The Heterogeneous Effects of Optional Oversight: Evidence from Police Body-Worn Cameras

Andrew Jordan Taeho Kim Soumendra Lahiri
Christopher Lucas Kyle Rozema

March 20, 2025

We study body-worn cameras (BWCs), the decision to activate them, and their effects on police behavior. Using novel data on BWC activation for specific officers and arrests in Chicago, we find large differences in activation rates between police officers that cannot be explained by the policing environment, suggesting that many officers strategically fail to activate. We develop a model of strategic activation that suggests that officers who underactivate will also make fewer arrests when given BWCs. We confirm this using the staggered rollout of body-worn cameras across districts between 2016 and 2017. Overall arrest rates are unaffected, but officers who activate less and officers with high prior complaint rates both reduce their arrests. Finally, using variation in daily assignments of police officers to managers, we find that managers who produce higher activation rates also produced lower arrests and use of force rates following BWC rollout.

Preliminary. Please do not distribute.

---

*We thank Hannah O'Neil and Sidney Nwakanma for excellent research assistance. This research is supported by a grant from the Incubator for Transdisciplinary Futures at Washington University in St. Louis and NSF Grant #2312383

Introduction

In an effort to improve police accountability, most police departments across the United States now require officers to wear body-worn cameras (BWCs) and activate them in most interactions with civilians. Evidence on how BWCs affect police behavior is mixed, ranging from large positive to large negative effects for outcomes including arrests, civilian complaints, and use of force (Lum et al., 2020). Despite more than 30 studies of BWCs, there is a lack of consensus on the effects of cameras and the mechanisms of those effects.

One explanation offered for why findings vary across studies is the amount of discretion officers have in activating (Lum et al., 2020). Based on surveys of officers, BWC activation rates vary significantly across departments (Katz and Huff, 2023), though cross-department comparisons are inevitably confounded by other differences.[1] Barbosa et al. (2021) find that estimates from BWC RCTs are attenuated when researchers aggregate across officers, indicating that officer-level heterogeneity is important to these estimates.

We consider understanding activation heterogeneity to be of first order importance because it has implications for the design of BWC policies, general oversight, and police productivity. Discretion in activation creates opportunities for strategic behavior by police in the activation decision itself but also in decisions about engagement and arrests. Absent this understanding, police departments lack a credible way to develop model policies when adopting or refining their BWC programs.

In this article, we study how body-worn cameras shape police behavior. To do so, we use novel administrative data from Chicago that identifies the BWC activation decisions of specific officers as they make specific arrests. This results in data on BWC activation of unprecedented scope and detail, focused on a set of policing events where officers have an unambiguous duty to activate. We first use these data to summarize patterns of BWC

[1]For example, Kim (2023) shows that the effect of BWCs differ in departments with more police-involved homicides.

activation. We then develop a simple model of policing to better understand these patterns, focusing on heterogeneity across officers. Finally, we investigate the impact of BWCs on police behavior in Chicago with an eye towards the heterogeneity emphasized by the model. This includes differences in officer-level behavior and in supervisors, whom we show have significant impacts on activation rates. We estimate effects of BWCs on arrests and force using both the staggered rollout of cameras across patrol districts.

This article's contribution to the understanding of body-worn cameras is threefold. First, we document patterns of BWC activation in rich data covering a major city. Second, we develop a framework that helps to explain those patterns in terms of officer heterogeneity and its effect on activation and arrest decisions. Third, guided by this model, we document important heterogeneity in the effect of a BWC rollout on police outcomes, finding reduced arrest rates specifically among the officers most resistant to BWCs.

To construct our data, we link arrests and activation records by exploiting a unique event identifier that Chicago's emergency services dispatch generates. This identifier is recorded in some (but not all) arrest and activation records. To link records that do not contain this unique event identifier, we first estimate a statistical model of activation and arrest times using arrest events with direct linkages. We then use this model to mark records that should be linked even in the absence of administrative identifiers. Arrests with no plausible nearby BWC activation by the arresting officer are marked as non-activations.

Across 1,018,593 arrest-officer pairs where the officer was equipped with a BWC between 2016 and 2020, we find an activation rate of 72%. Substantial heterogeneity in activation rates across both police districts and officers reflects inconsistent enforcement of activation policies in Chicago during this period. Fewer than 7% of officers activate in more than 90% of arrests they make. This suggests that idiosyncratic factors, like equipment failure or requests from civilians to not activate, play at least some role in preventing perfect activation. However, a long left tail of officers who activate very rarely—10% of officers

activate in less than 50% of arrests—suggests that at least some of these non-activations were strategic decisions made by officers.

Our investigation into BWC activation patterns reveals notable differences based on suspect demographics and officers' prior complaint histories. Specifically, we find that activation is more likely in arrests involving Black suspects, which we relate to officer anxieties about misconduct allegations from Black civilians. Conversely, officers with a higher number of complaints prior to the BWC implementation are generally less likely to activate their cameras, possibly due to apprehension about documentation that could be detrimental if scrutinized. Exploring interactions, when a Black officer arrests a Black suspect, activation rates are particularly high. However, when a complaint-prone officer conducts the arrest, activation rates are especially low. These findings suggest important heterogeneity across officers and arrest situations where officers seemingly desire records of their actions in some risky situations but not others.

To explain these patterns, we develop a simple model of policing and use it to understand how BWCs are expected to influence policing behavior for different types of officers. Officers consider the risk of misconduct accusations when deciding whether to make arrests, and BWC activation decreases the rate of both false negatives and false positives in misconduct investigations. Officers may therefore view BWCs as either protective or compromising depending on which type of error is more relevant. The model suggests that officers who find BWCs protective will activate frequently and make more arrests when they do activate. However, officers who find BWCs compromising will fail to activate in riskier arrests and even refuse to make arrests if forced to activate. We therefore expect BWCs to have an ambiguous effect on overall police behaviors like arrests, but they will decrease arrests among officers with more problematic histories and lower BWC activation rates.

We next use the staggered adoption of BWCs to estimate the heterogeneous impact of body cameras on policing behavior. Chicago has 22 geographic police districts, and BWCs

were implemented through a staggered rollout at the district-level between 2016 and 2017. Our primary outcomes are average arrests and uses of force per shift within a district-month. We use the difference-in-differences approach in Callaway and Sant'Anna (2021), which avoids negative weighting issues associated with two-way fixed effects approaches with staggered policy adoption (see, e.g., Wing et al., 2024). We are particularly interested in heterogeneity across different groups of officers, so we apply this approach to differenced outcomes that measure how BWCs changed arrests or force in one group relative to the other.

We find no overall effect of BWCs on arrests but important heterogeneity. Officers in the bottom quartile of BWC activation reduce their arrests relative to other officers. The same is true of officers with high prior complaint rates, but we find a positive differential effect for officers with high prior force rates. This underlines a recurring finding that high-complaint and high-force officers are not equivalent in their BWC behavior, with the former generally activating as if they find BWCs compromising and the latter as if they find them protective. When we turn our attention to force as an outcome, our estimates are noisy and not statistically significant. Overall, these findings are consistent with our theory. Officers who find BWCs compromising, represented by officers who activate infrequently or face a high rate of complaints, make fewer arrests with BWCs. This is offset in the aggregate by other officers who find BWCs protective and so increase their arrest output.

Next, we investigate whether differences in enforcement of activation rules affect activation patterns and policing behavior. In Chicago, the primary enforcer of activation rules in a particular shift and district is the watch operations lieutenant (WOLs; "manager"). These managers oversee day-to-day operations from the district's police station, and, according to CPD rules, they are required to review BWC footage for every arrest to confirm probable cause and review at least one randomly selected BWC recording per shift. Each shift an officer works has a single manager, and a typical police district in a typical month has 5

managers who oversee operations in 88% of shifts.

We first investigate the extent to which managers affect BWC activation. Drawing on within-officer and within-policing environment variation, we find that managers have a large impact on activation. For example, we find that the strictest 10 percent of managers increase activation rates by 14 percentage points compared to the most lenient 10 percent of managers. These results suggest that the managers play an important role in activation decisions. This is in itself an encouraging finding for future efforts to encourage more BWC activation. Next, we explore how managers influence other policing behaviors. Since strict managers may possess a bundle of characteristics that shape officer behavior beyond BWC activation, we account for non-BWC-related aspects of WOLs in the rollout design. We find that teams managed by strict WOLs reduce arrests more than those managed by lenient WOLs following BWC adoption. This reduction is primarily driven by non-felony arrests, while felony arrests remain largely unchanged. Additionally, we find suggestive evidence that strict managers reduce use-of-force incidents after BWC rollout. These findings highlight the role of management practices in oversight and help explain why previous studies have found varying BWC effects across departments. The significant reductions in use of force observed in departments with stringent activation policies seem to have stemmed from management practices rather than inherent department-level differences alone.

Finally, we conduct a structural estimation of our model. We fit parameters to match the empirical distribution of officer-level activation rates and conclude that a little over one third of officers at least occasionally neglect to activate their BWC to avoid scrutiny. However, we also find that many failures to activate are non-strategic. Officers would have activated if they could but were prevented by some idiosyncratic factor (e.g. equipment failure or forgetfulness). This suggests that CPD could increase both activation rates and arrest rates by curtailing activation failures of this kind.

The balance of the evidence collected in this paper leads us to conclude that some

CPD officers are making a conscious, strategic decision to not activate body-worn cameras. The significant variation in officer-level activation rates that we document can be justified by idiosyncratic factors only insofar as officers differ considerably in their exposure to them. Furthermore, the causal relationships we estimate between arrests and BWCs are consistent with the strategic behavior described in our theoretical model. Officers who activate infrequently respond to the introduction of BWCs by making fewer arrests. This type of systematic relationship between activation and arrest strongly suggests that officers are considering the consequences for how an arrest, and any related disciplinary actions, will be perceived when choosing whether to activate.

Recent studies have advanced our understanding of the factors influencing officers' BWC activation decisions. Activation is negatively correlated with complaints (Martain et al., 2020; Huff et al., 2024) and is less likely during violent calls than during property crime calls. Activation is more likely during incidents involving arrests or use of force (Martain et al., 2020; Huff et al., 2024; Katz and Huff, 2023). Officer attitudes also play a role, as skepticism toward BWCs or concerns about public scrutiny are associated with lower activation rates (Adams et al., 2021). Activation rates further vary by context, as demonstrated in an RCT in Brazil, where they were just 23.9% in calls without civilian contact (Barbosa et al., 2021). Relative to these studies, we have a substantially larger sample. We observe thousands of officers and over a million opportunities to activate, even when we restrict attention to arrests. This is in part due to the fact that we gather data from more officers than just those treated by an RCT, which gives us a picture of activation patterns once BWCs have ceased to be a novelty. We also have access to a variety of supplemental data including: demographic information of arrestees, historical force and complaint data, and assignment to managers. Finally, our data spans the staggered rollout of BWCs, which we use to estimate the causal effect of BWCs on police behavior. We leverage these data to enhance our analysis and develop theories that help explain activation patterns in terms of

officers' incentives and decisions.

Outside of the BWC literature, this paper relates to the literature that studies officers' responses to citizen oversight. A few recent studies have utilized the CPD officer complaint records obtained in Rivera and Ba (2019) to explore this topic. Ba (2017) finds that a greater distance from the citizen oversight center results in higher opportunity costs, which discourages individuals from filing complaints about misconduct, potentially leading to a reduction in the rates at which certain complaints are upheld. Rozema and Schanzenbach (2019) finds that previous civilian complaints predict future misconduct. Additionally, Rozema and Schanzenbach (2020) find that officers receiving complaints that are ultimately sustained are less likely to engage in subsequent misconduct. Jordan and Kim (2024) find that officers who have complaints handled by stricter investigators accumulate fewer complaints in normal times but that this trend reverses following a police scandal. Beyond the Chicago context, Devi and Fryer Jr. (2020) document instances of "de-policing" in the wake of scandals across several cities. Prendergast (2020) presents a theoretical model suggesting that small increases in scrutiny can improve officer conduct, while larger increases may lead to disengagement. Our study builds on this literature by using microdata to reveal heterogeneity in officers' responses to oversight. While prior research studies a general impact of oversight, we find that some officers perceive BWC oversight as protective, whereas others view it as compromising.

More broadly, this paper contributes to a substantial body of literature on agency issues. BWCs epitomize a new wave of intensive workplace monitoring technologies. Organizations make substantial investments to monitor their employees. While these technologies possess the capacity to mitigate misconduct, they may also infringe upon privacy, which could adversely affect productivity and employee retention (Bernstein (2012)). For instance, Pierce et al. (2015) find that theft-monitoring software significantly reduces theft incidents and enhances productivity in restaurant settings. This study focuses on a workplace monitoring

technology where agency issues are particularly salient. Motivating government employees is often more challenging than motivating private-sector workers due to difficulties in defining objective performance metrics and the limited contracts available to public-sector employees. Enforcement agencies, in particular, face severe challenges in this regard (Finan et al. (2017); Khan et al. (2016); Prendergast (2001)). Our study explores specifically a context where agents wield significant discretion, including the ability to undermine the oversight technology. We also investigate ways to strengthen these oversight mechanisms.

The remainder of the paper proceeds as follows. Section 1 describes how CPD is organized and the adoption of BWCs. Section 2 describes the data and reports descriptive statistics. Section 3 documents officer-level activation patterns and investigates the drivers of differences in activation rates between officers. Section 4 incorporates these patterns into a simple model of BWCs and policing behavior. Section 5 investigates how BWCs affect policing behavior. Section 6 investigates how managers affect activation and policing outcomes. Section 7 considers insights from a structural estimation of our model. Section 8 concludes.

1 Background

1.1 BWCs in Chicago.

The CPD introduced BWCs following significant public and governmental pressure that emerged after the shooting of Laquan McDonald. The release of the dash camera footage in November 2015 triggered widespread protests throughout Chicago. In response, the CPD initiated a pilot program in January 2015 and began a full-scale rollout across all 22 districts in June 2016, which was completed by December 2017.[2] The program provided BWCs to all field personnel within the districts, including district-assigned officers, tactical teams, and

[2]Cameras were rolled out in roughly three adoption periods: six events from June 2016 to February 2017 (Districts 2, 6, 9, 10, 14, and 15), 5 events from March 2017 to July 2017 (Districts 1, 4, 7, 11, and 18), and eleven events from October 2017 to December 2017 (Districts 3, 5, 8, 12, 16, 17, 19, 20, 22, 24, 25).

supervisors. Personnel assigned to city-wide roles (e.g. detectives, vice, and SWAT) were not provided cameras as part of this rollout or at any point in our analysis period. Officers received both hands-on training and presentation-based instruction to ensure proper camera usage.[3]

For officers with a BWC, CPD policy mandates that BWCs are worn on the front of the uniform and powered on at the start of each shift. CPD policy mandates that officers activate their BWCs at the start of any law-enforcement-related activity and record the entire duration of the incident.[4] This includes all arrests, which is our focus, but it also includes responses to calls for service, investigatory stops, and any citizen encounters that become adversarial. At the end of the officer's shift, BWC footage is uploaded to a centralized digital management system.

BWCs may malfunction or lose power, preventing an officer from using them. However, officers are required to verify that their BWC is working and charged at the beginning of each shift. Furthermore, throughout each officer's shift, an officer's superior checks in with the officer and document whether their BWC is functioning properly. If it is not, CPD has a dedicated help desk to get the BWC back functioning properly or replace the BWC with one that is functioning properly.[5] We will view non-activation during arrests due to a technological failure of the BWC itself as a possibility but not one that would affect the same officer persistently.

1.2 Supervisors and BWC Activation Enforcement.

Our analysis of managers will focus on watch operations lieutenants (WOLs) who oversee day-to-day operations from a district's police station and function as the highest

[3]Mayor Press Office, Mayor Emanuel, Chicago Police Department Announce Expedited Expansion of Body Worn Cameras Program, December 28, 2016

[4]In specific cases, such as when requested by a victim or witness, officers may deactivate their cameras, but they must first provide a verbal justification on the recording before ending it.

[5]Body Worn Cameras, Special Order 03-14

on-site supervisory level consistently present in a district during each shift.[6] These lieutenants ensure effective personnel management and review and approve arrest reports and complaints. From a police officer's perspective, each shift is managed by a single watch operations lieutenant, though officers will typically work with several different WOLs across different shifts.

WOLs are particularly important for our purposes because they are the primary enforcers of BWC activation rules. For example, according to CPD rules, part of their oversight responsibilities is to review at least one randomly selected BWC recording per shift, and, whenever operationally feasible, to review footage for every arrest to confirm probable cause. There is evidence that WOLs vary considerably in their adherence to these rules. In an audit of seven districts of the share of watch operations lieutenants who completed required random reviews of BWC footage, compliance was 85% of the one district, 39% in another, 19% in another, and 0% in the remaining four districts (City of Chicago Office of Inspector General, 2019).

2 Data

2.1 Description

We combine several datasets acquired from CPD with Freedom of Information Act requests. Across all datasets, we identify officers using a unique ID based on their full name and date of appointment to CPD. Each supplementary dataset covers, at minimum, the time period of our main analysis.

Arrests

The arrest data record arrests made by CPD from January 1999 to May 2021, though we focus on data beginning in 2016 when BWC rollout began. A typical record describes the

[6]Police officers, sergeants, and lieutenants work in three primary shifts, referred to as "watches." The "morning" shift runs from 6am to 2pm, the afternoon shift runs from 2pm to 10pm, and the "night" watch runs from 10pm to 6am.

officers making the arrest, demographic information of the arrested person, location, time, and charges. Records of arrests where the arrestee was a juvenile are heavily redacted, and we do not use them in our analysis. It is typical for multiple officers to be involved in an arrest, and we organize our data at the officer-arrest level. Unless otherwise specified, when we refer to an "arrest" we are referring to the interaction between an officer and a specific arrest.

Body-Worn Camera Metadata

The BWC metadata record every activation of a body-worn camera by a CPD officer from July 2016 to December 2021.[7] Each activation has a start time, end time, associated officer, and unique evidence ID. Many, but not all, metadata records also record a unique ID for the event the officer was responding to. These IDs are generated by Chicago's emergency services dispatch, either upon a civilian call for service or a report from the officer in the field.

Others

We use several additional datasets from CPD. Officer assignment data allows us to count shifts and assign them to managers. We also use data on complaints made to the civilian police oversight agency in Chicago. These complaints are primarily for excessive force and improper conduct. Complaints for violations of internal regulations and failure to provide service are handled internally by CPD. Our use of force data come from "Tactical Response Reports," which must be filed by CPD officers even after relatively minor uses of force, such as restraining a suspect with hands.

[7]Chicago initiated a small BWC pilot in January 2015, but CPD was only able to provide metadata beginning on July 31, 2016. When considering questions of activation behavior, we restrict our attention to data beginning on that date. When considering the effect of camera rollout, we use the actual rollout dates.

2.2 Labeling BWC Activation

A key step in our analysis is labeling the arrests where officers did and did not activate their BWCs. This is conceptually simple. If the BWC metadata have a record for an activation that corresponds to that arrest and officer, then the camera was activated. If no such record exists, then the camera was not activated. However, the match between the arrest and BWC metadata is not perfect. In both datasets, the event ID is sometimes missing. Officers are responsible for recording event IDs and properly associating them with BWC recordings and may neglect to do so.

We wish to avoid a situation where we label an arrest as a non-activation simply because we failed to match the appropriate metadata record to that arrest. We therefore make a particular effort to match even those records with missing event IDs and will err on the side of accepting too many matches rather than too few. In the absence of the event ID, we can still match arrests with BWC metadata using officer and timing data. When an officer makes an arrest, and the listed timestamp for that arrest falls near the timestamps of a BWC recording associated with that officer, we can reasonably conclude that the officer activated their BWC during the arrest.

We proceed by using the records with an explicit event ID match, of which we have 425,460, as a model against which we compare candidate matches without event IDs, of which we have 488,523. Appendix B sets out the details of our procedure. Figure B2 summarizes its results. It is a histogram of measured differences between arrest and BWC timestamps, topcoded at 100 minutes. It stacks arrests explicitly matched to the metadata by event numbers, accepted proximity matches, and rejected proximity matches. Observations where the arrest timestamp fell between the BWC activation and deactivation timestamps, of which there are 397,785, are omitted to preserve scale. Up to 20 minutes, event number matches are common, and most proximity matches are accepted.[8] By 30 minutes, event number matches

[8]The rule for accepting proximity matches is not a simple cutoff in time because our model allows for arrest characteristics to explain some of the observed discrepancy in timestamps.

are nearly unheard of, and most proximity matches are rejected.

3 Activation Patterns

When studying activation patterns, we focus on officers who plausibly had an opportunity to activate a body-worn camera. We include only officers working in CPD's 22 patrol districts because officers in special units did not have BWCs at this time. We further restrict attention to the period after BWC rollout in the officer's district and to officers whom we observe to activate a camera at least once in the BWC metadata.

The overall activation rate in this sample is 71.7%. Even under a substantially more lenient rule for marking BWC activations, which accommodates timestamp differences in excess of 100 minutes, the overall activation rate is 81.7%. A typical arrest involves more than one officer, often three or more. A single officer failing to activate, therefore, does not preclude the possibility that the arrest was recorded by one of the other officers present. Across our analysis sample, 7.1% of arrests have no BWC activation associated with them. When officers who were not issued cameras are included in the data, the officer-level BWC activation rate falls to 62.3% and 18.1% of arrests have no associated footage.[9] This statistic does not speak directly to each officer's decision to activate their camera, but it does more accurately capture the fraction of arrests in Chicago that went unrecorded even after body-worn cameras rolled out to patrol districts.

Activation may be expected to differ in different policing environments. For example, if the type of crimes and combativeness of suspects differ across neighborhoods and at different times in the week, activation may be expected to differ across them as well. Along these lines, we find that activation rates vary considerably across police districts. Figure 1 presents a map of police districts and their activation rates, documenting cross-district differences exceeding 10 percentage points. Activation rates are typically higher on the North side of the city, which has a population that is on average whiter and wealthier, than the

[9] We calculate these statistics in 2018 and beyond, when all patrol officers had cameras.

South and West side of the city, which have a populations that are on average less white and poorer. However, there are notable exceptions, such as District 7, which has the highest district-level activation rate despite covering some of the poorest neighborhoods in Chicago. This suggests that activation rates could depend on factors beyond those associated with the policing environment, including the income and demographics of people living in the area and the prevalence and severity of the crime officers face. For example, activation rates may depend on culture within a group of officers working together and respond to leadership styles.

We also find that activation rates are lower for more serious charges. Figure 2 presents activation rates by charge type and severity: felony or misdemeanor. Overall activation rates are 73% for misdemeanor charges and 71% for felony charges. These patterns consistently exist when comparing misdemeanor to felony charges within a different type of charge, with only vice charges (e.g., prostitution, gambling) having higher activation rates for felonies. Activation rates tend to be higher for more volatile arrest types, including those involving weapons and violent behavior, and lower for the least volatile arrest types, including arrests for property and nuisance crimes. For example, activation rates range from as low as 66% for felony warrant arrests to 80% for misdemeanor domestic violence arrests. The 3 percentage point bump between misdemeanor domestic violence cases and the misdemeanor arrest category with the next highest activation rate is notable because special emphasis has been placed on the potential for BWCs to provide an objective record of domestic violence even when the injured party withdraws their testimony (Owens et al., 2014; Morrow et al., 2016). The fact that this effect appears to dominate any effect of requests to deactivate cameras while in a victim's home indicates that such requests are likely rare.

To compare activation rates between officers after accounting for differences in policing environment, we regress activation on officer, district by calender month, and charge fixed effects and plot the officer fixed effects. Figure 3 plots a histogram of these fixed ef-

fects.[10] We find that while nearly no officers can boast a perfect activation rate, many have activation rates in the 70-90% range. However, we also find a long left tail in activation rates, with a small but notable number of officers, 10%, activating their BWCs in less than half of their arrests. Compared to the officer at the 25th percentile, the officer at the 75th percentile is 17.9 percentage points more likely to activate. And compared to the officer at the 10th percentile, the officer a the 90th percent is 37.6 percentage points more likely to activate.

For most officers in Figure 3, it is plausible that all non-activations are due to idiosyncratic factors such as forgetfulness, witness/victim requests, equipment failure, or match failure during our labeling routine. However, the low activation rate of officers in the left tail can be explained by these idiosyncratic factors only insofar as one expects them to vary substantially across officers. Otherwise, these left tail officers can reasonably be said to be engaging in some level of intentional or strategic non-activation, if only by maintaining a level of inattentiveness to BWCs that makes them an outlier among their peers.[11] The wide variance in officer-level activation rates suggests important heterogeneity in how officers approach BWCs. The model we develop below thinks through that heterogeneity.

3.1 Subject Race and Officer Complaint History

We next investigate two key dimensions of heterogeneity: subject race and officer complaint histories. We expect that subject race will matter because of the racialized nature of policing and calls for reform after the high-profile incidents involving Black civilians.

[10]To verify the extent that these officer activation fixed effects capture true differences in activation propensities across officers rather than differences in officer environments, we assess how our measure of officer activation predicts activation in out of sample arrests (Appendix Figure A1). We find strong relationships between the activation in one arrest and the officer's activation rate in other arrests (leaving out the arrest in question) (Panel A), activation rates within an officer over time (Panel B), and activation rates within an officer between districts (Panel C). For example, a 1 percentage point increase in an officer's leave-out activation rate is associated with a 0.94 percentage point increase in activation (Panel A).

[11]One may reasonably ask if officers with low activation rates face removal from policing duties. We do not expect that officers may be fired for low activation rates, but they may be moved to duties that do not involve either BWCs or arrests. However, we find that activation rates in 2018 are uncorrelated with an indicator for making no arrests as a patrol officer in 2019.

A 2016 Pew Research survey of nearly 8,000 officers found that officers have race-specific anxieties, with three-quarters of officers saying "as a result of high-profile fatal police-black encounters, officers in their department have been more reluctant to use force when it is appropriate" (Pew Research Center, 2017). This concern about both legitimate and false allegations is reflected in their perceptions of citizen relationships, which have worsened particularly with the Black community. Three-quarters of officers report increased tension between Black civilians and the police, while nine out of ten officers agree that relations between the police and white civilians are generally positive (Pew Research Center, 2017). This evidence suggests that officers making their activation decision may expect, rightly or wrongly, that suspects from historically disadvantaged groups will be more likely to allege bias or excessive force, and perhaps even file a formal complaint.

Meanwhile, officers with a high rate of complaints in their past may be more likely to engage in conduct that merits a complaint, or at least have more reason to suspect that one will arise. Both cases drive towards more complaints, but in the first case this is due to the suspect and in the second case it is due to the officer.

We first require a measure of the propensity of officers to attract complaints that is calculated prior to our rollout period. We calculated the total number of complaints filed against each officer from 1999 (the start of reliable complaint data) up to 2015. We then divided this total by the number of years each officer was active during this period and obtained an average annual complaint rate per officer. We define complaint-prone officers as those whose average complaint rate is in the top quarter of the distribution.

Table 1 reports patterns by subject race and officer complaint history. We find that officers are 1.2 percentage points more likely to activate their cameras when arresting Black suspects (Column 1). This pattern aligns with officers' views that general policing is less error-prone than what the protests and social environment suggest. In the Pew survey cited above, 92% believed that long-standing anti-police bias motivates the protests following high-

profile deaths of Black individuals at the hands of police, and two-thirds of officers mentioned "the fatal encounters that prompted the demonstrations as isolated incidents and not signs of broader problems between police and the black community." However, this relationship is flipped for high-complaint officers. We find that overall complaint-prone officers are 6.8 percentage points less likely to activate cameras (Column 2), and they are 3.1 percentage points *less* likely to activate cameras when arresting a Black suspect (Column 3) relative to a non-Black suspect.

Table A1 repeats this exercise with an indicator for officers with a high prior propensity for force.[12] Officers with a high propensity for using force have higher activation rates, with an additional increase when arresting Black suspects. This indicates that officers frequently involved in use-of-force incidents exhibit a different pattern of behavior compared to complaint-prone officers.

Combined, these results reveal a fascinating pattern. In arrests where there is a greater risk of complaints, whether real or in the officer's perception, officers *generally* respond by activating their BWCs more often. This is consistent with the intuitive story where BWCs benefit officers. They can protect the officer from complaints and accusations that might be harder to disprove in the absence of a video recording of their interaction with a suspect. However, within the subgroup of high-complaint officers, this pattern reverses. Not only are high-complaint officers less likely to activate BWCs overall, but they are particularly unlikely to activate when dealing with Black suspects. This is consistent with the intuitive story where officers may resist BWCs if there is some threat that they will record improper conduct or the appearance of such. A major goal of the model developed below will be to better understand this heterogeneity.

We also assess potential differences in activation rates due to the racial concordance of suspects and officers. Because an officer will have the same race as the suspect for only some

[12]The "high-force officer" measure is calculated in the same way as the "high-complaint officer" measure but with use of force incidents in place of civilian complaints.

of the officer's arrests, we estimate how racial concordance influences activation by drawing on within-officer variation. Table 2 reports the results. We find no racial concordance effects overall (Column 1). However, we find that this overall estimate masks differences between races, with fairly large positive concordance effects for black officers (a 2.1 percentage point higher activation rate) and smaller negative concordance effects for officers of other races (around -0.5 percentage point lower activation rate) (Columns 2-4).

Finally, we document the extent to which officer activation patterns change during ongoing investigations of their misconduct. For each officer, i, and calender month, t, we count the number of investigations previously opened against that officer that had not yet been closed, C_{it}.[13]

Then, we regress:

$$a_{it} = \alpha + \beta C_{it} + \eta_i + \gamma_{dgt} + \epsilon_{it} \tag{1}$$

where a_{it} is the activation rate, η_i are officer fixed effects, and γ_{dgt} are district by officer appointment cohort by calender month fixed effects. Through officer fixed effects, we estimate differences in activation rates within an officer by the number of ongoing complaints. The district by appointment cohort by calender month fixed effects control for changes in activation rates within generations of officers in particular districts over time.[14] Whether an officer is under an ongoing investigation varies by officer, so we cluster standard errors at the officer level.

Table 3 reports the results. Column 1 shows that an ongoing investigation increases the activation rate by 2.2 percentage points. Column 2 includes officer fixed effects, restricting our attention to within-officer comparisons. The effect of an ongoing complaint is reduced

[13]For the small share of complaints where we do not observe the closing date, we assume it closed after 13 months, which is the average length of investigations in our sample.

[14]To avoid dropping observations due to singular officers of an appointment cohort in a district in a given calender month, we define appointment cohort in five year spans.

but remains positive at 0.9 percentage points. Columns 3 and 4 instead consider the *number* of ongoing investigations. The effect on activation rates is 1 percentage point without officer fixed effects and null with them. Finally, columns 5 and 6 consider an indicator for any investigation and a count of active investigations simultaneously. This reveals that officers with more than one ongoing investigation are actually less likely to activate.

The fact that we generally find that ongoing complaint investigations encourage BWC activation may seem at odds with the results earlier in this section, which showed that complaint-prone officers activate less frequently. However, an officer need not be complaint-*prone* to receive a complaint. An officer who tries hard to avoid misconduct may reasonably react to a single complaint investigation by more carefully documenting his activities. Conversely, we find that officers who are the subject of *multiple* complaint investigations respond instead by reducing BWC activation. This exemplifies the heterogeneity in how officers approach BWCs and complaints we model in the next section.

4 A Model of BWC Activation

To capture the features of the data presented in Section 3, we require a model of BWC activation considerations that can accommodate two dualities. First, most officers are more likely to activate their BWCs as the risk of a complaint increases, but some are instead less likely to activate. Second, nearly all officers sometimes fail to activate for idiosyncratic reasons, but these coexist with apparently strategic non-activations.

Appendix C develops a model of officers who must decide whether to attempt an arrest while balancing the direct payoff of doing so with the risk of a misconduct accusation. The mechanism that assigns blame for misconduct is imperfect and subject to both false positives, incidents where officers did not commit misconduct but are blamed regardless, and false negatives, incidents where officers did commit misconduct but get away with it. As the probability of actual misconduct rises, officers eventually decide that engagement is too risky and decline to make an arrest. The arrest decision follows a threshold rule in

misconduct probability.

In this model, body-worn cameras are a monitoring technology that reduce errors of *both* types. For arrests with a low risk of actual misconduct, the biggest threat to officers is a false positive. Activation is beneficial to officers because it reduces this threat. Conversely, for arrests with a high risk of actual misconduct, officers hope for a false negative, but BWC activation reduces the chance of this. Thus, BWC activation also follows a threshold rule in misconduct probability.

The model shows that how officers activate and how other police activity changes in response to BWCs depends crucially on the relative order of the arrest and activation thresholds. When the arrest threshold is below the activation threshold, we say that officers find activation "protective." Absent idiosyncratic activation failures, officers activate in every arrest, and in fact, the ability to activate gives them the security to engage in arrests they would have avoided without BWCs. However, if idiosyncratic activation failures, which we model as a random shock to activation payoffs, are common, these benefits will be only partially realized.

When the arrest threshold is above the activation threshold, we say that officers find activation "compromising." These officers will neglect to activate in the riskiest arrests that they make. If pushed to activate by outside forces (e.g. stringent enforcement of activation rules), they will even opt out of arrests they would have made previously.

The model thus captures the two dualities we set out to explain. Protective activation officers will activate frequently and make more arrests but only if idiosyncratic factors do not get in the way. Compromising activation officers will neglect to activate or make arrests, especially when the risk of misconduct is high. Reminders to activate cameras, insofar as they primarily reduce idiosyncratic non-activations, will encourage arrests and activation among protective activators but have little impact on compromising activators. Conversely, stringent enforcement of BWC rules, insofar as this primarily reduces strategic non-activations,

will encourage activation among compromising activators but discourage arrests. The remainder of this paper is thus focused on heterogeneity in how police responded to the BWC rollout in Chicago and the role of leadership in that response.

5 The Impact of BWCs on Policing Behavior

We now use the staggered rollout of BWCs to police districts in Chicago to consider the effect of a discretionary monitoring technology on officer behavior measured in the form of arrests and use of force. The model in Section 4 suggests that this effect will depend on the balance of officers who find BWC activation to be protective rather than compromising. If BWCs provide cover for officers to engage in risky situations, they may increase arrest rates and uses of force. However, if officers are primarily concerned about being "caught on camera," BWCs may make them less inclined engage.

Because we expect heterogeneity to be important, we will consider both the effect of the BWC rollout on arrests overall and the differential effect by officer characteristics we think may be predictive of finding BWCs compromising. We pay particular attention to officers with high prior complaint and force propensities (both as defined in Section 3) and to officers in the bottom quartile of BWC activation rates, measured only in 2018 and beyond after cameras had rolled out to all districts.

5.1 Empirical Strategy

To use the variation as police districts adopted BWCs, we conduct our analysis at the district-calender month level. We begin with monthly data on each officer that counts the number of shifts worked, S_{ot} and arrests made, A_{ot}, in that month where o indexes officers and t indexes months. For clarity, we will focus on our arrest outcome when describing our approach, but all calculations for our force outcome are similar. We continue to focus on district patrol officers but no longer require that officers be observed activating a BWC. We

assign officers to a district using the modal district of their shifts in that month.[15]

Within each district-month, we calculate the number of arrests made per shift worked by officers assigned to that district.

$$y_{dt} = \frac{\sum_{d(o)=d} A_{ot}}{\sum_{d(o)=d} S_{ot}}$$

where $d(o)$ is the modal unit of officer o. We exclude from this calculation any officer-months with fewer than 5 shifts. This drops 3.2% of officer-months in our data.

Given the negative weighting issues associated with two-way fixed effects (TWFE) approaches with staggered policy adoption (see, e.g., Wing et al., 2024), our primary specification uses the approach in Callaway and Sant'Anna (2021). This approach estimates treatment effects that are both district and time specific using only comparisons between treated and not-yet-treated districts. We report both aggregated treatment effects across all districts for each event time and a simple aggregate across all districts and times. We estimate effects using data from 12 months before the first rollout to 12 months after the last rollout. The identifying assumption in our design is that the policing outcomes in a district that adopts BWCs would have evolved similarly to policing outcomes in other districts that had not yet adopted BWCs. We will present pre-trends that do not show any violation of this assumption.

Beyond the overall effect of BWCs on arrests, we take particular interest in heterogeneity across types of officers. We wish to understand, for example, whether complaint-prone officers react to BWCs differently than others. In a TWFE framework, we could estimate this kind of heterogeneity in effects using a triple-difference specification. The Callaway and Sant'Anna (2021) framework does not have an immediate counterpart to triple-

[15]Officers are assigned to one district at any time and typically work exclusively in that district until their assignment changes. However, officers occasionally work in another district, such as if they work overtime or on a special detail. We treat all arrests made by an officer as if they were made in this modal district. Modal districts account for 99.5% of all assignments within a given officer-calender month.

differences, but we follow the guidance in Roth et al. (2023) to approximate it. Following their recommended specification, we replace our original outcome, average arrests per officer in a district-month, with the difference in outcomes calculated in different subsamples of officers. Formally, define an outcome conditional on a binary indicator $c(o)$ as:

$$y_{dt}^{c} = \frac{\sum_{d(o)=d,\, c(o)=1} A_{ot}}{\sum_{d(o)=d,\, c(o)=1} S_{ot}}$$

then our new outcome of interest is

$$y_{dt}^{\Delta\, c} = y_{dt}^{c} - y_{dt}^{(1-c)}$$

So for example, when studying heterogeneity by prior complaint rates, y_{dt}^{c} is estimated only within officers in the top quartile of prior complaint rates and $y_{dt}^{(1-c)}$ is estimated among all other officers.

5.2 Results

Our estimates of the effect of the BWC rollout on arrests are in Table 4. The "All" column shows that the overall effect on arrests per shift is statistically insignificant.[16] This is a rather precise null with a confidence interval of $[-0.0178, 0.0142]$. Compared to the outcome mean of 0.22 arrests per shift, any true effect of BWC rollout on arrests must be small. Figure A2 shows event study results for this specification and those discussed below.

However, the remaining columns in Table 4 show that this initial null result disguises the significant heterogeneity we set out to measure. We find that complaint-prone officers

[16]Ferrazares (2024) examines uses of force and arrests following BWC rollout at the unit-month level. Using TWFE as the main specification, he finds an imprecise effect estimate of -0.69% for the mean for overall use of force, while strike use of force declines by 34%. For arrests, the TWFE estimate indicates a 7% reduction ($p < 0.1$). All of these effects attenuate toward null when using the Callaway and Sant'Anna (2021) specification. Our overall results are broadly similar, except for the TWFE estimates on overall use of force, where we find an increase. The primary reasons for this difference are: (1) we analyze at the officer-month level rather than the unit-month level, and (2) our TWFE specification restricts the sample to rollout months between -6 and +12 months, whereas Ferrazares (2024) uses a more extended panel from 2014 to 2019.

make 0.041 fewer arrests per shift with BWCs than other officers (Column 2). This aligns with our earlier regression results and with the model. Insofar as officers with many prior complaints are more likely to find BWC activation compromising, we should expect that the rollout of cameras to their district will cause their arrest rates to fall relative to their peers, and this is what we find.

Moreover, our activation data can be used to consider heterogeneity by activation rates directly. We calculate officer-level BWC activation rates from 2018 onward to avoid overlap with the rollout period, then label officers in the bottom quartile as low activators. The model predicts that these low-activation officers find activation compromising, and they should therefore reduce their arrest rate once BWCs are rolled out. Indeed, we find that low-activation officers make 0.03 fewer arrests per shift with BWCs (Table 4 Column 3).

The fourth column of Table 4 repeats this exercise, dividing officers by use of force. We now find an increase of 0.027 arrests per shift. Force and complaints thus appear to have different implications for BWC activation. The strong positive effect of a high force rate on activation rates observed in Table A1, in conjunction with the positive effect observed here, indicate that officers who use force more often tend to view BWCs as protective. If most officers who often use force adhere stringently to department regulations, it makes sense that they would prefer to create an objective record of this. Consider also that we base our measure of high-force officers on official reports kept by CPD. Insofar as some officers are using force improperly, they may be less likely to file these reports.

Table 5 and Figure A3 repeat this exercise but with use of force in place of arrests. We do not find any significant effects of the BWC rollout on use of force, either in the aggregate or across any of the dimensions of heterogeneity we consider. Force is a sparser outcome than arrests, with the average officer only using force 0.003 times a shift. These estimates are thus noisier, and the confidence intervals admit large effects relative to the outcome mean.

Tables A2 and A3 present difference-in-difference and triple-difference coefficients using a regression-based TWFE design. The first column controls for officer fixed effects, district fixed effects, and year-by-month fixed effects. The remaining columns control for officer fixed effects and district-by-year-by-month fixed effects, which absorb the main effect of the rollout. The results in these tables agree with those from the Callaway and Sant'Anna (2021) method in all but two cases. First, they find no significant effect on arrest rates of being a high force officer. Second, they find a significantly positive overall effect on use of force. We do not wish to speculate too much on why the TWFE specifications sometimes differ. It may be because the TWFE specification incorporates more of the effects from CPD districts that adopted cameras relatively late. Alternatively, it may be because the y_{dt}^{Δ} estimators we use are not a perfect analogue for traditional triple-differences.

6 Effects of Enforcement of Activation Rules

As discussed above, watch operations lieutenants, which we refer to as an officer's "manager," oversee day-to-day operations from the district's police station and are the primary enforcers of BWC activation rules. For each officer-shift, there is a single manager. We use the assigned manager in a shift as a way to study how BWC activation enforcement affects activation and policing outcomes.[17]

We first want to understand the assignment of managers to shifts, as this drives the variation in officer-level exposure to different managers who may have different enforcement propensities. Figure 4 plots the mean number of shifts per month for the top 5 managers in each district along with the total shifts covered by all other managers. The large majority of shifts are covered by the top 5 managers in that district (across districts, an average of 88% of shifts). Within a district-month, the manager covering the most shifts covers 26% of shifts and the next cover 13, 11, 8, and 6 percent of shifts.

[17]Two officers working in the same district are exposed to roughly the same mix of managers, so activation rates between officers are unlikely to be explained by managers. However, we do expect that managers will play a role in overall activation rates on the shifts they supervise.

6.1 Manager Contribution to Activation

We first assess the extent to which managers influence BWC activation. Figure 5 plots estimates of manager fixed effects after controlling for officer, district by month, and charge by felony fixed effects and adding back the mean activation rate. We calculate these manager effects in 2018-2020, after the BWC rollout completed. Compared to managers at the 25th percentile of this distribution, managers at the 75th percentile have fixed effects 5 percentage points higher. The corresponding 90-10 difference is 14 percentage points. Net of charge by felony and district by month fixed effects, officer fixed effects explain 12.9 percent of the variation in BWC activation. Manager fixed effects explain an additional 1.4 percent of the variation.

To assess whether the manager fixed effects capture unbiased estimates of differences in manager contribution to activation, we investigate the stability of manager contribution to activation over different arrests, within the same manager over time, and within the same manager in different districts. We plot our results in Figure 6. We find strong relationships between activation in one arrest and the activation rate of officers where the manager was in charge in other arrests (Panel A), between the activation rate of officers where the manager was in charge in one month and activation rates of officers where the manager was in charge in every other month (Panel B), and between the activation rate of officers where the manager was in charge in one district and the activation rate of officers where the manager was in charge in another district (Panel C). For example, a 1 percentage point increase in a manager's leave-out activation rate is associated with a 0.92 percentage point increase in activation.

In summary, we find evidence for a manager effect on BWC activation that is substantial and consistent over time. This in itself is an interesting finding that indicates an important role for leadership. Unfortunately, we do not have data on specific practices that managers may be using to encourage BWC activation, so we can only assess the overall

impact of managers rather than that of management practices.

6.2 Impact of Managers on Policing Behavior

The results above show that managers vary considerably in terms of how they enforce the activation rules. This suggests that officers may have more discretion with activation on workdays with lenient managers than they do on workdays with stricter managers, leading to a greater potential for strategic behavior in terms of other policing decisions. As such, differences in the amount of discretion over activation due to how much the activation rules are enforced may be expected to to influence officer policing behavior beyond activation.

In our final analysis, we leverage the rollout design to investigate how officer behaviors changed under strict versus lenient managers around the rollout period. Strict managers may have management styles that influence team outcomes independently of BWC oversight, so we cannot simply compare policing outcomes across the two. By examining the differences in behaviors before and after the rollout, we can isolate the effect of BWC stringency by comparing the same managers and officers before and after BWCs were available.

We divide managers into two groups: lenient and strict. We identify strict managers as those with FEs more than the median in the distribution of a manager's contribution to activation rates net of officer effects and policing environment.

We then examine how team outcomes under the same set of strict and lenient managers evolved around the rollout. For each district and month, we focus on watches managed by managers for whom we can estimate FEs. Since the rollout of BWCs occurred at the unit-month level, we construct an aggregate outcome measure at the unit-month level to capture differences in team performance under strict versus lenient managers. Specifically, we begin by computing the daily team-level outcome, which is averaged across all officers within a team. For each watch k on day $\tilde{t}$ in unit d, under manager j, the team-level outcome is given by

$$y_{d\tilde{t}kj} = \frac{1}{N_{d\tilde{t}kj}} \sum_{i \in T_{d\tilde{t}kj}} y_{id\tilde{t}kj}$$

where $y_{id\tilde{t}kj}$ denotes the individual officer-level outcome for officer i in watch k, on day $\tilde{t}$, in unit d, under manager j. The total number of officers in the team is represented by $N_{d\tilde{t}kj}$. Next, we average these daily team-level outcomes over the course of a month, distinguishing between manager types. For a given manager type m (strict or lenient), the monthly outcome is computed as

$$y_{dt}^{m} = \frac{1}{\sum_{j \in m} \sum_{\tilde{t} \in t} K_{d\tilde{t}j}} \sum_{j \in m} \sum_{\tilde{t} \in t} \sum_{k \in K_{d\tilde{t}j}} y_{d\tilde{t}kj},$$

where the numerator sums over all team-level outcomes $y_{d\tilde{t}kj}$ across all managers of type m, across all days $\tilde{t}$ in the month t, and across all teams k, and the denominator denotes the total number of team-days managed by all managers of type m across all days in the month. Finally, to capture the difference in monthly team-level performance between strict and lenient managers, we create the following unit-month level differenced outcome variable:

$$\Delta y_{dt} = y_{dt}^{\text{Strict Manager}} - y_{dt}^{\text{Lenient Manager}}.$$

We then apply this outcome to the difference-in-difference design described in Section 5. There are two empirical considerations worth noting. First, this manager analysis focuses only on officers in regular watches, as their schedules closely overlap with those of the managers. Tactical officers, whose shift schedules do not align with the managers' schedules, are excluded from the main analysis. Second, some unit-month observations are dropped. For the DID design to be valid, we require that each unit-month has exposure to both types of managers at some point during the month and that the unit is observed both before and after the rollout. Figure A4 illustrates that four districts are excluded; however, for each included unit, the sample provides reasonable coverage of the outcome variable. Figure 7 and Table 6

report the results. Before BWCs, teams managed by strict and lenient managers exhibited similar trends in arrests and use of force. With BWCs, we find that strict mangers decrease arrests by 0.055 per month (p-value 0.002). The estimate represents a 39% reduction relative to the mean daily team-level average. This reduction is primarily driven by non-felony arrests, which decline by 0.052, whereas felony arrests remain largely unchanged.

Use of force also declines visually following BWC adoption, with an estimated reduction of 0.0068 (p-value = 0.055). While this marginally significant estimate is large compared to the mean daily team-level average of 0.002, it has a wide confidence interval, ranging from a 6.95% increase to a decrease of 6.8 times the daily mean. This variability likely stems from the rarity of use-of-force incidents, which are mostly recorded as zero on a given day.

Table 6 further investigates these estimates. Panels B and C separately estimate the difference-in-differences (DID) effect for teams managed by strict and lenient WOLs. We find that while arrests and use of force tend to increase under lenient managers, they decline under strict WOLs. In Panel D, we reintroduce officers in tactical teams who were initially excluded from the sample. While their inclusion makes the estimates less precise, the overall results remain consistent. Finally, we examine whether the effects differ between officers with low and high numbers of prior complaints. We find that both groups exhibit similar declines, with no statistically significant differences between them.

7 Structural Estimation

To gain more insight into the interplay of officer incentives, arrests, and activation, we structurally estimate the model described in Section 4. Our goals in this estimation are twofold. First, we want to better understand the prevalence of officers who find activation protective and compromising and how they combine to generate the distribution of officer-level activation rates shown in Figure 3. To aid this goal, we prioritize fitting this distribution in our estimation. Second, we wish to consider counterfactual BWC enforcement scenarios and in particular to understand the differential effect of policies that reduce idiosyncratic non-

activations (e.g. ensuring cameras are always charged, neutral reminders to activate) vs. those that curtail strategic non-activations (e.g. pressure to follow activation rules, automated activation systems). Appendix C presents the technical details of the structural estimation.

We address the first goal of the structural estimation with Figure 8. This histogram mirrors Figure 3, except it plots the distribution of 200,000 officers simulated using our estimated parameters. Figure 8 is a stacked histogram that plots officers for whom activation is protective ($\theta_{arr} > \theta_{act}(0)$) in white and officers for whom activation is compromising ($\theta_{arr} < \theta_{act}(0)$) in red. About 63.5% of officers find activation protective, and they tend to have the highest activation rates, with rates above 90% being almost entirely driven by this kind of officer. Conversely, officers who find activation compromising have lower activation rates, making up nearly all officers with rates less than 50%. However, the significant dispersion in activation rates among officers of both types means that for moderate activation rates, the rate alone is not sufficient to identify an officer's type. In this context, our officer-level indicator for low activation, which divides officers at an activation rate of about 64%, captures mostly compromising officers (85%) with a comparison sample that is mostly protective officers (83%).

We address the second goal of the structural estimation with counterfactual specifications that adjust the prevalence of cases in which officers are either forced not to activate, even if they want to, or forced to activate even if they would rather not. Table 7 presents average arrest and activation rates at the officer level for our simulated officers. The arrest rate should be thought of as the percentage of *possible* arrests made under the assumption that the 95th percentile of empirical arrest rates among pre-BWC officers is 100%. Relative to the pre-BWC environment, our model shows arrest rates increasing substantially from 37% to 46% once cameras are introduced. This is at odds with our reduced-form estimates, which find no important change in aggregate arrest rates. This is because our model captures increases in arrests from officers who find BWC activation protective, but the estimated

decrease in arrests from officers who find BWC activation compromising is quite small.

Continuing in Table 7, the rows labeled "High p_h" present a counterfactual scenario where the probability that an officer is forced to activate his BWC, even against his better judgment, rises from 0.001 to 0.5. This has the effect of increasing overall activation by about 5 percentage points, primarily among officers who find activation compromising, but those officers also make substantially fewer arrests, dropping their arrest rate from 68% to 60%. The rows labeled "Low p_l" present a different counterfactual scenario where the probability that an officer is unable to activate his BWC, even when it would be beneficial, falls from 0.217 to 0.05. This has the effect of increasing both arrests (from 46% to 48%) and activations (from 75% to 85%). Our results on WOLs are qualitatively consistent with the "High p_h" scenario wherein BWC activation rules are enforced more stringently, increasing the overall activation rate but decreasing arrests.

8 Conclusion

Using data from Chicago, we have shown that body-worn camera activation differs across several dimensions. Activation rates differ considerably across officers, districts, and managers, which suggests a lack of consistent enforcement of BWC policies. Our model of activation and arrest decisions highlights an important dimension of heterogeneity in theory: whether officers find BWC activation protective or compromising. Regression evidence on activations and difference-in-differences evidence on arrests show some practical import for this theoretical distinction.

Future work should make further efforts to distinguish protective activators from compromising activators, whether through survey responses or careful analysis of policing behavior. There is also an imperative to explore enforcement policies that are sensitive to this distinction. Initiatives to ensure officers are always prepared, both physically and mentally, to activate can reduce idiosyncratic non-activations, improving both activation and arrest rates for protective activators while having minimal effects on compromising

activators. Conversely, policies that assign strict punishments for any failure to activate may cause compromising activators to reduce their arrest productivity. Given that this reduction in arrests would be focused on those most likely to generate a misconduct complaint, the welfare implications of such a shift are beyond the scope of this study.

References

Adams, I. T., S. M. Mourtgos, and S. H. Mastracci (2021). High-Stakes Administrative Discretion: What Drives Body-Worn Camera Activations? *Public Administration Review 81*(4), 691–703.

Ba, B. A. (2017). Going the Extra Mile : the Cost of Complaint Filing, Accountability, and Law Enforcement Outcomes in Chicago.

Barbosa, D., T. Fetzer, C. Soto, and P. Souza (2021). De-escalation technology: the impact of body-worn cameras on citizen-police interactions. Working Paper.

Bernstein, E. S. (2012). The transparency paradox: A role for privacy in organizational learning and operational control. *Administrative Science Quarterly 57*(2), 181–216.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics 225*(2), 200–230.

City of Chicago Office of Inspector General (2019). Evaluation of the Chicago Police Department's Random Reviews of Body-Worn Camera Recordings. Technical Report July.

Devi, T. and R. G. Fryer Jr. (2020). Policing the Police: The Impact of "Pattern-or-Practice" Investigations on Crime. *National Bureau of Economic Research (NBER)*, 63.

Ferrazares, T. (2024). Monitoring Police with Body-Worn Cameras: Evidence from Chicago. *Journal of Urban Economics 141*(October 2021), 103539.

Finan, F., B. A. Olken, and R. Pande (2017). The Personnel Economics of the State. *Handbook of Field Experiments* (December).

Huff, J., M. D. White, A. E. Malm, and C. M. Katz (2024). An Officer-Level Examination of the Prevalence and Correlates of Police Body-Worn Camera Activation. *Police Quarterly*, 1–33.

Jordan, A. and T. Kim (2024). Strengthening Police Oversight: the Impacts of Misconduct Investigators on Police Officer Behavior. Working Paper.

Katz, C. M. and J. Huff (2023). The Achilles Heel of Police Body-Worn Cameras: Understanding the Factors That Influence Variation in Body-Worn Camera Activation. *Justice Quarterly 40*(3), 315–336.

Khan, A. Q., A. I. Khwaja, and B. A. Olken (2016). I. Introduction Tax systems in developing countries collect substantially less revenue as a share of GDP than do their counterparts in. *The Quarterly Journal of Economics*, 219–271.

Kim, T. (2023). Facilitating Police Reform: Body Cameras, Police-Involved Homicides, and Law Enforcement Outcomes.

Kucherenko, S. and Y. Sytsko (2005). Application of deterministic low-discrepancy sequences in global optimization. *Computational Optimization and Applications 30*.

Lum, C., C. S. Koper, D. B. Wilson, M. Stoltz, M. Goodier, E. Eggins, A. Higginson, and L. Mazerolle (2020). Body-worn cameras' effects on police officers and citizen behavior: A systematic review. *Campbell Systematic Reviews 16*(3), 1–40.

Martain, B. R., V. Harinam, and B. Ariel (2020). Linking body worn camera activation with complaints: The promise of metadata. *Journal of Criminology* (October 2020).

Morrow, W. J., C. Katz, and D. E. Choate (2016). Assessing the impact of police body-worn cameras on arresting, prosecuting, and convicting suspects of intimate partner violence. *Police Quarterly 19*(3), 303–325.

Nelder, J. A. and R. Mead (1965, 01). A simplex method for function minimization. *The Computer Journal 7*(4), 308–313.

Owens, C., D. Mann, and R. McKenna (2014). The essex body worn video trial the impact of body worn video on criminal justice outcomes of domestic abuse incidents.

Pew Research Center (2017). Behind the Badge.

Pierce, L., D. C. Snow, and A. McAfee (2015). Cleaning house: The impact of information technology monitoring on employee theft and productivity. *Management Science 61*(10), 2299–2319.

Prendergast, C. (2001). Selection and Oversight in the Public Sector, with the Los Angeles Police Department as an Example. *NBER Working paper series* (January), 43.

Prendergast, C. (2020). "Drive and Wave": The Response to LAPD Police Reforms After Rampart.

Rivera, R. G. and B. A. Ba (2019). The Effect of Police Oversight on Crime and Allegations of Misconduct: Evidence from Chicago.

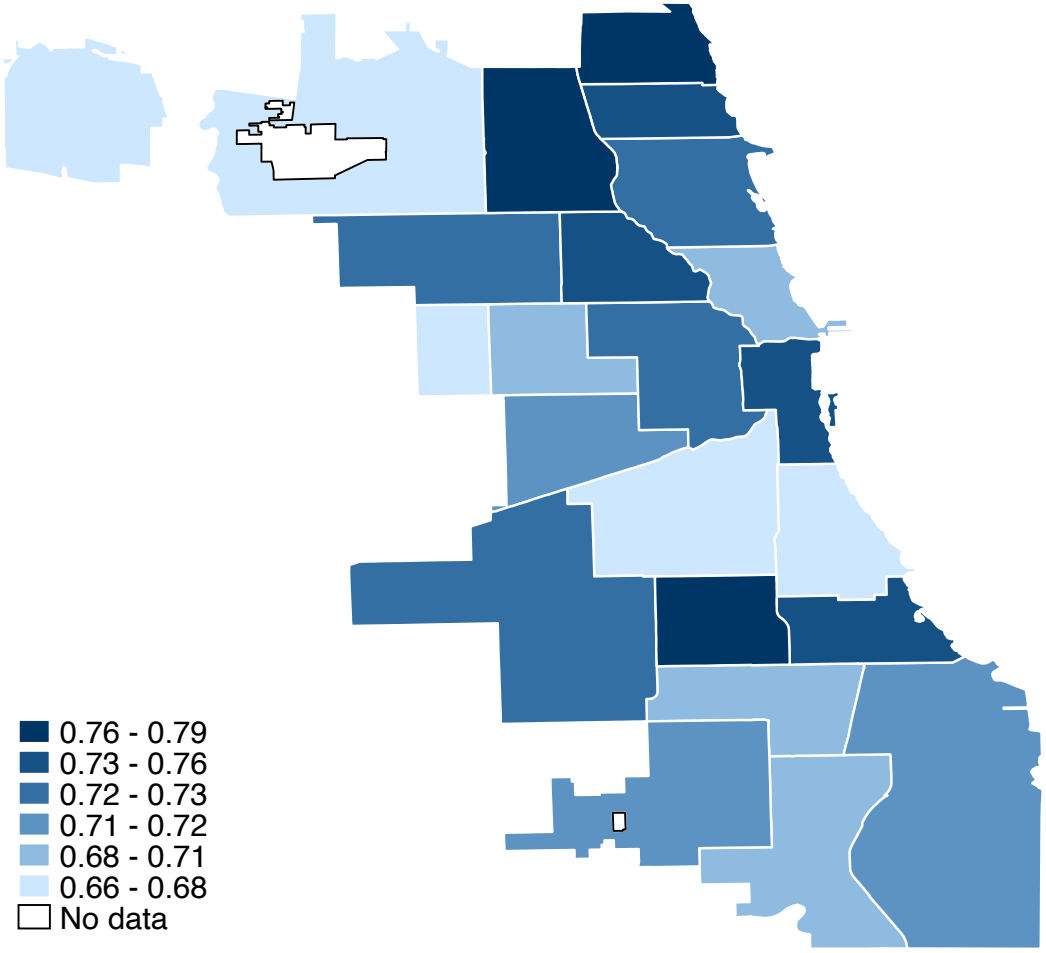
Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics 235*(2), 2218–2244.

Rozema, K. and M. Schanzenbach (2019). Good Cop, Bad Cop: Using Civilian Allegations to Predict Police Misconduct. *American Economic Journal: Economic Policy*.

Rozema, K. and M. M. Schanzenbach (2020). Does Discipline Decrease Police Misconduct? Evidence from Chicago Civilian Allegations.

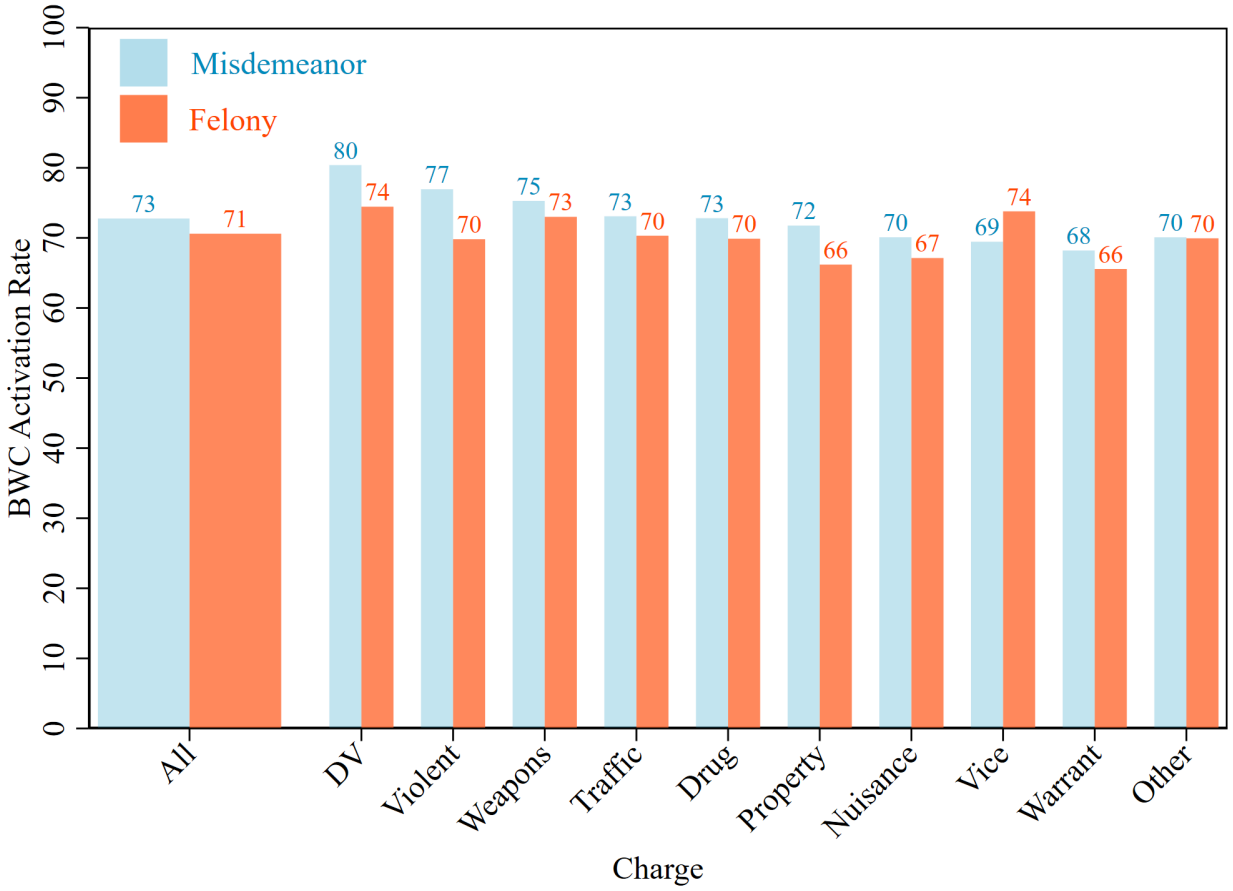
Wing, C., S. M. Freedman, and A. Hollingsworth (2024). Stacked Difference-in-Differences. *NBER Working Paper # 32054*.

Figure 1: Activation by Arrest Characteristic: District



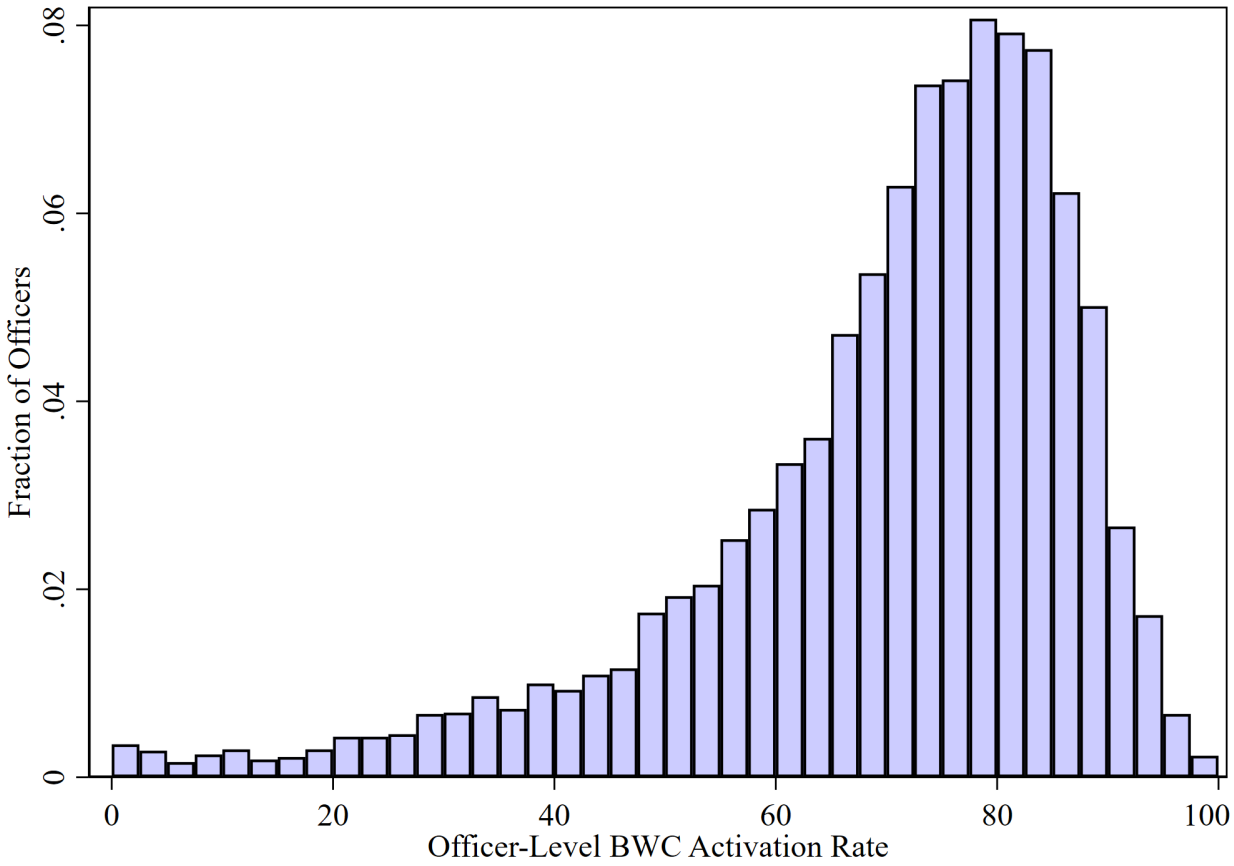
Note: This figure reports the body-worn camera activation rate for arrests by Chicago police district among officers who had cameras assigned at the time of the arrest.

Figure 2: Activation by Arrest Characteristic: Charge



Note: This figure reports the overall body-worn camera activation rate for arrests by the type of charge, separately by whether the charge was a felony (orange bars) or a misdemeanor (blue bars).

Figure 3: Activation Rates of Individual Officers



Note: The figure reports the distribution of estimates for the officer-level body-worn camera activation rates. The estimated coefficients are from officer fixed effects in a regression that also includes charge by felony and district by calender month fixed effects.

Table 1: Differential Activation by Subject Race and Officer Type: Complaints

	(1)	(2)	(3)
Black Arrestee	.0117***		.0183***
	(.00245)		(.00307)
High Complaints Officer		-.0677***	-.0454***
		(.00687)	(.00949)
High Complaints Officer x Black Arrestee			-.0309***
			(.0089)
Y mean	.698	.698	.698
N	649,989	649,989	649,989
Time FEs	Y	Y	Y
District FEs	Y	Y	Y
Arrest Categories FEs	Y	Y	Y

Notes: This table presents descriptive regressions of the relationship between BWC activation, suspect race, and prior officer complaint rate. Prior complaints are counted from 2000 and divided by the number of years of active service and data availability. We label an officer as high complaint if their calculated complaints is in the top quartile among officers in our analysis sample. We estimate these regressions in the sample of arrests where officers were working in patrol districts that had already rolled out cameras, and we observe that officer to have activated a BWC at least once. All regressions further condition on the high complaint measure being available, even when that regression does not include the high complaint indicator. We do not calculate the high complaint measure for officers who began working after 2015. Standard errors are clustered by officer and are in parentheses. $^{+}$ p<0.1, * p<0.05, ** p<0.01, *** p<0.001.

Table 2: Officer-Suspect Racial Concordance and BWC Activation

	Officer Race			
	All	Black	Hispanic	White
	(1)	(2)	(3)	(4)
Same Race of Officer and Suspect	-0.002	0.021***	-0.006*	-0.005*
	(0.002)	(0.005)	(0.003)	(0.003)
Mean	0.722	0.722	0.722	0.722
Observations	969,125	969,125	969,125	969,125

Note: This table estimates the relationship between racial concordance of the officer and suspect and BWC activation. Column 1 estimates concordance effects for all officers, and Columns 2-4 estimate concordance effects separately by race. The estimates draw on within-officer variation in activation by regressing an indicator for BWC activation for an arrest on the racial concordance indicator (the same race of the officer and suspect) and officer, charge by felony, and district by calender month fixed effects. In Columns 2-4, all officers are included in the regression but the racial concordance indicator is zero for all races other than the one in question, as indicated at the top of the table (for example, in Column 2, the racial concordance indicator is zero for all officers other than black officers). Standard errors are clustered by officer and are in parentheses. $^{+}$ p<0.1, * p<0.05, ** p<0.01, *** p<0.001.

Table 3: Activation Patterns During Ongoing Investigations of Complaints

	(1)	(2)	(3)	(4)	(5)	(6)
Any Ongoing Investigation	0.022***	0.009**			0.026***	0.017**
	(0.004)	(0.003)			(0.006)	(0.005)
Number of Ongoing Investigations			0.010***	0.002	-0.003	-0.007+
			(0.002)	(0.002)	(0.004)	(0.004)
Covariates						
District by Cohort by Time FE	✓	✓	✓	✓	✓	✓
Officer FE		✓		✓		✓
Mean Activation Rate	0.701	0.702	0.701	0.702	0.701	0.702
Observations	192,667	192,351	192,667	192,351	192,667	192,351

Note: This table presents estimates of the relationship between BWC activation rates and ongoing investigations of civilian complaints. The level of analysis is an officer-month. The first row is an indicator for that officer having any open investigation in that month. The second row counts open investigations. All specifications control for district-by-offficer-cohort-by-calendar-month fixed effects. Even-numbered specifications add officer fixed effects and so capture variation just within officers. + p<0.1, * p<0.05, ** p<0.01, *** p<0.001.

Table 4: Effects of BWC Rollout: Arrests

	(1) All	(2) High Complaint	(3) Low BWC	(4) High Force
ATT	-0.00183	-0.0409*	-0.0301*	0.0273*
	(0.00816)	(0.0162)	(0.0145)	(0.0118)
Observations	528	528	528	528

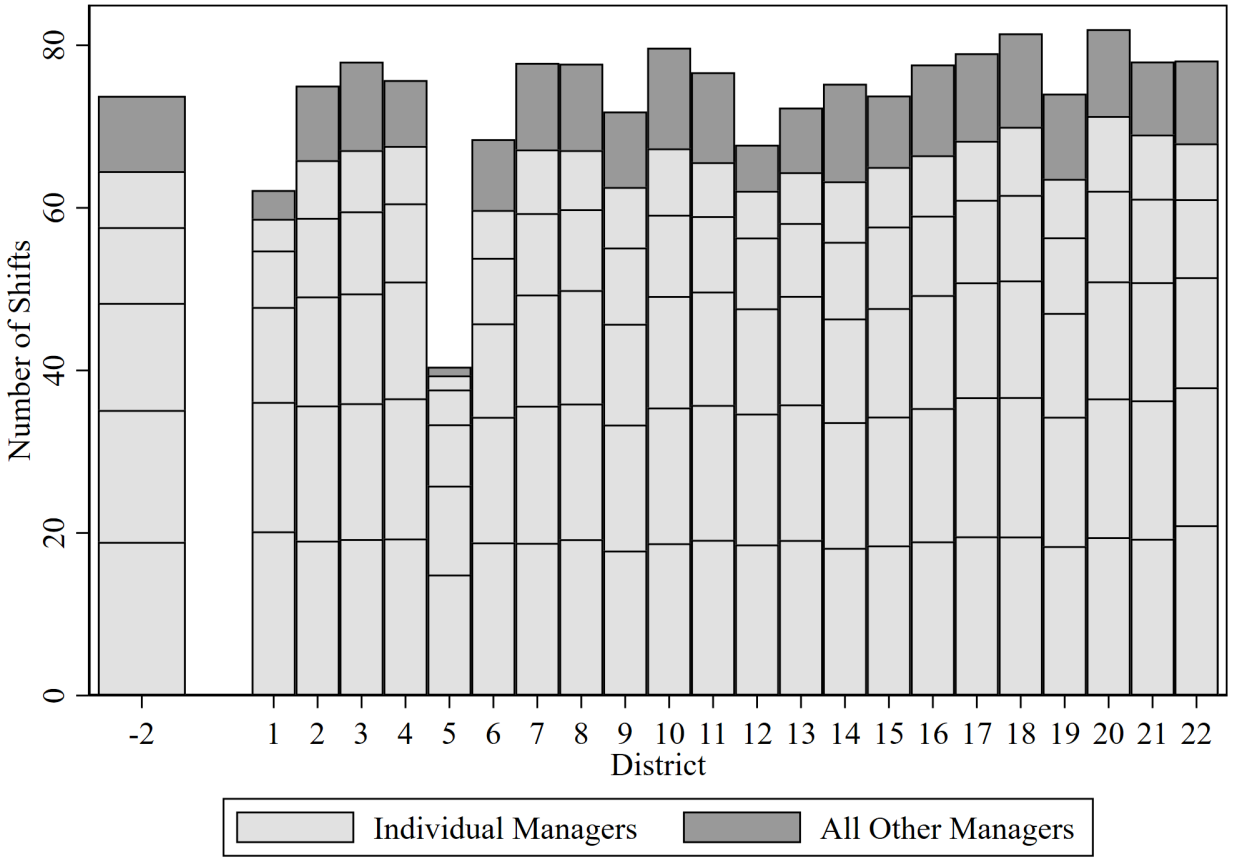
Note: This table presents estimates of the effects of BWC rollout on average monthly arrests. The data used are officers-months with at least 5 shifts and up to 6 months before or 12 months after any rollout event. The first column applies this outcome to the difference-in-differences procedure in Callaway and Sant'Anna (2021) and reports the ATT aggregated across groups and events. Figure A2 presents corresponding event study estimates. The remaining columns calculate the difference in the monthly arrest outcome between officers in the listed category (e.g. those with high prior complaint rates) and officers not in that category. They then repeat the Callaway and Sant'Anna (2021) procedure with this differenced outcome. These categories are not defined for some officers, in which case they are excluded from the estimation. Standard errors are clustered by district and are in parentheses. $^{+}$ $p<0.1$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 5: Effects of BWC Rollout: Force

	(1) All	(2) High Complaint	(3) Low BWC	(4) High Force
ATT	0.000705	0.000637	-0.000680	0.00130
	(0.000487)	(0.000743)	(0.000889)	(0.00111)
Observations	528	528	528	528

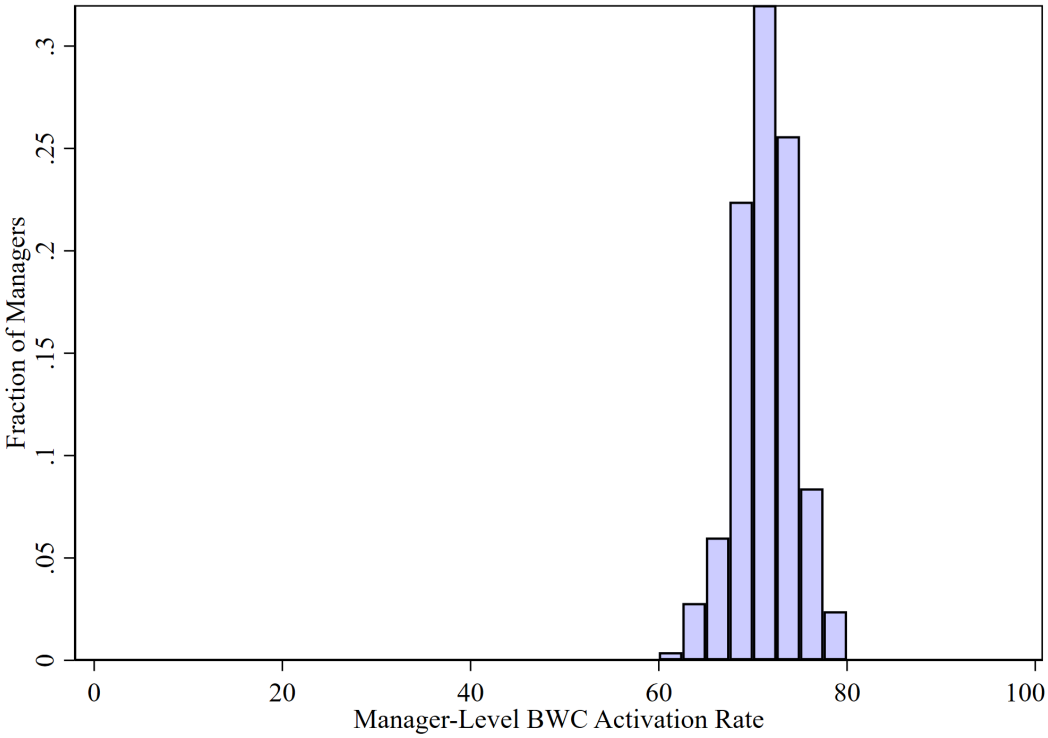
Note: This table presents estimates of the effects of BWC rollout on average monthly uses of force. The data used are officers-months with at least 5 shifts and up to 6 months before or 12 months after any rollout event. The first column applies this outcome to the difference-in-differences procedure in Callaway and Sant'Anna (2021) and reports the ATT aggregated across groups and events. Figure A3 presents corresponding event study estimates. The remaining columns calculate the difference in the monthly force outcome between officers in the listed category (e.g. those with high prior complaint rates) and officers not in that category. They then repeat the Callaway and Sant'Anna (2021) procedure with this differenced outcome. These categories are not defined for some officers, in which case they are excluded from the estimation. Standard errors are clustered by district and are in parentheses. $^{+}$ p<0.1, * p<0.05, ** p<0.01, *** p<0.001.

Figure 4: Managers Shift Coverage



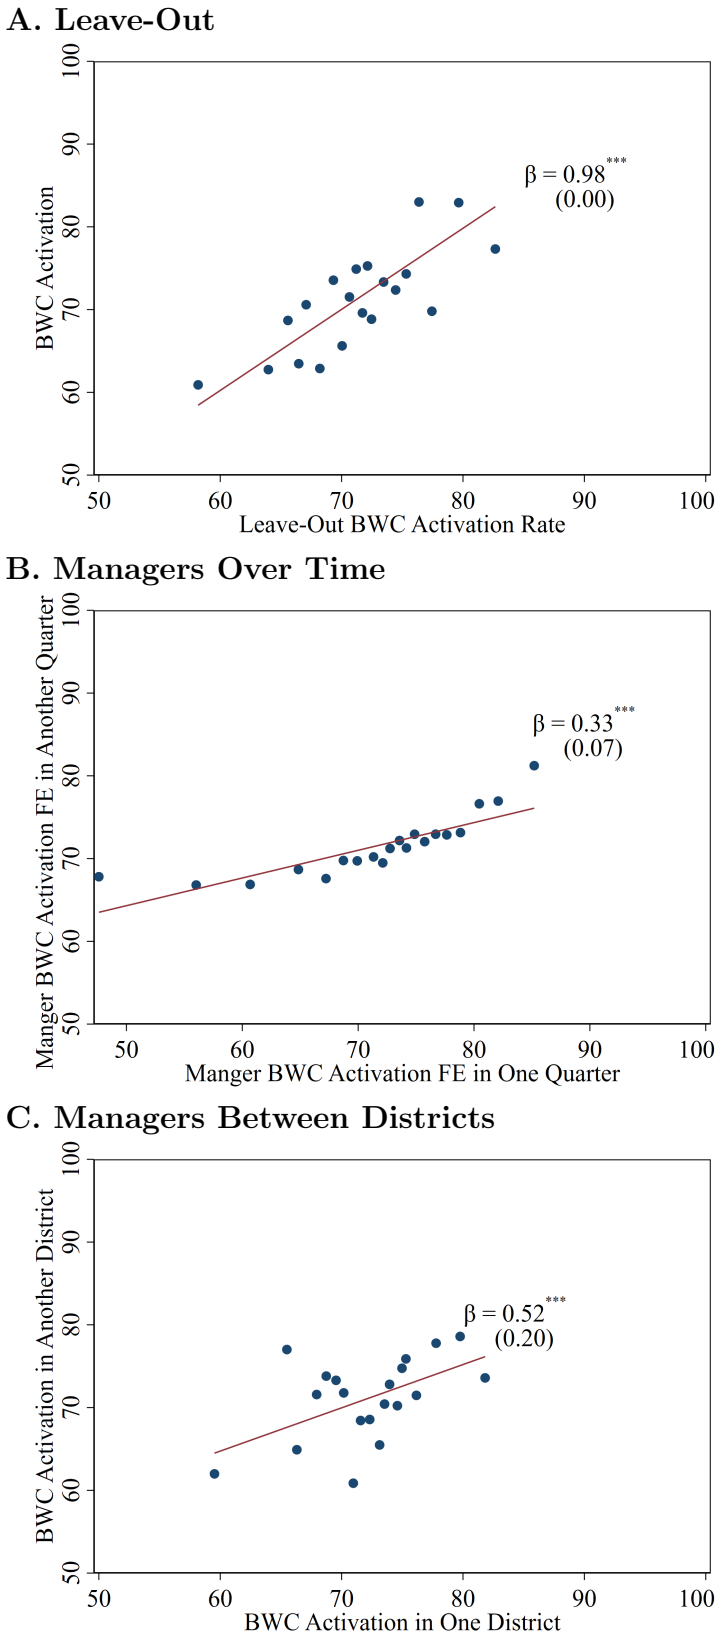
Note: This figure reports the mean number of shifts per month for the top 5 managers in a district in the calender month (light gray bars) along with the total shifts covered by all other managers (dark gray bars), overall (on the left) and separately by Chicago police district (on the right).

Figure 5: Manager Contribution to BWC Activation



Note: This figure reports the distribution of coefficients on officer fixed effects in a regression that controls for district by calender month and charge fixed effects.

Figure 6: Validation: Manager Contribution to BWC Activation



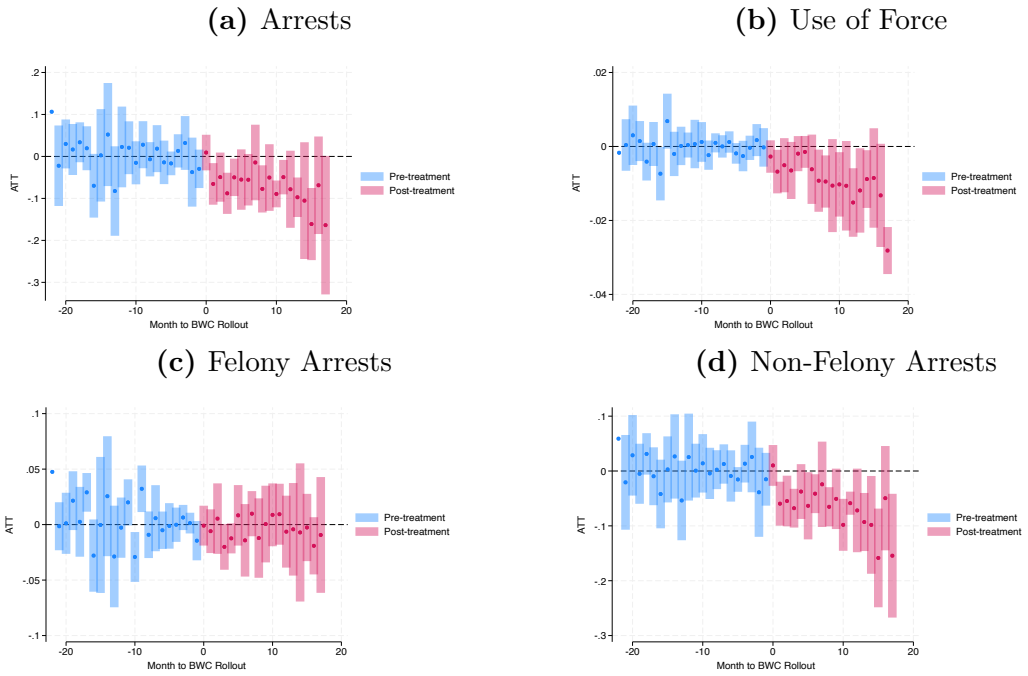
Note: This figure assess how various measures of manager contribution to officer activation rates in some arrests predict manager contribution to officer activation in other arrests. Panel A reports the relationship between the activation rate in one arrest where a manager is in charge and the leave out activation rate in other arrests where the same manager was in charge. Panel B reports the relationship between the activation rate of all arrests where a manager was in charge in one quarter against the activation rate of all arrests where the same manager was in charge in every other quarter (restricting to manager-quarters with at least 200 arrests when they were in charge). Panel C reports the relationship between the activation rate of all arrests where a manager was in charge in one district against the activation rate of all arrests where the same manager was in charge in every other district (restricting to manager-quarters with at least 200 arrests when they were in charge).

Table 6: Heterogeneity in DID Estimates by WOL Strictness

	(1)	(2)	(3)	(4)
Panel A: Effects of Strict WOLs				
	Arrests	Use of Force	Felony Arrests	Non-Felony Arrests
ATT	-0.055**	-0.007+	-0.004	-0.052**
	(0.018)	(0.004)	(0.010)	(0.012)
Observations	355	355	355	355
Panel B: DD under WOLs with Low Strictness				
	Arrests	Use of Force	Felony Arrests	Non-felony Arrests
ATT	0.020	0.003**	-0.005	0.024**
	(0.017)	(0.001)	(0.010)	(0.009)
Observations	355	355	355	355
Panel C: DD under WOLs with High Strictness				
	Arrests	Use of Force	Felony Arrests	Non-felony Arrests
ATT	-0.031**	-0.004	-0.005	-0.027**
	(0.010)	(0.003)	(0.003)	(0.008)
Observations	355	355	355	355
Panel D: All Officers				
	Arrests	Use of Force	Felony Arrests	Non-Felony Arrests
ATT	-0.054**	-0.004+	-0.001	-0.044+
	(0.020)	(0.003)	(0.011)	(0.023)
Observations	355	355	355	355
Panel E: Officers with Low Previous Complaints				
	Arrests	Use of Force	Felony Arrests	Non-Felony Arrests
ATT	-0.056**	-0.008*	-0.006	-0.049**
	(0.018)	(0.004)	(0.010)	(0.013)
Observations	355	355	355	355
Panel F: Officers with High Previous Complaints				
	Arrests	Use of Force	Felony Arrests	Non-Felony Arrests
ATT	-0.023	-0.006	0.022	-0.044+
	(0.031)	(0.004)	(0.015)	(0.023)
Observations	355	355	355	355

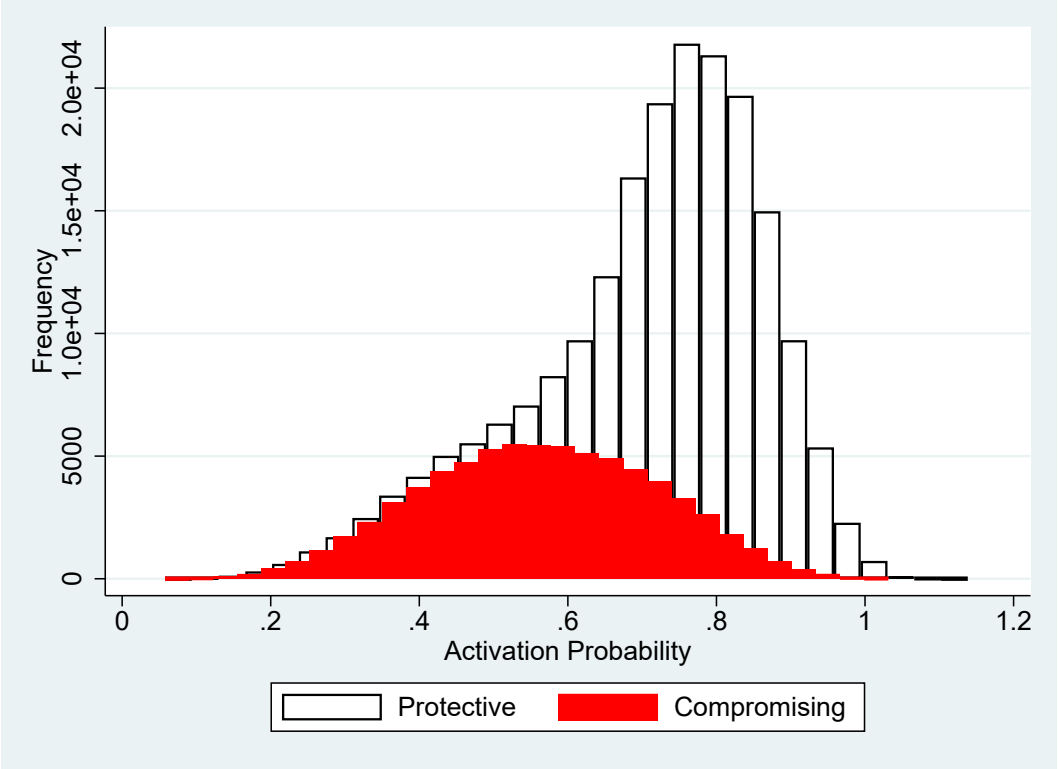
This table presents estimates of the effects of BWC rollout on average monthly arrests of officers supervised by a given manager. The data used are manager-months up to 6 months before or 12 months after any rollout event. The first columns calculate the difference in the monthly arrest outcome between strict and lenient managers, then applies the Callaway and Sant'Anna (2021) procedure with this differenced outcome. The next columns repeats with with force and different types of arrests in place of arrests. Some managers are not defined as strict or lenient because they were not present in post-BWC period, in which case they are excluded from the estimation. Panel A contains main estimates, Panel B and C estimate DD separately for teams under strict and lenient WOLs. Panel D contains the sample of tactical officers who do not neatly map onto regular watches. Panel E and F separately estimate heterogeneity by offices who had low and high previous complaints. Standard errors are clustered by district and are in parentheses. + p<0.1, * p<0.05, ** p<0.01, *** p<0.001.

Figure 7: The Effects of Stringent BWC Oversight



Notes: In this figure, we examine how team outcomes under strict and lenient managers about BWC activation policies evolved around the rollout. For each district and month, we focus on watches managed by managers for whom we can estimate FEs. To construct the outcome variable, within each unit and month, we group watch-days based on whether they were managed by strict or lenient managers. We calculate the average team performance for watch-days under strict managers and for those under lenient managers, and then take the difference to have a unit-month level outcome variable.

Figure 8: Simulated Distribution of Officer Activation Probabilities



Notes: This histogram plots the activation rate of 200,000 officers simulated using the estimated parameters of our structural model. The white bars show the distribution of officers with $\theta_{arr} < \theta_{act}(0)$. The red bars show the distribution of officers with $\theta_{arr} < \theta_{act}(0)$. The aggregate distribution in any bin is the sum of the white and red bars.

Table 7: Counterfactual Results

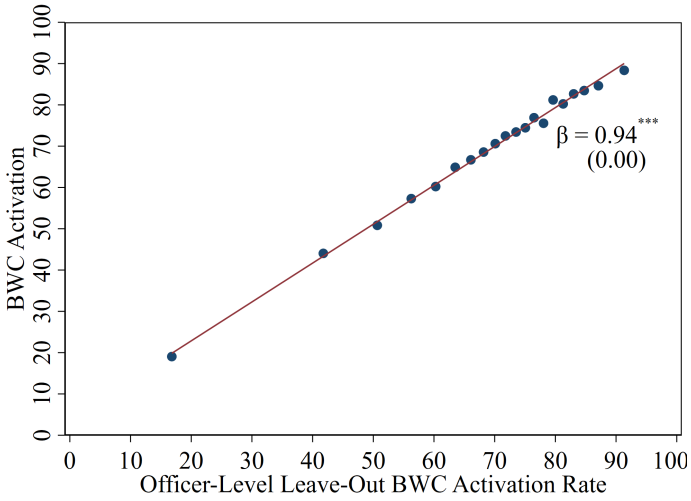
		All	Protective	Compromising
Pre-BWC	Arrests	37%	20%	68%
	Activation	0	0	0
Post-BWC	Arrests	46%	33%	68%
	Activation	70%	79%	65%
High p_h	Arrests	43%	33%	60%
	Activation	75%	78%	70%
Low p_l	Arrests	48%	37%	68%
	Activation	85%	95%	68%

Notes: This table presents averages of officer-level arrest and activation rates in simulated data for some counterfactual scenarios. The Pre-BWC scenario uses our estimated parameters except with $p_h = 0$ and $p_l = 1$. The post-BWC scenario uses our estimated parameters as-is. The High p_h scenario sets $p_h = 0.5$. The Low p_l scenario sets $p_l = 0.05$. Averages are presented for all officers, officers who find activation protective (those with $\theta_{arr} < \theta_{act}(0)$), and those who find activation compromising (those with $\theta_{arr} < \theta_{act}(0)$).

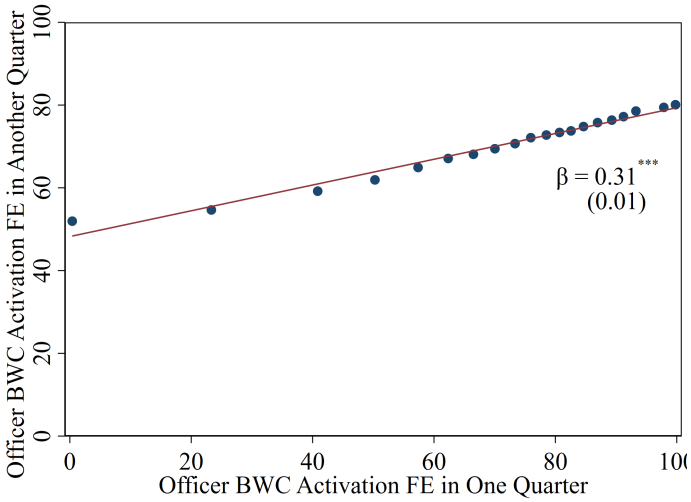
Supplemental Appendix A: Additional Results

Figure A1: Validation of Officer Contribution to BWC Activation

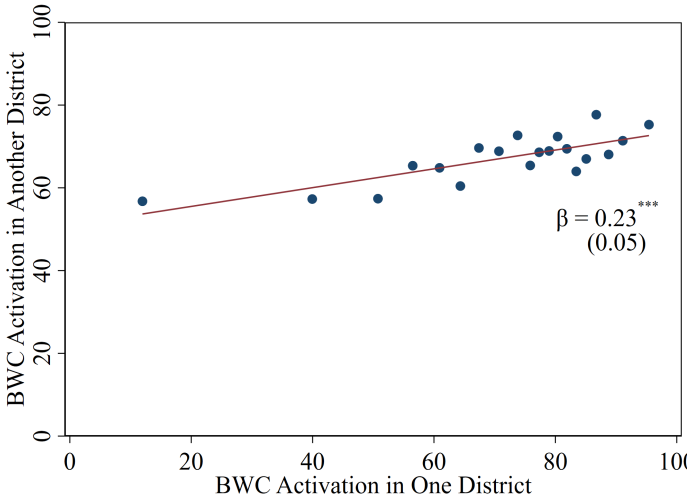
A. Leave-Out Arrest



B. Officers Over Time



C. Officers Between Districts



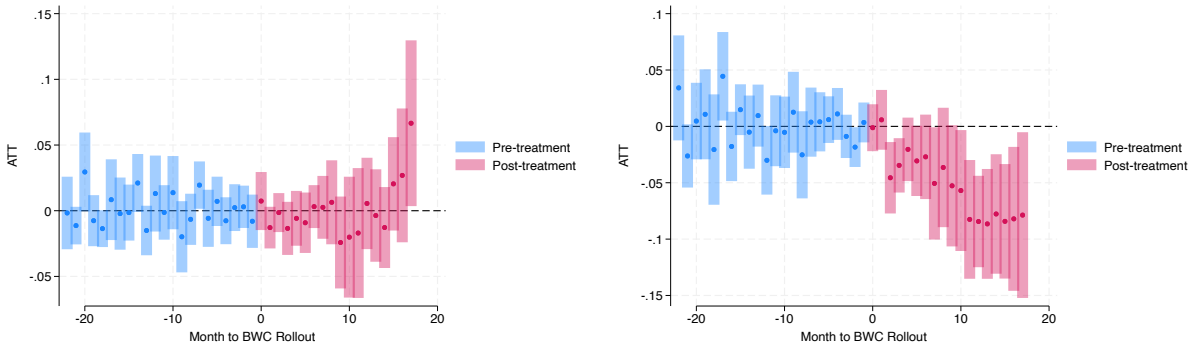
Note: This figure assesses how officer activation rates in some arrests predict the officer's activation in other arrests. Panel A reports the relationship between activation in one arrest and the officer's leave out activation rate. Panel B reports the relationship between an officers activation rate in one quarter against the same officer's activation in every other quarter (restricting to officer-quarters with at least 10 arrests in the quarter). Panel C reports the relationship between an officers activation rate in one district against the same officer's activation rate in every other district (restricting to officer-districts with at least 20 arrests in the district).

Table A1: Differential Activation by Subject Race and Officer Type: Use of Force

	(1)	(2)	(3)
Black Arrestee	.0117***		-.000467
	(.00245)		(.00415)
High Force Officer		.0792***	.0639***
		(.00545)	(.007)
High Force Officer x Black Arrestee			.0207***
			(.0068)
Y mean	.698	.698	.698
N	649,989	649,989	649,989
Time FEs	Y	Y	Y
District FEs	Y	Y	Y
Arrest Categories FEs	Y	Y	Y

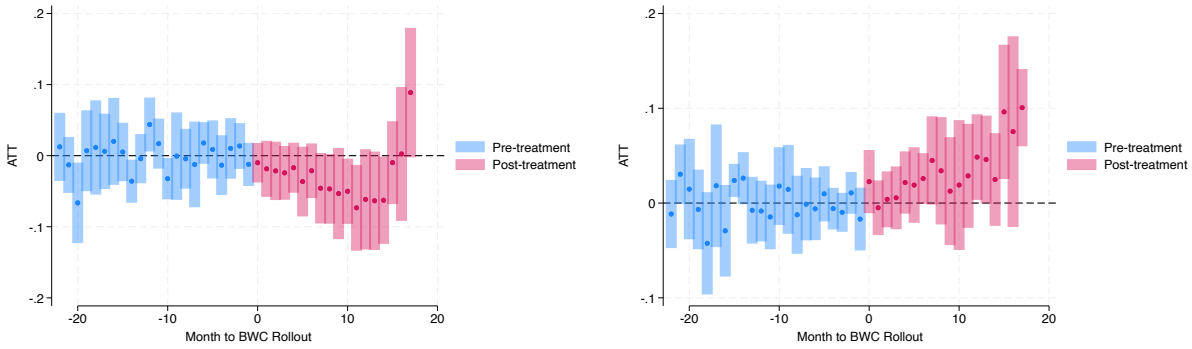
Notes: This table presents descriptive regressions of the relationship between BWC activation, suspect race, and prior officer force rate. Prior uses of force are counted from 2004 and divided by the number of years of active service and data availability. We label an officer as high force if their use of force rate is in the top quartile among officers in our analysis sample. We estimate these regressions in the sample of arrests where officers were working in patrol districts that had already rolled out cameras, and we observe that officer to have activated a BWC at least once. All regressions further condition on the high force measure being available, even when that regression does not include the high force indicator. We do not calculate the high force measure for officers who began working after 2015.

Figure A2: BWC Rollout Event Study: Arrests



(a) All

(b) High Complaint

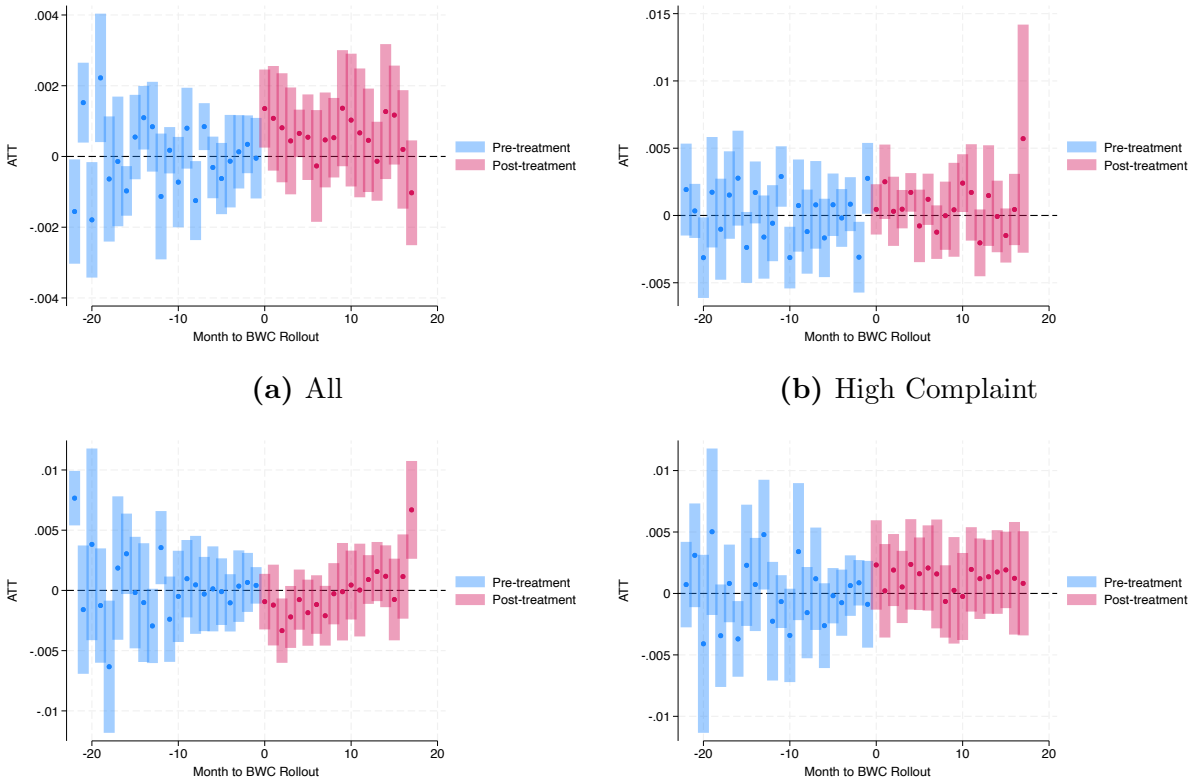


(c) Low BWC

(d) High Force

Notes: This figure presents estimates of the effects of BWC rollout on average monthly arrests. The data used are officers-months with at least 5 shifts and up to 6 months before or 12 months after any rollout event. The first figure applies this outcome to the difference-in-differences procedure in Callaway and Sant'Anna (2021) and reports event study estimates of the event-time-specific ATT. The remaining figures calculate the difference in the monthly arrest outcome between officers in the listed category (e.g. those with high prior complaint rates) and officers not in that category. They then repeat the Callaway and Sant'Anna (2021) procedure with this differenced outcome. These categories are not defined for some officers, in which case they are excluded from the estimation.

Figure A3: BWC Rollout Event Study: Force



(a) All **(b)** High Complaint

(c) Low BWC **(d)** High Force

Notes: This figure presents estimates of the effects of BWC rollout on average monthly uses of force. The data used are officers-months with at least 5 shifts and up to 6 months before or 12 months after any rollout event. The first figure applies this outcome to the difference-in-differences procedure in Callaway and Sant'Anna (2021) and reports event study estimates of the event-time-specific ATT. The remaining figures calculate the difference in the monthly force outcome between officers in the listed category (e.g. those with high prior complaint rates) and officers not in that category. They then repeat the Callaway and Sant'Anna (2021) procedure with this differenced outcome. These categories are not defined for some officers, in which case they are excluded from the estimation.

Table A2: TWFE Effects of BWC Rollout: Arrests

	(1) All	(2) High Complaint	(3) Low BWC	(4) High Force
Treated	-0.179			
	(0.182)			
Triple-Difference		-0.629***	-0.485***	0.317
		(0.148)	(0.124)	(0.199)
Observations	106240	94285	53382	94285

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table presents TWFE estimates of the effects of BWC rollout on average monthly arrests. The data used are officer-months with at least 5 shifts and up to 6 months before or 12 months after any rollout event. The first column applies this outcome to a regression with officer, time, and district fixed effects. The remaining columns present a triple-differences specification that interacts the rollout treatment with the category of interest and includes officer and district-by-time fixed effects. These categories are not defined for some officers, in which case they are excluded from the estimation.

Table A3: TWFE Effects of BWC Rollout: Force

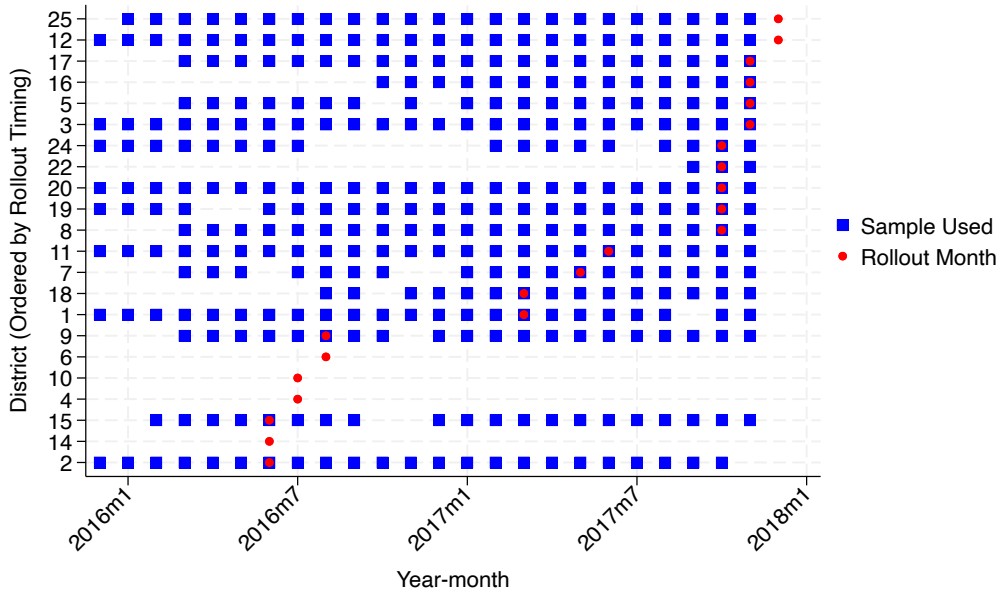
	(1) All	(2) High Complaint	(3) Low BWC	(4) High Force
Treated	0.0265*			
	(0.0111)			
Triple-Difference		0.00572	-0.00194	-0.00881
		(0.00581)	(0.00585)	(0.00912)
Observations	106240	94285	53382	94285

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table presents TWFE estimates of the effects of BWC rollout on average monthly uses of force. The data used are officer-months with at least 5 shifts and up to 6 months before or 12 months after any rollout event. The first column applies this outcome to a regression with officer, time, and district fixed effects. The remaining columns present a triple-differences specification that interacts the rollout treatment with the category of interest and includes officer and district-by-time fixed effects. These categories are not defined for some officers, in which case they are excluded from the estimation.

Figure A4: Sample Used in WOL Analysis



Notes: In this figure, we present the sample used for analyzing WOL strictness. Unit-month observations are excluded if the unit was managed exclusively by one manager type during that month or if the unit is not observed in both the pre- and post-periods. The blue squares indicate the retained sample, while the red marker denotes the rollout month.

Supplemental Appendix B: Activation Labels

In this appendix, we document the details of the procedure we use to label activations by matching arrest records to BWC metadata. Recall that this process is made difficult by the fact that only some arrests can be matched directly to the metadata using explicit event IDs recorded by CPD. To establish the remaining matches, we will compare the timestamps associated with arrests and BWC activations.

The set of arrests that do match cleanly to the BWC data by event ID provide a model for the range of these timing discrepancies. We call these the "event number matches." The event number matches demonstrate the necessity of considering timestamp-based matches where the arrest time falls outside the BWC activation window. Figure B1 plots the distribution of "gaps" between the arrest timestamp and BWC activation window. We calculate the gap in each case as the absolute value of the distance between the arrest timestamp and the endpoints of the BWC activation window, assigning 0 to timestamps that fall inside the window.[18] Positive values indicate the arrest timestamp was recorded before the BWC recording ended. Negative values indicate the arrest timestamp was recorded after the BWC recording ended.

We see that in 45.2% of 425,460 event number matches, the arrest timestamp is outside the BWC activation window. Positive gaps are more common than negative gaps. This may be driven by arrest timestamps that default to when an officer was assigned to or arrived at a call for service, rather than the arrest itself. There are long tails on both ends of the distribution, which we collect in Figure B1 into those that exceed 20 minutes on either side. We therefore find that non-zero gaps are common and occasionally large in the event number matches, and this gives us reason to believe they will also be common among arrests that cannot be matched using the event number.

Our next step is then to use the event number matches to model these timing

[18] Occasionally, more than one BWC activation is associated with a single event. In this case, the activation window is a disjoint interval, but the gap is still calculated as the distance to the nearest endpoint.

discrepancies in a more formal sense by building a statistical model of the difference between arrest time and activation window. We model the distribution of gaps as a mixed Gaussian where each component c has weight π_c. Within each component, the size of the gap is determined by a linear model in characteristics of the arrest:

$$Y_i = \alpha_c + \beta_c X_i + \epsilon_i$$

where i indexes the known matches, Y_i is the gap recorded in a given event number match, α_c and β_c are a component specific constant and coefficient vector, X_i is a vector of arrest characteristics, and ϵ_i is drawn from a component specific Gaussian with standard deviation σ_c. X_i consists of: defendant age and sex, recorded officer role (first officer, second officer, or assisting), year and month FEs, charge category FEs, and the size of the BWC recording in MBs.

We estimate the Gaussian mixture using the flexmix package in R, allowing it to choose the optimal number of components and π_c, α_c, β_c, and σ_c for each. Table B1 presents the estimated parameters for the four estimated components. Of these, three are associated with short gaps of no more than 15 minutes. Only the fourth component suggests the possibility for especially long gaps.

We then apply this model to the set of 488,523 candidate "proximity matches" from our primary analysis sample, discussed below. These are arrests that do not match to the BWC metadata by event ID but do have a BWC activation associated with that officer within 4 hours of the arrest timestamp on either side. Another 145,948 arrests see no BWC activations within this range. Note that a single arrest can be prospectively matched to multiple activations. Taking $F(X_i)$ as the CDF of our estimated mixture, we mark an activation in an arrest if $1 - F(X_i) > 0.05$ for at least one of its candidate proximity matches.

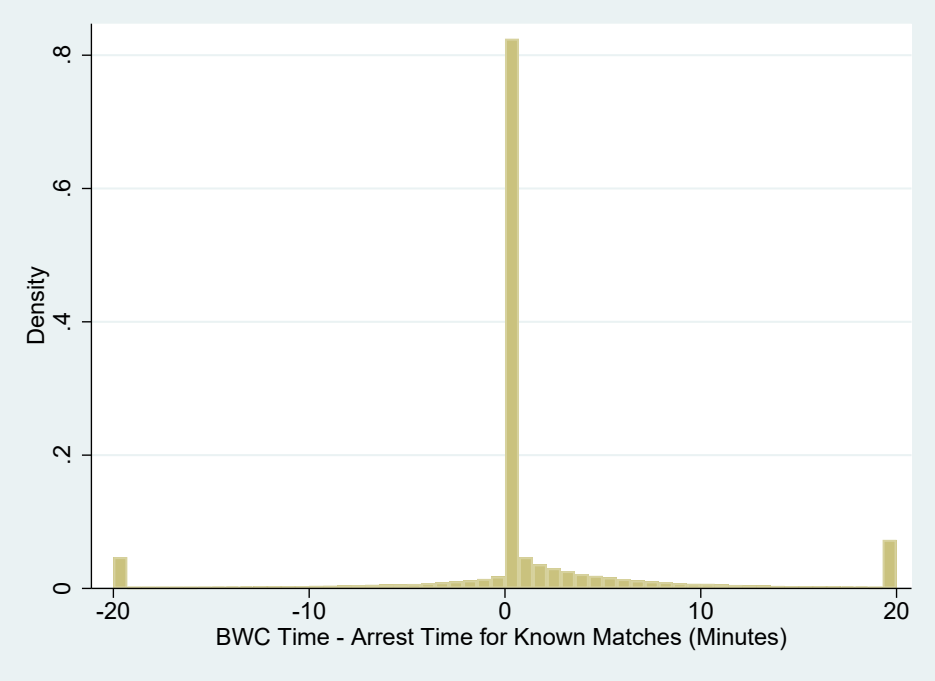
Figure B2 summarizes the results of this procedure. It is a histogram of measured gaps that stacks event number matches, accepted proximity matches, and rejected proximity matches. Observations with gaps of 0, of which there are 397,785, are omitted to preserve

scale, and gaps are topcoded at 100 minutes. Up to 20 minutes, event number matches are common, and most proximity matches are accepted. The rule for accepting proximity matches is not a simple cutoff in time because our model allows for arrest characteristics to explain more or less of the observed gap. By 30 minutes, event number matches are nearly unheard of, and most proximity matches are rejected.

Figure B3 plots the distribution of arrests by the probability of a BWC activation assigned by this procedure. The vertical red line shows our standard 0.05 cutoff for accepting a match. From this graph, it is apparent that the distribution of probabilities created by our procedure is bimodal, with most arrests being assigned either a very high or very low probability. Our 0.05 cutoff divides a sizable set of arrests with a very low, but nonzero, activation probability from those with a more moderate probability. Importantly, loosening our standard to e.g. 0.01 would include many more of these low-probability matches, while tightening it to e.g. 0.1 would exclude only a small number of moderate-probability matches.

These moderate probability matches correspond to the 2nd and 3rd components in the mixed Gaussian model described in Table B1 with mean gaps of about 150 and 850 seconds, respectively. Our mixed Gaussian model puts a high weight on its 1st component, describing near perfect matches, so the 2nd and 3rd components, even though they represent plausible gap lengths, can have no more than a score of approximately 0.4. Our threshold of 0.05, in essence, marks gaps that could have been generated by the first 3 components as activations and gaps from the 4th component as non-activations. Using the score directly in place of this binary application label has the effect of down-weighting activations with gaps in the 2nd and 3rd components.

Figure B1: Distribution of Arrest-Activation Gaps in Event Number Matches



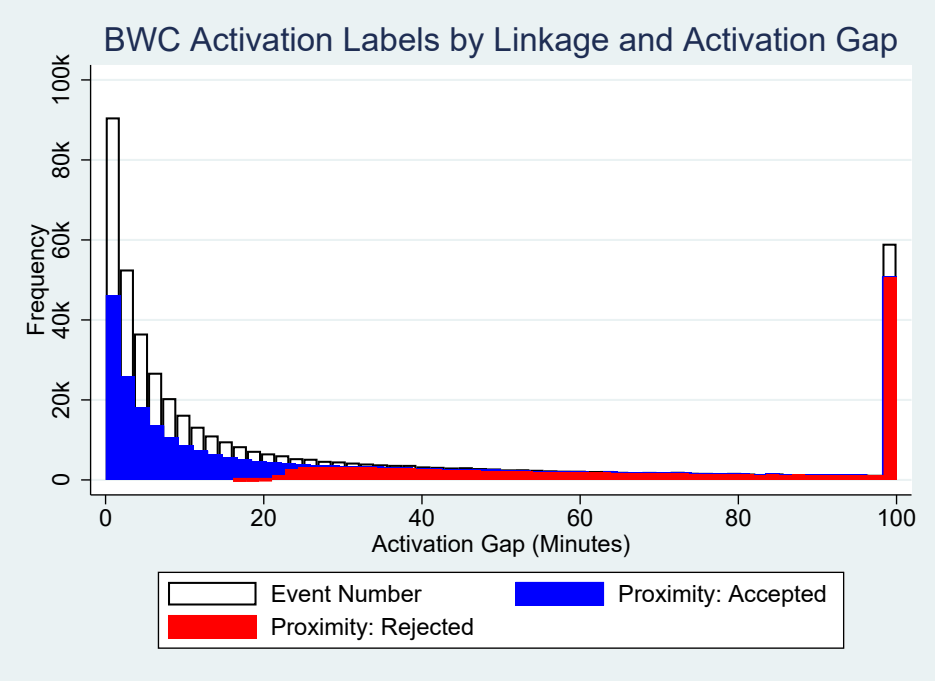
Notes: This figure plots a histogram of arrests by the difference between their timestamp and the closest endpoint of a BWC recording logged by the officer who made the arrests. Only arrest-BWC matches confirmed by event ID numbers are included. Gap magnitudes are topcoded at 20 minutes. 0 indicates the arrest timestamp fell within the BWC activation window. Positive values indicate the arrest timestamp was recorded before the BWC recording ended. Negative values indicate the arrest timestamp was recorded after the BWC recording ended.

Table B1: Components of Activation Gap Model

Component	Mean Gap (seconds)	SD	Weight
1	0.10	1.01	0.56
2	149.26	98.57	0.24
3	845.68	407.80	0.15
4	5419.20	3119.50	0.05

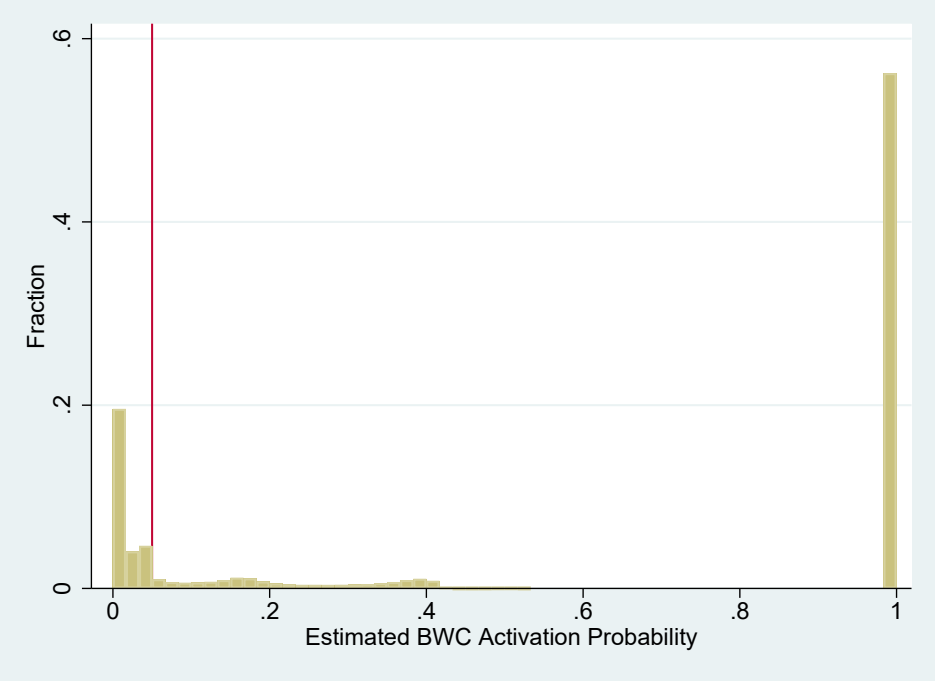
Notes: This table presents information about the four components estimated in our Gaussian mixture model of the gaps between arrest and BWC timestamps. The first two columns capture the mean and standard deviation of each component. In addition to these main effects, we estimate but do not present effects from arrest characteristics: defendant age and sex, recorded officer role (first officer, second officer, or assisting), year and month FEs, charge category FEs, and the size of the BWC recording in MBs. The last column is the weight placed on that component in the mixture model.

Figure B2: Distribution of Arrest-Activation Gap Sizes by Match Type



Notes: This figure plots a histogram of arrests by the absolute difference between their timestamp and the closest endpoint of a BWC recording logged by the officer who made the arrests. Arrests with timestamps within a BWC interval are excluded. Activation gaps are topcoded at 100 minutes. The bars are a stacked histogram by type of match. Event number matches directly use the unique ID number assigned to arrests and the corresponding footage to match arrest and BWC record. Proximity matches look for nearby recordings by the same officer. A proximity match is accepted if it was plausibly generated by a model fit to the activation gaps seen among event number matches and rejected otherwise.

Figure B3: Distribution of Estimated BWC Activation Probabilities



Notes: This figure shows the distribution of BWC activation probabilities estimated for arrests in our sample. For arrests matched to an activation via event numbers, this probability is mechanically 1. For arrests with no activation within 4 hours on either side, this probability is mechanically 0.For all other arrests, this probability is calculated using the CDF of the mixed Gaussian model fit to the event number matches. The vertical red line at 0.05 shows our cutoff for labeling an activation.

Supplemental Appendix C: Model Details

Setup and Arrest Decision

Consider an officer who must decide whether to make an arrest. The reward for making the arrest is $w > 0$, relative to a normalized payout of 0 for not making the arrest (or deferring to colleagues). If the officer makes the arrest, they may be forced to pay a cost $c > 0$ if a misconduct investigation finds them to be in the wrong. In each potential arrest, i, The ex-ante probability that the officer actually commits misconduct is $\theta_i \sim F(\theta)$. If the officer actually commits misconduct, this is discovered with probability q_T. If the officer *did not* actually commit misconduct, an investigation may still conclude that they did (or otherwise compromise their career or reputation) with probability q_F. Define the difference in these probabilities as $\Delta = q_T - q_F$. We will assume that $\Delta > 0$. The probability of being found culpable is higher if the officer actually engaged in misconduct.

Putting this together, the officer makes the arrest if

$$w - c(\theta_i q_T + (1 - \theta_i)q_F) > 0$$

Rearranging, this gives the condition:

$$\theta_i < \frac{\frac{w}{c} - q_F}{\Delta} \equiv \theta^*_{arr}$$

This analysis leads to a reasonable conclusion. Officers will only make an arrest up to a certain threshold of misconduct risk. This threshold is higher if the punishment for misconduct is small relative to the cost, or if the probability of being caught conditional on misconduct is low.

Monitoring Technology and the Decision to Activate

Let us now consider a monitoring technology, m. For mathematical simplicity, we will think of this monitoring technology as continuous even though the actual technology we have in mind, body-worn cameras, is discrete.

Activating this monitoring technology affects both misconduct investigation prob-

abilities. First, the probability of a misconduct finding conditional on actual misconduct rises, $\frac{\partial q_T}{\partial m} = q_T^{'} > 0$. This is because investigations with a recording are more accurate, and hence more likely to correctly detect the misconduct. Perhaps more importantly, the ease of gathering evidence when BWC recordings are available may increase the probability of any investigation happening at all, which also contributes to an increase in q_T.

Conversely, the probability of a misconduct finding conditional on no actual misconduct falls $\frac{\partial q_F}{\partial m} = q_F^{'} < 0$. Investigations which occur are more accurate, decreasing the chance of finding misconduct when there is none to find. We assume this effect is greater than any impact of more frequent investigations. This aligns with accounts about BWC adoption in police departments, where police officers were convinced to approve adoption by their potential to decrease false accusations and misconduct findings. Combining these effects of BWC activation, we can define $\Delta^{'} = q_T^{'} - q_F^{'} > 0$.

Finally, consider a direct, idiosyncratic cost or benefit of BWC activation in each arrest, $\epsilon_i \sim G(\epsilon)$. When $\epsilon_i < 0$, the officer faces some cost to activating his BWC, which we will interpret expansively to include forgetfulness and equipment failures. When $\epsilon_i > 0$, the officer has a particular incentive to activate his BWC, perhaps because another person on the scene specifically requested it.

The timing in our model is as follows:

1. The officer observes θ_i and ϵ_i
2. The officer chooses whether to activate their BWC
3. The officer chooses whether to make an arrest

In principle, the decisions (2) and (3) can be treated as simultaneous because no information is revealed between them, but they will be easier to analyze in this sequence. This is helped by the fact that we are not interested in the officer's activation behavior in cases where he chooses not to make an arrest. We can therefore analyze the activation decision as if the

officer will make an arrest, then if the officer does not arrest given that activation decision, we will categorize that case simply as "No Arrest."

We can now ask when officers choose to activate body-worn cameras. Specifically, this is about when activating the camera is in an officer's personal best interest, leaving aside other rules that would compel activation. This is when activation increases the officer's expected utility from making an arrest:

$$\frac{\partial}{\partial m}[w - c(\theta_i q_T + (1-\theta_i)q_F) + \epsilon_i m] > 0$$

Evaluating the derivative and rearranging yields:

$$\theta_i < \frac{\frac{\epsilon_i}{c} - q_F^{'}}{\Delta^{'}} \equiv \theta^*_{act}(\epsilon_i)$$

Once again, the officer's decision can be cast as a threshold rule in the probability of actual misconduct. When this is low, the decrease in false positives offered by BWC activation is attractive because these pose the main threat of a misconduct finding. As θ_i increases, the increase in true positives becomes more salient until eventually BWC activation is no longer optimal.

We can also ask how BWC activation affects whether an officer is inclined to make an arrest. That is, taking activation as given, does having the camera running make an officer more or less likely to make an arrest. We can get at this by examining how θ^*_{arr} changes with m:

$$\frac{\partial \theta^*_{arr}}{\partial m} = \frac{-\Delta q_F^{'} - (\frac{w}{c} - q_F)\Delta^{'}}{\Delta^2} = \frac{-\Delta q_F^{'}}{\Delta} - \frac{\theta^*_{arr}\Delta^{'}}{\Delta} = \frac{\Delta^{'}}{\Delta}(\theta^*_{act}(0) - \theta^*_{arr})$$

This expression reveals how the change in the arrest threshold depends on the relationship between θ^*_{arr} and $\theta^*_{act}(0)$. Depending on which is larger, it may be positive or negative. Importantly, the relevant activation threshold is the one where $\epsilon_i = 0$, meaning the officer faces no special incentive to activate one way or the other. Finally, it is useful to define the

officer's arrest threshold while monitored θ^*_m:

$$\theta^*_m = \theta^*_{arr} + dm\frac{\partial\theta^*_{arr}}{\partial m} = dm\frac{\Delta^{'}}{\Delta}\theta^*_{act}(0) + (1 - dm\frac{\Delta^{'}}{\Delta})\theta^*_{arr}$$

Where dm captures the the total change in monitoring due to BWC activation.

Protective Activation

We will first consider the case where $\theta^*_{arr} < \theta^*_{act}(0)$, shown in Panel A of Figure C1. When $\epsilon_i = 0$, then $\theta^*_{act}(0)$ accurately reflects the officer's activation threshold. In this case, the set of encounters where the officer would make an arrest is a subset of the encounters where the officer would activate a body-worn camera. Consequently, the officer will activate his camera in every observed arrest. Furthermore, the officer's monitored arrest threshold, θ^*_m, exceeds his original arrest threshold, so he will make additional arrests relative to the pre-BWC baseline.

Intuitively, the officer in this scenario finds the body-worn camera to be protective. He always finds it beneficial to be recorded due to the reduction in the probability that he will be punished for a spurious complaint. Furthermore, the ability to create a recording empowers the officer to engage in some scenarios he would have previously avoided for being too risky. When $\epsilon_i > 0$, the analysis is largely unchanged.

Panel B of Figure C1 depicts an ϵ_1 that is negative but small, such that $\theta^*_{arr} < \theta^*_{act}(\epsilon_1) < \theta^*_m$. In this case, the officer still activates in every arrest he makes, but he neglects to activate for some of the riskiest encounters and consequently also declines to make an arrest in them.

Panel C of Figure C1 depicts an ϵ_2 that is negative and large, such that $\theta^*_{act}(\epsilon_2) < \theta^*_{arr} < \theta^*_m$. In this scenario, the officer now neglects to activate even in some encounters where he makes an arrest. He makes no more arrests than he did in the pre-BWC baseline.

The case where $\theta^*_{arr} < \theta^*_{act}(0)$ thus captures some of the key patterns in the data we are seeking to explain. Officers who find body-worn cameras protective may neglect to activate their BWC simply due to a low draw of ϵ_i. Importantly, this is relevant only to

encounters with $\theta_i < \theta^*_{arr}$. For the riskiest encounters in which these officers may attempt an arrest, those with $\theta^*_{arr} < \theta_i < \theta^*_m$, the camera will always be activated if the arrest occurs. This has the potential to create a positive correlation between complaint risk and activation rates.[19]

Compromising Activation

Next we consider the case where $\theta^*_{act}(0) < \theta^*_{arr}$, shown in Panel D of Figure C1. When $\epsilon_i = 0$, then the set of cases where the officer would activate his BWC is a subset of those cases where he would make an arrest. Consequently, he does not activate his BWC during all arrests, specifically declining to do so for arrests with the highest θ_i. In addition, the officer's monitored arrest threshold, θ^*_m, is lower than his original arrest threshold. So long as $\epsilon_i \leq \theta^*_m - \theta^*_{act}(0)$, this does not matter to the officer's actual arrest behavior because he will already be declining to activate at the frontier of his arrest decision

Panel E of Figure C1 depicts an ϵ_3 that is positive and large enough such that $\theta^*_m < \theta^*_{act}(\epsilon_3) < \theta^*_{arr}$. In this case, the set of arrests where the officer chooses to make an arrest but not activate shrinks to just those $\theta^*_{act}(\epsilon_3) < \theta_i < \theta^*_{arr}$. For all $\theta_i < \theta^*_m$, the officer now activates his BWC while making an arrest. For $\theta^*_m < \theta_i < \theta^*_{act}(\epsilon_3)$, the officer opts to make no arrest at all. In these encounters, the officer is compelled to turn on his BWC, but he feels that the footage it captures would be compromising to him should he attempt to make an arrest. Officers with $\theta^*_{act}(0) < \theta^*_{arr}$ walk a fine line between activating their BWC when pressed to do so while also ensuring the camera does not capture any compromising behavior.

This dynamic generates yet more of the key patterns we are trying to explain in the data. Officers who find body-worn cameras compromising will neglect to activate in some arrests even when they receive a neutral draw of ϵ_i. Furthermore, the cases in which they

[19]This is true especially if extremely low values of ϵ_i (those associated with a total inability to activate) are common relative to slightly low values of ϵ_i (those associated with a disincentive to activate). In those circumstances, the officer will fail to activate given a bad draw of ϵ_i regardless of θ_i, and there will be no negative correlation between θ_i and activation among cases with $\theta_i < \theta^*_{arr}$

choose not to activate will be those with the highest probability of generating a complaint. Furthermore, if pressed to activate their camera, they may opt to forego some arrests they would have otherwise made in the pre-BWC baseline.

Shifting Incentives

Given this framework, we can next consider the impact of policies that incentivize BWC activation. These affect officers in the same way as ϵ_i, so we can simply think of them as a shift towards higher values of ϵ_i. The exact nature of this shift depends on the policy considered. An initiative to ensure that cameras were always charged and operational, or to get officers into the habit of activating them, would replace some very low values of ϵ_i with 0. An initiative to discipline officers who fail to activate could uniformly shift the distribution of ϵ_i to the right.

For officers who find BWC activation protective, the main impact of higher values of ϵ_i would be to reduce idiosyncratic non-activations and encourage more arrests in the upper reaches of θ_i. For officers who find BWC activation compromising, higher values of ϵ_i will wipe out some strategic non-activations but create some scenarios where officers decline to make an arrest they would have made in the pre-BWC baseline. In a population that combined both types of officers, we should therefore expect activation rates to rise unambiguously, but the predicted effect on arrests is ambiguous.

Structural Estimation
Pre-BWC Environment

We first describe the arrest and discipline environment without BWCs. When an officer makes an arrest and commits misconduct, there is a probability q_T that this misconduct is detected and punished. When an officer makes an arrest and does not commit misconduct, there is a probability q_F that the officer is punished anyway. To relate the theory to observable statistics of complaints, we need a third possibility, which is the probability that misconduct is charged but not sustained when there is no misconduct.[20] Call this

[20]We assume all actual misconduct is charged even if it is not sustained

probability ψ.

Therefore, given a rate of complaints per arrests, C:

$$C = E[(1-\theta)(q_T + \psi) + \theta | \theta < \theta_{arr}]$$

and given a rate of sustained complaints S:

$$S = \frac{E[(1-\theta)q_T + \theta q_F | \theta < \theta_{arr}]}{C}$$

We assume that θ is uniformly distributed between 0 and A. Furthermore, we restrict our calculation of C and S to officers at or above the 95th percentile in arrests per shift to minimize selection from the requirement that θ fall below the arrest threshold θ_{arr}. Removing that conditioning entirely, we can calculate:

$$E[\theta] = \frac{CS - q_F}{q_T - q_F}$$

and furthermore

$$A = 2E[\theta]$$

and while it is not crucial to our analysis, ψ is given by:

$$\psi = C - E[\theta] - (1 - E[\theta])q_F$$

Next we turn our attention to the officer-level distribution of arrest incentives, $\frac{w}{c}$. We assume this is normally distributed with mean μ and standard deviation σ. Given R, the mean of the distribution of observed arrests per opportunity and V, the variance of that distribution, we can make use of the following facts:

$$Pr\{\theta < \theta_{arr}\} = \frac{\frac{c}{w} - q_F}{q_T - q_F}$$

$$E[Pr\{\theta < \theta_{arr}\}] = \frac{\mu - q_F}{q_T - q_F} = R$$

$$Var[Pr\{\theta < \theta_{arr}\}] = \frac{\sigma^2}{(q_T - q_F)^2} = V$$

Where the last two equations can be solved for μ and σ.

Thus, given data for C, S, R, and V, only q_T and q_F are left as parameters that must be estimated.

Post-BWC Environment

In the model, body-worn camera activation increases the probability of punishment conditional on misconduct by an amount q'_T and decreases the probability of punishment conditional on no misconduct by q'_F. These combine in the activation threshold: $\theta^*_{act}(0) = \frac{-q'_F}{q'_T - q'_F}$. With activation data alone, these two parameters are not separately identified, and the first phase of estimation will focus on estimating $\theta^*_{act}(0)$ directly. We will later differentiate between q'_T and q'_F using what we know about how arrest patterns differentially changed following the introduction of BWCs.

The model allows for an arbitrary ϵ shifter to activation incentives. For the structural estimation, we allow only extreme values of ϵ. It is $-\infty$ with probability p_l, ∞ with probability p_h and 0 with probability $1 - p_l - p_h$. In essence, officers are randomly forced to not activate in proportion p_l of arrests and forced to activate in proportion p_h of arrests.

Given a draw from the distribution of $\frac{w}{c}$ and the parameters $\{q_T, q_F, \theta^*_{act}(0), p_l, p_h\}$ it is possible to estimate a resulting arrest and activation rate for a hypothetical officer. To the estimated activation rate, we add an ad-hoc noise term, u, which is normally distributed with mean 0 and standard deviation σ_u. This ad-hoc noise term accounts for variation in incentives across arrests within the same officer, non-extreme values of ϵ, and general finite sampling error. The final vector of 6 parameters to estimate is thus $\{q_T, q_F, \theta^*_{act}(0), p_l, p_h, \sigma_u\}$.

Estimation Procedure

Some aspects of the pre-BWC environment can be estimated directly in data. We use data from the period 2013-2015. First, we estimate the distribution of arrests per shift among officers who worked at least 100 shifts in that period. We use the 95th percentile of this distribution (0.783 arrests per shift) as a baseline rate of *available* arrests per shift. This assumes that the 5% of officers who are most prolific are foregoing no arrests. Given this,

we calculate that the average officer in 2013-2015 makes proportion $R = 0.354$ of available arrests with a variance of $V = 0.108$. Restricting attention to just those 5% most prolific officers, to avoid selection into less risky arrests, we find a rate of $C = 0.0036$ complaints per arrest and $S = 0.0098$ sustained complaints per complaint.

We estimate the remainder of the model using the method of simulated moments, generating 200,000 observations per simulation. Each observation corresponds to an officer with their own draw from the distribution of $\frac{w}{c}$. Within an officer, individual arrest opportunities have different draws from the distribution of θ, but we do not simulate these draws directly, handling them instead with analytic expressions. We model sampling variance in activation rates in the following way: given an officer with analytic activation rate, r_i, and who makes a number of arrests, n_i drawn from $U(50, 100)$, the observed activation rate for that officer is $b_i + u_i$. Here, b_i is drawn from $Binomial(n_i, r_i)$ and u_i is drawn from $N(0, \sigma_u)$.

Our goal is to select parameters in $\Omega = \{q_T, q_F, \theta^*_{act}(0), p_l, p_h, \sigma_u\}$ in order to match the distribution of officer-level activation rates shown in Figure 3. We select 9 moments of that distribution: the 1st, 5th, 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles and weight them equally in the objective function. Call these p_1 through p_{99} and their simulated equivalents $\tilde{p}_1$ through $\tilde{p}_{99}$. We emphasize extreme moments of the distribution to match the behavior in the tails of the officer-level activation rate distribution, which we find most distinctive. To this objective function, we add a quadratic penalty term for having the 99th percentile of the activation distribution exceed 1 and another for having $p_h + p_l > 1$. Thus, our final objective function is:

$$O = \sum_{i=1...99} (p_i - \tilde{p}_i)^2 + \max(100(\tilde{p}_{99} - 1), 0) + \max(100(p_h + p_l - 1), 0)$$

We then minimize O with respect to Ω, first with a global optimization using multi-level single-linkage Kucherenko and Sytsko (2005), then "polishing" the result with a local simplex algorithm Nelder and Mead (1965).[21].

[21]In both cases, the algorithm is as implemented in the NLopt R package

Finally, we do a second stage of estimation to pin down q'_F vs. q'_T. We consider values of q'_F from 0 to -1 times our estimate of q_F and calculate: average arrests among officers who find activation compromising prior to BWCs (which we impose by fixing p_l at 1 and p_h at 0), average arrests among officers who find activation protective prior to BWCs, and the same averages in our estimated BWC environment. We then calculate the relative change in arrest rates for compromising activators when BWCs are introduced and compare this to 0.152 (the coefficient on low activation in Table 4, normalized by 3.55 arrests per month).

Sketch of Identification

This model can be seen as fitting an augmented Gaussian mixture to the distribution of officer-level activation rates with each component in the mixture corresponding to a type of officer, either those who find activation protective or those who find it compromising.

The distribution of θ_{arr} across officers is normal with mean $\frac{R-q_F}{q_T-q_F}$ and standard deviation $\frac{\sqrt{V}}{q_T-q_F}$. Therefore, it can be fully parameterized by q_T and q_F, subject to their domain restrictions. Given any pair (q_T, q_F), $\theta^*_{act}(0)$ determines the fraction of officers who consider activation protective rather than compromising. Within the group of officers who find activation protective, all officers activate at rate $1 - p_l$, and the noise term σ_u - along with the natural binomial variation - controls the spread of observed activation rates for this group.

The component corresponding to officers who find activation compromising is augmented to allow a long tail in this group. The θ_{arr} distribution induced by (q_T, q_F) produces officers who find activation compromising and who have a range of effective activation rates. Furthermore, there is a mass of officers at the lowest activation rate. (q_T, q_F) together control the spread of this range and the relative size of the mass at the lowest activation rate. p_h can then shift this entire distribution further to the right if necessary.

The model as it is parameterized estimates $\theta^*_{act}(0)$ directly. Given that $\theta^*_{act}(0) = \frac{-q'_F}{q'_T-q'_F}$, the changes in misconduct detection induced by BWC activation are not separately

identified by the distribution of activation. However, $q_F^{'}$ also appears in the definition of θ_m, which in turn controls arrest rates and how they react to the introduction of BWCs. We can therefore recover $q_F^{'}$ by matching our result on the differential arrest effect of BWC rollout from Section 5.

Estimation Output

Table C1 presents the actual moments of the officer-level activation distribution, our simulated moments, and the associated parameter estimates. The primary issue with the fit is that we cannot match the extent of the left tail. Our simulated 1st percentile, 0.268, is substantially above the actual 1st percentile, 0.184. Making either the distribution of $\frac{w}{c}$ or θ more flexible in a way that allows some officers to have particularly high complaint risk could address this problem at the cost of degrees of freedom.

We include our estimate of A in Table C1, even though it is not a free parameter, to emphasize the fact that misconduct is rare. Even the arrest with the highest θ has only a 0.02% chance of misconduct. Furthermore, our estimated values of q_T and q_F are consistent with the fact that punishment for misconduct is even rarer. Officers who commit misconduct have an 18.1% chance of facing a sustained complaint as a consequence. The probability of punishment given no misconduct is four orders of magnitude smaller, though it is still relevant to officer behavior because arrests without misconduct are quite common.

Our estimate of $\theta_{act}^*(0)$ is difficult to interpret directly, but it drives the division of officers between those who find BWC activation protective, 63.5% of our simulated sample, and those who find it compromising, the remaining 36.5%. In the former sample, the primary driver of activation rates is p_l. Officers who find activation protective will activate in every encounter unless there is some outside intervention. In the latter sample, p_h has the opposite effect, driving activation rates up. We find that $p_l = 0.217$, consistent with the fact that very few officers activate in all or even 90% of arrests. We find a much smaller value of $p_h = 0.1\%$, which indicates that officers are very rarely compelled to activate in situations

where they otherwise would not.[22] The ad-hoc noise parameter $\sigma_u = 0.08$ then smooths out the distribution of activation rates.

Finally, the secondary estimation to match our findings in Table 4 hits a corner. We cannot match the changes observed in our data even when q'_F is nearly $-q_F$, i.e. BWC activation nearly eliminates incorrect punishment when officers do not commit misconduct. The values for q'_F and q'_T in Table C1 reflect this. $q'_F = -0.0000141$, nearly the additive inverse of q_F. We estimate a similarly large $q'_T = 0.131$, but this still leaves substantial room for actual misconduct to go unpunished, even when captured on video. At this corner, the differential change in arrest probability across protective and compromising officers when BWCs are introduced is .137. We are continuing to refine this model and hope to find a specification that both fits the data better and avoids this corner solution.

[22] We view the pull towards a low value of p_h to be a deficiency of the model as it is currently parameterized, and we are actively exploring alternative specifications, especially those that allow for a small mass of officers with very high $\frac{w}{c}$.

Figure C1: Theoretical Arrest and Activation Thresholds

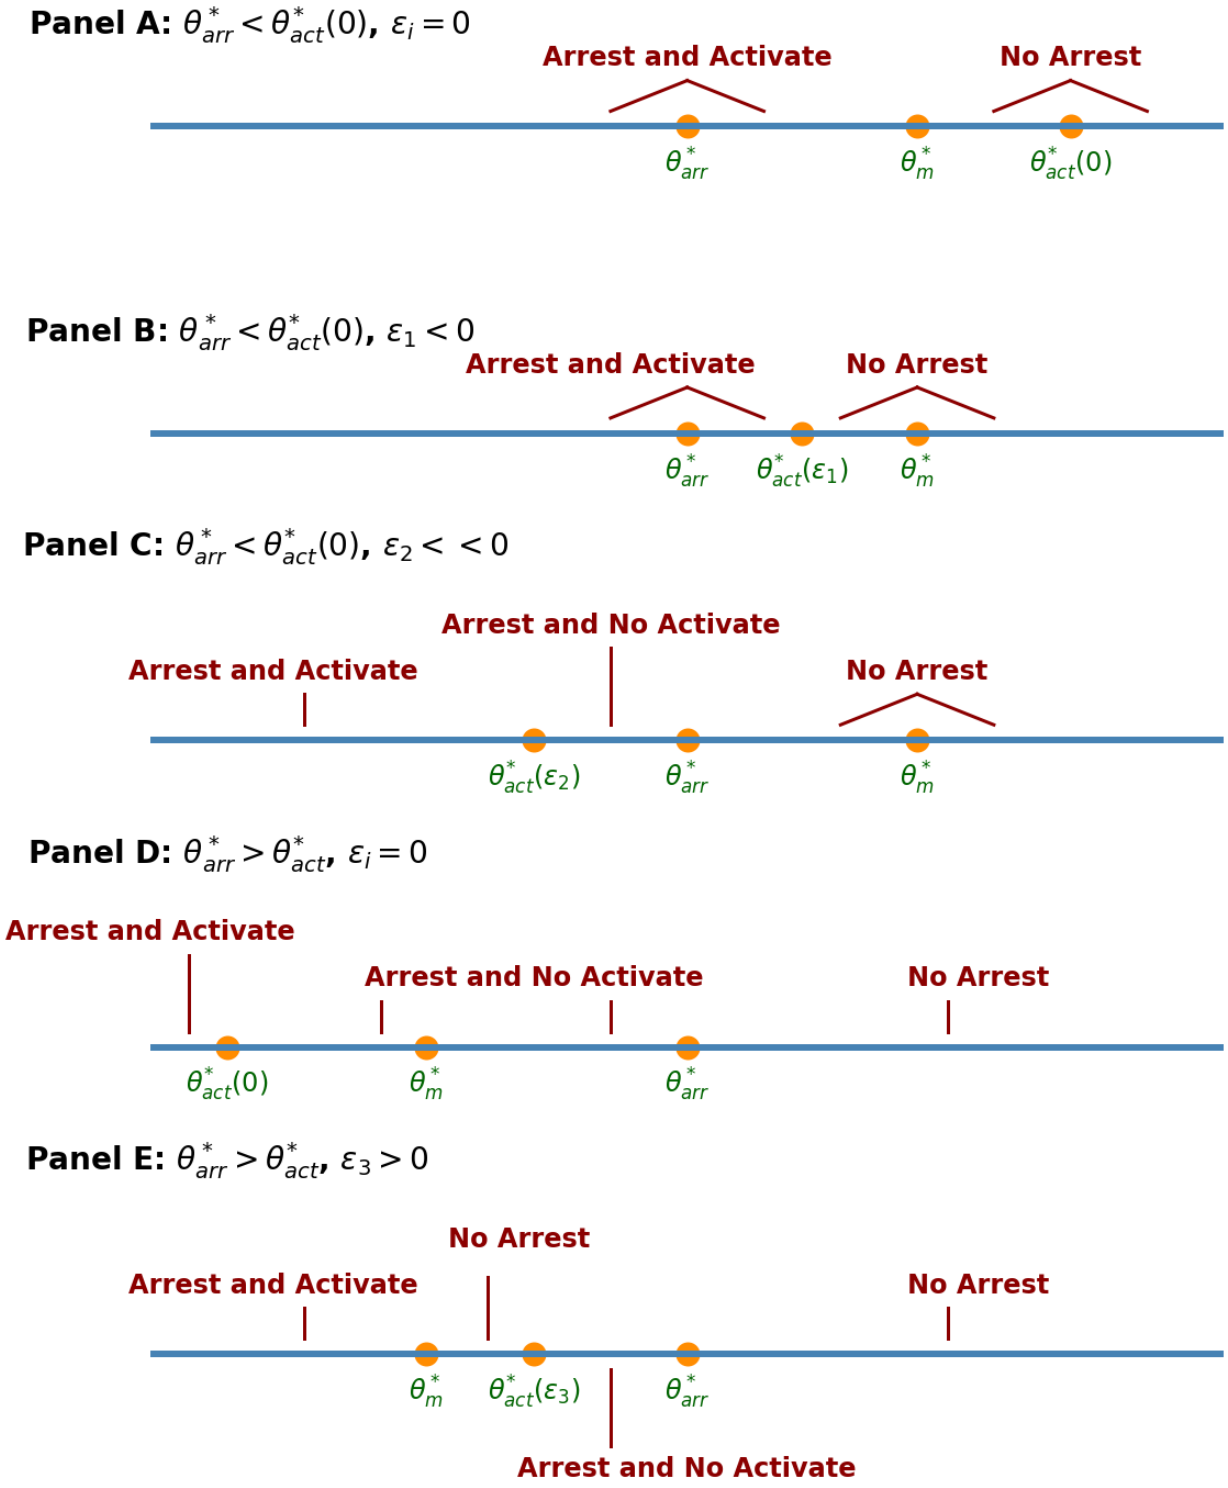


Table C1: Structural Estimation Results

Moment	Data	Simulated
1st Percentile	0.184	0.268
5th Percentile	0.395	0.384
10th Percentile	0.5	0.462
25th Percentile	0.642	0.614
50th Percentile	0.754	0.736
75th Percentile	0.83	0.819
90th Percentile	0.882	0.882
95th Percentile	0.906	0.915
99th Percentile	0.948	977
Parameter		**Estimate**
q_F		0.000014
q_T		0.181
$\theta_{act}(0)$		0.00011
p_h		0.001
p_l		0.217
σ_u		0.08
q'_F		-0.000014
q'_T		0.131
A		0.00023

Notes: This table shows the fitted output of our structural estimation. The upper panel shows the moments of the officer-level activation rate distribution in both our BWC data and the dataset simulated with our fitted parameters. The lower panel shows those fitted parameter values.