University of Virginia Department of Economics

Facing the Facts: The Efficacy of Police Facial Recognition Technology

Davis Taliaferro

April 28, 2025

Abstract: In this paper, I explore the efficacy of facial recognition technology (FRT) in policing. Facial recognition technology has potential to act as a tool to both find suspects and deter crime by increasing the risk of being caught by the police. However, it is also a technology that relies on large scale surveillance and which has the potential to exacerbate preexisting systems of injustice. Thus, the net impact of FRT on crime is unclear. Despite this, increasing numbers of police departments around the United States are implementing the technology. Using data on crime rates in Massachusetts and bans on facial recognition technology, I find that the impact of FRT is generally mixed, underscoring a lack of compelling evidence behind FRT implementation. Notably, there were some statistically and economically significant results with property damage increasing following bans and motor vehicle theft decreasing, but a cost benefit analysis suggests any associated benefits do not outweigh the societal costs of this technology.

Acknowledgments: I would like to thank Professor John Pepper for advising this Distinguished Majors Program thesis. His feedback and guidance were invaluable, and I have learned so much through his commitment to my education and this project. I would also like to thank Professor Kerem Coşar for directing the DMP this past year. This program has been a key part of my undergraduate experience. I would also like to thank my DMP cohort who were an amazing source of motivation and laughs during this extremely, extremely busy semester. Lastly, I would like to thank all my friends, family, and professors who made my undergraduate career truly special and inspired me to dive so deep into the world of research.

Contents

1 Introduction

For decades, economists have studied crime. Gary Becker first formalized models of crime as the behavior of rational agents who weigh the risks of a crime with its benefits (1968). Since then, crime and the tools used to fight it have developed significantly. One recent shift is the adoption of facial recognition technology (FRT) by police forces.

Although FRT is relatively new, Garvie et al found that as of October 18, 2016, approximately one in two Americans lived within the jurisdiction of FRT-using law enforcement (2016). FRT works by comparing pre-existing surveillance footage to a large dataset of faces, returning a list of the strongest matches. These matches are tied to individual IDs or records, allowing users to identity and locate people. In policing contexts, this technology is used to help find suspects and locate missing people. For this paper, I focus solely on the former application.

Becker-type logic would suggest that FRT impacts crime by increasing the risk of committing a crime through a heightened likelihood of being caught. A rational actor considering a criminal act would internalize that increased risk and be less likely to commit said act.

With this theoretical behavior, FRT has the potential to fulfill two goals for the police: increasing arrests and improving public safety. Police departments tend to emphasize the suspect-finding power of FRT, (NYPD, 2023), but not on the question of overall crime rates. In this paper, I examine that latter goal closer and ask the question: does FRT reduce or increase crime rates?[1]

[1] An important disclaimer before digging deeper into the study of crime is that crime is socially constructed. Resultantly, a decrease in crime does not inherently mean people are becoming safer. Many acts, deemed moral by some people (or even many people), are criminalized in today's America. Simply because they are criminalized does not make them immoral. In this study, I look at property crime and violent crimes. For the sake of property crime, consider vandalism. Some people appreciate a lot of the art created by vandalism, and styles in vandalism have informed a lot of other art forms. Other crimes are more clearly immoral in many cases, including many cases of violent crimes. Regardless, looking at flatout levels of crimes oversimplifies the complexity of crime. For this reason, the decrease of crime is not necessarily tied to an increase in public safety or societal well being.

Still, this paper focuses on changes in crime as reported by police departments. This is because that is the goal of the police–the agents choosing to adopt this technology. Understanding how the police's motivating goals as they are measured by the police are impacted by that choice is deeply important in understanding whether this technology is justified.

This question is particularly important because the technology is deeply controversial. FRT relies on large scale tracking of biometric information, which is a security risk for any included individuals[2]. Further, FRT can be extremely biased, with one study showing that the rate at which FRT misidentifies darker-skinned women is 30%[3] higher than lighter-skinned men, working to further exaggerate pre-existing biases in the criminal justice system (Buolamwini & Gebru, 2018). And, to back up these claims, there have been some high-profile controversies regarding FRT's use.

In 2020, a Black man in Michigan was wrongfully arrested after FRT mistakenly identified him as a suspect in a robbery case, who happened to look nothing like him (Allyn, 2020). This garnered national attention, as other mistaken arrests have since, and contributed to a strong opposition of the technology (Allyn, 2020).

This strong opposition becomes apparent in some survey data. Rainie et al surveyed adults in 2022, finding that 27% of adults think that police use of FRT is a bad idea and 46% think it is a good idea, and Bragias et al found that a majority of online discussions about FRT were negative in 2021 (2022; 2021). And, numerous political organizations have sprung up to oppose its existence and use (Ozer et al, 2021).

Even when ignoring concerns over the accuracy of the technology, there are still a number of other negative externalities. For one, bad interactions with police can limit citizen involvement with other surveilling industries, such as hospitals or educational systems (Brayne, 2014). Further, in regards to FRT specifically, Beraja et al find that its implementation can limit "political unrest"[4] (2021). If FRT can substantially change behavior, then it may dampen economic conditions, providing a conducive environment for economically-motivated increases in crime.

Despite the controversy of this technology, little-to-no empirical studies of its efficacy

[2] In 2024, an Australian facial recognition firm had a large scale data leak (Pearson, 2024). This data can be used for a variety of nefarious reasons ranging from identity theft to stalking.

[3] Darker skinned women have an error rate of 34.7% while the maximum error for lighter skinned men is 0.8% (Buolamwini & Gebru, 2018)

[4] "Political unrest" includes standard political protests, meaning a decrease is likely a negative.

have been conducted. Instead, police departments rely on anecdotal evidence of when FRT is useful in a criminal case (NYPD, 2023).

In this paper, I work to expand our understanding of facial recognition by focusing on the impact of facial recognition bans on crime. Police tend to be vague or intentionally deceitful to the public about their use of facial recognition technology (Jarmanning, 2022). In Boston, for example, when a FRT ban was enacted the city claimed it was never used before; however, the ACLU revealed documents through the Freedom of Information Act that detailed numerous instances of FRT use by Boston detectives (ACLU of Massachusetts (ACLU), 2019). In many cases, police have no mandate to disclose FRT use, and at some police departments, leaders are not even aware of its use themselves (Jany, 2022). In this unclear landscape, bans provide clear and precise policy shifts.

In this paper, I look primarily at bans within Massachusetts municipalities. Massachusetts had 8 municipal bans within the state from June 2019 to December 2021, which is by far the most per state (Fight for the Future, n.d.). Further, due to the ACLU documents mentioned above, there is evidence of access and use of the technology prior to the bans (ACLU, 2019). This data is paired with crime data from the National Incident Based Reporting System which includes detailed information at the incident level, allowing us to study crime before and after bans are enacted. This is further complemented with a before and after analysis of crime in 14 cities nationally that have banned the technology.

With the data that is available, difference-in-difference and before and after analyses uncover a relationship between FRT and crime. Most analysis on individual crimes showed no statistically significant relationship. However, property destruction tended to show a positive and significant (or near significant) relationship across studies (implying FRT decreases property destruction). As will be discussed later in the paper, this is reasonable given the theoretically, more risk-reactive nature of property crime. And, there is some evidence that certain crimes may decrease in reaction to a ban. The paper concludes with a discussion of these results, showing how these findings underscore the misapplication of this technology and explore further research options.

The remainder of this paper is as follows. The remainder of Section 1 describes how the technology functions. Section 2 discusses the relevant literature. Section 3 details the data and empirical methods. Section 4 discusses the results. Section 5 analyzes those results and contextualize them within the relevant policy landscape. Section 6 closes the paper with a brief conclusion.

1.1 How the technology works

Facial recognition as a technology has become deeply entrenched in society. Our phones unlock at the sight of our face. Airport security now validate our IDs not with their own eyes but with FRT-enabled cameras. It is beginning to be used in malls, sports venues, and a variety of other public spaces (Mall of America, n.d; Tenbarge, 2024). And important to this paper, police use the technology to identify suspects and locate missing people.

Facial recognition technology is a software rather than a physical invention. Necessary for its function are databases of photos tied to known identities. These can come from state DMVs (Department of Motor Vehicle) which possess the ID photos from every ID-holding citizen. At the same time, these databases may be collected from scraping publicly available data online[5] (Scarcella, 2025).

Some public organizations, including police departments and state DMVs have their own facial recognition algorithms tied to their databases. Additionally, private companies will sell access to their databases and an algorithm to match new photos within the database. In either case, when a police officer wants to find the identity of a person based on footage, they utilize one of these databases and their associated software to try and find a match. For example, they may upload some frames from surveillance footage or photos from a witness's phone, and the software would return similar images and associated people.

Figure 1 contains an example of what one of these experiences may look like. The

[5] Clearview AI, a giant in the space of private facial recognition, has been found liable in a class action for this process (Scarcella, 2025).

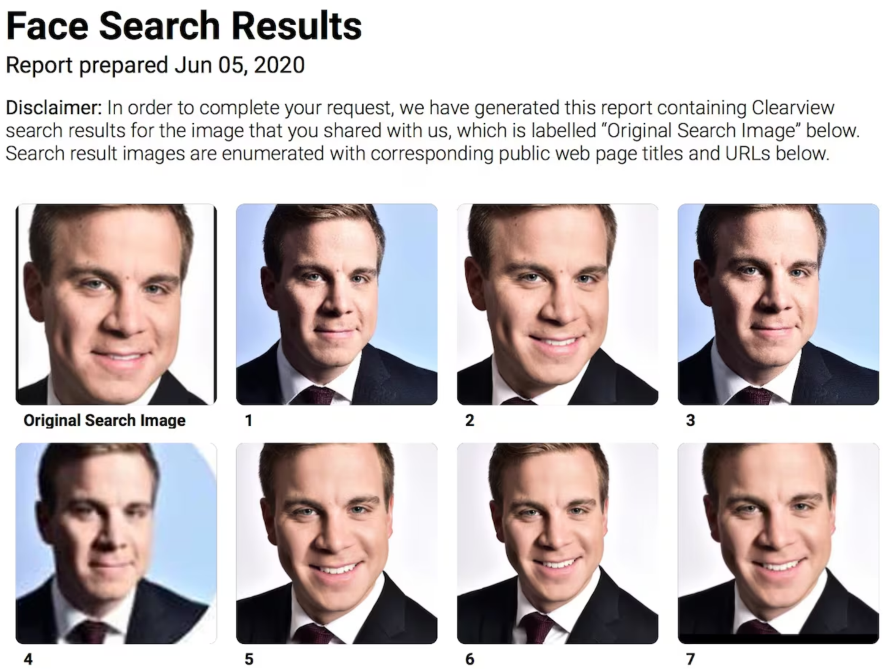

Figure 1: Clearview AI Search Results (Daigle, 2020)

included photo is taken from Clearview AI – a private company that scrapes the internet to create a database of people's image and personal information. Law enforcement can license out the software, which would allow them to submit a search image like in the photo (captioned "Original Search Engine"). The other photos numbered 1-7 are photos from Clearview AI's database, selected using their facial recognition algorithm. This example is for a news publication so it is somewhat unreasonable in that the submitted photos are nearly identical to the ones that are returned. A more reasonable example would likely show photos in multiple different contexts (in different settings, with the person wearing different clothes, different lighting, etc.). Additionally, in a real example any information relevant to the subject of photos 1-7 (such as their name, age, address, etc) would be available.

So, FRT works to enhance pre-existing surveillance devices. In theory, CCTV cameras, body cameras, and even cell phones become more powerful crime fighting tools with the addition of facial recognition software.

A last but important fact about facial recognition algorithms is that they are statistical. They are "trained" on vast amounts of images, continually trying to predict whether or not a face is the same as a face in another photo. Depending on if it is wrong or right and how confident it was, it shifts its parameters, gradually improving its ability to match new photos of a person to other photos of that person. Because this data is trained on pre-existing data, any new data is exactly that: new. So if a FRT algorithm is trained on a non-representative dataset (say it includes a disproportionately low amount of people of color), then bias can arise which impacts policing outcomes when applied in the real world. This is discussed in more detail in Section 2.1.

2 Relevant Literature

While little focus has been placed on FRT in the literature, there have been a variety of studies exploring related crime deterrents employed by police. While each are distinct, they carry potential relevant takeaways.

Di Tella and Schargrodsky (2004) found that police presence has a large deterrent effect on observed crime. In particular, they find that "police protection[6] induces a decline in auto theft of approximately 75 percent." They show this causally through a natural experiment of somewhat randomly assigned police following a terrorist attack. Lin (2009) corroborates this using random variation in state funding as an instrument for police size. Police presence, as primary law enforcers, likely serves as a gold standard for deterrent policing practices. Arrest can be immediate given a crime is committed in front of an officer and police increasingly also capture surveillance footage through body cameras. In many ways active police presence could be interpreted as an upper limit in terms of immediate crime deterrents [7]

The most similar question asked in the literature is whether or not surveillance cameras alone deter crime. The results generally show a negative relationship between cameras and crime. However, some papers find that the effect is small or nonexistent in certain

[6] police protection refers to the presence of police. For example, one can imagine a police stationed on a corner to sit and watch an intersection all day.

[7] "Immediate crime deterrents" as opposed to "crime deterrents" more generally is an important distinction here as some scholars see more potential in more systemic solutions, such as UBI or cheaper education, which target the roots of crime as better deterrents.

circumstances. Jung and Wheeler (2019) study the clearance[8] of cases given the existence of CCTV cases. They find that, in general, CCTV cameras do have an impact in increasing the clearance of cases, although not necessarily large enough to justify their cost. Welsh and Farrington (2007) conducted a meta-analysis on the studies of surveillance cameras. Their "results suggest that CCTV caused a modest (16%) but significant decrease in crime in experimental areas compared with control areas." But these results tended to be limited to specific areas and in most public areas they found small and insignificant results. And Gomez-Cardona et al (2017) find that camera presence carries an associated decrease in crime and arrests.

FRT modifies camera productivity at solving crimes, which could alter these already mixed results. Nonetheless, cameras are a common tool for crime deterring and crime solving, suggesting that they are popularly thought of as productive crime deterrents. The fact that there is this potential misalignment between the evidence and this popular belief, suggests that similarly mixed results for FRT is not unreasonable. Regardless, surveillance cameras without FRT, provide a useful basis of comparison for our final results.

This study extends the literature on surveillance by incorporating a systematic study of facial recognition technology. This furthers the study of surveillance cameras and surveillance more generally, as this technology only functions given pre-existing, large scale collection of data. Lastly, this study furthers the study of policing by underscoring that at least some policing decisions are not evidence-based. In particular, it adds more context to help us understand how police choose to adopt certain technologies and whether or not they use evidence-based reasoning to make decisions.

2.1 The Societal Costs of FRT

Facial recognition technology has several negative externalities that have been highlighted by both researchers and activists.

[8] Clearances are what happens within the justice system after a criminal incident. Essentially, clearances are whether or not a crime is solved.

Firstly, FRT has the potential to be biased depending on the identity of the suspect. Preexisiting research largely focuses on the race, color, and gender of people. One study found that in some facial recognition algorithms, darker skinned females were misidentified at a rate of 38% while lighter skinned males were misidentified at a rate of just 0.8% (Buolamwini and Gebru, 2018). These errors mean that women and people of color are more likely to be falsely targeted by the justice system than white people and men when this tech is employed. With preexisting systems of bias against these groups within the justice systems, this bias in facial recognition can exacerbate harm.

Even when unbiased, the possibility for false identification is a cost in and of itself. There have been multiple reports of misuse of the technology. For instance, NYPD police have run celebrities against FRT because they happen to look like a suspect (Emerson, 2019). Other police departments have been caught using DNA to predictively generate a face, and then use FRT on that – an approach that is not backed by science (Collings and Guariglia, 2024). And, large police departments don't always track this usage, which could allow for abuses of this technology (Jany, 2022). By using potentially faulty tech, these cities risk costly lawsuits, and when mistakes do happen, society learns to trust the government less and less, another potentially high cost impact of the technology.

Additionally, for FRT to be an effective tool to police departments, it must utilize large datasets of our personal information. Collection of this data is more than a picture and a name. It is a complex set of data points that hold the information of people's facial structure. That same data can be used to track people's location using public cameras, unlock face ID scanners, and more easily steal someone's identity. There have been several large data breaches of FRT data, such as in May 2024 when an Australian firm that used facial recognition at bars and clubs had a data leak (Pearson, 2024). With a limited amount of regulations globally on the storage and transfer of facial recognition data, such data leaks are bound to continue (Garvie et al, 2016). Resultantly, the associated collection of people's facial data for FRT continues to be dangerous.

Lastly, facial recognition technology furthers the existence of our surveillance-based economy. Increasingly, every action a person makes has the potential to be recorded. This is thanks to more and more cameras, drones, and mobile phones. With FRT, those recordings can instantly be matched to a unique person. Calculating the economic impacts of surveillance is certainly difficult; however, the vast majority of Americans value less surveillance suggesting the cost of surveillance is high (Rainie & Madden, 2015).

2.2 Theoretical Background

For this paper, I look at some high level theoretical underpinnings relevant to facial recognition technology in policing[9]. Becker considered crimes as rational acts that are committed by rational agents who balance the risks of the crime with the potential reward (1968). They maximize potential return. With the introduction of FRT, this paper assumes that the expected cost of a crime gets higher as the likelihood of getting caught increases. Under this model, a removal of FRT undeniably increases propensity for crimes to be committed.

But, we can expand this model by considering the reward of crimes. We know that the marginal utility of money someone has decreases as the amount increases. If FRT has the potential to destabilize society in any meaningful way that contributes to economic decline, the expected utility of a crime increases.

Further, I can complicate these results by considering a potential substitution between crimes. Katya discusses this behavior in his 1997 paper, stating that "the criminal law can be seen as setting the prices for crimes, and these prices may cause substitution." Some crimes tend to be impacted by FRT more than others. For instance, property crimes may be more exposed. These crimes often happen outside (ex. vandalism) or inside a stranger's house (ex. burglary), making them more open to surveillance. On the other hand, crimes that may occur inside and within ones own household or building would be less exposed to

[9]This paper does not dive particularly deep into the theoretical underpinnings of this topic. This is largely due to the fact that multiple theoretical assumptions easily push the impact of FRT in one direction or the other as discussed below. Additionally, rooting any theory in Becker-type logic creates a relatively simple model of risk versus reward. Therefore, for the sake of this is paper I consider the theory at at high level.

facial recognition technology. When it comes to a technology like FRT that impacts crimes differently, we may see the substitution that Katya mentions take place.

3 Data and Methods

3.1 Data

The data for this study is taken from the National Incident Based Reporting System (NIBRS) (U.S. Bureau of Justice Statistics; 2022, 2023, 2024). NIBRS contains incident-level data, which allows for a highly granular level of analysis. For the sake of this study, I aggregate to the city level of incidents. In particular, I look at the years from 2017 to 2023. In some analysis, I truncate the data from 2017 to March of 2020 to ensure the pandemic is not substantially impacting results. For the Massachusetts analysis, 293 cities are considered[10]. 60 or so cities were excluded solely because they did not report crimes in the NIBRS system and therefore did not have data. A list of the Massachusetts cities is included in Appendix A. And for the nationwide bans analysis, 14 cities are considered. These cities were the total list of cities that had bans and also had NIBRS data available for analysis.

The NIBRS database consists only of individual incidents and their characteristics. So, I collapsed the data by city and date to create crime counts, which were then scaled by population and multiplied by 100,000 (the crime rate units are "instances per 100,000 people"). For any periods without an incident of a particular type of crime, zero-counts were entered for that period. When looking at the specific crimes of focus and the particular cities studied, the number of zeroes was minimal when aggregated into large time periods[11]. When daily totals of crimes were used, zeros were more common. This is to be expected as crime is a relatively rare activity. Multiple periods were used to check for robustness, including daily,

[10] In total, there are 351 cities and towns in Massachusetts, so this includes the vast majority of municipalities (Mass.gov, 2025).

[11] For instance, when looking at yearly property damage, there were 2,065 city-year entries before inputting zeroes. Only 196 city-year entries had to be added with zeroes. With the amount of small towns in Massachusetts, this amount of zero-counts is not unreasonable, and likely did not impact results.

monthly, quarterly, and yearly periods.

An additional and important fact about NIBRS data is that it is not a universal system. Instead, individual states and localities opt into the use of the program. For that reason, this paper focuses mainly on Massachusetts which has almost all departments using NIBRS during this time period, meaning I have data of incidents from most departments. In the Before and After analysis, I do look at some states with lower coverage of NIBRS (i.e. there are more departments that don't use the reporting system); however, since I focus on cities that do use NIBRS reporting, this is less concerning, and there was reporting (at least some coverage) within each city during the time period studied.

3.2 Massachusetts Data

In a difference-in-difference setup, I look at data from Massachusets and compare cities with facial recognition bans and without facial recognition bans. Massachusetts data was selected for a variety of reasons. Firstly, Massachusetts has the most cities with facial recognition bans at a total of 8. Additionally, facial recognition technology is and has been readily available to any police department in the state. The Massachusetts Department of Transportation takes requests for processing photos through their facial recognition database of Massachusetts ID holders (ACLU, 2019).

Additionally, through the Freedom of Information Act, the ACLU of Massachusetts was able to uncover a database of email requests, which include requester names and emails (ACLU, 2019). Through this, I was able to validate that most cities with bans did, in fact, have at least some officers utilize the technology prior to any bans. Table 1 includes some of the cities with FRT usage and cities without FRT usage. It shows, for instance, that Boston, Brookline, Cambridge, and Somerville, which all banned facial recognition technology did have some record of using it before their bans. This ACLU reporting was important as some cities, including Boston, banned facial recognition technology while proclaiming it was never utilized (Jarmanning, 2020). This FOIA request proves the contrary, suggesting that facial recognition bans did serve as meaningful policy shifts.

Some FRT	No Reported FRT
Brookline	Northampton
Revere	Easthampton
Boston	Springfield
Cambridge	Worcester
Somerville	Everett
Chelsea	Watertown
Newton	Quincy
Dedham	Milton
Needham	Winthrop

Table 1: Selected cities sorted by DMV FR usage according to ACLU (2019)

To collect the ban dates, I began by referencing banfacialrecognition.com – a site dedicated to tracking the use of and banning police use of facial recognition technology (Fight for the Future, n.d.). This site works through crowdsourcing bans of facial recognition and confirms them using news reports. To confirm the findings, I checked each news source for the cities I have included in the treatment group. Table 2 provides a list of the banning cities and their policy announcements.

City	Policy Announcement
Boston	June 24, 2020
Brookline	December 11, 2019
Cambridge	January 13, 2020
Easthampton	July 15, 2020
Northampton	December 19, 2019
Worcester	December 14, 2021
Springfield	February 4, 2021
Somerville	June 27, 2019

Table 2: Policy Announcement Dates by City

The sample statistics for this data is summarized in Table 3 to give a sense of the levels of each crime type as reported in NIBRS. The most common crime is simple assault at an average of 37.76 instances per month. The highest property crime is the "All Other Larceny" category, which is general theft not captured by the other property crime categories, and has an average of 30.48 instances per month. The next highest property crime is vandalism

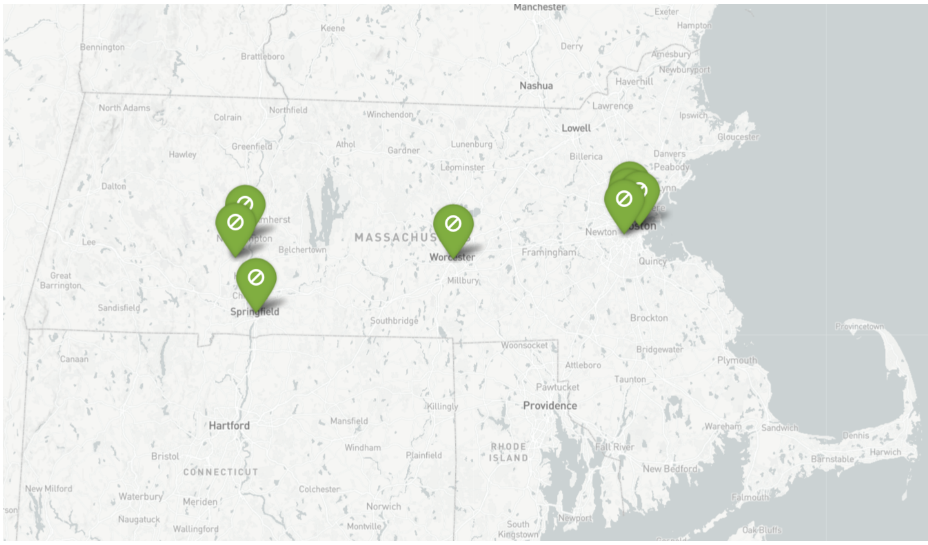

Figure 2: Locations of Facial Recognition Bans in Massachusetts (Fight for the Future, n.d.)

at 27.84 crimes per month. Each crime has a really high variance compared to the average, showing that crime varies highly from city to city. Table 4 gives a sense of how the sample statistics differ in cities with bans prior to and after bans are implemented. In general, crime decreases in all categories in the periods after bans. However, note that if crime rates were going down universally, this would be an expected result, which is why we turn to the causal analysis used in this paper.

Crime Category	**Monthly Mean***	**Standard Deviation**
Burglary/Breaking and Entering	13.33	13.58
Pocket-picking	0.65	1.51
Shoplifting	15.04	17.16
Theft From Building	7.77	9.04
Theft From Motor Vehicle	12.53	13.00
All Other Larceny	30.48	23.99
Destruction/Damage/Vandalism of Property	27.84	21.07
Motor Vehicle Theft	8.00	8.75
Robbery	3.30	4.83
Aggravated Assault	15.83	14.08
Simple Assault	37.76	32.00
n=293		
*** instances per 100,000**		

Table 3: Monthly Per Capita Crime Statistics in Massachusetts

Crime Category	**Pre-Treatment***	**Post-Treatment***
Burglary/Breaking and Entering	4.80	3.22
	(8.84)	(7.26)
Pocket-picking	0.23	0.07
	(0.48)	(0.18)
Shoplifting	5.34	4.94
	(10.71)	(10.30)
Theft From Building	8.03	7.48
	(11.96)	(7.63)
Theft From Motor Vehicle	7.20	5.17
	(14.17)	(12.60)
All Other Larceny	8.24	8.22
	(11.78)	(13.51)
Destruction/Damage/Vandalism of Property	10.13	7.20
	(18.04)	(15.06)
Motor Vehicle Theft	2.09	2.66
	(4.99)	(6.32)
Robbery	2.58	1.57
	(5.07)	(3.88)
Aggravated Assault	6.70	5.09
	(12.84)	(12.51)
Simple Assault	14.97	9.91
	(28.75)	(23.18)

n=8
***instances per 100,000**

Table 4: Pre-Treatment and Post-Treatment Crime Rates (Mean and Std. Dev.)

3.3 National Data

The second model of analysis is regression discontinuity in time. This setup required more cities with bans, so I expanded to several cities outside of Massachusetts. These cities are all cities that implemented full facial recognition bans. There are some important notes however. Firstly, King County, which contains Seattle, banned the technology for their police. In all other cities, there were bans at the city-level. However, Seattle contains an overwhelming majority of the population within King County, so although Seattle did not ban the tool for their police, they were included in the analysis because they were impacted by the county-wide ban. Additionally, some cities were left out if they did not have substantial NIBRS coverage in the state. The full list of cities and their ban dates can be seen in Table 5. These dates and cities come from the same website used for Massachusetts, banfacialrecognition.com (Fight for the Future, n.d.).

I compared these cities across the U.S. before and after their bans using the days before and after a ban as my time variable (with "0" as the date of treatment and days after being positive and days before being negative). The sample statistics for this group of cities is shown in Table 6. Crime tended to be more prevalent in these cities on average than in the Massachusetts-only analysis. For instance, there were approximately 8 more instances of simple assault per month, 15 more instances of the "All Other Larceny" category of crime. The nationwide analysis still had large variance across cities.

City	**Treatment Date**
Portland, OR	September 9, 2020
Seattle	June 1, 2021
Minneapolis	February 10, 2021
Madison	December 1, 2020
Portland, ME	August 3, 2020
Jackson	August 20, 2020
Baltimore	August 9, 2021
Boston	June 24, 2020
Brookline	December 18, 2019
Cambridge	June 13, 2020
Easthampton	July 1, 2020
Northampton	February 27, 2020
Springfield	February 24, 2020
Somerville	June 28, 2019
Worcester	December 14, 2021

Table 5: Treatment dates by city

3.4 Empirical Strategy

Two general empirical strategies were employed in this project: difference in difference and a before and after analysis. These two strategies were applied to check for the robustness of the results and to get a better understanding of how immediate the effects of a FRT ban may be.

Crime Category	**Daily Mean***	**Standard Deviation**
Burglary/Breaking and Entering	0.95	1.29
Pocket-picking	0.03	0.18
Shoplifting	0.72	1.03
Theft From Building	0.51	0.86
Theft From Motor Vehicle	1.25	1.83
All Other Larceny	1.51	1.74
Destruction/Damage/Vandalism of Property	1.43	1.62
Motor Vehicle Theft	1.00	1.54
Robbery	0.32	0.55
Aggravated Assault	0.71	0.95
Simple Assault	1.50	1.67
n=14		
*** instances per 100,000**		

Table 6: Monthly Per Capita Crime Statistics in Massachusetts

3.5 Difference In Difference

The DID appraoch compares cities with FRT bans to cities without FRT bans. It uses difference in difference to attempt to find a causal link between FRT and crime. Here is the estimating equation:

$$y_{i,t} = \beta_0 + \beta_1 \cdot treated + \beta_2 \cdot post + \beta_3(treated \cdot post) + \beta_3 t + \beta X + \epsilon$$

Here, i denotes the county and t denotes time. $y_{i,t}$ is our outcome variable, which is always a crime rate for a specific crime or set of crimes. *treated* represents the county being treated with FRT and *post* denotes a time indicator for whether or not the county has adopted FRT yet. X is a vector of covariates. In this particular study, the only covariate used is population. Given the data construction process, adding more covariates would be highly labor intensive, as there are not many city level datasets with relevant information. I leave the addition of more covariates for future research. Lastly, ϵ represents the unobserved determinants of y. I assume that $E[\epsilon|X] = 0$.

This approach also has an important characteristic in that it will naturally be biased in the direction of FRT decreasing crime. This is because rational actors move from counties

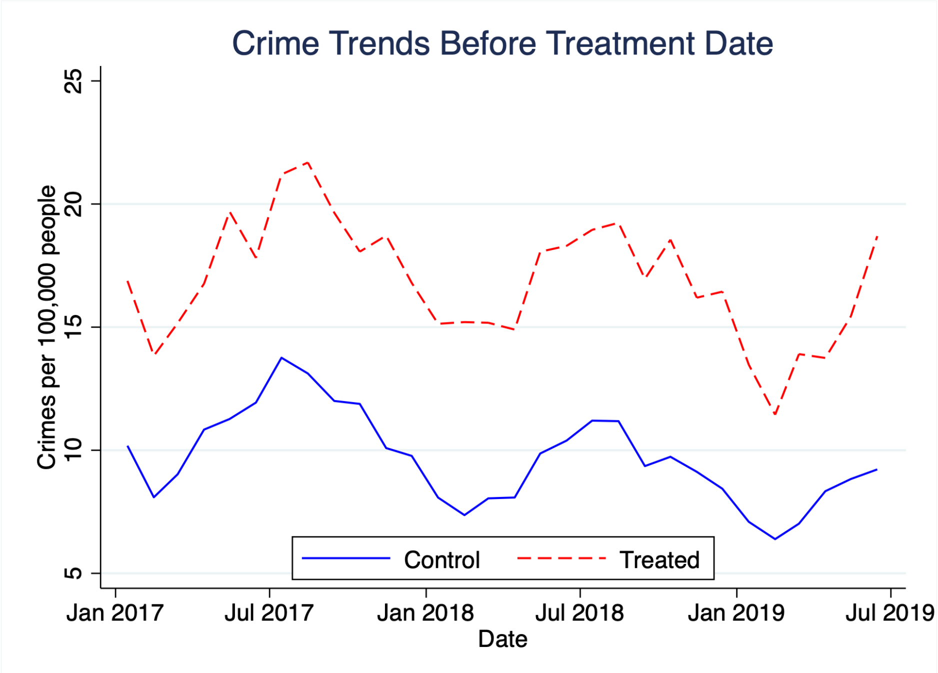

Figure 3: Crime trends leading up to the first ban

with FRT to bordering counties without FRT which would show an increase within treated counties but not across counties (due to an assumed decrease within nontreated counties). So given this bias, if I find little increase or no effect in cities with bans, that finding is further justified. Similarly, if I found strong positive results those may not necessarily reflect the imapct of FRT on crime overall but rather a movement of people to a place where crime is less risky.

Additionally, graphical analysis shows that the parallel trends assumption seems to hold for this crime data. Figure 3 shows this with plots of property crime over time comparing treated groups to untreated groups. As there were multiple treatment times, the graph tracks all property crime until the very first ban, which occured in Somerville, MA on June 27, 2019. The treated and untreated plots are relatively similar, highlighting the suitability of this data for a DID approach.

For a more complicated approach, I also employ the difference in difference with multiple treatment periods from Callaway and Sant'Anna (2021). This approach better accounts for the multiple treatment periods.

3.6 Before & After Analysis

For the Before & After (BA) analysis, I slice out the day on which the public is made aware of the city council's vote to ban facial recognition in policing. Then the post period is the following period of days and the pre period is the same number of days before the ban. For this study, I look at a variety of periods, including 1 day, 2 days, 3 days, 14 days, and 30 days. These varying period widths surrounding the ban allowed us to check for robustness. Moreover, the enactments were relatively disbursed throughout the year and days of week, making this approach more robust[12] The covariates were population size and the estimating equation was as follows:

$$crime\_per\_capita_{t,i} = \sigma + \delta t + \beta\{AfterBan\} + \beta X + \epsilon$$

Here, t is the days before (negative) and after (positive) the ban. $\{AfterBan\}$ is an indicator function for whether or not the date is in the second period or not. X is the covariates, which is currently constrained to population. As mentioned above, given the data constraints, the addition of other covariates will be left for future research. Additionally, as before, it is assumed that $E[\epsilon|X] = 0$.

The thinking behind this study is that the crime landscape would be very similar in the days leading up to and the days after a ban, meaning any changes would likely be attributable to the FRT ban.

[12] If policy enactments were all on the same day of the week, regular weekly cycles could impact the estimated effect of the policy.

4 Results

Each model was run for each type of crime. Most crimes showed an insignificant relationship to a ban but not all crimes. The results were somewhat mixed and sensitive to model specifications. However, the strongest causal approaches, which included the difference in difference and the before and after comparison with a one day window showed that there was a positive, significant relationship between the FRT ban and property destruction. Further, most signs were consistent across the 1 day BA approach and the DID approach.

Crime Type	**Coefficientˆ**	**95%**
All Property Crimes	2.023 (7.024)	[-11.800, 15.848]
Burglary/Breaking and Entering	-1.873 (1.516)	[-4.857, 1.110]
Pocket-picking	-0.262 (0.426)	[-1.101, 0.577]
Shoplifting	-1.485 (2.865)	[-7.124, 4.155]
Theft From Building	-0.916 (1.952)	[-4.757, 2.926]
Theft From Motor Vehicle	0.519 (2.148)	[-3.709, 4.747]
All Other Larceny	2.274 (5.115)	[-7.794, 12.342]
Destruction/Damage/Vandalism of Property	4.235* (2.523)	[-0.731, 9.201]
Motor Vehicle Theft	-0.449 (0.504)	[-1.441, 0.543]
Robbery	-0.019 (0.536)	[-1.073, 1.035]
Aggravated Assault	-1.070 (2.058)	[-5.121, 2.981]
Simple Assault	0.002 (2.650)	[-5.213, 5.217]

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
n=293
ˆ instances per 100,000

Table 7: Difference-in-Differences Estimates for Crime Outcomes

Number of Days Surrounding Ban	1	2	3	14	30
Burglary	0.238	0.508	0.316	0.017	0.169
	(0.253)	(0.52)	(0.306)	(0.183)	(0.105)
Pocket-Picking	-	-	-0.024	0.035	0.005
	-	-	(0.107)	(0.059)	(0.041)
Shoplifting	-0.187	-0.394	-0.016	-0.27*	-0.054
	(0.112)	(0.411)	(0.321)	(0.131)	(0.102)
Theft from Building	0.039	0.192	0.41	-0.294	-0.132
	(0.141)	(0.254)	(0.56)	(0.24)	(0.132)
Theft from Motor Vehicle	-2.526	-4.539	-3.595	-0.262	0.165
	(2.765)	(5.542)	(3.6)	(0.21)	(0.177)
All Other Larceny	0.325	0.811	1.105	-0.058	-0.279
	(0.416)	(0.785)	(0.89)	(0.285)	(0.299)
Destruction/Vandalism/Danger to Property	0.458	1.075	1.055*	-0.385	-0.379*
	(0.3)	(0.608)	(0.531)	(0.281)	(0.206)
Motor Vehicle Theft	-0.469*	-1.001*	-0.479	0.24	0.229
	(0.247)	(0.516)	(0.28)	(0.399)	(0.214)
Robbery	-0.036	-0.152	-0.099	0.194***	0.094***
	(0.11)	(0.23)	(0.199)	(0.063)	(0.029)
Aggravated Assault	-0.131	-0.043	-0.198	-0.293*	-0.146*
	(0.186)	(0.318)	(0.248)	(0.153)	(0.081)
Simple Assault	-0.214	-0.069	-0.088	-0.319	-0.212
	(0.579)	(1.14)	(0.812)	(0.383)	(0.215)

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Before And After Analysis

4.1 Results: Difference-in-Difference

Table 7 displays the results from the Massachusetts city DID. The results show that there is no significant effect of FRT on crime except for an increase in property destruction (titled "Destruction/Damage/Vandalism of Property"). This had an estimated increase of 4.235 instances per month. With the average of 10.13 instances per month, this is a substantial increase.

Generally, there are both positive and negative relationships shown when I consider the insignificant findings. A fairly close to significant relationship exists with motor vehicle thefts that shows a negative reaction to the FRT ban. Here the estimate sits around -0.449 crimes per month while the monthly average sits around 2.

Additionally, the results from the Callaway and Sant'Anna (2021) method were similarly ambigious with similar signs across crimes that only differ when there are high standard errors. It also shows a positive trend for property crime, although it is not as statistically significant. See Appendix B for a table of results.

4.2 Results: Before & After Analysis

Table 8 displays the results from the city regression discontinuity. Notably, the results for the 1, 2, and 3 day setups all show the same signs and they share most signs with the DID approach. When there is a difference in signs, it is always true that one approach's estimate or the other (or both) have a large standard error.

The surprising results come when I expand the interval before and after the ban in the BA approach For 14 and 30 days, I find that robbery and burglary increases following a ban. I also find that property destruction, shoplifting, and aggravated assault decrease. The coefficient signs of aggravated assault, shoplifting, and are consistent across period choice, while the sign of robbery and property destruction flip. These results are interesting, but given the longer time window, the larger period results should be considered cautiously. Furthermore, the results from this analysis are largely statistically insignificant. Resultantly, although this behavior is interesting and should be explored further, this approach is too imprecise to draw general conclusions at the moment.

5 Discussion

Current models and data were not able to reveal substantial statistically significant estimates for most crimes. These results are reasonable given similar results with CCTV cameras; however, the lack of significance is most likely due to the lack of data. Only a small number of cities have implemented bans at this point in time. However, this also underscores an important yet not inherently clear finding of this study: there is little evidence supporting the theory that the implementation of this technology deters crime, and the many police departments who have implemented this technology have no basis for believing it improves public safety.

Police often rely on anecdotal evidence–times that using FRT has identified potential suspects or relevant people (NYPD, 2023). However, the mere fact that this is helpful in finding suspects does not necessarily support the idea that FRT reduces crime. If it did, we would see a significant shift in the number of crimes committed following the banning of the technology. Rather, I only find some consistent evidence of increases in highly specific crimes, property damage, and some associated decreases in other types of crimes following bans.

There are significant potential costs to widespread use of this technology ranging from distrust in the government to privacy violations. The fact that there is not evidence shown by police and not substantial evidence to be found by researchers suggests that police are not properly internalizing these costs if their goal is to decrease crime. It does not necessarily even align with their goal of making money for the city, which comes largely from arrests that provide funding from processing fees.

5.1 A Preliminary Cost Benefit Analysis on FRT

As mentioned in the results, I see an increase in property damage following the banning of facial recognition, which suggests that FRT limits property damage. While these crimes carry inherent monetary costs to the victims, the societal cost of facial recognition technology likely far outweighs the cost of most property crimes.

For an anecdotal example, the man in Michigan who was wrongfully arrested following a false identification using FRT was awarded \$300,000 alongside a commitment to technology reforms within the police. Additionally, the court and legal fees of the plaintiff were covered. This process likely costs orders of magnitude more than a single act of property damage. Its difficult to find systematic estimates of property damage costs, however in 2019, the FBI did estimate the cost of a single act of burglary at \$2,661.00 (FBI, 2019). We can use this as an upper bound on property damage, which is likely less costly. The current estimate from this paper implies that facial recognition limits crime by 30% (or in levels, it implies a decrease of approximately 4.5 crimes per 100,000 people per month, or approximately 54 crimes a year). This implies an estiamted savings of \$122,094 per year in 2019 USD. In 2024, a man was awarded \$300,000 in a lawsuit after police falsely arrested him following FRT wrongly tagging him as a person in surveillance footage. Adding in the associated legal fees, if there is a 33% chance of a high level mistake regarding FRT, then the savings and costs of facial recognition technology are approximately equal. This is, of course, a non-systematic approach to a cost-benefit analysis, based on a singular legal case.

However, this setup also fails to internalize the broader societal costs of both facial recognition technology and crime. Though these are hard to quantify as mentioned earlier, the public clearly values non-surveillance. Incorporating this into the estimate would clearly make the expected costs of FRT far outweigh the benefits for police, and convincingly show that police certainly should not be implementing this technology given the current evidence available.

5.2 Explaining Negative Coefficients

It is difficult to explain the negative coefficients that arise in the longer periods of the BA model. In some senses, in tandem with the positive coefficients it suggests a reallocating of crime where people commit less of crimes like vandalism and assault which are not tied to any sort of payoff inherently and instead commit crimes like robbery and burglary that carry higher sentences and typically a higher likelihood of getting caught. This idea is cooperative with the idea of substitutability of crime mentioned in Section 2.2.

Although this is an interesting interpretation of this behavior, a more important note is that almost all of these results with the sign flips were insignificant. So we may not need to pay too much attention to this behavior unless future research reveals more statistically significant instances of it occurring.

5.3 Sign Flips in the BA Approach

The before and after analysis shows sign flips in several of the crimes as the time period expands. This is difficult to explain. Perhaps this is the result of some outside factors that confound results as you expand the time period. Hausman studied this type of setup in a 2017 paper and argued that there are number of issues regarding this approach, in particular when the time bandwidths are large. Large bandwidths allow for other trends to impact the results. The fact that the short term results (1, 2, 3 day setup) corroborate the results from the DID, which does incorporate longer time horizons, suggests that those results are more robust.

5.4 Why Property Damage?

This paper shows evidence that property damage is more sensitive to facial recognition technology. This is reasonable given that property damage has little reward and resultantly, the tradeoff calculation would be more sensitive to any change in risk. Similarly, property damage can be done as a political act, suggesting that changes in politics, such as an anti-surveillance move from a city council (like a FRT ban), may potentially shift behavior (Lai, 2020).

5.5 A potential source of insignificance

As discussed in the Introduction, police are not regularly mandated to publicize use of facial recognition technology. And, it is unclear whether the general public is aware of its usage. Documents from the ACLU revealed that certain police departments, such as Boston, had detectives using the technology and still maintained that the technology was never used (ACLU, 2019; Jany, 2022). This may be a a source of the insignificance between FRT and crime rates. Following a Becker-type logic, people need to be aware of FRT to properly introduce it to their risk analysis for comitting crimes. So if they are not aware of the technology's use, then the lack of information about the technology would imply limited-to-no relationship between the tech and crime rates under Becker type logic.

5.6 Further Research

In order to better understand this topic, there must be more data collection. Use of the technology tends to be highly individual, employed on a detective-by-detective basis. Police departments have access to the casework of these detectives and can compare performance of users before and after FRT bans. If this data were made public through a FOIA request, such research could be conducted.

Additionally, the NIBRS data includes some information regarding clearance. As higher clearance rates are the most immediate benefits in the eyes of police, confirming or invalidating that perception through further research would be a worthwhile endeavor.

Lastly, as more governments ban FRT and more data from NIBRS becomes available, the number of cities ready for comparison increases. This exact same project is ready to be repeated as soon as that data becomes available.

6 Conclusion

This study provides an early look into the effectiveness of FRT. This paper shows that while police have little evidence to cite in support of the technology, researchers also lack clarity as to the potential impacts of FRT on crime. The lack of strong positive results in tandem with the potentially monumental societal costs of FRT underscore that the implementation of this technology should be stopped until more evidence is available. At best, FRT is doing little to prevent crime. All the while, FRT's externalities are enacting significant and irreversible harms on society, including large scale unsecure collection of personal data, wrongful arrests, and an automated form of systemic racism. This paper shows that some property crimes may be somewhat sensitive to the technology, with associated increases in property damage and decreases in car theft following FRT bans. However, the overall conclusion is that most crime rates showed a statistically insignificant relationship with FRT bans. Undeniably, more research should be conducted, but this paper represents a first step in a new and important area of study.

References

[ACLU of Massachusetts, 2019] ACLU of Massachusetts (2019). The data for justice project: Briefcam at the city of boston.

[Allyn, 2020] Allyn, B. (2020). "the computer got it wrong": How facial recognition led to false arrest of black man. NPR.

[Becker, 1968] Becker, G. S. (1968). Crime and punishment: An economic approach. In *The Economic Dimensions of Crime*, pages 13–68.

[Benson et al., 1994] Benson, B. L., Kim, I., and Rasmussen, D. W. (1994). Estimating deterrence effects: A public choice perspective on the economics of crime literature. *Southern Economic Journal*, 61(1):161.

[Beraja et al., 2021] Beraja, M., Kao, A., Yang, D. Y., and Yuchtman, N. (2021). Ai-tocracy. Working Paper 29466, National Bureau of Economic Research.

[Bragias et al., 2021] Bragias, A., Hine, K., and Fleet, R. (2021). 'only in our best interest, right?': Public perceptions of police use of facial recognition technology. *Police Practice and Research*, 22(6):1637–1654.

[Brayne, 2014] Brayne, S. (2014). Surveillance and system avoidance: Criminal justice contact and institutional attachment. *American Sociological Review*, 79(3):367–391.

[Buolamwini and Gebru, 2018] Buolamwini, J. and Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 77–91.

[Callaway and Sant'Anna, 2021] Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230. Themed Issue: Treatment Effect 1.

[Clearview AI, 2020] Clearview AI (2020). Clearview ai provided cbc news reporter thomas daigle with these search results from their image database. CBC/Radio Canada.

[Collings and Guariglia, 2024] Collings, P. and Guariglia, M. (2024). Cops running dna-manufactured faces through face recognition is a tornado of bad ideas. Accessed: 2025-04-10.

[Daigle, 2020] Daigle, T. (2020). Clearview ai facial recognition offers to delete some faces - but not in canada. CBC News.

[Di Tella and Schargrodsky, 2004] Di Tella, R. and Schargrodsky, E. (2004). Do police reduce crime? estimates using the allocation of police forces after a terrorist attack. *American Economic Review*, 94(1):115–133.

[Emerson, 2019] Emerson, S. (2019). Police are feeding celebrity photos into facial recognition software to solve crimes. Accessed: 2025-04-10.

[FBI, 2019] FBI (2019). Burglary. https://ucr.fbi.gov/crime-in-the-u.s/2019/crime-in-the-u.s.-2019/topic-pages/burglary#:~:text=Victims%20of%20burglary%20offenses%20suffered,per%20burglary%20offense%20was%20%242%2C661. Accessed April 18, 2025.

[Fight for the Future,] Fight for the Future. Ban facial recognition. Accessed: 2025-04-10.

[Garvie et al., 2016] Garvie, C., Bedoya, A., and Frankle, J. (2016). The perpetual line-up.

[Gomez-Cardona et al., 2017] Gomez-Cardona, S., Mejia, D., and Tobon, S. (2017). The deterrent effect of public surveillance cameras on crime. *SSRN Electronic Journal*.

[Jany, 2022] Jany, L. (2022). Lapd doesn't fully track its use of facial recognition, report finds. Los Angeles Times.

[Jarmanning, 2020] Jarmanning, A. (2020). Boston bans use of facial recognition technology. it's the 2nd-largest city to do so. WBUR News.

[Jung and Wheeler, 2019] Jung, Y. and Wheeler, A. (2019). The effect of public surveillance cameras on crime clearance rates.

[Katyal, 1997] Katyal, N. K. (1997). Deterrence's difficulty. *Michigan Law Review*, 95(8):2385–2446.

[Klose, 2024] Klose, S. (2024). Re-defining evidence-based policing. *Policing: A Journal of Policy and Practice*, 18:paad095.

[Lai, 2020] Lai, T. W. (2020). Political vandalism as counter-speech: A defense of defacing and destroying tainted monuments. *European Journal of Philosophy*, 28(3):602–616.

[Lin, 2009] Lin, M.-J. (2009). More police, less crime: Evidence from us state data. *International Review of Law and Economics*, 29(2):73–80.

[Mall of America,] Mall of America. Facial recognition technology at mall of america for person of interest detection. Accessed: 2025-04-10.

[Mass.gov,] Mass.gov. Massachusetts city and town websites. Accessed: 2025-04-10.

[New York City Police Department, 2023] New York City Police Department (2023). Facial recognition: Impact and use policy. Accessed: 2025-04-10.

[Ozer et al., 2021] Ozer, N., Ruane, K., and Cagle, M. (2021). Grassroots activists are leading the fight to stop face recognition. it's time for congress to step up, too. ACLU.

[Pearson, 2024] Pearson, J. (2024). A face recognition firm that scans faces for bars got hacked—and that's just the start. Wired.

[Rainie et al., 2022] Rainie, L., Funk, C., Anderson, M., and Tyson, A. (2022). Public more likely to see facial recognition use by police as good, rather than bad for society. Pew Research Center.

[Rainie and Madden, 2015] Rainie, L. and Madden, M. (2015). Americans' views on government surveillance programs. https://www.pewresearch.org/internet/2015/03/16/americans-views-on-government-surveillance-programs/. Accessed: 2025-04-28.

[Scarcella, 2025] Scarcella, M. (2025). Us judge approves "novel" clearview ai class action settlement. Reuters.

[Tenbarge, 2024] Tenbarge, K. (2024). Facial recognition technology use at stadiums across u.s. sparks protests. NBC News.

[United States. Bureau of Justice Statistics, 2022a] United States. Bureau of Justice Statistics (2022a). National incident-based reporting system, 2017: Extract files.

[United States. Bureau of Justice Statistics, 2022b] United States. Bureau of Justice Statistics (2022b). National incident-based reporting system, 2018: Extract files.

[United States. Bureau of Justice Statistics, 2022c] United States. Bureau of Justice Statistics (2022c). National incident-based reporting system, 2019: Extract files.

[United States. Bureau of Justice Statistics, 2022d] United States. Bureau of Justice Statistics (2022d). National incident-based reporting system, 2020: Extract files.

[United States. Bureau of Justice Statistics, 2023] United States. Bureau of Justice Statistics (2023). National incident-based reporting system, 2021: Extract files.

[United States. Bureau of Justice Statistics, 2024] United States. Bureau of Justice Statistics (2024). National incident-based reporting system, 2023: Extract files.

[Welsh and Farrington, 2007] Welsh, B. C. and Farrington, D. P. (2007). Closed-circuit television surveillance. In *Preventing Crime*, pages 193–208.

A Appendix: List of cities for Massachusetts DID

Abington
Acton
Acushnet
Adams
Amesbury
Amherst
Andover
Aquinnah
Arlington
Ashburnham
Ashby
Ashfield
Ashland
Athol
Attleboro
Auburn
Avon
Ayer
Barre
Becket
Bedford
Belchertown
Bellingham
Belmont
Berkley
Berlin
Bernardston
Beverly
Billerica
Blackstone
Bolton
Boston
Bourne
Boxborough
Boxford
Boylston
Braintree
Brewster
Bridgewater
Brimfield
Brockton
Brookfield
Brookline
Burlington
Cambridge
Canton
Carlisle
Carver
Charlton
Chatham
Chelmsford
Chelsea
Cheshire
Chester
Chesterfield
Chicopee
Chilmark
Clarksburg
Clinton
Cohasset
Concord
Dalton
Danvers
Dartmouth
Dedham
Deerfield
Dighton
Douglas
Dover
Dracut
Dudley
Dunstable
Duxbury
East Bridgewater
East Brookfield
East Longmeadow
Eastham
Easthampton
Edgartown
Erving
Essex
Everett
Fairhaven
Fall River
Falmouth
Fitchburg
Framingham
Franklin
Gardner
Georgetown
Gill
Gloucester
Goshen
Grafton
Granby
Granville
Great Barrington
Greenfield
Groton
Groveland
Hadley
Halifax
Hamilton
Hampden

Hanover
Hanson
Harvard
Harwich
Hatfield
Haverhill
Hingham
Hinsdale
Holbrook
Holden
Holland
Holliston
Holyoke
Hopedale
Hopkinton
Hudson
Hull
Ipswich
Kingston
Lakeville
Lancaster
Lanesborough
Lawrence
Lee
Leicester
Lenox
Leominster
Leverett
Lexington
Lincoln
Littleton
Longmeadow
Lowell
Ludlow
Lunenburg
Lynn
Lynnfield
Malden
Mansfield
Marblehead
Marion
Marlborough
Marshfield
Mashpee
Mattapoisett
Maynard
Medfield
Medford
Medway
Melrose
Mendon
Merrimac
Methuen
Middleton
Milford
Millbury
Millis
Millville
Milton
Monson
Monterey
Nahant
Nantucket
Natick
Needham
New Bedford
New Braintree
Newbury
Newburyport
Norfolk
North Adams
North Andover
North Reading
Northampton
Northborough
Northfield
Norton
Norwell
Norwood
Oak Bluffs
Oakham
Orange
Orleans
Otis
Oxford
Palmer
Paxton
Peabody
Pelham
Pembroke
Pepperell
Pittsfield
Plainville
Plymouth
Plympton
Princeton
Provincetown
Quincy
Randolph
Raynham
Reading
Rehoboth
Revere
Rochester
Rockland
Rockport
Rowley
Royalston

Rutland
Salem
Salisbury
Sandwich
Saugus
Scituate
Seekonk
Sharon
Shelburne
Sherborn
Shirley
Shrewsbury
Shutesbury
Somerset
Somerville
South Hadley
Southampton
Southborough
Southbridge
Southwick
Spencer
Springfield
Sterling
Stockbridge
Stoneham
Stoughton
Stow
Sturbridge
Sudbury
Sunderland
Sutton
Swampscott
Swansea
Taunton
Templeton
Tewksbury
Tolland
Topsfield
Townsend
Truro
Upton
Uxbridge
Wakefield
Wales
Walpole
Waltham
Ware
Wareham
Warren
Watertown
Wayland
Webster
Wellesley
Wellfleet
Wenham
West Boylston
West Bridgewater
West Brookfield
West Newbury
West Springfield
West Tisbury
Westborough
Westfield
Westford
Westminster
Weston
Westport
Westwood
Weymouth
Whately
Whitman
Wilbraham
Williamstown
Wilmington
Winchendon
Winchester
Winthrop
Woburn
Worcester
Worthington
Wrentham

B Appendix: Results Table for Callaway and Sant'Anna (2021) Approach

Quartlery Analysis

Crime Type	**Coefficient**	**95% CI**
Burglary/Breaking and Entering	3.726	[-8.290, 15.743]
	(6.131)	
Pocket-picking	-1.172	[-3.724, 1.379]
	(1.302)	
Shoplifting	-8.570	[-26.913, 9.773]
	(9.359)	
Theft From Building	14.070	[-4.988, 33.128]
	(9.724)	
Theft From Motor Vehicle	6.130	[-5.573, 17.832]
	(5.971)	
All Other Larceny	-2.064	[-34.466, 30.338]
	(16.532)	
Destruction/Damage/Vandalism of Property	5.538	[-11.142, 22.219]
	(8.511)	
Motor Vehicle Theft	8.299	[-3.236, 19.833]
	(5.885)	
Robbery	-0.567	[-5.165, 4.031]
	(2.346)	
Aggravated Assault	0.662	[-8.636, 9.961]
	(4.744)	
Simple Assault	6.381	[-6.794, 19.557]
	(6.722)	

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Difference-in-Differences Estimates for Crime Outcomes

Monthly Analysis

Crime Type	**Coefficient**	**95% CI**
Burglary/Breaking and Entering	-6.573	[-15.978, 2.832]
	(4.799)	
Pocket-picking	-0.307	[-2.021, 1.408]
	(0.875)	
Shoplifting	1.981	[-5.788, 9.749]
	(3.963)	
Theft From Building	-7.307	[-16.846, 2.232]
	(4.867)	
Theft From Motor Vehicle	-2.284	[-18.693, 14.124]
	(8.372)	
All Other Larceny	7.242	[-0.742, 15.226]
	(4.074)	
Destruction/Damage/Vandalism of Property	1.566	[-1.762, 4.895]
	(1.698)	
Motor Vehicle Theft	-4.216	[-10.306, 1.874]
	(3.107)	
Robbery	-3.290***	[-5.524, -1.056]
	(1.140)	
Aggravated Assault	0.880	[-2.002, 3.762]
	(1.471)	
Simple Assault	-7.402	[-18.879, 4.075]
	(5.856)	

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Average Treatment Effects on the Treated (ATT) by Crime Type

Note that the results in this are fairly different in terms of signs compared to the quarterly analysis. However, there are, again, really high standard errors in these results, underscoring the idea that there is not evidence to justify FRT implementation. They do confirm a trend not too deeply discussed in the body of the paper that robbery seems to decrease following bans. This only works to underscore the idea that the impact of FRT on crime is ambiguous.

Yearly Analysis

Crime Type	**Coefficient**	**95% CI**
Burglary/Breaking and Entering	-37.877	[-99.290, 23.535]
	(31.334)	
Pocket-picking	-5.824	[-15.198, 3.549]
	(4.782)	
Shoplifting	-23.036	[-83.796, 37.725]
	(31.001)	
Theft From Building	-5.715	[-80.978, 69.548]
	(38.400)	
Theft From Motor Vehicle	-7.912	[-57.229, 41.406]
	(25.162)	
All Other Larceny	56.251*	[-9.733, 122.235]
	(33.666)	
Destruction/Damage/Vandalism of Property	2.914	[-19.180, 25.007]
	(11.272)	
Motor Vehicle Theft	13.757	[-26.944, 54.458]
	(20.766)	
Robbery	-6.584	[-19.287, 6.118]
	(6.481)	
Aggravated Assault	-8.008	[-27.564, 11.548]
	(9.978)	
Simple Assault	35.549**	[5.406, 65.691]
	(15.379)	

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Average Treatment Effects on the Treated (ATT) by Crime Type – Alternative Model