Crime Prevention Through Intelligence and Information Sharing: An Evaluation of an Information Intervention at the Philadelphia Police Department

Research Team

Aaron Chalfin
Associate Professor
Department of Criminology

Greg Ridgeway
Professor
Department of Criminology
Department of Statistics and Data Science

John MacDonald
Professor
Department of Criminology
Department of Sociology

Rachel Ryley
Lead Data Scientist
Department of Criminology

July 2022

This evaluation research was made possible by a grant from the Neubauer Family Foundation. The opinions expressed in this report are those of the authors and do not necessarily reflect the views of the Neubauer Family Foundation. The research was conducted with collaboration from the Philadelphia Police Department. The opinions expressed in this report are those of the authors and do not necessarily reflect the views of the Philadelphia Police Department.

Findings in brief

The Philadelphia Police Department began distributing 435 mobile smartphones to officers in police districts 22, 24, and 25 in February 2021. At the same time PPD established Crime Information Centers (CICs) to facilitate analysis and information sharing. We compared changes in police-related outcomes in districts 22, 24, and 25 with six districts (12, 14, 15, 19, 35, 39) that received no phones and had similar levels of serious crime.

The smartphones provided officers with improved access to information and a convenient technology to receive requests for intelligence crucial to investigations, report street-level intelligence, and communicate directly with members of the community.

Mobile phones/CICs have public safety benefits

- An increase in the violent crime clearance rate from 24% to 30%
- An increase in the likelihood that a stop resulted in an arrest — from 10% to 28% — suggesting more surgical policing, without increasing the number of stops conducted

Mobile phones substantially ease regular PPD officers' tasks

- Greatly increased the amount and variety of evidence collected
  - Weekly uploads increased 40% after at least one SIG detective received a phone
- Facilitated 311 reports to address physical disorder in districts
- Made officers more willing and able to create intelligence reports
- Enhanced basic communication between police and community members through calls and text messages, including direct contact about the location of illegal firearms
- Can improve the completeness and timeliness of NCIC/PCIC checks, patrol logs, and court notices

PPD has more to gain from mobile IT and CICs

- More incentives are needed to promote smartphone use among officers. A few officers in each of the pilot districts were more active users of the smartphones
  - 5/7 squads use the phones a lot, others were infrequent users

  - 3 officers submitted half of the 311 requests
  - 86% of officers submitted no intelligence reports at all
  - Usage has essentially ended in District 22

- Regular use of phones among officers could support mission-directed patrol
  - Monitoring the time spent in mission areas
  - Documenting mission-related business checks and home visits
  - Promote intelligence reports in mission areas

- PPD could encourage additional phone usage
  - Encourage officers to share information and give feedback on how their intel reports and 311 reports are solving community problems
  - Emphasize phone usage in CompStat by tracking key metrics
    * Time spent in strategic areas
    * Number of leads connected to shootings or priority incidents
    * Volume of direct calls/messages from community leading to crime clearances
  - Establish a clear policy on expected phone use, monitor use, and provide feedback to command staff and officers on the successful uses of technology

Executive Summary

Can information technology improve policing and promote public safety? The shift to more data-driven policing under the CompStat model in the 1990s suggests that the answer to this question is a resounding "yes." Yet relative to private firms, police departments continue to invest relatively few resources in information technology innovations. In a number of cities, the majority of officers continue to handle many of their daily tasks using pencil and paper in much the same way that things were done a generation ago. Until 2021, police officers in the Philadelphia Police Department (PPD) did not have department-issued mobile phones, limiting their ability to send and receive law-enforcement sensitive information and discouraging community contact through calls and messaging.

Without mobile phones, there are constraints on the flow of information within PPD. Information flows more slowly from the agency's analytical teams to street-level officers, including analysis of recent crime patterns and requests for specific intelligence useful for solving a crime or building a criminal case. Similarly, street-level officers lack a convenient way to pass information on to analysts. This is an especially challenging problem for officers who patrol on foot or bicycle. Without mobile devices, PPD analysts do not have geolocation of officers on foot or bicycle patrol and these officers do not have access to information rich computers that are available in every squad car. Since these officers are tasked with responding to fewer service calls and have more time to engage with the community while out on routine patrols, it is precisely these officers who can add the most value when equipped with better information technology.

In order to improve the flow of information within PPD, beginning in February 2021, PPD began distributing mobile phones to officers in the 22nd, 24th, and 25th districts. The initiative was not intended to be merely an IT upgrade, simply making looking up records more efficient. Instead, in concert with the establishment of CICs, the mobile devices were intended to be a part of an information intervention that would result in getting new questions answered, gathering better intelligence, pushing relevant information, and ultimately improving decision-making for everyone from the front-line officer to the crime analyst to the police management and leadership. By July 2022, PPD would acquire 970 phones with 790 of them distributed to officers assigned to District 22, District 24, District 25, Homicide Division, Shooting Investigation Group (SIG), Criminal Intelligence Unit, Central Detectives, East Detectives, Taskforce, and various other PPD positions (e.g. Intel Bureau Social Media). Our analysis focuses on the 435 mobile phones deployed in Districts 22, 24, and 25.

Across many different areas of public life, research has documented that there is wisdom in crowds, including wisdom that is not easily tapped into by experts working on their own.

While experts bring their own brand of knowledge to problem-solving, information is typically costly to acquire and is often distributed diffusely among many different people. In trying to solve serious crimes, criminal investigators are often at a loss and lack the information necessary to make an arrest. Without a witness, camera footage, or a trusted informant who is in the know, it is very unlikely that a crime will be solved. But the information to clear many of these cases is probably available. Community members may hold some of that information but may be reluctant to step forward. In other cases, PPD personnel may hold the key to solving a serious crime, but they simply do not know that they are holding information that would be vital to an investigation. The problem is that those in possession of useful information may not always recognize its utility. Compounding these problems, few formal mechanisms are available to collect that information and deliver it to criminal investigators.

In addition to communication within PPD, department-issued mobile phones provide a mechanism for street-level officers to communicate more frequently and more organically with members of the community. Many people are uncomfortable to call 911 to report criminal activity or do not wish to call attention to themselves. Citizens also sometimes do not have a sense for how PPD will respond, who will respond or when a response, if any, will happen. When officers meet with community members and share their direct PPD mobile number, people can communicate directly with an officer who they know and have some relationship. Community members have contacted PPD officers on their new department-issued mobile devices about specific crimes, including locations of guns, and arrests have been made on this basis. While such stories are just examples, they underscore the tremendous promise that information technology has to improve policing and make it more responsive to the community once the technology is widely adopted.

While department-issued mobile phones hold great promise for more efficient and effective policing, research on this topic is scant. Many interventions work well in theory but run into unanticipated roadblocks upon implementation. Leveraging the staggered roll-out of department-issued mobile phones to PPD officers in three police districts — the 22nd, 24th, and 25th — beginning in February 2021, we study the extent to which officers have used the phones and ultimately the degree to which the phones have led to improvements in public safety and police practice.

The evidence on the effectiveness of the phones is mixed. On the one hand, it is clear that a year after the phones were deployed, relatively few officers are using their phones regularly. While some officers have taken advantage of the phones to interact with members of the community, perform business checks, make 311 reports to address physical disorder, and pass intelligence reports on to analysts at the city's Crime Information Centers (CICs),

these activities are driven by a small number of enthusiastic adopters rather than broad interest among a large group of patrol officers. Use of the phones has, if anything, declined over time suggesting that use of the phones must continue to be emphasized by supervisors and leadership within PPD.

Overall, we do not see evidence that the provision of department-issued mobile phones has led to crime reduction in the three pilot districts relative to other districts serving similar communities. However, the evidence suggests that the provision of phones, alongside the CICs, have made policing more effective in clearing violent crimes and initiating a small number of arrests during vehicle stops in the three Philadelphia districts in which the intervention was deployed.

Despite PPD's difficulties in facilitating take-up among pilot district officers, in the year since the phones have been adopted there are some noteworthy successes which suggest that the phones are operating well below their potential to improve the efficiency and the effectiveness of policing in Philadelphia. First, as more officers within a district have been issued a phone, violent crime clearance rates in the pilot districts have risen relative to other districts in which no phones were available. Importantly, improvements in clearance rates are driven by investigators linked to the pilot districts, indicating that information sharing from officers and the CICs have been critical in facilitating this success. Second, the rate at which stops, particularly vehicle stops, have been successful — that is, yielded an arrest — has increased in these districts. This increase came without an increase in the number of stops. Overall, this finding suggests that the phones have allowed officers in the field to tap into more up-to-date and tactical information in order to police more surgically and more effectively.

In the future, the degree to which the phones will add value depends on whether their use is emphasized by key stakeholders within PPD and if PPD is able to augment the functionality of the phones with additional features that are useful to officers in the course of carrying out their daily duties. We regard these results as proof-of-concept of the promise of investments in information technology, understanding that PPD might either build upon these successes or fail to replicate them, depending upon the future path that is taken.

1 Introduction

Fueled by the "gales of creative destruction," and the incredible dynamism of market capitalism, information technology has transformed the global economy over the last fifty years. Nearly 9 in 10 Fortune 500 companies in 1955 are no longer around today. Among the five most valuable companies in the world in 2020 — Microsoft, Apple, Amazon, Google and Facebook — none were in existence as recently as 1975. The previous generation's most valuable firms — in particular GE and ExxonMobil — provided fuel, literally and figuratively, for an economy that was driven by the manufacture and transport of consumer goods. Notably, each of today's most valuable companies has risen to global prominence by fueling the information economy, revolutionizing the way that information is collected, disseminated, stored, and analyzed by businesses and individuals alike.

Policing is notoriously slow to change and has not adopted information technology to the same degree and vigor as society at large. Police departments, like other government agencies, do not collect profits and do not have shareholders and so are relatively insulated from the competitive pressures faced by private firms. Police officers, many of who are members of big city police unions, have bargained successfully for institutional arrangements that prioritize job security over opportunities for advancement and adoption of new technology. American founding principles also lean toward limiting police collection of information and restricting surveillance capabilities, essentially sacrificing police efficiency in order to limit government intrusion. While technologies like DNA databases, body-worn cameras, and predictive policing software have changed the way that police do their jobs, none of these changes has transformed the way in which police services are delivered. The model of the cop on the beat, patrolling in a police car, responding to service calls, and filling out paperwork remains just as salient today as it did in 1980. Indeed, in many large police departments, forms continue to be filled out using paper and pens and many police officers continue to lack something as basic as a department-issued mobile phone.

While information technology has not transformed police departments to the same degree as it has transformed the work of most private firms, it has nevertheless been adopted successfully to some degree by a number of large municipal police departments. An early example may be found in the development of CompStat by the New York City Police Department in the mid-1990s. Under CompStat, police officials collected data in a more systematic way and used it to empower and incentivize police managers of geographic areas (patrol sectors, precincts, or districts) to improve police performance in their command areas. Today, nearly every major police department has a version of CompStat in which management decisions are guided by the collection and dissemination of data to key stakeholders, including precinct

or district commanders and investigators.

Improvements to police practice under CompStat were made possible, in large part, because the police departments collect copious amounts of data including information on where and when crimes happen, the characteristics of crimes and the criminal histories of high-volume offenders. In recent years, investments in information technology have allowed police departments to augment these items with more extensive information on crime patterns and offender networks. Even so, CompStat and its successors have only scratched the surface with respect to fully utilizing the information that police departments possess. While creating crime maps has allowed police departments to allocate patrol strength more efficiently, today's challenges — perpetually low (and falling) clearance rates for serious crimes and tenuous and deteriorating police-community relations — require marshaling different types of information. By improving the process through which information is gathered, curated, and disseminated, police managers can potentially make policing more effective and more precise, thus maximizing the considerable benefits of policing while minimizing its inevitable costs.

A critical factor in data-driven, intelligence-led policing strategies is the ability to plan, collect, analyze, disseminate, and evaluate information. To do this effectively and efficiently, data must be turned into intelligence that provides meaningful information at all levels of the organization, informing decision-making for executive staff, commanders, detectives, and street-level officers. Ultimately, information sharing is a two-way street. On the one hand, it is crucial that data-driven intelligence be shared in a way that is salient and timely for beat officers, the "street level bureaucrats" who, in practice, enforce much of the law. On the other hand, information should flow upwards as well. In particular, it is likely that beat officers, those officers who are closest to the community, possess an enormous reservoir of information that would be helpful to criminal investigations. In many or even most cases, officers may not be aware that the information they possess is valuable to an investigator or an intelligence analyst who is looking to make connections among a number of people or places in order to identify a crime pattern or solve a crime. Likewise, in the absence of a convenient and readily accessible way to share information, it is likely that information sharing will be underutilized by officers.

In the last few years, the PPD has adopted a suite of investments in information technology which promote the transmission of information to officers in a targeted, concise, actionable, and timely format, and allow those officers, in return, to share their observations with the analysts who might be able to put those observations to work. Following other large cities, including New York, Los Angeles, and Chicago, Philadelphia has developed decentralized crime information centers within key police districts and has placed the deployment of

secure mobile devices in the hands of every officer working in those districts.

Staffed by new analysts hired through the city's investment in the Philadelphia Roadmap for Safer Communities and Operation Pinpoint Strategies, these new crime information centers work in coordination with the PPD Real-Time Crime Center (RTCC), putting camera monitoring and data analysis capabilities directly in the hands of units operating in the field who require it the most. In order to interface more directly with beat officers, PPD has, with the assistance of private philanthropy, issued mobile phones to officers in three of the city's high-priority districts. The phones allow officers to search various PPD systems independently, provide GPS tracking for foot beat and bike deployments, and allow officers to share critical information they receive from their communities in an efficient and secure manner, providing analysts the access they need to the latest information.

Critically, the phones also allow officers to interface more directly with community members. By passing out their contact information, officers are beginning to develop and curate networks of community members who supply them with critical information about the going on in their community. This has allowed community members to report local crime problems without making a 911 call and exposing themselves to physical jeopardy or judgment from disapproving neighbors. Through discussions with officers, we have learned that community members have contacted officers on their department-issued mobile phones to let them know where drugs and being dealt and even where guns have been stashed. This information can potentially facilitate beat officers making arrests themselves, but the information may ultimately be most valuable in the hands of intelligence analysts and investigators who are charged with building criminal cases. As the phones have only been in the field for 18 months at the time of this writing, this capacity is still being developed and figures to become more prominent as the department-issued phones are taken to scale and integrated more organically into PPD policy and practice.

This research evaluates the degree to which the provision of phones to beat officers, alongside the creation of crime information centers that facilitate the exchange of information within PPD, has changed police officer behavior or an improvement in public safety. While police officers have, on the whole, used their phones sparingly during the pilot period, we are able to document that they have used the phones to interact with citizens and communicate with other police officers, crime analysts, and investigators. Have the phones made a difference? While we do not find evidence of changes in crime rates in the first year in which phones were deployed in the field, critically, the provision of phones to police officers is associated with an increase in the rate at which violent crimes are cleared by investigators as well as an increase in the rate at which stop and searches in a motor vehicle are substantiated by an arrest. The evidence suggests that the provision of phones, alongside the

crime information centers, have made policing more effective in clearing violence crimes and initiating a small number of arrests during vehicle stops in the three Philadelphia districts in which the intervention was deployed.

In the future, the degree to which the phones will continue to add value figures to be tightly linked to whether their use is evangelized by key stakeholders within PPD as well as the extent to which PPD is able to augment the functionality of the phones with additional features that are useful to officers in the course of carrying out their daily duties. We regard these results as proof-of-concept of the promise of investments in information technology, understanding that PPD might either build upon these successes or fail to replicate them, depending upon the future path that is taken.

The remainder of this report is organized as follows. In Section 2, we provide a brief summary of how the phones were disseminated to officers in the city's three pilot districts and document the extent to which the officers engaged with — or failed to engage with — the phones in practice. In Section 3, we introduce our formal evaluation of the effect of the provision of the phones on several key outcomes including public safety, the extent to which the police clear violent crimes and the efficiency of the police in conducting street stops. In Section 4, we provide a qualitative overview of barriers to successful adoption of the phones and identify several avenues through which the use of the phones can be improved.

2 Institutional Setting

Department-issued phones were distributed to PPD officers in three districts — the 22nd, 24th, and 25th districts — beginning in February 2021. Delayed by the arrival of the COVID-19 pandemic and resulting disruptions to police service in 2020, the launch was approximately one year after the intended start. Initially, the phones were seen by PPD as a vehicle for pushing specific types of location-specific information to officers including reminders as to the location of priority areas within the city as well as information on high-priority crime patterns in and around those areas.

Given that information was to be pushed to officers via the phones, the research team had initially proposed a randomized experiment to evaluate whether improving the supply and utility of information flowing to officers, detectives, and police managers results in more informed decision-making, lower crime rates, and higher case clearance rates in three of the city's highest crime districts. The timing of messages was to be randomized to see if a squad of officers increased their effectiveness during time periods in which more information was pushed to them via the phones. However, after the research team observed the phones in the field and gathered data on their use by pilot district officers, it became apparent that

officers were not responding to push notifications sent by the intelligence bureau and were, in fact, rarely opening these messages at all.

To adapt the research to the ways in which the phones were being used by officers in practice, the research team shifted strategies and focused on studying the overall impact of the phones — and all of the features that the phones bring — rather than a specific feature of the phones such as push notifications. Leveraging the quasi-random roll-out of the phones to different squads within the city's three pilot districts as well as the fact that similar districts in the city have not yet received the phones, we are able to learn how the provision of phones has changed policing in Philadelphia. While the phones have not had a broad impact on crime at this point, we do find evidence that the phones, in conjunction with the crime information centers in these districts, have led to higher clearance rates for violent crimes and better targeted vehicle searches.

In the following section, we provide a brief summary of the parameters of the intervention and review evidence on the extent to which officers have used the phones during the pilot period. In Section 3, we provide a systematic analysis of key outcomes such as crime clearance rates and "hit rates" from pedestrian and vehicle stops as well as public safety.

2.1 Staffing and Schedule Information

We begin with a description of how the phones were rolled out in the three pilot districts, beginning in February 2021. Each district has four squads, numbered 1, 2, 3, and 5/7. In each district, squads 1, 2 and 3 are "line squads" and consist of roughly 40-50 officers. Officers in these squads primarily respond to calls for service and engage in routine preventative patrol only when they are not en route to a service call. Squad 5/7 is a "flex" squad, made up of a tactical team of 40-60 officers in patrol cars, on bike, or on foot. Squad 5/7 officers are more proactive and spend less time responding to service calls. They, therefore, have more discretionary time and likely have greater latitude to gather information from the community and to respond to information pushed by PPD's Intelligence Bureau. It is likewise important to note that foot and bike patrol officers do not have access to mobile desktop computers (MDCs) while on patrol as they are not in squad cars. Hence the provision of a phone that provides at least some of the features of an in-car computer, is a potential important technology to officers on foot or bike patrol.

Table 1 presents staffing numbers for each squad in the 22nd, 24th and 25th police districts — the numbers in the table indicate the number of officers in each of the pilot district's squads. In total, 507 officers were due to receive a phone. Each squad is further divided into three platoons corresponding with an officer's scheduled work shift. Barring

absences or temporary changes in an officer's work schedule, officers in the same platoon work the same schedule and report to the same sergeant. Figure 1 provides a schematic of a squad's work life, focusing on the dates and times when each squad is on the job. Officers in each squad generally work four or five days on and two days off. Among squads 1, 2, and 5, start times rotate throughout the year — sometimes officers work the shift beginning at 8:00am (labeled as "8" in the figure); at other times they will rotate to the shift beginning at 4:00pm (labeled as "16" in the figure). Squad 3 is permanently assigned to work the overnight shift which begins at midnight (labeled as "0" in the figure). We use the features of this schedule in our research design which leverages the timing of the roll-out of phones to each squad.

Table 1: Staffing by District and Shift

	District		
Squad	22	24	25
1	42	43	48
2	41	38	45
3	43	36	40
5	39	37	55
Total	165	154	188

Figure 1: Work Schedule by Platoon and Squad

Day	1A	1B	1C	2A	2B	2C	3A	3B	3C	5A	5B	5C
1	8		16	16		8	0		0			
2		16	16		8	8		0	0	18	18	8
3		16	16		8	8		0	0	18	18	8
4	16	16		8	8		0	0		18	18	8
5	16	16		8	8		0	0		18	18	8
6	16		16	8		8	0		0	18	18	8
7	16		16	8		8	0		0			
8		16	16		8	8		0	0			
9		16	16		8	8		0	0	8	18	18
10	16	16	16	8	8	8	0	0	0	8	18	18
11	16	16		8	8		0	0		8	18	18
12	16	16		8	8		0	0		8	18	18
13	16		8	8		16	0		0	8	18	18
14	16		8	8		16	0		0			
15		8	8		16	16		0	0			
16		8	8		16	16		0	0	18	8	18
17	8	8		16	16		0	0		18	8	18

This figure contains a sample work schedule for each platoon and squad combination in our pilot districts.

2.2 Phone Deployment and Use

Phones were deployed to each squad between late February 2021 and August 2021. Figure 2 shows the timing of phone deployment by district and squad during the pilot period. For all three districts, the phones were issued first to squads 5 and 7 as these squads have the greatest contact with members of the community. As officers may have not been present on the day that the phones were deployed either because they were on leave, out sick, or dealing with other responsibilities such as testifying in court, we also use the first time at which a phone reports a GPS location to indicate phone reception at the officer level.[1]

Next, in Figure 3, we show the time path of phone deployment based on the first GPS connection made by each phone as of April 1, 2021.[2] In total, 435 mobile phones were both deployed to pilot districts and successfully connected to PPD's network during the pilot. The majority of phones were deployed between Spring and Summer 2021.

Figure 2: Phone Deployment Over Time

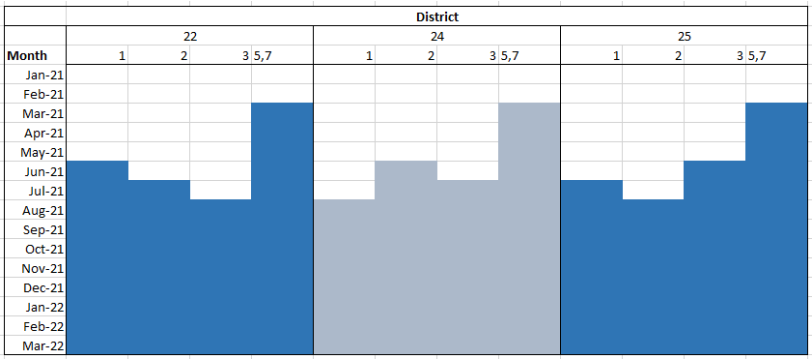

This figure shows phone deployment over time by district and squad. Shaded areas mark the months that officers in that district/squad had assigned phones.

2.3 Information Provision to Officers

One of the ways in which the phones were used in the first year of deployment is to push specific types of information to officers. While the research team assisted PPD in facilitating the dissemination of information, the content of the messaging was the sole responsibility of PPD. In this section, we briefly describe the content of the push notifications and assess officers' responsiveness to the notifications.

[1] This also allows us to account, to some extent, for technical difficulties as phones that are never fully registered to an officer for whatever reason do not report GPS location.

[2] There were issues with GPS data prior to this date.

Figure 3: Phone Deployment by Earliest Connection Date

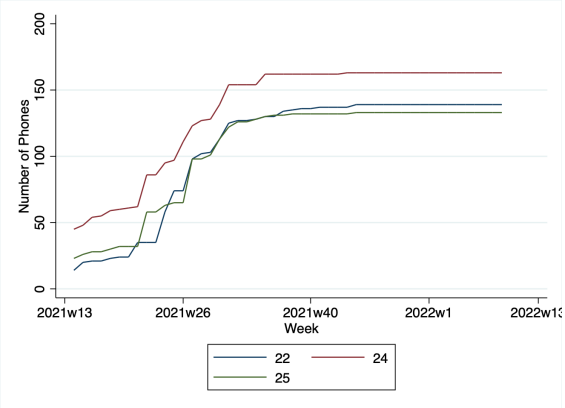

This figure shows the count of phones deployed in each pilot district based on the earliest date at which each phone connected to PPD's network and emitted a GPS signal.

2.3.1 Push Notifications

Information was pushed because it was thought to be both i) aligned with PPD policy and practice and ii) concise, actionable, and timely for officers. A short description of each notification type is provided below.

1. Intelligence: notifications include a mobile version of an intelligence report and solicit an officer's assistance in addressing intelligence gaps; notifications are delivered to officers based on their proximity to situational awareness grids

2. Grid awareness: notifications that alert officers to the locations of situational awareness (SA) grids[3] and Pinpoint (PP) grids[4]

3. Philly311: notifications which remind officers of the Philly311 functionality of the phone through which they can report quality of life issues such as abandoned cars and overgrown lots

4. Control: Virtual Roll Call content only – officers receive a daily debrief on notable incidents and other happenings in their districts that may be relevant to the tour that they are about to start

[3]SA grids are short term areas created by intelligence reports that are generally related to shootings or group violence

[4]PP grids are longer term priority areas determined by the PD's existing Operation Pinpoint which aims to focus deployment and PPD resources on the city's worst neighborhood attractors for crimes.

2.3.2 Acting on Notifications

Our evaluation was originally designed to study the effects of changes in the intensity of information provision via push notifications among platoons in a given work spell. This means that, in order to detect a treatment effect, a sufficient number of officers in each platoon would have had to interact with phones on a regular basis and react to the information they are sent in ways that are measurable to the research team.

Although officers appear to have interacted with their phones in a number of meaningful ways, we observed little evidence that they engaged with the push notifications sent to their phones as part of the pilot. In particular, we noticed that as of January 2022 when the randomization of push notifications ended, officers opened geofenced notifications and intelligence products less than 1% of the time. Although we are not able to directly observe receipt of all other notification types as they are sent via Teams after Virtual Roll Call, the evidence suggests that officer behavior cannot be attributed to these notifications

To test officer responsiveness to the push notifications formally, we ran the following ordinary least squares regression, studying officers' propensity to make 311 requests or file intelligence reports as a function of push notifications.

$$Y_{tm} = \beta_0 Philly311_{tm} + \beta_1 Intel_{tm} + \beta_2 Grid_{tm} + \gamma_t + \lambda_m + \epsilon_{tm} \tag{1}$$

The equation identifies the effect of a given type of push notification on a given outcome of interest at the squad t by month m level. Y_{tm} is either the total number of Philly311 requests or intel reports completed by officers at the district by month level. $Philly311_{tm}$, $Intel_{tm}$, and $Grid_{tm}$ are indicator variables that are 1 if team t received related reminders during month m and 0 otherwise. γ_t and λ_m are team and month fixed effects, respectively. Standard errors are clustered at the team level.[5]

Figure A1 shows parameter estimates from (1) for 311 requests (Panel A) and intel reports (Panel B) based on the schedule of notifications sent via Teams. 95% confidence intervals for our parameters of interest span zero. This evidence, in addition to the fact that officers rarely open push notifications, indicates that it is unlikely that the notifications have resulted in changes in officer behavior. Given these findings, randomized push notifications were discontinued on January 1, 2022 to free up civilian analyst time in the crime information centers so that they could pursue alternative efforts.

[5]There are 45 unique squads among the three pilot districts.

Table 2: Pilot District Intel Reports by Type

Report Method	Count	Percent
Law Enforcement	1006	69.67
Surveillance	277	19.18
Tip	113	7.83
Other	48	3.32

This table contains data on pilot district IIR submissions from February 28, 2021 through April 30, 2022.

2.4 Intelligence Reports

Pilot district officers use the Survey123 app that is pre-loaded on their phones to share intelligence with analysts working in the Crime Information Centers. Those analysts collect the information submitted by officers and store a portion of it in PPD's Infoshare system.[6] The Intelligence Bureau classifies each piece of information based on its source, indicating one of four different sources of information[7]

- Law Enforcement: information received from a law enforcement entity and/or an intelligence analyst working for a law enforcement entity.
- Surveillance: Observations made while operating on authorized surveillance activities
- Tip: Information where the source will not supply a name and there is no way of validating a confidence level in the source
- Other: human source, debriefing, gang vetting report, social media

Table 2 presents the composition of intelligence reports sent by PPD officers using their phones during the pilot period. Law enforcement and surveillance make up approximately 90% of the information shared by pilot district officers. Panel A of Figure 5 shows the volume of these submissions, which have decreased over time for each of the three pilot districts. While we are unable to determine why this is the case, one possibility is that officers may not feel incentivized to share information either because they are not consistently encouraged to do so or because of a lack of feedback from the CICs.

[6] Criminal intel that meets certain standards is stored in Infoshare, a CJIS-compliant information storage system.

[7] The research team learned in March 2022 that officers often select the incorrect source.

2.5 Philly311 Requests and Other Community-Related Activity

In addition to responding to service calls and performing routine patrols, PPD officers are also charged with engaging with the community in ways that are intended to prevent crime. Officers perform checks on local business and knock on doors to inform community residents of local crime problems and provide advice on how to keep themselves and their property safe. They can also call in 311 requests when they identify signs of disorder — for example an abandoned vehicle or a lack of trash collection — that can have implications for community well-being or public safety.

Pilot district officers use Survey123 to record community related activities such as making 311 requests and performing business checks.[8] Table 3 shows the number of these submissions by type. Business checks comprise 87% of all entries. This shows that phones are adopted most readily when they have a function that unequivocally makes the officer's work easier. Prior to phone deployment, officers who conducted business checks had to document these checks on paper. The phones allow them to complete this function via the click of a single button.

Table 3: Pilot District Community Activity by Type

Type	Count	Percent
Business Check	3282	87.01
Philly311 Request	297	7.87
Other	193	5.12

This table contains data on pilot district police activity submissions from February 28, 2021 through April 30, 2022. Categories included in other are Observed Gun/Drug Activity, Community Meeting Attended, Home Visit, Poster/Leaflet Distribution, Victim Follow-Up, Homeless Count.

2.6 Call and Text Activity from AT&T

Apart from the ability to receive information that is pushed to them, the department-issued phones naturally have the same functionality that other phones do and can be used to make and receive phone calls and text messages. While these features of the phones continue to remain highly underutilized, there is promise with this basic functionality once it is adopted

[8] A different application called QuickCapture was used for these submissions at the time of initial phone deployment through November 2022. The research team has combined entries from both applications for the purposes of this report.

at scale.

Across many different areas of public life, research has documented that there is wisdom in crowds, including wisdom that is not easily tapped into by experts working on their own. While experts bring their own brand of knowledge to problem-solving, information is typically costly to acquire and is often distributed diffusely among many different people. In trying to solve serious crimes, investigators are often lack the information necessary to make an arrest. Without a witness, camera footage or a trusted informant who is in the know, it is very unlikely that a crime will be solved. To wit, fewer than 20% of non-fatal shootings were cleared by PPD in 2021.[9] But the information to clear these cases is probably out there. Community members may hold some of that information but may be reluctant to step forward. In other cases, PPD personnel may hold the key to solving a serious crime but they simply do not know that they are holding information that would be vital to an investigation. The problem is that those in possession of useful information may not always recognize that it is useful nor do they necessarily know who the information would be useful to. Compounding these problems, few formal mechanisms are available to collect that information and deliver it to investigators.

Prior to being equipped with a phone, officers did not have a convenient or direct pipeline to receive information from people in the community. Many people are afraid to call 911 to report criminal activity or do not wish to call attention to themselves. Citizens also sometimes do not have a sense for how PPD will respond, who will respond or when a response, if any, will happen. When officers meet with community members and share their direct PPD mobile number, people can communicate directly with an officer who they know and trust. We have been told by officers that community members have told them about street-level criminal activity, including the location of firearms, and arrests have been made on this basis. While these stories remain a small number of examples they underscore the tremendous promise that this technology has to improve policing and make it more responsive to the community once the technology is widely adopted.

In order to understand whether officers are using their phones to procure information from contacts within the community, we analyze the volume of text messages exchanged between pilot officers and members of the community. The research team has access to the universe of call and text activity (only usage, not call or message content) from phones for the deployment portion of the study period — these data are especially useful because

[9]Here we present the clearance rate as the fraction of the number of non-fatal shootings that both occurred in 2021 and were cleared in 2021 over the number of non-fatal shootings that occurred in 2021. This metric is not what the FBI uses, shootings cleared in 2021 (regardless of when they occurred) / shootings in 2021. We believe our measure is more appropriate here because it ensures that all investigations started around or after when phones were deployed.

officers have been able to use these features since the beginning of the phone deployment.[10] Mobility phones, like any other phone, are subject to the mundane annoyance of both robo-calls and robo-texts. Because we do not have information on the identity of the person (or robot) on the other end of the activity we see, we attempt to limit activity to those of true community members by requiring bilateral communication between mobility phones and potential community members. More specifically, for us to count a communication as a genuine community contact we require that a candidate community number send at least one text message or make at least one phone call lasting one minute or more to a PPD mobile phone number AND that a PPD mobile phone *also* sends at least one text message or makes at least one phone call lasting one minute or more to that same candidate community number.

The data are presented in Table 4 which documents phone use at the officer level in each pilot district and Figure 4 which presents the full distribution of phone use for officers that we call "power users." On average, each officer used their department-issued phone to communicate with approximately four unique contacts over the study period. The average number of text messages exchanged varied between 5 in District 22 and 13 in District 24. Underneath these averages, there is wide variation — the standard deviation for the number of text messages is 27 indicating that a small number of officers used their phones quite frequently to message with members of the community. Figure 4 shows the variation in use among the most frequent users in the project. We call an officer a "power user" if he or she connected with at least four unique community members during the project.[11] As Figure 4 shows, officers who took the initiative to use their mobile phones to communicate with community members made a large number of contacts (up to 38) and communicated with those contacts via call or text frequently. We believe use among this subset of officers is indicative of the potential of the mobile phones to directly engage in crime prevention and community problem-solving with community members.

In Panel D of Figure 5 we present text message and call activity per deployed phone over time. The figure shows that use was quite popular at the start of the project and, on average, dropped as additional squads came online. Since the 5/7 squads were first to receive phones, the decline suggests that the 5/7 squad officers in each district communicated most often with community members, a finding consistent with PPD practice.

[10] This period covers the start of the project through August 1, 2021.

[11] This represents the top 25% of users in terms of total community contacts.

Table 4: Officer-Level Phone Use Summary

District	Number of Unique Contacts	Number of Calls	Number of Texts
22	3.43	6.43	4.94
	(4.60)	(12.53)	(14.37)
24	4.09	5.09	13.48
	(5.56)	(9.55)	(26.55)
25	3.74	3.64	11.94
	(4.31)	(5.83)	(11.94)

This table was produced using data from February 28,2021 to August 8, 2021. The table contains means at the officer-level for each district with standard deviations in parentheses. Activity is limited to only that generated by contacts that engaged in bilateral communication with at least one mobility phone.

Figure 4: Distribution of Phone Use for Power Users

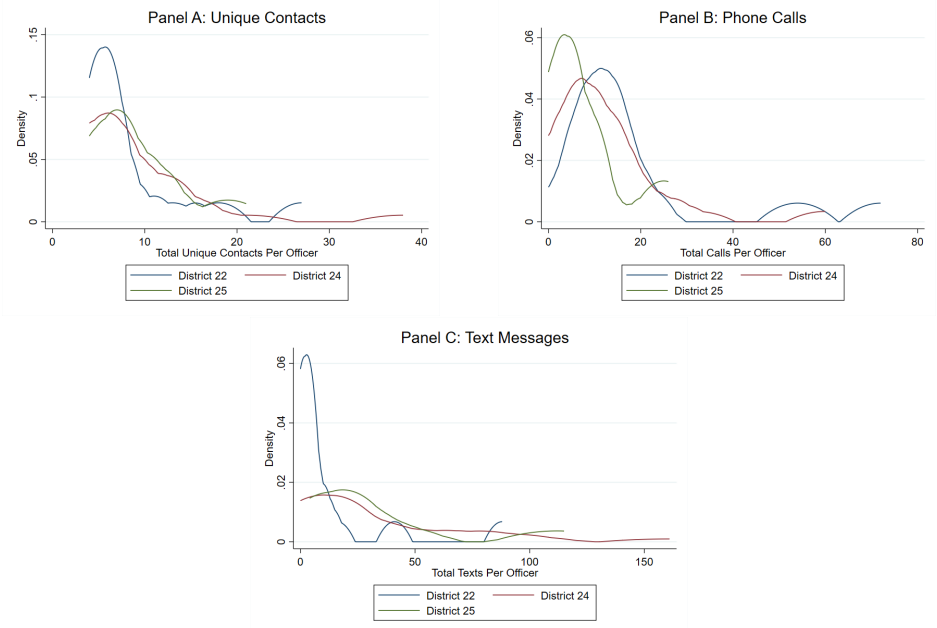

This figure was produced using data on calls and texts from February 28, 2021 through August 1, 2021. The distribution of contacts (Panel A), phone calls (Panel B), and text messages (Panel C) are shown for power users. Power users are defined as officers who made at least four unique community contacts during the period for which the research team has data.

Figure 5: Interactive Phone Use Over Time By District

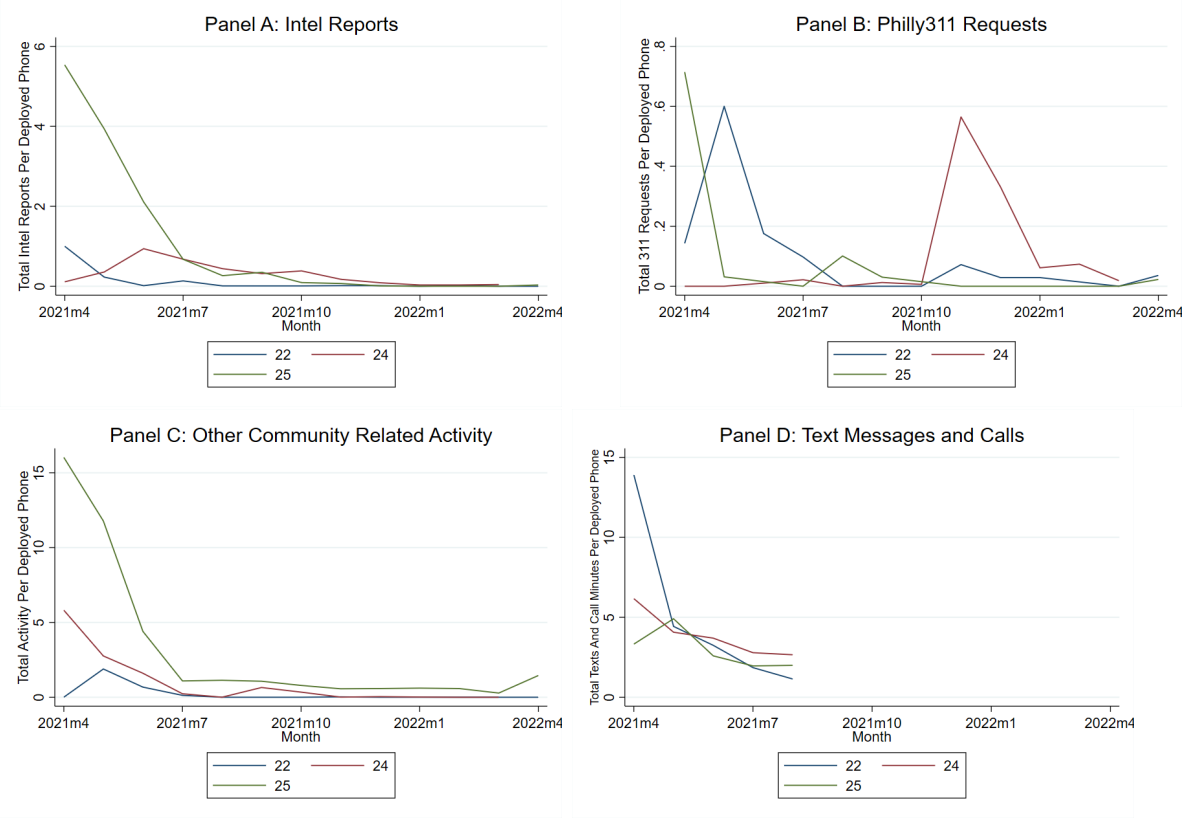

This figure was produced using data on interactive phone activity from February 28, 2021 through April 30, 2022. Note that the research team has access to AT&T data from February 28, 2021 through August 1, 2021. Y-axes in all sub-plots represent the number of entries per deployed phone.

2.7 A Comprehensive Measure of Phone Usage

To provide a convenient summary measure of phone use, we aggregated each of the measures described above into two more global measures of phone use. We call the first measure "interactive activity" and consider an officer to have used the phone on a given day if the phone sent a text to or made a call to a community member, submitted an intelligence report, or submitted community-related activity including Philly311 requests. Put plainly, this is the best measure we have for whether or not an officer interacted with the phone on a given day. Our second measure, "any activity," is meant to measure whether or not we think that a phone was powered on during a given day. An officer is considered active here if the officer's phone reports a GPS point[12] *or* submits any interactive activity from the first

[12]Prior to opening GIS to the internet, a phone must be powered on and connected to the city's proxy to report a GPS location. Phones that migrated to Zscaler must fully complete re-registration for DVIC to receive GPS points from phones. Many officers may not be aware they did not fully move to Zscaler and allow the Mobility application the necessary permissions to share GPS information.

measure.

Figure 6: Comprehensive Phone Use Over Time by District

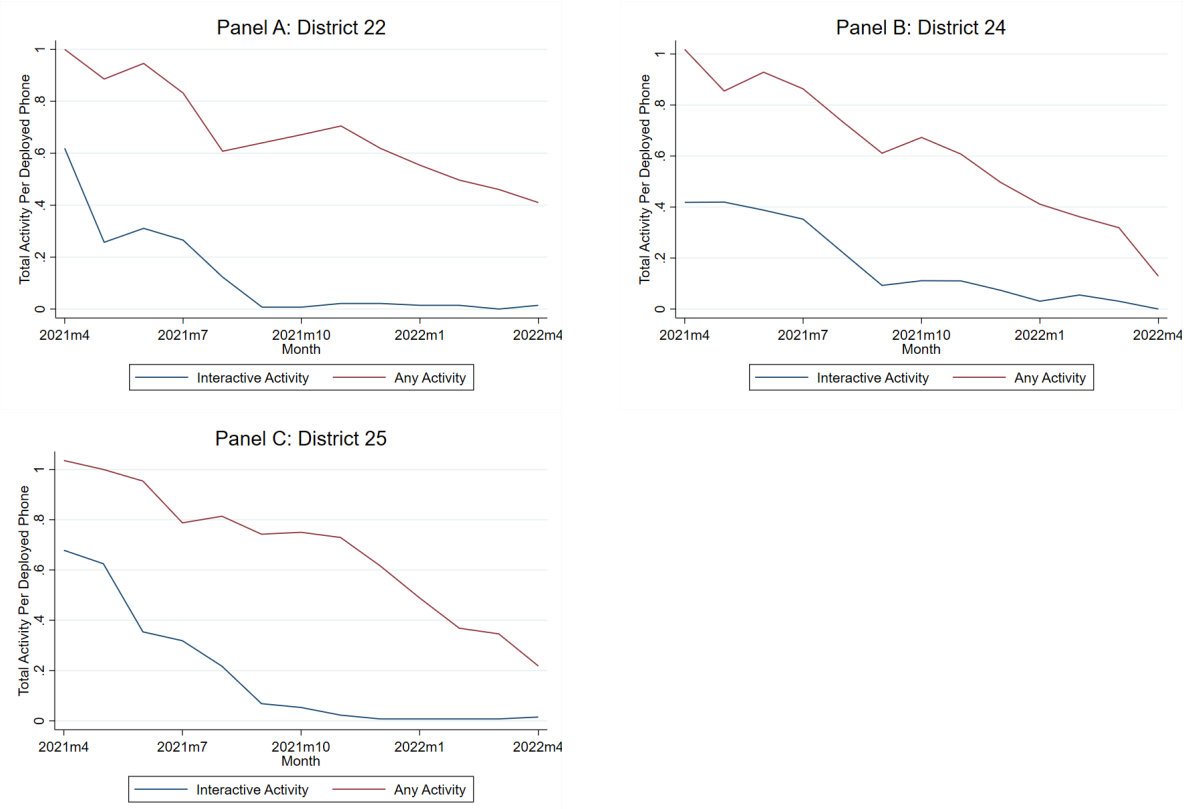

This figure was produced using data on interactive phone activity from February 28, 2021 through April 30, 2022. Note that the research team has access to AT&T data from February 28, 2021 through August 1, 2021. Y-axes in all sub-plots represent the portion of deployed phones that were active according to either measure.

Referring to Figure 6, we see wide variation in phone use by district. Our call and text activity data does not go past August 1, 2021. This data truncation is at least partially responsible for the drop in use around that time. The fourth round of deployment also occurred in August 2021, meaning that the data on the graph after August 2021 is generated by more officers and excludes data on calls and text messages. Over the course of the project, approximately two-thirds of officers appear to have powered on their phones on a given day, though the share of officers powering on their phones appears to be falling over time. With respect to interacting with the phones, phone use is lower. In all three districts, phones were used between 10% and 60% of officer-days throughout the majority of 2021. However, to an extent, the phones have fallen out of use in 2022, especially in District 22 where officers have interacted with the phones very infrequently.

3 Evaluation of the Effect of Mobile Phones

Although officers have not been responsive to the push notifications, we are nevertheless able to evaluate the extent to which the provision of phones along with the work of the crime information centers have had an impact on police officer behavior and public safety. We use the stepped-wedge rollout of the phones (see Figure 2) in each of our pilot districts along with a set of comparison districts that did not receive mobile phones to study the net effect of the Mobility Project on several key outcomes of interest. In each of the pilot districts, phones were rolled out on a squad-by-squad basis over time. As officers in each of a district's squads were issued phones, the number of phones in the hands of officers increased, creating the opportunity to study if policing practices or public safety changed as a function of the number of squads which had access to the phones.

We study the combined effect of the provision of phones and CICs, using a difference-in-differences framework. Differences-in-differences (DiD) methods are well-suited to estimate the effect of sharp changes in policy or practice such as those caused by the Mobility Project. These methods compare outcomes of similar groups (in our case, patrol districts) over time where some groups receive an intervention and others do not.[13] More specifically, DiD designs compare group level averages (treatment districts or officers within them versus comparison districts or officers within them) before and after a treatment to uncover the effect of that intervention on outcomes of interest, such as the number of pedestrian stops made in a district or the share of violent crimes that are cleared in a given district in a month.

3.1 Data and Sample Selection

Our analysis focuses on the three pilot districts that received phones in 2021 — the 22nd, 24th, and 25th districts — as well as six other comparison districts (Districts 12, 14, 15, 19, 35, 39) selected because of their historically high rates of shooting victimization — as well as other crimes — are similar to the pilot districts. Figure A2 presents this information graphically, with each panel presenting the monthly average for all districts using only data from *before* any phones were deployed. In all panels, districts that were not selected for analysis are denoted in gray. Our set of comparison districts are denoted in black, and the three pilot districts denoted using red. As is clear in Panel A, our comparison districts are quite similar to pilot districts with respect to the number of shooting victims. Although not selected specifically for similarities in shooting victim clearances, violent crimes, or violent clearances, Panels B-C show that it also the case that our comparison districts are similar

[13] We compare outcomes in our three pilot districts with those in 6 similar districts that were selected based on violent crimes and shootings.

to pilot districts along those dimensions as well.

Tables 5 and 6 present descriptive statistics on our outcomes of interest both citywide (Column 1) and for the pilot and comparison districts (Column 2). The unit of analysis is the district-month. Districts in the study, on average, have more crime and more activity than the city as a whole. For example, pilot and comparison districts combined average roughly 45% *more* stops and 48% *more* arrests per district-month than do all districts in the city. Referring to Table 6, we see that although our project districts experience significantly *more* violent crimes, including homicides and shootings, they tend to have *lower* clearance rates. Our pilot and comparison districts have a homicide rate that is 1.8 times homicide rate for the city but has a homicide clearance rate that is two percentage points lower. Figure A3 shows district by month level outcome-data averages for pilot districts (black) compared to the similar subsample (gray). The figure shows that, while pilot and similar districts may differ in levels, trends in outcomes are quite similar across the two groups.

3.2 Statistical Methods

To study the effect of the provision of phones, we turn to a series of Poisson regression models. Our empirical strategy is motivated by the following difference-in-differences regression model:

$$\log(E(Y_{dm})) = \beta_0 + \beta_1 Treat\_Level_{dm} + \alpha_d + \gamma_m \tag{2}$$

where d is district and m is month. In (2), $Treat\_Level$ ranges from 0 to 4 and counts the number of squads in district d who have received mobile phones by month m. β_1 represents the public safety benefit of each additional squad receiving phones in pilot districts relative to our subsample of comparison districts. We fit the model using a Poisson regression in order to estimate relative rates among the groups.

We present incidence rate ratios (IRRs), computed as $\exp(\beta_1)$, in all subsequent tables. The incidence rate ratio can be interpreted as a the multiplicative change in a given outcome that results from an additional unit of treatment. For example, $\beta_1 = 1.1$ in a regression with pedestrian stops as the outcome would indicate that pedestrian stops increased 1.1 times, or equivalently by 10%, for each additional squad that receives mobile phones in pilot districts relative to comparison districts. The model includes district (α_d) and month (γ_m) fixed effects. The former control for district-level differences that are fixed over time and the latter control for city-wide crime trends that may vary at the month level. We use robust standard errors for all regressions.

Table 5: Summary Statistics – Policing Outcomes

	(1) Full Sample	(2) Similar Subsample
All Stops	402.11	582.31
	(346.65)	(350.10)
All Stop Hit Rate	0.094	0.10
	(0.057)	(0.051)
Pedestrian Stops	54.77	83.40
	(51.65)	(57.40)
Pedestrian Stop Hit Rate	0.35	0.36
	(0.13)	(0.12)
Vehicle Stops	330.05	479.12
	(299.57)	(273.97)
Vehicle Stop Hit Rate	0.043	0.057
	(0.039)	(0.039)
All Infoshare Entries	41.30	63.72
	(38.80)	(42.08)
Officer-Initiated Infoshare Entries	31.37	49.12
	(32.12)	(36.09)
All Arrests	108.50	160.27
	(74.42)	(81.50)
Violent Arrests	21.08	30.80
	(12.10)	(10.00)
District by Month Observations	546	234

This table was produced using data from 3/1/2020 to 4/30/2022. The table contains means at the district by month level with standard deviations in parentheses. All hit rates are defined as the percentage of stops that resulted in an arrest. Arrests for homicide, aggravated assault, and robbery are considered violent arrests.

Table 6: Summary Statistics – Public Safety Outcomes

	(1) Full Sample	(2) Similar Subsample
Homicide	2.04	3.62
	(2.20)	(2.29)
Homicide Clearance Rate	0.22	0.20
	(0.32)	(0.27)
Shooting Victims	6.12	10.91
	(5.89)	(5.49)
Shooting Victim Clearance Rate	0.13	0.11
	(0.22)	(0.12)
Violent Crimes	57.10	88.43
	(34.12)	(22.76)
Violent Crime Clearance Rate	0.26	0.24
	(0.11)	(0.064)
District by Month Observations	546	234

This table was produced using data from 3/1/2020 to 4/30/2022. The table contains means at the district by month level with standard deviations in parentheses. All clearance rates are defined as the percentage of incidents that are cleared by arrest within 30 days of report. Clearances for homicide, aggravated assault, and robbery are considered violent clearances.

3.3 Results

3.3.1 Main Findings

The primary results of our evaluation are presented in Table 7. We begin with investigatory stops, a key metric of police activity. The addition of newly-issued mobile phones did not lead to a change in the number of investigatory stops made by officers but it did lead to a significant increase in the quality of those investigatory stops as measured by the stops' "hit rate" — the proportion of stops that led to an arrest. As Panel (A), Column (2) shows, the addition of phones to a squad was associated with a 30% increase in the likelihood that a stop was yielded an arrest. Given that the hit rate for vehicle stops in the comparison districts is 10%, the addition of phones to the first squad is estimated to increase the hit rate from 10% to 13%. Considering the rollout of phones to all four squads, we estimate that the hit rate increased from 10% to 28% as a function of the availability of phones.

Notably, the increase in hit rates is being driven entirely by increases in hit rates for vehicle stops. Column (6) shows that each additional squad coming online is associated with an increase in vehicle stop hit rates of 43%. It is feasible that officers in pilot districts, especially those not assigned to patrol cars, received additional information on vehicles and passengers via mobile phones that allowed those officers to make higher quality vehicle stops. The threshold to make a pedestrian stop may have been significantly higher than the threshold to make a vehicle stop during the study period, so it is perhaps unsurprising that the increase in hit rates is significant for vehicle stops only.

Panel (B) of Table 7 presents estimates of the effect of phone provision on violent crimes as well as the number of violent crimes that are cleared through an arrest. We do not observe a significant change in violent crimes in pilot districts as a function of the intervention. However, referring to column (8), we see that each additional squad receiving mobile phones was associated with a 6% increase in the clearance rate for violent crimes. Given a violent crime clearance rate of 24% in the comparison districts, we estimate that the provision of phones to all four squads in the pilot districts increased the violent crime clearance rate to 30% (computed as $24\% \times 1.06^4$). With respect to shootings and shooting clearance rates, we do not observe a significant impact as a result of the phones, though the comparative rarity of shootings means that this outcome is difficult to detect.

In column (9) we study the length of the investigative process and find that the time until clearance *increased* by 6% per squad receiving department-issued mobile phones, which is likely a mechanical artifact of a higher clearance rate.[14] That is, if information from mobile

[14]We allow for up to 30 days for an incident to be cleared. Investigation length is the number of days between incident report and arrest for the subset of incidents that are cleared in 30 days or less.

Table 7: Differenes-In-Differences IRRs: Main Results

			Panel A: Stops			
	(1)	(2)	(3)	(4)	(5)	(6)
	All Stops	All Stop Hits	Ped Stops	Ped Stop Hits	Veh Stops	Veh Stop Hits
IRR	0.947	1.300***	1.000	1.051	0.947	1.430***
	(0.0390)	(0.1000)	(0.0525)	(0.0367)	(0.0390)	(0.142)
Mean of DV	582.3	0.104	83.40	0.356	479.1	0.057
			Panel B: Crimes and Clearances			
	(7)	(8)	(9)	(10)	(11)	(12)
	Viol Crimes	Viol Clearances	Viol Days Clear	Shooting Vics	Shooting Clearances	Shooting Days Clear
IRR	0.977	1.061**	1.058*	1.019	1.121	1.034
	(0.0213)	(0.0296)	(0.0351)	(0.0275)	(0.0804)	(0.107)
Mean of DV	88.43	0.240	3.541	10.91	0.110	6.763
Observations	234	234	234	234	234	234

District and month fixed effects are included in all regressions. Hits are defined as stops that result in an arrest. The exposure is set to the total number of stops for all regressions where the outcome is stop hits. A crime is considered cleared if an arrest is made within 30 days of the report of the incident. The exposure is set to the total number of crimes for all regressions where the outcome is clearances. IRRs; robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

phones increased the likelihood that crimes that ordinarily require longer investigations were cleared, we would expect a mechanical increase in the average time to clear an incident in pilot districts relative to comparison ones.

Figure 7 presents these results graphically. The x-axis in each panel represents the number of squads in each pilot district that have received phones. The y-axis represents the average of our outcome of interest at the district-by-month level. Each plot has a horizontal, gray dashed line at the pre-deployment, pilot district mean. Rather than presenting incidence-rate ratios, we present estimates of the increase in a given outcome. For example, in B2, each district-month has an average of around 350 vehicle stops with hits or vehicle stops that result in an arrest. After the first round of deployment, pilot districts experienced roughly 150 additional vehicle stops with hits compared to similar districts. By the final round of deployment, pilot districts experienced over 1000 additional vehicle stops with arrests compared to similar districts.

3.3.2 Robustness

Finally, we probe the robustness of our results to alternative models which we could have reasonably selected or even preferred. We present results for two other difference-in-differences models below. Both models are of the form

$$\log(E(Y_{dm})) = \beta_0 + \beta_1 Treat\_Ind_{dm} + \alpha_d + \gamma_m \tag{3}$$

where d is district and m is month. In (3), $Treated\_Ind_{dm}$ is an indicator variable. In our first supplemental model, (S1), $Treated\_Ind_{dm}$ is equal to 1 if *any* squad in pilot districts has received mobile phones and 0 otherwise. In our second alternative specification, (S2), $Treated\_Ind_{dm}$ is equal to 1 when *all* squads in pilot districts has received mobile phones and 0 otherwise. IRRs for these supplemental models are in columns (2) (the "Any Phone" column) and (3) (the "All Phones" column) in Appendix Tables B1, B2, B3, and B4. As is clear in the tables, our results are robust to both of these operationalizations of the treatment.

3.3.3 Event Study and Falsification Test

Next, we provide several additional tests to probe the validity of the quasi-experimental research design. First, we conduct an event study in order to test the parallel trends assumption that is required for the validity of a difference-in-differences research design. This assumption requires that, all else equal, outcomes in pilot and comparison districts *would* have followed similar trends absent treatment. In order to investigate this assumption, we estimate an event study model in which we estimate effects for each relative time period

for our entire subsample as well as for the pilot districts. If it is the case that any of our month-level estimates for pilot districts are significantly different from 0 in the pre-period, we would be concerned that that the assumption of parallel trends assumption does not hold for a given outcome. Appendix Figure A4 presents the results of these tests graphically for our primary outcomes of interest. As is clear in the figure, we pass the parallel trends test for all stop hits, vehicle stop hits, and violent clearances by non-district officers. For violent crime clearances, there are multiple pilot-month-specific point estimates that are significantly different from 0. While this finding raises the possibility of confounding variation, the absence of a particular trend as well as the absence of a violation of parallel trends for other outcomes reduces our anxiety that this is an important source of confounding.

Finally, we conduct non-parametric falsification tests that aim to measure the likelihood that our findings could be due to chance rather than to the intervention itself. More specifically, we are interested in understanding if it could be the case that, using our preferred statistical models, we find similarly compelling treatment effects elsewhere in the city where there was no intervention. If so, we would lose faith in our models as there is no treatment effect in these areas by definition. In order to do this, we conduct 1,000 simulations, evaluating 1,000 fake interventions that are due to chance, and compare our actual findings with the simulated distribution of placebo findings.

In each simulation, we randomly select nine districts from all PPD patrol districts *excluding* pilot districts as they were in fact treated. We then choose three of those nine districts to be "treated" in simulation world and run our Poisson models on the data from those districts *as if* our simulated pilot districts received phones at the same time as our true pilot districts did. Appendix Figure A5 presents simulated z-score distributions for our falsification tests. In each sub-plot, a red vertical line indicates the true z-score. It is clear from the figure that our true z-scores are far to the right of each simulated distribution, suggesting that it is very unlikely that the treatment effects we found are purely due to chance.

3.3.4 Extension – Violent Crime Clearances

In Appendix Table B5 we further investigate our principal result for violent crime clearances by dividing incidents based on the composition of officers credited on an arrest. The purpose of this analysis is two-fold: (1) to better understand the mechanisms through which the provision of phones led to an increase in clearance rates and (2) to ensure that our principal findings accord with a commonsense notion of how cases are cleared by PPD.

We study whether the increase in clearances was driven by arrests made by patrol officers or investigators, separating arrests into those made by at least one patrol officer and those

which involved no patrol officers. The table presents IRRs for the following groups of arresting officers: (1) any clearance, (2) at least 1 patrol officer listed as an arresting officer, and (3) no patrol officers listed as arresting officers. Consistent with our understanding of PPD practice, the increase in clearances is driven, to a large extent, by arrests made by investigators rather than patrol officers, in particular investigators who investigate crimes in the pilot districts. This is an important finding for two reasons. First, it suggests that the arrests were facilitated by information sharing, either directly by officers or via information processed by the crime information centers. Second, given that violent crime arrests in our pilot districts are infrequently made by patrol officers and should not be made by investigators working on crimes in other districts, the result is sensible and fits with our knowledge of PPD practices.

4 Discussion

This research evaluates the degree to which the provision of phones to beat officers, alongside the creation of crime information centers that facilitate the exchange of information within PPD, has changed police officer behavior and improved public safety. While police officers have, on the whole, used their phones sparingly during the pilot period, we are able to document that they have used the phones to interact with citizens and communicate with other police officers, crime analysts, and investigators. Have the phones made a difference in public safety? While we do not find evidence of changes in crime rates in the first year in which phones were deployed in the field, critically, the provision of phones to police officers is associated with an increase in the rate at which violent crimes are cleared by investigators and an increase in the rate at which stop and searches in a motor vehicle are substantiated by an arrest. The evidence suggests that the provision of phones, alongside the crime information centers, have made policing more effective in clearing violent crimes and initiating a small number of arrests during vehicle stops in the three Philadelphia districts where the phones were deployed.

The fact that the phone use was sparse and yet still showed a positive effect on some outcomes suggests that Philadelphia has more to gain from greater use of mobile phone technology. PPD will soon expand the use of mobile devices to all officers, creating an opportunity for greatly improved use of information but also carries the risk of nothing changing.

The mobile devices offer the potential to

- improve officer safety
- encourage better connections with the community
- push important information to officers
- get investigative questions answered, particularly in stalled investigations

In order to gain these benefits, PPD needs to prepare incentives for officers to use their PPD phones. One clear message from our analysis is that phone usage was low, far lower than we anticipated. In a culture where we regularly work with our mobile phones, we expected PPD officers to all become power users. However, we found that few officers made mobile phone usage a regular part of their work.

To encourage mobile device usage PPD will need to

- ensure that the phones are easily accessible and easily connected to PPD servers

- make the phones an indispensable part of the job
- monitor officers' phone usage and respond to under-utilization through problem-solving and accountability measures

Lastly, PPD should continue to evaluate this important information intervention. PPD has much to learn from the deployment of new mobile devices and those lessons will be important for other law enforcement agencies as well.

PPD should continue an evaluation program by

- Continue to collect data on phone usage, including the volume of calls, text messages, IIRs, and 311 requests
- Distribute new phones or new mobile applications in a manner that allows for rigorous evaluation

4.1 Potential of Mobile Devices for Policing

Knowing where officers are located and communicating critical information to them are fundamental components of officer safety. Prior to phone deployment, PPD managers could not determine the real-time location of officers assigned to foot and bike beats except by asking them to report their location over the radio. When foot and bike officers responded to an incident, they only had information reported over the radio. Officers in cars would have access to substantially more information including the ability to run people and cars through their MDTs. Mobile devices can solve both of these problems *if* officers carry their phones, keep them charged and on, and PPD provides the information services currently available on MDTs to the phones as well.

Mobile devices can also help improve connections with community members. First, officers do not need to remain tied to their car MDTs to stay connected. Second, we heard from several officers who were sharing widely their newly issued PPD mobile number to community members. They were (reasonably) reluctant to distribute their personal mobile number, but comfortable giving their direct contact information when it involved a department-issued mobile device. Members of the community, business owners, school officials, and others would directly call or text PPD officers who they knew personally, even if they felt uncomfortable dialing 911 or calling the local district police station.

Push notifications can be a valuable part of an information intervention, one that we had anticipated (incorrectly) would be an important feature of the initiative. We thought that officers would receive orders to spend extra time in Pinpoint areas (targeted small geographic areas with the most severe crime problems), reminders to report disorder concerns through a

mobile 311 request, and questions from investigators and analysts who need the street level expertise of PPD's front line officers. We designed geofenced notifications, notifications about specific crime problems in specific locations that would only be triggered when an officer passed through the relevant area. Although we worked with PPD to push such notifications, few officers ever noticed the notifications. Nonetheless, the potential remains.

Notifications may be particularly useful for connecting the neighborhood knowledge of patrol officers with investigations. If officers become actively engaged with notifications, then investigators can use that pathway to query relevant officers about people, vehicles, places, relationships, and other features that officers might know well. Investigators can ask "Has anyone seen X lately?" or "Does this vehicle look familiar?" or "Has anyone had a recent dispute with [name of person] recently?" Officers we met indicated that they know a lot about the people and places in their districts, but may not know that an active investigation needs their street-level expertise. A mobile connection between investigators and patrol would flatten the organization and improve the exchange of critical information.

Some of these issues require technical solutions. For example, initially phones had to be turned on and connected to a proxy to report GPS data. The proxy connection would regularly disconnect so that officers needed to log on multiple times during a shift. The proxy eventually was discontinued when phones were migrated to ZScaler. However, it is possible for a phone to enroll in ZScaler without completing the final step that allows the Mobility Application on the phones to send PPD real-time location information. PPD needs to ensure that phones fully complete enrollment such that all phones, if turned on, report GPS information. In addition, a "mobile MDT" function is not yet fully functional.

Technology, however, is only part of the problem and solution. Much of the topics discussed here (officer safety, connections with community, notifications) are organizational issues that deal with the culture of a police department. When it becomes the norm for officers to share their mobile number, PPD will get better contact with concerned community members. When analysts push highly relevant notifications and PPD officers look for them, then PPD can focus better on key problems. Achieving this potential depends on officers using their mobile devices.

4.2 Encourage Mobile Device Usage

Phone usage will remain light until they become a regular part of PPD business. Moving mundane administrative tasks to mobile applications may be an effective way to encourage mobile phone usage. As noted in Section 2.5, officers used their phones most readily when the phones made a required task easier to perform. Officers tasked with completing business

checks, once a paper-based process, were among the most frequent users of QuickCapture and Survey123. Because of this response, PPD began developing an electronic patrol log to replace the department's current paper-based log. Such transitions from paper to mobile devices will make officers' daily work easier, allow the department to track activity that typically goes unrecorded, and encourage the regular practice of using PPD mobile devices. Foot and bike-beat patrol officers in the Kensington special district also requested a "mobile MDT" capable of displaying information on calls for service and other incidents in the same way that the MDT in a patrol car does. If this were to be developed and made available, phone use among officers who patrol without access to an MDT would increase dramatically.

Increasing the availability of administrative applications on the mobile devices is a soft method of encouraging usage, but PPD will need to have clear policies and expectations on mobile phone use. PPD has effectively deployed body-worn cameras to officers and is ensuring that those are used according to policy. The department may benefit from having individuals in each unit (similar to body-worn camera technicians) who are able to help officers with issues and ensure compliance with policy at the unit level. PPD should also monitor the use of phones and respond appropriately to officers not using their phones. Perhaps, the lack of phone use is purely technical and requires a technical fix. Any other reason for lack of usage is important for PPD to understand. Perhaps the officer needs additional instruction on using the application, or is hesitant to distribute their mobile number to community members, or does not find the applications useful. All of these responses require a reaction either to improve the program or improve officer training.

In one district in which we found very high phone utilization, that utilization was due to a command structure that insisted on officers using their phones and encouraging their officers to use the phones in specific ways to address crime problems. The degree to which the phones will continue to add value figures to be tightly linked to whether their use is evangelized by key stakeholders within PPD as well as the extent to which PPD is able to augment the functionality of the phones with additional features that are useful to officers in the course of carrying out their daily duties. Ultimately, the adoption of technology in any profession benefits from individuals directly involved in the work evangelizing the benefits to their colleagues. The best advocacy for using the phones will likely come from officers themselves. Officers who have used the phones to gather key intelligence and facilitate high-value arrests should be invited to share their stories with their colleagues. Even a small number of examples are likely to be highly memorable.

4.3 Evaluation

Our evaluation benefited from a staged rollout of the mobile devices to different districts and squads. It was already not possible to hand out and activate all mobile devices simultaneously. By rolling out the distribution slowly over time, we were able to compare squads who received phones to squads who had not received them yet. We could also explore how the "dosage" of phones affected outcomes as a district had more and more officers receive phones. Although our experiments with notifications had no impact, we were able to randomize their distribution so that we could test whether those officers receiving notifications responded differently from those who did not receive notifications (e.g. like whether a reminder about 311 requests resulted in more 311 requests). Such evaluation designs depend on PPD not delivering all phones and all information to all officers at the same time.

As PPD expands phone usage, they should continue the practice of rolling out in stages over time, in part because it is not feasible to do otherwise but also so that they can evaluate the effect of the new phone expansion. The same applies as new mobile applications become available. For example, if a mobile MDT app becomes available, PPD should randomly select a set of officers to have access and delay access to all officers until later. Partly, this will be worthwhile to test the new mobile app, but also PPD will be able to evaluate how officers are using the MDT information and whether it is producing the desired effects (e.g. more surgical policing, officer safety). Randomization and slow rollouts of new devices and applications set up PPD to be a learning organization, one that can produce rigorous evaluations with evidence with value that accrues to the residents of Philadelphia and to law enforcement agencies elsewhere who can take the lessons learned.

In conclusion, we found that PPDs deployment of mobile devices and CICs resulted in an increase in the violent crime clearance rate from 24% to 30% and an increase in the likelihood that a stop resulted in an arrest — from 10% to 28%. This finding was driven by the few squads and few officers that heavily used the mobile devices. These results are promising. They show that PPD can gain from an information intervention and potentially has more to gain from expanded usage of mobile technology.

Figure 7: Estimates

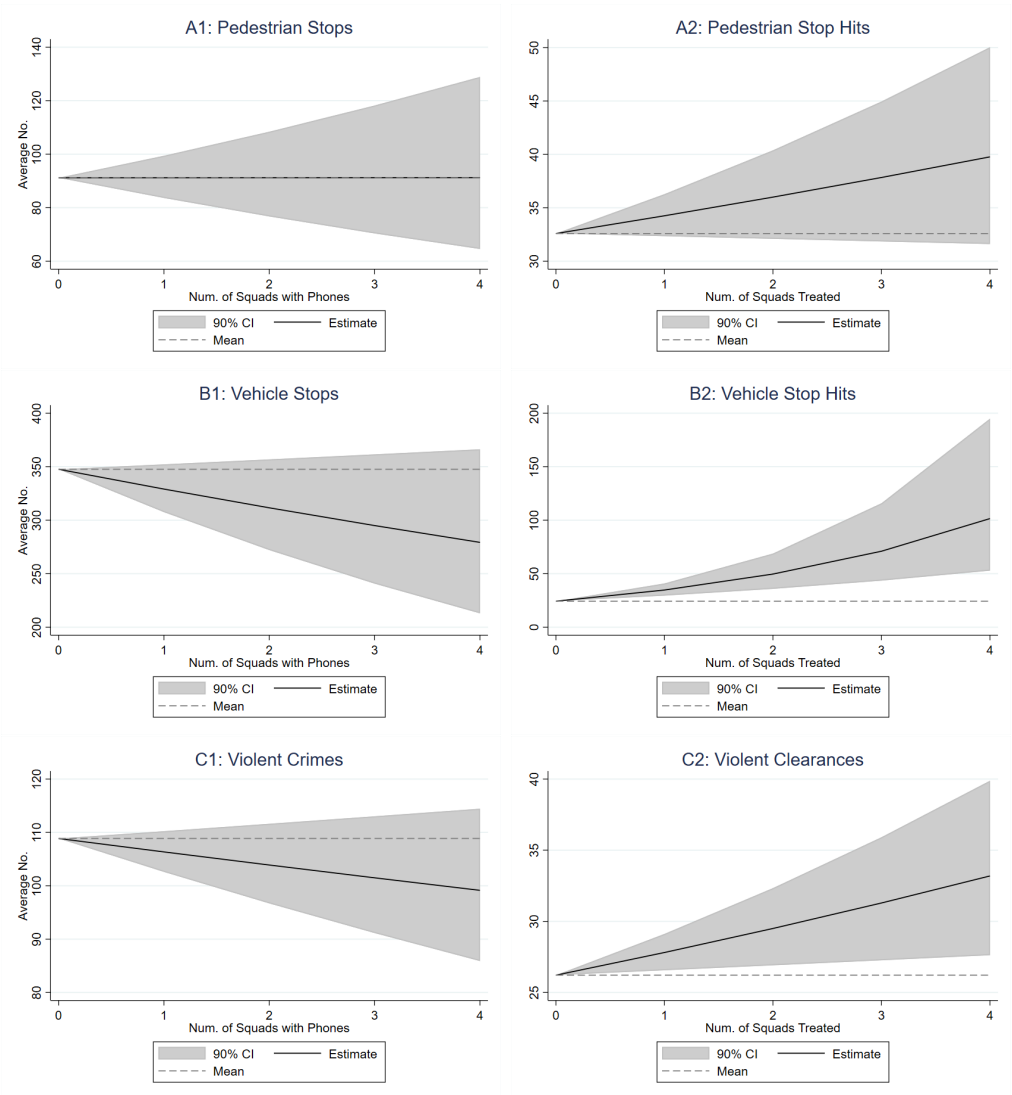

This figure presents estimates and related 90% confidence intervals for pedestrian stops (A1), pedestrian stop hits (A2), vehicle stops (B1), vehicle stop hits (B2), violent crimes (C1), and violent crime clearances (C2). The gray dashed lines in each sub-plot is the pilot district mean.

Appendix A: Additional Figures

Figure A1: Regression Results: Acting on Notifications

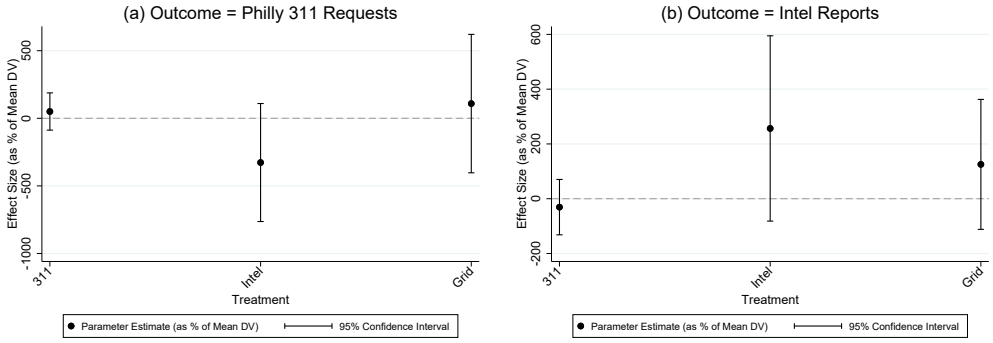

This figure contains parameter estimates and 95% confidence intervals, both expressed as a percent of the mean of the dependent variable, for the difference-in-differences regressions described in Section 2.6.

Figure A2: Sample Selection Visualization

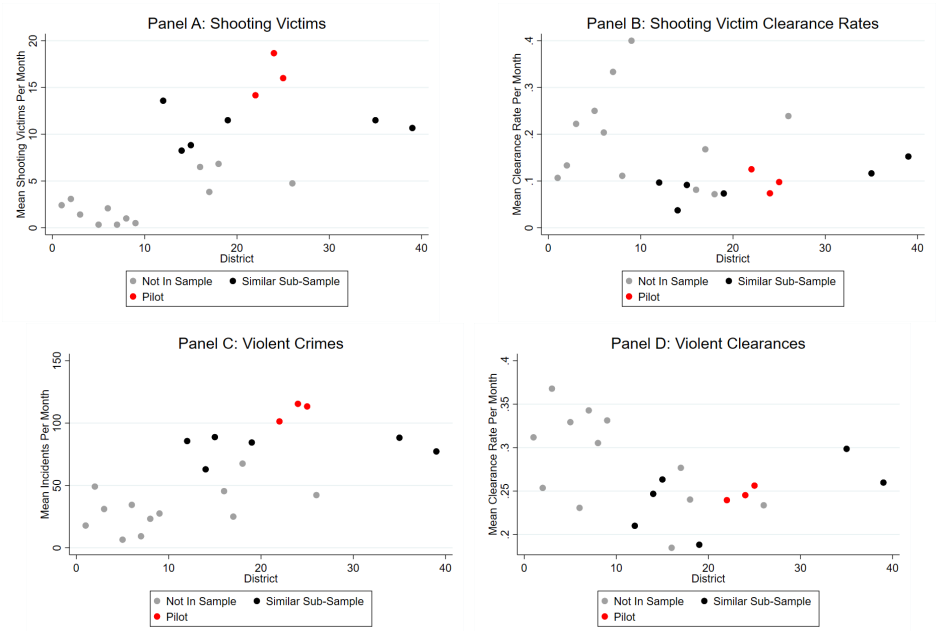

This figure was produced using data from March 1, 2020 to February 28, 2021. Means are taken at the district level using only pre-phone deployment data. Police district is listed on the x-axis and pre-deployment means are shown on the y-axis.

Figure A3: Outcome Data Time Series

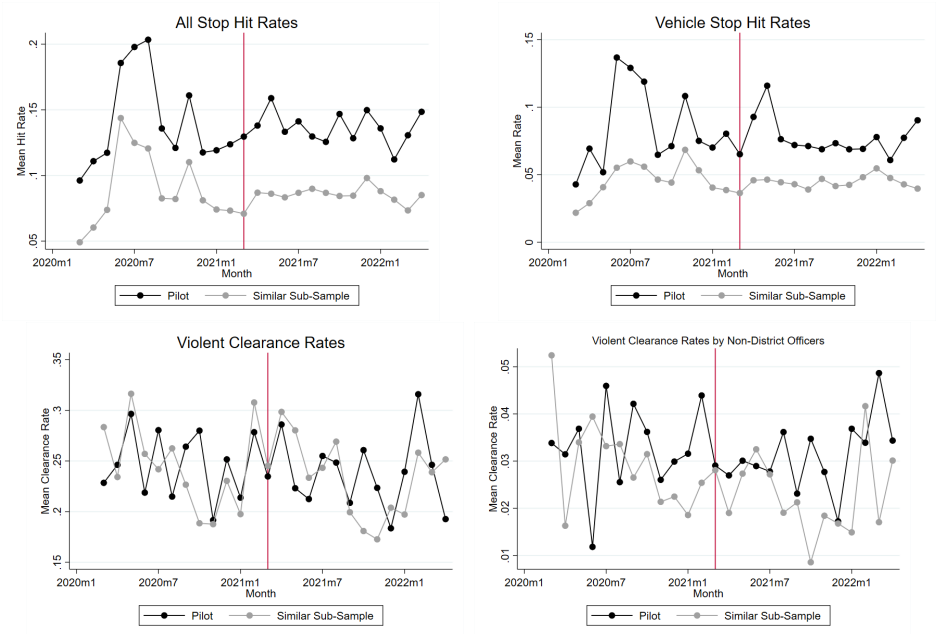

This figure was produced using data from March 1, 2020 to April 30, 2022. Means are taken at the group by month level. Pilot districts and other similar districts each make up one group. The red vertical line is at March 2021, the month the first set of phones were deployed.

Figure A4: Event Study Figures

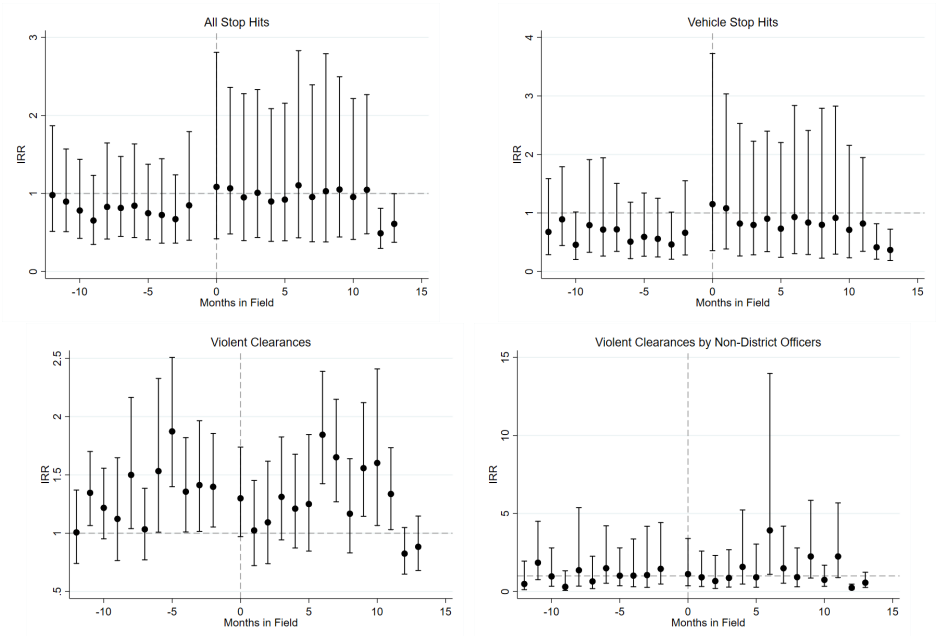

This figure was produced using data from March 1, 2020 to April 30, 2022. Estimates of β from described in 3.3.3 are shown with associated 95% confidence intervals.

Figure A5: Simulation Z Score Distributions

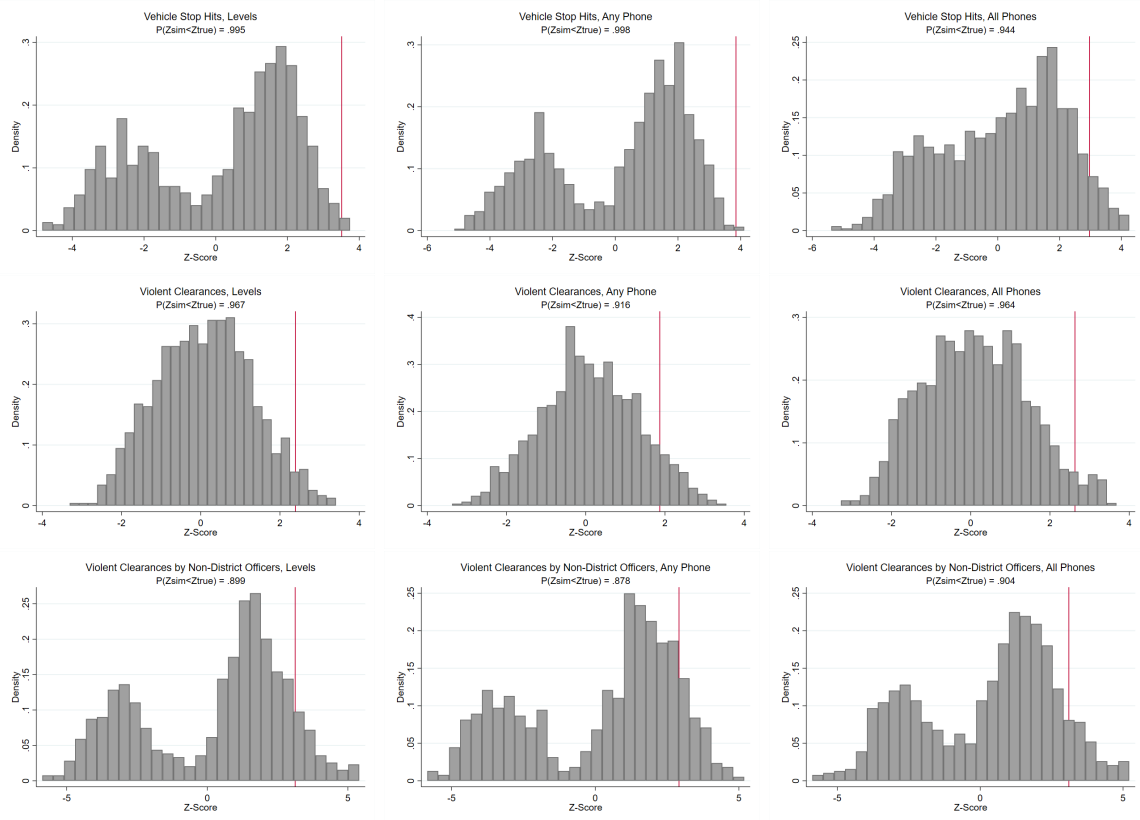

The first column of figures presents simulation z-scores for the model in (2). The second column uses a model that considers a district month treated if *any* squad in that district by month bin has received mobile phones. The third column uses a model that considers a district month treated if *all* squads have received mobile phones. Red vertical lines are at each model's true Z-Score.

Appendix B: Additional Tables

Table B1: Differences-In-Differences IRRs: Stops

	(1) Levels	(2) Any Phone	(3) All Phones
All Stops	0.947	0.806	0.824
	(0.0390)	(0.129)	(0.114)
Mean of DV	582.3	582.3	582.3
All Stop Hits	1.300***	2.834***	2.335***
	(0.1000)	(0.812)	(0.669)
Mean Hit Rate	0.104	0.104	0.104
Pedestrian Stops	1.000	0.990	0.986
	(0.0525)	(0.204)	(0.167)
Mean of DV	83.40	83.40	83.40
Pedestrian Stop Hits	1.051	1.206	1.161
	(0.0367)	(0.163)	(0.139)
Mean Hit Rate	0.356	0.356	0.356
Vehicle Stops	0.947	0.805	0.826
	(0.0390)	(0.128)	(0.115)
Mean of DV	479.1	479.1	479.1
Vehicle Stop Hits	1.430***	4.323***	3.180***
	(0.142)	(1.616)	(1.187)
Mean Hit Rate	0.057	0.057	0.057
District by Month Observations	234	234	234

District and month fixed effects are included in all specifications. Hits are defined as stops that result in an arrest. The exposure is set to the total number of stops for all regressions where the outcome is stop hits. IRRs; robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B2: Differences-In-Differences IRRs: Crimes and Clearances

	(1) Levels	(2) Any Phone	(3) All Phones
Violent Crimes	0.977	0.923	0.909
	(0.0213)	(0.0818)	(0.0644)
Mean of DV	88.43	88.43	88.43
Violent Clearances	1.061**	1.206*	1.243**
	(0.0296)	(0.133)	(0.119)
Mean Clearance Rate	0.240	0.240	0.240
Violent Crime Days to Clearance	1.058*	1.311**	1.177
	(0.0351)	(0.168)	(0.153)
Mean of DV	3.541	3.541	3.541
Shooting Victims	1.019	1.031	1.076
	(0.0275)	(0.105)	(0.105)
Mean of DV	10.91	10.91	10.91
Shooting Victim Clearances	1.121	1.588*	1.485
	(0.0804)	(0.422)	(0.394)
Mean Clearance Rate	0.110	0.110	0.110
Shooting Victim Days to Clearance	1.034	1.302	1.016
	(0.107)	(0.479)	(0.414)
Mean of DV	6.763	6.763	6.763
District by Month Observations	234	234	234

District and month fixed effects are included in all specifications. A crime is considered cleared if an arrest is made within 30 days of the report of the incident. The exposure is set to the total number of crimes for all regressions where the outcome is clearances. IRRs; robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B3: Differences-In-Differences IRRs: Intel

	(1) Levels	(2) Any Phone	(3) All Phones
All Infoshare Entries	0.891**	0.628**	0.726**
	(0.0408)	(0.114)	(0.105)
Mean of DV	63.72	63.72	63.72
Officer-Initiated Infoshare Entries	0.851***	0.430***	0.737**
	(0.0374)	(0.0727)	(0.113)
Mean of DV	49.12	49.12	49.12
District by Month Observations	234	234	234

District and month fixed effects are included in all specifications. IRRs; robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B4: Differences-In-Differences IRRs: Arrest Incidents

	(1) Levels	(2) Any Phone	(3) All Phones
All Arrest Incidents	0.991	0.972	0.955
	(0.0230)	(0.0905)	(0.0709)
Mean of DV	131.1	131.1	131.1
Violent Arrest Incidents	0.988	0.923	0.974
	(0.0234)	(0.0861)	(0.0776)
Mean of DV	27.08	27.08	27.08
Violent Gun Arrest Incidents	1.015	0.971	1.061
	(0.0281)	(0.105)	(0.107)
Mean of DV	10.62	10.62	10.62
District by Month Observations	234	234	234

District and month fixed effects are included in all specifications. Violent arrests are those made for homicide, aggravated assault, and robbery. IRRs; robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B5: Differences-In-Differences IRRs: Violent Crime Clearances by Officer Type

	(1) Any Clearance	(2) At Least 1 Patrol	(3) No Patrol
Number of Squads with Phones	1.061**	1.057*	1.301***
	(0.0296)	(0.0325)	(0.104)
Mean Clearance Rate	0.240	0.214	0.028
At Least 1 Squad with Phones	1.206*	1.200	2.521***
	(0.133)	(0.144)	(0.778)
Mean Clearance Rate	0.240	0.214	0.028
All Squads with Phones	1.243**	1.210*	2.686***
	(0.119)	(0.129)	(0.773)
Mean Clearance Rate	0.240	0.214	0.028
District by Month Observations	234	234	234

District and month fixed effects are included in all specifications. A crime is considered cleared if an arrest is made within 30 days of the report of the incident. The exposure is set to the total number of crimes for all regressions. Column (1) uses the full sample of violent crime clearances. Column (2) counts only clearances made where at least one arresting officer was assigned to a patrol district. Column (3) counts only clearances made where no arresting officers were assigned to a patrol district. IRRs; robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$