Worker Behavior Under Optional Oversight: Theory and Evidence from Police Body-Worn Cameras

Andrew Jordan Taeho Kim Soumendra Lahiri
Christopher Lucas Kyle Rozema

January 27, 2026

Police officers with body-worn cameras may activate to prove good conduct, but they may strategically fail to activate in risky situations that tend to lead them to violate department standards. We model this optional oversight setting and estimate it using Chicago data on activation during arrests. We find that roughly half of officers sometimes strategically fail to activate and respond to cameras by making fewer arrests as a way to avoid being recorded in risky situations. The mix of officers who sometimes and never strategically fail to activate leads to important policy tradeoffs that are not present in non-discretionary oversight settings.

*Jordan: Washington University, e-mail: awjordan@wustl.edu. Kim: University of Toronto, email: thk.kim@utoronto.ca. Lahiri: Washington University, e-mail: s.lahiri@wustl.edu. Lucas, Washington University, e-mail: christopher.lucas@wustl.edu. Rozema: Northwestern University Pritzker School of Law, email: kyle.rozema@law.northwestern.edu. For helpful comments, we thank David Abrams, Bocar Ba, Jennifer Doleac, Felipe Goncalves, Ben Feigenberg, Paul Heaton, Jens Ludwig, Aurelie Ouss, and workshop participants at the Conference on the Economics of Crime and Justice, Duke Structural Microeconomics Conference, Virtual Crime Economics seminar series, Washington University, and Western Economics Association International. We thank Hannah O'Neil and Sidney Nwakanma for excellent research assistance. This research is supported by a grant from the National Science Foundation (#2312383) and the Incubator for Transdisciplinary Futures at Washington University.

1 Introduction

When employers cannot directly observe worker behavior, they may invest in monitoring technologies to mitigate the principal-agent problem. Standard economic models of monitoring assume employers choose the level of oversight through their investment (Holmstrom and Milgrom, 1991). Although this assumption likely holds in settings like manufacturing or retail, it is less plausible where privacy, legal, or practical constraints prevent employers from fully controlling when and how monitoring occurs. In some settings, workers retain discretion over when their behavior is observed. As a result, employers in these "optional oversight" settings choose not only the level of investment in monitoring but also monitoring rules and how to enforce them. Given that these choices can affect worker compliance and productivity, these settings may create different tradeoffs for employers seeking to monitor workers than those suggested under traditional economic models of monitoring.

With the growth of remote work and new monitoring technologies, optional oversight settings are becoming increasingly common in private and public organizations. For example, remote workers decide whether to avoid employer-managed platforms that automatically track their behavior by completing tasks offline, and university faculty decide whether to record their classes. Although optional oversight settings abound, perhaps the canonical and most widespread example is policing. There, officers are equipped with body-worn cameras (BWCs) that must be activated at their discretion. A rich theoretical and empirical literature studies the impacts of BWCs and the role of monitoring and incentives in policing more generally (e.g., Ba, 2017; Chalfin and Gonçalves, 2025; Prendergast, 2022; Ba and Rivera, 2024; Bollman, 2024; Harrington and Shaffer, 2025). However, no study of policing considers the optional oversight nature of BWCs. This is in part because of a lack of systematic evidence on officer- and arrest-level activation behavior.

In this article, we study the impact of optional oversight in the context of police BWCs. We model officers with BWCs as workers who can choose to activate a monitoring technology at their own discretion. This discretion creates opportunities for strategic behavior, not only in the activation decision but also in decisions about engagement and arrests.

Officers consider the risk of misconduct accusations when deciding whether to make arrests, and activation decreases the rate of both false negatives and false positives in misconduct investigations. Depending on which type of error officers care more about, they can view BWCs very differently—as either protective or compromising. "Protective activators" use BWCs to document good conduct when claims of misconduct arise, while "compromising activators" sometimes strategically fail to activate if they anticipate a high risk of misconduct.

The model has contrasting predictions for how BWCs affect activation and arrest behavior for these two types of officers. With respect to activation, officers who view BWCs as protective activate whenever possible, whereas officers who view BWCs as compromising are less likely to activate in riskier arrests. With respect to arrests, officers who find BWCs protective make more arrests because the recording protects them from false accusations, whereas officers who find BWCs compromising refrain from making some arrests when activation is required. These contrasting officer-level predictions about activation and arrest behavior provide the empirical tests we use to assess the validity of the model, and the resulting estimates of these predicted effects provide the key moments that we use to identify and calibrate the model.

Our model captures activation decisions for individual arrests. As such, to connect the model to data, we need to observe activation decisions at the arrest level—data that has not been used in academic research across a department before. To this end, we use novel administrative data from the Chicago Police Department (CPD) that identifies the activation of specific officers as they make specific arrests—encounters where officers have an unambiguous duty to activate. To construct our activation data, we link arrests to metadata that is automatically generated by BWCs when they are activated. We do this using a unique event identifier that Chicago's emergency services dispatch generates. This identifier is recorded in most activation and arrest 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 link activation and arrest

records that do not have unique identifiers. Our final sample consists of 939,707 arrest-officer pairs between 2016 and 2020 in which the officer was equipped with a BWC.

Our empirical analysis consists of three parts. In the first part, we investigate the model predictions on activation decisions—if some officers view BWCs as protective but others view BWCs as compromising, there will be large differences in activation rates across officers that reflect an officer-level component. Compared with the overall activation rate of 72%, we find that officers differ considerably in how often they activate during arrests, with one in five activating in more than 87% of arrests and one in five activating in less than 63% of arrests. To assess how much of these differences reflect a stable officer-level component rather than differential exposure to situations, we estimate how well an officer's activation behavior in one arrest, time period, or assignment predicts their activation behavior in other settings. We also control for a wide range of situational factors and aspects of the oversight environment—characteristics of the arrest, shift, managers, other officers involved in the arrest, and pending complaint investigations. We find that officer activation rates are stable and strongly predictable in different settings and that situational factors explain no more than 11% of officer-level differences in activation rates. Most of this modest explanatory power is driven by whether the officer was called as backup, where they may have a less legitimate need to record the portion of the arrest they are present for.

In the second part of the analysis, we investigate the model predictions on arrest behavior—compared to how the officers who view BWCs as protective respond to the introduction of BWCs, officers who view BWCs as compromising should respond by making fewer arrests. To do so, we estimate how BWCs differentially affect the arrest behavior of compromising and protective activators by leveraging the staggered adoption of BWCs across Chicago's 22 geographic police districts between 2016 and 2017. In this analysis, we use activation rates as a proxy for whether an officer finds BWCs compromising. We find that BWCs decrease the number of arrests that compromising activators make compared to protective activators by 12%. We also find that the effects of BWCs on arrests are concentrated among the officers who likely find BWCs the most compromising, as predicted by the

model. We also explore the drivers for the relative decrease and find suggestive evidence that our results reflect changes in the decision of particular officers to engage in high-risk situations, as predicted by the model. In particular, we find that effects are primarily driven by the officers who initiate the arrest rather than those called as backup and by serious felony cases rather than lower-level offenses.

In the third part of the analysis, we consider how frequently officers strategically fail to activate and the resulting consequences for the design of BWC policies. To do so, we structurally estimate the theoretical model by matching the distribution of officer activation rates reflecting an officer-level component from the first analysis and the estimated effects of BWCs on arrests from the second analysis. We find that 55% of officers are compromising activators who strategically fail to activate in anywhere from a few arrests to over 40% of arrests. These differences in the rate of strategic non-activation between officers explain as much as 62% of the overall differences in activation rates between officers.

To assess the policy implications of optional oversight, we study counterfactual policies designed to increase activation and contrast their effects with settings in which officers lack discretion over monitoring. The central difference is that discretion expands the set of policy tools available to departments. When activation is non-discretionary, monitoring expands to all officers and situations equally, which can cause some officers to disengage in some situations. Under optional oversight, departments must choose how to address non-activation and can do so along multiple policy margins. We consider two such tools: an "enforcement" policy that is designed to detect and punish strategic non-activation through stricter enforcement of activation rules, and a "reminder" policy that aims to ensure that officers remember to activate even in the heat of the moment. These policies affect compromising activators differently than protective activators. On the one hand, stricter enforcement policies have no effect on protective activators and reduce arrests by compromising activators, but the marginal arrests they eliminate are those with the highest probability of misconduct. On the other hand, stronger reminder policies embolden protective activators to make riskier arrests while doing nothing to curb the arrest behavior of compromising activators. A per-

haps surprising implication is that optional oversight can lead to more arrests compared to when activation is non-discretionary. With optional oversight, police departments can mix enforcement and reminder policies to increase activation and decrease misconduct without decreasing arrests. Tradeoffs between activation, misconduct, and arrests are unavoidable under non-discretionary oversight, so police departments may *prefer* an optional oversight regime.

This article contributes to a growing literature studying police officer responses to oversight. Research examining officer activation decisions rely on limited data (Martain et al., 2020; Huff et al., 2024; Katz and Huff, 2023; Adams et al., 2021), and research studying the effects of police oversight on arrests, misconduct, and use of force has found mixed results (e.g., Prendergast, 2022; Ba, 2017; Ba and Rivera, 2024; Rozema and Schanzenbach, 2019, 2023; Jordan and Kim, 2024; Devi and Fryer Jr., 2020; Ferrazares, 2024; Lum et al., 2020; Kim, 2023; Barbosa et al., 2021). We advance this literature by building optional oversight into models of policing and using data on activation decisions when officers have an unambiguous duty to activate. Our results suggest that aggregate changes in arrest behavior following BWC rollout will depend on the mix of officer types in a department. This may help to explain the mixed evidence of the overall effect of BWCs on arrest behavior.

More generally, this article contributes to literatures on monitoring and the incentives of bureaucrats. The monitoring literature largely abstracts away from how the impacts of employer investment in oversight can be influenced by worker discretion over whether specific events are monitored (Holmstrom and Milgrom, 1991; Adhvaryu et al., 2022; Guiteras and Jack, 2018; Cai et al., 2025; Adhvaryu et al., 2022; Friebel et al., 2025; Hubbard, 2000; Kelley et al., 2024; Bernstein, 2012; Pierce et al., 2015). Our model and evidence suggest that in such optional oversight settings, the level of enforcement affects not only worker compliance but also productivity.

The literature on the incentives of bureaucrats starts with the premise that motivating government workers is more challenging than motivating private-sector workers due to difficulties in defining objective performance metrics and the limited contracts available

(Prendergast, 2003, 2007; Khan et al., 2015). This is especially true in policing and other enforcement agencies, where intrinsic motivations to make arrests rather than direct incentives like overtime pay are key to understanding police behavior (Prendergast, 2001; Chalfin and Gonçalves, 2025). Our model and evidence extend this literature into a system of optional oversight that allows intrinsic arrest motivations to interact with intrinsic activation motivations.

The remainder of the paper proceeds as follows. Section 2 presents a model of BWCs and policing behavior under optional oversight. Section 3 describes how CPD is organized and the adoption of BWCs. Section 4 describes the activation data and reports descriptive statistics. Section 5 documents officer-level activation patterns and investigates the drivers of differences in activation rates between officers. Section 6 investigates how the introduction of BWCs affected arrest behavior. Section 7 structurally estimates the model to investigate the extent to which officers strategically fail to activate and the resulting consequences for the design of BWC policies. Section 8 concludes.

2 A Model of Worker Behavior Under Optional Oversight

2.1 Setup

An officer, j, must decide whether to make a given arrest, i. The officer-specific reward for making the arrest is $w_j > 0$ relative to a normalized payout of 0 for not making the arrest. If the officer makes the arrest and an investigation finds that they engaged in misconduct, they pay a cost $c > 0$. The arrest-specific probability that the officer engages in misconduct is $\theta_i \sim F(\theta)$.[1]

Misconduct investigations produce both false negatives and false positives. If the officer engages in misconduct, this is discovered with probability q_T (the true positive rate). If the officer did not engage in misconduct, an investigation nonetheless finds that they did

[1]We assume that misconduct is always unintentional on the part of an officer. We treat misconduct risk as an irreducible consequence of the complex and fast-moving circumstances that officers work in. Although there may in fact be some situations where officers intentionally engage in misconduct due to animus or corruption, we do not expect BWCs to strongly affect this kind of misconduct in settings where oversight is optional.

with probability q_F (the false positive rate). We assume investigations are still informative: $\Delta = q_T - q_F > 0$, meaning actual misconduct is more likely to be detected than improper findings of misconduct are to occur.

Officers wear BWCs that must be manually activated to record arrests.[2] Activation (denoted by m) improves investigation accuracy by increasing detection of actual misconduct ($\frac{\partial q_T}{\partial m} = q_T^{'} > 0$) and decreasing the rate of improper findings of misconduct: ($\frac{\partial q_F}{\partial m} = q_F^{'} < 0$).[3] Define the change in probabilities as $\Delta^{'} = q_T^{'} - q_F^{'} > 0$.

Finally, there is a direct shock to utility from activation: $\epsilon_i \sim G(\epsilon)$. If $\epsilon_i < 0$, the officer faces some additional cost to activating, which we will interpret to include non-activation due to forgetfulness and equipment failures. If $\epsilon_i > 0$, the officer faces a cost of non-activation. These may include anticipated punishment for non-activation or a specific request from someone on the scene. To relate these shocks to policies discussed below, we will say that when ϵ_i is negative, the officer is in need of reminders, and when ϵ_i is positive, the officer faces BWC enforcement.

Officers in the model manage their overall risk of being found guilty of misconduct solely by choosing which arrests to engage in and whether to activate in them. To the extent that officers also exert costly effort to reduce misconduct risk in each arrest, this will reduce both θ_i and w_j. The values we work with in the model can be treated as the end products of this initial optimization step. That is, the officer has already chosen effort so as to maximize his payoff in the event he makes the arrest.[4]

---

[2]We model activation as a binary decision that is made at some fixed point early in the encounter. Officers cannot delay activation and cannot de-activate in response to an evolving arrest situation. This restriction stems from our data, which are such that we can only label whether any kind of activation occurred during an arrest and not whether it captured all important aspects of the arrest.

[3]For notational simplicity, we will mark changes due to activation as if continuous, even though officers can only choose between two levels: off and on.

[4]In Appendix B, we consider an extension that models this effort decision directly. The implications of the model are unchanged, but it is now necessary to keep track of the benefits (or costs) of re-optimizing effort upon activation.

2.2 Solution

We solve the model by looking first at the officer's decision to make an arrest when they do not have a BWC. In this case, the officer makes an arrest if the utility from doing so is greater than the utility from refraining:

$$w_j - c(\theta_i q_T + (1 - \theta_i) q_F) > 0 \tag{1}$$

Rearranging yields a threshold rule under which the officer arrests only when the misconduct risk is below:

$$\theta_i < \frac{\frac{w_j}{c} - q_F}{\Delta} \equiv \theta^*_{arr} \tag{2}$$

Officers are willing to make riskier arrests (higher θ^*_{arr}) if the punishment for misconduct is small relative to their benefit from the arrest, w_j, or if the probability of being caught conditional on misconduct is low.

Next, we consider the decision to activate. Officers activate if it increases expected utility from making an arrest:[5]

$$\frac{\partial}{\partial m}[w_j - c(\theta_i q_T + (1 - \theta_i) q_F) + 1\{\epsilon_i \leq 0\}\epsilon_i m - 1\{\epsilon_i > 0\}\epsilon_i (1 - m)] > 0 \tag{3}$$

Here, the first ϵ_i term captures the need for reminders. $\epsilon_i < 0$ may sometimes cause officers to not activate even though it would be beneficial for them at a neutral $\epsilon_i = 0$. The second ϵ_i term captures enforcement of activation rules. Officers must pay ϵ_i in the event that they make the arrest without activating.

Evaluating the derivative and rearranging yields another threshold rule in the prob-

[5]We focus on the effect on an officer's arrest utility because activation does not change non-arrest utility, and we are not interested in activation behavior in cases when officers do not make an arrest.

ability of actual misconduct:

$$\theta_i < \frac{\frac{\epsilon_i}{c} - q_F^{'}}{\Delta^{'}} \equiv \theta_{act}^*(\epsilon_i) \tag{4}$$

If the misconduct risk, θ_i, is low, the decrease in false positives offered by 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 activation is no longer optimal.

Finally, how activation changes arrest behavior depends on how θ_{arr}^* changes with m:

$$\frac{\partial \theta_{arr}^*}{\partial m} = \frac{-\Delta q_F^{'} - (\frac{w_j}{c} - q_F)\Delta^{'}}{\Delta^2} = \frac{-\ q_F^{'}}{\Delta} - \frac{\theta_{arr}^* \Delta^{'}}{\Delta} = \frac{\Delta^{'}}{\Delta}(\theta_{act}^*(0) - \theta_{arr}^*) \tag{5}$$

The effect on the arrest decision is therefore dictated by the relationship between two thresholds: θ_{arr}^* (the riskiest arrest the officer makes without BWCs), and $\theta_{act}^*(0)$ (the riskiest arrest where activation is attractive when $\epsilon_i = 0$). If $\theta_{act}^*(0) > \theta_{arr}^*$, then activation makes an officer willing to arrest in riskier situations than they otherwise would. Conversely, if $\theta_{act}^*(0) < \theta_{arr}^*$, then activation makes the officer unwilling to arrest in situations where they would have previously.

Another way to express this is to contrast the officer's arrest threshold when the BWC is not activated, θ_{arr}^*, with the officer's arrest threshold when the BWC is activated, which we denote θ_m^*. θ_m^* is the weighted average of θ_{arr}^* and $\theta_{act}^*(0)$, where the weights depend on the impact of BWCs on misconduct investigations: $\theta_m^* = \frac{\Delta^{'}}{\Delta}\theta_{act}^*(0) + (1 - \frac{\Delta^{'}}{\Delta})\theta_{arr}^*$. When $\theta_{act}^*(0) > \theta_{arr}^*$, the new threshold θ_m^* moves towards $\theta_{act}^*(0)$. In contrast, when $\theta_{act}^*(0) < \theta_{arr}^*$, θ_m^* moves away from $\theta_{act}^*(0)$.

2.3 Analysis of Officer Behavior

We separate officers into two different types based on θ_{arr}^* and $\theta_{act}^*(0)$. First, officers with $\theta_{arr}^* < \theta_{act}^*(0)$ are "protective activators" who activate whenever they can and where the ability to generate a record of their actions emboldens them to make additional arrests in risky situations. This attitude is illustrated by a post on Chicago's unofficial anonymous message board from around the time BWCs were introduced: "I welcome these cameras

whole heartedly. I am a pro active police officer. I know for a fact that if I do everything by the book then the Department HAS to back me up. I will continue to be proactive no matter what" (Bradley, 2021).

Second, officers with $\theta^*_{act}(0) < \theta^*_{arr}$ are "compromising activators" who strategically fail to activate in at least some arrests and where an obligation to generate a record of their actions will dissuade them from making some arrests. This attitude is illustrated by another post on the message board: "Handle jobs properly and do the bare minimum and you won't have any problems. Start being proactive with these cameras on and now you leave yourself open for lots of interpretations, second guessing and conflicts with offenders that will always look bad. Stay Fetal."

Appendix Figure A1 summarizes officer behavior under various conditions for θ^*_{arr}, $\theta^*_{act}(0)$, and ϵ_i. Panel A illustrates protective activators when $\epsilon_i = 0$. These officers activate in every arrest they make. Compared to having no BWCs, they make more arrests because activation expands the set of situations where they are comfortable arresting. Panels B and C illustrate the effect of a need for reminders ($\epsilon_i < 0$). Protective activators may not activate in all arrests, but, at a minimum, they still make all the arrests they would make without BWCs. For protective activators, enforcement ($\epsilon_i > 0$) has no impact on their decision because they activate whenever possible.

Panel D turns to compromising activators when $\epsilon_i = 0$. These officers make the same arrests they would without BWCs but strategically fail to activate in their riskiest arrests to avoid creating records of those encounters. Panel E shows that enforcing activation rules with compromising activators ($\epsilon_i > 0$) will cause them to withdraw from some arrests entirely. A need for reminders ($\epsilon_i < 0$) has no effect on compromising activators, since they already avoid activation in risky situations.

Overall, the model has two primary implications. First, protective activators have higher activation rates and receive fewer complaints than compromising activators.[6] Second, the arrest behavior of compromising activators responds to BWCs differently than that of

[6]We show this formally in Appendix B using variation in officer arrest incentives, $\frac{w_j}{c}$.

protective activators. Compromising activators reduce arrests, and protective activators increase them.

3 Background

In November 2015, the release of police footage of the shooting of Laquan McDonald triggered widespread protests throughout Chicago. In response to significant public and governmental pressure, CPD rapidly expanded its ongoing small-scale BWC pilot program and began a full-scale rollout of BWCs in June 2016. Chicago is broken up geographically into 22 police districts, and BWCs were rolled out at the district level between June 2016 and December 2017. Once BWCs were adopted in a district, all officers in the district were provided with one. Officers assigned to city-wide roles (e.g., detectives, vice, and SWAT) were never provided BWCs in our analysis period. Officers with BWCs received both hands-on training and instruction on how to use them (CPD, 2016).

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, and officers are required to verify that their BWC is working and charged at the start of each shift. Officers work in three primary shifts, referred to as watches, with the morning shift running from 6am to 2pm, the afternoon shift running from 2pm to 10pm, and the night shift running from 10pm to 6am.

CPD policy also mandates that officers activate at the start of any law-enforcement-related activity and record the entire duration of the incident. A law-enforcement related activity includes all arrests regardless of the officer's role in the arrest, responses to calls for service, investigatory stops, and any civilian encounter that becomes adversarial. If a victim or witness requests that an officer deactivate the BWC, officers may do so after providing a verbal justification on the recording before deactivating. 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. If a BWC is not functioning properly, officers must contact a dedicated help desk that troubleshoots how to get a functioning BWC (Special Order 03-14).

According to CPD rules, BWC oversight falls to the watch operations lieutenants

("lieutenants") on duty in each shift. Their responsibilities include reviewing footage for every arrest to confirm probable cause and reviewing at least one randomly selected BWC recording per shift. Audits of this oversight system have found that some lieutenants perform these oversight duties rarely, if ever.[7]

If an officer does not follow activation rules, they can be punished through what is called a Summary Punishment Action Request. These requests are less serious and do not trigger the lengthier misconduct review process as in complaints for other types of misconduct. During our sample period (2016–2020), violation of BWC activation rules was not recorded in CPD complaints data, but the annual number of all equipment-related requests ranged from 55 to 99, suggesting limited oversight response to BWC violations. This interpretation is consistent with contemporary investigations of CPD's BWC practices. For example, the Chicago Civilian Office of Police Accountability's 2021 report on non-compliance with BWC regulations called for a re-evaluation of disciplinary protocols for BWC compliance failures.

4 Data and Descriptive Statistics

4.1 Data

We combine several datasets acquired from CPD through Freedom of Information Act requests. Across all datasets, we identify officers using a unique identifier based on their full name, year of birth, and date of appointment to CPD. Each dataset covers at least the time period before and after BWCs were adopted in all districts (2015-2020).

Arrest Data. We obtained data on adult arrests made by CPD officers containing the identity of the officer(s) making the arrest, the role of the officer in the arrest (e.g., assisting officer), the time and location of the arrest, the charges, and demographic information of the person arrested.[8] Many records contain a unique "event identifier" on the specific

[7]In an audit of random reviews in seven districts, compliance was 85% in one district, 39% in another, 19% in another, and 0% in the remaining four districts (City of Chicago Office of Inspector General, 2019).

[8]We keep arrests made by assisting officers in our primary analysis for several reasons. First, CPD regulations dictate that officers should activate their BWCs when participating in an arrest regardless of their role. Second, assisting officers must have been involved in the arrest itself; transport and booking are

event the officer is responding to that is generated by Chicago's emergency services dispatch after receiving a civilian call for service or a report from the officer in the field. Because all officers involved in an arrest have an obligation to activate their BWC, we organize the data at the officer-arrest level. We drop arrests made by officers working in special units without BWCs.

BWC Activation Data. We obtained data on every activation from July 2016 to December 2021. Each record contains an identifier for the officer and the recording start and end time. Many records contain the same unique event identifier as in the arrest data.

Assignment Data. We obtained data on shifts worked for officers and lieutenants. Each record contains the date, the district, and the watch. We use this data to count the number of shifts an officer works in each calendar month. We also use this data to link each officer-shift to the lieutenant overseeing the shift. We exclude arrests made while the officer was not on an identifiable shift.

Complaints Data. We obtained data on each complaint made to the civilian police oversight agency since 2000. Each record contains an identifier for the officer, the date the complaint was made, the date the investigation ended, and whether the complaint was sustained. We use this data to construct officer-level complaint propensities before BWCs were adopted and to identify arrests that officers made during and after misconduct investigations.

BWC Help Desk Data. We obtained data on BWC help desk requests that we use to identify BWC equipment failures. Each record contains the identity of the officer who made the request and the date the request was made.

Labeling BWC Activation in Arrests. Our first goal is to label each officer-arrest as either an activation or a non-activation. For officer-arrests where event identifiers are available in both the arrest data and BWC data, we can directly link arrests to activation and thus label the arrest as having an activation. Unfortunately, although officers are supposed to record event identifiers for each arrest and BWC recording, they often fail to do so.

---

separate roles. Third, officer-level activation rates are highly correlated across their primary and assisting arrests, and the activation patterns we present below are not significantly impacted by dropping assisting officers from the sample.

For an arrest or a BWC that has a missing event identifier, we link records using timestamps for the arrest and activation. In the set of arrests that *can* be matched by event identifiers, arrest timestamps frequently fall significantly outside activation and de-activation timestamps. These gaps can arise when the arrest timestamp is given not as the actual moment of arrest but instead as the time a call was received, an officer was dispatched, or an arrest was booked instead of the actual moment of arrest.

To label whether any activation can be associated with an arrest, we use the arrests linked by event identifiers to estimate a Gaussian mixture model of the relationship between arrest and activation timestamps. We then use the estimated model and a threshold for the proximity between arrest and activation timestamps to label unlinked arrests. We accept most such matches up to a 20-minute gap between the arrest timestamp and the activation window.[9] Appendix C sets out the details of our procedure, estimation, and validation exercises that reveal that the model performs well.

Our matching procedure inevitably produces some false negatives, cases where the BWC was truly activated but we label it as having not been activated, and false positives, cases where the BWC was not activated but we label it as having been activated. It is possible that officers are differentially exposed to the risk of these errors if they have different habits about writing down the event identifiers we use to assemble our set of more definite matches. However, our main interest is in each officer's activation rate, and this will be unaffected by false negatives and positives insofar as failing to write down event identifiers or accurate arrest timestamps generates these problems in roughly equal proportion. Within the sample of known matches using event identifiers, our model fails to label 6.7% of them. This provides an estimate for the false negative rate. We cannot easily estimate the false positive rate because we do not have a sample of arrests that we know are not activations. To understand the extent that false positives are likely to bias our estimates of officer activation, we calculated the extent of the potential bias for false positive rates from 0 to 20%. For an officer with a typical activation rate, 72.4%, who is at the 95th percentile of event identifier

[9]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.

failure, 40%, the error in measured activation rate ranges from -1.9 percentage points when the false positive rate is 0 to 0.3 percentage points when the false positive rate is 20%.

One potential concern is that not all activations are guaranteed to capture all of the relevant parts of an arrest. Officers may sometimes start the recording too late or end it too early. Our arrest timestamps are measured with error, so we cannot distinguish these deficient recordings from recordings that capture the entire arrest. We will therefore mark an activation in some arrests where officers did not generate a useful recording. Our activation indicator should therefore be thought of as an upper bound on BWC usage.

4.2 Descriptive Statistics

For arrests made by officers with BWCs on patrol in the 22 districts, the activation rate at the officer-arrest level is 72.4%. Given that most arrests are made by multiple officers, a much higher share of arrests had at least one BWC activated (92.8%). And given that officers in units without BWCs also make arrests, the activation rate across all arrests in Chicago is much lower (81.9% at the arrest level following full BWC adoption in 2018).

Figure 1 reports activation rates across police districts, and Figure 2 reports activation rates across police districts over time. Activation rates differ between districts by up to 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 South and West sides of the city, but there are notable exceptions. For example, District 9, on the South side, has among the highest district-level activation rates. Activation rates are lower and more variable during the period when BWCs were rolled out but stabilized by 2018. Moving forward, we focus on data from 2018 and later after activation rates have stabilized.

Finally, Figure 3 investigates the model predictions that compromising officers will lower activation rates, receive more complaints, and make more arrests on the other. Panel A reports the relationship between activation rates and complaints and reveals a strong negative relationship, consistent with the model predictions. Panels B and C report the relationships between arrests and measures associated with officer type, where Panel B uses our preferred measure of officer activation rates and Panel C uses our proxy for officer type

that combines activation rates and complaints rate.[10] We find strong negative relationships but a stronger one if both proxies are used, suggesting that using both measures improves the proxy for officer type.

5 Differences in Officer Activation Rates

Our model divides officers into protective and compromising activators and predicts that protective activators activate BWCs at higher rates than compromising activators. Later sections will use officer-level variation in activation rates, first to test implications of the model using the rollout of BWCs and then to structurally estimate the model. In this section, we investigate the extent to which differences between officer activation rates reflect a fixed officer-level component in two ways. First, we estimate the persistence of activation rates within officers across arrests, time, and situations and how much of it can be explained by noise and officers sorting to certain environments.

Second, we decompose the effect of observables on the variation in officer activation rates. Unlike many domains with a well-established baseline empirical model (e.g., the Mincer equation for wages), little is known about the basic drivers of BWC activation. We therefore devote more space than usual to this type of analysis, describing at some length the motivation for including different situational factors and reporting results on how well these factors explain both activation and variation in officer-level activation rates.

The question of whether activation has a strong officer-specific component bears on more than just the validity of our model of officer-level heterogeneity. Even in the absence of a particular theory of why some officers activate less than others, there is policy insight in understanding how much circumstances explain activation rates. If officers with particularly low activation rates were just disproportionately in situations that lead to less activation, then the apparent lack of discipline or remediation of officers with low activation rates would

[10]For this analysis, we control for officer demographics, which is important because older officers have much lower activation rates. Without the controls for officer demographics, the officers with the lowest 10 percent of activation rates, who are on average much older than other officers, have lower arrests per shift, and the negative relationship between activation and arrests is only observed for officers above the bottom decile.

represent proper enforcement of a reasonable policy. However, if meaningful differences remain even after accounting for situational factors, this calls into question the policy or its enforcement.

5.1 Officer Effects

Figure 4 reports the distribution of officer activation rates (and an adjusted distribution discussed below). In this distribution and subsequent analyses in this section, we drop officers who made 10 or fewer arrests in the period studied. We find large differences in activation rates between officers. For high-activating officers, while the activation rate is not perfect for any officers, it is more than 90% for 12% of officers and more than 75% for 58% of officers. In contrast, for low-activating officers, the activation rate is less than two-thirds for 20% of officers and less than half for 10% of officers. The result is that the magnitude of the differences between officers is large. For example, the difference in activation rates between the 25th and 75th percentile officers is 19 percentage points.

To investigate the extent to which these differences in officer activation rates reflect a fixed officer-level component, we first document the within-officer stability of activation rates across policing environments. To do so, we regress each officer's activation in a particular environment on their leave-out activation rate elsewhere. We do this in terms of individual arrests, time (a quarter), geographic environment (district and beat), different times of the day (morning, day, and night watch), arresting role (primary, secondary, assisting officer), and under different lieutenants. Note that each resulting regression will be aggregated at a different level and thus have a different sample size.[11] For example, for the analysis across quarters, we regress officer activation rates in a given quarter on the officer's leave-out activation rate in other quarters; and for the analysis across police districts, we regress an officer's activation rate in one district on their activation rate in other districts.

Second, we repeat the above two exercises but replace each officer's leave-out activation rate with an Empirical Bayes leave-out activation rate that shrinks noisy estimates

[11] In the lieutenant regression, we also drop 14% of officer-arrests where we do not observe lieutenant identifiers for the officer-shift.

towards the mean activation rate (Weisburst, 2024). For example, for the officer-level Empirical Bayes, we estimate $\hat{Y}_i^{\text{EB}} = \frac{y_i+a}{n_i+a+b}$, where y_i is the number of arrests by officer i with activation, n_i is the total number of arrests by officer i, and a, b are the tuning parameters that define the priors on activation rates.[12] Measurement error from noise in the leave-out activation rates could attenuate the estimates towards zero, but the shrunken activation rate should correct for some of the noise and reduce attenuation bias. This shrinkage approach has been used in the policing setting (e.g., Rozema and Schanzenbach, 2019) and is common in teacher value-added literature (Chetty et al., 2014).

Third, we assess the extent to which the persistence of activation rates within officers across arrests, time, and situations can be explained by officers sorting to certain environments. One concern with assessing the stability of officer activation rates across different environments is that officers can select into environments, and it could be the environments rather than a fixed officer-level component that drives the stability of activation rates. If an officer's activation rate is fixed, then an officer switching into a new beat where they have a higher activation rate than officers already in the beat should increase the activation rate in the new beat. Importantly, this is true even if the type of beat they switch to is endogenous to their activation rate. The same intuition holds in reverse for an officer switching into a new beat where they have lower activation rates than officers already in the beat, and the same intuition holds for officers switching out of beats as well. Switchers tests in this spirit were introduced in the teacher value-added literature (e.g., Chetty et al., 2014) but have been also applied in policing (Rozema and Schanzenbach, 2019; Rivera, 2025). Appendix D describes the setup of the switchers test in detail.

Table 1 reports the results on the within-officer stability of activation rates across policing environments, and Appendix Table D1 reports the results of the switchers test. We find strong evidence of a fixed officer component in each of the tests. Officers with a

[12]Specifically, $a = \overline{p_i}\left(\frac{\overline{p_i}(1-\overline{p_i})}{\text{var}_i(p_i)-\overline{p_i}(1-\overline{p_i})\,\overline{1/n_i}} - 1\right)$, is a parameter that added to each officers activation rate that pulls officers with fewer observations toward the overall activation rate at higher levels, and $b = (1-\overline{p_i})\left(\frac{\overline{p_i}(1-\overline{p_i})}{\text{var}_i(p_i)-\overline{p_i}(1-\overline{p_i})\,\overline{1/n_i}} - 1, \right)$ is a parameter that allows the shrinkage to incorporate the overall activation rate and spread of activation rates across officers.

1 percentage point higher activation rate are 0.71 percentage points more likely to activate in the current arrest. We likewise find a strong persistence of activation rates across each of the environments we consider. The shrunken activation rates are even more predictive, suggesting that noise attenuates the estimates towards zero. In the switchers test, we find a 1:0.94 relationship between the actual impact of officers entering and departing a beat and the impact that would be expected if officers had a fixed activation rate. The confidence interval of this estimate includes a 1:1 relationship. Combined, these results provide strong evidence that there is persistence of activation rates within officers across arrests, time, and situations that is not explained by officers sorting to certain environments.

5.2 The Role of Situational Factors

We next consider how three types of observable situational factors influence activation and how much they explain the variance of officer-level activation rates. First, we consider the conditions of the arrest—such as the time, location, and offense—that may affect activation. Second, we consider operational constraints that may prevent the officer from recording the arrest, such as arriving later in an assisting role or equipment malfunctions. Finally, we consider other aspects of the institutional and oversight environment—such as lieutenants, peers, and complaint investigations—that may affect how officers weigh the oversight implications of activation.

In the remainder of this section, we first describe the situational factors in more detail and then present our estimation strategy and results.

5.2.1 Measuring Situational Factors

(A) Conditions Surrounding the Arrest. Work Environment. Work environment can contribute to differences in activation rates due to differences in place and time-specific conditions, prevailing norms, and different types of interactions. We measure this with district by calendar month and hour of the week fixed effects.

Work Schedule. Work schedules can contribute to differences in activation rates between officers due to fatigue, disruptions to routine, and the challenges of adapting to

unfamiliar assignments. We assess the contribution of work schedules by including indicators for whether the officer had a high recent workload (defined as working more than 5 of the last 7 days), for a double shift in a day, for transition shifts immediately before and after time off (defined as 5 days of not working), for atypical shifts where the officer works a different watch than the officer's recent pattern (defined as shifts where 0 or 1 of the officer's past 10 shifts were on the same watch), and beat transition for the first week and first two weeks following assignment to a new beat.

Encounter. Encounter characteristics can contribute to differences in activation rates between officers because certain encounters may carry greater perceived risk, complexity, or scrutiny. We assess the contribution of encounter characteristics by including indicators for violent crimes, police-related crimes, felony crimes, and suspect race, sex, and age ranges. Because we find that there are not large differences in activation rates between crime types and severity (Appendix Figure A2), our primary approach uses these three measures of crime type, but we find similar results if we include crime category fixed effects.

(B) Operational Constraints. Officer Role. Officers involved in an arrest are either the first arresting officer, the second arresting officer, or an assisting officer. The first two may be likely to activate before making the physical arrest, whereas assisting officers may not help make the physical arrest but instead may help with securing the area and interviewing witnesses. We assess the contribution of arresting roles by including indicators for second arresting officer and assisting officer. We also assess the contribution of the number of officers involved with the arrest by including indicators for a sole officer, 5 or more officers, and 10 or more officers.

Equipment Issues. Officers with equipment issues may not be able to activate their BWCs. We assess the contribution of potential equipment issues by including indicators for the officer filing a BWC help desk request the day the arrest was made and for the officer having made fewer than 20 and fewer than 50 arrests with a BWC. The former captures the inability of the officer to activate due to BWC equipment failure, and the latter captures potential inactivation due to inexperience with BWCs.

(C) Institutional Pressures. Lieutenants. The lieutenant presiding over a shift may be expected to influence officer activation rates because they may set different expectations for activation and vary in how closely they supervise and discipline officers for failing to activate. We assess this contribution by including lieutenant fixed effects.

Peer Effects. Officer activation decisions could be influenced by other officers involved in the arrest through peer effects. We assess the contribution of peer effects by including indicators for low and high-activating peers, measured as being in the bottom and top quartiles of the peer activation rates within a given district and calendar month.[13]

Enhanced Oversight. Officers who have experienced recent misconduct investigations may feel a sense of enhanced oversight. We assess the contribution of enhanced oversight by including indicators for the officer having any ongoing misconduct investigations, having multiple ongoing misconduct investigations, and having a complaint sustained in the past 3 and 6 months.

5.2.2 Research Design

To investigate how much of the differences in officer activation rates can be explained by situational factors, we begin by estimating a series of regressions of the form:

$$Y_{ij} = \beta X_{ij} + \theta_j + \epsilon_{ij} \tag{6}$$

where Y_{ij} is activation in arrest i by officer j, X_{ij} is some situational factor, θ_j are officer fixed effects, and ϵ_{ij} is a noise term.

Defining $\hat{\theta}_j^X$ and $\hat{\theta}_j$ as the estimated fixed effects when X_{ij} is and is not included in the regression, the difference in their variances is given by:

$$\text{Var}(\hat{\theta}_j) - \text{Var}(\hat{\theta}_j^X) = \beta^2 \text{Var}(\bar{X}_j) + 2\beta \text{Cov}(\bar{X}_j, \theta_j) \tag{7}$$

Equation 7 shows that the amount of variance explained by any given X_{ij} depends on how

[13] If the officer is the only officer involved with the arrest, we set both indicators to 0.

much X_{ij} directly affects activation (β), how much it varies across officers $\big(\text{Var}(\bar{X}_j)\big)$, and how much it covaries with officer fixed effects $\big(\text{Cov}(\bar{X}_j, \theta_j)\big)$. Situational factors that explain a meaningful amount of variation in activation rates across officers must have both a large effect and a large variance or covariance. Below, we will plot β, $\text{Var}(\bar{X}_j)$, and $\text{Cov}(\bar{X}_j, \theta_j)$ for several X_{ij} in order to better understand how much officer-level activation variance they explain.[14]

To summarize the explanatory power of a vector of situational factors $\mathbf{X}_{ij}$, we estimate the fraction of the variance they explain collectively:

$$\gamma^{\mathbf{X}} = 1 - \frac{\text{Var}(\hat{\theta}_j^{\mathbf{X}})}{\text{Var}(\theta_j)} \tag{8}$$

5.2.3 Results

We report the results of the variance decomposition first for individual situational factors and then for their combination.

Figure 5 reports β^X, the variance of X_{ij}, and the imputed covariance effect for several covariates grouped according to the categories above. Each β^X comes from a regression that only includes X_{ij}. For this figure, covariates with more than two categories (lieutenant, district by calendar month, and crime type) are reported as an indicator for categories in the bottom quartile of activation rates. Note that bars in Figure 5 filled with black dots reflect factors that *decrease* activation.

Figure 6 reports the total impact of each category of situational factors by plotting $\gamma^{\mathbf{X}}$, the percentage of the overall officer variance explained by the vector $\mathbf{X}$. For this figure, covariates with more than two categories (lieutenant, district by calendar month, and crime type) are included as fixed effects. The group of situational factors that explain the most variance of officer activation rates is officer role at 8.3%. Some situational factors, such as schedule, actually increase the measured officer-level variance due to the covariance term in

[14] We cannot directly calculate $\text{Cov}(\bar{X}_j, \theta_j)$ because the true officer fixed effect, θ_j, is unobserved, but we can calculate it as a residual using $\frac{\text{Var}(\hat{\theta}_j) - \text{Var}(\hat{\theta}_j^X) - \text{Var}(\hat{\theta}_j - \hat{\theta}_j^X)}{2\beta}$.

Equation 7.

The light gray bar on the far right of Figure 6 sums the other bars while the dark gray estimates $\gamma^{\mathbf{X}}$ for our full set of covariates. The former is 14.9% while the latter is 10.8%. Some of the characteristics that explain officer-level activation rates are themselves correlated, so the model that includes all factors explains less than the sum of individual factors.

From Equation 7, a covariate explains a substantial fraction of the cross-officer variation in activation rates only if both β^X and either the variance or covariance term is high. We find that this is rare. Some covariates have large effects on activation but do not vary substantially across officers and covariance effects are typically quite small. Examples of this include working a double shift and having made a help desk request that day. These factors affect the force-wide activation rate, but they cannot explain variation across officers.

Other covariates meaningfully vary across officers but have only small effects on activation. Examples include experience with BWCs and oversight measures. For example, some officers are much more likely than others to be under a complaint investigation, but being under investigation is only associated with a 1.5 percentage point higher activation. Of course, some covariates could be endogenous (e.g., selection into peer groups), but we suspect that a design that adequately controlled for selection would still be unlikely to find activation effects of sufficient size to explain much of the overall variance in officer-level activation rates.

Finally, a small number of covariates have meaningful activation effects and variation across officers. This includes whether the officer is in an assisting role (16 percentage points less likely to activate, and a variance of the share of arrests in assisting roles of 0.028) and whether the officer is in a low-activation work environment (16 percentage points less likely to activate, and a variance of the share of arrests that are for violent crimes of 0.020).

In Figure 6, we find that very little of the officer-level variation in activation rates can be explained by observables. Although the situational factors do explain some differences in activation between officers, they do not change officer activation rates very much, with the

unadjusted rates explaining 95.1 percent of the variation in the adjusted rates (Appendix Figure A4). This leaves a large potential role for the strategic considerations described in our model. We turn our attention to this in the remaining sections.

For this analysis, we considered an expansive set of covariates in order to paint a clear picture of what did and did not explain variation in activation rate across officers. Of course, we observe only an incomplete set of imperfect proxies for situational circumstances that could explain differences in activation rates between officers. Importantly, however, whereas some factors explain differences in activation rates between officers (e.g., arresting roles), other factors actually amplify differences for some officers and reduce differences for others (e.g., unusual shifts). This implies that the true differences in activation rates between officers that are attributable to officer disposition or strategic non-activation if all situational circumstances could be fully controlled could actually be greater than the officer activation rates adjusted for the situational factors we can observe.

Moving forward, however, we must be sensitive to the possibility that certain controls will be endogenous to the activation and arrest decisions that we are interested in. The situational factors that we believe are exogenous to arrest decisions include those related to work environment, encounter characteristics, and officer role. We therefore use adjusted activation rates that only control for these factors as an input into our analyses in Sections 6 and 7. These adjusted activation rates are reported in Figure 4. With that said, adding the additional controls has little impact on the adjusted activation rates (Appendix Figure A3), with the adjusted activation rates without the endogenous controls explaining 93 percent of the variation in the adjusted activation rates with all the controls.

6 Body-Worn Cameras and Arrest Behavior

Our model predicts that compromising activators respond to the introduction of BWCs by making fewer arrests but protective activators respond by making more. This section investigates this prediction by leveraging the staggered rollout of BWCs to different police districts.

6.1 Research Design

We are interested in how BWCs differentially affect arrest behavior for compromising activators compared to protective activators. Although whether a specific officer is a compromising or protective activator is an unobserved latent variable, our model implies that officer type is correlated with observable activation rates. We therefore use activation rates as our primary proxy to identify compromising officers. The choice of where to set the threshold is admittedly arbitrary, and we will show robustness to a wide range of thresholds below. In our primary specification, we separate officers at the 25th percentile of adjusted activation rates. We compute this percentile based on 2018–2020 activation rates as an out-of-sample proxy for officer type (that is, it is not constructed from the same period as the arrests we study during BWC rollout).

BWCs were rolled out in Chicago at the district-level in roughly three adoption periods: 6 districts from June 2016 to February 2017, 5 districts from March 2017 to July 2017, and 11 districts from October 2017 to December 2017.[15] We use this staggered adoption of BWCs across police districts in a difference-in-differences design. Given the negative weighting issues associated with two-way fixed effects (TWFE) approaches with staggered policy adoption, our primary specification uses the approach in Callaway and Sant'Anna (2021). This estimates treatment effects that are both district and time-specific using only comparisons between treated and not-yet-treated districts. We use data from 12 months before the first district adopted BWCs up to the last adoption in our data.

Callaway and Sant'Anna (2021) only allows for comparisons of the treated and control districts, but we are interested in how BWCs differentially affect arrest behavior by officer type within the same district in a triple-differences setting. To estimate the differential impact of BWCs by officer type, we follow guidance in Roth et al. (2023) and define the outcome as the difference in the average arrests per shift per month between officer types:

$$y_{dt} = \frac{\sum_i A^c_{jdt}}{\sum_i S^c_{jdt}} - \frac{\sum_i A^p_{jdt}}{\sum_i S^p_{jdt}} \tag{9}$$

[15] Appendix Figure A5 reports the rollout dates for each district and the full panel of data we use.

for officer j, district d, and calendar month t, where A^c_{jdt} and A^p_{jdt} are arrests and S^c_{jdt} and S^p_{jdt} are shifts worked by compromising and protective activators, respectively. We restrict to full-time officers assigned to a district.[16] The identifying assumption is that the differential arrest measure in a district that adopts BWCs would evolve similarly to the districts that have not yet adopted BWCs in the absence of BWCs. We assess this assumption by reporting results from the event study specification from Callaway and Sant'Anna (2021).

6.2 Primary Results

Figure 7 reports the event study results and reveals that the difference of the differential arrest measure between the treated and control districts in the pre-period is stable and not statistically different from zero, consistent with parallel pre-trends.

Table 2 reports the regression results. In Column 1, we find that BWCs decrease arrests per shift by compromising activators relative to protective activators by 12% (0.026 relative to a baseline of 0.224 arrests per shift). Columns 2-3 report results by officer type and reveal results consistent with the model—a positive point estimate for protective activators and a negative one for compromising, but only the former is statistically significant.

Column 4 reports results on the overall effect of BWCs on arrests. We find that the changes in the two groups offset, combining to produce no significant overall effect. Because the fraction of compromising activators will determine the overall effect of BWCs on arrests, these results illustrate how the officer heterogeneity highlighted by our model may explain the mixed evidence of the overall effect of BWCs on arrest behavior documented in the literature (e.g., Lum et al., 2020).

6.3 Heterogeneity

Persistent Officers. Our main results allow officers to move between districts over time. If officers respond to the presence of BWCs by seeking to move into or out of districts,

[16]For the full-time restriction, we exclude officer-months where the officer works fewer than 5 shifts. This drops 3.3% of officer-months in our data. 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 their modal district for a month. Modal districts account for 99.4% of all assignments within a given officer-calendar month.

and if these preferences are correlated with baseline arrest rates, this could drive the results but would be a different mechanism than that suggested by our model.

We assess the extent to which this mechanism drives the main results by focusing on officers who have a consistent assignment throughout the rollout period. Specifically, in one specification we focus on the 41% of officers who are in the same district in 12 of the 13 months around BWC adoption in their district (the month of rollout and 6 months on either side), and in another we focus on the 38% of officers who are in the same district in 20 of the 25 months considered in our primary specification. Columns 1-2 of Table 3 report the results and reveal consistent results.

Initiating Arrests. Next, our model is about the decision to initiate an arrest. We therefore want to verify that the overall effect we observe is not driven by assisting officers who may have been called to the scene only after the decision to make the arrest was made by a primary arresting officer. We therefore consider alternate specifications that focus on either only primary or only assisting arrests. Columns 3-4 of Table 3 report the results and reveal that relative activation rates in both groups decline by approximately 11.5% (0.011 primary arrests per shift from a baseline of 0.096 and 0.014 assisting arrests per shift from a baseline of 0.119). This shows that primary officers are indeed changing their arrest behavior in response to BWCs in the way our model predicts.

Hassle Costs. We next consider the possibility that BWCs change officer behavior primarily because they introduce added "hassle costs" from a perceived need to be more formal in the encounter or increased paperwork associated with logging the video itself. BWCs may cause officers who are particularly sensitive to these hassle costs to reduce their arrest effort. For some marginal arrests, BWCs may leave officers with two less-than-ideal options: activate and pay the hassle cost, or do not activate and risk discipline. Officers in this position may activate less often and also opt not to make some arrests that they would have made in the absence of BWCs.

We test this possibility in two ways. First, note that if activation behavior is primarily driven by hassle costs rather than misconduct considerations, then we should expect that

marginal arrests will be those with the lowest perceived return, such as misdemeanors and public order arrests. In these low-return arrests, the fixed hassle costs will be more likely to deter officers from making the arrest. Columns 5-7 of Table 3 thus report results separately for felony, misdemeanor, and public order arrests. We find that, relative to baseline arrests per shift of that kind, the estimated effects are largest for felony arrests (29%), followed by misdemeanor arrests (7%) then public order arrests (not statistically significant, but with a point estimate of 5%). These results are inconsistent with a hassle cost story and consistent with our model, where marginal arrests are those with the highest risk of misconduct accusations.

Second, we divide officers by a characteristic that our model also relates to the latent officer type but that is unrelated to hassle costs: complaints. Complaints are not related to BWC hassle costs because we measure them prior to BWC rollout and because we associate compromising activators with *high* complaint rates, an unlikely condition for officers who primarily wish to avoid hassle. We specifically identify compromising officers as those in the top quartile of the monthly complaint rate for each month an officer is active from 2015 to the beginning of the rollout.[17] Column 8 of Table 3 reports the results and reveals that, if anything, using complaints to identify officer type leads to an even greater change in behavior, a differential decrease of 16% compared to 12% in the primary specification.

Intensive Margin. Our primary result focuses on the extensive margin between compromising and protective activators, that is, it treats all compromising activators the same. However, compromising activators should differ across an intensive margin in terms of how much they respond to BWCs. We therefore investigate the extent to which changes in arrests differ more for the compromising officers at the extreme.

We do this in two ways. First, we compare a constant comparison group of high-activating officers above the 50th percentile with two different sets of low-activating officers. We specifically use a +/- 10 percentile range around each percentile from the 10th to the 40th. For example, the 10th percentile estimates compute the differential arrests per shift

[17]We cannot calculate this rate for some officers who were appointed around the time of the BWC rollout, so they are excluded from this analysis.

between officers in the 1st to 20th percentile to that of officers with above median activation rates, and the 40th percentile estimates compute the differential arrests per shift between officers in the 31st to 50th percentile to that of officers with above median activation rates, and so on.

Second, we study whether the power of this comparison can be enhanced with information about officers' prior complaint records. Specifically, we repeat this same exercise with a constant comparison group but instead impose rules on both the percentile of officer activation rates and the percentile of officer complaints. For example, the 10th percentile estimates compute the differential arrests per shift between officers in the 1st to 20th percentile for activation rates and between the 81st and 100th percentile of complaints to that of officers with both above median activation rates and below median complaints.

Figure 8 reports the results. We find evidence that the effects of BWCs are concentrated among the officers who likely find BWCs the most compromising. When we define compromising activators as those centered around the 10th percentile of activation, we find a 2.3 percentage point differential decrease in arrests. This falls to 1.5 percentage points when our definition of compromising activators is centered at the 40th percentile of activation. The differences are even more pronounced when using both activation rates and complaints. For officers with activation rates below the 20th percentile *and* complaint rates above the 80th percentile, BWC rollout induces a differential decrease in arrests of 7.9 percentage points. These results are strongly consistent with the model predictions of which officers have a higher incentive for strategic non-activation.

6.4 Robustness

We next assess the sensitivity of the primary estimates to alternative estimators and specifications.

Alterative Estimators. We consider three types of alternative estimators. First, using the same basic setup in the primary specification, we use various alternative estimators from the recent difference-in-differences literature in settings with staggered treatment adoption (Butts and Gardner, 2022; Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille,

2020; Sun and Abraham, 2021; Borusyak et al., 2024; Dube et al., 2023; Arkhangelsky et al., 2021). Second, using a dataset at the officer-arrest level, we extend a simple two-way fixed effects (TWFE) specification to incorporate our triple-differences setting by including fixed effects for the interaction of the presence of BWCs and compromising activator. To draw on changes in arrest behavior within officers, we include officer fixed effects. This specification may have negative weighting issues but it has the advantage of estimating a triple-difference specification directly. Finally, we use a stacked event study design (Cengiz et al., 2019; Deshpande and Li, 2019). For each BWC adoption event, we create a sub-dataset with treated districts that adopted BWCs from 6 months before the event to 12 months after and control districts that did not have BWCs in the event window.[18] Following the specification in Deshpande and Li (2019), we include event by district and event by calendar month fixed effects. To draw on within-officer variation, we conduct the analysis at the officer-arrest level and include officer fixed effects. Following Wing et al. (2024), we cluster standard errors at the district level and with their suggested weights, but the results are very similar without the weights. Appendix Table A1 reports the results and reveals largely consistent estimates.[19]

Choice of Activation Threshold. Recall from Section 4 that we are only able to directly link about half of arrests with BWC metadata via event identifiers and that some arrests and BWC activations in our data do not include these identifiers. We therefore estimated a statistical model of the arrests and activations we were are able to directly link and used it to label BWC activation for the arrests we are unable to directly link. This method requires specifying a threshold for labeling activations in terms of which arrest-activation pairs we accept as a match. Our primary approach accepts matches where the probability that the observed arrest time-activation time gap was drawn from our estimated

[18]We shorten our pre-period to 6 months to increase the number of controls that do not roll out in that period.

[19]The two exceptions are Goodman-Bacon (2021) and de Chaisemartin and D'Haultfœuille (2020), both of which find small and statistically insignificant effects. In the former case, this is due to the inclusion of "reverse" comparisons where a district that has already rolled out is used as a comparison for one that has not yet rolled out. When we remove these comparisons, the point estimate rises to -0.027. We think that de Chaisemartin and D'Haultfœuille (2020) is a poor fit for our setting due to its focus on changes immediately around treatment and the fact that treatment typically happens midway through a period in our data.

distribution of gaps is at least 5%. We assess the robustness of our results to the threshold used. Appendix Table A2 reports results using alternative threshold rules that are less demanding (a threshold at double the primary specification, 10%) and more demanding (a threshold at half of the primary specification, 2.5%) and find consistent results.

Noise in Officer Activation Rates. We use officers' observed activation rates as a proxy to identify officer type. However, observed activation rates are noisier for officers who make fewer arrests. This can lead us to misclassification of officer type, as both the group of high- and low-activating officers contains some officers in the other group. This in turn introduces measurement error in our outcome and attenuates estimates towards zero.

One potential solution to address noise would be to use Empirical Bayes methods to shrink noisy estimates towards the mean as we did in Section 5 when studying only activation rates. If noise is uncorrelated with officer type, this approach would likely reduce measurement error. However, we expect noise to be correlated with officer type because different types of officers make arrests at different rates. Therefore, through the shrinkage, the activation rates of protective activators/high-activating officers who make fewer arrests will be pulled down toward the mean more than the activation rates of compromising activators/low-activating officers who make more arrests will be pulled up to the mean. This misclassification will be worse for low-arrest officers, causing the middle of the shrunken distribution to mix more protective activators with the true compromising activators. Thus, the shrinkage approach could introduce non-classical measurement error and bias the estimates.

Given this concern with shrunken activation rates, we prefer to use unshrunken activation rates as our primary specification. However, the concern with the shrinkage approach is not as strongly present at the tails of the distribution, as the tails will consist of high- and low-activating officers with many arrests. We therefore use the shrunken approach but focus on the tails by repeating the intensive margin exercise from Figure 8. Appendix Figure A6 reports the results and reveals slightly larger estimates at the tails that are not statistically different from our primary specification.

7 Differences in Officer Strategic Non-Activation

This section investigates strategic non-activation and the resulting consequences for the design of BWC policies. To do so, we estimate the model from Section 2 to match the distribution of officer-level activation rates from Section 5 and the arrest responses to the presence of BWCs estimated in Section 6. The results in this section are, of course, specific to CPD in the late 2010s and should not be generalized to all police departments with BWCs or to other optional oversight settings. But with that said, our setting is an important one—the second largest police department in the US—and our results illustrate the novel tradeoffs involved with optional oversight settings.

7.1 Estimation

Appendix E describes in detail the estimation procedure that we summarize here. We begin by assuming distributions for key parameters. The arrest-level probability of misconduct, θ_i, is drawn from a uniform distribution bounded by 0 and A. Each officer's relative reward for an arrest, $\frac{w_j}{c}$, is drawn from a log-normal distribution. Idiosyncratic activation incentives, ϵ_i, are drawn from a distribution parameterized by p_l and p_h. With probability p_l, $\epsilon_i = -\infty$ and activation is impossible. With probability p_h, $\epsilon_i = \infty$ and activation is unavoidable. Otherwise, $\epsilon_i = 0$. This parameterization will make it easier to characterize different potential explanations for activation failure or disengagement due to a perceived need to activate.

Next, we compute summary statistics of the arrest and discipline environment to eliminate some parameters in the model. Following this, we are left with a vector of 7 parameters: $\{q_T, q_F, \theta^*_{act}(0), \delta_F, p_l, p_h, \sigma_u\}$. Here, $\delta_F = -\frac{q'_F}{q_F}$ captures the relative reduction in the probability of false positives achieved from activation, and σ_u adds mean-zero noise to officer-level activation rates.

Finally, we choose these parameters to match 11 moments: 9 moments of the distribution of adjusted officer-level activation rates (the 1st, 5th, 10th, 25th, 50th, 75th, 90th, 95th, and 99th percentiles), and the estimated effect of BWCs on arrests differentially by officer type and overall (Columns 1 and 4 from Table 2).

7.2 Strategic Non-Activation

Having found that differences in officer activation rates cannot be explained by observable covariates but are consistent with our model of optional oversight, we now leverage the estimated model to investigate the prevalence of compromising activators and how much of the variance in officer-level activation can be explained by strategic non-activation.

Panel A of Figure 9 reports simulated officer activation rate data that divides officers by their latent compromising or protective type (those with $\theta^*_m < \theta^*_{arr}$ and $\theta^*_{arr} < \theta^*_m$). Panel B reports the share of non-activations due to strategic behavior among the compromising activators.[20] We find that 55% of officers are compromising activators, meaning they strategically fail to activate at least sometimes. These compromising activators activate less overall and have a greater spread of activation rates. This is because their activation decisions depend not just on idiosyncratic factors but also on their willingness to make risky arrests in which they prefer to keep the BWC off.

To estimate how much of the differences in activation rates between officers is explained by strategic non-activation, we calculate $\gamma^{\mathbf{x}}$ as we did with situational factors in Equation 8 but using simulated distributions with and without strategic non-activation. We find that strategic non-activation accounts for as much as 62% of officer-level activation variation after adjusting for environment, encounter, and role.[21]

These estimates assume that officer incentives for strategic non-activation are unrelated to idiosyncratic activation failures. It is not obvious that this assumption will hold. For example, the assumption may not hold if compromising activators who strategically fail to activate can select into certain districts, assignments, or arrest roles where expected activation is lower. We use officer activation rates that adjust for situational factors in order to account for this and related concerns, but there could still be unobserved factors that would cause officer strategic non-activation to be correlated (or inversely correlated) with inability

[20] Appendix Figure A7 reports the simulated distribution of officer activation rates overall and compares it against the data.

[21] The simulated distribution with strategic non-activation has a variance of 0.033. The simulated distribution without strategic non-activations has a variance of 0.012.

to activate for idiosyncratic reasons. We assess the robustness of the results to a failure of this assumption by allowing idiosyncratic activation failure, p_l, to directly covary with the arrest incentives that shape whether an officer is a protective or compromising activator, $\frac{w_j}{c}$. Under conservative assumptions that allow a standard deviation for p_l of 0.1 and a correlation coefficient with $\frac{w_j}{c}$ of up to 0.5, we find that strategic non-activation still explains up to 24% of the variance across officers.

7.3 Policy Design with Optional Oversight and Strategic Non-Activation

Next, we study how BWC policies operate in a system of optional oversight and contrast it with a counterfactual regime of non-discretionary oversight. Throughout, we will consider effects on three outcomes: activation, arrests, and misconduct.[22] Police departments may care directly about a high rate of activation, both because it improves evidence gathering for downstream criminal justice proceedings and because it fulfills oversight promises made to the community. However, police departments also want officers to make justified arrests and not engage in misconduct. Our goal is to study how those tradeoffs depend fundamentally on whether oversight is discretionary and characterize them under different policy regimes in the optional oversight setting.

We will consider two types of policies. The first type of "reminder" policy aims to ensure that officers maintain functioning BWCs and remember to activate even in the heat of the moment. In our model, this corresponds to a reduction in p_l, the probability that officers do not activate even if it would be in their interest to do so. The second type of "enforcement" policy is stricter enforcement of activation rules designed to detect and punish strategic non-activation. In our model, this corresponds to an increase in p_h, the probability that officers must activate even if they would prefer not to. For each pair of policies $(p_l, p_h) \in [0,1] \times [0,1]$, we use our model to simulate activation rates, arrest rates

[22]Note that we will rely on our model-based estimate of actual misconduct events, not reports of misconduct or sustained misconduct. The reason we do not estimate our model using moments of complaints or sustained complaints—common proxies for misconduct—is that BWCs are likely to directly impact the decision to file complaints and the investigation process. While we therefore do not place heavy trust in the magnitudes of these simulated misconduct figures, the comparisons of relative misconduct across policies are still informative.

per shift, and misconduct rates per shift for each simulated officer.

We proceed in two steps. First, we characterize the set of outcomes available under non-discretionary oversight and contrast it with the expanded outcome space that arises under optional oversight. Second, we examine how reminder and enforcement policies move departments within this outcome space.

Outcomes Space. Figure 10 reports activation rates (x-axis), average arrests per shift (y-axis), and misconduct events per shift (shading) separately for protective activators and compromising activators as well as department-wide outcomes. The colored region is feasible for some (p_l, p_h) such that $p_l + p_h \leq 1$, and the white region is not feasible.

To isolate the role of officer discretion over monitoring, we begin with a benchmark in which optional oversight is shut down. In our model, these are settings where $p_l + p_h = 1$, and they form the dashed line at the bottom boundary of each shape in Figure 10. Along this boundary, activation is non-discretionary: BWCs are either always on ($p_l = 0$ and $p_h = 1$), always off ($p_l = 1$ and $p_h = 0$), or randomly on with probability p_h. In the intermediate settings, officers know when they are being monitored, but activation may occasionally fail for idiosyncratic reasons that affect all officers and arrests in the same way.

Figure 10 shows that under non-discretionary oversight, departments face a sharp tradeoff. Departments that adopt a new technology or enact procedures that ultimately increase monitoring decrease not only misconduct but also arrests. Compromising activators respond to increased monitoring by reducing both arrests and misconduct, while protective activators increase both. In our setting, compromising activators make up roughly half of officers and react more strongly to oversight than protective activators, so they dominate the department-wide pattern. In a setting where compromising activators are less prevalent, it is not necessarily the case that increasing monitoring increases both arrests and misconduct. In either case, lower misconduct can only be achieved by fewer arrests.

Optional Oversight and Expanded Policy Possibilities. Optional oversight fundamentally alters the tradeoff between arrests, monitoring, and misconduct. When $p_l + p_h < 1$, the discretion over monitoring expands the set of feasible outcomes, as shown by

the full shapes in Figure 10. However, this is only driven by compromising activators, as protective activators are unaffected by optional oversight because they activate whenever possible.

The circles in Figure 10 mark the outcome estimated from the data ($p_h = 0.31,\ p_l = 0.18$). The gray line tangent to this point indicates all points with the same total of p_h and p_l, which replicates the tradeoff shown by the $p_h + p_l = 1$ line. However, optional oversight allows departments to move off this line by adjusting enforcement and reminder policies independently. For example, the square point maximizes arrests without decreasing activation rates, although this comes at the cost of more misconduct. The triangle highlights an option that holds arrests constant while decreasing misconduct and increasing activation rates. Reaching this point requires the strongest reminder policy where officers never forget to activate (by driving p_l to 0). Reaching this point may be expensive or logistically difficult, but this nonetheless demonstrates why police departments might *prefer* optional oversight to non-discretionary oversight, especially if they do not want to curb arrests. Furthermore, the circular point that we observe is on precisely the same left-right line as the diamond point, which marks outcomes with no BWCs. One possible implication of this is that CPD has already selected enforcement and reminder policies in such a way as to hold arrests constant while simultaneously reducing misconduct.

Mapping Policy Space to the Outcomes Space. Finally, we investigate what policies interact with optional oversight to expand the space of outcomes available in favorable directions. To do so, Figure 11 reports the movement around the outcome space from changes in p_l and p_h.

We find stark differences in the effects of strengthening an enforcement policy as compared to a reminder policy. As a department strengthens its enforcement policy (through an increase in p_h), it has no effect on protective activators because they already always activate whenever they are able, but for compromising activators it increases activation rates but decreases arrest and misconduct rates, a pattern that carries over into the department-wide outcomes.

In contrast, as a department strengthens its reminder policy (through a decrease in p_l), activation rates increase for both officer types, but the effect on arrests and misconduct differs by officer type. Protective activators make more arrests because their actions are more likely to be recorded, which increases their willingness to make arrests. This also produces a marginal increase in misconduct resulting from the additional arrests. Compromising activators make the same number of arrests because they were already willing to make their riskiest arrest without the BWC on, but this comes with no change in misconduct. At the department level, strengthening the reminder policy increases activation and arrests but marginally increases misconduct because protective activators make more arrests.

7.4 Takeaways: Optional versus Non-Discretionary Oversight

The results highlight a fundamental distinction between non-discretionary and optional oversight regimes. Under non-discretionary oversight, increases in monitoring applies to all officers and arrests equally, thus forcing a tradeoff between arrests, monitoring, and misconduct. Optional oversight relaxes this constraint and expands the set of outcomes departments can achieve.

With optional oversight, enforcement and reminder policies operate on different dimensions of discretion and therefore have qualitatively different effects. Enforcement policies increase activation while curbing misconduct but reduce arrests. Reminder policies are effective at increasing activation while maintaining or increasing arrests, with limited effects on misconduct. Under optional oversight, departments can therefore independently strengthen reminder policies, allowing activation to increase for compromising officers in some encounters without suppressing arrests overall, while simultaneously increasing activation and arrests for protective officers.

A combined policy of enforcement and reminders can reallocate arrests from compromising activators to protective activators, holding total arrests constant while reducing misconduct. This mechanism explains how departments can increase activation and decrease misconduct while holding arrests constant, a flexibility that is not available under non-discretionary monitoring. That said, this strategy is ultimately limited: once accidental

non-activation is eliminated, further reductions in misconduct require fewer arrests.

Finally, optimal policy choices depend on how departments weigh activation, enforcement activity, and misconduct. Adopting BWCs with relaxed enforcement—consistent with the CPD experience we study—can preserve arrest rates but delivers limited improvements in misconduct. Stricter enforcement of activation rules may substantially reduce misconduct but risks placing oversight goals in tension with enforcement priorities. More broadly, our results suggest that the central design question for oversight technology is not how intensively to monitor but instead what benefits can be reaped from making monitoring discretionary.

8 Conclusion

Privacy, legal, or practical constraints prevent many employers from fully controlling when and how they can monitor worker behavior. We studied a setting where workers retain discretion over when their behavior is observed: policing. Officer discretion creates opportunities for strategic behavior, not only in the activation decision but also in decisions about engagement and arrests. We modeled officers as protective activators who use BWCs to document good conduct when claims of misconduct arise or compromising activators who sometimes strategically fail to activate if they anticipate a high risk of misconduct. The model predicts that officers' responses to BWCs vary by officer type, with distinct changes in activation and arrest behavior.

To connect the model to data, we used Chicago data on activation of specific officers as they make specific arrests where they have an unambiguous duty to activate, which required linking arrests to BWC metadata using a statistical model of activation and arrest times. Using this setting and the staggered rollout of BWCs to police districts across Chicago, we found reduced-form evidence consistent with the two major model predictions—there are large differences in activation rates across officers that reflect officer-level components, and compromising activators respond to the introduction of BWCs by making fewer arrests compared to protective activators.

Using these estimates as key moments to identify and calibrate our model of optional oversight, we built insight into how and why the optional oversight setting differs from simpler

oversight regimes and the resulting implications for policy design. We found that policies designed to remove barriers to activation behave differently from policies designed to punish strategic non-activation, and the targeted combination of the policies has the potential to break tradeoffs between activation rates, arrest rates, and misconduct rates in ways that are not possible in a simpler model of non-discretionary oversight.

Future research should extend the optional oversight framework to other settings. These settings are becoming increasingly common, from university faculty recording classes to remote workers working offline to avoid being monitored. Although the classification of workers into protective and compromising activators will be a useful theoretical framing in other settings, the specific tradeoffs associated with monitoring technologies may differ for several reasons. For example, the detection of misconduct even in the absence of a recording may be considerably higher than with policing, such as in a university class where students can corroborate the professor's statements. Moreover, not all optional oversight is on a case-by-case basis as in policing, such as with university classes either being automatically recorded each session or not recorded at all. Finally, worker protections may be weaker, so strategically avoiding monitoring can have more severe employment consequences, such as a non-unionized remote worker being fired for repeatedly working offline.

With police BWCs and in other optional oversight settings, measurement of oversight preferences of individual workers is likely to be a limiting factor in understanding the impacts of monitoring policies. Future research on BWCs and in other settings should focus on measuring worker oversight preferences directly rather than relying, as this study does, on proxies based on other behaviors. More direct measurement can increase the precision of estimates and build evidence for the theory of optional oversight.

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.

Adhvaryu, A., A. Nyshadham, and J. Tamayo (2022). An Anatomy of Performance Monitoring. *Working Paper*.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic Difference-in-Differences. *American Economic Review 111*(12), 4088–4118.

Ba, B. and R. Rivera (2024). The Effect of Police Oversight on Crime and Allegations of Misconduct: Evidence from Chicago. *Review of Economics and Statistics*.

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

Barbosa, D., T. Fetzer, C. Soto, and P. Souza (2021). De-Escalation Technology: The Impact of Body-Worn Cameras on Citizen-Police Pnteractions. *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.

Bollman, K. (2024). The Effects of Body-worn Cameras on Policing and Court Outcomes: Evidence from the Court System in Virginia. *Working Paper*.

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting Event-Study Designs: Robust and Efficient Estimation. *Review of Economic Studies 91*(6), 3253–3285.

Bradley, B. (2021). Here's Why Chicago's Unofficial Police Blog 'Second City Cop' Went Dark.

Butts, K. and J. Gardner (2022). did2s: Two-Stage Difference-in-Differences. *The R Journal 14*(3), 162–173.

Cai, J., S. Luo, and S.-Y. Wang (2025). Money or Monitoring: Evidence on Improving Worker Effort. *NBER Working Paper #33977*.

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

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics 134*(3), 1405–1454.

Chalfin, A. and F. M. Gonçalves (2025). Professional Motivations in the Public Sector: Evidence from Police Officers. *NBER Working Paper #31985*.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014). Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates. *American Economic Review 104*(9), 2593–2632.

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

CPD (2016). Mayor Press Office, Mayor Emanuel, Chicago Police Department Announce Expedited Expansion of Body Worn Cameras Program.

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

Deshpande, M. and Y. Li (2019). Who is Screened Out? Application Costs and the Targeting of Disability Programs. *American Economic Journal: Economic Policy 11*(4), 213–248.

Devi, T. and R. G. Fryer Jr. (2020). Policing the Police: The Impact of "Pattern-or-Practice" Investigations on Crime. *NBER Working Paper #27324*.

Dube, A., D. Girardi, O. Jorda, and A. M. Taylor (2023). A Local Projections Approach to Difference-in-Differences. *NBER Working Paper #31184*.

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

Friebel, G., M. Heinz, M. Hoffman, T. Kretschmer, and N. Zubanov (2025). Is This Really Kneaded? Identifying and Eliminating Potentially Harmful Forms of Workplace Control. *NBER Working Paper #34122*.

Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics 225*(2), 254–277.

Guiteras, R. P. and B. K. Jack (2018). Productivity in Piece-Rate Labor Markets: Evidence from Rural Malawi. *Journal of Development Economics 131*, 42–61.

Harrington, E. and H. Shaffer (2025). Surveilling the Police: Do Body-Worn Cameras Prevent Police Bias from Becoming Systemic? *Working Paper*.

Holmstrom, B. and P. Milgrom (1991). Multitask Principal–Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *Journal of Law, Economics, and Organization 7*, 24–52.

Hubbard, T. N. (2000). The Demand for Monitoring Technologies: The Case of Trucking. *Quarterly Journal of Economics 115*(2), 533–560.

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.

Kelley, E. M., G. Lane, and D. Schönholzer (2024). Monitoring in Small Firms: Experimental Evidence from Kenyan Public Transit. *American Economic Review 114*(10), 3119–3160.

Khan, A. Q., A. I. Khwaja, and B. A. Olken (2015). Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors . *Quarterly Journal of Economics 131*(1), 219–271.

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

Kucherenko, S. and Y. Sytsko (2005). Application of Deterministic Low-Discrepancy Sequences in Global Optimization. *Computational Optimization and Applications 30*, 297–318.

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*.

Nelder, J. A. and R. Mead (1965). A Simplex Method for Function Minimization. *Computer Journal 7*(4), 308–313.

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 #8664*.

Prendergast, C. (2003). The Limits of Bureaucratic Efficiency. *Journal of Political Economy 111*(5), 929–958.

Prendergast, C. (2007). The Motivation and Bias of Bureaucrats. *American Economic Review 97*(1), 180–196.

Prendergast, C. (2022). "Drive and Wave": The Response to LAPD Police Reforms After Rampart.

Rivera, R. (2025). Are "Bad" Cops Better Police? The Tradeoff Between Officer Aggression and Public Safety. *Working Paper*.

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 11*(2), 225–268.

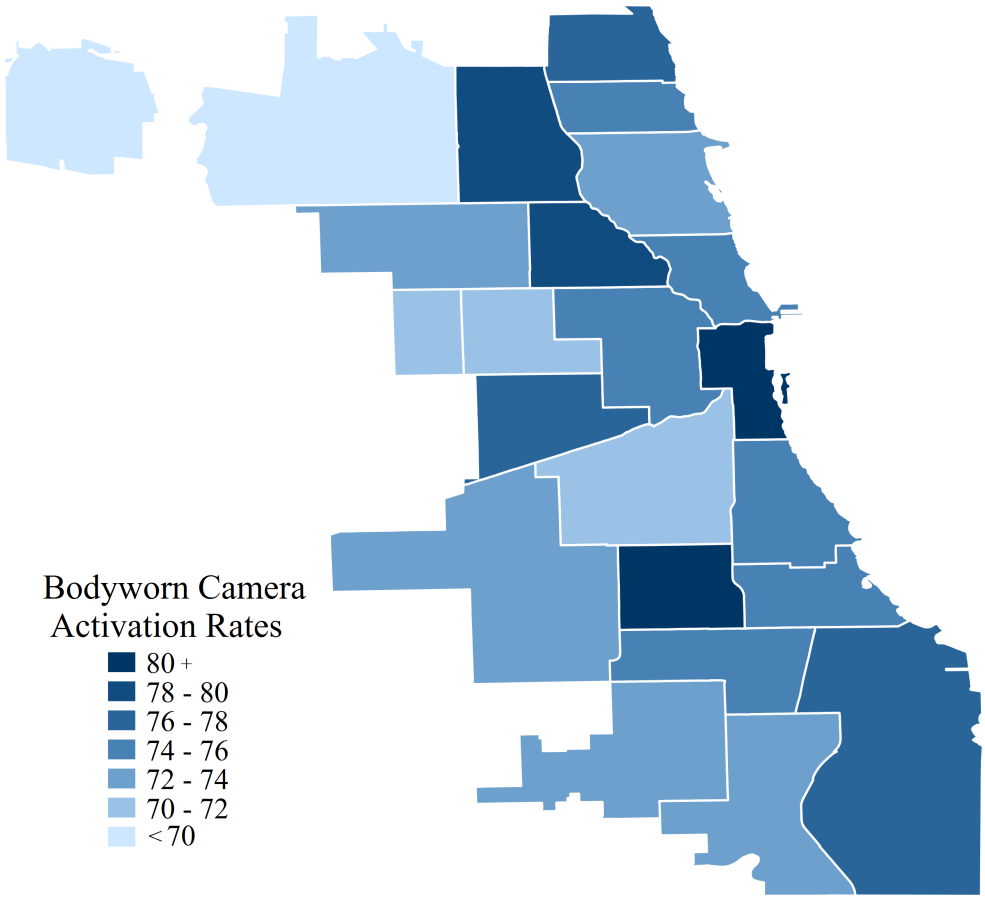
Rozema, K. and M. Schanzenbach (2023). Does Discipline Decrease Police Misconduct? Evidence from Chicago Civilian Allegations. *American Economic Journal: Applied Economics 15*(3), 80–116.

Sun, L. and S. Abraham (2021). Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics 225*(2), 175–199.

Weisburst, E. K. (2024). Whose Help Is on the Way? *Journal of Human Resources 59*(4), 1122–1149.

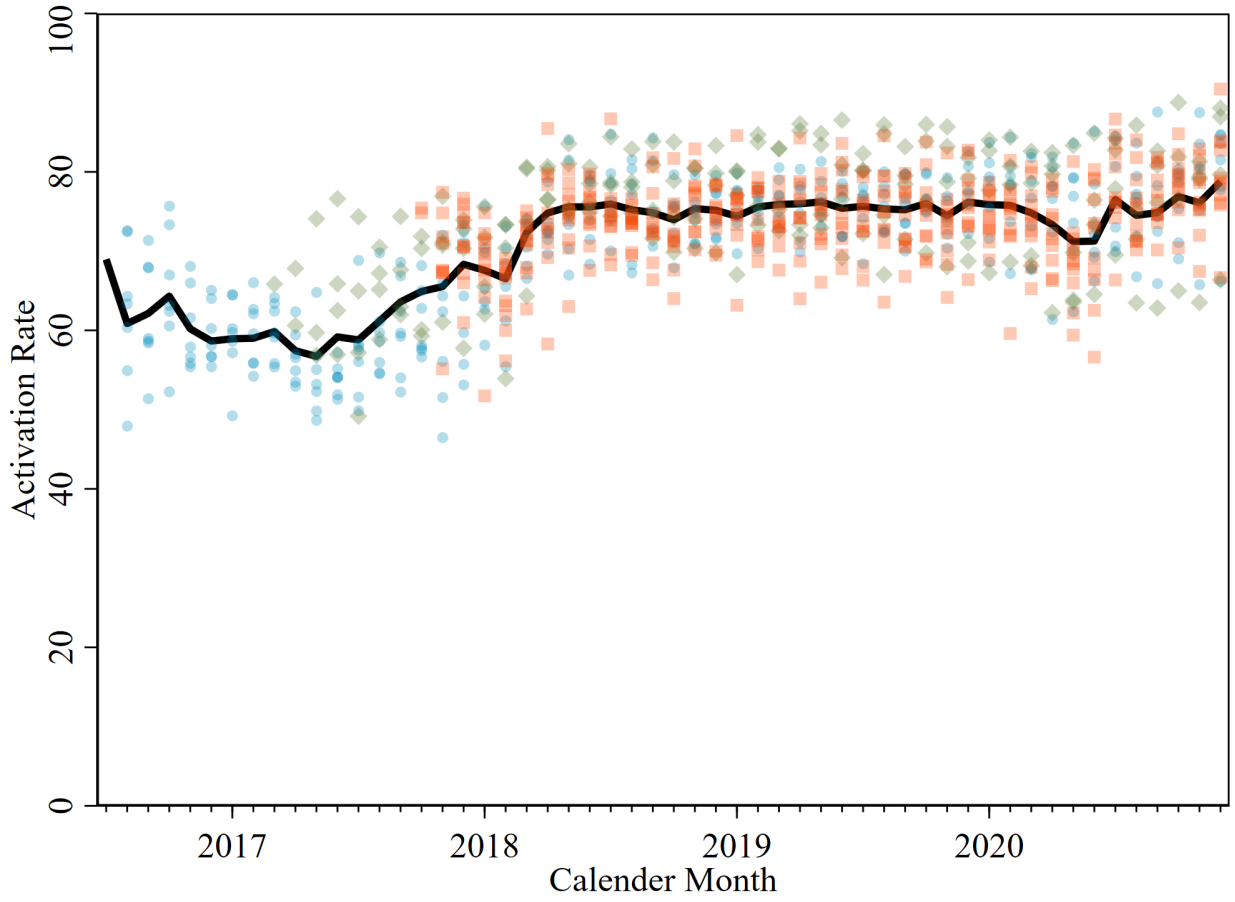
Wing, C., S. M. Freedman, and A. Hollingsworth (2024). Stacked Difference-in-Differences. *NBER Working Paper #32054*.

Figure 1: Body-worn Camera Activation Rates Across Police Districts



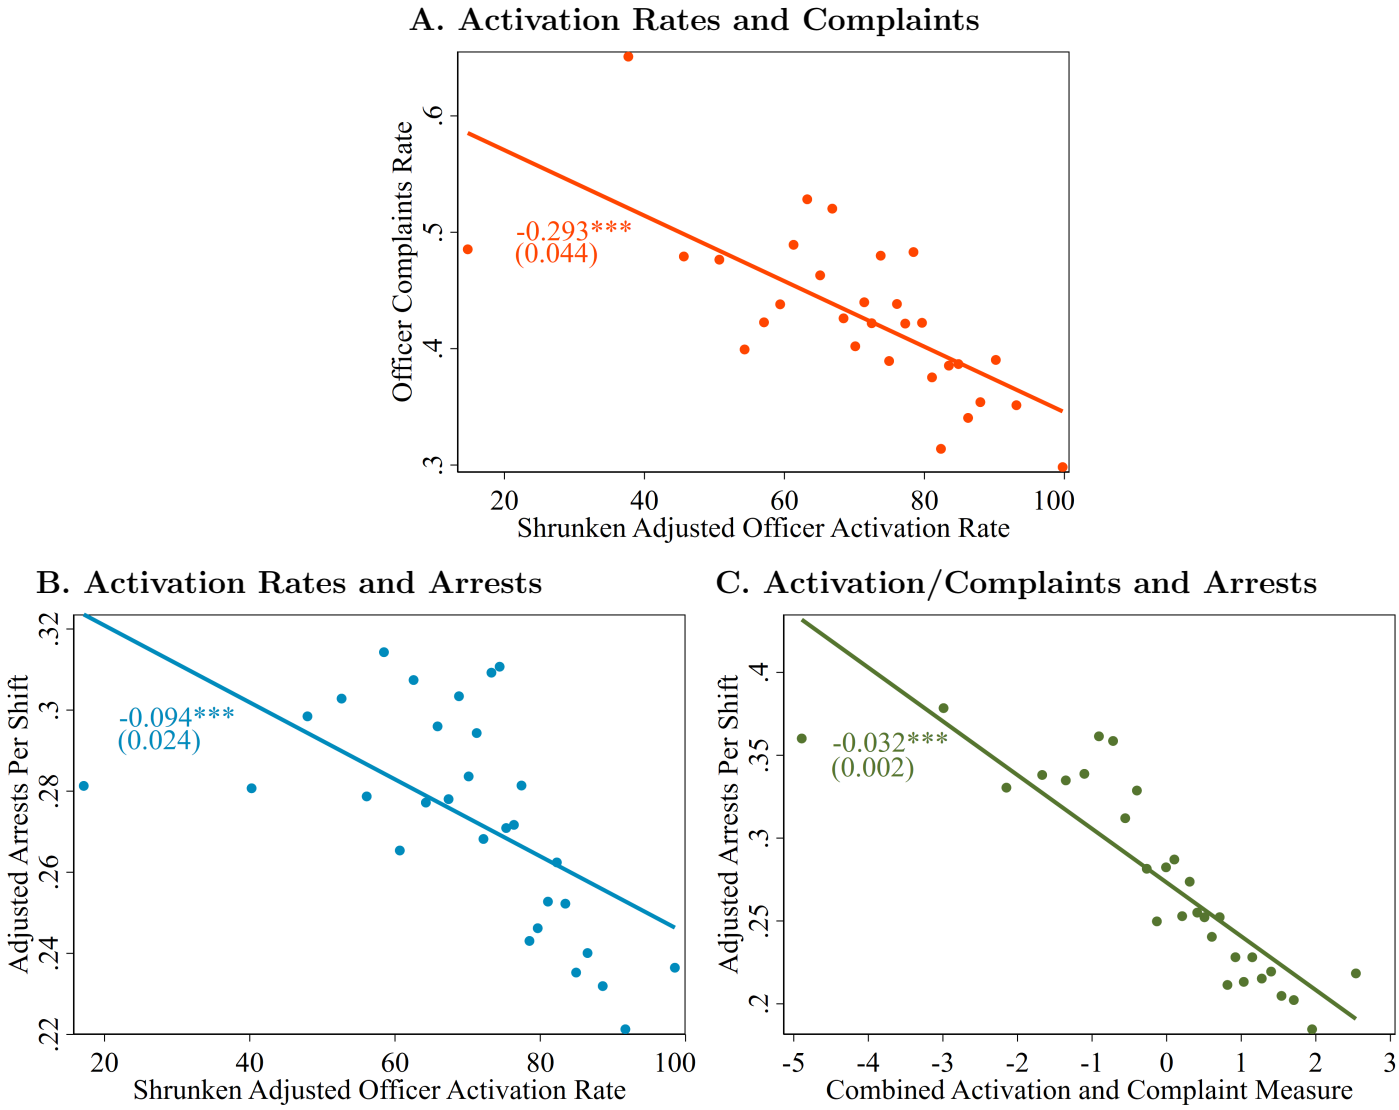
Note: This figure reports a heat map of activation rates across police districts. The rate is computed at the officer-arrest level for all arrests made after the district had body-worn cameras.

Figure 2: Body-worn Camera Activation Rate Over Time



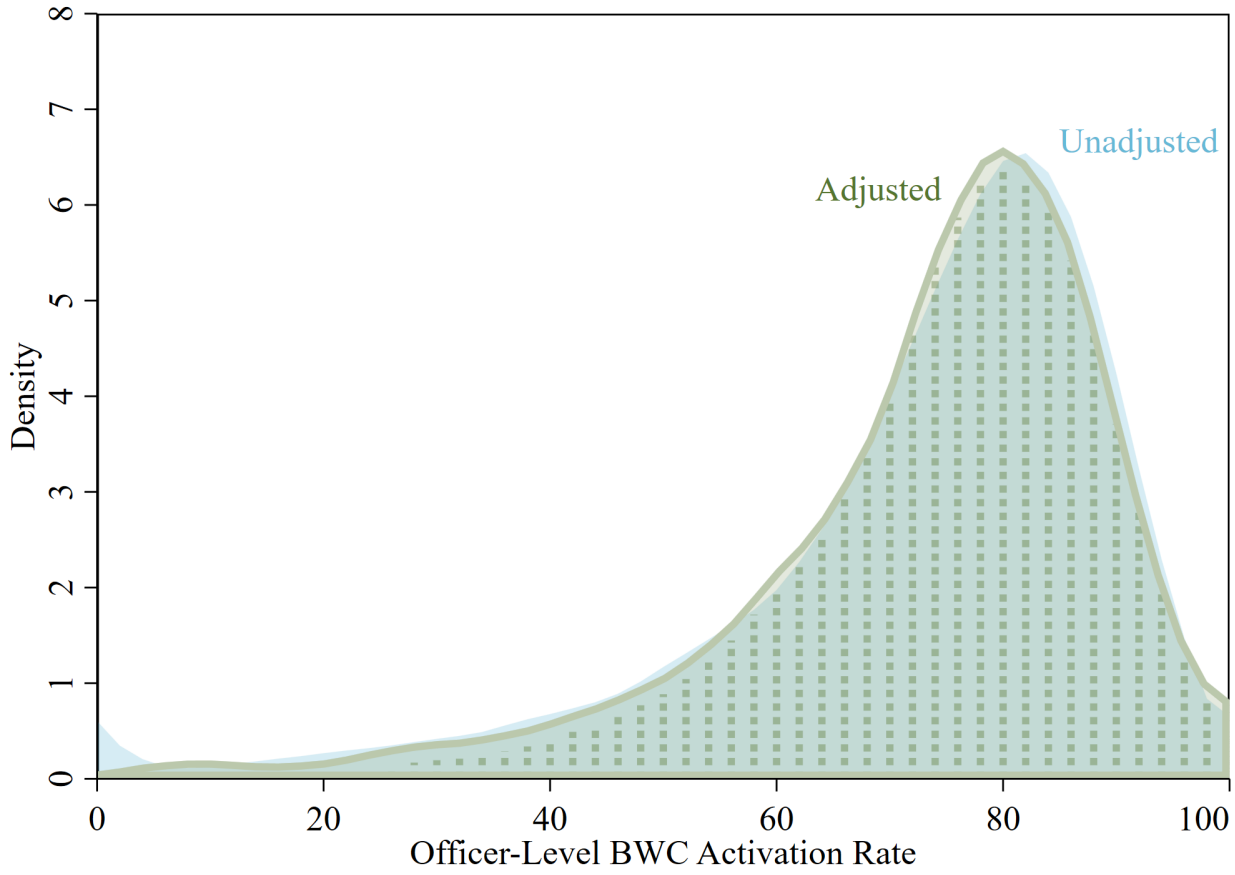
Note: This figure reports the activation rate over time across all districts that had body-worn cameras as of that month (black line) and separately by district (colored markers). The blue circle markers are for the six districts that adopted body-worn cameras from June 2016 to February 2017, the green diamond markers are for the five districts that adopted body-worn cameras from March 2017 to July 2017, and the square orange markers are for the eleven districts that adopted body-worn cameras from October 2017 to December 2017.

Figure 3: Relationship between Activation, Complaints, and Arrests



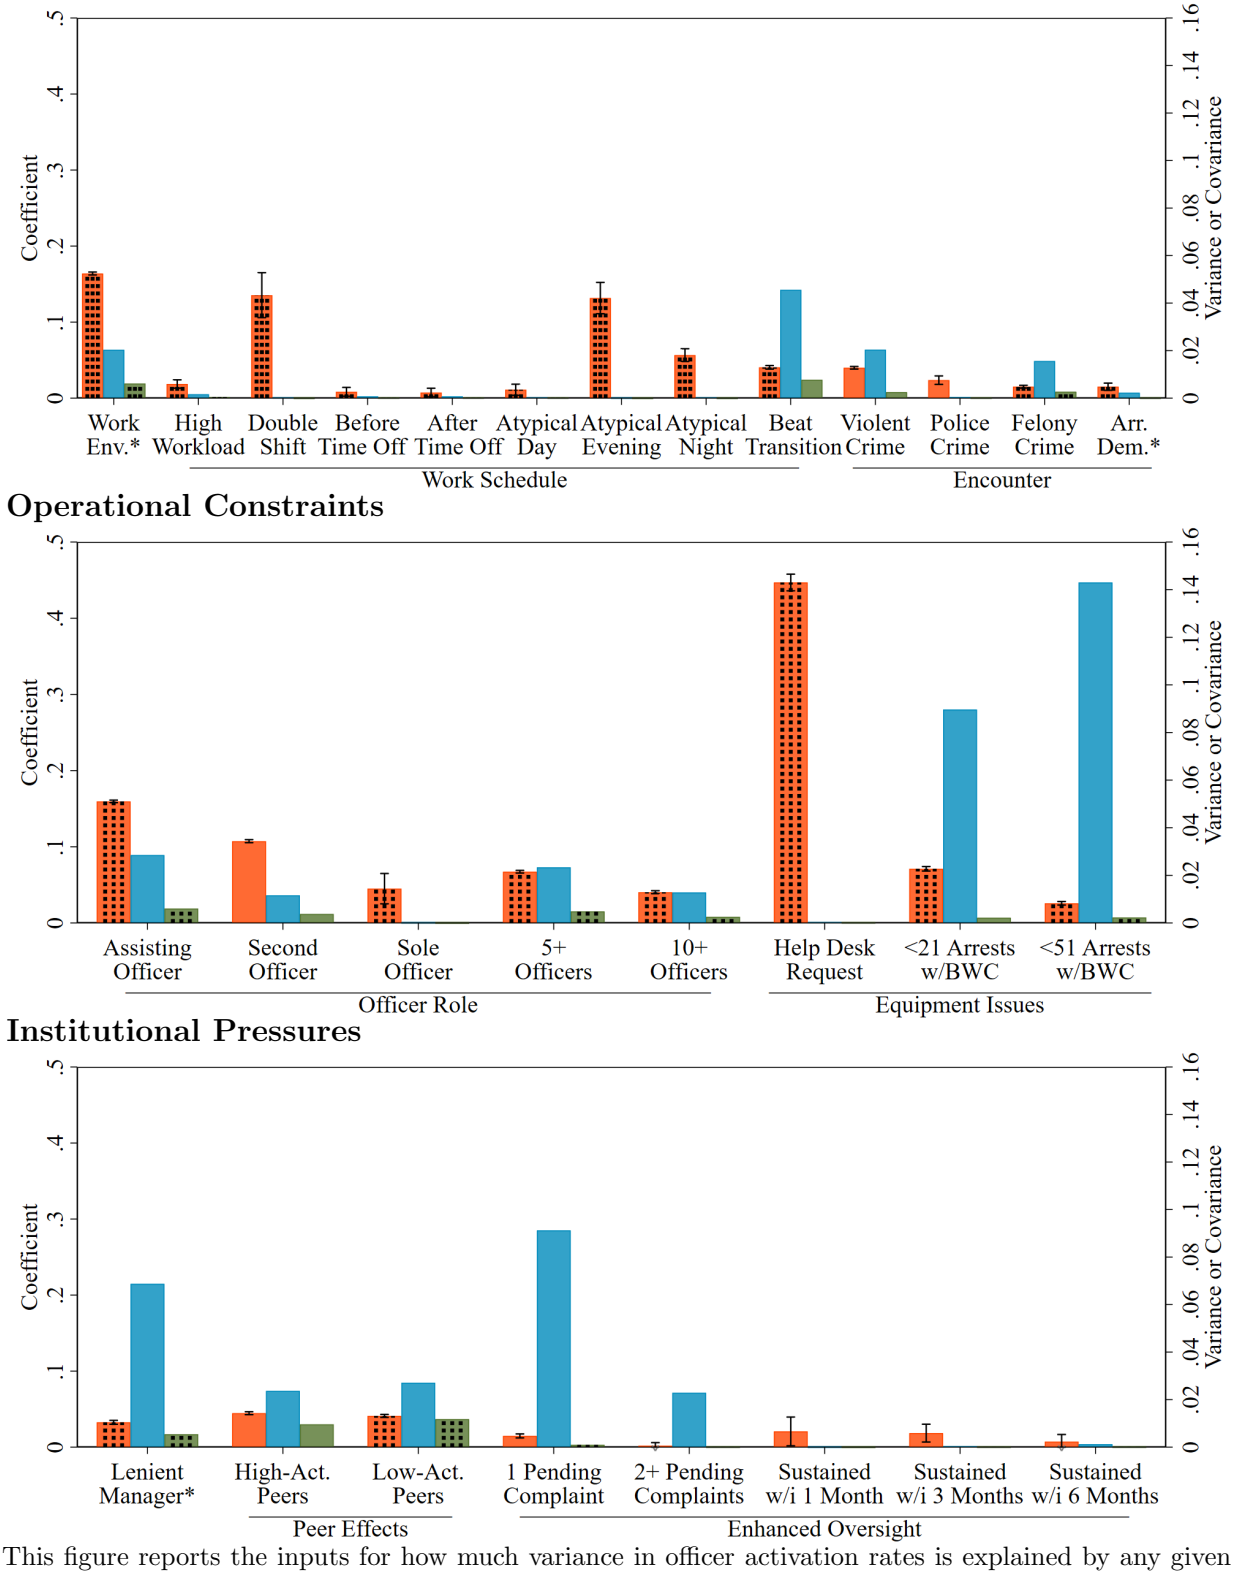
Note: This figure investigates the observational relationship between activation, complaints, and arrests. Each panel reports a binned scatterplot at the officer-level of 20 equally sized bins. Panel A reports the relationship between our preferred activation rate measure (that controls for situational factors and officer demographics and that shrinks noisy fixed effects to the mean through Empirical Bayes) and complaints. Panels B and C report the relationships between arrests and measures associated with officer type. Panel B uses our preferred activation rate measure, and Panel C uses a proxy that combines our preferred activation rate measure with a measure of officer complaints. The measure of officer complaints is the annual complaint rate computed from all years between 1999 and 2015 where the officer was on the force. In Panel C, the combined activation-complaint measure is calculated by subtracting each officer's z-score in the complaint distribution from their z-score in the activation distribution, and is higher when officers activate more and receive fewer complaints, both associated with protective activators. In Panels B and C, the arrests per shift measure controls for district-by-calender month fixed effects.

Figure 4: Distribution of Officer Activation Rates



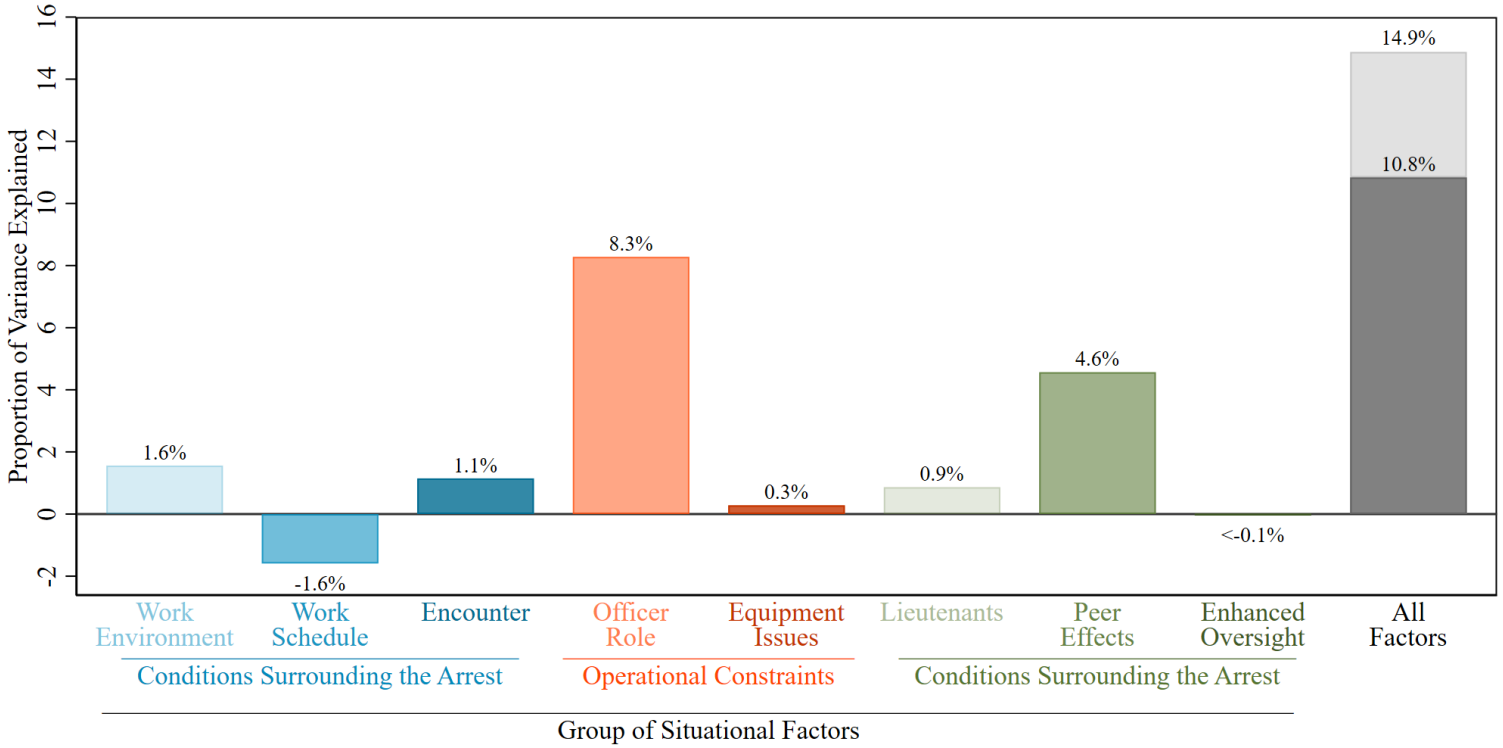
Note: This figure reports the distributions of officer activation rates (blue area labeled "Unadjusted") and officer activation rates after accounting for environment, encounter, and role controls (green area with dots labeled "Adjusted"). Officer activation rates are estimated using arrests after all districts had body-worn cameras (2018-2020). Officers with 10 or fewer arrests in this period are excluded.

Figure 5: The Role of Exposure to Different Situations in Body-worn Camera Activation



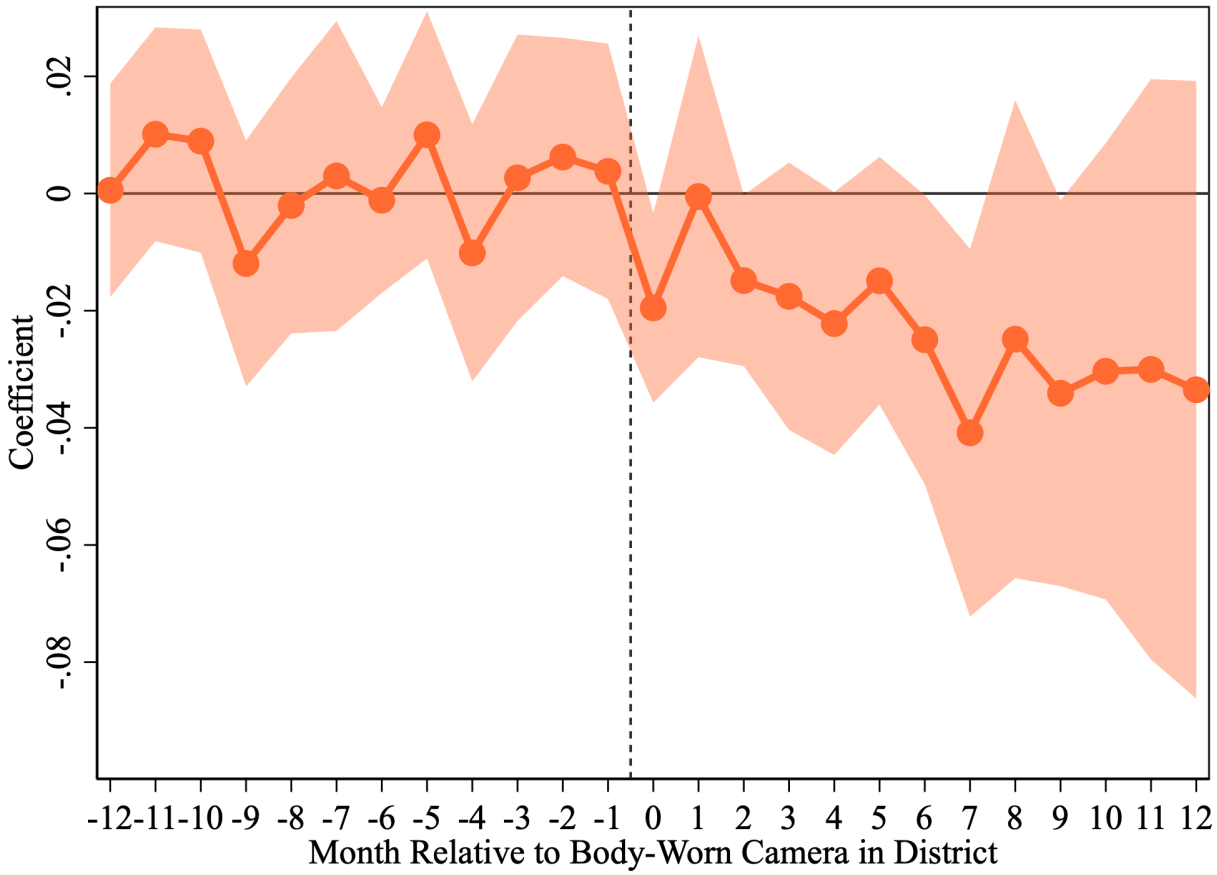
Note: This figure reports the inputs for how much variance in officer activation rates is explained by any given situational factor. In Equation 7, how much variance a factor X_{ij} explains depends on how much X_{ij} directly affects activation (β), how much X_{ij} varies across officers ($\text{Var}(\bar{X}_j)$), and how much X_{ij} covaries with officer fixed effects ($\text{Cov}(\bar{X}_j, \theta_j)$). The orange bars are β coefficients estimated from versions of Equation 6 that regress activation in an officer-arrest on different situational factors as indicated on the x-axis in the panels. The estimates for a given type of situational factors are estimated in a separate regression, but we draw horizontal lines underneath the labels on the x-axis to identify different categories of situational factors, such as "Work Schedule" and "Encounter" in Panel A. The blue bars labeled "Variance" captures how much the factor varies across officers, as measured by the variance of the variable ($\text{Var}(\bar{X}_j)$). The green bars labeled "Covariance" is the covariance of the factor with officer fixed effects ($\text{Cov}(\bar{X}_j, \theta_j)$). For the coefficients or the covariances, the bars have a black pattern if the estimates are negative. The estimated coefficients are reported on the left y-axis, and the variance and covariance estimates are reported on the right y-axis. The underlying variables for the factors that are labeled on the x-axis with stars (*) have multiple categories. The results shown compare categories with activation rates in the lowest quartile to all others.

Figure 6: Proportion of Variance in Officer Activation Rates Explained by Situational Factors



Note: This figure reports the results of the decomposition of differences in officer activation rates into three broad types of situational factors, with different groups within each type. The factors in the bars in shades of blue relate to the immediate conditions surrounding the arrest stemming from the officer's work environment, work schedule, and the encounter with the suspect which may affect activation. The factors in the bars in shades of orange relate to operational constraints during an arrest related to the officer's role and equipment which may affect activation. The factors in the bars in shades of green relate to institutional pressures due to officers' managers who enforce activation rules, the behavior of peer officers making the arrest, and enhanced oversight from the department may affect activation. Each bar summarizes the explanatory power of a group of situational factors by reporting the percentage of the overall variance in officer activation rates explained by those factors, as captured by $\gamma^{\mathbf{X}}$ in Equation 8. A group of situational factors with positive estimates implies that they explain some of the differences in officer activation rates, meaning that there would be smaller differences in officer activation rates once they are adjusted for those factors. A group of situational factors with negative estimates implies that they mask some of the true differences in officer activation rates, meaning that there would be larger differences in officer activation rates once they are adjusted for those factors. The bars differ by the group of situational factors included, where each of the colored bars estimates the portion of the variance in officer activation rates explained by those factors in isolation. The dark gray bar to the right labeled "All Factors" is the estimated portion of variance explained when all factors are considered, where the height of the light gray bar adds up the individual estimates.

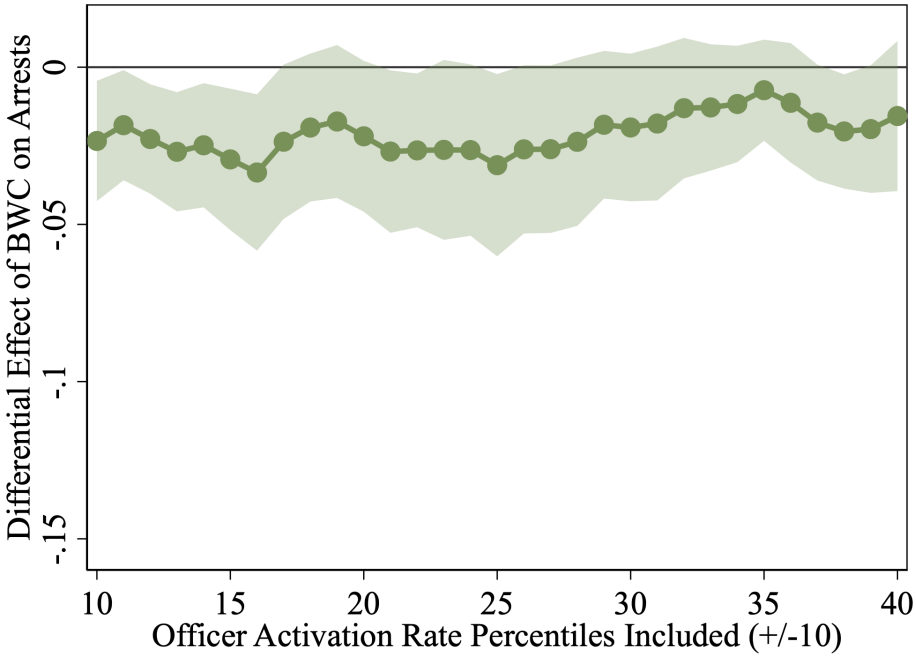
Figure 7: Differential Effect of Body-worn Cameras on Officer Arrest Behavior by Officer Type



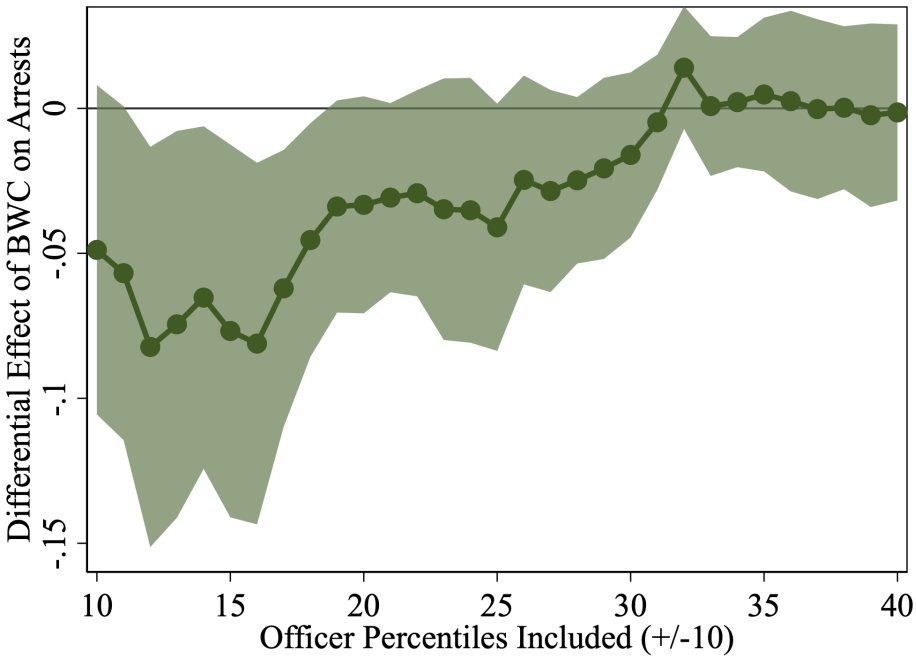
Note: This figure estimates the differential effect of body-worn cameras on arrests per shift between compromising and protective activators. We proxy for compromising activators using the 25th percentile of officer activation rates that account for exposure to different arrest situations using data after 2017. The estimates compare the differential number of arrests per shift in police districts that adopt body-worn cameras in the 12 months before and 12 months after adoption to the same outcome in districts that have not yet adopted in a difference-in-differences design in Callaway and Sant'Anna (2021). The line reports the estimated coefficients (under Callaway and Sant'Anna (2021), there is no leave-out period). The shaded region reports 95% confidence intervals. Standard errors are clustered at the district level.

Figure 8: Intensive Margin: Differential Effect of Body-worn Cameras on Arrests for Officers who Find Body-worn Cameras More or Less Compromising

A. Sorting Officers by Activation Rates



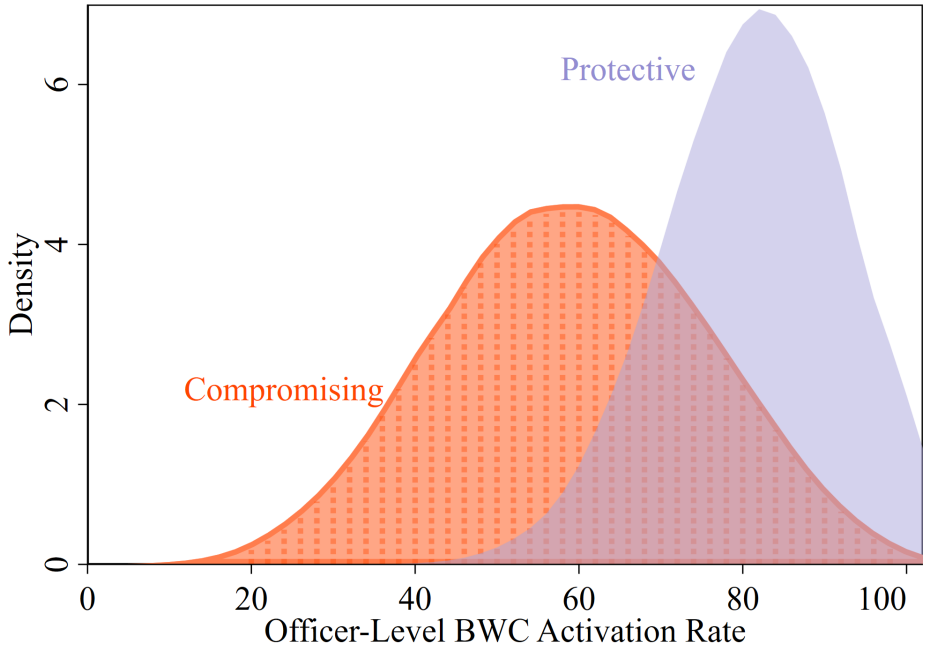
B. Sorting Officers by Activation Rates and Complaints



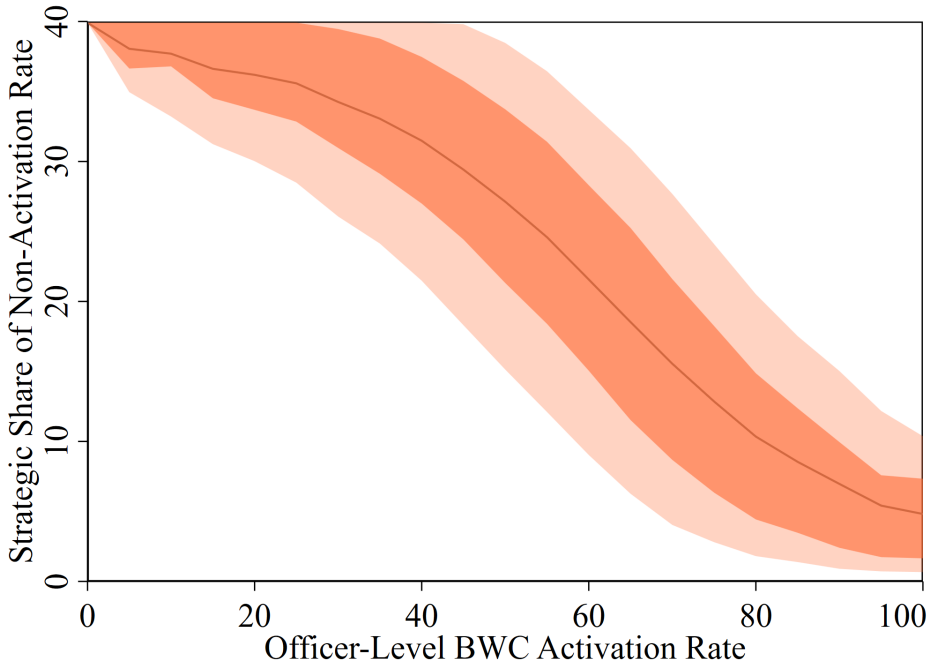
Note: This figure investigates the extent that the differential effect of body-worn cameras on arrests is stronger for the compromising officers who find cameras the most compromising by reestimating Column 1 of Table 2 using a constant comparison group of protective activators above the 50th percentile of the relative proxy for protective activators but with different sets of compromising activators. Panel A uses officer activation rates as a proxy for compromising versus protective activator. It reports a +/- 10 percentile range around each percentile of officer activation rates as indicated on the x-axis, from the 10th to the 40th percentile of activation. Panel B uses both officer activation rates and complaints as a proxy for compromising versus protective activator. It uses a +/- 10 percentile range around each percentile of officer activation rates and complaints as indicated on the x-axis, including only officers who meet both criteria. High complaints are associated with compromising activators, so we subtract the percentiles from 100 before making the restriction.

Figure 9: Strategic Non-Activation

A. Activation Rates by Officer Type

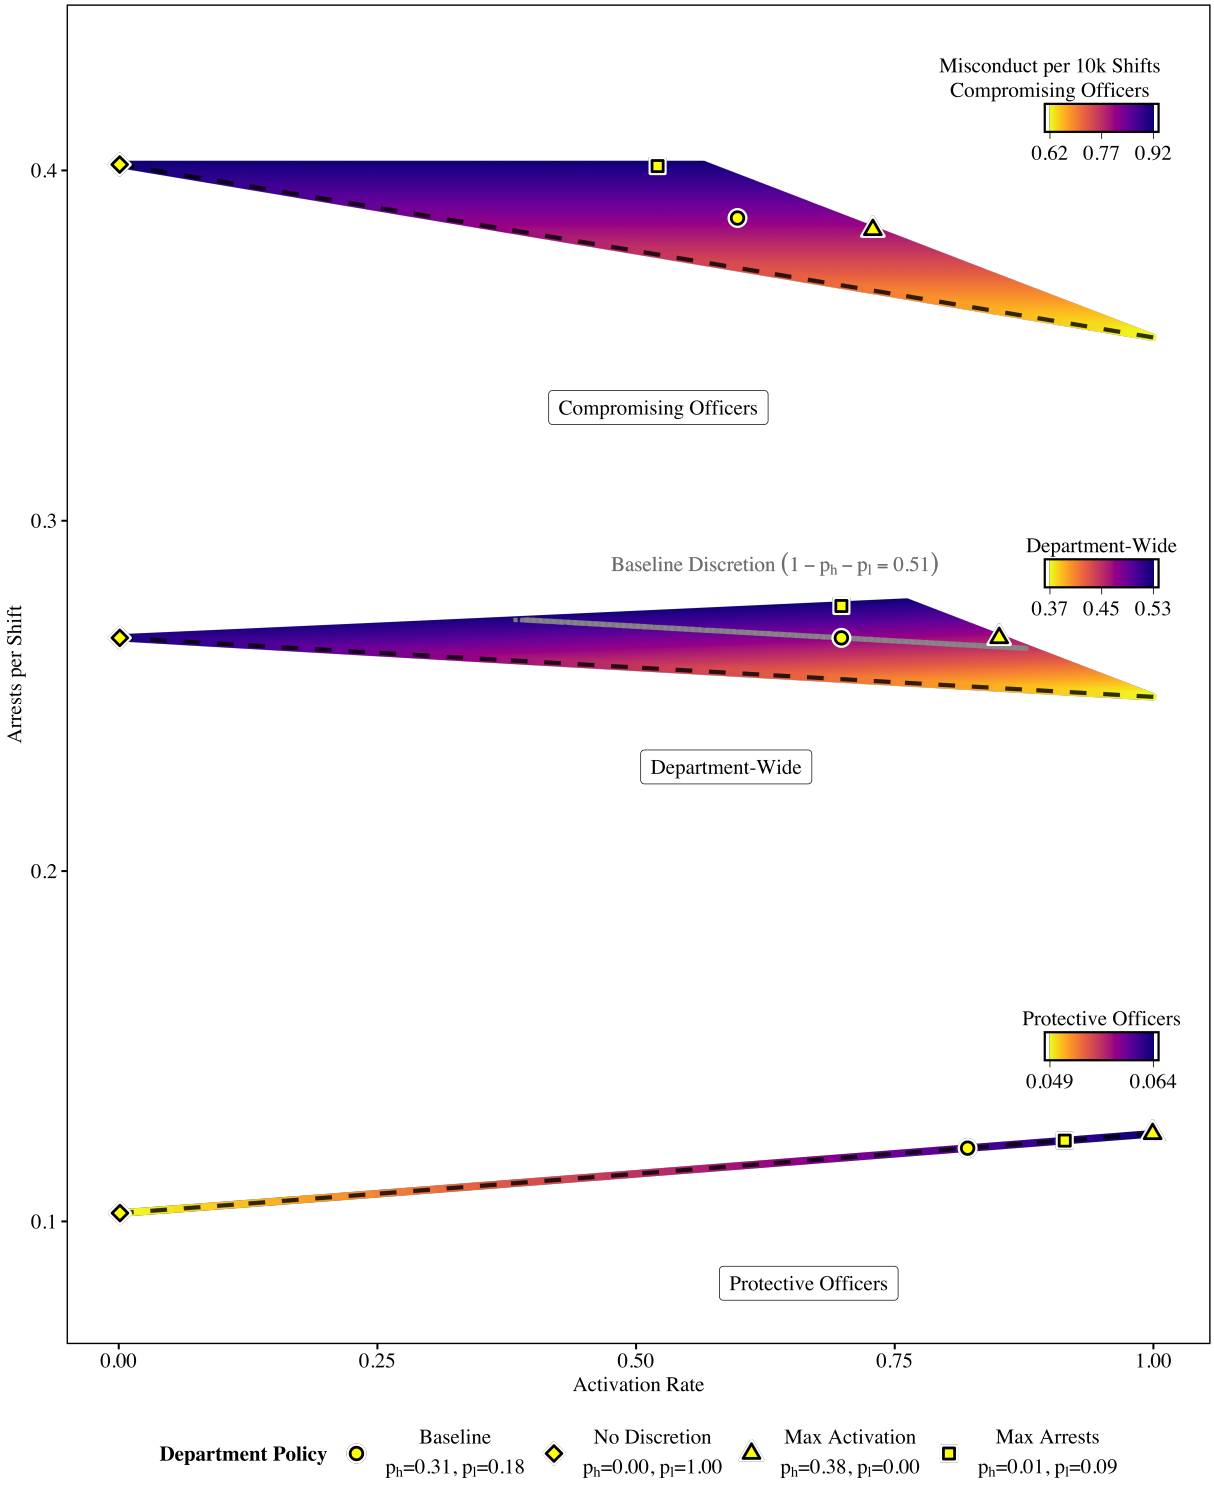


B. Strategic Non-Activation for Compromising Activators



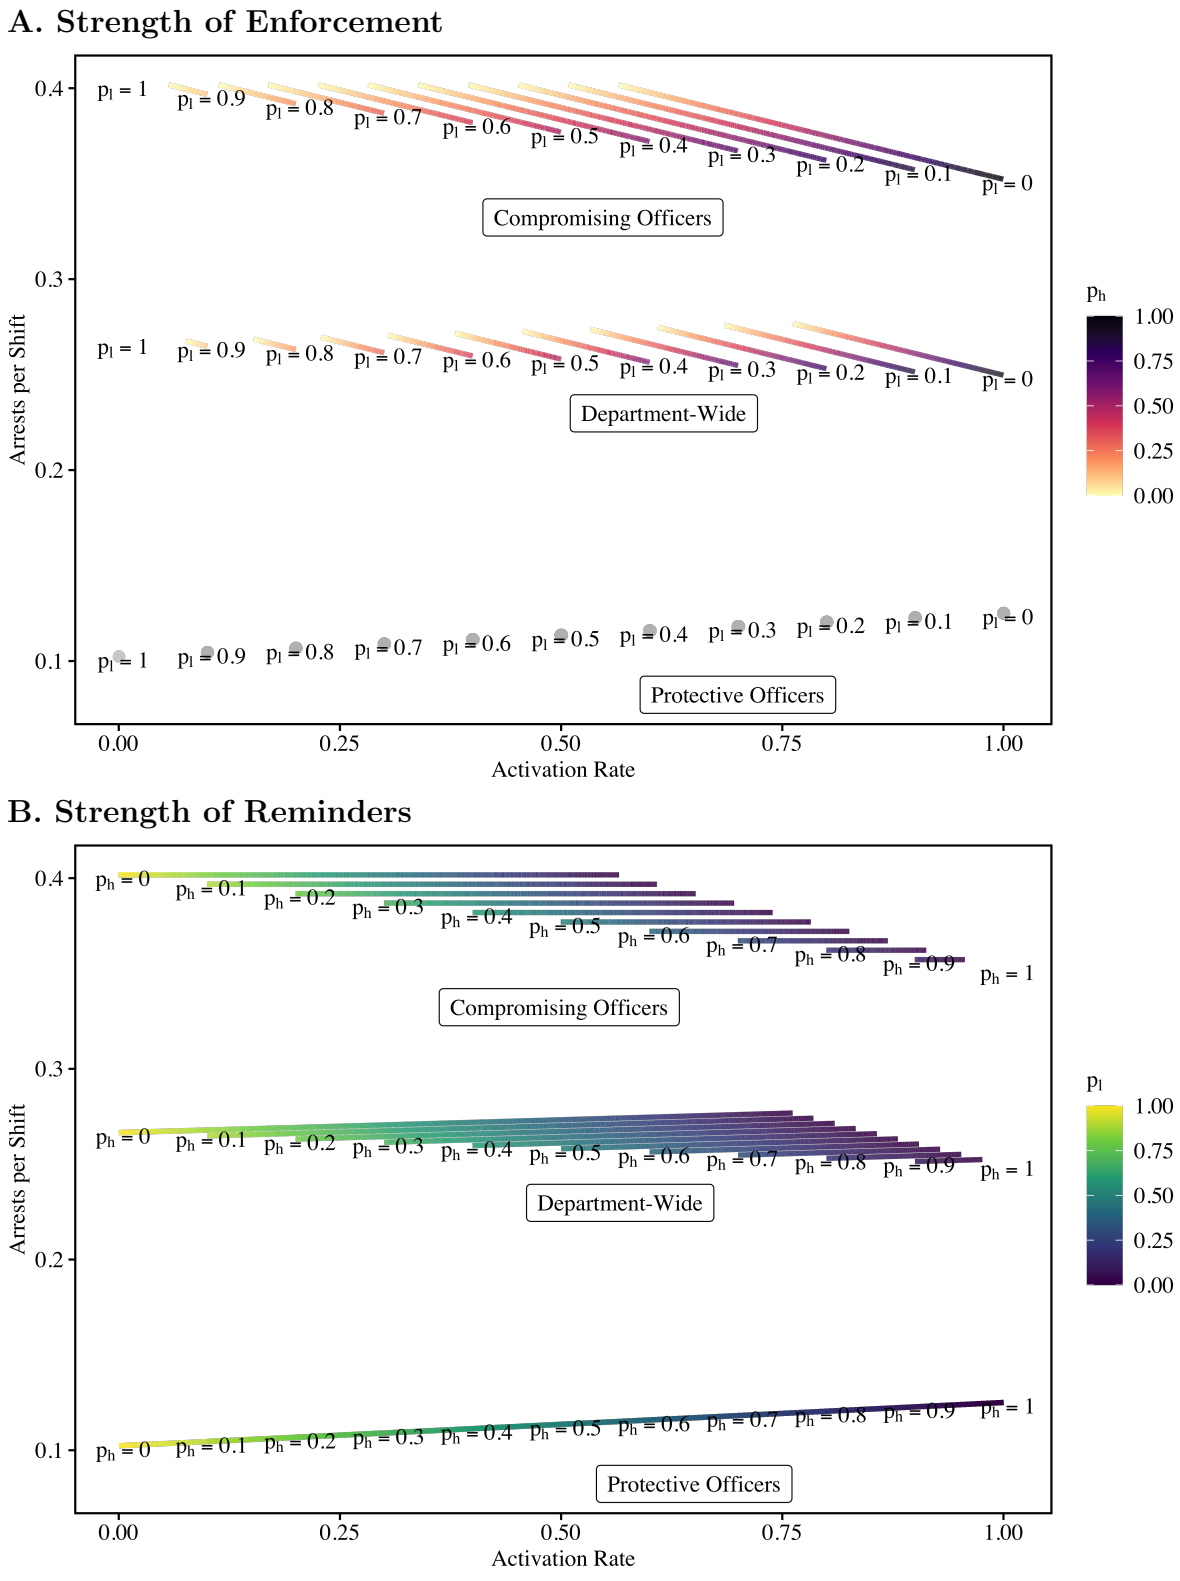
Note: This figure uses the structural model to document strategic non-activation. Panel A reports the distribution of simulated officer activation rates separately for compromising activators and protective activators. Using notation from the model, compromising activators are identified by $\theta_{arr} < \theta_{act}(0)$, and protective activators are identified by $\theta_{arr} < \theta_{act}(0)$. Panel B focuses on compromising activators and reports the percentage of non-activations that are strategic for officers with different overall activation rates. The line is the average, the dark orange area is the 25th and 75th percentiles, and the light shaded orange area is the 10th and 90th percentiles.

Figure 10: Counterfactuals: Activation, Arrests, and Misconduct with Optional Oversight



Note: This figure uses the structural model to construct counterfactual outcomes for different values of (p_l, p_h). The x-axis is the activation rate, the y-axis is the average arrests per shift, and the color of the shading is the number of misconduct events per 10,000 shifts. The department-wide outcomes are the combination of the compromising and protective activator outcomes.

Figure 11: Counterfactuals: Effects of Polices to Encourage Activation on Activation, Arrests, and Misconduct



Note: This figure uses the structural model to construct a counterfactual mapping of the strength of enforcement and reminder policies (p_l, p_h) onto the activation rate and arrests outcomes. Panel A reports the effect of changing p_h holding p_l fixed, where each line of points holds fixed p_l and the shading of each point reflects the value of p_h. Panel B reports the effect of changing p_l holding p_h fixed, where each line of points holds fixed p_h and the shading of each point reflects the value of p_l.

Tables

Table 1: Stability of Officer-Level Activation Rates

	Subset Dimension						
	Arrest	Time	District	Beat	Watch	Role	Manager
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. Raw Leave-Out Activation Rate							
Leave-Out Officer	0.71***	0.60***	0.45***	0.61***	0.46***	0.40***	0.55***
Activation Rate	(0.02)	(0.02)	(0.04)	(0.02)	(0.03)	(0.03)	(0.02)
R^2	0.25	0.24	0.19	0.20	0.20	0.13	0.21
Observations	6089	2608	697	4256	1138	1543	2617
B. Shrunken Leave-Out Activation Rate							
Shrunken Leave-Out	1.00***	0.69***	0.57***	0.79***	0.66***	0.57***	0.73***
Officer Activation Rate	(0.02)	(0.03)	(0.05)	(0.03)	(0.04)	(0.04)	(0.03)
R^2	0.31	0.30	0.25	0.27	0.27	0.17	0.31
Observations	6089	1229	404	1709	818	1169	1431

Note: This table investigates the extent to which differences in officer activation rates reflect a fixed officer-level component by documenting the within-officer stability of activation rates across policing environments. Panel A regresses each officer's activation rate in a particular environment on their leave-out activation rate elsewhere in terms of individual arrests (Column 1), time (a quarter, Column 2), geographic environment (district and beat, Columns 3 and 4), different times of the day (morning, day, and night shift, Column 5), arresting role (primary, secondary, assisting officer, Column 6), and under different lieutenants (Column 7, where the 14% of officer-arrests without lieutenant identifiers for the officer-shift are dropped). Panel B repeats Panel A but replaces the officer's leave-out activation rates with an Empirical Bayes leave-out activation rate that shrinks noisy estimates towards the mean activation rate. Standard errors are in parentheses and are clustered at the officer by subset dimension level. * p<0.1, ** p<0.05, *** p<0.01.

Table 2: Effect of Body-worn Cameras on Arrests

	Outcome			
	Difference	Arrests Per Shift		
	(1)	(2)	(3)	(4)
Body-worn Camera in District	-0.026***	-0.009	0.017*	-0.003
	(0.01)	(0.011)	(0.009)	(0.008)
Officers in Sample				
Low Activation	✓	✓		✓
High Activation	✓		✓	✓
Outcome Mean	-0.046	0.193	0.237	0.224
Observations	660	660	660	660

Note: This table estimates the effect of body-worn cameras in a district on arrests per officer-shift using the staggered rollout of body-worn cameras to police districts in a difference-in-differences setting. The unit of observation is the district by calendar month. The sample is from 6 months before the first district adopted body-worn cameras up to the month that the last district adopted body-worn cameras. Each column uses the procedure from Callaway and Sant'Anna (2021) not-yet-treated districts as a control group for a given event. The columns differ by the outcome used and the sample. The outcome in Column 1 is the difference in arrests per officer-shift between officers by whether they are in the bottom quartile of activation rates. The outcome in Column 2 is the arrests per officer-shift for officers in the bottom quartile of activation rates. The outcome in Column 3 is the arrests per officer-shift for officers in the top three quartiles of activation rates. The outcome in Column 4 is the arrests per officer-shift for all officers. Standard errors are in parentheses and are clustered at the district level. * p<0.1, ** p<0.05, *** p<0.01.

Table 3: Heterogeneity of the Effects of Body-worn Cameras on Arrests

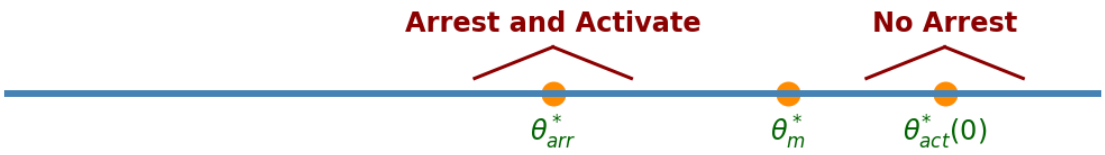
	Analysis							
	Mechanisms		Initiating Arrests		Hassle Costs			
	12/13 Months (1)	20/25 Months (2)	Primary Officer (3)	Assisting Officer (4)	Felony Arrest (5)	Misdem. Arrest (6)	Pub. Ord. Arrest (7)	Complaint Divide (8)
Body-worn Camera in District	-0.027**	-0.022*	-0.011***	-0.014*	-0.014**	-0.011*	-0.003	-0.036**
	(0.012)	(0.013)	(0.004)	(0.009)	(0.007)	(0.007)	(0.005)	(0.017)
Outcome Mean	0.216	0.218	0.096	0.119	0.049	0.167	0.058	0.221

Note: This figure investigates heterogeneity in the differential effect of body-worn cameras on arrests per shift between compromising and protective activators. Columns 1 and 2 assess the extent to which the primary underlying mechanism of interest—individual officers change their arrest behavior in response to body-worn cameras—is driving the result as opposed to officers responding to body-worn cameras by switching districts. To do so, the columns focus on officers who have a consistent assignment throughout the rollout period. Column 1 focuses on officers who are in the same district in 12 of the 13 months around body-worn camera adoption in their district (the month of rollout and 6 months on either side), and Column 2 focuses on officers who are in the same district in 20 of the 25 months considered in our analysis. Columns 3 and 4 assess the extent to which the results are likely driven by decisions to initiate an arrest by focusing on either only primary arresting officers (Column 3) or only assisting officers (Column 4). Columns 5-8 assess the possibility that body-worn cameras change officer behavior primarily because they introduce added "hassle costs" from a perceived need to be more formal in the encounter or increased paperwork associated with logging the video itself. Columns 5-7 report results separately for felony, misdemeanor, and public order arrests, and Column 8 divides officers into types by complaints which are likely unrelated to hassle costs. Standard errors are in parentheses and are clustered at the district level. * p<0.1, ** p<0.05, *** p<0.01.

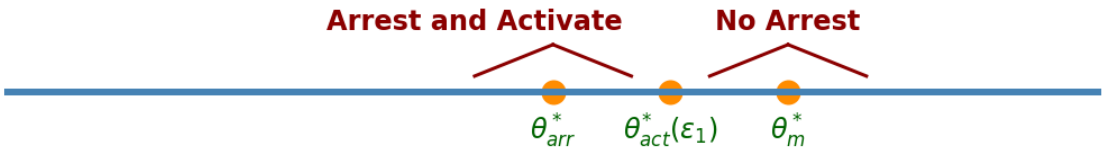
Supplemental Appendix A: Additional Results

Figure A1: Arrest and Activation Thresholds

A. Protective Activator with No Special Incentives



B. Protective Activator with Marginal Activations Disincentivized



C. Protective Activator with Inframarginal Activations Disincentivized



D. Compromising Activator with No Special Incentives



E. Compromising Activator with Marginal Activations Incentivized



Note: Panel A is for protective activators with no special incentives, defined as $\theta^*_{arr} < \theta^*_{act}(0), \epsilon_i = 0$. Panel B is for protective activators with marginal activations disincentivized, defined as $\theta^*_{arr} < \theta^*_{act}(0), \epsilon_1 < 0$. Panel C is for protective activators with inframarginal activations disincentivized, defined as $\theta^*_{arr} < \theta^*_{act}(0), \epsilon_2 < \epsilon_1$. Panel D is for compromising activators with no special incentives, defined as $\theta^*_{arr}(0) > \theta^*_{act}(0), \epsilon_i = 0$. Panel E is for compromising activators with marginal activations incentivized, defined as $\theta^*_{arr}(0) > \theta^*_{act}(0), \epsilon_3 > 0$.

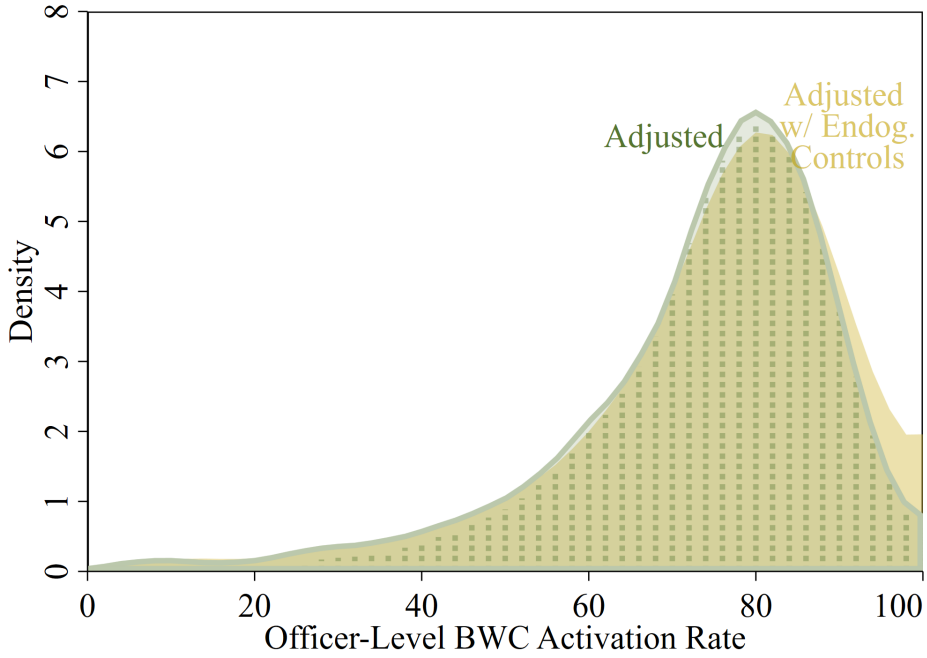
Figure A2: Activation Rates by Crime Type and Severity



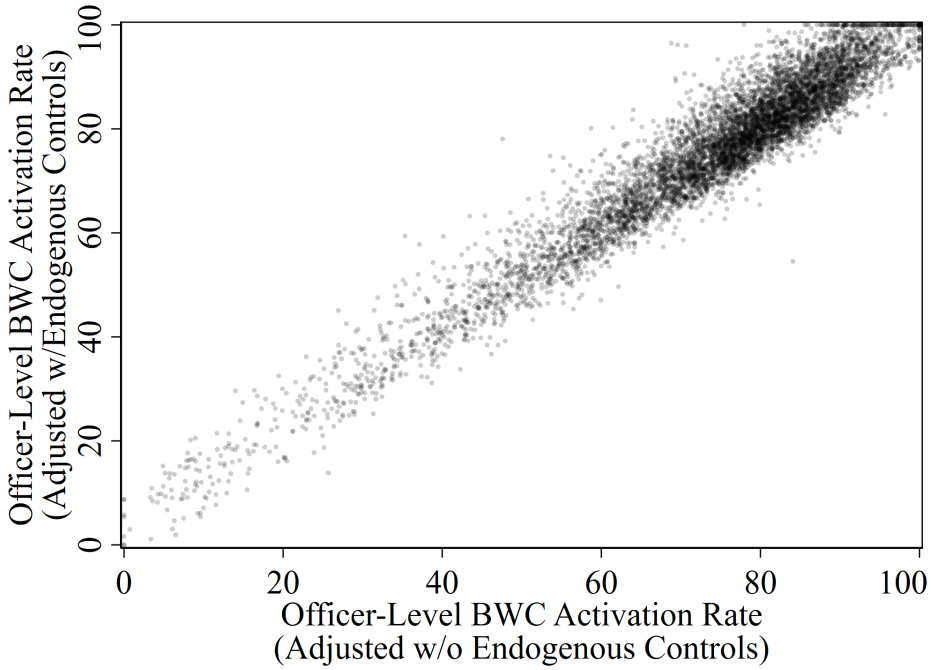
Note: This figure reports body-worn camera activation rates by crime types as listed on the x-axis and severity as indicated by the color of the bar. Misdemeanor is blue, and felony is orange.

Figure A3: Adjusted Officer Activation Rates with and without Endogenous Controls

A. Distributions

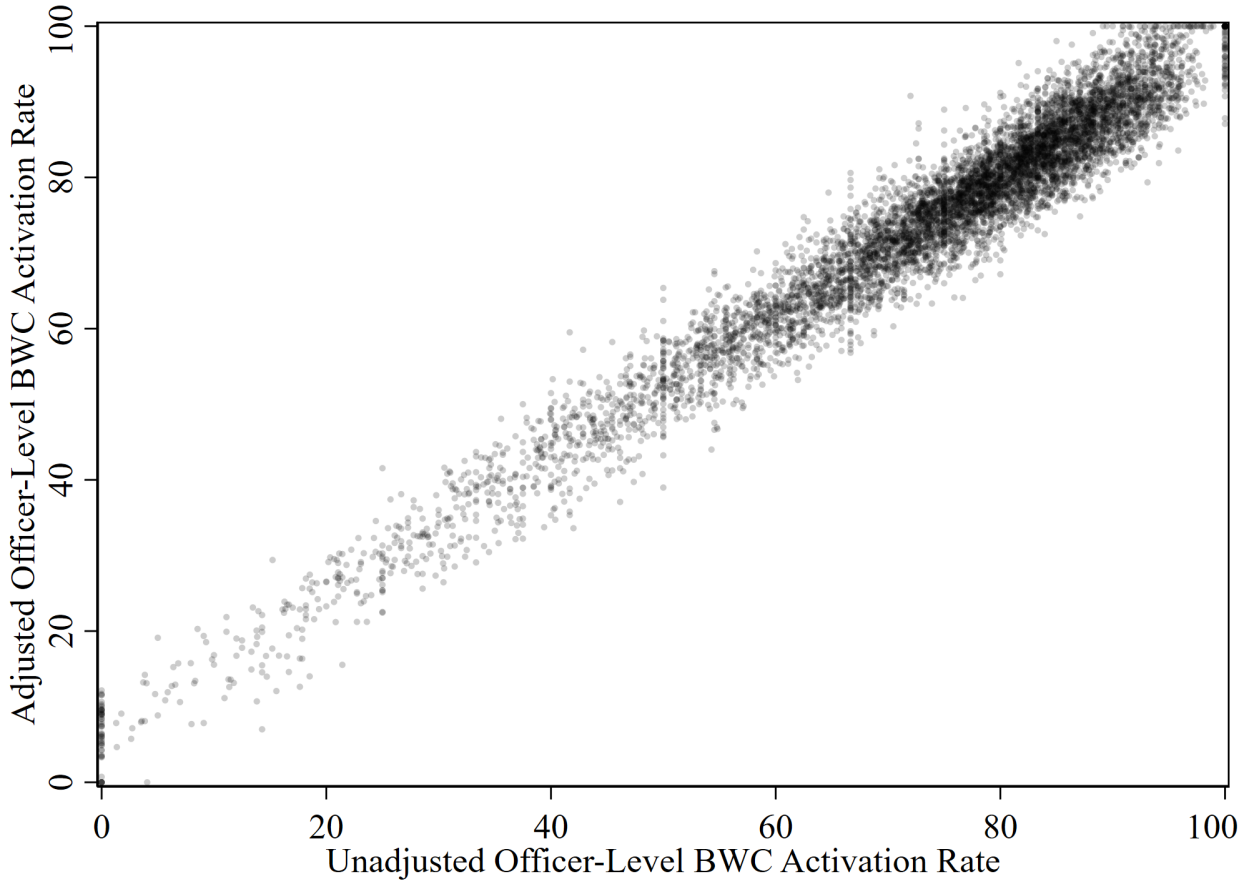


B. Relationship



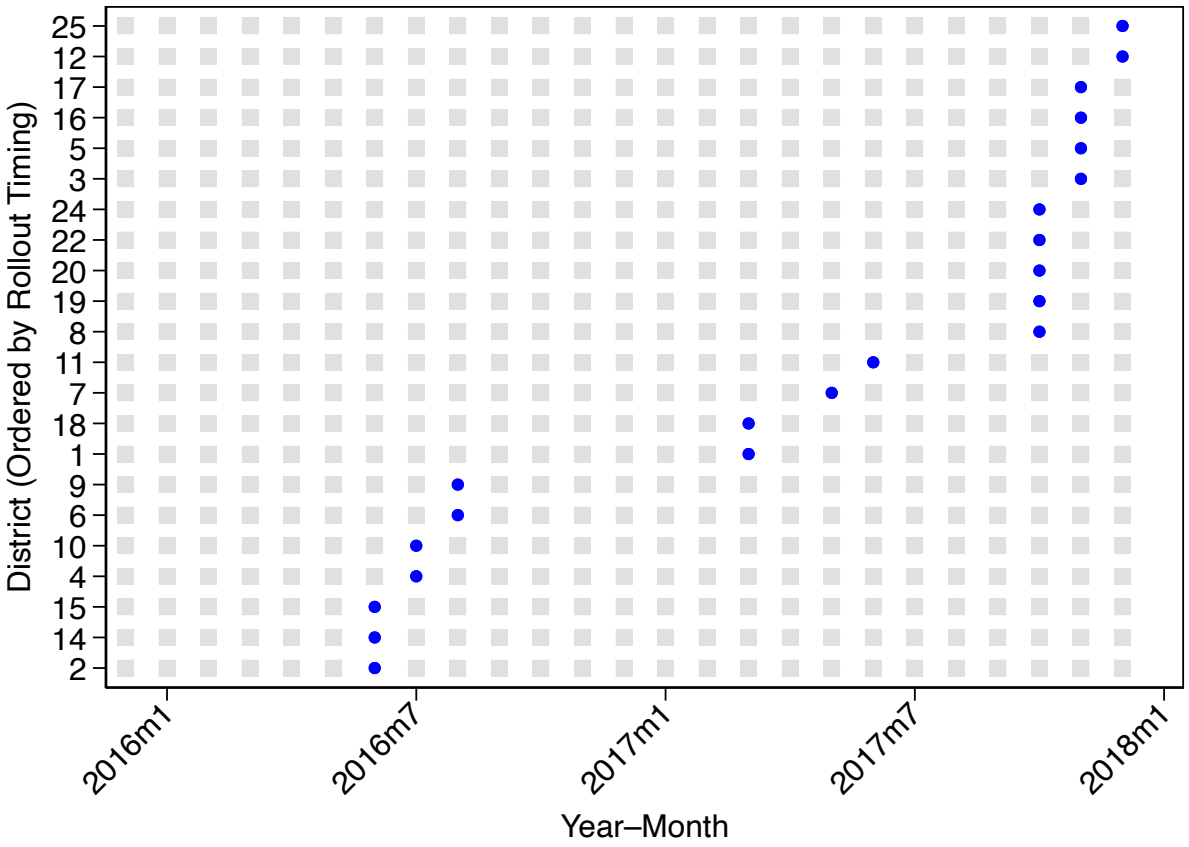
Note: This figure compares officer activation rates after accounting for environment, encounter, and role controls to officer activation rates that also control for factors that are potentially endogenous to activation and arrest decisions we are interested in (all other factors considered in Section 5). Panel A reports distributions of both rates, and Panel B reports a scatterplot of the fully adjusted rates on the rates adjusted without the endogenous controls.

Figure A4: Relationship between Unadjusted and Adjusted Officer Activation Rates



Note: This figure reports a scatterplot of officer activation rates after accounting for environment, encounter, and role controls on the y-axis and overall officer activation rates without any controls on the x-axis.

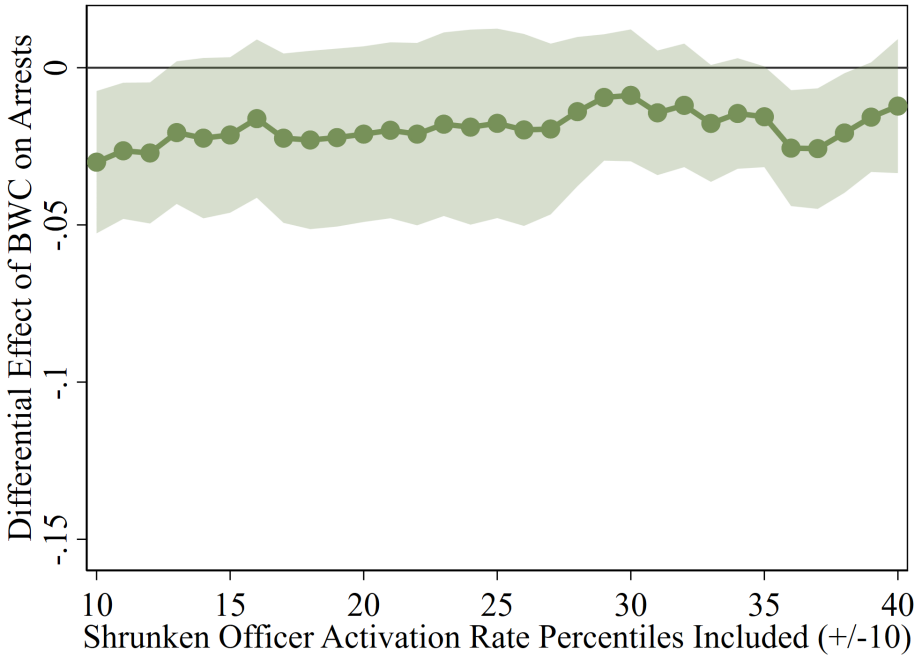
Figure A5: Body-worn Camera Rollout Across Police Districts in Chicago



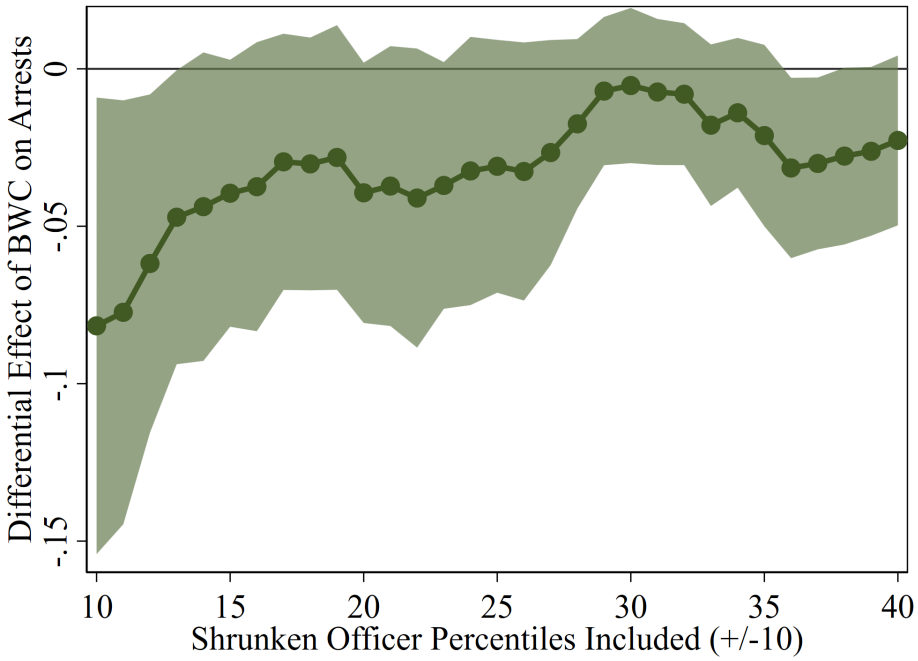
Note: This figure reports the month that each police district in Chicago adopted body-worn cameras.

Figure A6: Robustness: Addressing Noise in Officer Activation Rates by Estimating the Effect of Body-worn Cameras on Arrests using Shrunken Activation Rates

A. Sorting Officers by Activation Rates

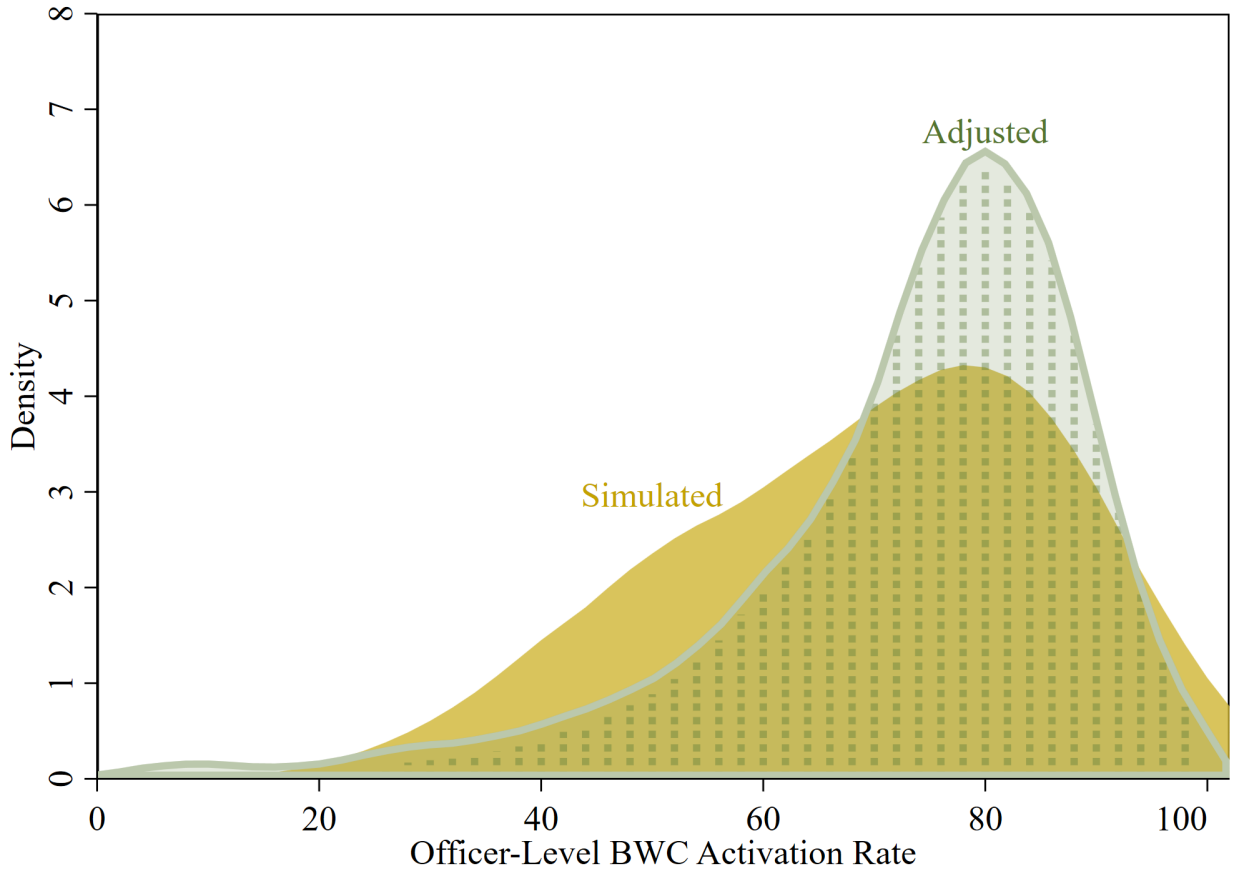


B. Sorting Officers by Activation Rates and Complaints



Note: This figure repeats the analysis in Figure 8 by sorting officers using shrunken adjusted officer activation rates instead of unshrunken adjusted officer activation rates.

Figure A7: Simulated Versus Actual Officer Activation Rates



Note: This figure reports the distribution of adjusted officer activation rates that is used to estimate the structural model and the resulting simulated distribution of officer activation rates from the structural model.

Table A1: Robustness of Effect of Body-worn Cameras on Arrests to Difference-in-Differences Specifications

	Design		
Estimator	DD	DDD	Estimate
(1)	(2)	(3)	(4)
Butts and Gardner (2022)		✓	-0.029***
			(0.01)
Goodman-Bacon (2021)	✓		-0.001
			(0.011)
de Chaisemartin and D'Haultfoeuille (2020)	✓		0.002
			(0.01)
Sun and Abraham (2020)	✓		-0.020***
			(0.007)
Borusyak et al (2024)	✓		-0.021*
			(0.011)
Dube et al (2023)	✓		-0.018
			(0.015)
Arkhangelsky et al (2021)	✓		-0.015
			(0.016)
TWFE		✓	-0.029***
			(0.008)
Stacked Event Study		✓	-0.017*
			(0.01)

Note: This table reports estimates similar to Columns 1 and 4 of Table 2 using alternate difference-in-differences specifications. Column 2-3 identifies the type of design used. "Triple Difference" is a design at the officer-month level where an indicator for body-worn camera in the district is interacted with either an indicator for officers who have activation rates at or below the 25th percentile with district by month fixed effects. "Differenced Outcome" is a design at the district-month level as in our primary specification. For specifications that require a never-treated group, we end the data in 10/2017, which leaves 11 districts never treated (Sun and Abraham, 2021; Arkhangelsky et al., 2021). Column 4 reports the estimated coefficient and standard errors clustered at the district level. * p<0.1, ** p<0.05, *** p<0.01.

Table A2: Alternative Activation Label Threshold

	Outcome			
	Difference	Arrests Per Shift		
	(1)	(2)	(3)	(4)
A. Accept When Match Probability Above 10%				
Body-worn Camera in District	-0.025***	-0.008	0.016*	-0.003
	(0.009)	(0.011)	(0.009)	(0.008)
Officers in Sample				
Low Activation	✓	✓		✓
High Activation	✓		✓	✓
Outcome Mean	-0.042	0.197	0.236	0.224
Observations	660	660	660	660
B: Accept when Match Probability Above 2.5%				
Body-worn Camera in District	-0.021**	-0.005	0.016*	-0.003
	(0.008)	(0.01)	(0.009)	(0.008)
Officers in Sample				
Low Activation	✓	✓		✓
High Activation	✓		✓	✓
Outcome Mean	-0.054	0.188	0.239	0.224
Observations	660	660	660	660

Note: This table investigates the robustness of the results to the threshold for labeling arrest-activation pairs that we accept as a match by reestimating Table 2 using different thresholds. Our primary approach accepts matches where the probability that the observed arrest time-activation time gap was drawn from our estimated distribution of gaps is at least 5%. Panel A reports results using an alternative threshold rule that is less demanding—specifically, a threshold at double the primary specification, 10%. Panel B reports results using a threshold that is more demanding—a threshold at half of the primary specification, 2.5%. Standard errors are in parentheses and are clustered at the district level. * p<0.1, ** p<0.05, *** p<0.01.

Supplemental Appendix B: Theoretical Model

Officer Effort

In this extension, we consider the case where officers can exert costly effort to reduce the probability of misconduct in an arrest.

Let this effort, $e \geq 0$, enter the officer's utility function both through the reward for an arrest, $w_j(e)$, and the probability of misconduct, $\theta_i(e) = \theta_i - \beta e$. We assume that $w_j'(e) < 0 \; \forall e$ and $w'(0) > -\beta c \Delta$ to ensure an interior solution to the effort problem. We also assume that the effect of e on θ_i is linearly separable in order to maintain threshold rules for arrest and activation that can be clearly expressed.

Now consider the problem of an officer choosing effort conditional on making the arrest:

$$\max_{e \geq 0} w_j(e) - c(\theta_i(e) q_T + (1 - \theta_i(e)) q_F) \tag{10}$$

This is solved by:

$$-\frac{w_j'(e^*)}{\beta} = c\Delta \tag{11}$$

Then, the officer chooses whether to make an arrest given the optimal e^*, which gives a threshold rule:

$$\theta^*_{arr} = \frac{\frac{w_j(e^*)}{c} - q_F}{\Delta} + \beta e^* \tag{12}$$

This is isomorphic to the original model when $w_j = w_j(e^*)$ and $\theta_i = \theta_i(e^*)$. That is, we can simply focus on the values of w_j and θ_i that result from the officer's effort maximization problem.

However, we should not expect the same isomorphism in the activation decision because BWC activation will change q_T and q_F and hence the optimal effort supplied by the officer. In particular, it will increase the LHS of Equation 11, so $\frac{\partial e^*}{\partial m} > 0$.

Turning to the activation decision, it is now:

$$\theta^*_{act}(\omega_j + \epsilon) = \frac{\frac{\omega_j+\epsilon}{c} - q'_F}{\Delta'} + \beta e^* \tag{13}$$

where we call ω the re-optimization shifter, $\omega_j = w'_j(e^*)\frac{\partial e^*}{\partial m} + \beta\frac{\partial e^*}{\partial m}c\Delta'$.[23] The sign of ω_j is ambiguous. Officers may appreciate the extra dividends paid by their increased effort being captured on camera, or they may be frustrated by the need to exert extra effort when monitored.

Finally, the change in the arrest threshold due to activation is given by:

$$\frac{\partial \theta^*_{arr}}{\partial m} = \frac{\Delta'}{\Delta}(\theta^*_{arr} - \theta^*_{act}(0)) + \beta\frac{\partial e^*}{\partial m} \tag{14}$$

Two things about this result are noteworthy. First, the change in the arrest threshold due to activation continues to include $\theta^*_{act}(0))$, but now setting $\epsilon_i = 0$ returns $\theta^*_{act}(\omega_j))$ instead. The reference incentive level is instead $\epsilon_i = -\omega_j$, just enough activation incentive to offset the re-optimization shifter. Second, the $\beta\frac{\partial e^*}{\partial m}$ puts upward pressure on $\frac{\partial \theta^*_{arr}}{\partial m}$ for all officers. This equation still allows the division of officers into protective and compromising on the basis of its sign, but that division no longer reflects the simple comparison of arrest thresholds to activation thresholds in the absence of additional incentives.

Overall, we expect that the ability to re-optimize in response to activation will make more officers protective activators because they can take better advantage of the changes in q_T and q_F that arise from activation. An increase in protective activators is a certainty when there is also a direct benefit of re-optimization, $\omega_j > 0$.

Correlations of Type with Activation and Complaints

At the end of Section 2, we claim that protective activators will have higher BWC activation rates and lower complaint rates relative to compromising activators. In this section, we show this using variation in the relative reward for an arrest, $\frac{w_j}{c}$, which we will

[23]For truly incremental changes in m, this re-optimization shifter will be 0 by the Envelope Theorem, but recall that we are primarily interested in discrete changes in m despite using notation for continuous changes.

hereafter call R_j and allow to differ across officers according to a distribution $H(R)$. Note that $\theta^*_{arr}(R_j)$ and $\theta^*_m(R_j)$ depend by R_j but not ϵ_i while $\theta^*_{act}(\epsilon_i)$ depends on ϵ_i but not R_j.

Given an R_j, define $\epsilon_{arr}(R_j)$ to be the value that sets $\theta^*_{act}(\epsilon_{arr}(R_j)) = \theta^*_{arr}$ and likewise define $\epsilon_m(R_j)$ to be the value that sets $\theta^*_{act}(\epsilon_m(R_j)) = \theta^*_m$. Any value R^p such that $\theta^*_{arr}(R^p) < \theta^*_{act}(0)$ belongs to a protective activator, while compromising activators have values higher values of R^c such that $\theta^*_{arr}(R^c) > \theta^*_{act}(0)$

For a generic R^p and R^c, activation rates for various ranges of ϵ_i are given by:

ϵ_i	R^p	R^c
$(-\infty, \epsilon_{arr}(R^p)$	$\frac{F(\theta^*_{act}(\epsilon))}{F(\theta^*_{arr}(R^p))}$	$\frac{F(\theta^*_{act}(\epsilon))}{F(\theta^*_{arr}(R^c))}$
$(\epsilon_{arr}(R^p), \epsilon_m(R^c))$	1	$\frac{F(\theta^*_{act}(\epsilon))}{F(\theta^*_{arr}(R^c))}$
$(\epsilon_m(R^c), \epsilon_{arr}(R^c))$	1	$\frac{F(\theta^*_m(R^c))}{F(\theta^*_m(R^c))+\theta^*_{arr}(R^c))-\theta^*_{act}(\epsilon))}$
$(\epsilon_{arr}(R^c), \infty$	1	1

Recalling that $\theta^*_{arr}(R^p) < \theta^*_{act}(0) < \theta^*_{arr}(R^c)$, the activation rate for R^p is greater than the activation rate of R^c for all draws of ϵ_i and all pairs of R^p and R^c. This establishes that protective activators have higher activation rates.

Finally, the facts that $\theta^*_{arr}(R^p) < \theta^*_{arr}(R^c)\ \forall R^p, R^c$ and $q_T > q_F$ imply that:

$$E[\theta q_T + (1-\theta)q_F | \theta < \theta^*_{arr}(R^c)] > E[\theta q_T + (1-\theta)q_F | \theta < \theta^*_{arr}(R^p)] \tag{15}$$

that is, in the absence of body-worn cameras, compromising activators will accumulate more complaints than protective activators.

Supplemental Appendix C: Activation Labels

This appendix documents 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 does match cleanly to the BWC data by event ID provides 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 C1 plots the distribution of "gaps" between the arrest timestamp and BWC activation window. We calculate the gap in each case as the distance between the arrest timestamp and the nearest endpoint of the BWC activation window, assigning 0 to timestamps that fall inside the window.[24] Positive values indicate that the arrest timestamp was recorded before the BWC recording started. 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 C1 into those that exceed 20 minutes on either side. We therefore find that nonzero 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 to use the event number matches to model these timing discrep-

[24] 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.

ancies 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:

$$Y_i = \alpha_c + \beta_c X_i + \epsilon_i \tag{16}$$

where i indexes the known matches, Y_i is the gap recorded in a given event number match, α_c and β_c are component specific constants and coefficient vectors, X_i is a vector of arrest characteristics, and ϵ_i is drawn from 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, which allows it to choose the optimal number of components and π_c, α_c, β_c, and σ_c for each. Table C1 reports 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 of especially long gaps.

We then apply this model to the set of 488,523 candidate "proximity matches" from our primary analysis sample. 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 have 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.

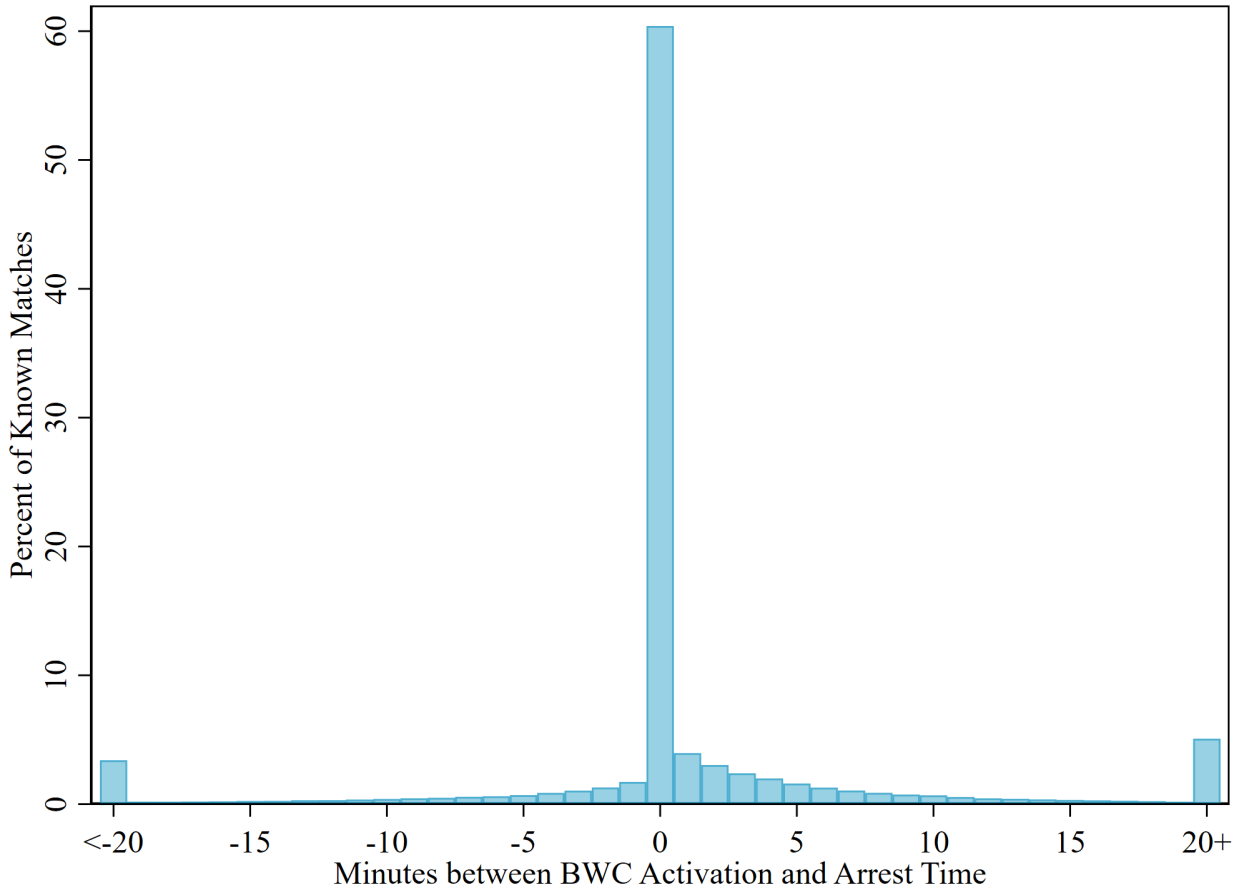
Figure C2 summarizes the results of this procedure. It is a histogram of measured gaps that stacks event number matches, accepted proximity matches, and rejected proximity matches. Observations with gaps of 0, of which there are 397,785, are omitted to preserve

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 C3 plots the distribution of arrests by the probability of a BWC activation assigned by this procedure. The vertical red line reports 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 C1 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 C1: Distribution of Arrest-Activation Gaps in Event Number Matches



Note: This figure reports the distribution of the number of minutes between the time of arrest and the closest endpoint of a BWC recording logged by the officer who made the arrest. Only arrest-BWC matches confirmed by event identifiers are included, where 0 indicates that 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.

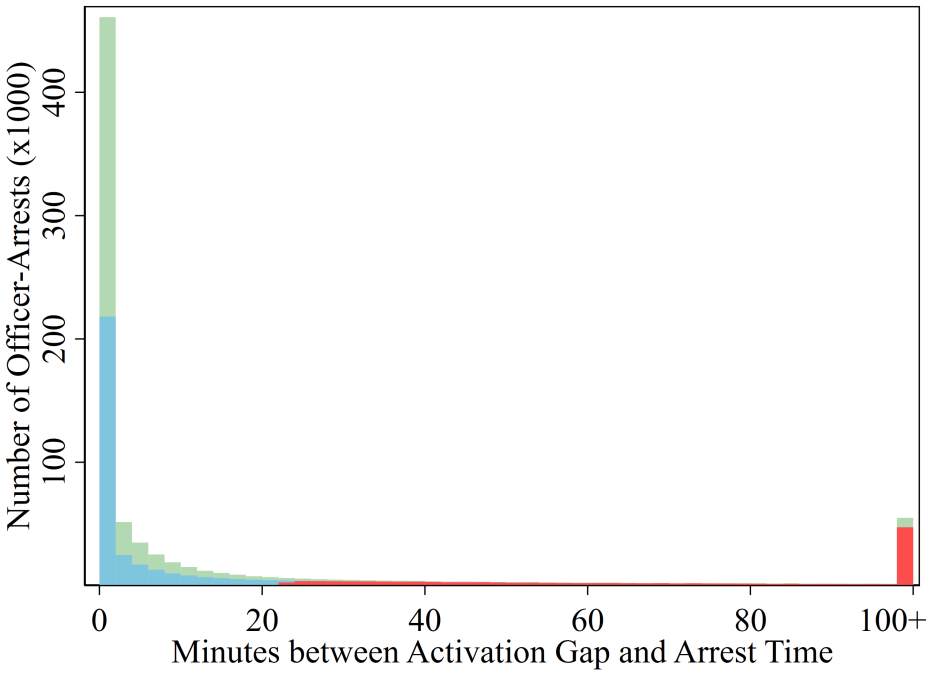
Table C1: Components of Activation Gap Model

Component (1)	Mean (2)	St. Dev. (3)	Weight (4)
1	0.3	1	0.56
2	146.9	99	0.24
3	844.7	410	0.15
4	5430	3121	0.05

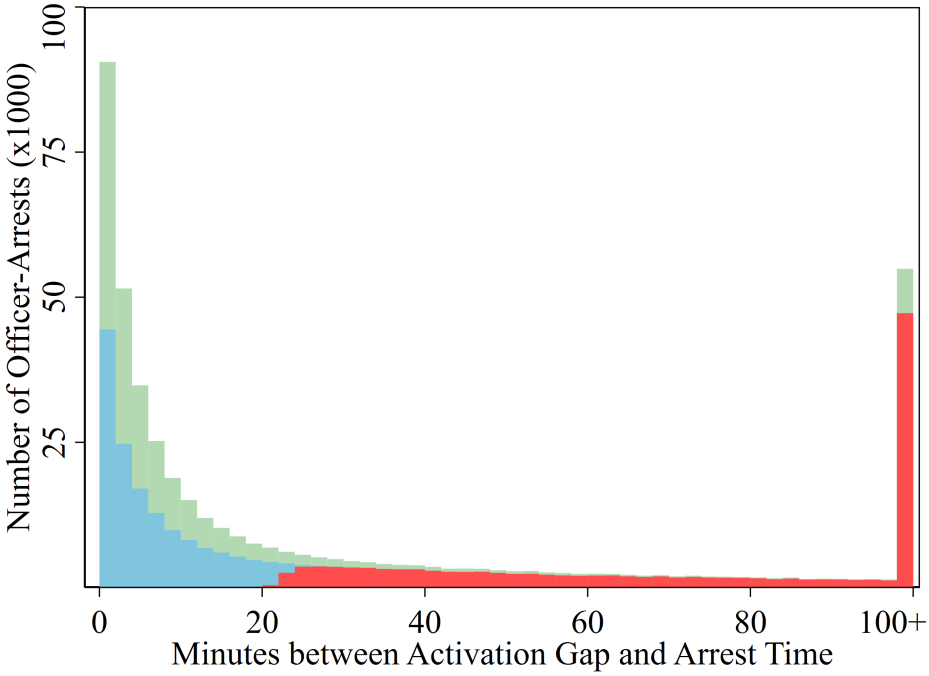
Note: This table reports information about the four components estimated in our Gaussian mixture model of the gaps between arrest and BWC timestamps. Column 1 identifies the component. Columns 2 and 3 report the mean and standard deviation of the component. Column 4 is the weight placed on that component in the mixture model.

Figure C2: Distribution of Arrest-Activation Gap Sizes by Match Type

A. All Officer-Arrests



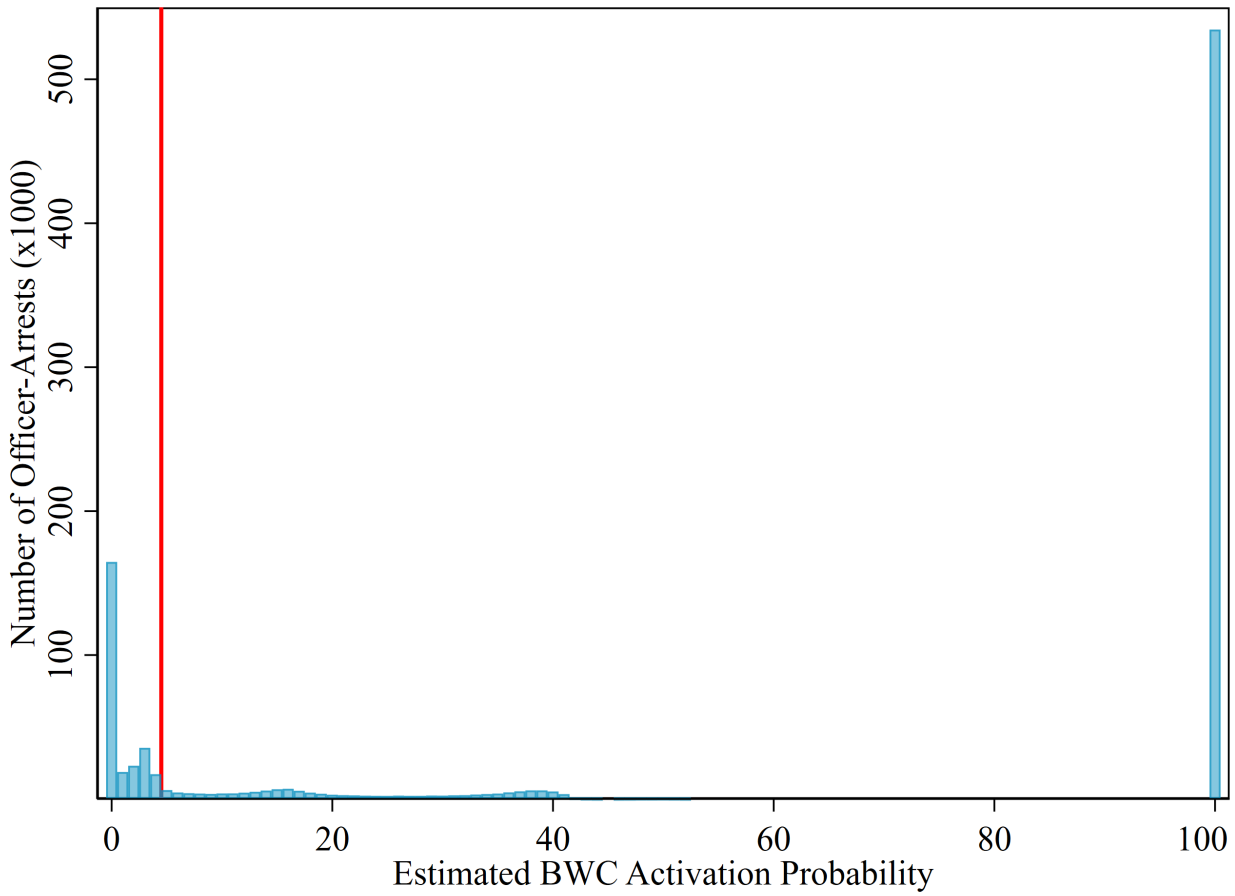
B. Officer-Arrests without Overlapping Activation and Arrest Times





Note: This figure reports the minutes between arrest time and the closest endpoint of a BWC recording logged by the officer who made the arrest. The unit of observation is the officer-arrest. Panel A includes all officer-arrests. Panel B excludes officer-arrests where the arrest time is within the time the BWC was recording. In each panel, the bars are stacked, so the total height of the bar is the total number of officer-arrests for a given bar. The green bars labeled "Confirmed Match" are officer-arrests where an event number matched the BWC activation directly to the arrest. The bars labeled "Proximity Match" are officer-arrests where an event number was not contained in either the activation data and/or the arrest data so that we cannot directly match the arrest to BWC activation and where we use the Gaussian mixture model of the gaps between arrest and BWC timestamps to label arrests as being associated with activation or not. The blue bars labeled "Proximity Match: Accepted" are officer-arrests that we label as having an activation, and the red bars labeled "Proximity Match: Rejected" are officer-arrests that we label as having no activation.

Figure C3: Distribution of Estimated BWC Activation Probabilities



Note: This figure reports the distribution of BWC activation probabilities estimated for arrests in our sample. BWC activation probabilities are 1 for arrests matched to an activation through event numbers, 0 for arrests with no activation within +/-4 hours of the arrest time, and are calculated using the CDF of the mixed Gaussian model fit to the event number matches for all other officer-arrests. The vertical red line at 0.05 reports our cutoff for labeling an officer-arrest as having an activation.

Supplemental Appendix D: Switchers Test

This appendix assesses the extent to which the persistence of activation rates within-officers across arrests, time, and situations can be explained by officers sorting to certain environments by using officers who switch beats in the spirit of switchers tests introduced in the teacher value-added literature (e.g., Chetty et al., 2014). The idea is that if officers' activation rates are fixed, then officers switching into beats who have higher activation rates than officers already in the beat should produce an increase in overall activation rates in that beat even if the type of beat they switch to is endogenous to their activation rate. The same holds in reverse for officers switching into beats who have lower activation rates.

Define an event, e, as an officer entering a new beat, b, and working in that new beat for at least 30 shifts over the next three months. Let Y_{et}^{Mover} be the activation rate of the officer entering the beat in event e in period t, where $t = 0$ is the three months before the move and $t = 1$ is the three months after the move. Let Y_{et}^{Peers} be the activation rate of the officers in the relevant beat but excluding the treated officer. Similarly, let $A_{et}^{\text{Mover}}, A_{et}^{\text{Peers}}$ be the number of arrests made. We estimate

$$\Delta_{et=1} = \alpha + \beta(\Delta_{et=0} s_{et=1}) + \epsilon_e \tag{17}$$

where$\Delta_{et} = \frac{Y_{et=1}^{\text{Mover}} A_{et=1}^{\text{Mover}} + Y_{et=1}^{\text{Peers}} A_{et=1}^{\text{Peers}}}{A_{et=1}^{\text{Mover}} + A_{et=1}^{\text{Peers}}} - Y_{et=0}^{\text{Peers}}$ is the change in the activation rate in the beat that the mover is entering, where the first term is the activation rate in the new beat after the move (a weighted average of activation rate for the mover and the peers); $\Delta_{et=0} = Y_{et=0}^{\text{Mover}} - Y_{et=0}^{\text{Peers}}$ is the pre-move difference in activation rates between moving officers and peer officers in the new beat in the three months before the move; and $s_{et=1} \frac{A_{et=1}^{\text{Mover}}}{A_{et=1}^{\text{Mover}} + A_{et=1}^{\text{Peers}}}$ is the moving officer's share of arrests in the new beat in the three months after the move. This scales the differential activation rate by the moving officer's contribution to the overall activation rate after the move. β captures the moving officer's impact on the activation rate in the beat. If moving officers have a fixed activation rate and there are no peer effects, we expect

the relationship between their arrest-portion of the activation rate to have a 1:1 relationship with the activation rate in the new beat. If peer effects exist, then the relationship should be more than 1:1.

Appendix Table D1 reports the results for the 1,699 entry and departure events combined (Column 1) along with 302 entry events (Column 2) and the 1,397 departure events (Column 3). In Column 1, we find that the confidence interval of the 1:0.94 relationship between the actual impact of officers entering and departing a beat and the impact that would be expected if officers had a fixed activation rate includes 1:1. In Columns 2 and 3, we find effects that are almost three times larger for the 302 entry events than in the 1,397 departure events.

Table D1: Switchers Test: Officers Moving Beats

	Event Included		
	All	Entry	Departure
	(1)	(2)	(3)
Leave-Out Excess Activation Rate	0.94***	2.28***	0.79***
	(0.12)	(0.38)	(0.12)
Activation Rate without Moving Officer	0.72	0.72	0.72
Observations	1,699	302	1,397

Note: This table reports results of a switcher's test for events where an officer enters a new beat and works in that new beat for at least 30 shifts over the next three months. For a given event, the outcome is the change in the activation rate in the beat that the mover is entering, and the independent variable labeled "Leave-Out Excess Activation Rate" is the pre-move difference in activation rates between moving officers and peer officers in the new beat in the three months before the move. The estimates come from estimating Equation 17 and capture the moving officer's impact on the activation rate in the beat. Column 1 includes events for officers entering and departing a beat. Column 2 includes events only for officers entering a beat. Column 3 includes events only for officers departing a beat. Standard errors are in parentheses and are clustered at the event level. * p<0.1, ** p<0.05, *** p<0.01.

Supplemental Appendix E: 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.[25] Call this probability ψ.

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

$$C = E[(1 - \theta_i)(q_T + \psi) + \theta_i | \theta_i < \theta^*_{arr}] \tag{18}$$

and given a rate of sustained complaints S:

$$S = \frac{E[(1 - \theta_i)q_T + \theta_i q_F | \theta_i < \theta^*_{arr}]}{C} \tag{19}$$

We assume that θ_i 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_i fall below the arrest threshold θ^*_{arr}. Removing that conditioning entirely, we can calculate:

$$E[\theta_i] = \frac{CS - q_F}{q_T - q_F} \tag{20}$$

and:

$$A = 2E[\theta_i] \tag{21}$$

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

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

$$\psi = C - E[\theta_i] - (1 - E[\theta_i])q_F \tag{22}$$

Next, we turn our attention to the officer-level distribution of arrest incentives, $\frac{w_j}{c}$. We assume this is log-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:

$$Pr\{\theta_i < \theta^*_{arr}\} = \frac{\frac{c}{w_j} - q_F}{q_T - q_F} \tag{23}$$

$$E[Pr\{\theta_i < \theta^*_{arr}\}] = \frac{\exp\left(\mu + \frac{\sigma^2}{2}\right) - q_F}{q_T - q_F} = R \tag{24}$$

$$Var[Pr\{\theta_i < \theta^*_{arr}\}] = \frac{(\exp(\sigma^2) - 1)\exp\left(2\mu + \sigma^2\right)}{(q_T - q_F)^2} = V \tag{25}$$

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, so we will also match what we know about how arrest patterns differentially changed following the introduction of BWCs. To respect the inherent bounds that $q_F, q_T \in [0, 1]$ even after BWC

activation, we estimate a parameter $\delta_F = -\frac{q'_F}{q_F}$ with domain $(0, 1)$.[26] Given $\theta^*_{act}(0)$ and δ_F, we can back out both q'_F and q'_T.

The model allows for an arbitrary ϵ_i shifter to activation incentives. For the structural estimation, we allow only extreme values of ϵ_i. 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 $\frac{w_j}{c}$ and the parameters $\{q_T, q_F, \theta^*_{act}(0), \delta_F, p_l, p_h\}$ 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_i, and general finite sampling error. The final vector of 7 parameters to estimate is thus $\{q_T, q_F, \theta^*_{act}(0), \delta_F, p_l, p_h, \sigma_u\}$.

Estimation Procedure

Some aspects of the pre-BWC environment can be estimated directly from the 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 average arrest rate of officers above 95th percentile of this distribution (0.852 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 a proportion $R = 0.327$ of available arrests with a variance of $V = 0.088$. Restricting attention to just those 5% most prolific officers, to avoid selection into less risky arrests, we find a rate of $C = 0.0035$ complaints per arrest and $S = 0.011$ sustained complaints per complaint.

[26] This does not per se guarantee that $q_T + q'_T < 1$, but we find in practice that the bound on false positives always binds first.

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 $\frac{w_j}{c}$. Within an officer, individual arrest opportunities have different draws from the distribution of θ_i, but we do not simulate these draws directly, handling them instead with analytic expressions. For each officer, we estimate a body-worn camera activation rate as well as an arrest rate with and without body-worn cameras. We model sampling variance in activation rates in the following way: given an officer with an analytic activation rate, r_j, and who makes a number of arrests, n_j, which we draw from $U(50, 100)$, the observed activation rate for that officer is $b_j + u_j$. Here, b_j is drawn from $Binomial(n_j, r_j)$ and u_j is drawn from $N(0, \sigma_u)$. We are only interested in aggregate changes in arrest rates, so we do not model sampling variation in them.

Our goal is to select parameters in $\Omega = \{q_T, q_F, \theta^*_{act}(0), \delta_F, p_l, p_h, \sigma_u\}$ in order to match the distribution of adjusted officer-level activation rates shown in Figure 4 and our results on differential changes in arrest rates following the introduction of BWCs. We select 9 moments of the activation 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: our estimate in Column 1 of Table 2 of the differential change in arrests among the lowest activators, d_{diff}; and our estimate in Column 4 of Table 2 of the aggregate change in arrests following BWC rollout, d_{agg}, which we simulate as $\tilde{d}_{agg}$, the difference in average arrest rates in the BWC and no BWC regimes. In the simulation, we calculate this as, $\tilde{d}_{diff}$, the differential change in arrest rates between the BWC and no BWC regimes of compromising activators (those with $\theta_{act}(0) < \theta_{arr}$) and protective activators (those with $\theta_{arr} < \theta_{act}(0)$).[27] Finally, we include a penalty term for

[27]We weight these two moment conditions higher than the others because even small deviations imply

having $p_h + p_l > 1$. Thus, our final objective function is:

$$O = \sum_{k=1...99} (p_k - \tilde{p}_k)^2 + 100(d_{agg} - \tilde{d}_{agg})^2 + 100(d_{diff} + \max(100(p_h + p_l - 1), 0)) \quad (26)$$

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).[28]

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 log-normal and 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 a 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. This distribution of protective activators centered on p_l is the first component of the Gaussian mixture.

The second component of the Gaussian mixture is made up of the compromising activators. It 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 the mean of this distribution as 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

arrest responses that are qualitatively different from what we estimate.

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

identified by the distribution of activation alone. 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. Including d_{agg} and d_{diff} in the model therefore helps to separately identify the changes in error rates via δ_F.

Estimation Results

Tables E1 and E2 report 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, 0.126. Meanwhile, we undershoot the 10th and 25th percentiles of the activation distribution because our simulated distribution does not drop off quite as quickly as in the low-but-long left tail in the data. Making either the distribution of $\frac{w_j}{c}$ or θ_i more flexible in a way that allows some officers to have particularly high complaint risk could address these problems at the cost of degrees of freedom.

We include our estimate of A, even though it is not a free parameter, to emphasize the fact that misconduct is rare. Even the arrest with the highest θ_i has only a 0.05% 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 a 12.4% chance of facing a sustained complaint as a consequence. The probability of punishment given no misconduct is five 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 compromising, 54.8% of our simulated sample, and those who find it protective, the remaining 45.2%. 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 $p_l = 17.5\%$, consistent with the fact that very few officers activate in all or even 90% of arrests. We find a value of $p_h = 31.2\%$,

which indicates that officers are not infrequently compelled to activate in situations where they otherwise would not. The ad-hoc noise parameter $\sigma_u = 0.103$ then smooths out the distribution of activation rates.

Finally, we find that BWC activation has a large effect on the probability of a false positive investigation, reducing it by 95%. On the other hand, the increase in the rate of true positive investigations, 3.3 percentage points, is meaningful relative to the baseline rate but still leaves the probability of accurate punishment substantially below 0.

Table E1: Structural Estimation Results for the Moments of Officer Activation Rates

Moment	Data	Simulated
1st Percentile	0.126	0.268
5th Percentile	0.391	0.381
10th Percentile	0.515	0.448
25th Percentile	0.672	0.576
50th Percentile	0.778	0.724
75th Percentile	0.852	0.84
90th Percentile	0.905	0.923
95th Percentile	0.936	0.967
99th Percentile	0.999	1.04
Arrest Change	-0.004	0
Low Activator Arrest Differential	-0.031	-0.034

Note: This table reports the fitted output of our structural estimation for the moments of the officer-level activation rate distribution in both our BWC data and the dataset simulated with our fitted parameters.

Table E2: Structural Estimation Results for Fitted Parameter Values

Parameter	Estimate
q_F ($\times$100k)	0.422
q_T ($\times$100	0.124
$\theta_{act}(0)$ ($\times$1000)	0.121
p_h	0.312
p_l	0.175
σ_u	0.103
q'_F ($\times$100k)	-0.399
q'_T	0.033
A ($\times$1000)	0.545

Note: This table reports the fitted output of our structural estimation for the fitted parameter values.