis denotes the 95% confidence intervals and point estimates (in seconds). The red error-bars/points represent confidence intervals/point estimates from OLS estimation, while the blue are from Gardner (2021) two-stage difference-in-difference estimators which are robust to heterogeneous treatment effects in staggered adoptions. All pre-/post-periods are normalized by the month before ShotSpotter adoption. Twelve periods are estimated, but only 11 pre-periods and 23 post-periods are reported as the -12 and +24 are binned endpoints. Controls match the preferred specification in addition to SDSC rollout. Standard errors are clustered at the district level.

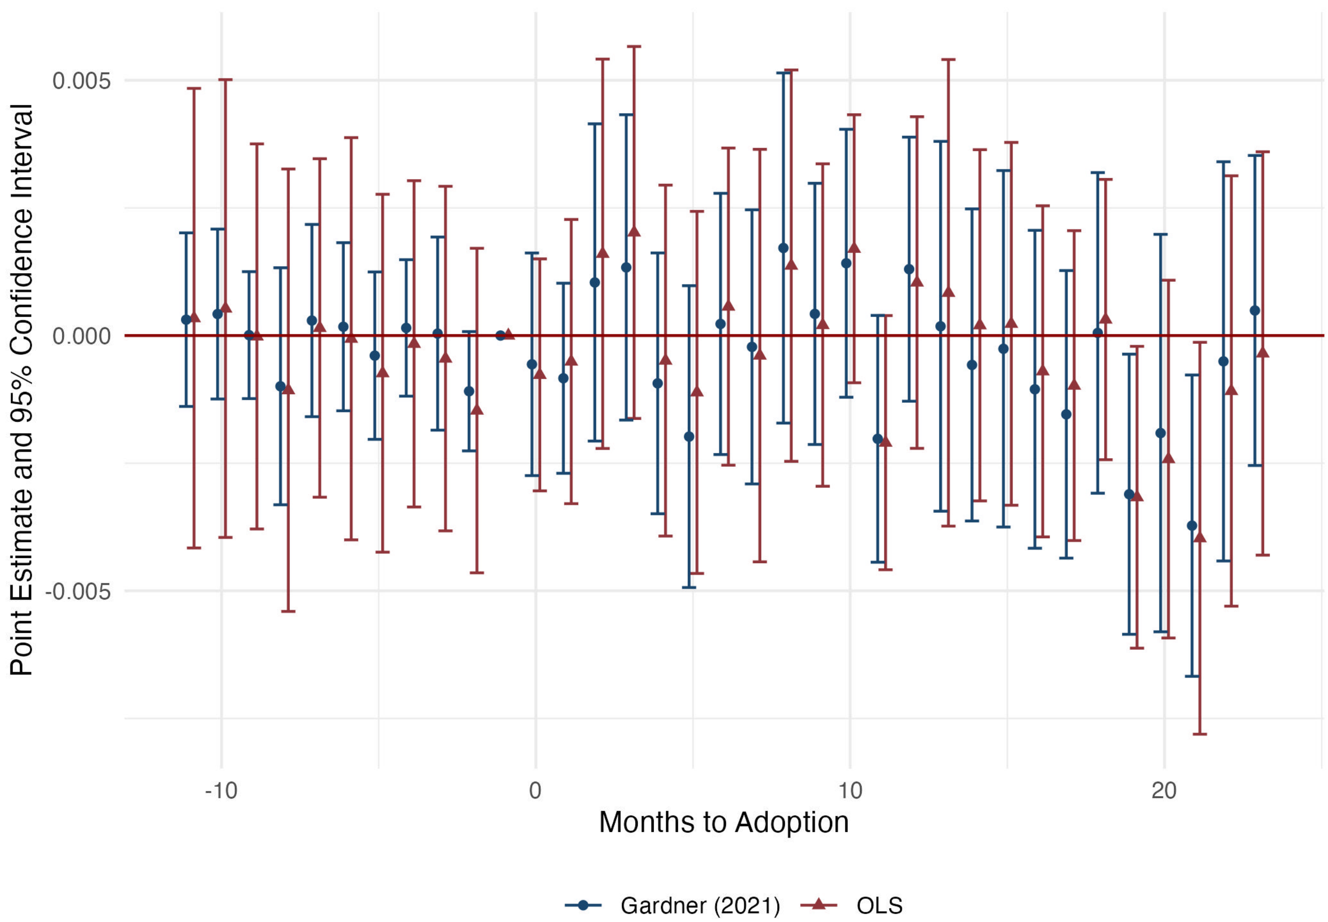

Figure C3: Event Study w/ SDSC Controls (Arrest Made)

Note: This figure shows the event study as specified in Equation 2 for whether an arrest is made. The outcome variable is a binary variable equal to one if an arrest is made during the 911 call, and zero otherwise. The x-axis denotes the number of months pre-/post-adoption of ShotSpotter technology. The y-axis denotes the 95% confidence intervals and point estimates (in percentage points). The red error-bars/points represent confidence intervals/point estimates from OLS estimation, while the blue are from Gardner (2021) two-stage difference-in-difference estimators which are robust to heterogeneous treatment effects in staggered adoptions. All pre-/post-periods are normalized by the month before ShotSpotter adoption. Twelve periods are estimated, but only 11 pre-periods and 23 post-periods are reported as the -12 and +24 are binned endpoints. Controls match the preferred specification in addition to SDSC rollout. Standard errors are clustered at the district level.

Appendix D Sensitivity Analysis of Event Studies

In this appendix, we conduct analysis following Rambachan and Roth (2023) on the OLS event study specifications in Figures 3 and 4 to illustrate the sensitivity of the estimates to possible violations of parallel trends. Specifically, we evaluate the degree of nonlinearity we can impose on a linear extrapolation of the pre-treatment trend. We adopt the notation used in Rambachan and Roth (2023) and define M as the maximum amount that the pre-treatment trend can change across consecutive periods. As an example, $M = 0$ implies no change in the post-treatment trends—the counterfactual difference in trends is exactly linear. Conversely, as M increases ($M > 0$), we allow for more nonlinearity in the pre-treatment trend and therefore greater uncertainty in the treatment effect estimates.

Since we are most interested in the average effect of ShotSpotter post-implementation, rather than one particular post-period, we perform the sensitivity analysis on the average of all post-implementation estimates obtained from Equation 2. Appendix Figures D1 and D2 report two important features: the confidence interval of the average of all post-period estimates (Original) and the corresponding robust fixed-length confidence intervals (FLCI) which show the average post-period effect under the assumption that the difference in pre-period trends can differ by up to M across consecutive periods. For both outcomes, the average of all post-implementation periods maintain their statistical significance under both a linear extrapolation of the pre-period ($M = 0$) and increasing amounts of non-linearity ($M > 0$) for both the Call-to-Dispatch and Call-to-On-Scene time.

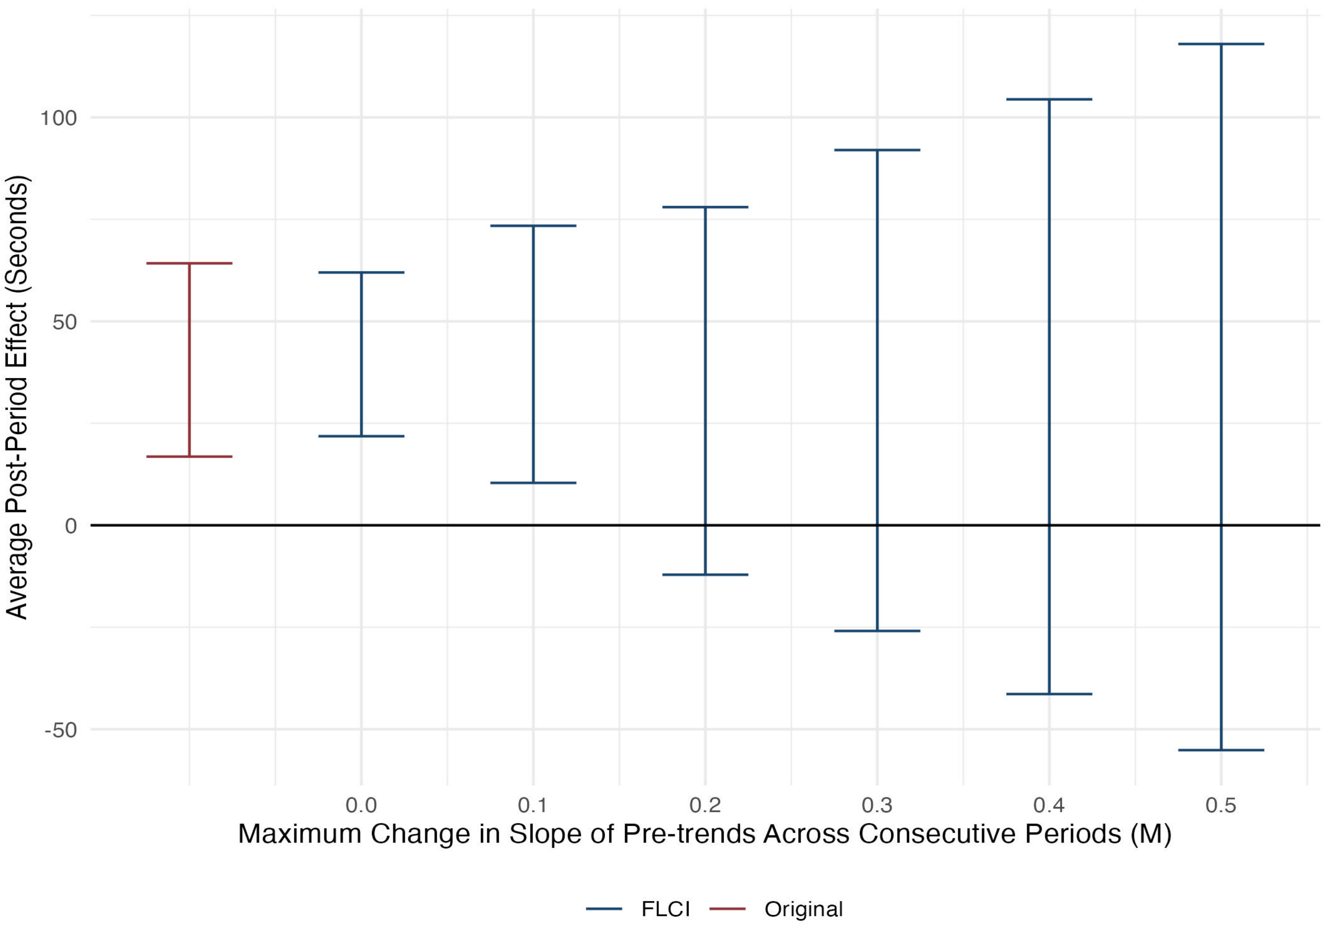

Figure D1: Sensitivity Analysis of Pre-Trends

Note: This figure shows sensitivity analysis of the event study plot in Figure 3. The x-axis shows the maximum change in slope of pre-trends across consecutive periods (M). We gradually increase M where M = 0 corresponds to allowing a linear trend and M > 0 allows for increasingly more varied nonlinear trends. In red, the average of the post-implementation periods are plotted. In blue, alternative Fixed-Length Confidence Intervals (FLCI), averaged over all post-implementation periods, that are proposed by Rambachan and Roth (2023) are plotted which relaxes the parallel trends assumption and requires only that differential trends evolve smoothly over time. Note that here, the breakdown value is 0.2 which means the significant effects observed in the post-implementation periods are only valid if we allow for the change in slope of the pre-period to change by no more than 0.2.

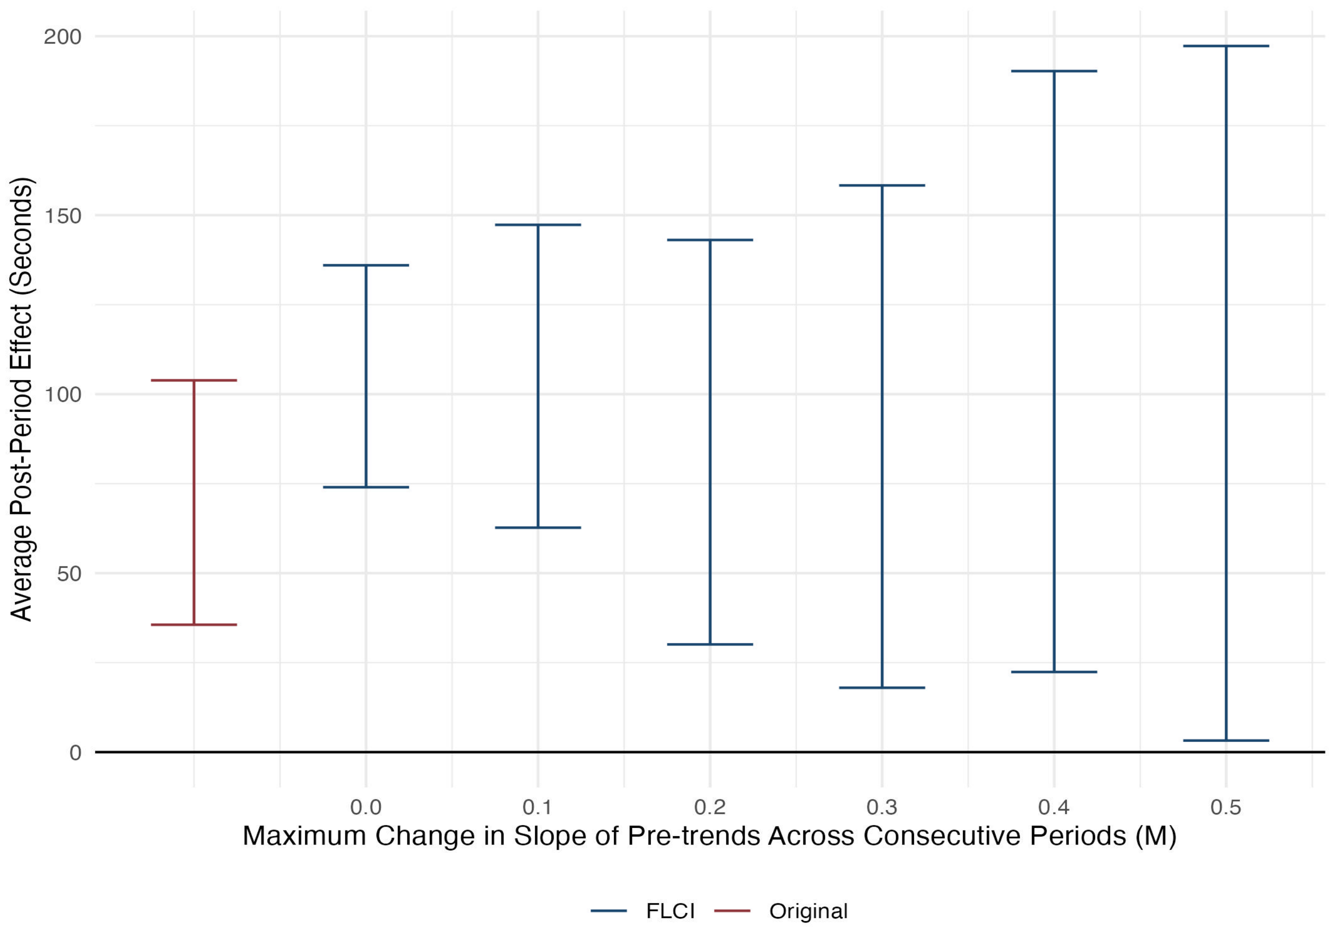

Figure D2: Sensitivity Analysis of Pre-Trends (Call-to-On-Scene)

Note: This figure shows sensitivity analysis of the event study plot in Figure 4. The x-axis shows the maximum change in slope of pre-trends across consecutive periods (M). We gradually increase M where M = 0 corresponds to allowing a linear trend and M > 0 allows for increasingly more varied nonlinear trends. In red, the average of the post-implementation periods are plotted. In blue, alternative Fixed-Length Confidence Intervals (FLCI), averaged over all post-implementation periods, that are proposed by Rambachan and Roth (2023) are plotted which relaxes the parallel trends assumption and requires only that differential trends evolve smoothly over time. Note that here, the breakdown value is larger than 0.5 which means the significant effects observed in the post-implementation periods are only valid if we allow for the change in slope of the pre-period to change by no more than a number larger than 0.5.

Appendix E Does ShotSpotter Affect Other Priority Calls?

Within this appendix, we pivot the analysis beyond response times for Priority 1 dispatches to the lower level priorities of Priority 2 (rapid dispatch) and Priority 3 (routine dispatch).[36] In doing so, we show implications that extend beyond Priority 1 dispatches, introducing trade-offs that dispatchers and officers face for lower-level reports of crime. Specifically, we find a 'trickle-down' effect, wherein time-sensitive lower-priority calls (Priority 2) are also impacted by ShotSpotter implementation. Interestingly, we find suggestive evidence that time-insensitive dispatches (Priority 3) may also be affected, further implying a potential strain on officers' responsibilities when ShotSpotter is implemented. Moreover, we separately analyze the five most frequent types of calls within each priority. Surprisingly, this analysis leads to suggestive health implications where ShotSpotter may be unintentionally costly for victims in need of medical services. In effect, we build upon the rapid-response literature related to health outcomes, as longer travel times and ambulance response times have been linked to higher mortality rates (Avdic, 2016; Wilde, 2013; Leslie and Wilson, 2020; DeAngelo et al., 2023).

First, Equation 1 is estimated by priority on Call-to-Dispatch and Call-to-On-Scene times in Figures E1 and E2, respectively. In each figure, the point estimates and confidence intervals are divided by the mean of the dependent variable to show percentage changes. As an example, the top rows of each corresponding priority, labeled "Pooled Estimate," represent the 95% confidence intervals for the percentage change from the mean. Moreover, within each priority, the five most frequent call types are uniquely estimated and plotted in descending order of their mean response time. For instance, in the Priority 1 panel of Figure E1, the call description Battery in Progress has the lowest average

[36] A Priority 2 dispatch is defined as a response in which timely police action which has the potential to affect the outcome of an incident. A Priority 3 dispatch is defined as a response to a call for service that does not involve an imminent threat to life, bodily injury, or major property damage/loss, and a reasonable delay in police action will not affect the outcome of the incident.

Call-to-Dispatch time, while Suspicious Person and Check Well Being have the second and third lowest. Using this ranking, we find that the Priority 1 call-types that have the fastest response times exhibit the largest effects for both outcomes after ShotSpotter implementation.

As shown in the first row labeled Pooled Estimate of both Figure E1 and Figure E2, Priority 2 response times for both outcomes show significant increases. Priority 2 calls are categorized as incidents that are non-life-threatening, but where police intervention may affect the outcome of the event. This significant increase in Priority 2 response times suggests a 'trickle down' effect from delays in Priority 1 dispatches. Intuitively, an officer that is delayed for a higher priority call, may also be delayed for less important tasks. However, for Priority 3 calls, which are time insensitive, we find only suggestive evidence of increased response times as Call-to-Dispatch is not statistically significant and Call-to-On-Scene is significant at the 10% level. Despite this, the point estimates for Priority 3 calls are positive, and the insignificant estimates may be a result of the large average response times for Priority 3 call types. As shown in the first row of Figures E1 and E2, the average response times for Priority 3 Call-to-Dispatch and Call-to-On-Scene are 16 minutes and 31 minutes, respectively. Given that these averages are substantially larger than Priority 1 and Priority 2, the estimated change in average time may not be large enough to detect. Despite this limitation, the positive coefficient estimates support the notion that officers' responsibilities are strained in the presence of ShotSpotter, creating further delays in responding to time-insensitive calls.

Second, as mentioned, Equation 1 is estimated for each of the five most frequent call types by priority. The results of these estimations are also plotted in Figures E1 and E2 below the Pooled Estimate. For Priority 1 and Priority 2 calls, we find consistent evidence of increased delays for both response times for nearly all call-types, thus showing that the effects are wide-spread across different emergency situations. Of notable importance, Figure E1 reports longer Call-to-On-Scene times for Emergency Medical Services (EMS), which may have significant health implications. In particular, the point estimate reports a 69-second increase in the response time for EMS calls. According to the Chicago EMS System Policies and Procedures, treatment and transport of injured civilians should be delayed pending police arrival if the safety of the EMS personnel could be jeopardized. Therefore, this observed delay in police response may postpone critical medical services. Specifically, Wilde (2013) find that a minute increase in response times increases mortality between 8-17%. Given the additional minute increase we find in Call-to-On-Scene times, ShotSpotter may have significant social costs beyond a lower likelihood of arresting perpetrators, and may hinder injured civilians from receiving timely care.

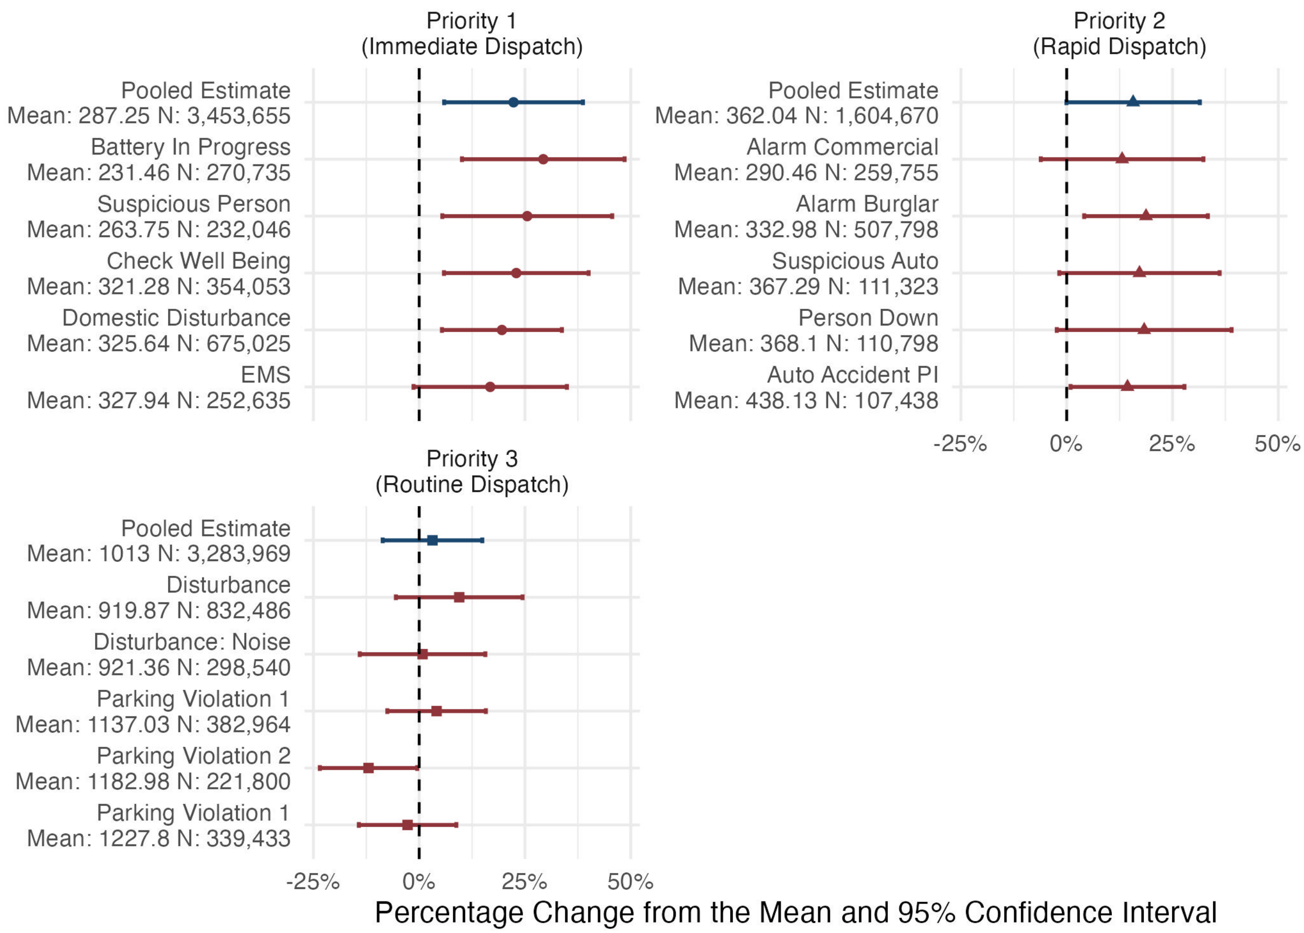

Figure E1: Effect of ShotSpotter by Priority (Call-to-Dispatch)

Note: This figure plots the effects of ShotSpotter on Call-to-Dispatch times by priority and by most frequent call-type. In the first row of each panel, the pooled estimate combining all respective call types is reported. The subsequent rows report estimates for the most frequent call-types, ranked by their average Call-to-Dispatch time. For instance, in Priority 1, Battery in Progress has the lowest average Call-to-Dispatch time, while Suspicious Person has the second lowest. The x-axis shows the percent change from the mean (i.e., the point estimate divided by the mean of the outcome), as well as the corresponding 95% confidence interval using the specification from Equation 1. The number of observations and means are shown in the y-axis for each call-type. All estimations are estimated using OLS and the preferred specification.

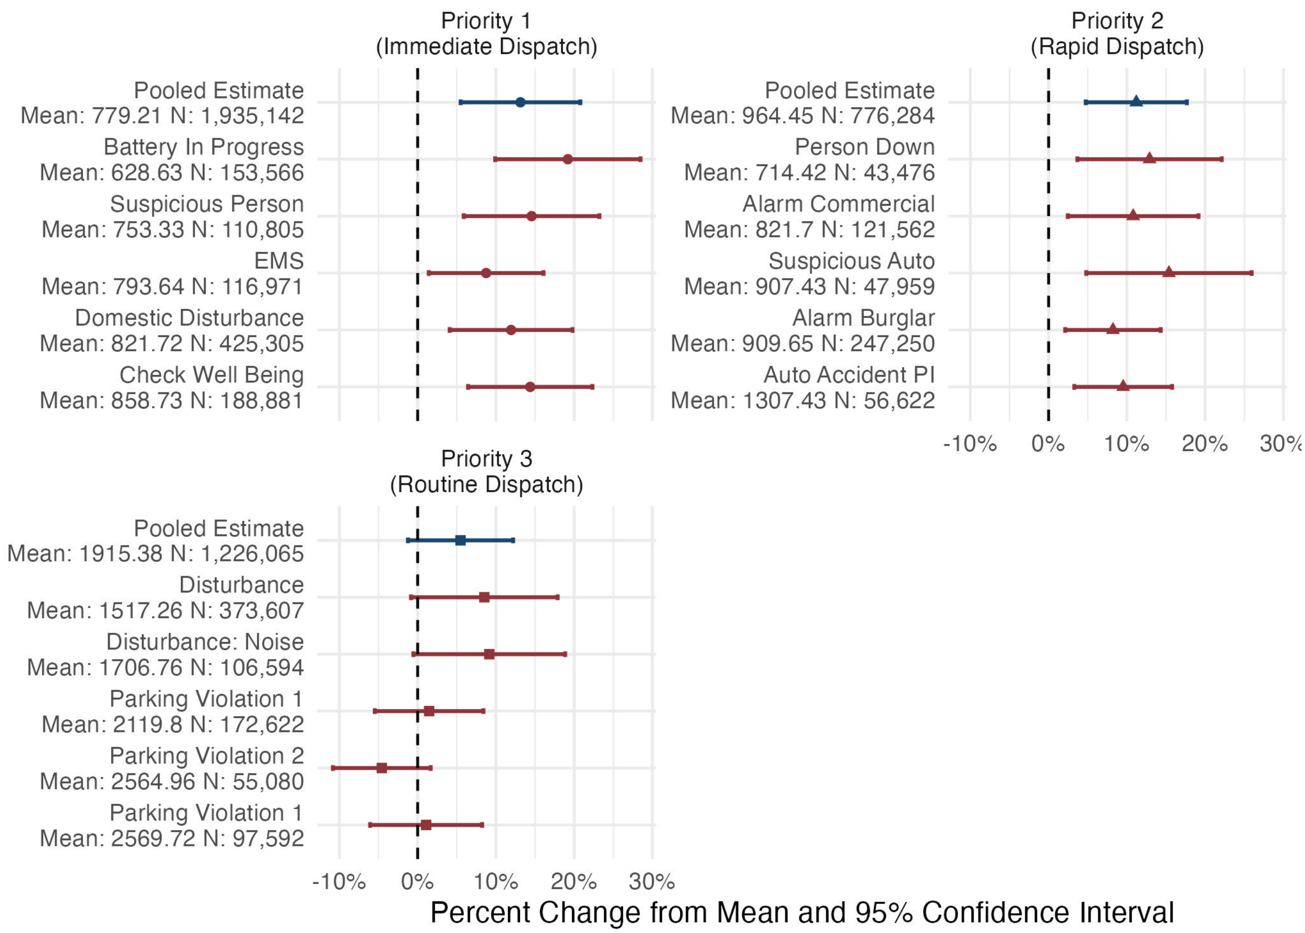

Figure E2: Effect of ShotSpotter by Priority (Call-to-On-Scene)

Note: This figure plots the effects of ShotSpotter on Call-to-On-Scene times by priority. In the first row of each panel, the pooled estimate combining all respective call types is reported. The subsequent rows report estimates for the most frequent call-types, ranked by their average Call-to-On-Scene time. For instance, in Priority 1, Battery in Progress has the lowest average Call-to-On-Scene time, while Suspicious Person has the second lowest. The x-axis shows the percent change from the mean (i.e., the point estimate divided by the mean of the outcome), as well as the corresponding 95% confidence interval using the specification from Equation 1. The number of observations and means are shown in the y-axis for each call-type. All estimations are estimated using OLS and the preferred specification.

Appendix F Supplemental Figures and Tables

Table F1: Effect of ShotSpotter Implementation on Confounding Controls (OLS)

	(1)	(2)
Panel A: Number 911 Dispatches		
ShotSpotter Activated	-1.885	-1.889
	(2.194)	(2.516)
Mean of Dependent Variable	149.550	149.550
Observations	55,792	55,792
Panel B: Officer Availability		
ShotSpotter Activated	-23.949	-42.806*
	(22.709)	(25.534)
Mean of Dependent Variable	1,277.860	1,277.860
Observations	55,792	55,792
FE: Day-by-Month-by-Year	X	X
FE: District	X	X
Gardner (2021) Robust		X

Note:
* p < 0.1, ** p < 0.05, *** p < 0.01
Standard errors are clustered by district. Coefficient estimates are reported in seconds. This table shows estimations on two outcome variables, Number of 911 Dispatches and Officer Availability, which are not included in the main specification due to the possibility of being confounding controls. Each panel refers to a distinct outcome variable. Since each outcome variable is at the district-day level, we aggregate the call-level data to the district-day. Hence, in these models, we cannot control for call-type nor hour of the day. Number 911 Dispatches is the number of 911 dispatches. Officer Availability is the number of police officer hours within a district. ShotSpotter Activated refers to the timing in which each district receives ShotSpotter technology. The Gardner (2021) estimator is robust to the heterogeneous treatment effects in staggered two-way-fixed-effects designs. January 1, July 4, and December 31 are omitted due to their correspondance with potential celebratory gunfire.

Table F2: Effect of ShotSpotter on 911 Call Resolutions (Logit)

		Most Frequent Final 911 Dispositions		
	Arrest Made	Other Police Service	No Person Found	Peace Restored
	(1)	(2)	(3)	(4)
ShotSpotter Activated	-0.086***	0.046*	0.048**	-0.101***
	(0.023)	(0.026)	(0.021)	(0.033)
Mean of Dependent Variable	0.026	0.410	0.214	0.063
Observations	3,446,996	3,453,582	3,453,235	3,451,182
FE: Day-by-Month-by-Year	X	X	X	X
FE: District	X	X	X	X
FE: Call-Type	X	X	X	X
FE: Hour-of-Day	X	X	X	X

Note:
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Standard errors are clustered by district. The dependent variable in Column 1 is an indicator equal to one if a 911 call resulted in an arrest. The dependent variable in Columns 2-4 is an indicator equal to one if a 911 call resulted in Other Police Service (Column 2), No Person Found (Column 3), or Peace Restored (Column 4). Columns 2-4 report the three most frequent 911 final dispositions: Other Police Service, No Person Found, and Peace Restored. The final disposition is the final result of what happened on the 911 call.

Table F3: Effect of ShotSpotter on 911 Call Resolutions (Gardner 2021)

		Most Frequent Final 911 Dispositions		
	Arrest Made	Other Police Service	No Person Found	Peace Restored
	(1)	(2)	(3)	(4)
ShotSpotter Activated	-0.271***	1.342**	1.101***	-0.479*
	(0.061)	(0.672)	(0.398)	(0.274)
Observations	3,453,623	3,453,623	3,453,623	3,453,623
Mean of Dependent Variable	2.553	41.047	21.367	6.283
FE: Day-by-Month-by-Year	X	X	X	X
FE: District	X	X	X	X
FE: Call-Type	X	X	X	X
FE: Hour-of-Day	X	X	X	X
Gardner (2021)	X	X	X	X

Note:

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Standard errors are clustered by district. All coefficient estimates and means are in percentages. All estimates are computed using the Gardner (2021) estimator. The dependent variable in Columns 1 is an indicator equal to one if a 911 call resulted in an arrest. The dependent variable in Columns 2-4 is an indicator equal to one if a 911 call resulted in Other Police Service (Column 4), No Person Found (Column 5), or Peace Restored (Column 6). Columns 2-4 report the three most frequent 911 final dispositions: Other Police Service, No Person Found, and Peace Restored. The final disposition is the final result of what happened on the 911 call.

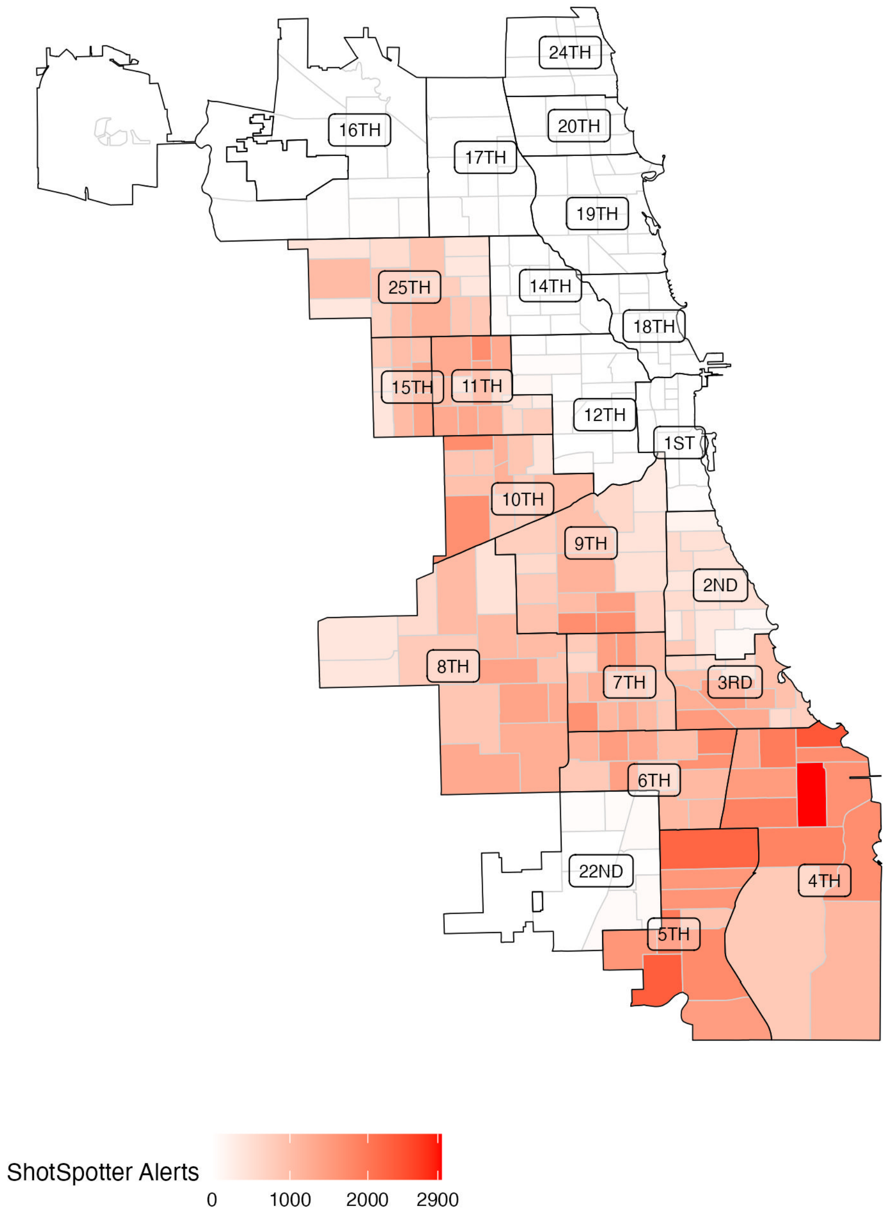

Figure F1: Map of ShotSpotter Districts in Chicago

Note: There are 22 police districts in Chicago, and 12 are equipped with ShotSpotter technology. Each district contains beats which are designated by the boxes within the district lines. ShotSpotter implementation began in January 2017 and ended in May 2018.

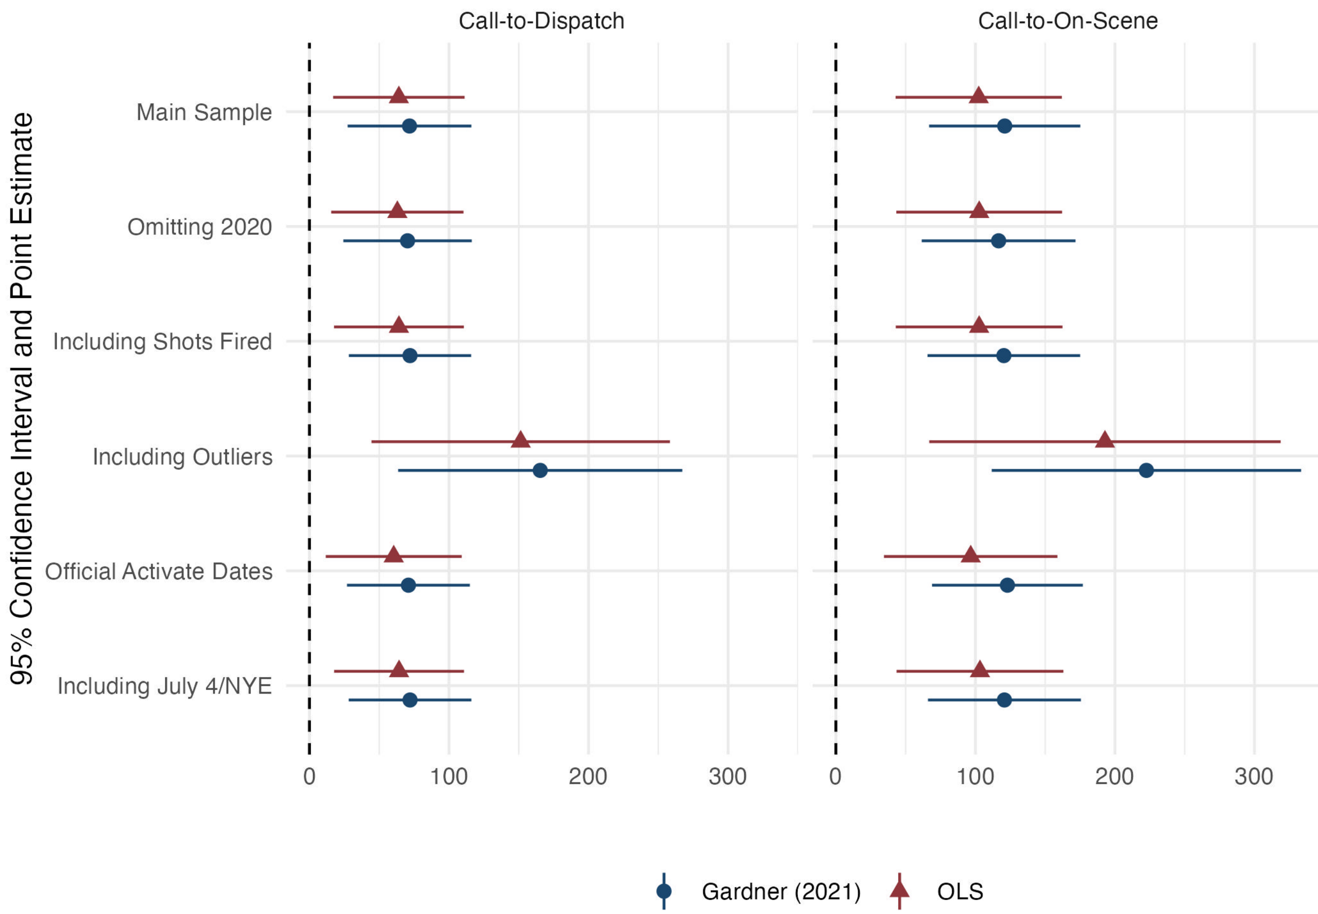

Figure F2: Robustness of Main Results

Note: This figure shows the results from estimation of Equation 1 with five different samples for both Call-to-Dispatch and Call-to-On-Scene. These restrictions are outline in Section 3.2. Main Sample refers to the main sample used in the paper. Omitting 2020 omits the year 2020 due to Covid-19. Including Shots Fired includes any 911 call dispatches related to the description of 'Shots Fired'. Including Outliers includes all outliers that are removed from the main analysis (+3 standard deviations from the mean). Official Activate Dates uses the official ShotSpotter activation dates as received from a Freedom of Information Request from the Chicago Police Department. These dates are similar, but not exact, to the dates we use due to what we observe in the data. Last, we include July 4th, New Year's Eve, and New Year's Day, which are excluded from the preferred sample since there may be many false-positive reports of gunfire.

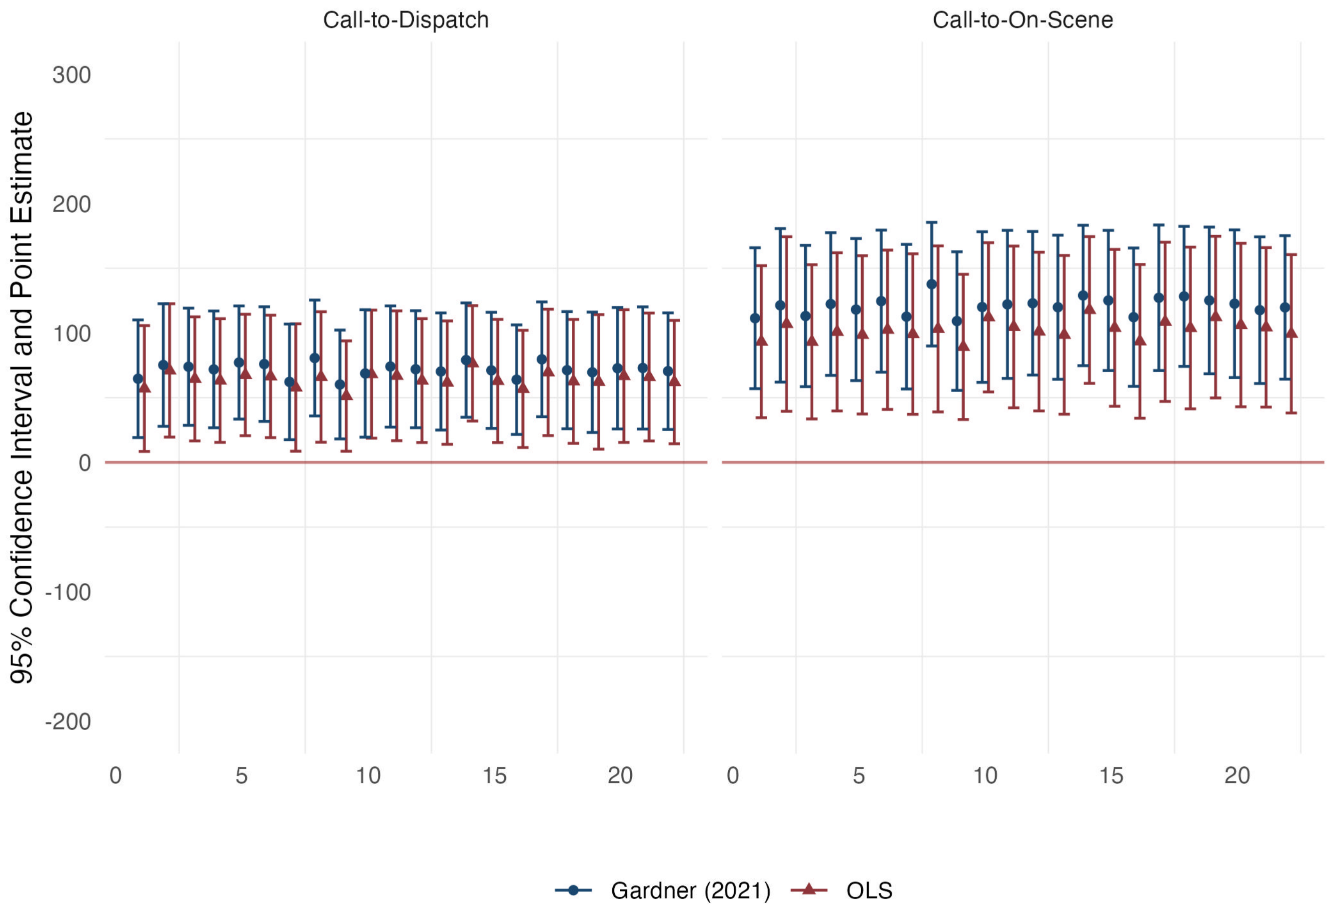

Figure F3: Leave-One-Out Analysis

Note: This figure shows the results from 22 distinct OLS and Gardner (2021) regressions using Equation 1. Both outcomes of Call-to-Dispatch and Call-to-On-Scene are pictured. In each iteration, one police district is removed from estimation to ensure that the effects of ShotSpotter are not driven by one district. The blue points and error-bars represent Gardner (2021) point estimates and 95% confidence intervals, which are robust to heterogeneous treatment effects in staggered designs. The red points and lines denote point estimates and 95% confidence intervals from OLS estimates. Standard errors are clustered at the district level.

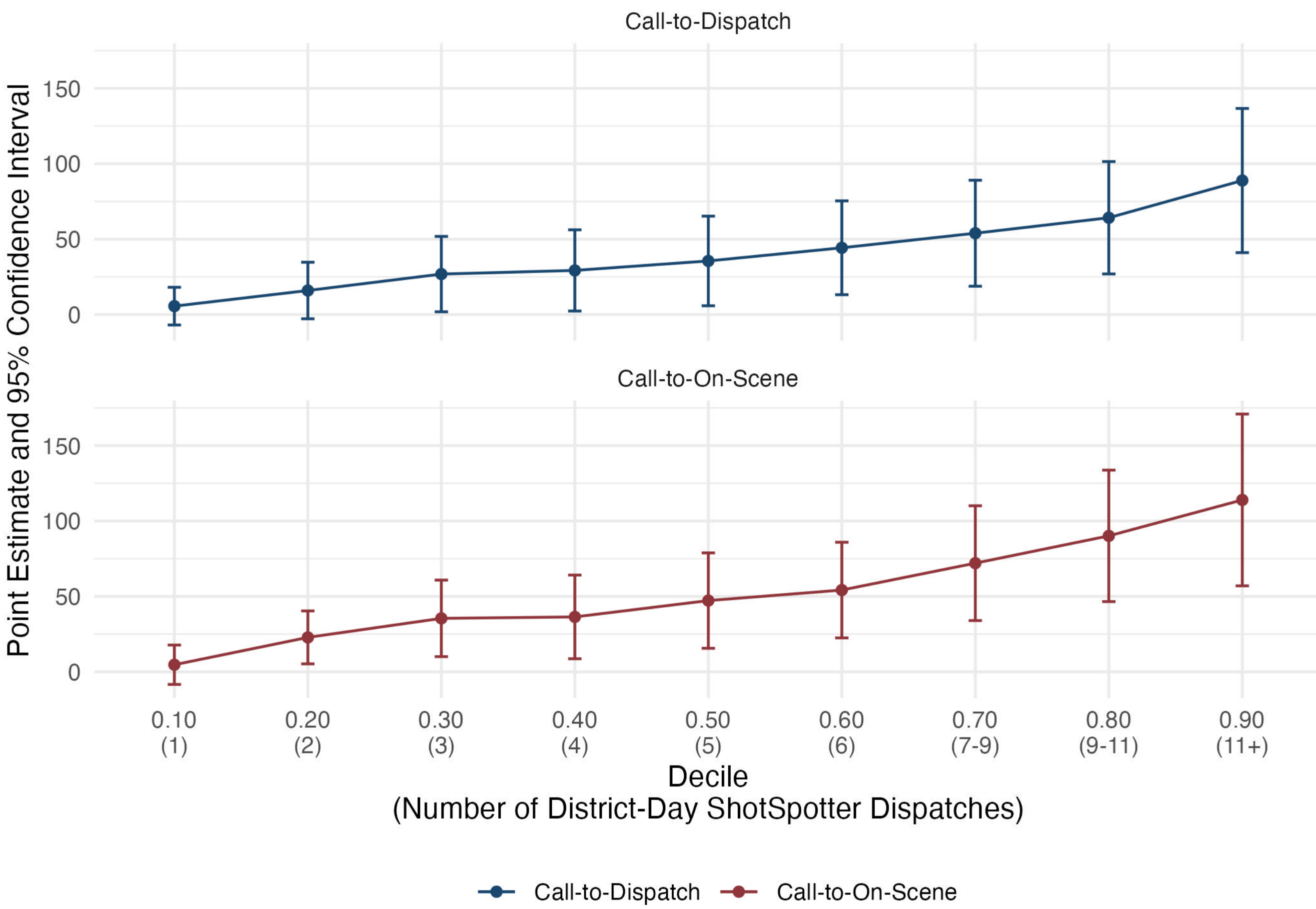

Figure F4: Marginal Effect of ShotSpotter Dispatches on Response Times (OLS)

Note: This figure shows the marginal effect of ShotSpotter dispatches using the model outlined in Section 5.2. However, the number of ShotSpotter dispatches is split into deciles to show the linear relationship between number of ShotSpotter dispatches and response times. In this figure, 9 deciles are plotted, with the reference decile being when the number of ShotSpotter dispatches is zero. All coefficient estimates are in seconds. Deciles are on the x-axis, and the number of ShotSpotter dispatches corresponding to each decile is in parentheses.

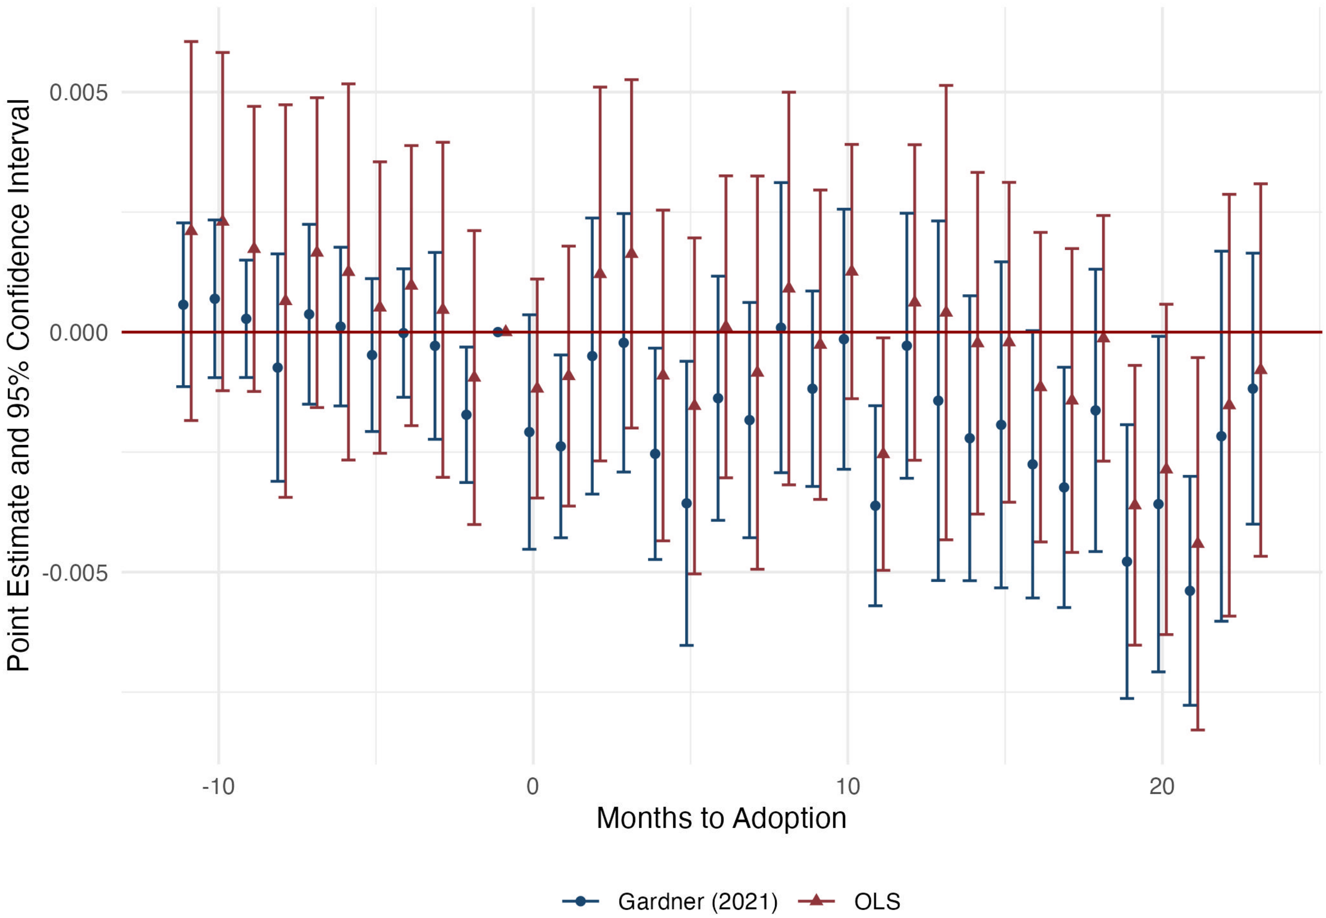

Figure F5: Event Study (Arrests)

Note: This figure shows the event study as specified in Equation 2 for the probability of a 911 call ending in an arrest. The outcome variable is a binary variable equal to one if a 911 call resulted in an arrest being made. The x-axis denotes the number of months pre-/post-adoption of ShotSpotter technology. The y-axis denotes the 95% confidence intervals and point estimates. The red error-bars/points represent confidence intervals/point estimates from OLS estimation while the blue are using the Gardner (2021) two-stage difference-in-difference estimator, which is robust to heterogeneous treatment effects in staggered adoptions. All pre-/ post-periods are normalized by the month before ShotSpotter adoption. Twelve pre-periods (24 post-periods) are estimated, but only 11 pre-periods (23 post-periods) are reported, as the -12 (+24) is a binned endpoint. Controls match the preferred specification. Standard errors are clustered at the district level.

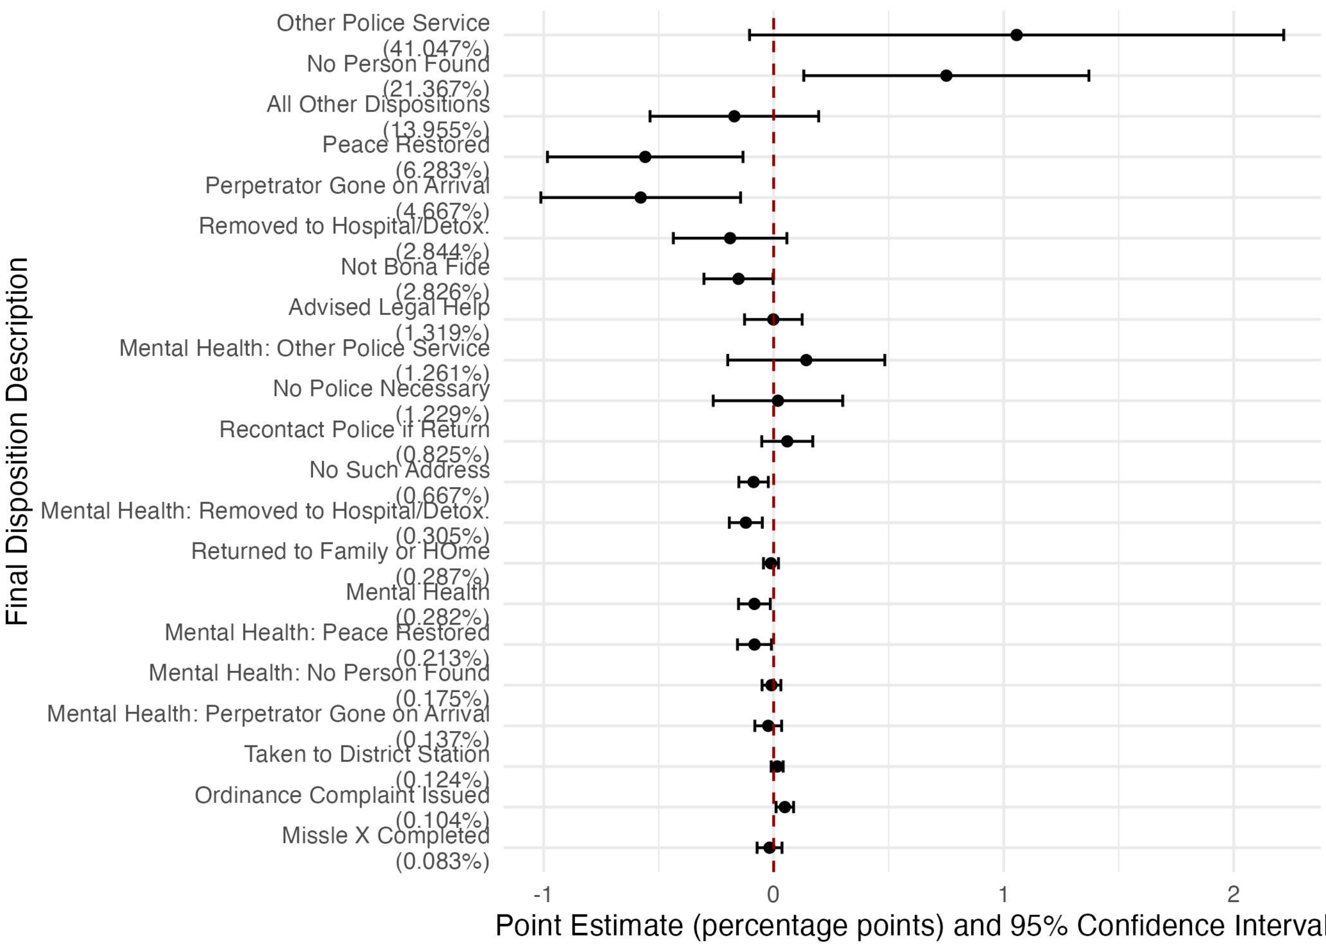

Figure F6: Effect of ShotSpotter on Final Disposition Codes (OLS)

Note: This figure plots the results of Equation 2 for each Final Disposition Code. Final Disposition Codes are the codes that an officer uses to describe the action they took on a 911 call. Each outcome variable is a binary variable equal to 1 for each corresponding code, multiplied by 100 to give the coefficient estimates a percentage point interpretation. Numbers in parentheses correspond to the fraction of 911 calls that have the disposition code. The code All Other Dispositions is the aggregation of all final disposition codes that were not in the twenty most frequent. Controls match the preferred specification. Standard errors are clustered at the district level.