

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Crime is Terribly Revealing: Information Technology and Police Productivity

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1763397> since 2020-11-30T17:36:29Z

Published version:

DOI:10.1093/restud/rdaa009

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Crime is Terribly Revealing: Information Technology and Police Productivity*

Giovanni Mastrobuoni[†]

October 2019[‡]

Abstract

An increasing number of police departments use information technology (IT) to optimize patrolling strategies, yet little is known about its effectiveness in preventing crime. Based on quasi-random access to “predictive policing,” this study shows that IT improves police productivity as measured by crime clearance rates. Thanks to detailed information on individual incidents and offender-level identifiers it also shows that criminals strategies are predictable. Moreover, the introduction of predictive policing coincides with a large negative trend-discontinuity in crime rates. The benefit-cost ratio of this IT innovation appears to be large.

Keywords: predictive policing, IT, police, crime, quasi-experiment
JEL classification codes: O33, K42, L23, H1, H41

*I would like to thank the Police Chief of Milan (*Questore di Milano*) for providing the data, as well as Mario Venturi and his staff for sharing their knowledge on robberies and policing with me. I would also like to thank the editor Uta Schönberg and the anonymous reviewers for their time and valuable remarks. I would also like to thank Phil Cook, Arie Kapteyn, Mark Kleinman, Theodore Koutmeridis, Ignacio Monzon, Daniel Nagin, Santiago Oliveros, Emily Owens, Justin Rao, David Rivers, and John Von Reenen for invaluable comments. Financial support from the Italian Ministry of Education, University and Research (MIUR) “Dipartimenti di Eccellenza” grant 2018-2022 is gratefully acknowledged.

[†]Collegio Carlo Alberto, University of Turin (ESOMAS) and University of Essex, Email: giovanni.mastrobuoni@carloalberto.org.

[‡]© 2015 by Giovanni Mastrobuoni. Any opinions expressed here are those of the author.

“But, yes, Hastings, I think it is almost certain there will be another. A lot depends on *la chance*. So far our *inconnu* has been lucky. This time the luck may turn against him. But in any case, after another crime, we shall know infinitely more. **Crime is terribly revealing.** Try and vary your methods as you will, your tastes, your habits, your attitude of mind, and your soul is revealed by your actions. There are confusing indications - sometimes it is as though there were two intelligences at work - but soon the outline will clear itself, I shall know.” (Agatha Christie, 1936)

1 Introduction

Over the past 30 years, organizations have dramatically increased their use of information technology (IT). Their purpose is often to predict individual behavior, of patients, consumers, firms, taxpayers, or criminals. E-commerce companies like Google, Amazon, and Netflix develop applications that use individual browsing and purchase history to customize search results.¹ Social network companies like Facebook and Twitter use individuals’ networks and posting behavior to suggest new connections.²

The public sector also develops predictive algorithms to improve resource allocation. Internal revenue services use data mining techniques to forecast fraudulent behavior, including tax evasion (see Bolton and Hand, 2002, for a review). Government agencies inspect buildings and businesses that are predicted to be at risk of fire or at risk of violating health regulations. Pattern recognition algorithms use satellite pictures to detect illegal buildings. The US Environmental Protection Agency and several agencies around the world use analytics to predict catastrophic events. Increasingly, agencies use predictive models of air and water pollution to regulate emissions. Departments of Defense apply

¹These are known as recommendation algorithms (see Adomavicius and Tuzhilin, 2005).

²For an overview on IT and productivity see Brynjolfsson and Saunders (2009), Van Reenen et al. (2009), and Goldfarb and Tucker (2019).

data science to predict a variety of different threats: from civil unrest to terrorism and cyber-terrorism, but also veteran suicides, and Army vehicle failures. Criminal justice systems use predictions for risk assessments in parole decisions, sentencing decisions, prison treatment programs, and decisions about bail and court appearances (for a review, see Desmarais and Singh, 2013).

The practice of using statistics to predict where crimes will occur and focusing police resources in these areas is spreading across law enforcement agencies. Between 1987 and 2003 the proportion of agencies that use IT for criminal investigations, dispatch and fleet management, went up from 11, 9 and 7 percent to 59, 58, and 34 percent. In 2013 more than 90 percent of agencies used IT to maintain and analyze criminal incident reports, and more than half of these were using exact crime addresses for mapping purposes.³

The adoption of IT by the public service may generate huge gains. However this adoption is typically not driven by market forces, and without market forces it is in the hands of analysts and researchers to identify best practice. While a large body of research focuses on the relationship between IT, work processes, and productivity,⁴ only a few studies show direct evidence of the role of IT in increasing productivity.⁵ Moreover, IT investments are often intangible and disproportionately difficult to measure and link to productivity (see Brynjolfsson and Hitt, 2000, David, 1990). Robert Solow’s oft-cited observation that one “can see the computer age everywhere except in the productivity statistics” is a good summary of this well known “Productivity Paradox.”

Even when detailed data are available, estimates of IT’s impact are usually based on cross-sectional or at best panel-data variation in IT use, where the organizations that use IT innovations may be those that benefit the most from such innovations or differ in ways

³See the 1987, 2003, and 2013 Law Enforcement Management and Administrative Statistics (LEMAS).

⁴See, among others, Acemoglu et al. (2007), Autor et al. (1998), Berman et al. (1994), Black and Lynch (2001), Bloom et al. (2012), Bresnahan et al. (2002), Doms et al. (1997), Stiroh (2002).

⁵See Angrist and Lavy (2002), Goolsbee and Guryan (2006) for schooling, and Garicano and Heaton (2010) for police management. Moreover, Doleac (2017) shows that the introduction of DNA databases reduces crime.

that are unobservable by the econometrician. For example, the adoption of IT might coincide with other new management practices that are unobserved by the researcher (see Bartel et al., 2007). A few papers circumvent these obstacles by focussing on specific applications of IT.⁶ In this paper I focus on a recently popularized IT innovation that is spreading quickly across police departments worldwide (see Section 2), called “predictive policing.” This collects and analyzes data on past criminal events to predict future events and allocate police resources more efficiently. Police patrols are given these predictions to allow them to plan their routes to increase their clearance or arrest rates, and reduce crime (see, for example, Ferguson, 2017).

Regardless of the growing interest and growing investments in predictive policing, several stakeholders have highlighted that very little is known about its effectiveness (Sengupta, 2013). The two main identification issues are the endogeneity of its use and possible displacement effects (criminals have an incentive to defy these predictions, and one way is to simply move away from locations that are predicted to see surges in crime). I address both issues, leveraging off individual offense-level randomization in the availability of predictive policing.

I estimate the empirical relationships between IT use and the productivity of police patrols, measured by the likelihood that individual crimes are cleared by arrest, using micro-level data on the universe of commercial robberies against businesses in Milan (Italy) over a two-and-a-half-year period.⁷ The estimation is based on difference-in-differences, where the first difference builds on the existence of two separate police forces that patrol the streets. The *Polizia* uses predictive policing, while the *Carabinieri* follow traditional patrolling strategies, which tend to be less proactive, and more random.

A peculiar rotating mechanism generates quasi-random assignment of investigations to

⁶See Athey and Stern (2002), McCullough et al. (2016), Hubbard (2003), and Dranove et al. (2014).

⁷Several economic studies have used clearances as a measure of police performance (see, among others, Blanes i Vidal and Kirchmaier, 2017, Garicano and Heaton, 2010, Mas, 2006).

these two separate police forces (see Mastrobuoni, 2019).⁸ The city is divided into three sectors and approximately every 6 hours, when the shifts are changing, the two police forces are assigned to different sectors.

Even though the two forces share similar staffing and equipment and have access to the same information (including the opportunity to interview the victims), this difference would not be sufficient to identify a productivity change. *Polizia* and *Carabinieri* may have a different productivity, irrespective of their use of IT.

I exploit three alternative second differences to control for separate productivity levels. The first “second” difference is based on the nature of predictive policing: the analysis of past criminal events. Any difference in clearance rates for the very first robbery of a sequence of robberies perpetrated by a criminal group would be evidence of differential productivity that is not based on IT. The second “second” difference is based on a procedural delay in producing the crime predictions. The *Polizia* requires time to collect and analyze the data. To improve the victims’ recollections, officers wait about one day before interviewing the victims. This means that predictions are not updated on the same day a robbery has taken place, generating a discontinuity in the availability of updated predictions.

While there is no evidence of a productivity differential between *Polizia* and *Carabinieri* for the very first robbery of a sequence, subsequent robberies that fall in the *Polizia* sector as opposed to the *Carabinieri* sector are 8 percentage points more likely to be solved (the overall clearance rate is 14 percent across the two police forces). Similar productivity differences emerge between robberies that happen before and after the predictions are updated, as long as the crime happened in an area surveilled by the *Polizia*. The results are robust to narrowing the sample to robberies that happen around the time the software is updated. Again, no differences emerge for the *Carabinieri* and between

⁸Mastrobuoni (2019) finds that the disruption in police patrolling, driven by shift changes of *Polizia* and *Carabinieri* patrols, reduces clearance rates. Excluding robberies that happen within 15 minutes of the shift changes strengthens the results shown in Section 5. These results are available upon request.

Polizia and *Carabinieri* before the data update takes place.

I also provide evidence on the mechanisms that drive the productivity effect. Individual criminal behavior shows clear signs of predictability.⁹ Over time criminal groups tend to select the same business types, around the same time of the day, and in the same city neighborhood, especially if previous robberies have been lucrative. Moreover, robbers tend to be very criminally active, which implies that at any given time the *Polizia* focuses on a small number of sequences.¹⁰ I show that the instructions distributed to the police patrols highlight these patterns and that conditioning on past behavior can improve the exact predictions of where and when the next robbery is going to happen by more than 10 percentage points, which is consistent with the size of the productivity effects.

The large productivity boost in terms of clearances is expected to translate into more incapacitation and lower crime rates. Evidence based on auxiliary monthly municipality-level bank robbery rates provided not by the *Polizia* but by the Italian Banking Association shows that around the beginning of 2008, when predictive policing was first introduced, Milan robbery rates compared to rates in *any* other major Italian municipality experienced a sharp and abrupt reversal of a previously increasing trend.

I conclude the analysis with a conservative cost benefit analysis where, even discounting deterrence effects, predictive policing appears to be very cost-effective.

As previously mentioned, this paper contributes to the literature on IT and productivity. A few studies have examined micro-level empirical relationships between IT use and productivity. Athey and Stern (2002) use a difference-in-differences setup to evaluate the effect of enhanced 911 emergency response systems that link caller identification to a location database on health outcomes. IT is shown to generate significant improvements

⁹Predictability does not necessarily mean that criminals are not choosing an optimal criminal strategy. Becoming more unpredictable seems costly: apart from the potential cost of travelling more, the data shows that targeting different types of businesses is associated with a lower haul.

¹⁰Sixty percent of matched offenders commit a new robbery within one week, 77 percent within two weeks, and 85 percent within one month. As a result, each month the average number of unique groups that are active and whose actions need to be predicted is around 13.

in the health status of cardiac patients. McCullough et al. (2016) use the same setup to show that emergency care patients with complex health conditions see a reduction in mortality when health IT is adopted. Hubbard (2003) uses a conditional independence assumption to test whether trucks that use on board computers are more productive. On-board computers (GPS, etc.) are shown to significantly increase the ability to predict the availability of trucks and therefore to improve their capacity utilization. This paper contributes to the growing literature on the mechanisms through which policing reduces crime, which I discuss in more detail in the next section. It also has implications about data collection for law enforcement agencies.

2 Predictive Policing

According to a 2012 survey disseminated to more than 500 US police agencies, 38 percent of agencies were using predictive policing, and an additional 22 percent were planning to use it by 2017 (Police Executive Research Forum, 2014). There are no comprehensive worldwide statistics, but according to different news media, police departments in several European countries (among others, Kent and Manchester Police in the United Kingdom, Baden-Wuerttemberg in Germany, Amsterdam in the Netherlands, Zurich in Switzerland, the Central Criminal Intelligence Service in France) and India use or plan to use predictive policing tools (Sharma, 2017). Another indicator of the growing interest in predictive policing is media attention. The number of news articles featuring the market leaders of predictive policing software (*Predpol*, *Hunchlab*, and *Precops*) is growing exponentially over time. A Nexis.com search shows that the time-series of citations from 2009 to 2018 was 2, 4, 6, 24, 54, 66, 91, 62, 68, and 126. A linear extrapolation based on the first 2 months of 2019 leads to 252 articles for 2019.¹¹ Not surprisingly, almost all news articles

¹¹The same trends emerge searching for the “predictive policing,” while it does not emerge if one searches for just “policing.”

are concentrated in the United States (59 percent), in Europe (26 percent) and in Asia (8 percent).

The use of IT to improve policing started about 25 years ago. The precursor of predictive policing is CompStat, a data gathering and accountability process developed by the New York Police Department in 1994 and since then adopted by most US police departments (Weisburd et al., 2003).¹² Due to a striking empirical regularity the data are often used to map patterns of crime and to plan police patrolling: a few intersections or city blocks, called crime ‘hot spots,’ often produce the majority of crime incidents (see, among others, Sherman et al., 1989, Weisburd and Eck, 2004, Weisburd and Green, 1995). These hot spots have prompted police departments to target police patrols in geographic areas (e.g., blocks or specific addresses) that show high levels of criminal activity.

Focussed policing has evolved from using data to simply identify high crime areas into predictive models that are based on the frequency rates of local crime.¹³ The criminology theory behind hot spots is called repeat and near repeat victimization, which highlights the importance of identifying those people, properties and places that are at a disproportionate risk of victimization.

The most advanced software predicts the most likely type, location, and time of future crimes, downweighting observations that are farther away in space and time (see, for example, Mohler et al., 2011). Several US Police Departments are currently using such software, often funded by National Institute of Justice grants (Pearsall, 2010).¹⁴

¹²Garicano and Heaton (2010) study the relationship between information technology, productivity, and the organization of police departments. Such investments are linked to improved productivity when they are complemented with programs like CompStat, which was developed by the New York’s Police Commissioner William Bratton. In 2002 Bratton moved to California and is credited with developing the first predictive policing model within the Los Angeles Police Department (Perry et al., 2013).

¹³A few studies have evaluated hot-spots policing strategies, and most criminologists believe that focused policing works (Braga, 2001, Cohen and Ludwig, 2003, Mohler et al., 2015, Sherman and Weisburd, 1995, Weisburd and Green, 1995). Levitt (2004) is more skeptical about the decline in crime that occurred during the 1990s that can be attributed to CompStat. While there is currently little evidence that hot-spots policing simply displaces crime to nearby locations, one potential limitation of these studies is in identifying the areas where crime might spill over. These areas are not necessarily contiguous to the area that is being targeted.

¹⁴The U.S. Department of Homeland Security note Space and Naval Warfare Systems Center Atlantic

The first evaluation happened in 2011. The Santa Cruz Police Department ran a city-wide six-month experiment of the software “Predpol,” later named one of Time Magazine’s 50 best innovations (Grossman et al., 2011). Like many police departments around the world, the Santa Cruz Police Department had a declining budget and shrinking police force. After an unprecedented crime wave at the beginning of 2011, the Department decided to work with researchers at UCLA to test a new method of modelling crime using data on burglaries, car thefts, and thefts from vehicles (Economist, 2010, Mohler et al., 2011). The experiment seemed to reduce crime, though the initial absence of a control group, the lack of an independent evaluation, and the possibility that crime was merely displaced make it difficult to draw any definite conclusions.¹⁵

The main drawback in most previous work is the lack of a proper comparison group. Criminals might move from treated to control regions, contaminating the experimental design. Reducing contamination by choosing larger regions would introduce additional heterogeneity between treated and control areas. Exploiting pure time-series variation would also be impractical. A spike in crime followed by the use of predictive policing might, just naturally, lead to reversion to the mean that is unrelated to the newly adopted technology. Moreover, part of the effect of predictive policing might be due to an incapacitation effect, which is dynamic in nature and thus hard to separate over time.¹⁶

(2013) contains a list of software providers, a list of cities that use such software, as well as an overview of how these different systems work.

¹⁵More than 50 police departments around the US now use PredPol to organize their patrolling. Predictive policing is also being evaluated in the UK where, in the single ward of the Greater Manchester area studied, burglary decreased by 26 percent versus 9 percent city-wide, which led to follow-up studies in Birmingham. A later study, Mohler et al. (2015), uses a randomized experiment to evaluate PredPol, where treated areas are allowed to change over time and across space, and patrolling in untreated areas is organized by conventional crime analysts. The study shows that even when compared to police analysts, PredPol is better at predicting crimes. There is also some evidence, though it is not experimental, that extra patrol time reduces crime more in PredPol “treated” missions compared to missions designed by crime analysts.

¹⁶The two mechanisms are often hard to separate when only aggregate data are available (Owens, 2014). See Durlauf et al. (2010) for additional issues that might arise from estimating aggregated crime regressions. Mastrobuoni (2019) uses the same crime level data used in this paper, in particular the variation in police presence that is driven by shift changes, to show that an increase in police patrolling leads to higher clearance rates. A related paper, Blanes i Vidal and Kirchmaier (2017), shows that exogenously driven reductions in response time increase the likelihood that the police clears the crime.

Later, in Section 4.3 I will describe how the Milan setup allows me to circumvent these issues. Before that let me describe the Milan predictive policing.

2.1 The Milan Police Predictive Policing

The Milan Police (*Polizia*) Department was one of the first police departments in the world to use predictive policing. The software, called KeyCrime, became fully operational in January 2008.¹⁷ KeyCrime collects and analyzes micro-level data on all commercial robberies that take place in the municipality of Milan (*Comune di Milano*).¹⁸

Mario Venturi, a former senior police officer (*Assistente Capo*), who had many years of active duty experience combatting commercial robberies, had noticed the high prevalence of repeat offenders with persistent *modi operandi*. Realizing that past data could help to predict future robberies, in his spare time he developed and later copyrighted KeyCrime. He later convinced the Milan police chief to use the software, which was made available for free to the Milan *Polizia*. Venturi, former head of *Polizia*'s predictive policing group and current CEO of KeyCrime, is currently planning to expand its use to both other Italian cities and internationally. This type of bottom-up innovation, driven by motivated public sector employees without managerial tasks, lowers the risk that other innovations may have taken place around the same time for the same police force.

Milan's KeyCrime is focused on commercial robberies, as robbers are believed to be mainly professional and hardened criminals who sooner or later are going to perpetrate

Di Tella and Schargrodsky (2004), Draca et al. (2011) and Klick and Tabarrok (2005) exploit exogenous variation in the deployment of "high deterrence" police officers following terrorist attacks, and find strong evidence in favor of a deterrent effect of police stationed at a circumscribed area. Blanes i Vidal and Mastrobuoni (2018), instead, find limited evidence of deterrent effects in the case of random mobile patrolling.

¹⁷For comparison, PredPol became operational in Los Angeles only in 2011.

¹⁸Appendix Figure A2) shows a map of the municipality of Milan with the distribution of commercial robberies. These are crimes of violence against businesses motivated by theft, and are quite prevalent in Italy. Bank robberies, which in Milan comprise between 10 and 15 percent of all commercial robberies, were until recently more prevalent in Italy than in the rest of Europe altogether (see Mastrobuoni and Rivers, 2019).

a crime. According to private conversations with Mario Venturi, the software’s aim is to improve the officers’ role as “apprehension agents” (see Nagin et al., 2015).¹⁹ The software is used to input and analyze large sets of individual characteristics of robbers and individual criminal strategies (*modus operandi*) collected from closed-circuit security cameras and victim reports to: i) identify robberies that share at least one offender or one vehicle (a “sequence” of a criminal group); and ii) predict when and where the offenders are going to strike next. Thus, rather than generating aggregate predictions, like most predictive policing software does, it generates individual predictions. It is important to highlight that the linkages across robberies are constructed regardless of whether an arrest is made.

KeyCrime equips police officers and prosecutors with well-summarized analytics that are used to help police patrols decide where to drive and what to look for.²⁰ As such, in our set-up ‘information technology’ is defined as “the application of computers to store, study, retrieve, transmit, and manipulate data” (Daintith, 2009). Since officers and machines are constantly interacting with each other, and algorithms augment human abilities rather than replace them, I argue that KeyCrime is the most interesting policy parameter.

This is how KeyCrime collects information: after a robbery takes place the *Polizia*’s predictive policing team, made up of four officers, collects the official *Polizia* or *Carabinieri* report (these reports are centralized and both forces can observe the universe of committed crimes, with a detailed description of the events), and later interviews victims and collects surveillance camera footage. The *Polizia* collects this information for the universe of reported commercial robberies that take place in Milan, even for those robberies that are

¹⁹For example, the *Polizia* would sometimes use undercover officers to “wait” for the perpetrators in front of predicted victims, using visible officers to displace perpetrators from alternative targets.

²⁰An additional advantage is that the individual data can assist the prosecutors once the perpetrators have been arrested and are put on trial. Thus not only clearance rates, defined as the likelihood of solving a specific crime before the offender’s next crime, are likely to respond to this IT innovation; conviction rates could improve as well. Unfortunately, the identification strategy used to estimate the causal effect of predictive policing on clearance rates cannot be extended to conviction rates. The reason is that all police forces share all information collected with the prosecutors, even when the competing police force, the *Carabinieri*, which later represents our control group, made the arrest.

investigated by the *Carabinieri*.²¹

The data collected about the crime (time, date, location, type of business, type of crime, etc.), the observed perpetrators (perceived age, height, body structure, skin, hair, eye color, clothing, etc.), the observed weapons (type, maker, model, color, etc.), and the observed mode of transportation used by the perpetrators (type, maker, model, license-plate, etc.) are then used to generate predictions. About 80 percent of businesses have closed-circuit security cameras (CCTV), and these videos are also uploaded to KeyCrime.²² Between January 2008 and June 2011 the *Polizia* recorded more than 2000 robberies, at a rate of 1.5 robberies per day.

The core of the information collected is a telephone interview with the victims the day after the robbery. The predictive policing team rather than the police patrols make the call. The purpose of the delay is to give the victim time to recover from the immediate trauma and improve their recollections.²³ After the data are entered into KeyCrime, the software aids the operators' job of matching robbers or groups of robbers over time. The software allows the operator to use one screen to compare the characteristics of different robberies, including their photographic evidence (Appendix Figure A4).²⁴

Once links are established (later I discuss the possibility that the links might be misclassified), the data are used to highlight and to predict criminal strategies. The predictions are based on a mix of statistical and psychological/criminological models.

The potential future targets (Appendix Figure A5) are then communicated to police patrols, together with the likely day of the week and time of day of the future offense.

²¹Given the monetary and non-monetary incentives to report these crimes (many businesses are insured and understand that future patrolling strategies may depend on their reporting behavior), reporting rates among commercial businesses are believed to be close to 100 percent.

²²According to the *Polizia* all banks, postal offices, pharmacies, and jewelers have at least one CCTV camera.

²³Later I exploit such delays to set up the second difference-in-differences strategy.

²⁴While I do not have access to the proprietary algorithm that predicts criminal behavior, I have been told that the current pattern recognition software is not capable of automatizing the matching of photographic evidence. Moreover, when the evidence is missing the operator can still use peculiar and rare physical appearances to establish these links.

These reports describe the offenders and their typical *modus operandi*, including the means of transportation, the typical time of the day and target type chosen (Appendix Figure A3). On the second page of the report a map indicates the neighborhoods where the criminals are likely to strike, while the final page collects all the photographic evidence.

Up until 2011, the last year covered by KeyCrime data, the Milan Police Department was the only one of the 103 provincial *Polizia* departments to use individual predictions to fight commercial robberies. Not one of the 103 *Carabinieri* provincial departments was using predictive software, at least until 2010 when prosecutors in Milan forced the Milan *Polizia* to share reports with the Milan *Carabinieri*.²⁵ Not developing a predictive system is the norm, but the sharing agreement might have pushed the *Carabinieri* to develop similar policing strategies.²⁶ After 2011 other Italian cities have started using predictive policing software, though none matches robbers over time to build individual predictions.

3 Evidence Based on Aggregate Crime Data

As already mentioned, a traditional approach that tracks the evolution of crime rates before and after the use of predictive policing is prone to reverse causality and reversion to the mean. Standing on its own, evidence of a reduction in crime would not be a powerful test.

However, at any given point in time KeyCrime covers two-thirds of the city of Milan (see Section 4.3), and if IT improves the productivity of the police it should also lead to a tangible reduction in crime. Therefore, lack of a reduction in crime would be evidence that KeyCrime is not improving the productivity of the police (or that displacement within the city is large).

²⁵The *Polizia* started sharing their predictions in January 2010 and by the end of the year the *Carabinieri* had access to 33 classified reports, which were at least as detailed as those shown in Appendix Figure A3 and that were regularly distributed to the *Polizia* patrols.

²⁶Later I use the sharing agreement to test whether it improves the productivity of the *Carabinieri*.

I start the analysis using the monthly city-level bank robbery data, which have the advantage of being independently collected by the Italian Banking Association.




Figure 1: Bank Robbery Rates in Milan

Notes: The solid line represents the Milan monthly bank robbery rates per 100,000 inhabitants. The dashed line smoothes the solid line using a local linear regression.




Figure 2: Bank Robbery Rates in Milan against its Synthetic LASSO

Notes: The solid thick line represents the difference in monthly bank robbery rates per 100,000 inhabitants between Milan and its synthetic city computed using a LASSO approach (with optimal λ using cross-validation). The thin lines represent the differences with respect to their synthetic lasso controls for the other cities with at least 200,000 inhabitants.

The monthly series of bank robberies (which represent about 10 percent of all commercial robberies) covers all major cities (*capoluoghi di provincia*) from 2004 to 2015. Another advantage of city-level aggregation is that crime displacement is internalized. Figure 1 shows the time series of bank robbery rates per 100,000 inhabitants (based on 2006 population estimates) in the city of Milan, $R_{1,t}$. There is a clear upward trend in bank robberies that reverses around the time KeyCrime becomes operational (early 2008). The changes are large. The subsequent evaluation based on micro-data covers the period 2008-2011, and during this time robbery rates fell from about 1.4 to about 0.5.

Some of the identification concerns can be addressed by comparing these changes to what happens in other cities that did not introduce any predictive policing software. It would be hard to ascribe the trend reversal to predictive policing if similar changes were

also visible in other Italian cities. Rather than comparing Milan to single cities, I use a synthetic control approach based on Italian cities that are large enough to compare well to Milan and also have a sufficient number of bank robberies. Using a threshold of 200,000 inhabitants—Milan is the second most populous city after Rome, with about 1.3 million inhabitants—I am left with 13 cities.²⁷

In Abadie et al. (2010)’s synthetic control approach, the outcomes of comparison cities are weighted to mimic the robbery rates in Milan before the 2008 introduction of predictive policing (see also Abadie and Gardeazabal, 2003), and the weights ω_i for comparison city i are constrained to be non-negative and to sum up to one:

$$\hat{R}_{1,t} = \sum_{i \neq 1} \omega_i R_{i,t} \quad .$$

One important condition for this method to work is that at least one city shows pre-trends in robberies that are increasing, or, stated more precisely, that Milan’s pre-intervention robbery rates are inside the convex hull of the control cities’ robbery rates. This is not the case.²⁸

This might be related to the fact that Milan is not the typical Italian city. For one, the province of Milan has the highest GDP per capita, comparable to those of Paris and London and is growing fast.²⁹, and this might attract a growing number of criminals. A synthetic control approach would have to allow for differential pre-trends

²⁷Naples and Turin have almost one million inhabitants and the next largest cities are Palermo and Genova with about 650,000 inhabitants. There are 9 more cities with a population between 200,000 and 400,000 inhabitants.

²⁸The two panels in Appendix Figure A6 show that this is not the case. In the left panel, to highlight Milan’s trend reversion, robbery rate differences between cities and their synthetic counterparts, are computed net of their December 2007 difference $R_{1,12/2007} - \sum_{i \neq 1} \omega_i R_{i,12/2007}$. I match average robbery rates for each pre-2008 semester. Using semesters, rather than years, quarters, or individual months minimizes the mean squared prediction error. Adding population size, latitude, and longitude would also slightly improve the prediction but does not change the results. For statistical inference “placebo” differences for each city in the donor pool are shown in light gray. The right panel shows pairwise differences for the 9 major cities.

²⁹See the *Conti Economici Territoriali* of the Italian Statistical Office.

to capture Milan’s anomaly.³⁰ Moreover, crime displacement across cities may generate a negative correlation in robbery rates across cities. To allow for both, negative weights and differential trends, I construct the synthetic Milan using a linear regression model

$$\hat{R}_{1,t} = \mu_0 + t\mu_1 + \sum_{i \neq 1} \omega_i R_{i,t}.^{31}$$

Given that the donor pool is large I use the Lasso shrinkage estimator (least absolute shrinkage and selection operator) to set weights with little predictive power equal to zero (see Tibshirani, 1996). Figure 2 shows that Lasso’s synthetic control approach adjusts for differential trends and improves the pre-treatment fit.³² The changes in robbery rates become larger, as they now include the pre-trend, and are close to -1.5.³³

Yearly crime data by the Italian Statistical Office that cover the period from 1983 to 2016 allow me to test for geographical displacement from the municipality to the rest of the province.³⁴ The size of Milan’s province is about 9 times as large as the municipality ($181km^2$ against $1,575km^2$). The figures show that Milan province experiences a similar trend-reversion around 2008 (see Appendix Figure A8). The increasing trend in bank robberies and commercial starts in 2004.³⁵ In relative terms these provincial-level changes are less pronounced than for the municipality of Milan, which is consistent with treatment being concentrated in the municipality of Milan, but is not consistent with large displacement effects.³⁶

³⁰See Doudchenko and Imbens (2016) for an overview of how to relax the conditions of the Synthetic Control approach.

³¹Dziuda and Mastrobuoni (2009) apply this method to generate counterfactual product-specific inflation rates in Euro countries.

³²I use an optimal shrinkage parameter using cross-validation methods.

³³Using in addition 17 donor cities with a population between 100,000 and 200,000 adds more noise to the randomization test but delivers similar results (see Appendix Figure A7).

³⁴See Appendix Section A.1 for a description of the data.

³⁵Over the long period the pre-trends are less severe and so I use Abadie et al. (2010)’s approach. The vertical grey lines indicate the period covered by the monthly robbery data of the Italian Banking Association.

³⁶The commercial robbery rate in Milan was 38 per 100,000 inhabitants in 2007, and dropped by about 10 within a couple of years. The bank robbery rate was 11 in 2007 and dropped by about 5.

Using province-level data allows me to consider other crime types. To challenge the hypothesis that KeyCrime is behind the trend reversion that is observable around 2008, I test whether other crimes follow the same patterns as commercial robberies. Alternatively, an unexpected increase in other crimes around the introduction of predictive policing could indicate displacement from robberies to other crime types.

Before the introduction of predictive policing Milan was not just experiencing differential trends in bank robberies and commercial robberies (see Appendix Figure A9). Total crime (net of robberies), thefts, violent crimes and shoplifting were all on the rise before the introduction of KeyCrime but, unlike commercial robberies, never reversed. This seems to be consistent with the idea that Milan is attracting a growing number of criminals. In the presence of displacement one would expect an even steeper trend in other crimes once KeyCrime is introduced. This does not seem to be taking place.

Summing up, KeyCrime was developed at a time where crime was on the rise. The introduction of KeyCrime coincided with a reduction in commercial robberies but not other crimes, and there is no evidence of displacement across space and crime types. Next, I use micro-data on individual robberies to better understand whether the results can be attributed to predictive policing.

4 Micro-level Evidence

4.1 The Data

The dataset I use was collected through KeyCrime between January 2008 and June 2011, with great deal of detailed information on the robberies (location, time, haul, arrest, number of offenders, weapons, type of business, etc.).³⁷

³⁷The data were shared before I decided to evaluate KeyCrime. Moreover, since the predictive policing software, which started to be introduced in 2008, is also used to collect the data, no data on clearances by police force are available before 2008.

Table 1: Summary statistics

Variable	Mean	Std. Dev.	Min	Max	Mean	Std. Dev.	Min	Max
	Full Sample (2008-2011)				Restricted Sample (2008-2009)			
Cleared robbery (0/1)	0.14	0.35	0	1	0.13	0.34	0	1
Number of the sequence	5.10	6.88	1	84	4.20	5.53	1	84
Police/ <i>Polizia</i> Intervention (0/1)	0.73	0.44	0	1	0.74	0.44	0	1
Days between subsequent	16.80	46.43	0	555	14.48	43.47	0	555
Subsequent robberies (0/1)	0.58	0.49	0	1	0.54	0.50	0	1
North-Western area (0/1)	0.35	0.48	0	1	0.38	0.48	0	1
North-Eastern area (0/1)	0.22	0.41	0	1	0.19	0.39	0	1
Year	2009.24	1.02	2008	2011	2008.47	0.50	2008	2009
Month	5.88	3.71	1	12	6.20	3.75	1	12
Day of the month	15.60	8.86	1	31	15.74	8.97	1	31
Day of the week	3.24	1.83	0	6	3.19	1.82	0	6
Daylight (0/1)	0.59	0.49	0	1	0.57	0.49	0	1
Average age	26.57	12.47	0	68	26.14	13.10	0	68
Amount stolen in euros ($\times 1000$)	2.86	11.18	0	206	2.11	7.90	0	100
Firearm (0/1)	0.23	0.42	0	1	0.21	0.41	0	1
At least one knife, but no firearm (0/1)	0.09	0.29	0	1	0.09	0.28	0	1
Some Italian involved (0/1)	0.79	0.41	0	1	0.77	0.42	0	1
Different nationalities (0/1)	0.14	0.35	0	1	0.12	0.32	0	1
Number of robbers	1.57	0.72	1	7	1.51	0.68	1	5
Obs	2167				1255			

Notes: The observations are the universe of commercial robberies perpetrated between January 2008 and June 2011. The restricted sample is based on the years when the *Polizia* was not sharing part of the patrolling reports with the *Carabinieri*.

The original data contain information on the physical characteristics of the offenders, and photographs, but the police did not share these data. The summary statistics of the available variables are shown in Table 1, both for the full sample and for the sample which restricts the data to the first two years (before the *Polizia* started sharing their predictions with the *Carabinieri*). Each observation represents a separate robbery. Over the period 2008-2011 there were more than 2000 separate robberies in Milan. According to the Milan police 70 percent of these robberies show some link with other robberies, meaning that at least one robber or one vehicle were seen in two different instances. The variable ‘Number of the sequence’ $\tau = 1, \dots, T_i$ counts the number of crimes that have been linked to a serial group of offenders i .³⁸ The criminal group with the largest number of offences organized 84 robberies.

The *Polizia* defines a robbery as cleared if an arrest is made before the same group of robbers re-offends. I do not have complete information on the exact date of arrest, but according to the *Polizia*, considerably more than half of all arrests happen *in flagrante*, meaning when a robbery is taking place, or when the robbers are fleeing. Yet, as shown in Blanes i Vidal and Kirchmaier (2017), even when the arrest does not happen on the spot, a quicker response time increases the likelihood that the crime is eventually cleared. A sequence i is defined as solved when all observed robbers have been arrested.

More than half of the robberies (1,221 robberies out of 2,164) belong to a sequence where at least one arrest has been made. Of these, 981 (80 percent) belong to a sequence that has presumably been fully cleared.³⁹

Table 1 shows that the individual clearance rate of robberies is 14.9 percent, which leads to 45 percent of the sequences being fully cleared by June 30, 2011, the day the data were extracted. The *Polizia* variable indicates whether the Milan Police Department

³⁸For the serial crimes that started in 2007 and continued in 2008 I have the number of robberies performed in 2007, which I added to the “Number of the sequence.”

³⁹Though, in principle it would still be possible for the sequence to proceed if new perpetrators were using the same vehicles used by those already arrested.

handled that particular robbery and in the following section is a description of how this assignment of investigations to the *Polizia* and the *Carabinieri* works. Since the police forces divide the city into three sectors and the *Polizia* is responsible for two of these at any given time, the fraction of robberies handled by the police is slightly larger than expected at 73 percent (as opposed to the expected two thirds based on geographical coverage). The main deviation from the expected number of robberies happens in the first year, where the *Carabinieri* appear to investigate 26 percent of the first robberies but only 19 percent of the subsequent ones (see Appendix Table A1). This might be driven by the fact that the *Polizia* is collecting the data and could have missed some commercial robberies, especially in the first year. The estimated effects will be similar when focussing on later years, where the *Carabinieri* cover close to one third of the investigations.

According to the victims' reports, the robbers appear to be on average 26 years old. The average haul is around €2,000, or US\$2,200. One quarter of robberies involve firearms, and in about 10 percent a knife is used. Robberies are mainly an "Italian job," meaning that in 80 percent of cases at least one Italian seems to be involved. Only in 12 percent of cases do the robbers seem to be of different nationalities. The average number of robbers involved in each robbery is about 1.5.

4.2 Predictability

Several mechanisms can rationalize the predictability of robbers, for example, superior information about targets, learning through experience, time constraints (legitimate work, darkness, etc.), or liquidity constraints. Robbers might thus choose to operate in certain parts of the city, concentrate on certain types of businesses, and even at certain times of the day, of the week, or operate at regular intervals, for completely rational reasons.

Predictive policing exploits such predictability to improve police patrolling (delivering a list of potential targets). The two most important predictions are *time* and *location* of

a robbery. Time can be broken down into day and time of day. Location can be broken down into businesses located in certain neighborhoods.




Figure 3: Persistence in *Modus Operandi*

Notes: Each bar represents the conditional (dark gray) or unconditional (light gray) likelihood of selecting a given characteristic. Conditioning refers to the characteristic being equal to the modal characteristic that has been selected before that robbery. The city has been divided into 22 equally sized squares (x, y) (see Appendix Figure A2). Only characteristics present in at least 1 percent of robberies are displayed.

Therefore, using information that was available to the police before a given robbery, I measure predictability or persistence based on six variables: week of month, day of the week, time of day (in particular I use the daily four police shifts), targeted businesses ('Targets'), and neighborhoods ('Space'). In addition, I test whether there is persistence in the mode of transportation, as police patrols are often given photographic evidence on motorcycles, helmets, mopeds, etc.. The six panels in Figure 3 correspond to the dimensions.

Each variable j can take K^j mutually exclusive values or features. For example, the day of the week variable can take values 1 to 7, while targets can be classified into banks, pharmacies, jewelers, etc. For the locations I divide the map of the municipality of Milan in 22 3km-wide squares (see Appendix Figure A2).⁴⁰

A criminal group i shows persistence in targets when the chosen target type is identical to their most frequently chosen target type before time t . Persistence for variable j is equal to one when $k_t^j = \text{mode}(k_0^j, \dots, k_{t-1}^j)$, and zero otherwise.⁴¹

Average persistence across criminal groups and over time measures the probability that the chosen feature k^j in period t is equal to the modal feature up until period $t - 1$, or $\Pr(k_t^j = \kappa | \text{mode}_{t-1}^j = \kappa)$. Rather than the mode, one can also condition on the features of specific robberies, for example the first $\Pr(k_t^j = \kappa | k_1^j = \kappa)$ or the last robbery $\Pr(k_t^j = \kappa | k_{t-1}^j = \kappa)$.

If there was no persistence, the conditional probabilities would be equal to the unconditional ones $\Pr(k^j = \kappa)$. Moreover, if there was little gain from gathering new information the very first robbery would already be predictive of future features. Furthermore, if robbers' choices were evolving based on those taken immediately before, as in a Markov chain, features at $t - 1$ would be a better predictor compared to the modal ones. Table 2 computes the average persistence across features for these different types of conditioning.

Figure 3 shows the conditional and the marginal probabilities of each feature as long as the latter are larger than 1 percent. The “Targets” and “Space” panels show that robbers' location choice is predictable based on past behavior. The marginal distributions are an order of magnitude smaller than the likelihood that a group of robbers targets the type of business they have been targeting most often in the past. Bank robberies, for example, represent less than 15 percent of robberies, but when conditioning on the mode

⁴⁰The size of the neighborhoods is arbitrary but persistence is visible even for smaller areas. Another way to show persistence in the choice of the location of a robbery is to plot these for the most prolific groups of robbers. Appendix Figure A10 shows the distribution of locations (by latitude and longitude) for groups of robbers with a total of at least 15 robberies.

⁴¹Whenever there is more than one mode, I randomly select one.

Table 2: Average Persistence

Conditioning feature	none	k_1	k_{t-1}	Mode
Targets (18)	0.055	0.653	0.686	0.703
Space (22)	0.045	0.303	0.282	0.342
Day of the week (7)	0.142	0.153	0.105	0.174
Hour of the day (24)	0.041	0.095	0.119	0.232
Shift (4)	0.250	0.457	0.491	0.501
Week of the month (4)	0.250	0.238	0.399	0.303
Transportation (6)	0.167	0.412	0.452	0.450
Weeks in between (70)	0.014	0.366	0.380	0.426

Notes: Each entry represents the average conditional probability that the corresponding characteristic (target, mode of transportation, etc.) is equal to the characteristic in the first robbery (k_1), in the previous robbery (k_{t-1}) or the modal characteristic. The number in parenthesis shows the number of different values each characteristic can take. For the “Weeks in between” one robbery and the next I need to use k_2 rather than k_1 . To mimic the functioning of KeyCrime, I exclude 121 out of 1,259 subsequent robberies that happen the same day as the previous one.

the likelihood goes up to almost 80 percent. Against most business types the conditional probability is at least 50 percent, meaning that by only looking at the past modal business type there is a 50 percent chance of correctly predicting the one that is going to be targeted next. Table 2 shows that overall the average conditional probability is 70 percent (against a marginal probability of 5.5 percent), and that even the very first target is predictive of future targets.

But predicting the business type is not enough to pin down the exact business. The *Polizia* need to at least focus on a small number of neighborhoods. For certain neighborhoods there is a 50 percent chance of correctly choosing which area to patrol, which combined with the type of business can be extremely informative compared to random patrolling. Overall the average conditional probability is 34 percent, against an average marginal probability of 7.7 percent. Again, even the very first location predicts future ones. Combining information about targets and space the chances for patrols to be at the right place is almost 25 percent ($0.70 \times 0.34 = 0.23$) as opposed to 2.4 percent with

random patrolling.⁴²

It is equally important to predict when the robbery is going to take place. 58 percent of repeat robberies happen within one week, 74 percent within 2 weeks, and 89 percent within one month. This implies that the *Polizia* can focus on trying to predict the exact day and exact time within the week.⁴³ Though there is some evidence of end of the month effects, possibly due to pay-cycles.

In terms of predicting the day of the week, there is some difference between the conditional and the marginal distribution for Sunday and Monday, and to a lesser degree Friday and Saturday, possibly because of working schedules.

The other important dimension is time of day. Dividing the day into patrol shifts, the overall average conditional probability is 50 percent, while the marginal probability is 25 percent. The night and morning shifts are those for which conditioning matters the most.⁴⁴ Table 2 shows that to predict the exact hour of the week, the modal hour chosen in the past is a much better predictor than the first or last hour chosen (about 20 percent against 10 percent likelihood). If police patrols choose to patrol a predicted target (in a specific neighborhood) in a given hour for several days they have an almost 6 percent chance of being in the right place at the right time. If they cover the entire shift the probability goes up to 12.5 percent. With random patrolling the likelihood is only 1 in 1000 or 6 in 1000, respectively. Later we will see that the difference between the conditional and the marginal probability is broadly consistent with the productivity difference attributable to KeyCrime.

Finally, robbers might be spotted while driving a specific vehicle. Panel 2 in Figure 3 shows that motorcycles and bicycles are the most predictable modes of transportation,

⁴²For simplicity, I assume independence between the different variables j . Estimating joint probabilities would improve the forecast at the expense of dimensionality problems.

⁴³Figure 3 changes very little when focussing on robberies that happen within two weeks.

⁴⁴A very similar pattern emerges when one classifies the time at which a robbery is committed into 60-minute periods (the length of the period does not matter). The most predictable robbers are those who target business when they either open (9am) or close (1pm for lunch or 6pm and 7pm). The results are available upon request.

though the most common one is to approach the target on foot. Conditioning on the first or the last robbery is equally helpful.

This section has shown that crimes are predictable. Characteristics of the very first robbery of a sequence are already highly predictive of all the following ones. This implies that KeyCrime can become effective as soon as data on the very first robbery has been collected. In the next section I show how to identify whether KeyCrime also improves clearance rates.

4.3 Identification

The *Polizia* uses KeyCrime to arrest repeat offenders, thus preventing subsequent crimes. Therefore, we can evaluate KeyCrime based on its impact on the likelihood that a case is cleared and perpetrators are arrested.⁴⁵

First Difference: Two Police Forces

The identification of the effect of KeyCrime on clearances rests on difference-in-differences strategies. The main feature that I exploit is that for historical reasons all major Italian cities have two police forces:⁴⁶ the *Carabinieri* is a military police force under the Italian Ministry of Defense and the *Polizia di Stato* is a civilian police force under the Ministry of the Interior.⁴⁷ As mentioned in Section 2.1, up until the end of 2009 only the *Polizia* had access to KeyCrime.⁴⁸ In all major cities the two police forces operate side by side,

⁴⁵Arrest of repeat offenders generate incapacitation effects. Such diffused benefits would spread over time and space, making it hard to infer from simple pre-post differences in crime rates the effect of predictive policing, or as a matter of fact any focused policing (see Cook, 1979, Nagin et al., 2015).

⁴⁶See Mastrobuoni (2019) for a discussion about the two forces.

⁴⁷The only difference between the two forces is that the *Polizia* operates exclusively in metropolitan areas, while the *Carabinieri* operate on the entire Italian territory. This difference does not influence this analysis as I am going to compare forces that operate within the boundaries of the municipality of Milan. While the *Carabinieri* might have an advantage when investigating criminal groups that operate both inside and outside of city, according to the *Polizia* the mobility of criminals in and out of the city is limited.

⁴⁸Any additional effort exerted by the *Carabinieri* in response to an increased productivity of the *Polizia* with respect to repeat offenders would bias the effects of predictive policing towards zero. Differences with respect to the update of the software would not be subject to this bias.

without communicating with each other.

Without random assignment to investigations, increased productivity (a higher clearance rate) might be driven by patrols cherry-picking the more predictable and potentially poorly organized crimes, overstating the effectiveness of predictive policing. This is avoided by the fact that cities are divided into three different areas (two falling under the responsibility of the *Polizia* and the third under the *Carabinieri*). Each force is solely responsible for keeping law and order in the assigned area, and they rotate approximately every 6 hours, counterclockwise (at 12am, 7am, 1pm, and 7pm).⁴⁹ Given that there are two forces, three areas, and four 6-hour shifts within a given day, patrols belonging to one police force cover the same area during the same 6-hour shift only every three days. This generates quasi-random variation in the days of the month, days of the week, and 6-hour shift in the coverage of police patrols.⁵⁰

Figure 4 shows the distribution of robberies in Milan based on the day triplet, where black squares and grey crosses indicate robberies that are under the responsibility of the *Carabinieri* and of the *Polizia*. Each panel represents a map of Milan (latitude vs. longitude) in different day/time combinations. In combination 1, the *Carabinieri* patrols the northwestern part of the city while the *Polizia* patrols the rest. In combination 2, the *Carabinieri* cover the northeastern part and in combination 3 the southern part.⁵¹

Given its predetermined pattern, criminals could target areas where predictive policing is not deployed. Such an endogenous response would lead to quantitative and qualitative sorting, with more crimes as well as more professionally organized crimes falling in the

⁴⁹Mastrobuoni (2019) shows that shift changes disrupt the productivity of police forces.

⁵⁰The assignment is also going to be informative about the underlying mechanism. KeyCrime might improve productivity through its prediction algorithm or simply because the improved data gathering delivers hidden clues about the crimes. For example, looking at the CCTV footage police officers might discover the criminals' identity. Yet, in this event the reductions would be equally spread over *Polizia* and *Carabinieri* assignments. In other words, any measurable change in clearances due to the assignment has to be linked to the way the data gathering and data analysis influences policing in the treated areas at times of treatment. Valuable predictions appear to be the only channel that is left.

⁵¹The handful of outliers may be driven by special police forces (e.g., motorcycles) or by measurement error, but are too few to influence any of the results.




Figure 4: Geographic Distribution of Robberies by Group

Notes: Each of the 12 panels corresponds to a map of Milan and each observation represents a robbery investigated by either the *Polizia* or the *Carabinieri*. Each group of police deployment is based on the exact day and time of a robbery. There is one panel for each shift (columns) and type of day (1-3). In each map the horizontal line separates the Southern area from the Northern one, while the vertical line separates the North-Western area from the North-Eastern one.

“untreated” areas. The endogenous response of criminals will be tested in Appendix Section A.3. There is no evidence of sorting.

Treatment (*Polizia*) and control (*Carabinieri*) patrols might, irrespective of predictive policing, differ in their productivity. The two forces share the same functions and objectives, which leads to considerable rivalry - and to surprising commonalities. Not only do the two forces share the same equipment (e.g. the Beretta 92 is their standard service weapon, and the Alfa Romeo 159, 2.4 JTDM 20v with 200 horsepower, is their standard service car, see Figure A11), they are almost the same size.⁵² Yet, a credible identification strategy would have to difference out any underlying productivity differences, which is the objective of the second difference. The balance tests for the first difference and the difference-in-difference are in Appendix Section A.3. In short, there is evidence that the quasi-experiment generates a well-balanced sample.

Second Differences: Data Availability

KeyCrime makes individual predictions, which are based on previous robberies. This implies that differences in clearances for the very first robbery can be used to control for underlying productivity differences between the two forces.

The identification rests on the assumption that differences in clearance rates between the two police forces that are not driven by predictive policing are the same for first and subsequent robberies within a sequence.⁵³ The probability of clearing a robbery might differ between first and subsequent robberies for other reasons too. There is likely to be a strong selection if the most inept robbers are caught immediately, and the ones that are not might also learn with experience. However, there is no reason why these differences

⁵²According to law, nationwide there are 57,336 police officers and 48,050 *Carabinieri* officers, both forces have 20,000 sergeants (*sovrintendenti*), they have similar numbers of inspectors (17,664 in the police and 16,031 in the *Carabinieri*), and the numbers of top-rank officials are also roughly the same.

⁵³In principle, when analyzing photographic evidence the *Polizia* might recognize individuals. But this would happen irrespective of the force that is operating on the ground and, thus, would not be able to explain differences in productivity. If the place and time of robberies for first time offenders were similar to the place and time for the repeat offenders, the estimates would be biased towards zero and vice versa.

should differ across the two police forces, unless smarter robbers choose the less productive police force (which is tested in Section A.3).

I model clearances using a linear probability model, where the dummy variable is equal to one when the τ -th robbery within a sequence i is cleared before the next robbery takes place:

$$Cleared_{i,\tau} = \alpha + \delta_0 Polizia_{i,\tau} + \delta_1 \mathbf{1}\{\tau_i > 1\} + \delta_2 Polizia_{i,\tau} \times \mathbf{1}\{\tau_i > 1\} + \gamma' X_{i,\tau} + \epsilon_{i,\tau}. \quad (1)$$

The coefficient δ_0 on *Polizia* measures the simple difference in clearance rates between the *Polizia* and the *Carabinieri* for the first robbery, while δ_2 measures the additional difference between the two forces when the *Polizia* has some data available for that sequence.

Not only should we expect there to be a difference between first and subsequent robberies, but as the police force keeps collecting information about the robbers, the difference in productivity should also increase. I use the variable τ_i instead of $\mathbf{1}\{\tau_i > 1\}$ to provide difference-in-differences estimates where the difference in productivity can differ as a function of the number of robberies.

There are two potential issues with the identification strategy based on τ , the number of the sequence. The first is that the *Polizia* officers might be improving as τ grows simply because they collect information across the entire city at all times, while the *Carabinieri* officers restrict their work to the assigned areas at the assigned times. A second issue is that the number of the sequence (τ) might be misclassified, potentially biasing the results.⁵⁴ Errors when linking different robberies over time, leading to errors in the measurement of first and subsequent robberies, would lower the accuracy of the predictions and the estimated efficacy of predictive policing, unless such errors are systematically and

⁵⁴As a robustness check in Appendix Section A.2 I focus on robberies where the definition of links has been aided by CCTV footage.

differentially linked to clearances. Since sequences are based on data collection by the *Polizia*, one might worry that such errors are artificially biasing the effects in favour of the *Polizia*. The next identification strategy addresses these concerns.

Alternative Second Differences: Delays and Information Sharing

Two alternative strategies restrict attention to subsequent robberies (there must be a linkage for both police forces) and exploit variation in the usage of predictive policing that does not depend on the measurement of τ . The first is based on the delay in data gathering and the second is based on the data sharing between *Polizia* and *Carabinieri* that started in January 2010.

To reduce victims' immediate distress and, consequently, avoid any recall bias about the robbery, the *Polizia* waits until the following day to interview victims. According to Mario Venturi, the inventor of KeyCrime, the interviews usually happen as soon as businesses are open (the exact time, which is not recorded, depends on when the victim becomes available). This means that when robbers perform two robberies in one day, patrols will not have an updated version of the prediction for the second robbery until the next day. Comparing clearance rates within the *Polizia* before and after KeyCrime has been updated eliminates the concern that the effects are driven by the differential coverage of the city (1/3 against 2/3). It also alleviates any concern about the measurement of τ .

The new equation is similar to Equation 1. The only difference is that I replace the binary variable $\mathbf{1}\{\tau_i > 1\}$ with a binary variable equal to one when KeyCrime has been updated. Moreover, trimming the window around the time the software has been updated I can compare the likelihood of clearing a robbery just before and just after the update for both police forces.

A final change in policy that I exploit is that prosecutors forced the *Polizia* to share their reports with the *Carabinieri*, starting at the beginning of 2010. The difference-in-differences compare the productivity of the *Polizia* and the *Carabinieri* in the two years

when the *Polizia* had exclusive access to KeyCrime (2008/2009) to the period after they started sharing the predictions.

Table 3: Clearances by Number of the Sequence and Police Force

Number of the Sequence	<i>Carabinieri</i>	<i>Polizia</i>	<i>Polizia-Carabinieri</i>
1	0.130 <i>161</i>	0.148 <i>420</i>	0.017 (0.032)
2	0.029 <i>34</i>	0.114 <i>105</i>	0.085** (0.043)
3	0.087 <i>23</i>	0.224 <i>76</i>	0.137* (0.077)
≥ 4	0.058 <i>104</i>	0.145 <i>332</i>	0.087*** (0.031)

Notes: Years 2008 and 2009. The number of observations are shown in italics.

5 Results

First vs. Subsequent Robberies

The simple difference between *Polizia* and *Carabinieri* for sequence numbers equal to 1, 2, 3 and 4 and more are shown in Table 3. Since the sample size naturally shrinks with the number of the sequence, the differences when the number is greater or equal to 4 are lumped together. The only difference that is not significantly different from zero is for the very first robbery of a sequence ($\tau = 1$). All the other differences are positive and significant. The productivity gap increases up until $\tau = 3$ and then drops by about 5 percentage points between $\tau = 3$ and $\tau > 3$. This could be because, as more predictable criminals are arrested, the actions of those remaining tend to be more unpredictable. The selection effect may dominate the effect of gathering additional information. There is also clear evidence that in the absence of predictive policing, between the first and the second robbery clearance rates drop considerably, which is consistent with additional selection and possibly learning.⁵⁵

⁵⁵The clearance rates sorted by police force, sequence, and year are shown in Appendix Table A1. Across all years there are few differences in clearance rates between *Polizia* and *Carabinieri* for first

Table 4: Difference in Differences by Number of the Sequences and Police Force

	(1)	(2)	(3)	(4)	(5)	(6)
	The robbery has been cleared (0/1)					
<i>Sample</i>	<i>First robbery</i>		<i>Subsequent robberies</i>		<i>All robberies</i>	
<i>Polizia</i> Intervention	0.016 (0.032)	0.018 (0.030)	0.100*** (0.025)	0.087*** (0.027)	0.018 (0.032)	0.021 (0.026)
Subsequent robberies					-0.078** (0.032)	
Number of the sequence						-0.005*** (0.002)
<i>Polizia</i> Intervention interacted with:						
Subsequent robberies					0.078* (0.040)	
Number of the sequence						0.009*** (0.003)
Constant	0.120*** (0.032)	- -	0.074*** (0.022)	- -	- -	- -
Other Xs		✓		✓		
Observations	581	581	674	674	1,255	1,255
R-squared	0.001	0.209	0.020	0.101	0.009	0.010

Notes: Linear probability models of clearances with simple differences estimated using $Cleared_{i,k} = \alpha + \delta_m Polizia_{i,k} + \gamma' X_{i,k} + \epsilon_{i,k}$; $m = \mathbf{1}\{k > 1\}$, and differences-in-differences estimated using $Cleared_{i,k} = \alpha + \delta_0 Polizia_{i,k} + \delta_1 \mathbf{1}\{k > 1\} + \delta_2 Polizia_{i,k} \times \mathbf{1}\{k > 1\} + \gamma' X_{i,k} + \epsilon_{i,k}$. The indices i and k correspond to the criminal group and the number of the sequence of robberies. All regressions control year fixed effects. Additional controls include the following fixed effects: month, day of the week, shift-turnover, morning, evening, and night shift, daylight, Western, North-eastern part of the city, firearm, knife, “some Italian,” “different nationalities,” pharmacy, other business, supermarket, bank, video rental, tobacco shop, as well as the following continuous variables: average age, haul, number of offenders, day of the month, number of the sequence. Robust clustered (by sequence) standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The estimated coefficients of Equation 1 are shown in Table 4. I start with simple differences. Columns 1 and 2 restrict the analysis to first robberies ($\tau_i = 1$), and columns 3 and 4 to subsequent incidents ($\tau_i > 1$). The difference in clearance rates is close to zero among first robberies and is equal to 10 percentage points (significant at the one percent level with standard errors clustered by sequence) for subsequent robberies.

Consistent with the quasi-random assignment of police forces to crimes, controlling for additional regressors listed at the bottom of the table (columns 2 and 4) leaves the coefficients almost unchanged. Relative to the *Carabinieri* these results mean that the *Polizia* officers are almost 3 times more likely to solve subsequent robberies compared to the *Carabinieri* officers. If this difference was driven by underlying differences in productivity, e.g. having wider control over the city (2 out of 3 areas), or, possibly, more efficient police officers, one would expect to find a similar difference among first robberies.

In Column 5, when computing the difference-in-differences between *Polizia* and *Carabinieri* for first and subsequent robberies, the effect of predictive policing is almost 8 percentage points.⁵⁶ In Column 6, allowing the effects of predictive policing to depend on the number of robberies increases the fit of the model as well as the precision of the estimates. When there is a *Polizia Intervention* the likelihood of clearing a case increases by 0.9 percentage points (more than 10 percent) for each additional robbery (*Number of the sequence*) the predictive policing software can analyze. The average number of subsequent robberies is 8, so the average predicted difference is about 0.07, which is close to the difference estimated in Column 5.

It is also important to note that for the *Carabinieri* the coefficients on “Subsequent robberies” and on the “Number of the sequence” are negative, indicating that, due either to selection or learning, successful robbers become increasingly difficult to arrest. Predictive

events in a sequence. For subsequent events, clearance rates converge only after 2010, when the *Polizia* started sharing predictive data with the *Carabinieri*.

⁵⁶The last two columns do not control for additional regressors, though as in the previous specifications the results do not change when doing so.

policing counteracts this pattern.




Figure 5: Clearance Rates by the Number of Days Since Previous Robbery

Notes: The numbers next to the bars indicate the number of observations. For lags 2 and 3 the estimates are based on a simple moving averages of order one ($\hat{r}_t = \frac{\bar{r}_t + \bar{r}_{t-1}}{2}$). The vertical lines show the 95 percent confidence intervals.

Delays in Data Collection

Next, I exploit the procedural lags in collecting the data from the previous robbery, which reduces the sample by about 50 percent. Figure 5 shows the respective clearance rate for the *Polizia* and the *Carabinieri* depending on whether the second robbery happens within the same day (a lag of 0 days), a few days later (1 to 3), or 4 and more days later. Due to the small sample size, for lags 2 and 3 I smooth the series using a moving average of order one⁵⁷ and use all the data (2008-2011).⁵⁸

The squares indicate the average clearance rates, the vertical bars the corresponding 95 percent confidence intervals. The horizontal line corresponds to the average when lumping together all “subsequent day” robberies that do not happen on the same day. For the *Polizia* clearances jump from less than 5 percent to more than 15 percent when the robberies happen one day later, rather than on the same day (there are 88 robberies in

⁵⁷The smoothed clearance rate is equal to the actual clearance rate c_t within the first two days $t \in \{0, 1\}$ and equal to a moving average $\tilde{c}_t = \frac{c_t + c_{t-1}}{2}$ for the other days.

⁵⁸Confidence intervals can be asymmetric and are constrained to be bound at 0.

the first category and 125 in the second). For the *Carabinieri* clearance rates are more noisy, as the sample size is smaller, but there is no evidence of an increase in clearances once a day has passed after the previous robbery.

Table 5: Difference in Differences by Same Day Robberies and Police Force

	(1)	(2)	(3)	(4)	(5)	(6)
	The robbery has been cleared (0/1)					
<i>Sample</i>	<i>Polizia in subsequent</i>		<i>Carabinieri in subsequent</i>		<i>Subsequent robberies</i>	
Police Intervention					-0.022 (0.057)	0.065* (0.035)
Different day	0.122*** (0.031)	0.129*** (0.038)	0.001 (0.061)	0.047 (0.075)	-0.011 (0.056)	0.026 (0.051)
Number of the series						-0.001 (0.003)
Number of the series $\times$ Different day robbery						-0.000 (0.003)
Polizia Intervention $\times$ Different day robbery					0.133** (0.064)	
Polizia Intervention $\times$ Number of the series						-0.006** (0.003)
Polizia Intervention $\times$ Number of the series $\times$ Different day robbery						0.012*** (0.004)
Constant	0.070** (0.031)	- -	0.059 (0.066)	- -	- -	- -
Other Xs		✓		✓		
Observations	510	510	160	160	670	670
R-squared	0.020	0.120	0.000	0.290	0.030	0.034

Notes: Linear probability models of clearances, where the estimated equations resemble Equation 1 (also shown in the previous table), with the binary variable

$$1(\text{different day robbery}_{i,k}) = \begin{cases} 1 & \text{if the robbery does not happen the same day;} \\ 0 & \text{otherwise.} \end{cases}$$

replacing $\mathbf{1}\{k > 1\}$. Robust clustered (by sequence) standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The additional controls are listed in the notes of Table 4.

The regression coefficients that measure differences based on software updates is in Table 5. Columns 1 and 2 show that the *Polizia*'s productivity is considerably larger when robbers do not carry out their subsequent robbery within the same day. For the *Carabinieri* no such difference emerges (columns 3 and 4). Conversely, within the same day the two police forces have very similar clearance rates. The difference-in-differences are shown in Columns 5 and 6 and are slightly larger than the difference-in-differences

Table 6: Difference in Differences by Police Force and Pre/Post the 2010 Data Sharing

	(1)	(2)	(3)	(4)
	Robbery has been cleared (0/1)			
<i>Sample</i>	<i>First robberies</i>		<i>Subsequent robberies</i>	
Polizia Intervention $\times$ Pre-2010	0.002 (0.056)	0.001 (0.056)	0.078** (0.037)	0.080** (0.039)
Pre-2010	-0.029 (0.047)	-0.048 (0.052)	-0.062** (0.031)	-0.041 (0.036)
Polizia Intervention	0.015 (0.046)	0.022 (0.047)	0.016 (0.029)	0.012 (0.028)
Constant	0.159*** (0.039)		0.118*** (0.025)	
Other Xs		✓		✓
Observations	904	904	1,263	1,263
R-squared	0.002	0.129	0.008	0.080

Notes: Linear probability models of clearances were the estimated equations resemble Equation 1 (also shown in the previous table), with the binary variable

$$1(\text{Pre-2010}_{i,k}) = \begin{cases} 1 & \text{if the robbery happened before 2010;} \\ 0 & \text{otherwise.} \end{cases}$$

replacing $1\{k > 1\}$. Robust clustered (by sequence) standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The additional controls are listed in the notes of Table 4.

based on subsequent vs. first robbery identification strategy.⁵⁹ ⁶⁰ The *Carabinieri*'s differences in clearance rates between same day and different day robberies are very close to zero.

Such a rapid response of predictive policing brings credit to the hypothesis that the effects are driven by updated police reports leading to better focused patrolling. If the effects were driven by investigations being aided by a more thorough data gathering, the timing of the effects would probably be less immediate.

Information Sharing with the *Carabinieri*

Table 6 shows that starting in 2010 the productivity of the *Carabinieri* does indeed converge to that of the *Polizia*.⁶¹ The *Carabinieri* did not disclose how they use the information from the KeyCrime, though the gap between the two forces is considerably smaller after 2010.

To summarize, estimates that exploit the discontinuity around the time the software is updated give very similar results: there are no productivity differences when predictive policing is either not available or has not yet been updated, and large productivity differences once predictive policing is fully operational. Moreover, once the *Carabinieri* have access to the KeyCrime reports their productivity goes up as well. Both results suggest that misclassification in the number of the sequence is not biasing the first set of results. Three results based on different identification strategies all point in the same direction and make it harder to come up with an alternative story—something that would drive productivity differences across “first” and “subsequent robberies,” “same day” and “different day” robberies, and between robberies that took place before and after information

⁵⁹As done earlier, the last two columns do not control for additional regressors, though again results do not change when doing so.

⁶⁰Consistently with the evidence shown in Figure 5, when I restrict the comparison to robberies that happen within a few days of the previous robbery, the results stay the same (see Appendix Table A2). The results in Columns 1 and 5 are not exactly the same as those in Table 5 because to increase power I am using all years, again potentially biasing the coefficients for the *Carabinieri* upwards.

⁶¹These results are in line with the raw numbers shown in Appendix Table A1.

sharing started.

Appendix Section A.2 outlines a series of robustness checks, which for brevity and for reasons of statistical power are based on the first set of results.

6 Policy Implications

As mentioned earlier, clearing a robbery means that at least one robber is arrested. Based on data collected by the *Polizia* close to 100 perpetrators were arrested between 2008 and 2009. Of these only one perpetrator was acquitted, while the rest received a total of 420 years of prison time (about 4 years per prisoner).⁶²

After their first robbery about 30 percent of robbers are linked to a second robbery. After that second robbery, almost all robbers repeat their crime until they are arrested.⁶³ For this reason differences in clearance rates lead to differences in the expected number of robberies criminal groups are able to organize before ending up in prison. The *Polizia* and the *Carabinieri* share these incapacitation effects (there is quasi-random assignment of crimes to the two police forces) so such effects cannot be measured directly, but it is possible to retrieve the effects using differences in clearance rates and some simple algebra.

The two difference-in-difference estimates were 7.8 and 13.3 percent. Using an average of 10 percent as an estimate (this is an approximation as clearance rates might also vary slightly with the number of the sequence) and setting the counterfactual clearance rate at 5.6 percent for subsequent robberies without predictive policing (this is the clearance rate for the *Carabinieri* in 2008 and 2009), the expected number of robberies each group of recurrent criminals commits drops from about 17.8 to about 6.4, a 2.8 to 1 ratio.⁶⁴

⁶²Four criminals were given alternative punishments instead of prison time.

⁶³The data available allow me to reconstruct the “survival table” of robbers (for simplicity I do not distinguish by police force). Appendix Table A3 shows the distribution of robberies based on the “Number of the series.”

⁶⁴When the criminal attempts are frequent and persistent, as happens to be the case for Milan, the expected number of robberies is approximately equal to $\sum_{\tau=0}^{\infty} (1-c)^{\tau} = 1/c$, where c is the clearance rate.

The analysis based on aggregate data (Section 3) showed that between 2008 and 2011 robbery rates fell by similar proportions. Any endogenous response of criminals as well as any mean-reversion would make it hard to compare the two estimates, and yet it is comforting to see that they are in the same ballpark.⁶⁵

Since there are about 255 successful first time robbers per year and about one-third re-offend, the reduction of 11 robberies per sequence leads to a reduction of 935 robberies over the lifetime of such sequences (in the long run deterrence might lead to even larger reductions). Multiplying 935 robberies by the overall average haul (€2,800), the direct costs that are prevented by the use of predictive policing are close to €2.5 million,⁶⁶ or about US\$ 2.8 million. Indirect costs of crime are likely to be an order of magnitude larger than the direct costs (Cook, 2009).

The labor cost of the four full-time police officers who collect the data and predict the crimes can hardly outweigh such benefits, as it is less than €100,000 per year.⁶⁷ The investments in capital (an office, computers, monitors, etc) are little more than few thousand euros a year.

But one needs to take into account the increased cost from incarcerating arrested criminals. Since almost all robbers eventually end up in prison, or in other words, since $(1 - c_t)^\tau$ converges to 0 reasonably quickly for clearance rates that are close to 10 percent, predictive policing merely brings forward the time of incarceration. Since most re-offending happens within a few weeks, predictive policing tends to bring forward arrests by a few months, at most a few years. The average time between robberies is 15 days and the number of prevented robberies is 11, so the average time to arrest falls by about

⁶⁵Notice that a gradual trend reversion is exactly what incapacitation would be generating, as more and more potential offenders are prevented from committing a sequence of bank robberies. Starting in 2012, probably as Milan transitions to an equilibrium with improved clearance rates, the reduction in robbery rates appears to slow down.

⁶⁶The standard deviation, derived using the delta method for the estimated clearances and the estimated haul, is close to €1.4 million, with an implied t-statistic of 1.79.

⁶⁷According to the decree D.P.R. 15/3/2018 n. 39 the before tax salary of police officers is between €17,680 and €19,819 depending on the experience.

6 months. With a reduction of 6 months, an interest rate of five percent (an upper bound of the yields on the Italian government bonds), and an average yearly cost per incarceration of €50,000 (see Barbarino and Mastrobuoni, 2014), the cost of bringing forward the incarceration expenses would be at most 2.5 percent of €1.25 million, or €31,250.⁶⁸

There are additional costs and benefits related to how the additional information collected through predictive policing helps the prosecutors to build a case in court. Unfortunately, there are no data (e.g. post-incarceration recidivism of convicted robbers) to evaluate such costs and benefits, though they are arguably smaller in magnitude than the direct cost of crimes, and would hardly overturn the cost/benefit findings. Overall, the cost of introducing predictive policing appears to be an order of magnitude lower than the benefit.

7 Conclusions

This study used the quasi-random allocation of two almost identical police forces to crimes, to test whether differences in police productivity can be attributed to the availability of advanced IT. Once the data to be analyzed become available, either because a history of criminal events is gathered or because the officers have enough time to process the new information, the differences in productivity are striking.

The micro-level information shows that these productivity differentials are consistent with the criminals' observed persistence in criminal behavior. Over time recurrent robbers tend to target similar businesses, around the same neighborhood, and at the same time of the day; together with the fact that robbers reoffend at a very high frequency (60 percent are "back in business" within one week), this makes robbers predictable. Predictive policing can increase the productivity of police patrols by allowing them to be closer to

⁶⁸Victimizations, as well as incarcerations, generate pain and suffering which I do not attempt to quantify, as both are extremely hard to measure.

the crime scene.

Moreover, the software’s micro-predictive ability has been shown to grow with the number of offenses collected, which helps in capturing the most prolific criminals. Since these criminals tend to be the most socially harmful, predictive policing leads to more selective incarcerations. A rough cost/benefit analysis suggests that micro-predictive policing represents a highly cost efficient IT investment. Predicting where a robbery is most likely to occur and being ready to apprehend the offenders is also very different from developing lists of potential future offenders without probable cause, a practice that would be vulnerable to misuse.

The experimental design allowed me to estimate the effect of predictive policing on the likelihood that a robbery is solved and a perpetrator is arrested. An open question is whether over time, as criminals perceive the productivity of policing to be increasing, additional deterrence is generated. It could either deter crime altogether, convincing criminals to switch to other crimes, or it might displace crime from Milan to other cities. Criminals may also exert more effort to try to be more unpredictable. Another open question is whether the *Polizia* is using the best possible prediction algorithm and whether predictions could be improved using more or even less detail about the robberies. Predictions based on such detail, where labor input is relatively high, may become unfeasible when the number of crimes grows larger. Whether there is a threshold level of complexity is another open question.

In conclusion, this paper adds to the limited micro-level evidence on the positive productivity effects of IT investments. It is also the first quasi-experimental evaluation of predictive policing. These IT investments can be highly effective in improving the productivity of the police officers in their role as apprehension agents.

References

- Alberto Abadie and Javier Gardeazabal. The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1):113–132, 2003.
- Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program. *Journal of the American statistical Association*, 105(490):493–505, 2010.
- Daron Acemoglu, Philippe Aghion, Claire Lelarge, John Van Reenen, and Fabrizio Zilibotti. Technology, information, and the decentralization of the firm. *The Quarterly Journal of Economics*, 122(4):1759–1799, November 2007.
- Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *Knowledge and Data Engineering, IEEE Transactions on*, 17(6):734–749, 2005.
- Agatha Christie. *The A.B.C. Murders*. Collins Crime Club, 1936.
- Joshua Angrist and Victor Lavy. New evidence on classroom computers and pupil learning*. *The Economic Journal*, 112(482):735–765, 2002.
- Susan Athey and Scott Stern. The impact of information technology on emergency health care outcomes. *Rand Journal of Economics*, 33(3):399–432, 2002.
- David H. Autor, Lawrence F. Katz, and Alan B. Krueger. Computing inequality: Have computers changed the labor market? *The Quarterly Journal of Economics*, 113(4):1169–1213, November 1998.
- Marzio Barbagli and Asher Colombo. *Rapporto sulla criminalità e la sicurezza in Italia*. Gruppo24Ore, 2011.

- Alessandro Barbarino and Giovanni Mastrobuoni. The Incapacitation Effect of Incarceration: Evidence from Several Italian Collective Pardons. *American Economic Journal: Economic Policy*, 6(1):1–37, 2014.
- Ann P. Bartel, Casey Ichniowski, and Kathryn L. Shaw. How Does Information Technology Really Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement and Worker Skills. *The Quarterly Journal of Economics*, 4(122), 2007.
- Eli Berman, John Bound, and Zvi Griliches. Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufactures. *The Quarterly Journal of Economics*, 109(2):367–97, May 1994.
- Sandra E. Black and Lisa M. Lynch. How to compete: The impact of workplace practices and information technology on productivity. *The Review of Economics and Statistics*, 83(3):434–445, August 2001.
- Jordi Blanes i Vidal and Tom Kirchmaier. The effect of police response time on crime clearance rates. *Review of Economic Studies*, 85(2):855–891, 2017.
- Jordi Blanes i Vidal and Giovanni Mastrobuoni. Police patrols and crime. 2018.
- Nicholas Bloom, Raffaella Sadun, and John Van Reenen. Americans Do IT Better: US Multinationals and the Productivity Miracle. *American Economic Review*, 102(1):167–201, 2012.
- Richard J Bolton and David J Hand. Statistical Fraud Detection: A Review. *Statistical Science*, pages 235–249, 2002.
- Anthony A. Braga. The effects of hot spots policing on crime. *Annals of the American Academy of Political and Social Science*, 578:pp. 104–125, 2001.

- Timothy F. Bresnahan, Erik Brynjolfsson, and Lorin M. Hitt. Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, 117(1):339–376, February 2002.
- Erik Brynjolfsson and Lorin M. Hitt. Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48, Fall 2000.
- Erik Brynjolfsson and Adam Saunders. *Wired for innovation: How information technology is reshaping the economy*. MIT Press, 2009.
- Jacqueline Cohen and Jens Ludwig. Policing crime guns. In *Evaluating Gun Policy: Effects on Crime and Violence*, page 217. Brookings Institution Press, 2003.
- Philip J Cook. The clearance rate as a measure of criminal justice system effectiveness. *Journal of Public Economics*, 11(1):135–142, 1979.
- Phillip Cook. Robbery. In Michael Tonry, editor, *Oxford Handbook on Crime and Public Policy*. Oxford University Press, 2009.
- John Daintith. A Dictionary of Physics (6 ed.). *Oxford reference online*: <http://www.oxfordreference.com/>, 2009.
- Paul A. David. The dynamo and the computer: An historical perspective on the modern productivity paradox. *American Economic Review*, 80(2):355–61, May 1990.
- Sarah L Desmarais and Jay P Singh. Risk assessment instruments validated and implemented in correctional settings in the United States. Report, Council of State Governments, Lexington, KY, 2013.
- Rafael Di Tella and Ernesto Schargrodsky. Do police reduce crime? Estimates using the

- allocation of police forces after a terrorist attack. *The American Economic Review*, 94(1):115–133, 2004.
- Jennifer L. Doleac. The effects of DNA databases on crime. *American Economic Journal: Applied Economics*, 9(1):165–201, 2017.
- Mark Doms, Timothy Dunne, and Kenneth R Troske. Workers, wages, and technology. *The Quarterly Journal of Economics*, 112(1):253–90, February 1997.
- Nikolay Doudchenko and Guido W. Imbens. Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. Technical report, National Bureau of Economic Research, 2016.
- Mirko Draca, Stephen Machin, and Robert Witt. Panic on the Streets of London: Police, Crime and the July 2005 Terror Attacks. *American Economic Review*, 101(5):2157–81, 2011.
- David Dranove, Chris Forman, Avi Goldfarb, and Shane Greenstein. The trillion dollar conundrum: Complementarities and health information technology. *American Economic Journal: Economic Policy*, 6(4):239–270, 2014.
- Steven N. Durlauf, Salvador Navarro, and David A. Rivers. Understanding aggregate crime regressions. *Journal of Econometrics*, 158(2):306 – 317, 2010. ISSN 0304-4076. Specification Analysis in Honor of Phoebe J. Dhrymes.
- Wioletta Dziuda and Giovanni Mastrobuoni. The Euro Changeover and its Effects on Price Transparency and Inflation. *Journal of Money, Credit and Banking*, 41(1):101–129, 2009.
- The Economist. The Aftershocks of Crime. *The Economist*, October 21 2010.

- Andrew Guthrie Ferguson. *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. NYU Press, New York, 2017.
- Luis Garicano and Paul Heaton. Information technology, organization, and productivity in the public sector: Evidence from police departments. *Journal of Labor Economics*, 28(1):167–201, 01 2010.
- Avi Goldfarb and Catherine Tucker. Digital economics. *Journal of Economic Literature*, 57(1):3–43, 2019.
- Austan Goolsbee and Jonathan Guryan. The Impact of Internet Subsidies in Public Schools. *The Review of Economics and Statistics*, 88(2):336–347, May 2006.
- Lev Grossman, Cleo Brock-Abraham, Nick Carbone, Eric Dodds, Jeffrey Kluger, Alice Park, Nate Rawlings, Claire Suddath, Feifei Sun, Mark Thomson, Bryan Walsh, and Kayla Webley. The 50 best inventions. *Time Magazine*, November 28 2011.
- Thomas N. Hubbard. Information, decisions, and productivity: On-board computers and capacity utilization in trucking. *American Economic Review*, 93(4):1328–1353, 2003.
- Jonathan Klick and Alexander Tabarrok. Using terror alert levels to estimate the effect of police on crime. *Journal of Law & Economics*, 48(1):267–79, April 2005.
- Steven D. Levitt. Understanding why crime fell in the 1990s: Four factors that explain the decline and six that do not. *Journal of Economic Perspectives*, 18(1):163–190, September 2004.
- Alexandre Mas. Pay, reference points, and police performance. *The Quarterly Journal of Economics*, 71(3):783–821, 2006.
- Giovanni Mastrobuoni. Police and Clearance Rates: Evidence from Recurrent Redeployment within a City. *Journal of Public Economics*, 176:18–31, 2019.

- Giovanni Mastrobuoni and David A. Rivers. Optimising criminal behaviour and the disutility of prison. *The Economic Journal*, 129(619):1364–1399, April 2019.
- Jeffrey S McCullough, Stephen T Parente, and Robert Town. Health information technology and patient outcomes: the role of information and labor coordination. *The RAND Journal of Economics*, 47(1):207–236, 2016.
- George O. Mohler, Martin B. Short, Jeffrey Brantingham, Frederick P. Schoenberg, and George E. Tita. Self-exciting Point Process Modeling of Crime. *Journal of the American Statistical Association*, 106(493):100–108, 2011.
- George O Mohler, Martin B Short, Sean Malinowski, Mark Johnson, George E Tita, Andrea L Bertozzi, and P Jeffrey Brantingham. Randomized Controlled Field Trials of Predictive Policing. *Journal of the American Statistical Association*, 110(512):1399–1411, 2015.
- Daniel S Nagin, Robert M Solow, and Cynthia Lum. Deterrence, criminal opportunities, and police. *Criminology*, 53(1):74–100, 2015.
- Emily G. Owens. Cops and cuffs. In Philip Cook, Stephen Machin, Olivier Marie, and Giovanni Mastrobouni, editors, *Lessons from the Economics of Crime: What Works in Reducing Offending?* Cambridge: MIT Press, 2014.
- Beth Pearsall. Predictive policing: The future of law enforcement? *National Institute of Justice*, 2010.
- Walter L. Perry, Brian McInnis, Carter C. Price, Susan C. Smith, and John S. Hollywood. *Predictive policing: The role of crime forecasting in law enforcement operations*. Rand Corporation, Santa Monica, CA, 2013.
- Police Executive Research Forum. Future trends in policing. Technical report, Washington, D.C.: Office of Community Oriented Policing Services, 2014.

- Somini Sengupta. In hot pursuit of numbers to ward off crime. *The New York Times*, June 19 2013.
- Vikram Sharma. Indian police to be armed with big data software to predict crime. *The New Indian Express*, 23rd September 2017.
- Lawrence W. Sherman and David Weisburd. General deterrent effects of police patrol in crime “hot spots”: A randomized, controlled trial. *Justice Quarterly*, 12(4):625–648, 1995.
- Lawrence W. Sherman, Patrick R. Gartin, and Michael E. Buerger. Hot spots of predatory crime: Routine activities and the criminology of place. *Criminology*, 27(1):27–56, 1989.
- Space and Naval Warfare Systems Center Atlantic. Geographic information systems and predictive policing application note. System assessment and validation for emergency responders (saver), SAVER Program, U.S. Department of Homeland Security, Science and Technology Directorate, Washington, DC, August 2013.
- Kevin J. Stiroh. Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say? *American Economic Review*, 92(5):1559–1576, December 2002.
- Robert Tibshirani. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
- John Van Reenen, Mirko Draca, and Raffaella Sadun. Productivity and ICTs: a review of the evidence. In Chrisanthi Avgerou, Robin Mansell, Fanny Quah, and Roger Siverstone, editors, *The Oxford Handbook of Information and Communication Technologies*, pages 100–147. Oxford University Press, Oxford and New York, 2009.
- David Weisburd and John E Eck. What can police do to reduce crime, disorder, and fear? *The Annals of the American Academy of Political and Social Science*, 593(1):42–65, 2004.

David Weisburd and Lorraine Green. Policing drug hot spots: The Jersey City drug market analysis experiment. *Justice Quarterly*, 12(4):711–735, 1995.

David Weisburd, Stephen D. Mastrofski, Ann McNally, Greenspan Rosann, and James J. Willis. Reforming to Preserve: Compstat and Strategic Problem Solving in American Policing. *Criminology & Public Policy*, 2(3):421–456, 2003.

A Online Appendix

A.1 Municipality and Province-level Crime Data

The Italian statistical office (ISTAT) started releasing municipality-level yearly (and not monthly) crime data only in 2010.⁶⁹ For earlier years, ISTAT’s crime data are only available at the province level (there are 107 provinces and more than 8000 municipalities).

Thanks to the Italian Statistical office ISTAT I put together province-level yearly crime data that cover the years 1983 to 2017. For the entire period ISTAT measures the total number of robberies, the number of bank robberies and the number of robberies against postal offices. Starting in 2010 ISTAT decided to discontinue measuring robberies against jeweleries, against armored cash transport vehicles, and against trucks, and started measuring robberies against all commercial businesses, and street robberies.

In Milan in 2010 there were 3,796 robberies. Of these, only 654, or 17 percent, were against commercial businesses. The main issue is that according to the *Polizia* most of the robberies that are not perpetrated against businesses are so called improper robberies, where the violence postdates a simple theft and is aimed at maintaining the loot. In 2010 there were 180 bank robberies, representing 27 percent of commercial robberies, and 19 robberies against postal offices. The prevalence of bank robberies at the province level suggests that professional robbers tend to target many banks that are located outside of the municipality.

The province of Milan has 3.7 million inhabitants. To balance population size, I select 24 provinces with at least 600000 inhabitants. To have a measure of commercial robbery rates for the entire period, I first use the years 2010 to 2017 to predict commercial

⁶⁹According to the US Uniform Crime Reports in 2009 robbery rates were 133 per 100,000 inhabitants, while they were 58.7 per 100,000 inhabitants in Italy (Barbagli and Colombo, 2011). About 25 and 42 percent of robberies reported to the police occur in businesses in Italy and in the US (Barbagli and Colombo, 2011, Cook, 2009). Appendix Figure A1 plots aggregate crime rates and clearance rates of Italian provinces between 1983 and 2003, several years before the introduction of KeyCrime. With respect to both rates the province of Milan appears to be in the middle of the distribution.

robberies based on bank robberies and postal offices and all remaining robberies. In a second step I predict commercial robberies for the entire period. All three robbery rates are good predictors, and the R-squared is 67 percent.

A.2 Robustness Regressions

This section contains some common robustness checks, including functional form assumptions, sampling, outliers, spillovers and heterogeneity of the effects. Column 1 of Appendix Table A4 indicates that the marginal effects of a probit model are in line with the linear probability estimates.⁷⁰

In the first year of KeyCrime the *Polizia* appears to have investigated more robberies than chance would predict given the 1/3-2/3 assignment rule. Between 2008 and 2009 the robberies that fall in the *Carabinieri* zones increase from 22.6 percent to 29 percent (see Appendix Table A1). This might just be chance or it might be that the *Polizia* omitted to collect data on robberies that occurred inside the areas assigned to the *Carabinieri*. Yet, when focussing on the year 2009 (Column 2) the effects are even larger.

To address the heterogeneity of results based on the availability of CCTV cameras, in Column 3 I only include robberies where it is known with certainty that current or past sites of commercial robberies have some CCTV cameras installed.⁷¹ The results are indeed a little larger, suggesting that the presence of CCTV cameras might benefit the predictions. Column 4 shows that the predictive policing effects are only slightly lower when focussing on robberies where the haul is above average. This may indicate that robberies with lower hauls are easier to predict. Excluding the single most targeted category, pharmacies, does not alter the results (Column 5), showing that the most targeted businesses are not driving the results. Finally, in the last column there is no evidence of

⁷⁰The same is true for the differences based on the software updates.

⁷¹I was not given any photographic evidence but was told that all banks, postal offices, pharmacies, and jewellery stores have CCTV systems that are running constantly.

biases due to spillover effects. When focussing on just the very first robbery of the day a police force has to deal with, the results are unchanged.

A.3 Testing for an Endogenous Response of Criminals

It is unlikely that robbers would know about the exact timing of when the data are collected and the software is updated, suggesting that the identification strategy based on the timing of the update is not prone to selection bias.

However, in the identification strategy that compares first vs. subsequent robberies, capable and experienced robbers might try to target victims when their business falls in the *Carabinieri* area. Starting with a balance test, Appendix Table A5 compares all the observable characteristics of the robbers and of the robberies, depending on whether the *Polizia* or the *Carabinieri* were covering the area. The single most striking difference is in the likelihood of clearing a robbery. The amount stolen, which is a measure of the ability of robbers (see Mastrobuoni and Rivers, 2019) does not show any differences. One cannot reject the hypothesis that not just the mean of the haul but its entire distribution is the same for the two forces (see the Appendix Figure A12). There are a few variables for which Appendix Table A5 measures small but significant differences between the two forces: year, day of the week, pharmacies and other businesses. Controlling for these small differences makes little difference. Balance tests that mimic the difference-in-differences strategy are in the Appendix Table A6. Again, for control variables only a few differences are significantly different from zero. Testing the joint significance of all difference-in-differences using a seemingly unrelated regression, one cannot reject that they are all equal to zero. For the strategy based on first vs. subsequent robberies the Wald test is 18.59 with a p-value of 72 percent, while for the strategy based on the timing of the update the Wald test is equal to 29 with a p-value of 17 percent. A few variables show a significant effect, especially pharmacies. Pharmacies are more likely to be targeted on a

different day (vs. the same day) when the *Polizia* is patrolling the neighborhood (vs. the *Carabinieri*). This by itself would not bias the results, as one can control for the business type. Yet, robbers targeting pharmacies more than once on the same day may be more likely to be intoxicated. In this case, assuming that robbers under the influence are more likely to be arrested, the results would be biased against finding a larger clearance rate for the *Polizia* in the following days. When I exclude pharmacies altogether, I find that the *Polizia* is still 11 percentage points more likely to clear a robbery.⁷²

If robbers knew about the productivity differences between the two forces there should be relatively more robberies that fall under the responsibility of the *Carabinieri* than of the *Polizia*, especially past the first robbery (when predictive policing might potentially aid the investigation).⁷³ In the Appendix Table A1 the curly brackets shows the fraction of robberies that fall in each area, and there is no evidence that robbers target areas that are patrolled by the *Carabinieri* more frequently. But the even more stringent test is that subsequent robberies, which are the ones for which the policing software predicts potential targets, time, etc., are not more likely to fall under the responsibility of the *Carabinieri*.⁷⁴

Despite the lack of an excess mass in the number of subsequent robberies that fall into the *Carabinieri* area, and the balance in the haul between the two forces, my final specification is again going to deal with the possibility that more able robbers sort into the *Carabinieri* zones. If such a selection was driving the results, sorting in the previous robberies would also be predictive of differences in the likelihood of a clearance as the unobserved ability of robbers that drives the selection would be a persistent trait. Appendix Table A7 shows that there is no evidence of this. Regressing the *Cleared robbery* dummy on the dummy whether the *Polizia* was patrolling the area as well as its first two lags reveals that only the last treatment matters. Moreover, Column 3 shows that

⁷²The results are available upon request.

⁷³The summary statistics Table 1 already showed that robberies do not seem to disproportionately target the *Carabinieri* patrol areas.

⁷⁴There is also no evidence that robbers tend to stay away from businesses with CCTV cameras as they get more experienced (results are available upon request).

the interaction between the current *Polizia* dummy and the past one has no impact on clearing the case. Moreover, a related implication is that the information collected by the *Polizia* officers when the *Carabinieri* officers covered the previous robbery is not inferior to that which would have collected.

Finally, in Column 5 there are controls for whether the last two robberies of a sequence have been investigated by the same force. Controlling for such potential informational advantage, which—covering two thirds of the city area—is more likely to be captured by the *Polizia*, does not alter the results.

A.4 Figures




Figure A1: Aggregate Robbery Rates and Clearance Rates

Notes: The dashed line simply plots the inverse of the clearance rate. Based on 103 Italian provinces between 1983 and 2003.




Figure A2: Map of the Municipality of Milan

Notes: The squares and crosses indicate the location of robberies assigned to the two forces. The grid defines the different neighborhoods.

SEARCH NOTE

Milan, XXXXXXX

REFERENCE: UPDATE OF SEARCH NOTE "XXXXXX"

TO THE HEAD OF GENERAL PREVENTION DEPT.
TO THE HEAD OF MOBILE FORCES

We would like to inform you that based on the analysis of robberies perpetrated against commercial businesses, particularly Supermarkets, SEQUENCE N. XXXXXXX appears to be still active.

The criminal group appears to be **responsible for 22 robberies**, starting February this year, in particular on day XXXXXX between time XXX and time XXXX they have robbed three business one after the other, XXXXX

Offender 1: Italian, aged between 25 and 35, approximately 185cm tall, medium build, masquerading with a white neck warmer, sunglasses and a motorcycle helmet.

Offender 2: Italian, aged between 25 and 35, approximately 170cm tall, skinny, with a black mole below his lip (left side) (it could be a piercing), masquerading with a dark coloured neck warmer, sunglasses and a motorcycle helmet.

Weapon: gun (in several occasions two guns).

Vehicles: IN RECENT CRIMES THEY HAVE USED A BLACK HONDA HORNET WITH PARTIAL PLATE NUMBER XXXXX

Look for vehicle:

- GREY Scooter with plate number XXXXXXX reported stolen on the XXXXX (used at least three times: the XXXXX, the XXXXX and the XXXXXX, not yet found).

Most likely victims: SUPERMARKETS.

Most likely times: between 5.30pm and 7pm

A likely area where the offenders and the vehicle can be found is the neighborhood ISOLA di Milano or the lower part of NIGUARDA.

Honda Hornet (black) plate n.. AK.....

Grey Scooter plate n. XXXXX reported stolen on the XXXXXX, still being used by the offenders.





Figure A3: Instructions for Police Patrols

Notes: This is a sample of instructions that have been used by the *Polizia*. The instructions have been translated from Italian. The original version is available upon request.




Figure A4: Comparison of Events

Notes: This is a screenshot of KeyCrime. The English text has been added to describe what is shown on the screen.




Figure A5: Predicted Targets

Notes: This is a screenshot of KeyCrime. The English text has been added to describe what is shown on the screen. Small blue dots indicate past victims, while little blue squares indicate potential victims.




Figure A6: Bank Robbery Rates in Milan against Synthetic LASSO Milan

Notes: Synthetic cities in the left panel are computed using Abadie et al. (2010)'s approach. The solid thick line represents the difference in monthly bank robbery rates per 100,000 inhabitants between Milan and its synthetic control city (cities without predictive policing that have more than 200,000 inhabitants). The thin lines represent the differences with respect to their synthetic control for all other cities. All lines represent differences-in-differences with respect to their January 2008 value. The right panel shows the simple difference in bank robbery rates between Milan and the largest 9 cities. Bank robbery data are collected by the Italian banking association (ABI).




Figure A7: Bank Robbery Rates in Milan against Synthetic LASSO Milan

Notes: Synthetic cities are computed using a LASSO procedure, which allows weights to be completely unconstrained (they do not sum up to one and can be negative). The solid thick line represents the difference in monthly bank robbery rates per 100,000 inhabitants between Milan and its synthetic control city (cities without predictive policing that have more than 100,000 inhabitants). The thin lines represent the differences with respect to their synthetic control for all other cities. All lines represent differences-in-differences with respect to their January 2008 value. Bank robbery data are collected by the Italian banking association (ABI).




Figure A8: Bank Robbery Rates and Commercial Robbery Rates in Milan against Synthetic Milan

Notes: Crime rates are based on ISTAT data (*Statistiche Giudiziarie Penali*). The predicted commercial robbery rates are based on a procedure described in Appendix Section A.1. The donor pool for the synthetic province of Milan is based on 24 provinces with at least 600,000 inhabitants in 2010 (the province of Milan has about 3.7 million inhabitants). The solid thick line represents the difference in yearly robbery rates per 100,000 inhabitants between Milan and its synthetic control city. The thin lines represent the differences with respect to their synthetic control for all other cities. The vertical lines indicate the time overlap between these data and the data on bank robberies collected by ABI (2004-2015), see Figure A6.




Figure A9: Differences in Crime Rates for Various Crime Types between Milan and Synthetic Milan

Notes: Crime rates are based on ISTAT data (*Statistiche Giudiziarie Penali*). The total crime rate is net of all the robberies. The donor pool for the synthetic province of Milan is based on 24 provinces with at least 600,000 inhabitants in 2010 (the province of Milan has about 3.7 million inhabitants). The solid thick line represents the difference in yearly rates per 100,000 inhabitants between Milan and its synthetic control city. The vertical lines indicate the time overlap between these data and the data on bank robberies collected by ABI (2004-2015), see Figure A6.




Figure A10: Geographic Distribution of Robberies by Criminal Group

Notes: Each panel corresponds to a map of Milan. Each panel compares the location of the universe of robberies (small grey dots) with the location of robberies organized by the most prolific offenders (large black circles).



Figure A11: *Carabinieri* and *Polizia*




Figure A12: Cumulative Distribution of the Haul

Notes: The haul is expressed in €1,000 and is truncated at 10,000 to focus where most the data are. The Kolmogorov-Smirnov test for equality of distribution functions cannot reject that null that the distributions are the same (irrespective of truncation).




Figure A13: Aggregate Robbery Rates against Bank Robbery Rates (in logs)

Notes: The dashed line simply plots the predicted values. The elasticity is equal to 1.39 with a standard deviation of 0.04 (clustered at the province level). Based on 103 Italian provinces between 1983 and 2003.

Table A1: Clearance Rates by Year and Police Force

	First event		Subsequent events	
	Carabinieri	Polizia	Carabinieri	Polizia
2008	0.124	0.160	0.049	0.121
	<i>89</i>	<i>256</i>	<i>61</i>	<i>257</i>
	<i>(0.26)</i>	<i>(0.74)</i>	<i>(0.19)</i>	<i>(0.81)</i>
2009	0.139	0.128	0.060	0.180
	<i>72</i>	<i>164</i>	<i>100</i>	<i>256</i>
	<i>(0.31)</i>	<i>(0.69)</i>	<i>(0.28)</i>	<i>(0.72)</i>
2008-2009	0.130	0.148	0.056	0.150
	<i>161</i>	<i>420</i>	<i>161</i>	<i>513</i>
	<i>(0.28)</i>	<i>(0.72)</i>	<i>(0.24)</i>	<i>(0.76)</i>
2010	0.136	0.152	0.116	0.135
	<i>66</i>	<i>158</i>	<i>129</i>	<i>288</i>
	<i>(0.29)</i>	<i>(0.71)</i>	<i>(0.31)</i>	<i>(0.69)</i>
2011	0.227	0.221	0.122	0.130
	<i>22</i>	<i>77</i>	<i>41</i>	<i>131</i>
	<i>(0.22)</i>	<i>(0.78)</i>	<i>(0.24)</i>	<i>(0.76)</i>
2008-2011	0.141	0.157	0.088	0.143
	<i>249</i>	<i>655</i>	<i>331</i>	<i>932</i>
	<i>(0.28)</i>	<i>(0.72)</i>	<i>(0.26)</i>	<i>(0.74)</i>

Notes: The number of observations are shown in italics, and the fraction of robberies assigned to *Polizia* and *Carabinieri* are shown in parentheses.

Table A2: Tighter Same Day vs. Different Day Differences

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	The robbery has been cleared (0/1)							
<i>Robberies happen within</i>	<i>whole period</i>	<i>4 days</i>	<i>3 days</i>	<i>2 days</i>	<i>whole period</i>	<i>4 days</i>	<i>3 days</i>	<i>2 days</i>
	<i>Polizia</i>				<i>Carabinieri</i>			
Different day	0.106*** (0.026)	0.109*** (0.032)	0.112*** (0.035)	0.107** (0.045)	0.030 (0.044)	0.017 (0.048)	0.028 (0.053)	-0.037 (0.047)
Constant	0.045* (0.023)	0.045* (0.023)	0.045* (0.023)	0.045* (0.023)	0.061 (0.041)	0.061 (0.041)	0.061 (0.041)	0.061 (0.041)
Observations	929	406	317	213	330	136	101	76
R-squared	0.008	0.018	0.023	0.029	0.001	0.001	0.002	0.009

Notes: Linear probability models of clearances. Simple differences based on Table 5 but focused on subsequent robberies that happen within a few days from the previous robbery. All regressions control for the four shift fixed effects (6 hour intervals). Robust clustered (by sequence) standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Table A3: Recurrence and Clearance

Number of the series	Cleared robbery		Clearance rate	Recurrence rate		
	No	Yes		estimate	upper bound	lower bound
1	771	136	0.15	-		-
2	215	29	0.12	0.32	0.42	0.27
3	153	23	0.13	0.82	0.84	0.72
4	111	19	0.15	0.85	0.87	0.74
5	91	8	0.08	0.89	0.91	0.76
6	77	7	0.08	0.92	0.93	0.85
7	64	8	0.11	0.94	0.94	0.86
8	53	5	0.09	0.91	0.92	0.81
9	44	5	0.10	0.92	0.93	0.84
10	39	4	0.09	0.98	0.98	0.88
11	32	4	0.11	0.92	0.93	0.84
12	29	3	0.09	1.00	1.00	0.89
13	24	5	0.17	1.00	1.00	0.91
14	21	1	0.05	0.92	0.93	0.76
15	20	0	0.00	0.95	0.95	0.91
16	17	2	0.11	0.95	0.95	0.95
17	17	0	0.00	1.00	1.00	0.89
18	13	3	0.19	0.94	0.94	0.94
19	11	2	0.15	1.00	1.00	0.81
20	10	1	0.09	1.00	1.00	0.85

Notes: The sample starts with 907 disjoint groups of robbers performing a robbery. Of these robberies 136 are cleared immediately (15 percent). Based on the remaining 771 groups given that 244 perform a second robbery, the recurrence rate is 32 percent. Depending on what one assumes about the recurrence of the 136 groups who were arrested after the first robbery, one can compute upper and lower bounds of the recurrence rate.

Table A4: Robustness Checks of Difference-in-Differences Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Sample</i>			The robbery has been cleared (0/1)			
			<i>With CCTV</i>	<i>Loot above</i>	<i>No pharmacies</i>	<i>First daily</i>
<i>Model</i>	<i>Probit</i>	<i>Only 2009</i>	<i>coverage</i>	<i>average</i>		<i>robbery</i>
			<i>DinD</i>	<i>DinD</i>	<i>DinD</i>	<i>DinD</i>
Police Intervention	0.016 (0.030)	-0.011 (0.049)	-0.010 (0.050)	0.002 (0.035)	0.010 (0.040)	0.003 (0.039)
Subsequent robberies	-0.101** (0.045)	-0.079 (0.048)	-0.074 (0.046)	-0.032 (0.037)	-0.111*** (0.041)	-0.084** (0.037)
Police Intervention × Subsequent robberies	0.106** (0.053)	0.131** (0.058)	0.114* (0.058)	0.081* (0.045)	0.110** (0.050)	0.095* (0.049)
Observations	1,255	592	702	642	787	746
R-squared		0.015	0.013	0.011	0.010	0.011

Notes: Columns 2 to 5 show linear probability model estimates based on the specification used in Table 4, while Column 1 shows the corresponding *probit* model marginal effects. Columns 3 to 6 restrict the analysis to the following subset of robberies: where CCTV cameras are definitely available (banks, postal offices, pharmacies, and jewelleries), where the haul is above average, where no pharmacy was targeted (the mode target), and where the robbery represents the first robbery of the day. Robust clustered (by sequence) standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. All regressions control for year effects.

Table A5: Balance Test

	Polizia		Carabinieri		Polizia-Carabinieri	
	Average	SE	Average	SE	Average	SE
Cleared robbery	0.149	0.012	0.093	0.017	0.056	0.020***
Number of the sequence	4.203	0.464	4.202	0.725	0.001	0.524
Days between subsequent robberies	14.978	2.271	12.887	2.542	2.091	3.013
Subsequent robberies	0.550	0.031	0.500	0.041	0.050	0.036
Shift change	0.160	0.014	0.146	0.020	0.014	0.023
Shift	3.055	0.041	2.963	0.056	0.092	0.057
North-Western area	0.375	0.027	0.376	0.037	-0.001	0.033
North-Eastern area	0.188	0.019	0.199	0.027	-0.011	0.026
Year	2008.450	0.036	2008.534	0.042	-0.084	0.035**
Month	6.151	0.246	6.351	0.315	-0.200	0.261
Day of the month	15.868	0.363	15.357	0.492	0.511	0.585
Sunday	0.054	0.008	0.071	0.014	-0.018	0.015
Monday	0.163	0.013	0.233	0.023	-0.070	0.025***
Tuesday	0.159	0.012	0.137	0.019	0.022	0.023
Wednesday	0.143	0.011	0.155	0.019	-0.013	0.022
Thursday	0.189	0.013	0.127	0.018	0.061	0.022***
Friday	0.167	0.013	0.149	0.024	0.018	0.028
Saturday	0.126	0.012	0.127	0.020	-0.001	0.023
Daylight	0.564	0.025	0.602	0.033	-0.039	0.034
Average age	26.080	0.655	26.308	1.205	-0.229	1.132
Amount stolen in euros	1.921	0.264	2.664	0.591	-0.743	0.536
Firearm	0.198	0.023	0.233	0.033	-0.035	0.029
At least one knife, but no firearm	0.084	0.016	0.090	0.020	-0.006	0.018
Some Italian	0.778	0.023	0.752	0.033	0.027	0.030
Different nationalities	0.120	0.011	0.112	0.019	0.008	0.020
Number of robbers	1.514	0.038	1.516	0.052	-0.001	0.046
Pharmacy	0.356	0.036	0.422	0.043	-0.067	0.034**
Other business	0.160	0.017	0.118	0.020	0.042	0.023*
Supermarket	0.152	0.021	0.165	0.025	-0.012	0.026
Bank	0.073	0.019	0.096	0.022	-0.023	0.021
Video rental	0.033	0.014	0.053	0.018	-0.020	0.013
Tobacco shop	0.025	0.007	0.016	0.010	0.009	0.010

Notes: Years 2008 and 2009. Standard errors are clustered by sequence. For the last two columns only: *** p<0.01, ** p<0.05, * p<0.1.

Table A6: Balance Test

	1st Strategy		2nd Strategy	
	Average	SE	Average	SE
Cleared robbery	0.077	0.040*	0.112	0.035***
North-Western area	0.010	0.067	0.004	0.069
North-Eastern area	-0.002	0.049	-0.060	0.064
Year	-0.065	0.067	0.057	0.163
Month	0.736	0.500	-0.012	0.609
Day of the month	0.921	1.113	1.474	1.204
Sunday	0.009	0.033	0.009	0.040
Monday	0.041	0.053	-0.166	0.066**
Tuesday	0.011	0.046	0.029	0.038
Wednesday	-0.002	0.044	0.027	0.042
Thursday	-0.047	0.043	0.028	0.047
Friday	-0.054	0.054	0.061	0.039
Saturday	0.042	0.048	0.013	0.046
Daylight	0.057	0.064	0.067	0.073
Average age	-2.115	1.927	-0.280	1.412
Amount stolen in euros	0.620	1.343	0.070	0.747
Firearm 0/1	0.021	0.053	0.019	0.044
At least one knife, but no firearm	0.007	0.034	-0.002	0.064
Some Italian	0.027	0.059	-0.060	0.037
Different nationalities	0.034	0.040	0.046	0.040
Number of robbers	0.044	0.089	0.063	0.101
Pharmacy	0.008	0.067	0.204	0.070***
Other business	-0.053	0.042	-0.098	0.048**
Supermarket	-0.011	0.056	-0.081	0.082
Bank	0.064	0.041	-0.011	0.032
Video rental	-0.038	0.028	0.027	0.022
Tobacco shop	0.012	0.014	-0.018	0.026
Joint difference (p-value)	0.72		0.17	

Notes: The 1st Strategy coefficient is the difference-in-differences between *Polizia* and *Carabinieri* for subsequent and first event. The 2nd Strategy coefficient is the difference between the robberies investigated by the *Polizia* that happen one day later and those that happen on the same day. The p-value for the joint test is based on a Wald test of a joint significance of all difference-in-differences using a seemingly unrelated regression. Standard errors are clustered by sequence. For the last two columns only: *** p<0.01, ** p<0.05, * p<0.1.

Table A7: Selection and Simple Difference Estimates

	(1)	(2)	(3)	(4)
	The robbery has been cleared (0/1)			
<i>Polizia</i> Intervention	0.087*** (0.028)	0.090*** (0.025)	0.095** (0.038)	0.105*** (0.025)
<i>Polizia</i> Intervention in $t - 1$	0.028 (0.033)	0.035 (0.028)	0.041 (0.036)	
<i>Polizia</i> Intervention in $t - 2$	-0.017 (0.032)			
<i>Polizia</i> Intervention in t and $t - 1$			-0.007 (0.051)	
Same force in t and $t - 1$				-0.012 (0.029)
Observations	529	666	666	674
R-squared	0.021	0.023	0.023	0.020

Notes: Linear probability model of clearances with simple differences between *Polizia* and *Carabinieri*, including lagged assignments. The analysis is restricted to subsequent robberies, where one can control for lagged “Polizia” intervention. The coefficients on the lags represent simple selection tests. All regressions control for year fixed effects. Robust clustered (by sequence) standard errors in parentheses: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.